

Contributing Authors

Gaetano Antinolfi

Department of Economics
Washington University in St. Louis
One Brookings Drive
St. Louis, MO 63130-4899
gaetano@artsci.wustl.edu

Costas Azariadis

Department of Economics
Washington University in St. Louis
One Brookings Drive
St. Louis, MO 63130-4899
cre@artsci.wustl.edu

Ravi Bansal

Fuqua School of Business
Duke University
Durham, NC 27708
ravi.bansal@duke.edu

James B. Bullard

Federal Reserve Bank of St. Louis
P.O. Box 442
St. Louis, MO 63166-0442
bullard@stls.frb.org

John H. Cochrane

Graduate School of Business
University of Chicago
5807 S. Woodlawn Avenue
Chicago, IL 60637
john.cochrane@chicagogsb.edu

Riccardo DiCecio

Federal Reserve Bank of St. Louis
P.O. Box 442
St. Louis, MO 63166-0442
Riccardo.DiCecio@stls.frb.org

Martin Fukáč

Reserve Bank of New Zealand
P.O. Box 2498
Wellington
New Zealand
martin.fukac@rbnz.govt.nz

Michael F. Gallmeyer

Mays Business School
Department of Finance
Texas A&M University
College Station, TX 77843-4218
mgallmeyer@mays.tamu.edu

Burton Hollifield

David A. Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3890
burtonh@andrew.cmu.edu

Peter N. Ireland

Department of Economics
Boston College
140 Commonwealth Avenue
Chestnut Hill, MA 02467-3895
irelandp@bc.edu

Narayana R. Kocherlakota

Department of Economics
University of Minnesota
271 19th Avenue South
Minneapolis, MN 55455
nkocher@econ.umn.edu

Pamela A. Labadie

Department of Economics
The George Washington University
1922 F Street, NW
Old Main, Suite 208
Washington, DC 20052
labadie@gwu.edu

Edward Nelson

Federal Reserve Bank of St. Louis
P.O. Box 442
St. Louis, MO 63166-0442
edward.nelson@stls.frb.org

Lee E. Ohanian

Department of Economics
University of California, Los Angeles
Box 951477
8283 Bunch Hall
Los Angeles, CA 90095-1477
ohanian@econ.ucla.edu

Adrian Pagan

Queensland University of Technology
School of Economics and Finance
Queensland University of Technology
GPO Box 2434
Brisbane QLD 4001
Australia
a.pagan@qut.edu.au

Francisco J. Palomino

David A. Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3890
fjp@andrew.cmu.edu

Glenn D. Rudebusch

Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94120-7702
Glenn.Rudebusch@sf.frb.org

Brian Sack

Macroeconomic Advisers
231 South Bemiston Avenue, Suite 900
St. Louis, MO 63105
bsack@macroadvisers.com

Thomas J. Sargent

Department of Economics
New York University
269 Mercer Street
New York, NY 10003
thomas.sargent@nyu.edu

Eric T. Swanson

Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94120-7702
eric.swanson@sf.frb.org

Stanley E. Zin

David A. Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3890
zin@cmu.edu

An Estimated DSGE Model for the United Kingdom

Riccardo DiCecio and Edward Nelson

The authors estimate the dynamic stochastic general equilibrium model of Christiano, Eichenbaum, and Evans (2005) on U.K. data. Their estimates suggest that price stickiness is a more important source of nominal rigidity in the United Kingdom than wage stickiness. Their estimates of parameters governing investment behavior are only well behaved when post-1979 observations are included, which reflects government policies until the late 1970s that obstructed the influence of market forces on investment. (JEL E31, E32, E52)

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 215-31.

In this paper we estimate a dynamic stochastic general equilibrium (DSGE) model with nominal rigidities for the U.K. economy. The model we estimate is due to Christiano, Eichenbaum, and Evans (2005; CEE) and has become a benchmark, matching important aspects of the U.S. data while also being derived from optimizing behavior.

Interest in DSGE modeling of the United Kingdom has been heightened in recent years with the introduction of the Bank of England quarterly model (BEQM) into the U.K. monetary policy process. This model is based to a considerable degree on explicit optimizing foundations; see Harrison et al. (2005) for the model and Pagan (2005) for a discussion. BEQM is, however, dissimilar in important respects from the CEE model of the United States and the variant of the CEE model that Smets and Wouters (2003) estimate for the euro area. These dissimilarities make it difficult to use BEQM to compare the structure of the U.K. economy with that of other economies. For example, the estimation procedure for BEQM is different from that used by CEE and by Smets

and Wouters; portions of the BEQM model are estimated over a considerably shorter sample than CEE consider for the United States, and there are deviations from explicit optimization in the dynamics of the BEQM model.

All in all, it is probably fair to say that there has been considerably less work done for the United Kingdom in terms of DSGE modeling with systems estimation than there has been for other economies. But U.K. data may contain a type of information that is ideal for estimation of a DSGE model—specifically, information on private sector responses to policy actions. As the present governor of the Bank of England, Mervyn King, observed some 30 years ago,

Maintenance of the existing order and existing rates produces no information, whereas more information can be obtained by making changes. In this respect the U.S....is at a disadvantage by comparison with the U.K. A good illustration of this is afforded by the excitement generated amongst American economists in the 1960s by the investment tax credit and the attempts to assess its effects. A British economist would have shrugged this off as a

Riccardo DiCecio is an economist at the Federal Reserve Bank of St. Louis, and Edward Nelson is an assistant vice president at the Federal Reserve Bank of St. Louis and a research affiliate at the Centre for Economic Policy Research. The authors thank Adrian Pagan and participants at the policy conference and in seminars at the Federal Reserve Bank of Atlanta for comments on an earlier draft and Nick Davey, Jennifer Greenslade, and Kenny Turnbull for helpful information. Justin Hauke provided research assistance.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Table 1**Parameters in the Model**

Parameter	Description	Parameter	Description
β	Discount factor	ξ_p	Degree of price stickiness
α	Production elasticity with respect to capital	ξ_w	Degree of wage stickiness
δ	Depreciation rate	$1/\kappa$	Elasticity of investment to \hat{p}_K
μ	Steady-state money-growth rate	$1/\sigma_a$	Elasticity of utilization to \hat{r}^k
ψ_l	Relative weight of leisure in utility	ρ	Interest rate smoothing parameter
ψ_q	Relative weight of real balances in utility	r_π	Interest rate response to inflation
λ_w	Wage markup	r_y	Interest rate response to output
$1/\sigma_q$	Interest semi-elasticity of money demand	$r_{\Delta\pi}$	Interest rate response to change in inflation
b	Habits parameter	$r_{\Delta y}$	Interest rate response to change in output
λ_f	Price markup	σ_ε	Standard deviation of monetary policy shock

mere trifle compared to the changes he had witnessed over the years. (King, 1977, p. 6)

This observation, though made with reference to the changes wrought in U.K. fiscal policy up to the 1960s, applies tenfold to monetary policy experience in the period since the 1960s. Over that period, the United Kingdom has undergone great variation in inflation, interest rates, and monetary regimes.¹ It is true that for estimation this is a mixed blessing because large regime changes make it problematic to estimate a structural model over a long sample. But Christiano, Eichenbaum, and Evans (1999) and Sims and Zha (2006) argue for the United States that constant-parameter policy reaction functions may be reasonable approximations even over long samples, a view also implicit in CEE's (2005) choice of a 1965-95 estimation period. In modeling the United Kingdom using a DSGE model, we make a compromise between these positions by treating the

¹ Various advantages of the U.K. data for testing macroeconomic hypotheses have been stressed by Ravn (1997) (evaluating a real business cycle model against the behavior of U.K. real aggregates), Nelson and Nikolov (2004) (using a small New Keynesian model to evaluate different U.K. policy regimes), and Benati (2004) (assessing the behavior of U.K. data moments over the postwar period).

period since 1979 as a single regime,² but also by presenting results for pre-1979 and a long sample covering 1962-2005.³

We present in the following sections our model, estimates for our main sample, and results for the longer sample, with a discussion of other regime-change issues.

MODEL

The model is the same as that in CEE (2005), by now standard in the DSGE literature. The model incorporates both nominal frictions (sticky prices and wages) and dynamics in preferences and production (habit formation in consumption, investment adjustment costs, and variable capital utilization). The pattern of timing in agents' decisions is consistent with the VAR identification

² Some work for the United Kingdom (e.g., Castelnuovo and Surico, 2006) focuses on 1992 as the start of the present policy regime. But our use of a baseline sample period that starts in the late 1970s matches the choices implied by some of the BEQM-equation estimation periods (see, e.g., Harrison et al., 2005, pp. 115-20).

³ The long-sample estimates are the U.K. analogue to the Del Negro et al. (2005) treatment of the U.S. sample 1954-2004 as a single regime (including an unchanged inflation target).

Table 2**Variable Definitions**

Variable	Description	Variable	Description
\hat{m}	Real money supply	$\hat{\pi}$	Inflation
\hat{q}	Households' demand for money	\hat{w}	Real wage
\hat{k}	Capital services	\hat{r}^k	Rental rate on capital
$\hat{\bar{k}}$	Physical capital	$\hat{\mu}$	Growth rate of nominal money stock
\hat{l}	Labor	\hat{p}_K	Price of capital
\hat{c}	Consumption	\hat{U}_C	Marginal utility of consumption
\hat{i}	Investment	\hat{R}	Nominal interest rate
\hat{y}	Output		

restriction that we use in the next section. In our outline here of the linearized version of the model, all variables are expressed in log-deviations from their steady-state values. For convenience, model parameters and variables are summarized in Tables 1 and 2, respectively.

Prices are governed by Calvo (1983) contracts, augmented by indexation to the previous period's inflation for those firms not allowed to reoptimize their pricing decision. The implied inflation dynamics are given by the following Phillips curve:

$$(1) \ E_{t-1} \left[\begin{array}{l} \left(\frac{\beta \xi_p}{(1-\xi_p)(1-\beta \xi_p)} \hat{\pi}_{t+1} - \frac{\xi_p(1+\beta)}{(1-\xi_p)(1-\beta \xi_p)} \hat{\pi}_t \right) \\ + \frac{\xi_p}{(1-\xi_p)(1-\beta \xi_p)} \hat{\pi}_{t-1} \\ + [\alpha \hat{r}_t^k + (1-\alpha)(\hat{R}_t + \hat{w}_t)] \end{array} \right] = 0.$$

Here, hats on variables indicate the log-deviations from steady-state values. For the nominal interest rate and inflation terms that appear in the model, the hatted variables are effectively the demeaned net inflation and interest rates, because the log-deviations are computed using gross rates.

Nominal wages are staggered along similar lines to prices,⁴ with a clause for indexation to

the preceding period's price inflation. This produces the nominal wage equation:

$$(2) \ E_{t-1} \left(\begin{array}{l} \varpi_1 \hat{w}_{t+1} + \varpi_2 \hat{w}_t + \varpi_3 \hat{w}_{t-1} + \varpi_4 \hat{\pi}_{t+1} \\ + \varpi_5 \hat{\pi}_t + \varpi_6 \hat{\pi}_{t-1} + \hat{U}_{C,t} - \hat{l}_t \end{array} \right) = 0,$$

where

$$\begin{aligned} \varpi_1 &= \frac{-\beta \xi_w}{(1-\beta \xi_w)(1-\xi_w)} \left(1 + \frac{\lambda_w}{\lambda_w - 1} \right), \\ \varpi_2 &= \frac{1 + \beta \xi_w^2}{(1-\beta \xi_w)(1-\xi_w)} \left(1 + \frac{\lambda_w}{\lambda_w - 1} \right) - \frac{\lambda_w}{\lambda_w - 1}, \\ \varpi_3 &= \frac{-\xi_w}{(1-\beta \xi_w)(1-\xi_w)} \left(1 + \frac{\lambda_w}{\lambda_w - 1} \right), \\ \varpi_4 &= \varpi_1, \quad \varpi_6 = \varpi_3, \\ \varpi_5 &= \frac{\xi_w(1+\beta)}{(1-\beta \xi_w)(1-\xi_w)} \left(1 + \frac{\lambda_w}{\lambda_w - 1} \right). \end{aligned}$$

Firms' optimality conditions imply that their total payments for capital services equal their total cost of hiring labor each period:

$$(3) \quad \hat{r}_t^k + \hat{k}_t = \hat{w}_t + \hat{R}_t + \hat{l}_t.$$

Underlying this condition is the assumption that firms finance their wage bill with funds borrowed one period earlier. Real unit labor costs are therefore (in log terms) equal to the sum of the real wage and the short-term nominal interest rate.

⁴ See Erceg, Henderson, and Levin (2000) for the development of this form of staggered wage contracts.

The typical household's intertemporal Euler equation for consumption and first-order condition for capital purchases are, respectively,

$$(4) \quad \hat{U}_{C,t} = E_t \left(\hat{R}_{t+1} - \hat{\pi}_{t+1} + \hat{U}_{C,t+1} \right)$$

$$(5) \quad E_{t-1} \left(\hat{p}_{K',t} + \hat{U}_{C,t} \right) = E_{t-1} \left(\hat{U}_{C,t+1} + \frac{r^k}{r^k + 1 - \delta} \hat{k}_{t+1}^k + \frac{1 - \delta}{r^k + 1 - \delta} \hat{p}_{K',t+1} \right).$$

Because of habit formation in preferences, the household marginal utility of consumption that appears in the above expressions is not a static function of consumption. Instead, it depends on the current, prior, and expected future levels of consumption:

$$(6) \quad (1 - b\beta)(1 - b)E_{t-1}\hat{U}_{C,t} - b\beta E_{t-1}\hat{c}_{t+1} + (1 + b^2\beta)E_{t-1}\hat{c}_t - b\hat{c}_{t-1} = 0.$$

The economy's technology allows additional productive services to be generated, at a cost, from an unchanged stock of physical capital. The degree of capital utilization—that is, the difference between the physical capital stock (denoted by an overbar) and capital services—is chosen by households to equate marginal cost with marginal benefit:

$$(7) \quad E_{t-1} \left(\hat{k}_t - \hat{\bar{k}}_t \right) = \frac{1}{\sigma_a} E_{t-1} \hat{t}_t^k,$$

where $1/\sigma_a$ is the elasticity of the utilization function.

The equilibrium condition for household investment choices can be written as

$$(8) \quad E_{t-1} \hat{p}_{K',t} = \kappa E_{t-1} (\hat{i}_t - \hat{i}_{t-1}) - \beta \kappa E_{t-1} (\hat{i}_{t+1} - \hat{i}_t).$$

This condition indicates that the price firms pay for capital services is a function of two parameters that emerge from the behavior of households (who are the producers and suppliers of capital services): the households' discount factor, β , and the elasticity of their investment adjustment cost function, $1/\kappa$.

The stock of physical capital obeys the law of motion:

$$(9) \quad \hat{\bar{k}}_{t+1} = (1 - \delta) \hat{\bar{k}}_t + \delta \hat{i}_t,$$

where δ denotes the depreciation rate. Though physical investment is subject to adjustment costs, equation (9) indicates that a unit of investment adds to the physical capital stock in a standard manner.

Households' money demand function is given by

$$(10) \quad \hat{q}_t = -\frac{1}{\sigma_q} \left[\frac{R}{R-1} \hat{R}_t + \hat{U}_{C,t} \right],$$

a condition that indicates the standard choice between holding money for the transaction services it provides⁵ or, instead, holding one-period securities for interest income.

The following identity relates growth of nominal money supply to inflation and changes in real money supply:

$$(11) \quad \hat{\mu}_{t-1} = \hat{m}_t - \hat{m}_{t-1} + \hat{\pi}_t.$$

The aggregate demand for money in the economy comes from two sources: demand by firms (to finance their wage bill) and by households as given by condition (10). In equilibrium, total money demand is equal to the aggregate money stock:

$$(12) \quad q\hat{q}_t + wl(\hat{w}_t + \hat{l}_t) = \mu m(\hat{\mu}_t + \hat{m}_t).$$

The resource constraint and the aggregate production function can be written as

$$(13) \quad \left(\frac{1}{\beta} + \delta - 1 \right) \frac{K}{Y} (\hat{k}_t - \hat{\bar{k}}_t) + \frac{C}{Y} \hat{c}_t + \delta \frac{K}{Y} \hat{l}_t = \hat{y}_t$$

$$(14) \quad \hat{y} = \alpha \lambda_f \hat{k}_t + (1 - \alpha) \lambda_f \hat{l}_t.$$

Equation (13) indicates that resources this period can be consumed, invested, or used to generate additional capital utilization. Equation (14) indicates that the two inputs in production are labor and capital services.

Monetary policy follows a dynamic version of the Taylor rule:

⁵ Because of habit formation, prior and expected future transactions create a demand for real balances—i.e., money over and above the demand generated by current transactions.

$$(15) \quad \hat{R}_t = \rho \hat{R}_{t-1} + (1 - \rho) \left[r_\pi \hat{\pi}_t + r_y \hat{y}_t \right] \\ + r_{\Delta\pi} (\hat{\pi}_t - \hat{\pi}_{t-1}) + r_{\Delta y} (\hat{y}_t - \hat{y}_{t-1}) - \varepsilon_t.$$

The short-term nominal interest rate is therefore a smoothed function of inflation, output, and changes in inflation and output. There is also a monetary policy shock, ε_t . This rule is similar to that in Smets and Wouters (2003) and Levin et al. (2005).

Other than our use of an interest rate rule, the model we use corresponds to the CEE benchmark. A limitation in our application to the U.K. data is that the CEE model describes a closed economy. But there are several reasons for using a closed-economy model when analyzing the United Kingdom; see Neiss and Nelson (2003) for a discussion. For the present paper, the main reasons why a closed economy of the DSGE model might be suitable for the United Kingdom are as follows: (i) Openness makes it difficult to model capital formation endogenously, whereas the presence of endogenous capital is a key feature of the CEE model. And (ii) the simplest open-economy models give counterfactual weight to the exchange rate in consumer price inflation dynamics; once the exchange-rate channel is “tamed” by such approaches as assuming incomplete pass-through, imported intermediates, etc., the model’s properties become more like those of a closed economy (see, e.g., Obstfeld, 2002).

ESTIMATION

To estimate the model, we first obtain data responses to a monetary policy shock from a vector autoregression (VAR) for the United Kingdom. Then, as in CEE, we match these impulse responses as closely as possible with the CEE model, using a minimum-distance estimation procedure.⁶ Our analysis here is limited to monetary policy shocks, but there is evidence for the United States that estimates of the CEE model are robust to incorporating technology shocks into the analysis (see DiCecio, 2005, and Altig et al., 2005).

⁶ This procedure was also used with a smaller VAR by Rotemberg and Woodford (1997).

VAR Estimates

We estimate our VAR on a U.K. dataset consisting of a subset of the variables studied in the U.S. case by CEE. Our VAR contains the logs of real gross domestic product (GDP), real consumption,⁷ real investment, and labor productivity, as well as the nominal Treasury bill rate and the quarterly (retail) inflation rate.⁸ As these choices imply, our focus is on the response of the policy rate, inflation, and aggregate demand to a monetary policy shock, as well as the split of aggregate demand among its components and the division of the output response between labor and other inputs.

The sample period is 1979:Q2–2005:Q4. The start date is the quarter corresponding to the period (May 1979) in which the Thatcher government first took office—and so an important monetary policy regime change.⁹ It also corresponds approximately to the date of some other significant changes in government policy that are important for the VAR responses, as we discuss in the next section.

Figure 1 plots the estimated VAR responses to a monetary policy shock and their bootstrapped confidence intervals, along with the match to each response made by our estimates of the CEE model; the model-based responses are the blue lines. Parameters fixed in estimation are given in Table 3.

Parameter Estimates

The parameter estimates resulting from this matching of impulse responses are given in Table 4. Standard errors for the parameter estimates appear in parentheses and are calculated by the asymptotic delta method.¹⁰

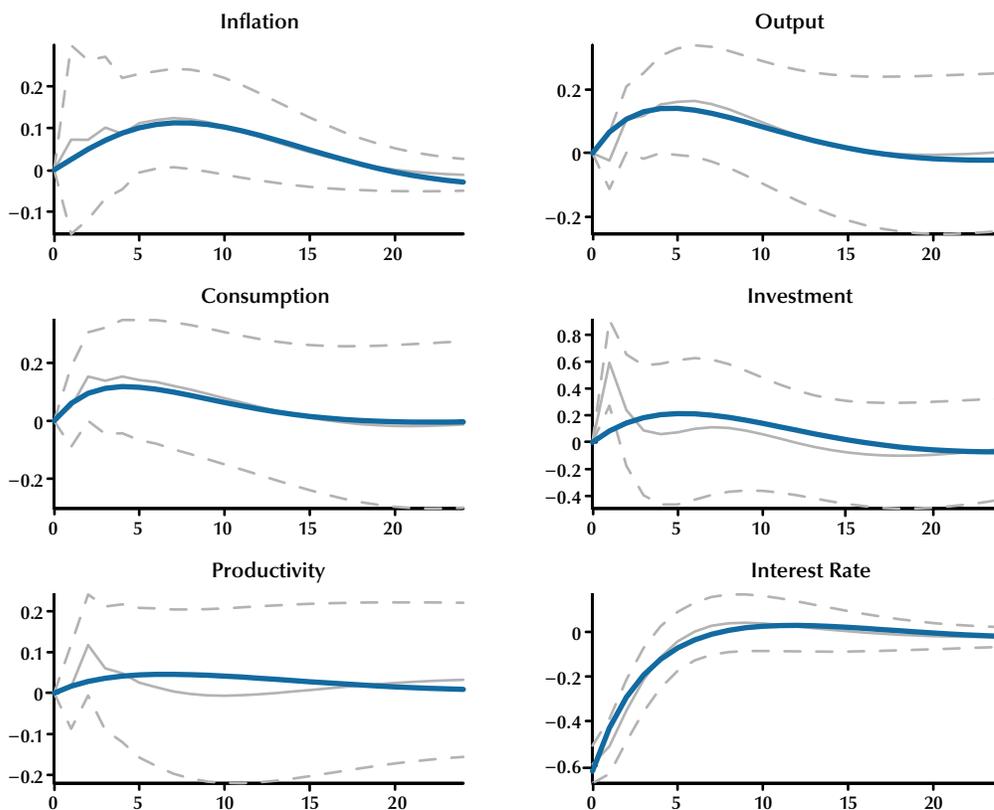
⁷ We have not split consumption between durables and non-durables. VAR impulse responses in Aoki, Proudman, and Vlieghe (2002, Chart 2), using a different VAR specification and sample period from ours, found similar response functions for the two types of consumption.

⁸ Our VAR does not include a time trend. Impulse responses look similar regardless of whether a linear trend is included in the VAR.

⁹ See, e.g., Goodhart (1989) for a perspective on this regime change.

¹⁰ See Newey and McFadden (1994, p. 2145).

Figure 1
Estimated and Model Responses to a Monetary Policy Shock



The parameter indexing habit formation in consumption is larger than that estimated by CEE and Altig et al. (2005) but is basically in line with Fuhrer (2000). So the degree of habit formation in the United Kingdom appears similar to U.S. estimates.

The markup estimate is, at somewhat above 2, high by the standards of calibrated and estimated DSGE models. It is, however, roughly in line with the estimate of the average U.K. gross markup (in manufacturing) by Haskel, Martin, and Small (1995, p. 30) of 2.0. Our high markup estimate appears more standard if it is regarded as the wedge between consumer prices and (principally) nominal wages,¹¹ including the impact of cost

elements we have not modeled explicitly.¹² It should be remembered that the model is an abstraction of a model with imported intermediate goods and indirect taxes. With these unmodeled elements built into the empirical price-level series, the estimated markup of retail prices on nominal wages is increased.¹³

The estimated interest rate policy rule has responses to both the level and growth rate of inflation as well as to the deviation of output from the steady state. Because the interpretation

¹¹ Interest on the nominal wage bill also enters the cost expression, with implications we discuss shortly.

¹² By contrast, Haskel, Martin, and Small's (1995) markup estimate allows for costs of materials, so our markup estimate should be higher than theirs, other things equal.

¹³ Therefore, our high estimate may be consistent with Britton, Larsen, and Small (2000) setting the U.K. steady-state markup value closer to 1.0 when calibrating a model with explicit imported intermediates.

Table 3**Parameters Fixed in Estimation**

Parameter	Value
β	$(1.04)^{-1/4}$
α	0.36
δ	0.025
μ	1.017
ψ_l	Such that $l = 1$
ψ_q	Such that $q/m = 0.25$
λ_w	1.05

of the inflation responses is affected by the output response, we deal with the latter response first. As technology shocks are held constant, any output movements reflect the opening of the output gap and so also inflationary pressure. It is precisely this type of output variation that a monetary authority will have greatest interest in stabilizing. This may account for the output response being larger than is usual in estimated interest rate rules, which typically do not remove from the output measure the variation that is due to technology shocks.

The monetary policy reaction to inflation consists of a standard level response and a negative response to the change in inflation. We find that the inflation-change response, although not very precisely estimated, is negative and economically sizable. Under some parameter values, an estimated negative response to the change in inflation implies that policymakers have lagged inflation rather than current inflation in their rule. Our estimated $r_{\Delta\pi}$ response is, however, too large (in absolute value) for this to be the case. Instead, policymakers actually make different-signed short-run responses to inflation, initially allowing a temporary reduction in the interest rate when inflation rises. To understand this response, one has to keep in mind the supply side of the CEE model. In the standard sticky-price model, the impulse responses of output and the output gap to a monetary policy shock are identical, because potential output depends on real shocks only. In the CEE model, however, this is not the case, because interest rates enter the production func-

Table 4**Baseline Model Estimates, 1979:Q2–2005:Q4**

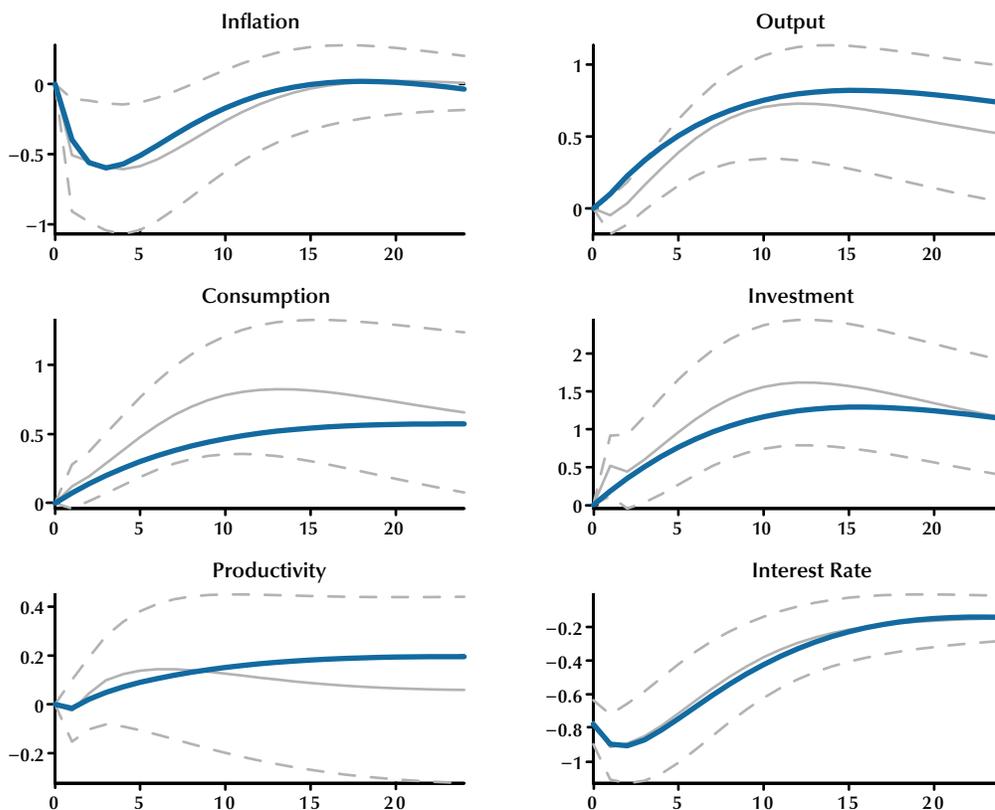
Private sector parameters	
b	0.7739 (0.0013)
λ_f	2.2681 (0.0323)
ξ_p	0.9371 (0.0004)
ξ_w	0.0000 (0.0083)
κ	16.4333 (0.1750)
σ_a	∞ (—)
Monetary policy rule parameters	
ρ	0.8720 (0.0907)
r_π	1.2813 (0.4977)
r_γ	0.3517 (0.6065)
$r_{\Delta\pi}$	-0.5129 (0.6971)
$r_{\Delta y}$	0.4259 (0.2759)
σ_R	0.1564 (0.0001)

NOTE: The number of impulse-response steps used, which ideally should be determined statistically (see Hall et al., 2007), is 25. A diagonal matrix is used to weight the responses.

tion, implying that potential output depends on the nominal interest rate (see Ravenna and Walsh, 2006, for further discussion). Holding constant its other effects, a cut in the interest rate stimulates potential output and helps inflation stabilization in the face of upward pressure on aggregate demand. Therefore, in the wake of a monetary policy shock, policymaker stabilization of the output gap and inflation takes a three-pronged approach: a large response to output to rein in incipient excessive aggregate demand ($r_\gamma > 0$, $r_{\Delta y} > 0$); a short-run cut in the interest rate as inflation rises to stimulate potential output ($r_{\Delta\pi} < 0$); and a sizable and durable

Figure 2

Estimated and Model Responses to a Monetary Policy Shock, 1962:Q3–2005:Q4



positive response of the interest rate to the level of inflation relative to the target ($r_\pi > 1$).

The estimates imply large investment adjustment costs, mainly driven by the matching of the smoother investment responses after the initial period; the model does not match the apparent initial spike in investment observed in the data.

Although the model allows for both wage stickiness and indexation of wages to price inflation, our parameter estimates imply that both these features are absent.¹⁴ On the other hand, price stickiness is substantial. Because full indexation of prices is superimposed on this price adjustment, it is not appropriate to infer from

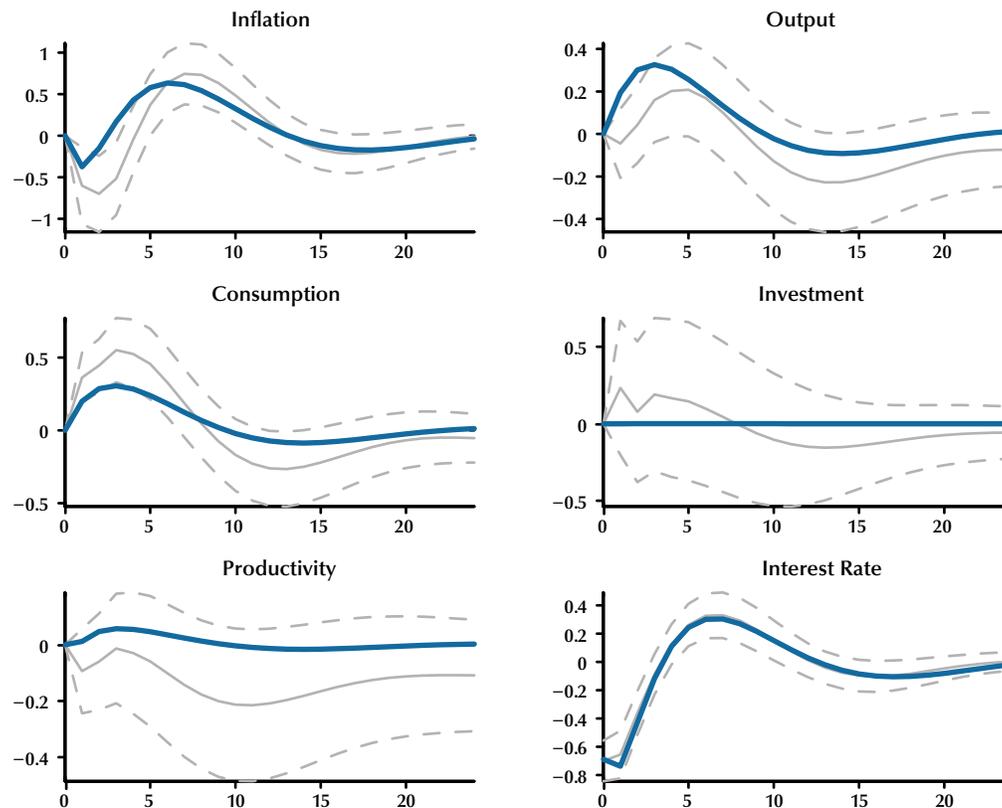
the low implied Calvo probability that prices are implausibly rigid; rather, the indexation implies substantial price movements every period, even with the underlying price stickiness. Empirical support for lagged inflation terms in the Phillips curve, when this parameter is estimated freely, is not universal (see, e.g., Ireland, 2001), so our assumption of full indexation may be restrictive. A lagged inflation term in the Phillips curve is, however, in line with the specification advocated by Blake and Westaway (1996) for U.K. monetary policy analysis.

The bottom line is that the estimates suggest that an emphasis on price stickiness as opposed to wage stickiness is appropriate in analyzing the United Kingdom. This emphasis is consistent with evidence for other European countries, such

¹⁴ That is, equation (2) collapses to the usual static labor-supply condition, equating real wages to the marginal rate of substitution between consumption and leisure.

Figure 3

Estimated and Model Responses to a Monetary Policy Shock, 1962:Q3–1979:Q1



as Coenen, Levin, and Christoffel's (2007) study of nominal rigidities in Germany.

Factor utilization is not found to be important, the relevant parameter being driven to the boundary of its admissible region. Our VAR productivity responses are not very precisely estimated. Consequently, the model can explain output variation in terms of input responses and therefore has little need to rely on the intensive margin to explain the data responses.

ESTIMATES INCLUDING PRE-1979 DATA

In this section we present results for the long sample 1962-2005 as well as a sample using only pre-1979 data. The long-sample impulse responses

and their matches are given in Figure 2, and those for the pre-1979 sample are given in Figure 3.

Parameter estimates for each sample are given in Table 5.

Turning to the policy rule first, the estimates deliver substantially lower inflation responses in the interest rate rule pre-1979, consistent with the assignment of inflation control to nonmonetary devices in the United Kingdom before 1979. But the response is large enough even in this sample period to deliver determinacy (i.e., a single-model equilibrium).¹⁵ The output response is “wrongly” (negatively) signed pre-1979. This may

¹⁵ Our estimation routine considers only parameter combinations that deliver a single solution. An alternative procedure, which we have not pursued here, would be to consider both determinacy and indeterminacy regions and select a solution in the latter case using the minimal state-variable procedure.

Table 5
Model Estimates for Samples that Include Pre-1970 Data

	1962:Q3–2005:Q4	1962:Q3–1979:Q1
Private sector parameters		
b	0.9410 (0.0003)	0.5806 (0.0015)
λ_f	1.3904 (0.0038)	1.2039 (0.0035)
ξ_p	0.4478 (0.0024)	0.1703 (0.0025)
ξ_w	0.9897 (0.0003)	0.6069 (0.0008)
κ	49.9261 (0.3288)	∞ (—)
σ_a	0.3900 (0.0062)	0 (—)
Monetary policy rule parameters		
ρ	0.9433 (0.0068)	0.7018 (0.0017)
r_π	1.2657 (0.1118)	0.9642 (0.0473)
r_γ	0.0321 (0.0125)	–0.1980 (0.0220)
$r_{\Delta\pi}$	0.4128 (0.0034)	0.6066 (0.0163)
$r_{\Delta\gamma}$	0.0879 (0.0104)	0.1618 (0.0182)
σ_R	0.1959 (0.0002)	0.1719 (0.0002)

be another reflection of the lack of monetary policy response to inflationary pressure because, as noted earlier, the output coefficient captures responses to the specific type of output increases that are likely to raise inflation. An additional departure from our baseline estimates is that both sets of estimates that include pre-1979 data have a more standard (i.e., positive) interest rate response to the change in inflation.

A major feature of the structural parameters when we move away from our baseline sample is that there is now sizable nominal wage rigidity. Another difference from our baseline structural parameter estimates pertains to investment behavior. In the pre-1979 sample, the model cannot

match the empirical investment impulse response; the best the model can do is to suppress investment altogether (and so generate a flat investment model response in Figure 3). Accordingly, the investment adjustment-cost parameter estimate becomes arbitrarily large.

It is tempting to suggest that our anomalous results for investment occur because the estimates including pre-1979 data are distorted by the existence of unmodeled breaks in monetary policy regime. But this does not really provide a satisfactory answer why we get these particular results. It is not obvious that estimated impulse responses over a sample that includes multiple regimes will be perverse in their shape; they are, more or less,

an average of the responses observed across each regime, we should expect them to be of standard shape. Instead of this, we get model estimates that appear to extinguish the investment portion of aggregate demand.

It is likely instead that U.K. government policy is indeed the culprit for the anomalies in the pre-1979 results, but the policy actions responsible were microeconomic interventions in the economy and not monetary policy. Before the 1980s, many large industries (e.g., steel and telecommunications) were principally government-owned. What is more, in a misguided effort to control inflation by nonmonetary means, governments frequently intervened in the pricing decisions of their enterprises. For example, George Brown, then Secretary of State for Economic Affairs, in 1965 said that the government was operating a price-control policy “in the field of government responsibility so far as charges for which they are responsible, prices which are their responsibility...” (*Glasgow Herald*, 1965). Ted Heath, prime minister 1970-74, said shortly before being elected that “we are going to see to it that the State does not put up its prices and charges with gay abandon” (Russell, 1970). The attempt to enforce this policy led to considerable interference in government enterprises’ operations, so much so that Anthony Crosland, a leading Labour Party figure, cited 1970-71 as a period that displayed a poor “balance...between Ministerial control and entrepreneurial freedom” (Crosland, 1974, p. 39).

From around 1978, however, it became much more standard for government-owned enterprises to base their pricing and investment decisions on market signals, with a government report on the subject in 1978 stating that the “Government intends that the nationalized industries will not be forced into deficits by restraints on their prices” (House of Commons, 1979). The stepping away by government from management of investment decisions was cemented by the privatization of many government enterprises in the 1980s.

Because the pre-1978 government interventions blocked investment from responding to

market signals, including those from monetary policy shocks, one can understand why investment responses might deviate greatly from those predicted by our model, in which investment behavior is based on optimal firm choices. Government prohibitions on a firm’s ability to raise prices might cut off funds to the firm, thus distorting investment decisions. On the other hand, for given monetary policy, government intervention in investment decisions might merely transfer aggregate demand pressure from investment to other categories of spending, rather than affect total demand. So impulse responses other than those for investment might still be compatible with the model, which is essentially what we find.

EXTENSIONS

In this section, we report further results regarding the robustness of our results to alternative data definitions (under “Alternative Investment Series”) and our choice of regime dates (under “VAR Stability and Regime Breaks”).

Alternative Investment Series

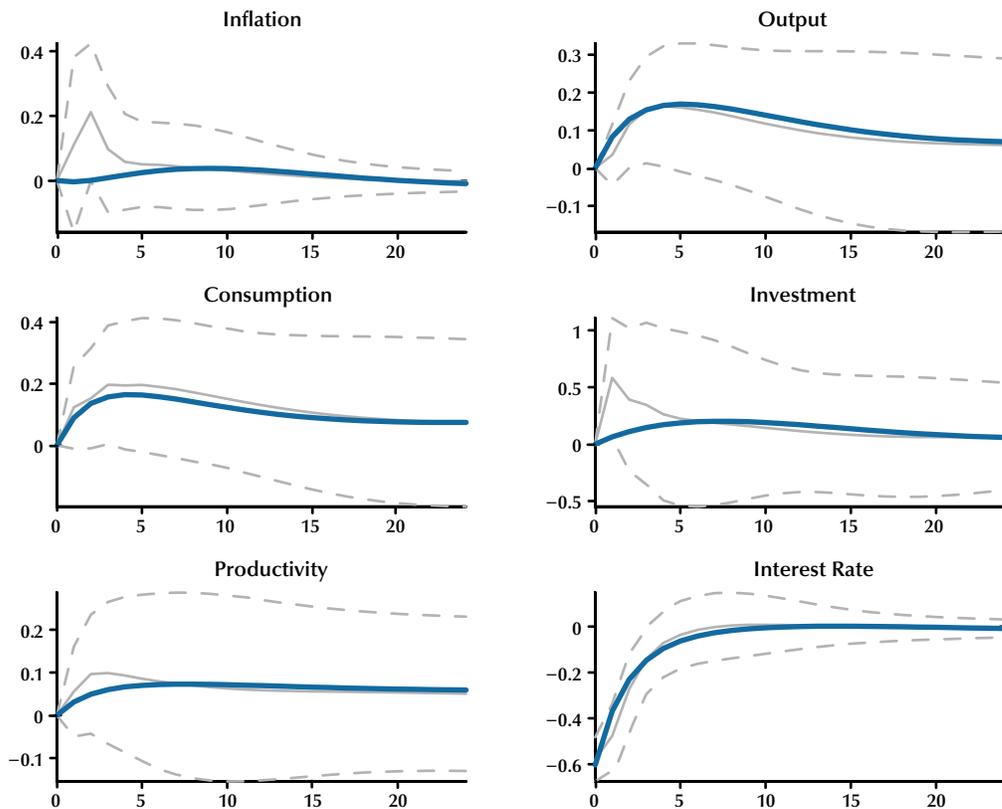
The BEQM model and such sources as the Bank of England *Inflation Report* use a slightly narrower definition of investment than we use in our analysis. This narrower definition is known as “business investment” (though, like our series, it includes investment by government enterprises).

We use this alternative investment series and examine the effect on our results. In Figure 4, we show that using this series in our VAR makes little difference in the empirical responses to a monetary policy shock.¹⁶ The parameter estimates using this series are given in Table 6. These are little changed from the baseline parameter estimates, with some important exceptions. Most notably, the policy rule estimates are much less precisely estimated and feature a much lower

¹⁶ The sample period we now use, 1980:Q1–2005:Q4, is slightly shorter than the 1979:Q2–2005:Q4 sample we used for our baseline results, owing to difficulty obtaining parameter estimates for 1979–2005 with the business investment series.

Figure 4

Estimated and Model Responses to a Monetary Policy Shock, with Alternative Investment Series, 1980:Q1–2005:Q4



estimated response to output; and the alternative estimates give roughly equal importance to price and wage stickiness, whereas in the baseline estimates prices were much stickier than wages.

VAR Stability and Regime Breaks

As noted above, our data cover the whole 1962-2005 period, but our baseline structural estimates are based on a sample covering only the period since the late 1970s, reflecting the changes in U.K. industrial and monetary policies that took place around that time. To investigate further the issue of regime break dates, in Table 7 we follow Boivin and Giannoni (2002, p. 99) by investigating the stability of the VAR when it is estimated for the long sample 1962:Q1–2005:Q4.

We report the *p*-values for the constancy of the coefficients associated with each group of regressors in the VAR. The break dates suggested by the test are also reported, and those highlighted in bold achieve statistical significance at the 10 percent level or better.

The results for the baseline VAR specification—that is, the VAR specification underlying the Figure 2 impulse responses—occupy the top half of the table. The results suggest a significant break in the inflation equation around 1975:Q2. This date, however, does not constitute a monetary policy regime break; instead, the mid-1975 instability reflects a one-time shock to industrial policy. The previously described U.K. government policy of holding down nationalized industries’

prices underwent an adjustment in this period, with prices allowed to rise to eliminate the accumulated discrepancy with costs. The large effect on consumer prices that resulted has sometimes been categorized as tantamount to a substantial increase in indirect taxes (Wilson, 1984, p. 50). The test statistics mechanically take this large shock as evidence of regime change, though, economically, it does not amount to the sort of change in systematic policy responses that truly qualify as a policy regime shift. The remaining stability rejections are spread over 1977-81 and so are roughly in line with our assumption of a 1979 break date.

Recall that our VAR does not use detrended real data, nor does it include a trend as a regressor. Some work on U.S. data—e.g., Rotemberg and Woodford (1997), Boivin and Giannoni (2002), and Giannoni and Woodford (2005)—detrends real variables before putting them in the VAR. We report in the bottom half of Table 7 the stability results for our VAR when our specification is modified by replacing the four real variables with their detrended counterparts. The detrending assumes that the log real variables are driven by a broken linear trend, with constant and trend breaks in both 1973:Q4 and 1981:Q4.

Besides continuing to show a break in 1975 in some of the inflation coefficients, these stability results largely reaffirm a focus on a regime break around the early 1980s (specifically, 1980 or 1981). Two of the significant rejections of stability do suggest a break in 1984 in GDP and productivity behavior, but these rejections can be discounted as reflecting the temporary disturbances to output from the coal-mining strike of that year.

One puzzling aspect of the results with detrended variables is that the interest rate equation no longer registers any significant regime break. This, however, is not decisive evidence against the importance of monetary policy regime change. For one thing, relatively minor and statistically insignificant changes in the VAR coefficients can imply large changes in the implied “long-run response” of the interest rate to endogenous variables. This is the case here because, despite the lack of rejection of stability, the VAR

Table 6

Estimates Using Alternative Investment Series, 1980:Q1–2005:Q4

Private sector parameters	
b	0.6621 (0.0015)
λ_f	2.4922 (0.0198)
ξ_p	0.8465 (0.0020)
ξ_w	0.9409 (0.0008)
κ	14.2684 (0.2340)
σ_a	∞ (—)
Monetary policy rule parameters	
ρ	0.8426 (1.2144)
r_π	2.0747 (21.0751)
r_γ	0.0546 (1.0372)
$r_{\Delta\pi}$	-0.2935 (0.0368)
$r_{\Delta y}$	0.4140 (2.1327)
σ_R	0.1509 (0.0001)

equation for the interest rate underlying the final row of the table has a long-run solution with an interest rate response to (annualized) inflation of about 0.3; but this response rises to 1.0 on restricting the sample to 1979-2005.¹⁷ Furthermore, the inflation VAR equation now exhibits a significant early-1980s break, which is indirect support for a monetary policy regime change around that time.

¹⁷ In some contexts (see, e.g., Rotemberg and Woodford, 1997, and Rudebusch, 1998) the VAR equation for the interest rate coincides with the interest rate policy rule. This is not the case in our analysis, as the policy rule that we use in estimation differs from the VAR equation. But solving the reduced-form VAR interest rate equation for its long-run solution nevertheless provides a means of cross-checking the stability test results.

Table 7
VAR Stability Tests, 1962:Q3–2005:Q4

Dependent variable	Regressor					
	π	y	c	i	$y - h$	r
1. Baseline VAR specification						
π	0.002 1975:Q2	0.001 1977:Q1	0.713 1975:Q3	0.153 1971:Q1	0.510 1970:Q4	0.574 1973:Q2
y	0.127 1972:Q3	0.682 1975:Q3	0.373 1990:Q2	0.358 1976:Q1	0.948 1975:Q3	0.347 1977:Q4
c	0.469 1990:Q2	0.444 1990:Q2	0.001 1977:Q1	0.012 1980:Q2	0.001 1977:Q1	0.008 1980:Q2
i	0.553 1973:Q1	0.844 1968:Q4	0.788 1976:Q1	0.195 1979:Q2	0.524 1975:Q3	0.711 1975:Q3
$y - h$	0.382 1973:Q2	0.943 1975:Q3	0.922 1973:Q2	0.869 1973:Q2	0.361 1990:Q2	0.199 1974:Q1
r	0.275 1977:Q4	0.000 1981:Q2	0.000 1981:Q2	0.006 1977:Q4	0.000 1981:Q2	0.471 1973:Q1
2. Baseline VAR specification with $y, c, i, y - h$ detrended						
π	0.002 1975:Q2	0.072 1975:Q2	0.040 1981:Q4	0.004 1976:Q1	0.567 1978:Q4	0.741 1988:Q4
y	0.110 1972:Q3	0.036 1984:Q1	0.121 1981:Q2	0.224 1987:Q2	0.296 1981:Q4	0.867 1977:Q4
c	0.738 1970:Q1	0.063 1981:Q3	0.065 1980:Q2	0.095 1980:Q2	0.105 1975:Q2	0.001 1980:Q2
i	0.812 1968:Q4	0.269 1987:Q1	0.239 1985:Q1	0.014 1976:Q4	0.049 1980:Q1	0.662 1990:Q2
$y - h$	0.325 1972:Q3	0.147 1984:Q1	0.078 1984:Q1	0.187 1976:Q4	0.476 1981:Q4	0.665 1970:Q1
r	0.425 1977:Q4	0.813 1980:Q1	0.714 1986:Q1	0.801 1990:Q3	0.865 1998:Q3	0.460 1976:Q3

NOTE: Values reported are the p -values for the Andrews (1993) *sup*-Wald test, computed using the procedure of Diebold and Chen (1996). The null hypothesis assumes no structural breaks, whereas the alternative hypothesis has breaks in the constant and group of lag coefficients on the indicated regressor. Each panel also gives the break-date associated with the p -value.

CONCLUSIONS

In this paper, we have estimated the Christiano, Eichenbaum, and Evans (2005) model on U.K. data. Although CEE found plausible estimates on U.S. data when treating the period since the 1960s as a single regime, for the United Kingdom it appears that more satisfactory estimates emerge if pre-1979 data are excluded; otherwise, the estimates imply degenerate behavior of

investment. This result is consistent with policy regime changes being an important factor in the postwar U.K. economy. These regime changes include not only changes in the role assigned to monetary policy but also shifts toward making investment decisions more closely related to market forces. Another important implication of our estimates is that price stickiness, rather than wage stickiness, is the major source of nominal rigidity in the United Kingdom.

REFERENCES

- Altig, David; Christiano, Lawrence J.; Eichenbaum, Martin and Linde, Jesper. "Firm-Specific Capital, Nominal Rigidities and the Business Cycle." NBER Working Paper No. 11034, National Bureau of Economic Research, 2005.
- Andrews, Donald W.K. "Tests for Parameter Instability and Structural Change with Unknown Change Point." *Econometrica*, July 1993, 61(4), pp. 821-56.
- Aoki, Kosuke; Proudman, James and Vlieghe, Gertjan. "House Prices, Consumption and Monetary Policy: A Financial Accelerator Approach." Working Paper No. 169, Bank of England, 2002.
- Benati, Luca. "Evolving Post-World War II U.K. Economic Performance." *Journal of Money, Credit, and Banking*, August 2004, 36(4), pp. 691-717.
- Blake, Andrew P. and Westaway, Peter F. "Credibility and the Effectiveness of Inflation Targeting Regimes." *Manchester School of Economic and Social Studies*, June 1996, 64(Suppl 1), pp. 28-50.
- Boivin, Jean and Giannoni, Marc P. "Assessing Changes in the Monetary Transmission Mechanism: A VAR Approach." Federal Reserve Bank of New York *Economic Policy Review*, May 2002, 8(1), pp. 97-111.
- Britton, Erik; Larsen, Jens D.J. and Small, Ian. "Imperfect Competition and the Dynamics of Mark-ups." Working Paper No. 110, Bank of England, 2000.
- Calvo, Guillermo A. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, September 1983, 12(3), pp. 383-98.
- Castelnuovo, Efram and Surico, Paolo. "The Price Puzzle: Fact or Artefact?" Working Paper No. 288, Bank of England, 2006.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L. "Monetary Policy Shocks: What Have We Learned and to What End?" in J.B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*. Volume 1A. Amsterdam: Elsevier/North Holland, 1999, pp. 65-148.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, February 2005, 113(1), pp. 1-45.
- Coenen, Günter; Levin, Andrew T. and Christoffel, Kai. "Identifying the Influences of Nominal and Real Rigidities in Aggregate Price-Setting Behavior." *Journal of Monetary Economics*, 2007 (forthcoming).
- Crosland, Anthony. *Socialism Now and Other Essays*. London: Jonathan Cape, 1974.
- Del Negro, Marco; Schorfheide, Frank; Smets, Frank and Wouters, Rafael. "On the Fit and Forecasting Performance of New-Keynesian Models." ECB Working Paper No. 491, European Central Bank, 2005.
- DiCecio, Riccardo. "Comovement: It's Not a Puzzle." Working Paper 2005-035B, Federal Reserve Bank of St. Louis, 2005.
- Diebold, Francis X. and Chen, Celia. "Testing Structural Stability with Endogenous Breakpoint: A Size Comparison of Analytic and Bootstrap Procedures." *Journal of Econometrics*, June 1996, 70(1), pp. 221-41.
- Erceg, Christopher J.; Henderson, Dale W. and Levin, Andrew T. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics*, October 2000, 46(2), pp. 281-313.
- Fuhrer, Jeffrey C. "Habit Formation in Consumption and Its Implications for Monetary Policy Models." *American Economic Review*, June 2000, 90(3), pp. 367-90.
- Giannoni, Marc P. and Woodford, Michael. "Optimal Inflation Targeting Rules," in B. Bernanke and M. Woodford, eds., *The Inflation-Targeting Debate*. Chicago: University of Chicago Press, 2005, pp. 93-162.
- Glasgow Herald*. "Brown's Pay Policy Approved." October 1, 1965, p. 5.

DiCecio and Nelson

- Goodhart, Charles A.E. "The Conduct of Monetary Policy." *Economic Journal*, June 1989, 99(396), pp. 293-346.
- Hall, Alastair; Inoue, Atsushi; Nason, James M. and Rossi, Barbara. "Information Criteria for Impulse Response Function Matching Estimation of DSGE Models." Working Paper 2007-10, Federal Reserve Bank of Atlanta, May 2007.
- Harrison, Richard; Nikolov, Kalin; Quinn, Meghan; Ramsay, Gareth; Scott, Alasdair and Thomas, Ryland. *The Bank of England Quarterly Model*. London: Bank of England, 2005.
- Haskel, Jonathan; Martin, Christopher and Small, Ian. "Price, Marginal Cost and the Business Cycle." *Oxford Bulletin of Economics and Statistics*, February 1995, 57(1), pp. 25-41.
- House of Commons, Parliament of the United Kingdom. *House of Commons Debates*, December 5, 1979, p. 563 (quoting U.K. Government, *White Papers on Nationalized Industries*, London: HMSO, 1978).
- Ireland, Peter N. "Sticky-Price Models of the Business Cycle: Specification and Stability." *Journal of Monetary Economics*, February 2001, 47(1), pp. 3-18.
- King, Mervyn A. *Public Policy and the Corporation*. London: Chapman and Hall, 1977.
- Levin, Andrew T.; Onatski, Alexei; Williams, John C. and Williams, Noah. "Monetary Policy under Uncertainty in Micro-Founded Macroeconomic Models." *NBER Macroeconomics Annual*, 2005, 20(1), pp. 229-87.
- Neiss, Katharine S. and Nelson, Edward. "The Real-Interest-Rate Gap as an Inflation Indicator." *Macroeconomic Dynamics*, April 2003, 7(2), pp. 239-62.
- Nelson, Edward and Nikolov, Kalin. "Monetary Policy and Stagflation in the U.K." *Journal of Money, Credit, and Banking*, June 2004, 36(3), pp. 293-318.
- Newey, Whitney K. and McFadden, Daniel. "Large Sample Estimation and Hypothesis Testing," in Robert Engle and Dan McFadden, eds., *Handbook of Econometrics*. Volume 4. Amsterdam: Elsevier/North Holland, 1994, pp. 2111-45.
- Obstfeld, Maurice. "Exchange Rates and Adjustment: Perspectives from the New Open-Economy Macroeconomics." Bank of Japan *Monetary and Economic Studies*, December 2002, 20(4), pp. 23-46.
- Pagan, Adrian. "Addendum to *Report on Modelling and Forecasting at the Bank of England*." Bank of England *Quarterly Bulletin*, Summer 2005, 45(2), pp. 190-93.
- Ravenna, Federico and Walsh, Carl E. "Optimal Monetary Policy with the Cost Channel." *Journal of Monetary Economics*, May 2006, 53(2), pp. 199-216.
- Ravn, Morten O. "Permanent and Transitory Shocks and the U.K. Business Cycle." *Journal of Applied Econometrics*, January/February 1997, 12(1), pp. 27-48.
- Rotemberg, Julio J. and Woodford, Michael. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." *NBER Macroeconomics Annual*, 1997, 12(1), pp. 297-346.
- Rudebusch, Glenn D. "Do Measures of Monetary Policy Shocks in a VAR Make Sense?" *International Economic Review*, November 1998, 39(4), pp. 907-31.
- Russell, William. "'Pushover' Jenkins Surrendered to Unions, Says Heath." *Glasgow Herald*, June 3, 1970, p. 8.
- Sims, Christopher A. and Zha, Tao. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, May 2006, 96(1), pp. 54-81.
- Smets, Frank and Wouters, Raf. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, 2003, 1(5), pp. 1123-75.
- Wilson, Thomas. "Monetarism in Britain," in T. Wilson, ed., *Inflation, Unemployment, and the Market*. Oxford: Clarendon Press, 1984, pp. 41-78.

APPENDIX

Data Sources and Definitions

Nominal interest rate: U.K. Treasury bill rate, quarterly average. Source: Haver-*IFS*, quarterly average series.

Output: real GDP, seasonally adjusted, quarterly, series abmi.q.

Source: U.K. Office of National Statistics (ons.gov.uk), downloaded May 2006.

Private household consumption: seasonally adjusted, quarterly, series abjr.q.

Source: ons.gov.uk, downloaded May 2006.

Investment: gross fixed capital formation, seasonally adjusted, quarterly, series npqt.q.

Source: ons.gov.uk, downloaded May 2006.

Alternative investment series: business investment at 2003 prices, seasonally adjusted, quarterly, series npel.q.

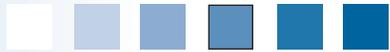
Source: ons.gov.uk, downloaded August 2006.

Productivity: Y/H , where H = hours worked.

Source for H : series ybus.q (source: U.K. Office of National Statistics [ONS]), with splice into Ravn (1997) U.K. hours worked series to obtain pre-1971:Q1 data.

Inflation: log difference of P , where P is a seasonally adjusted consumer price series. P was constructed as follows: A quarterly average of the retail price index (RPI) was spliced into a quarterly average of RPIX (RPI excluding mortgage interest payment) after 1973, the series was seasonally adjusted, then tax-related spikes of 4 percent (in 1979:Q3) and 2 percent (1990:Q2) were removed from the series. The seasonal regressions underlying the seasonal adjustment used the log-change as the dependent variable, and seasonal patterns were allowed to differ across 1955-76, 1976-86, and 1987-2005.

Source for the monthly RPI underlying the quarterly averages: ONS (ons.gov.uk). The ONS, however, provides RPIX data only from January 1987. An unofficial RPIX series starting in 1974 has, however, been constructed at the Bank of England, and this series underlies studies such as Nelson and Nikolov (2004) and Benati (2004). The OECD-Haver service also provides an RPIX series (though beginning only in 1975) that closely matches this series. We used the quarterly average of the unofficial RPIX series for 1974-87 and spliced it into the quarterly average of the official RPIX series that begins in 1987. Splicing this RPIX series at 1974:Q1 with RPI delivered the RPI/RPIX quarterly average series underpinning P .



Commentary

Martin Fukač and Adrian Pagan

The paper by DiCecio and Nelson (2007; DCN) considers the estimation of the parameters of a dynamic stochastic general equilibrium (DSGE) model for the United Kingdom that is virtually the same as that which Christiano, Eichenbaum, and Evans (2005; CEE) estimated for the United States. The CEE model is much larger than existing academic DSGE models of the United Kingdom, such as Lubik and Schorfheide (2007). It is not as large as the Bank of England Quarterly Model (BEQM), which has both DSGE elements and data-imposed dynamics; however, because the BEQM is used for policymaking, there is a much greater imperative to match the data than found in most academic work. There are a few other DSGE models that have been applied to the United Kingdom—for example, Leitemo (2006)—but, in general, these are often used to examine some particular question and are also rather restricted in their mode of operation. Often they use a standard open-economy New Keynesian model rather than a straight DSGE model like CEE's. Moreover, the authors of these models are often not that familiar with the U.K. context and data; the current authors, however, are experts in this area, and it certainly shows in their discussion of alternative data sources. So, given the paucity of studies, any new one would be welcome.

Now, as the Chinese proverb says, a journey of a thousand miles starts with a single step. What we have here is more than single step but well

short of a thousand miles. Reading it, one longed for a fully fleshed-out model along the lines of Smets and Wouters's (2003) work on the United States and the euro area (which is very similar to CEE's), where a complete set of shocks is described and estimated. Because the DCN model identifies only a money shock, there are few questions one can ask about the model. So it was disappointing that the authors were not a bit more adventurous. But we presume that this will be part of a broader piece of research and look forward to seeing a more complex model that recognizes the open-economy characteristics of the United Kingdom. Of course, one does have to acknowledge that DSGE models have not had a good record of producing useful models of the open economy. One reason DCN point to is the prediction of stronger exchange effects than seen in the data. We agree with this, and it was a central conclusion about the mini-BEQM model that was calibrated to the U.K. economy in Kapetanios, Pagan, and Scott (2007). Moreover, as Justiniano and Preston (2006) argue, it has been very hard to find much of an influence of the foreign economy on a small open economy, and this is contrary to evidence we have from vector autoregressions (VAR). So there is quite a bit to be done both on the broader front of developing useful open-economy models and in getting a U.K. model that is in a more complete state than this one. Because the model is not fleshed out that much, we will restrict comments to what DCN do rather than alternatives that might have been tried.

Martin Fukač is an economist at the Reserve Bank of New Zealand. Adrian Pagan is a professor of economics at Queensland University of Technology and the University of New South Wales. This research was supported by ARC Grant DP0449659.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 233-40.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

ESTIMATION STRATEGY

The methodology used for estimation is that of CEE. It has four steps:

1. Identify monetary policy shocks using a structural equation for the interest rate and a VAR(2) to represent the rest of the system. The identification condition used is that monetary shocks have no contemporaneous effect on any variables of the system except the interest rate. The monetary policy rule depends on all variables in the VAR, and these enter the rule both with lags and contemporaneously.
2. Compute the monetary impulse responses C_j^D from this structural VAR (SVAR) (j indexes the j th impulse response).
3. Choose some values for the DSGE model parameters θ and use them to compute the DSGE model monetary impulse responses C_j^M .
4. Find the value of θ that minimizes

$$\sum_{j=1}^M (C_j^D - C_j^M)' W (C_j^D - C_j^M),$$

where W is a diagonal matrix of weights.

They apply this to U.K. data. The original DSGE model they employ has 15 variables, whereas the SVAR(2) has 6. Estimates of the parameters are presented, and some standard errors are given along with plots of the monetary impulses implied by the SVAR and the DSGE model calibrated with the estimated θ .

ESTIMATION PROBLEMS

What could go wrong with this methodology? We discuss three issues in the following subsections.

How Many Impulses To Use?

There is a maximum useful choice for M because the C_j^D are simply functions of the SVAR coefficients. Let there be n variables in the SVAR (in DCN $n = 6$ and it is an SVAR(2)). Then the total number of coefficients in the DCN SVAR(2)

is 77: Of these 77, 72 are from the 2 lags of the six variables in the six equations plus the 5 possible coefficients attached to contemporaneous coefficients in the interest rate rule. Because M seems to be 25, that would mean that 150 impulses are used. This is much larger than the number of parameters determining them. Hence there are many redundant impulse responses and the covariance matrix of C_1^D, \dots, C_{25}^D must be singular. This might be a problem when one uses the δ method to compute standard errors. Indeed, the standard errors of $\hat{\theta}$ found by moment matching in DCN seem to be incredibly small. Thus they have an estimate of the markup λ_f parameter of 2.27 with a standard error of 0.03. It's hard to believe that one could ever get that degree of precision with just 26 years of quarterly observations. Because the SVAR coefficients have a t ratio above 5 in only one case (lagged productivity), it's hard to see how we can end up with t ratios above 400 for $\hat{\theta}$, which are fundamentally derived from the SVAR coefficients.

Approximating the DSGE Model with an SVAR

There is a generic problem here in that the DSGE model often determines m variables and $m > n$; that is, the SVAR is fitted to a smaller number of variables than appear in the DSGE model. This is true of the DCN model, where it appears that $n = 6$, $m = 15$. Now the DSGE model will be an SVAR in the m variables but is unlikely to be an SVAR in the n variables. An old literature, due to Zellner and Palm (1974) and Wallis (1977), has noted that, when a system that is a VAR(p) in n variables is reduced to a smaller system with $m < n$ variables, the smaller system will generally be a vector autoregression moving-average (VARMA) process. Because the CEE procedure involves such compression of variables, it might be expected that a VARMA process is needed rather than a VAR; so, the use of a VAR could lead to specification bias. It might be expected that a VAR of very high order could compensate for this misspecification—and this is generally true—but the order of VAR needed to deliver a good approximation may in fact be far too high for the

data sets one is normally faced with. For example, Kapetanios, Pagan, and Scott (2007) find that reducing a model that is BEQM-like (but half the size) would require a VAR(50) to capture the effects of some shocks (and this with 30,000 observations). The problem has been analyzed in a DSGE context by Ravenna (forthcoming) and Fry and Pagan (2005). We adopt the exposition of the latter.

Suppose that the DSGE model followed a VAR(1) solution (assuming that u_t is i.i.d.):

$$z_t = Pz_{t-1} + Gu_t.$$

Now consider what happens if we model only a subset of the variables. We will call the modeled subset z_{1t} and the omitted variables z_{2t} . We can decompose the VAR above as

$$(1) \quad z_{1t} = P_{11}z_{1t-1} + P_{12}z_{2t-1} + G_1u_t,$$

and we will assume that the following relation holds between z_{1t} and z_{2t} :

$$z_{2t} = \bar{D}_0z_{1t} + \bar{D}_1z_{2t-1} + \bar{D}_2u_t.$$

Substituting this in we get

$$z_{1t} = (P_{11} + P_{12}\bar{D}_0)z_{1t-1} + P_{12}\bar{D}_1z_{2t-2} + G_1u_t + P_{12}\bar{D}_2u_{t-1},$$

so that the sufficient conditions for there to be a finite-order VAR in z_{1t} will be that either $P_{12} = 0$ (i.e., z_{2t} does not Granger cause z_{1t} ; see Lütkepohl, 1993, p. 55, and Quenouille, 1957, p. 43-44) or $\bar{D}_1 = 0, \bar{D}_2 = 0$ (i.e., the variables to be eliminated must be connected to the retained variables through an identity and there can be no “own lag” in the omitted variables in the relation connecting z_{1t} and z_{2t}).

This observation looks trivial, but in fact many of the problems that have arisen where a finite-order VAR does not obtain occur because the omitted variables are connected with the retained variables through an identity, but one that contains an “own lag.” The classic example is in the basic real business cycle (RBC) model where, after log linearization around the steady state, we would get

$$(2) \quad l_t = y_t - c_t$$

$$C^*c_t + K^*k_t = Y^*y_t + (1 - \delta)K^*k_{t-1}$$

$$(3) \quad c_t = E_t(c_{t+1} - \alpha\gamma(y_{t+1} - k_{t+1}))$$

$$y_t = a_t + \alpha k_{t-1} + (1 - \alpha)l_t,$$

where c_t is the log of consumption, a_t is the log of the technology shock, k_t is the log of the capital stock, l_t is the log of labor input, and y_t is the log of output. An asterisk denotes steady-state values, and α is the steady-state share of capital in output. When a_t is an AR(1), the solution to this system can be made a VAR(1) in c_t , l_t , y_t , and k_t . It's clear that we could eliminate any of c_t , l_t , or y_t because none appear as a lagged variable in the system. Equally clearly, k_t cannot be eliminated unless we can find an identity that relates it to other variables but does not involve k_{t-1} . Thus, the identity (3) shows that this is not possible. Most of the literature that seeks to establish that a SVAR cannot approximate a DSGE model (Chari, Kehoe, and McGrattan, 2004; Erceg, Guerrieri, and Gust, 2005; Cooley and Dwyer, 1995) substitute out k_t and so end up with a non-finite-order VAR.

The implication of this for DCN's work is that the reduction of the system from 15 to 6 variables might necessitate a very long VAR and not the VAR(2) they adopted. They used statistical criteria to determine the order of the VAR. Kapetanios, Pagan, and Scott (2007) did this as well, and the tests produced a VAR of order six, far below what was needed (50th order) to produce the correct impulse responses. The reason is that the tests proceed on the assumption that the number of variables in the VAR is correct and it is only the order that needs to be found. So it seems that DCN might be matching impulses that are not strictly comparable. The appropriate procedure would be to (i) simulate a long history of data from the 15-variable DSGE model, incorporating just monetary shocks; (ii) fit a VAR(2) in just 6 of the variables; and then (iii) find the impulse responses from such an approximating VAR, being careful to note that some of the lagged values will be perfectly correlated with others and that it will be necessary to combine variables together to overcome that problem. These are then matched

up with the empirically observed VAR(2) impulse responses in the six variables. We have assessed this by examining a variant of their model, where information is dated at t rather than the combination of t and $t-1$ that is in their paper. However, we used the same model parameters as DCN. Although there are some differences between the true impulse responses and those delivered by a VAR(2), it seems that the approximations are quite good, except at longer horizons. So there does not seem to be an approximation issue here, although in any application one should check that there is no problem, as it is not very difficult to do. However, the δ method used by CEE to compute standard errors is correct only if the approximation is satisfactory. Basically, the estimator of the DSGE-model parameters is an indirect estimator, being derived from functions of the SVAR coefficients represented by the impulse responses. The covariance matrix of such an estimator requires that derivatives of the model-implied VAR impulse responses be computed with respect to the θ parameters, and not the derivatives with respect to the model impulses, as done by CEE. These are only the same if there is no approximation error.

Multiple Solutions

Ignoring the problem identified in the previous section, estimators such as the maximum likelihood estimator basically attempt to match the VAR coefficients from the data with those from the model, rather than attempting to match impulse responses. To see the problems you might encounter with the latter, let us look at the simple model

$$y_t = \beta E_{t-1}(x_{t+1}) + \varepsilon_{yt},$$

$$x_t = \rho x_{t-1} + \varepsilon_{xt}.$$

The VAR will be

$$y_t = a_1 x_{t-1} + \varepsilon_{yt}$$

$$x_t = a_2 x_{t-1} + \varepsilon_{xt},$$

where $a_1 = \beta\rho^2$, $a_2 = \rho$, and the impulse responses are

$$c_{1,y\varepsilon_x}^M = \beta\rho^2, c_{2,y\varepsilon_x}^M = \beta\rho^3, c_{1,x\varepsilon_x} = \rho, c_{2,x\varepsilon_x} = \rho^2.$$

If we would try to find β and ρ by matching the first two impulse responses, we would be minimizing (assuming that the weights in W are equal)

$$(c_{1,y\varepsilon_x}^D - \beta\rho^2)^2 + (c_{2,y\varepsilon_x}^D - \beta\rho^3)^2$$

$$+ (c_{1,x\varepsilon_x}^D - \rho)^2 + (c_{2,x\varepsilon_x}^D - \rho^2)^2.$$

Clearly, such an approach has the problem of producing an order-six polynomial in ρ , so that we may get multiple solutions. This would not arise if we were matching to the VAR coefficients, because then $\hat{\rho} = \hat{a}_2$, $\hat{\beta} = \hat{a}_1/\hat{\rho}^2$. Bearing in mind the first point as well, it seems better to match the VAR to get estimates of θ and then to show the impulse response correspondence.

LOOKING AT SOME OF THE EULER EQUATIONS

Now it would seem useful to develop a method that uses the same information as impulse-response matching but that is a bit simpler, provides ready ways of computing standard errors, emphasizes the economics, and can be used to tell us something about the ability of the DSGE model to match the data. Basically the proposal is to work with the Euler equations and estimate the model parameters directly from them with a single-equation estimator. Of course this is an old idea, but it has fallen out of favor, possibly because of the literature claiming that systems estimators of parameters of the New Keynesian system performed better than the single-equation estimators because of weak instruments. But, in many DSGE models, enough parameters are prescribed that weak instrument problems are not present, and we will see this in the DCN context.

The Euler equations of DSGE models have the generic form¹

$$E_{t-1}z_t = \eta_1 z_{t-1} + \eta_2 E_{t-1}(z_{t+1}) + \eta_3 E_{t-1}w_t.$$

In this equation, z_t is the endogenous variable whose coefficient is normalized to unity, w_t are

¹ The dating of expectations here comes from DCN and reflects the assumption that interest rates have no effect on contemporaneous variables.

either exogenous or other endogenous variables, and the parameters η_j are functions of some of the DSGE model parameters θ . Now this can be written as

$$z_t = \eta_1 z_{t-1} + \eta_2 E_{t-1}(z_{t+1}) + \eta_3 E_{t-1} w_t + \varepsilon_t,$$

and the right-hand-side regressors are uncorrelated with the error ε_t . If we had these conditional expectations we could run a regression. We note that this equation holds for any subset of information used by the economic agents. Hence, let us define the information used in the estimation as that of the DCN VAR(2), that is, two lagged values of $y_t, c_t, i_t, y_t - h_t, r_t$, and Δp_t .² Call these the vector ζ_{t-1} . Then, if we can estimate $E_{t-1}(z_{t+1})$ and $E_{t-1} w_t$, we could simply fit a regression to this equation and thereby measure η_j . Because the model is linear, we can indeed estimate $E_{t-1}(z_{t+1})$ and $E_{t-1} w_t$ as the predictions from the regression of z_{t+1} and w_t against ζ_{t-1} . Basically, this estimation method uses the same information as moment matching, that is, the information contained in the VAR. Notice that standard errors are easily found from this by estimating the Euler equation parameters with an instrumental variables estimator. As we will see later, in most cases the instruments are very good and so there is no reason to doubt the standard errors of η_j found in this way.

This is a relatively simple way to estimate the η_j . Whether the DSGE model parameters θ can be estimated is a different question, because there may be a nonlinear mapping between the η and θ and so we may not be able to recover θ uniquely. This shouldn't concern us unduly because, fundamentally, the impact of monetary policy depends on the η_j ; but there may be some cases where we want to think about changing θ and so would then need to identify it. Ma (2002) pointed out that there was an identification problem like this in strictly forward-looking New Keynesian Phillips curves, and we will see that it comes up in the CEE model as well.

Let us look at the above principles in the context of some of the equations in DCN. First

we look at the Phillips curve. After normalizing on π_t , the Euler equation becomes

$$(4) \quad -E_{t-1}\pi_t + \frac{1}{1+\beta}\pi_{t-1} + \frac{\beta}{1+\beta}E_{t-1}\pi_{t+1} + \frac{(1-\beta\xi)(1-\xi)}{(1+\beta)\xi}E_{t-1}s_t.$$

We can write this as an equation of the form

$$\pi_t = \frac{1}{1+\beta}\pi_{t-1} + \frac{\beta}{1+\beta}\pi_{t+1} + \frac{(1-\beta\xi)(1-\xi)}{(1+\beta)\xi}E_{t-1}\left[\left(\alpha r_t^k + (1-\alpha)(R_t + w_t)\right)\right] + \varepsilon_t$$

or

$$\pi_t - \frac{1}{1+\beta}\pi_{t-1} - \frac{\beta}{1+\beta}\pi_{t+1} = \eta_j E_{t-1}\left[\left(\alpha r_t^k + (1-\alpha)(R_t + w_t)\right)\right] + \varepsilon_t,$$

where $E_{t-1}(\varepsilon_t) = 0$. We note that, because $\beta = 0.99$ is imposed, we are not trying to estimate the coefficients attached to π_t and π_{t+1} . Now, using data, one can form $\alpha r_t^k + (1-\alpha)(R_t + w_t)$. Because DCN pre-set α to 0.36, this can then be regressed against the information represented by the VAR lagged variables to get

$$E_{t-1}\left[\left(\alpha r_t^k + (1-\alpha)(R_t + w_t)\right)\right].^3$$

The regression of this variable against ζ_{t-1} (the VAR(2) lagged variables) gives an R^2 of 0.99, so it is a very good instrument for $\alpha r_t^k + (1-\alpha)(R_t + w_t)$. If we fit a nonlinear regression to this equation, we get an estimated coefficient for ξ_p of 0.988, which is reasonably close to what is reported in the paper from impulse response matching. But the standard deviation is 0.092, which is nowhere near the 0.0004 given in the paper—although, if one makes it robust to serial correlation, it halves. Clearly, the estimate here implies very low frequency of price adjustment, as does DCN's.

² We work with data that are not deviations from steady-state values and so will have to include intercepts in any equation we estimate.

³ Because DCN conclude that $\sigma_a = \infty$, there is no difference between the capital stock and services. We compute the capital stock recursively, but this means the estimates are inaccurate until the initial condition disappears. Because we start the recursion in 1955:Q2, but use data only after 1979:Q2, we feel that the effects of the initial condition have died away, as it will be multiplied by the term $(0.975)^{104}$.

Although this estimate seems implausible, the interpretation would seem to be that there are some problems in the specification of the Phillips curve (indeed the serial correlation in the residuals is consistent with that).

Another equation in DCN we consider estimating is the production function with the form

$$y_t = \lambda_f(\alpha k_t + (1 - \alpha)l_t).$$

Given that DCN prescribe α , we can form $\zeta_t = \alpha k_t + (1 - \alpha)l_t$ and treat this as a regressor to estimate λ_f . There will of course be technology in this relation and, by treating it as an AR(1) process, the equation will have an AR(1) error term. Because the regressor will generally be correlated with the white noise shock driving the AR in technology, we need instruments to estimate λ_f ; for this we use y_{t-1} and ζ_{t-1} . We also include a constant to reflect the fact that we are not using variables that are deviations from a constant steady state and that technology should have a constant mean. Then we get an estimate of λ_f of 1.25, with a standard error of 0.05. This seems more reasonable than the value of 2.27 that DCN obtain, although they give a defense of it. Again, the standard errors are very different.

The interest rate rule parameter values are somewhat puzzling. Under the assumptions in force here, one should be able to fit this rule by ordinary least squares (OLS) regression because it is assumed that the regressors are all uncorrelated with the interest rate shock. If we run the regression, the fit we would get is

$$R_t = 0.883R_{t-1} + (1 - 0.833)(0.0001y_t + 1.28\pi_{t-1}) + 0.05\Delta y_t + 0.10\Delta\pi_t,$$

versus the estimated equation of the paper,

$$R_t = 0.872R_{t-1} + (1 - 0.872)(0.352y_t + 1.28\pi_{t-1}) + 0.43\Delta y_t - 0.62\Delta\pi_t.$$

The standard deviation on y_t from the OLS regression is very small, so these estimates are very different. DCN note that the rule they fit is not the one in the VAR(2), because that would include other lags in the variables. But if we just fit a VAR(1), then it should be comparable to what they claim the estimated money rule is. In fact,

there is not much difference if we add on extra lags. Notice also that the negative sign on $\Delta\pi_t$ that perturbed them has gone. Because this seems a logical way to estimate the money rule, given the assumptions made about the structure of the model, one is puzzled about the results that come from matching impulse responses.

What explains this? One possibility is that the DSGE model implies a particular value for the intercept of the equation, whereas we have just subsumed this into a constant term that is freely estimated. However, the steady-state values used in the model for variables seemed quite close to the sample means over the estimation horizon; so, it would seem that one would get much the same intercepts (provided of course the slope coefficients were correct).

There are some problems with multiple parameter values in both the Phillips curve and the wage equation. Because

$$\pi_t - \frac{1}{1 + \beta}\pi_{t-1} - \frac{\beta}{1 + \beta}\pi_{t+1} = \eta_1 E_{t-1} \left[(\alpha R_t^k + (1 - \alpha)(R_t + w_t)) \right] + \varepsilon_t$$

and

$$\eta_1 = \frac{(1 - \beta\xi_p)(1 - \xi_p)}{(1 + \beta)\xi_p},$$

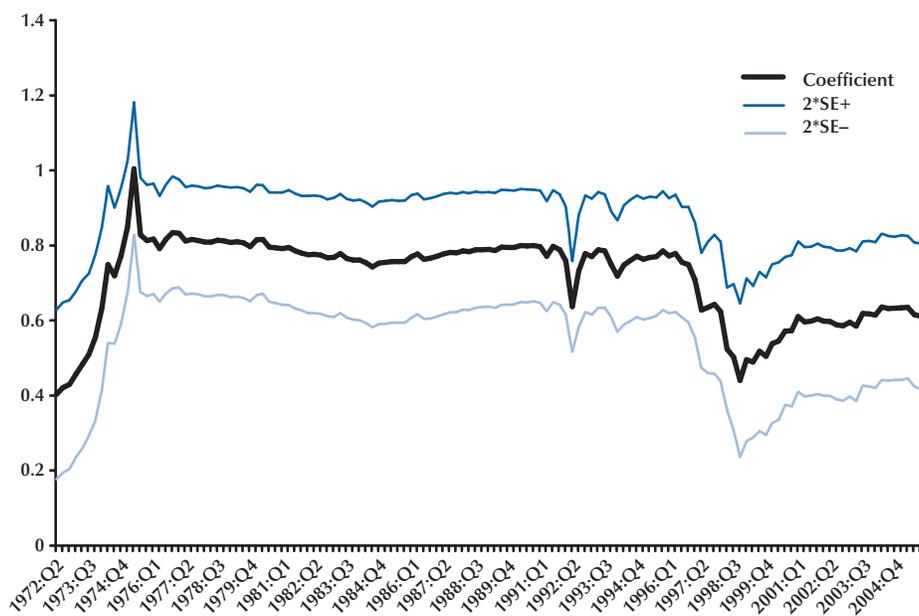
we see that the solution for ξ_p involves a quadratic. There are two estimates of ξ_p that produce exactly the same likelihood—the value of 0.988 given above and 1.02. A similar situation exists for the wage equation. Perhaps this is one reason why Bayesian methods might work better in these models—they would impose the restriction that ξ_p and ξ_w lie between 0 and 1.

STRUCTURAL CHANGE IN THE MODEL

The authors look at structural change in the SVAR and conclude that there was some back in the 1970s, but this was due to industrial issues and not monetary policy regime changes. But it's always difficult to learn something about the sta-

Figure 1

Coefficients of Lagged Inflation with Two Standard Errors on Each Side from a Rolling Regression (68 periods)



bility of the parameters in a VAR. One might also want to ask where to place a possible monetary policy regime change. Is it when Thatcher came in, when inflation targeting was adopted, or when there was a formal change to the institution with the formation of the Monetary Policy Committee (MPC)? Pagan (2003) argued that there had been a change in the level of persistence in inflation in the United Kingdom after the formation of the MPC. This is still evident in the data: See Figure 1, which gives estimates of the coefficient of π_{t-1} using a rolling horizon of 68 quarters.

So this looks like structural change in the dynamics, and perhaps the VAR stability tests should have focused more around the post-1997 point, although this means a very short post-break sample. At the end of the day, graphs like this have to make one wonder about applying a constant-parameter DSGE model to such data. It would seem one might need to use only the post-1997 period to estimate the DSGE model, although with such small samples one might need to use

some sort of Bayesian approach. Perhaps one could use the estimated values of this paper to produce priors.

REFERENCES

- Chari, V.V.; Kehoe, Patrick J. and McGrattan, Ellen R. "A Critique of Structural VARs Using Real Business Cycle Theory." Working Paper No. 631, Federal Reserve Bank of Minneapolis, 2004.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, February 2005, 113(1), pp. 1-45.
- Cooley, Thomas F. and Dwyer, Mark. "Business Cycles Without Much Theory: A Look at Structural VARs." *Journal of Econometrics*, March/April 1995, 83(1-2), pp. 57-88.

Fukač and Pagan

- DiCecio, Riccardo and Nelson, Edward. "An Estimated DSGE Model for the United Kingdom." Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 215-31.
- Erceg, Christopher J.; Guerrieri, Luca and Gust, Christopher. "Can Long-Run Restrictions Identify Technology Shocks?" *Journal of the European Economic Association*, December 2005, 3(6), pp. 1237-78.
- Fry, Renee and Pagan, Adrian R. "Some Issues in Using VARs for Macroeconometric Research." Centre for Applied Macroeconomic Analyses, CAMA Working Paper 2005/18, Australian National University, 2005.
- Justiniano, Alejandro and Preston, Bruce. "Can Structural Small Open Economy Models Account for the Influence of Foreign Disturbances." Centre for Applied Macroeconomic Analyses, CAMA Working Paper, 2006/12, Australian National University, 2006.
- Kapetanios, George; Pagan, Adrian R. and Scott, Alasdair. "Making a Match: Combining Theory and Evidence in Policy-Oriented Macroeconomic Modeling." *Journal of Econometrics*, February 2007, 136(2), pp. 505-94.
- Leitemo, Kai. "Targeting Inflation by Forecast Feedback Rules in Small Open Economies." *Journal of Economic Dynamics and Control*, March 2006, 30(3), pp. 393-413.
- Lubik, Thomas A. and Schorfheide, Frank. "Do Central Banks Respond to Exchange Rate Movements? A Structural Investigation." *Journal of Monetary Economics*, May 2007, 54(4), pp. 1069-87.
- Lütkepohl, Helmut. *An Introduction to Multiple Time Series*. Berlin: Springer-Verlag, 1993.
- Ma, Adrian. "GMM Estimation of the New Phillips Curve." *Economics Letters*, August 2002, 76(3), pp. 411-17.
- Pagan, Adrian R. "Report on Modelling and Forecasting at the Bank of England." Bank of England *Quarterly Bulletin*, Spring 2003, 43(1), pp. 1-29.
- Quenouille, M.H. *The Analysis of Multiple Time Series*. Griffin's Statistical Monographs and Courses No. 1. London: Griffin, 1957.
- Ravenna, Federico. "Vector Autoregressions and Reduced Form Representations of Dynamic Stochastic General Equilibrium Models." *Journal of Monetary Economics* (forthcoming).
- Smets, Frank and Wouters, Raf. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, 2003, 1(5), pp. 1123-75.
- Wallis, Kenneth F. "Multiple Time Series and the Final Form of Econometric Models." *Econometrica*, September 1977, 45(6), pp. 1481-97.
- Zellner, Arnold and Palm, Franz. "Time Series Analysis and Simultaneous Equation Econometric Models." *Journal of Econometrics*, May 1974, 2(1), pp. 17-54.



Macroeconomic Implications of Changes in the Term Premium

Glenn D. Rudebusch, Brian P. Sack, and Eric T. Swanson

Linearized New Keynesian models and empirical no-arbitrage macro-finance models offer little insight regarding the implications of changes in bond term premiums for economic activity. This paper investigates these implications using both a structural model and a reduced-form framework. The authors show that there is no structural relationship running from the term premium to economic activity, but a reduced-form empirical analysis does suggest that a decline in the term premium has typically been associated with stimulus to real economic activity, which contradicts earlier results in the literature. (JEL E43, E44, E52, G12)

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 241-69.

From June 2004 through June 2006, the Federal Reserve gradually raised the federal funds rate from 1 percent to 5¼ percent. Despite this 425-basis-point increase in the short-term rate, long-term interest rates remained at remarkably low levels, with the 10-year Treasury yield averaging 4¼ percent in both 2004 and 2005 and ending September 2006 at just a little above 4½ percent. The apparent lack of sensitivity of long-term interest rates to the large rise in short rates surprised many observers, as such behavior contrasted sharply with interest rate movements during past policy-tightening cycles.¹ Perhaps the most famous expression of this surprise was provided by then-Chairman of the Federal Reserve Alan Greenspan in monetary policy testimony before Congress in February 2005, in which he noted that “the broadly unanticipated behavior of world bond markets remains a conundrum.”

¹ For example, from January 1994 to February 1995, the Federal Reserve raised the federal funds rate by 3 percentage points and the 10-year rate rose by 1.7 percentage points.

The puzzlement over the recent low and relatively stable levels of long-term interest rates has generated much interest in trying to understand both the source of these low rates and their economic implications. In addressing these issues, it is useful to divide the yield on a long-term bond into an expected-rate component that reflects the anticipated average future short rate for the maturity of the bond and a term-premium component that reflects the compensation that investors require for bearing the interest rate risk from holding long-term instead of short-term debt. Chairman Greenspan’s later July 2005 monetary policy testimony suggested that the conundrum likely involved movements in the latter component, noting that “a significant portion of the sharp decline in the ten-year forward one-year rate over the past year appears to have resulted from a fall in term premiums.” This interpretation has been supported by estimates from various finance and macro-finance models that indicate that the recent relatively stable 10-year Treasury yield reflects that the upward revisions

Glenn D. Rudebusch is a senior vice president and associate director of research and Eric T. Swanson is a research advisor at the Federal Reserve Bank of San Francisco. Brian P. Sack is a vice president at Macroeconomic Advisers. The views expressed in this paper are the authors’ and do not necessarily reflect official positions of the Federal Reserve System. The authors thank John Cochrane, Monika Piazzesi, and Jessica Wachter for helpful comments and suggestions. Michael McMorro, Vuong Nguyen, and David Thipphavong provided excellent research assistance.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

to expected future short rates that accompanied the monetary policy tightening were offset, on balance, by a decline in the term premium (e.g., Kim and Wright, 2005, and Rudebusch, Swanson, and Wu, 2006).²

It is this recent experience of a declining term premium in long-term rates that motivates our paper. We examine what is known—both in theory and from the data—about the macroeconomic implications of changes in the term premium. This topic is especially timely and important because of the practical implications of the recent low term premium for the conduct of monetary policy. Specifically, as noted by Federal Reserve Governor Donald Kohn (2005), “the decline in term premiums in the Treasury market of late may have contributed to keeping long-term interest rates relatively low and, consequently, may have supported the housing sector and consumer spending more generally.” Furthermore, any such macroeconomic impetus would alter the appropriate setting of the stance of monetary policy, as described by Federal Reserve Chairman Ben Bernanke (2006):

To the extent that the decline in forward rates can be traced to a decline in the term premium...the effect is financially stimulative and argues for greater monetary policy restraint, all else being equal. Specifically, if spending depends on long-term interest rates, special factors that lower the spread between short-term and long-term rates will stimulate aggregate demand. Thus, when the term premium declines, a higher short-term rate is required to obtain the long-term rate and the overall mix of financial conditions consistent with maximum sustainable employment and stable prices.

Under this “practitioner” view, which is also prevalent among market analysts and private sector macroeconomic forecasters, the recent fall in the term premium provided a boost to real economic activity and, therefore, optimal mone-

tary policy should have followed a relatively more restrictive path as a counterbalance.³

Unfortunately, this practitioner view of the macroeconomic and monetary policy implications of a drop in the term premium is not supported by the simple linearized New Keynesian model of aggregate output that is currently so popular among economic researchers. In that model, output is determined by a forward-looking IS curve:

$$(1) \quad y_t = \beta E_t y_{t+1} - \frac{1}{\gamma} (i_t - E_t \pi_{t+1}) + e_t,$$

where y_t denotes aggregate output and $i_t - E_t \pi_{t+1}$ is the one-period ex ante real interest rate. Solving this equation forward, output can be expressed as a function of short-term real interest rates alone:

$$(2) \quad y_t = -\frac{1}{\gamma} E_t \sum_{j=0}^{\infty} \beta^j (i_{t+j} - \pi_{t+1+j}) + e_t.$$

According to this equation, it is the expected path of the *short-term* real interest rate that determines the extent of intertemporal substitution and hence current output. Long-term interest rates matter only because they embed expectations of future short-term interest rates (as in McGough, Rudebusch, and Williams, 2005). Taken literally, this simple analytic framework does not allow shifts in the term premium to affect output; therefore, according to this model, the recent decline in the term premium should be ignored when constructing optimal monetary policy, and the only important consideration should be the restraining influence of the rising expected-rate component.

Given these contradictory practitioner and New Keynesian views about the macroeconomic implications of changes in the term premium, this paper considers what economic theory more generally implies about this relationship as well as what the data have to say. We start in the next section by examining a structural dynamic stochastic general equilibrium (DSGE) framework

² Of course, as we discuss in detail below, such decompositions of the long rate into expected rates and a term premium are subject to considerable uncertainty.

³ For example, in a January 2005 commentary, the private forecasting firm Macroeconomic Advisers argued that the low term premium was keeping financial conditions accommodative and “would require the Fed to ‘do more’ with the federal funds rate to achieve the desired rate of growth.”

that can completely characterize the relationship between the term premium and the economy. In this framework, unlike its linearized New Keynesian descendant, there are important connections between term premiums and the economy. Unfortunately, given theoretical uncertainties and computational complexities, the model cannot be taken directly to the data, so it provides only qualitative insights about the macroeconomic implications of changes in term premiums, not quantitative empirical assessments.

To uncover such empirical assessments, the third section surveys the recent empirical macrofinance literature, which links the behavior of long-term interest rates to the economy, with varying degrees of economic structure (e.g., Ang and Piazzesi, 2003, and Rudebusch and Wu, 2003, denoted RW). However, although this new literature has made interesting advances in understanding how macroeconomic conditions affect the term premium, it has made surprisingly little progress toward understanding the reverse relationship. Indeed, restrictions are typically imposed in these models that either eliminate any effects of the term premium on the economy or require the term premium to affect the economy in the same way as other sources of long-rate movements. Accordingly, as yet, this literature is not very useful for investigating whether there are important macroeconomic implications of movements in the term premium.

In contrast, as reviewed in the fourth section, several papers have directly investigated the predictive power of movements in the term premium on subsequent gross domestic product (GDP) growth (e.g., Favero, Kaminska, and Söderström, 2005, and Hamilton and Kim, 2002), but because these analyses rely on simple reduced-form regressions, their structural interpretation is unclear. Nevertheless, taken at face value, the bulk of the evidence suggests that decreases in the term premium are followed by *slower* output growth—clearly contradicting the practitioner view (as well as the simple New Keynesian view). However, we reconsider such regressions and provide some new empirical evidence that supports the view

taken by many central bankers and market analysts that a decline in the term premium typically has been associated with *stimulus* to the economy.

The final section concludes by describing some practical lessons for monetary policymakers when confronted with a sizable movement in the term premium.

A STRUCTURAL MODEL OF THE TERM PREMIUM AND THE ECONOMY

In this section, we use a standard structural macroeconomic DSGE framework to study the relationship between the term premium and the economy. In principle, such a framework can completely characterize this relationship; however, in practice the DSGE asset-pricing framework has a number of well-known computational and practical limitations that keep it from being a useful empirical workhorse. Nevertheless, the framework can provide interesting qualitative insights, as we will now show.

An Asset-Pricing Representation of the Term Premium

As in essentially all asset pricing, the fundamental equation that we assume prices assets in the economy is the stochastic discounting relationship:

$$(3) \quad p_t = E_t[m_{t+1}p_{t+1}],$$

where p_t denotes the price of a given asset at time t and m_{t+1} denotes the stochastic discount factor that is used to value the possible state-contingent payoffs of the asset in period $t+1$ (where p_{t+1} implicitly includes any dividend or coupon payouts).⁴ Specifically, the price of a default-free n -period zero-coupon bond that pays one dollar at maturity, $p_t^{(n)}$, satisfies

⁴ Cochrane (2001) provides a comprehensive treatment of this asset-pricing framework. As Cochrane discusses, a stochastic discount factor that prices all assets in the economy can be shown to exist under very weak assumptions; for example, the assumptions of free portfolio formation and the law of one price are sufficient, although these do require that investors are small with respect to the market.

$$(4) \quad p_t^{(n)} = E_t[m_{t+1}p_{t+1}^{(n-1)}],$$

where $p_t^{(0)} = 1$ (the price of one dollar delivered at time t is one dollar).

We can use this framework to formalize the decomposition of bond yields described in the introduction, with the term premium defined as the difference between the yield on an n -period bond and the expected average short-term yield over the same n periods.⁵ Let $i_t^{(n)}$ denote the continuously compounded n -period bond yield (with $i_t \equiv i_t^{(1)}$); then the term premium can be computed from the stochastic discount factor in a straightforward manner:

$$(5) \quad \begin{aligned} i_t^{(n)} - \frac{1}{n} E_t \sum_{j=0}^{n-1} i_{t+j} &= -\frac{1}{n} \log p_t^{(n)} + \frac{1}{n} E_t \sum_{j=0}^{n-1} \log p_{t+j}^{(1)} \\ &= -\frac{1}{n} \log E_t \left[\prod_{j=1}^n m_{t+j} \right] + \frac{1}{n} E_t \sum_{j=1}^n \log E_{t+j-1} m_{t+j}. \end{aligned}$$

Of course, equation (5) does not have an easy interpretation without imposing additional structure on the stochastic discount factor, such as conditional log-normality. Nonetheless, even in this general form, equation (5) highlights an important point: The term premium is not exogenous, as a change in the term premium can only be due to changes in the stochastic discount factor. Thus, to investigate the relationship between the term premium and the economy in a structural model, we must first specify why the stochastic discount factor in the model is changing.

In general, the stochastic discount factor will respond to all of the various shocks affecting the economy, including innovations to monetary policy, technology, and government purchases. Of course, these different types of shocks also have implications for the determination of output and other economic variables. Thus, we would expect the correlation between the term premium and output to depend on which structural shock

⁵ This definition of the term premium (given by the left-hand side of equation (5)) differs from the one used in the theoretical finance literature by a convexity term, which arises because the expected log price of a long-term bond is not equal to the log of the expected price. Our analysis is not sensitive to this adjustment; indeed, some of our empirical term-premium measures are convexity-adjusted and some are not, and they are all highly correlated over our sample.

was driving the change in the term premium. We next elaborate on this point using a simple structural model.

A Benchmark DSGE Structural Model

The expression for the term premium described by equation (5) is quite general but not completely transparent, because it does not impose any structure on the stochastic discount factor. Thus, to illuminate the structural relationship between the term premium and the macroeconomy, we introduce a simple benchmark New Keynesian DSGE model.

The basic features of the model are as follows. Households are representative and have preferences over consumption and labor streams given by

$$(6) \quad \max E_t \sum_{t=0}^{\infty} \beta^t \left(\frac{(c_t - bh_t)^{1-\gamma}}{1-\gamma} - \chi_0 \frac{l_t^{1+\chi}}{1+\chi} \right),$$

where β denotes the household's discount factor, c_t denotes consumption in period t , l_t denotes labor, h_t denotes a predetermined stock of consumption habits, and γ , χ , χ_0 , and b are parameters. We set $h_t = C_{t-1}$, the level of aggregate consumption in the previous period, so that the habit stock is external to the household. There is no investment in physical capital in the model, but there is a one-period nominal risk-free bond and a long-term default-free nominal consol that pays one dollar every period in perpetuity (under our baseline parameterization, the duration of the consol is about 25 years). The economy also contains a continuum of monopolistically competitive firms with fixed, firm-specific capital stocks that set prices according to Calvo contracts and hire labor competitively from households. The firms' output is subject to an aggregate technology shock. Furthermore, we assume there is a government that levies stochastic, lump-sum taxes on households and destroys the resources it collects. Finally, there is a monetary authority that sets the one-period nominal interest rate according to a Taylor-type policy rule:

$$(7) \quad i_t = \rho_i i_{t-1} + (1 - \rho_i) \left[i^* + g_y (y_t - y_{t-1}) + g_\pi \pi_t \right] + \varepsilon_t^i,$$

where i^* denotes the steady-state nominal interest rate, y_t denotes output, π_t denotes the inflation rate (equal to $P_t/P_{t-1}-1$), ε_t^i denotes a stochastic monetary policy shock, and ρ_i , g_y , and g_π are parameters.⁶ This basic structure is very common in the macroeconomics literature, so details of the specification are presented in the appendix.

In equilibrium, the representative household's optimal consumption choice satisfies the Euler equation:

$$(8) \quad (c_t - bc_{t-1})^{-\gamma} = \beta \exp(i_t) E_t(c_{t+1} - bc_t)^{-\gamma} P_t / P_{t+1},$$

where P_t denotes the dollar price of one unit of consumption in period t . The stochastic discount factor is given by

$$(9) \quad m_{t+1} = \frac{\beta(c_{t+1} - bc_t)^{-\gamma} P_t}{(c_t - bc_{t-1})^{-\gamma} P_{t+1}}.$$

The nominal consol's price, $p_t^{(\infty)}$, thus satisfies

$$(10) \quad p_t^{(\infty)} = 1 + E_t m_{t+1} p_{t+1}^{(\infty)}.$$

We define the risk-neutral consol price, $p_t^{(\infty)rn}$, to be

$$(11) \quad p_t^{(\infty)rn} = 1 + \exp(-i_t) E_t p_{t+1}^{(\infty)rn}.$$

The implied term premium is then given by⁷

$$(12) \quad \log\left(\frac{p_t^{(\infty)}}{p_t^{(\infty)rn} - 1}\right) - \log\left(\frac{p_t^{(\infty)rn}}{p_t^{(\infty)rn} - 1}\right).$$

Having specified the benchmark model, we can now solve the model and compute the responses of the term premium and the other variables of the model to economic shocks. Parameters

of the model are given in the appendix. We solve the model by the standard procedure of approximation around the nonstochastic steady state, but because the term premium is zero in a first-order approximation and constant in a second-order approximation, we compute a third-order approximation to the solution of the model using the n th-order approximation package described in Swanson, Anderson, and Levin (2006), called perturbation AIM.

In Figures 1, 2, and 3, we present the impulse response functions of the term premium and output to a 1-percentage-point monetary policy shock, a 1 percent aggregate technology shock, and a 1 percent government purchases shock, respectively. These impulse responses demonstrate that the relationship between the term premium and output depends on the type of structural shock. For monetary policy and technology shocks, a rise in the term premium is associated with current and future weakness in output. By contrast, for a shock to government purchases, a rise in the term premium is associated with current and future output strength. Thus, even the sign of the correlation between the term premium and output depends on the nature of the underlying shock that is hitting the economy.

A second observation to draw from Figures 1, 2, and 3 is that, in each case, the response of the term premium is quite small, amounting to less than one-third of 1 basis point, even at the peak of the response! Indeed, the average level of the term premium for the consol in this model is only 15.7 basis points.⁸ This finding foreshadows

⁶ Note that the interest rate rule we use here is a function of output growth rather than the output gap. We chose to use output growth in the rule because definitions of potential output (and hence the output gap) can sometimes be controversial. In any case, our results are not very sensitive to the inclusion of output growth in the policy rule. For example, if we set the coefficient on output growth to zero, all of our results are essentially unchanged. We also follow much of the literature in assuming an "inertial" policy rule with gradual adjustment and i.i.d. policy shocks. However, Rudebusch (2002 and 2006) argues for an alternative specification with serially correlated policy shocks and little such gradualism.

⁷ The continuously compounded yield to maturity of the consol is given by $\log[p/(p-1)]$. To express the term premium in annualized basis points rather than in logs, equation (12) must be multiplied by 40,000. We obtained qualitatively similar results using alternative term-premium measures in the model, such as the term premium on a two-period zero-coupon bond.

⁸ From the point of view of a second- or third-order approximation, this result is not surprising, because only under extreme curvature or large stochastic variances do second- or third-order terms matter much in a macroeconomic model. Some research has arguably employed such model modifications to account for the term premium. For example, Hördahl, Tristani, and Vestin (2006b) assume that the technology shock has a quarterly standard deviation of 2.5 percent and a persistence of 0.986. Adopting these two parameter values in our model causes the term premium to rise to 141 basis points. Ravenna and Seppälä (2006) assume a shock to the marginal utility of consumption, with a persistence of 0.95 and a quarterly standard deviation of 8 percent. A similar shock in our model boosts the term premium to 41 basis points. Wachter (2006) assumes a habit parameter (b) of 0.961, which in our model boosts the term premium to 22.3 basis points. Thus, we are largely able to replicate some of these authors' findings; nonetheless, we believe that our benchmark parameter values are the most standard ones in the macroeconomics literature (e.g., Christiano, Eichenbaum, and Evans, 2005).

Figure 1

Impulse Responses to a 1-Percentage-Point Federal Funds Rate Shock

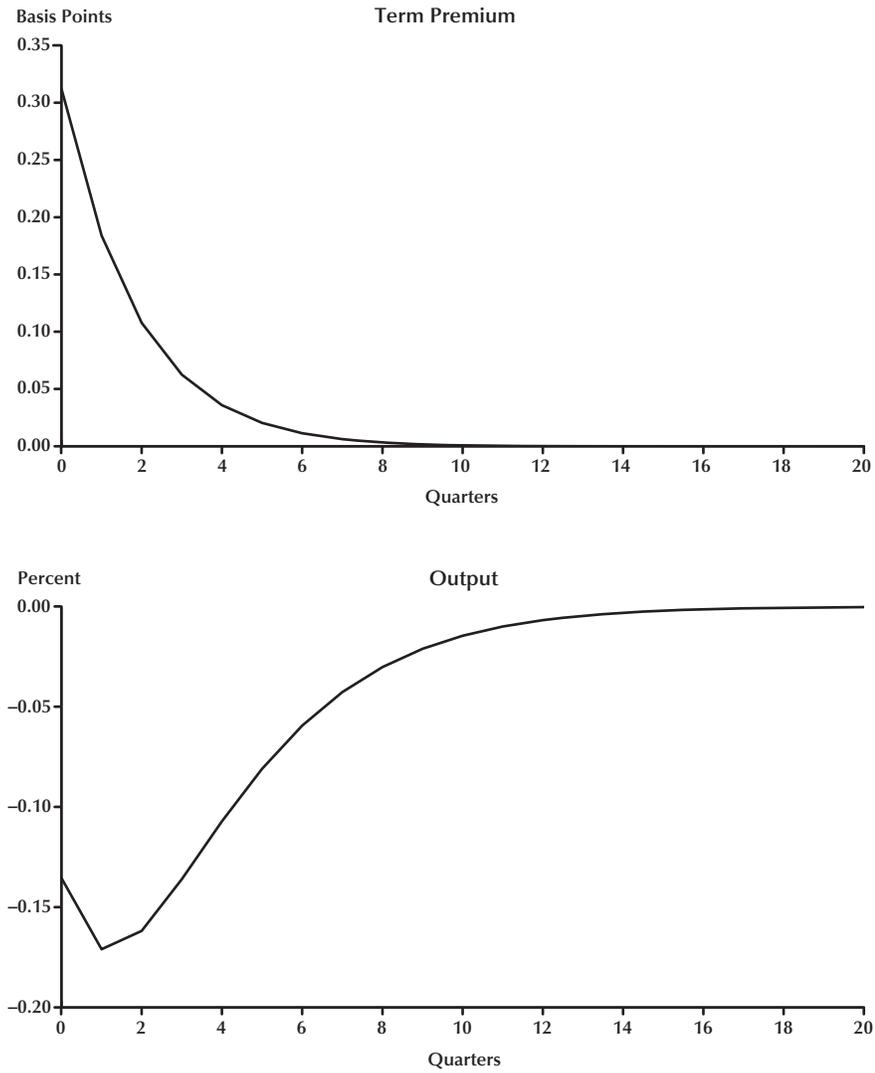


Figure 2

Impulse Responses to a 1 Percent Technology Shock

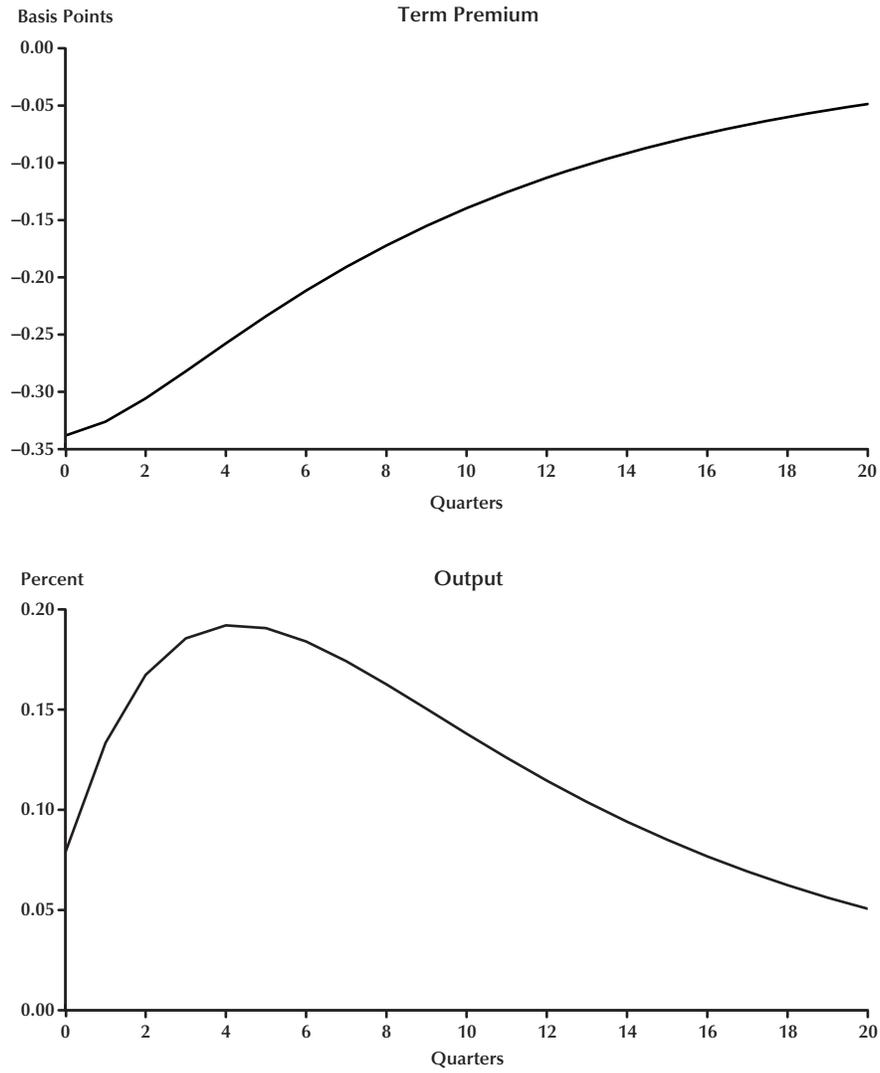
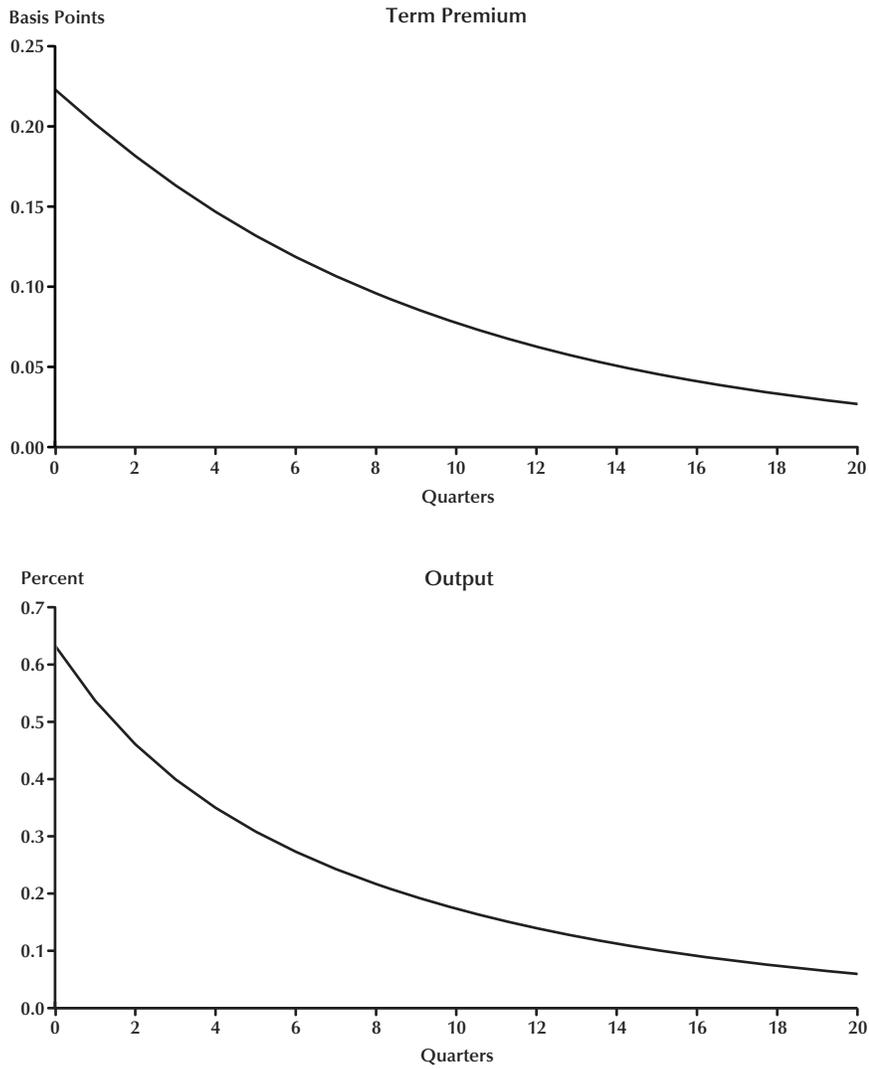


Figure 3

Impulse Responses to a 1 Percent Government Purchases Shock



one of the primary limitations of the structural approach to modeling term premiums, which we will discuss in more detail below.

Finally, we note that, although this structural model is very simple, in principle there is no reason why the same analysis cannot be performed using larger and more realistic DSGE models, such as Smets and Wouters (2003), Christiano, Eichenbaum, and Evans (2005), or the extensions of these in use at a number of central banks and international policy institutions.⁹ Even with these larger models, we can describe the term-premium response to any given structural shock and the broader implications of the shock for the economy and optimal monetary policy.

Limitations of the DSGE Model of the Term Premium

Using a structural DSGE model to investigate the relationship between the term premium and the economy has advantages in terms of conceptual clarity, but there are also a number of limitations that prevent the structural-modeling approach from being useful at present as an empirical workhorse for studying the term premium. This remains true despite the increasing use of structural macroeconomic models at policy-making institutions for the study of other macroeconomic variables, such as output and inflation. These limitations generally fall into two categories: theoretical uncertainties and computational intractabilities.

Regarding the former, even though some DSGE models—sometimes crucially augmented with highly persistent structural shocks—appear to match the empirical impulse responses of macroeconomic variables, such as output and inflation, researchers do not agree on how to specify these models to match asset prices. For example, a variety of proposals to explain the equity-premium puzzle include habit formation

in consumption (Campbell and Cochrane, 1999), time-inseparable preferences (Epstein and Zin, 1989), and heterogeneous agents (Constantinides and Duffie, 1996, and Alvarez and Jermann, 2001). This lack of consensus implies that there is much uncertainty about the appropriate DSGE specification for analyzing the term premium.

The possibility that a heterogeneous-agent model is necessary to understand risk premiums poses perhaps the most daunting challenge for structural modelers of the term premium. In the case of heterogeneous agents with limited participation in financial markets, different households' valuations of state-contingent claims are not equalized, so determining equilibrium asset prices can become much more complicated than in the representative-household case. Although a stochastic discount factor still exists under weak assumptions even in the heterogeneous-household case, it need not conform to the typical utility functions that are in use in current structural macroeconomic models.¹⁰

The structural approach to asset pricing also faces substantial computational challenges, particularly for the larger-scale models that are becoming popular for the analysis of macroeconomic variables. Closed-form solutions do not exist in general, and full numerical solutions are computationally intractable except for the simplest possible models.¹¹ The standard approach of log-linearization around a steady state that has proved so useful in macroeconomics is unfortunately not applicable to asset pricing, because by construction it eliminates all risk premiums in the model. Some extensions of this procedure to a hybrid log-linear log-normal approximation (Wu, 2006, and Bekaert, Cho, and Moreno, 2005)

⁹ Some notable extensions include Altig et al. (2005) to the case of firm-specific capital, Adolfson et al. (2007) to the case of a small open economy, and Pesenti (2002) and Erceg, Guerrieri, and Gust (2006) to a large-scale (several hundred equations) multicountry-block context for use at the International Monetary Fund and the Federal Reserve Board, respectively.

¹⁰ One might even question the assumptions required for a stochastic discount factor to exist. For example, if there are large traders and some financial markets are thin, then it is no longer the case that all investors can purchase any amount of a security at a constant price, contrary to the standard assumptions.

¹¹ See Backus, Gregory, and Zin (1989), Donaldson, Johnsen, and Mehra (1990), Den Haan (1995), and Chapman (1997) for examples of numerical solutions for bond prices in very simple real business cycle models. Gallmeyer, Hollifield, and Zin (2005) provide a closed-form solution for bond prices in a simple New Keynesian model, under the assumption of a very special reaction function for monetary policy.

and to a full second-order approximation around the steady state (Hördahl, Tristani, and Vestin, 2006b) are only moderately more successful, because they imply that all risk premiums in the model are constant (in other words, these authors all assume the weak form of the expectations hypothesis). Obtaining a local approximation that actually produces time-varying risk or term premiums requires a full third-order approximation, as in our analysis above and in Ravenna and Seppälä (2006). Even then, the implied time variation in the term premium is very small, due to the inherently small size of third-order terms, unless one is willing to assume very large values for the curvature of agents' utility functions, very large stochastic shock variances, and/or very high degrees of habit persistence (which goes back to the theoretical limitations discussed above). Thus, the challenges in computing the asset-pricing implications of DSGE models, while becoming less daunting over time, remain quite substantial.

MACRO-FINANCE MODELS OF THE TERM PREMIUM

Because of the significant limitations in applying the structural model discussed above, researchers interested in modeling the term premium in a way that can be taken to the data have had no choice but to pursue a less-structural approach. Although one can model “yields with yields” using a completely reduced-form, latent-factor, no-arbitrage asset-pricing model, as in Duffie and Kan (1996) and Dai and Singleton (2000), recent research has focused increasingly on hybrid macro-finance models of the term structure, in which some connections between macroeconomic variables and risk premiums are drawn, albeit not within the framework of a fully structural DSGE model (see Diebold, Piazzesi, and Rudebusch, 2005). The approaches employed in this macro-finance literature have generally fallen into two categories: vector autoregression (VAR) macro-finance models and New Keynesian macro-finance models. We consider each in turn.

VAR-Based Macro-Finance Models

The first paper in the no-arbitrage macro-finance literature was Ang and Piazzesi (2003).¹² They assume that the economy follows a VAR:

$$(13) \quad X_t = \mu + \Phi X_{t-1} + \Sigma \varepsilon_t,$$

where the state vector, X_t , contains output, inflation, the one-period nominal interest rate, and two latent factors (discussed below). The stochastic shock, ε_t , is i.i.d. over time. In this model, the one-period nominal interest rate, i_t , is determined by a Taylor-type monetary policy rule based on X_t , so that the model-implied expected path of the short-term interest rate is known at any point in time.

The VAR, however, does not contain any information about the stochastic discount factor. Ang and Piazzesi simply assume that the stochastic discount factor falls into the essentially affine class, as in standard latent-factor finance models, so it has the functional form

$$(14) \quad m_{t+1} = \exp\left(-i_t - \frac{1}{2} \lambda_t' \lambda_t - \lambda_t' \varepsilon_{t+1}\right),$$

where ε_t is assumed to be conditionally log-normally distributed and the prices of risk, λ_t , are assumed to be affine in the state vector, X_t :

$$(15) \quad \lambda_t = \lambda_0 + \lambda_1 X_t.$$

Estimation of this model is complicated by the inclusion of two unobserved, latent factors in the state vector, X_t , which are typical of no-arbitrage models in the finance literature. To make estimation tractable, Ang and Piazzesi impose the restriction that the unobserved factors do not interact at all with the observed macroeconomic variables (output and inflation) in the VAR. Because of this very strong restriction, the macro-

¹² A number of papers before Ang and Piazzesi (2003) investigated the dynamic interactions between yields and macroeconomic variables in the context of unrestricted VARs, including Evans and Marshall (2001) and Kozicki and Tinsley (2001). Diebold, Rudebusch, and Aruoba (2006) and Kozicki and Tinsley (2005) provide follow-up analysis. As with the no-arbitrage papers discussed below, however, none of these papers has explored whether the term premium implied by their models feeds back to the macroeconomy, the question of interest in the present paper.

economic variables in the model are determined by a VAR that essentially excludes all interest rates (both short-term and long-term rates). Thus, while the Ang-Piazzesi model can effectively capture the extent to which changes in macroeconomic conditions affect the term premium, it cannot capture any aspects of that relationship running in the reverse direction.¹³ In this regard, their model falls short of addressing the topic of interest in the present paper.¹⁴

Bernanke, Reinhart, and Sack (2004, denoted BRS), employ a similar model but assume that the state vector, X_t , consists entirely of observable macroeconomic variables, which determine both short-rate expectations (through the VAR) and the prices of risk (15). By eliminating the use of latent variables, the empirical implementation of the model is simplified tremendously. Of course, as in Ang and Piazzesi, the BRS framework will capture effects of movements in the term premium driven by observable factors included in the VAR, but it does not empirically separate the role of the term premium from that of lagged macroeconomic variables. Note that the BRS specification, as in the Ang and Piazzesi model, does not include longer-term interest rates in the VAR (but in this case does include the short-term interest rate), implying that movements in the term premium not captured by the included variables are assumed to have no effect on the dynamics of the economy.

¹³ Even when movements in the term premium are driven by the observed macroeconomic variables (output and inflation) rather than the latent factors, the Ang-Piazzesi model fails to identify effects of the term premium on the macroeconomy. For example, suppose higher inflation is estimated to raise the term premium and lead to slower growth in the future. We cannot ascribe the slower growth to the term premium, because the higher inflation may also predict tighter monetary policy or other factors that would be expected to slow the economy. Note that the VAR does at least partially address the issue that not all movements in the term premium are created equal, because the predictive power of a change in the term premium will depend on the specific combination of economic factors driving it.

¹⁴ Cochrane and Piazzesi (2006) also focus on the interaction between macroeconomic conditions and the term premium. They use the predictable component of the ex post returns from holding longer-term securities as a measure of the term premium. Their findings support the case that the term premium varies importantly over time, and they link those movements to macroeconomic conditions. However, they do not address whether the term premium itself affects economic activity.

Ang, Piazzesi, and Wei (2006, denoted APW), also estimate a no-arbitrage macro-finance model based on a VAR of observed state variables. However, in contrast to BRS, APW explicitly include the five-year Treasury yield as an element of the state vector. Thus, to a very limited extent, their model begins to address the types of effects that are the focus of the present paper. However, their VAR does not distinguish between the risk-neutral and term-premium components of the five-year yield, so it is only able to capture distinct effects from these two components if they are correlated (in different ways) with the other variables in the VAR (which are, specifically, the short-term interest rate and GDP growth). Even then, it would not be possible in their model to disentangle the direct effects of the short-term interest rate and GDP growth on future output from the indirect effects that changes in those variables have on the term premium; it is in this respect that the APW model cannot help answer the question we are interested in, even though it allows a separate role for longer-term yields in the VAR.¹⁵

Finally, Dewachter and Lyrio (2006a,b) consider a model that is very similar in spirit to APW and BRS, only they work in continuous time and allow for a time-varying long-run inflation objective of the central bank, as argued for by Kozicki and Tinsley (2001) and Gürkaynak, Sack, and Swanson (2005). However, just as with the other papers discussed above, Dewachter and Lyrio do not allow changes in the term premium to feed back to the macroeconomic variables of the model.

New Keynesian Macro-Finance Models

A separate strand of the macro-finance literature has attempted to bridge the gulf between DSGE models and VAR-based macro-finance models by incorporating more economic structure into the latter. Specifically, these papers replace the reduced-form VAR in the macro-finance models with a structural New Keynesian macro-

¹⁵ APW also present some related reduced-form results on the forecasting power of the term premium for future GDP growth, which we discuss in more detail in the next section.

economic model that governs the dynamics of the macroeconomic variables.

An early and representative paper in this literature was written by Hördahl, Tristani, and Vestin (2006a, denoted HTV). They begin with a basic New Keynesian structural model in which output, inflation, and the short-term nominal interest rate are governed by the equations

$$(16) \quad y_t = \mu_y E_t y_{t+1} + (1 - \mu_y) y_{t-1} - \zeta_i (i_t - E_t \pi_{t+1}) + \varepsilon_t^y,$$

$$(17) \quad \pi_t = \mu_\pi E_t \pi_{t+1} + (1 - \mu_\pi) \pi_{t-1} + \delta_y y_t - \varepsilon_t^\pi,$$

$$(18) \quad i_t = \rho_i i_{t-1} + (1 - \rho_i) [g_\pi (E_t \pi_{t+1} - \pi_t^*) + g_y y_t] + \varepsilon_t^i.$$

Equation (16) describes a New Keynesian curve that allows for some degree of habit formation on the part of households through the lagged output term; equation (17) describes a New Keynesian Phillips curve that allows for some rule-of-thumb price setters through the lagged inflation term; and equation (18) describes the monetary authority's Taylor-type short-term interest rate reaction function. Equations (16) and (17) are structural in the sense that they can be derived from a log-linearization of household and firm optimality conditions in a simple structural New Keynesian DSGE model along the lines of our benchmark model (although HTV modify this structure by allowing the long-run inflation objective, π_t^* , to vary over time).

In contrast to a DSGE asset-pricing model, however, HTV model the term premium using an ad hoc affine structure for the stochastic discount factor, as in the VAR-based models above. Although this approach is not completely structural, it makes the model computationally tractable and provides a good fit to the data while allowing the term premium to vary over time in a manner determined by macroeconomic conditions that are determined structurally (to first order). The true appeal of this type of model is that it is parsimonious and simple while allowing for expectations to influence macroeconomic

dynamics and for the term premium to vary non-trivially to macroeconomic developments.

However, as was the case in the VAR-based models, the HTV model does not allow the term premium to feed back to macroeconomic variables. As discussed in the introduction, the structure of the IS curve in the HTV model assumes that economic activity depends only on expectations of the short-term real interest rate and not on the term premium. Thus, this approach is also unable to address the issue considered in the current paper.

RW develop a New Keynesian macro-finance model that comes a step closer to addressing the topic of this paper by allowing for feedback from the term structure to the macroeconomic variables of the model. In particular, RW incorporate two latent term-structure factors into the model and give those latent factors macroeconomic interpretations, with a level factor that is tied to the long-run inflation objective of the central bank and a slope factor that is tied to the cyclical stance of monetary policy. Thus, the latent factors in the RW model can affect economic activity, and the term structure does provide information about the current values of those latent factors. However, RW make no effort to decompose the effects of long-term interest rates on the economy into an expectations component and a term-premium component, so there is no sense in which the term premium itself affects macroeconomic variables.

Wu (2006) and Bekaert, Cho, and Moreno (2005) come closer to a true structural New Keynesian macro-finance model by deriving the stochastic discount factor directly from the utility function of the representative household in the underlying structural model. Thus, like a DSGE model, their papers impose the cross-equation restrictions between the macroeconomy and the stochastic pricing kernel that are ignored when the kernel is specified in an ad hoc affine manner. However, these analyses also suffer from the computational limitations of working within the DSGE framework (discussed above), because both papers are unable to solve the model as specified.

Instead, those authors use a log-linear, log-normal approximation, which implies that the term premium in the model is time-invariant.¹⁶ Thus, their papers do not address the question we have posed in this paper.¹⁷

REDUCED-FORM EVIDENCE ON THE EFFECTS OF THE TERM PREMIUM

Because of the limitations discussed above, the models in the previous two sections do not provide us with much insight into the empirical economic implications of changes in the term premium. The benchmark structural model is largely unable to reproduce the magnitude and variation of the term premium that is observed in the bond market, and, although the macro-finance models are more successful at capturing the observed behavior of term premiums, they typically impose very restrictive assumptions that eliminate any macroeconomic implications of changes in term premiums. A separate literature that has provided a direct examination of these implications is based on reduced-form empirical evidence. Specifically, in the large literature that uses the slope of the yield curve to forecast subsequent GDP growth, several recent papers have tried to estimate separately the predictive power of the term premium. In this section, we summarize these papers and contribute some new evidence on this issue.

An important caveat worth repeating is that there is only a reduced-form relationship—not a structural one—between the term premium and future output growth, so even the sign of their pair-

wise correlation over a given sample will depend on which types of shocks are most influential. Nevertheless, it may be of interest to consider the average correlation between future output growth and changes in the term premium over some recent history. If the mixture of shocks is expected to remain relatively stable, then the average estimated reduced-form relationship between the term premium and future economic growth could be useful for forecasting. For this reason, the historical relationship may provide useful information to a policymaker who has to decide whether and how to respond to a given change in the term premium.

Evidence in the Literature

Recent research relating the term premium to subsequent GDP growth has been part of a much larger literature on the predictive power of the slope of the yield curve. A common approach in this literature is to investigate whether the spread between short-term and long-term interest rates has significant predictive power for future GDP growth by estimating a regression of the form

$$(19) \quad y_{t+4} - y_t = \beta_0 + \beta_1(y_t - y_{t-4}) + \beta_2(i_t^{(n)} - i_t) + \varepsilon_t,$$

where y_t is the log of real GDP at time t and $i_t^{(n)}$ is the n -quarter interest rate (usually a longer-term rate such as the 10-year Treasury yield).¹⁸ The standard finding is that the estimated coefficient β_2 is significant and positive, indicating that the yield-curve slope helps predict growth.

Note that equation (19) is a reduced-form specification that has no economic structure. However, it can be motivated by thinking of the long-term interest rate as a proxy for the neutral level of the nominal funds rate, so that the yield-curve slope captures the current stance of monetary policy relative to its long-run level. For example, a steep yield-curve slope (with short rates unusually low relative to long rates) would indicate that policy is accommodative and would

¹⁶ Indeed, the term premium would be zero except for the fact that Wu (2006) and Bekaert, Cho, and Moreno (2005) allow some second- and higher-order terms to remain in these models. In particular, they leave the log-normality of the stochastic pricing kernel in its nonlinear form, which implies a nonzero, albeit constant, risk premium. A drawback of this approach is that it treats some second-order terms as important, while dropping other terms of similar magnitude.

¹⁷ A related paper by Gallmeyer, Hollifield, and Zin (2005) provides a full nonlinear solution to a very similar model. However, they are able to solve the model only under the assumption of an extremely special reaction function for monetary policy; thus, their method has no generality and is invalid in cases in which that policy reaction function is not precisely followed.

¹⁸ This equation assumes that the dependent variable is future GDP growth (a continuous variable). Other papers in this literature use a dummy variable for recessions (a discrete variable). In either case, the motivation for the approach is the same and the results are qualitatively similar.

be associated with faster subsequent growth, thus accounting for the positive coefficient.

In this respect, the use of the long-term interest rate in the regression (19) is motivated entirely by the component related to the expected long-run level of the short rate. But the long-term rate also includes a term premium; hence, any variation in this premium will affect the performance of the equation. Indeed, it is useful to decompose the yield-curve slope into these two components, as follows:

$$(20) \quad i_t^{(n)} - i_t = \left(\frac{1}{n} \sum_{j=0}^{n-1} E_t i_{t+j} - i_t \right) + \left(i_t^{(n)} - \frac{1}{n} \sum_{j=0}^{n-1} E_t i_{t+j} \right).$$

The first term captures the expectations component, or the proximity of the short rate to its expected long-run level. The second component is the term premium, or the amount by which the long rate exceeds the expected return from investing in a series of short-term instruments. For notational simplicity, we will denote the first component in (20) as $exsp_t$, the expected-rate component of the yield spread, and the second component as tp_t , the term premium.

With this decomposition, the prediction equation (19) can be generalized as follows:

$$(21) \quad y_{t+4} - y_t = \beta_0 + \beta_1(y_t - y_{t-4}) + \beta_2 exsp_t + \beta_3 tp_t + \varepsilon_t.$$

The standard equation (19) imposes the coefficient restriction $\beta_2 = \beta_3$. Loosening that restriction allows the term premium to have a different implication for subsequent growth than the expected-rate component.¹⁹ Several recent papers have considered this issue, as we will briefly summarize.

The first paper to examine the importance of the above decomposition for forecasting was

Hamilton and Kim (2002), which forecasts future GDP growth using a spread between the 10-year and 3-month Treasury yields in equation (19). The innovation of their paper is that it then separates the yield spread into the expectations and term-premium components considered in equation (21). The authors achieve this separation by considering the ex post realizations of short rates, using instruments known ex ante to isolate the expectations component. They find that the coefficients β_2 and β_3 are indeed statistically significantly different from one another, although both coefficients are estimated to be positive. Note that a positive value for β_3 implies that a decline in the term premium is associated with *slower* future growth.

A second paper that decomposes the predictive power of the yield spread into its expectations and term-premium components is Favero, Kaminska, and Söderström (2005). These authors differ from Hamilton and Kim (2002) by using a real-time VAR to compute short-rate expectations rather than a regression of ex post realizations of short rates on ex ante instruments. As in Hamilton and Kim (2002), they find a positive sign for the coefficient β_3 , so that a lower term premium again predicts slower GDP growth.

A third relevant paper is by Wright (2006), who touches on this issue in the context of a probit model for forecasting recessions. Wright considers the predictive power of the yield slope, and then he investigates whether the return forecasting factor from Cochrane and Piazzesi (2005) also enters those regressions significantly. Since this factor is correlated with the term premium, he is implicitly controlling for the term premium, as in equation (21). He finds that this factor is insignificant for predicting recessions over horizons of two or four quarters but has a significant negative coefficient for predicting recessions over a six-quarter horizon; that is, a lower term premium raises the odds of a recession, consistent with the findings of the other papers that it would predict slower growth.

A final reference is APW. As noted above, they use a VAR that includes long rates, GDP growth, and a short rate, but they cannot separate out the effects of the term premium from other

¹⁹ Because this equation is intended to capture the effects on output from changes in interest rates, it is not far removed from the literature on estimating IS curves. Most empirical implementations of the IS curve, however, assume that output is related to short-term interest rates rather than long-term interest rates. Or, as seen in Fuhrer and Rudebusch (2004), these papers focus on the component of long rates tied to short-rate expectations, following the New Keynesian output equation very closely. As a result, even this literature is more closely tied to estimating the parameter β_2 than the parameter β_3 .

movements in long-term interest rates. However, the authors perform an additional exercise in which they calculate the expected-rate and term-premium components of the long rate as implied by the VAR and then estimate the forecasting equation (21), allowing for different effects from these two components. In contrast to the previously discussed papers, APW find that the term premium has no predictive power for future GDP growth; that is, the coefficient β_3 is zero.

Overall, the handful of papers that have directly tackled the predictive power of the term premium have produced results that starkly contrast with the intuition that Chairman Bernanke expressed in his March 2006 speech (see the introduction). The empirical studies to date suggest that, if anything, the relationship has the *opposite* sign from the practitioner view. According to these results, policymakers had no basis for worrying that the decline in the term premium might be stimulating the economy and instead should have worried that it was a precursor to lower GDP growth.

Empirical Estimates of the Term Premium

Estimation of equation (21) requires a measure of the term premium, and there are a variety of possibilities in the literature. We begin our empirical analysis by collecting a number of the prominent term-premium measures and examining some of the similarities and differences among them.

Specifically, we consider five measures of the term premium on a zero-coupon nominal 10-year Treasury security²⁰:

1. *VAR measure*: The first of these measures, which we label the “VAR” measure, is based on a straightforward projection of the short rate from a simple but standard three-

variable macroeconomic VAR comprising four lags each of the unemployment rate, quarterly inflation in the consumer price index, and the 3-month Treasury bill rate. At each date the VAR can be used to forecast the short rate over a given horizon, and the average expected future short rate can be used as an estimate of the risk-neutral long-term rate of that maturity.²¹ The difference between the observed long-term rate and the risk-neutral long-term rate then provides a simple estimate of the term premium. This approach has been used by Evans and Marshall (2001), Favero, Kaminska, and Söderström (2005), Diebold, Rudebusch, and Aruoba (2006), and Cochrane and Piazzesi (2006).

2. *Bernanke-Reinhart-Sack measure*: A potential shortcoming of using a VAR to estimate the term premium is that it does not impose any consistency between the yield curve at a given point in time and the VAR’s projected evolution of those yields. Such pricing consistency can be imposed by using a no-arbitrage model of the term structure. As discussed in the previous section, a no-arbitrage structure can be laid on top of a VAR to estimate the behavior of the term premium, as in BRS. Here, we consider the term-premium estimate from that paper, as updated by Rudebusch, Swanson, and Wu (2006).
3. *Rudebusch-Wu measure*: No-arbitrage restrictions can also be imposed on top of a New Keynesian macroeconomic model. Here we take the term premium estimated from one such model, Rudebusch and Wu (2003 and 2007), discussed earlier. As with the Bernanke-Reinhart-Sack measure, this term-premium measure was extended to a

²⁰ Note that some of these term-premium measures are adjusted for convexity (e.g., Kim-Wright, Bernanke-Reinhart-Sack, and Rudebusch-Wu), and some are not (e.g., our VAR-based measure and our extension of the Cochrane-Piazzesi measure). The adjustment for convexity has little or no impact on our results, however; for example, the correlation between the VAR-based term-premium measure and the Kim-Wright and Bernanke-Reinhart-Sack measures are 0.94 and 0.96, respectively.

²¹ Of course, there are several reasons for not taking these VAR projections too seriously as good measures of the actual interest rate expectations of bond traders at the time. Rudebusch (1998) describes three important limitations of such VAR representations: (i) the use of a time-invariant, linear structure, (ii) the use of final revised data and full-sample estimates, and (iii) the limited number of information variables. We examined several rolling-sample estimated VARs as well and obtained similar results.

longer sample by Rudebusch, Swanson, and Wu (2006), and we use this extended version below.

4. *Kim-Wright measure*: One can also estimate the term premium using a standard no-arbitrage dynamic latent-factor model from finance (with no macroeconomic structure underlying the factors). In these models, risk-neutral yields and the term premium are determined by latent factors that are themselves linear functions of the observed bond-yield data. We use the term-premium measure from a three-factor model discussed by Kim and Wright (2005), which we extend back to 1961.²²
5. *Cochrane-Piazzesi measure*: Cochrane and Piazzesi (2005) analyze excess returns for a range of securities over a one-year holding period. Their primary finding is that a single factor—a particular combination of current forward rates—predicts a considerable portion of the excess returns from a one-year holding period for Treasury securities. For our purposes, however, we are interested in the *term premium* on a 10-year security, or the (annualized) excess return expected over the 10-year period. Sack (2006a) provides a straightforward approach for converting the Cochrane-Piazzesi one-year holding-period results into a measure of the term premium. Specifically, the expected one-period excess returns implied by the Cochrane-Piazzesi estimates, together with the one-year risk-free rate, imply an expected set of zero-coupon yields one year ahead (because the only way to generate expected returns on zero-coupon securities is through changes

in yield). Those expected future yields can then be used to compute the expected Cochrane-Piazzesi factor one year ahead and, hence, the expected excess returns over the one-year period beginning one year ahead. By iterating forward, one can compute the expected excess return for each of the next 10 years, thereby yielding a measure of the term premium on the 10-year security.

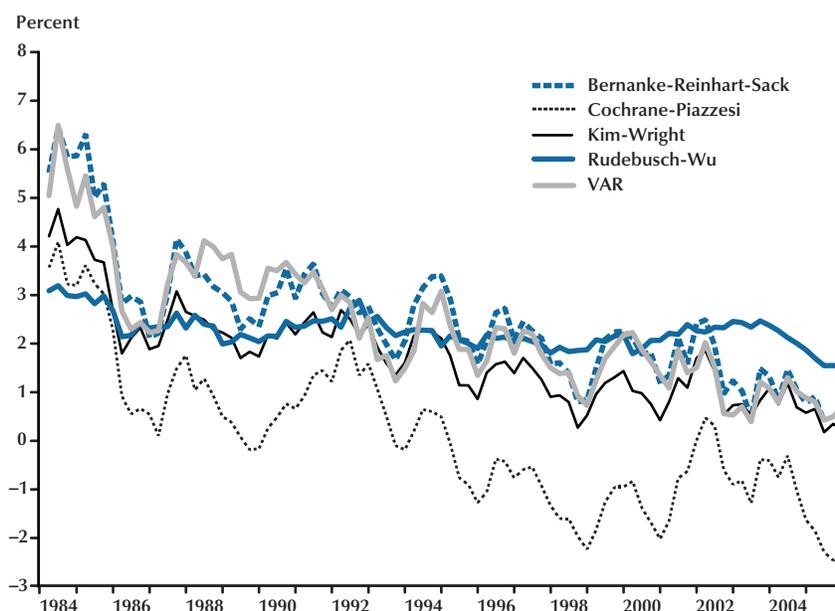
As is clear from the above descriptions, the approaches used to derive the five term-premium measures differ considerably in the variables included and the theoretical restrictions incorporated. Nevertheless, the measures show many similar movements over time, as can be seen in Figure 4, which plots the five measures of the term premium for the 10-year zero-coupon Treasury yield back to 1984.

Three of the measures, in particular—the VAR, Bernanke-Reinhart-Sack, and Kim-Wright—are remarkably highly correlated over this period.²³ As shown in Table 1, the correlation coefficients among these measures range from 0.94 to 0.98. The other two measures—Rudebusch-Wu and Cochrane-Piazzesi—are less correlated with the others. For example, the correlation coefficients with the VAR measure are 0.68 for Rudebusch-Wu and 0.88 for Cochrane-Piazzesi. These lower correlations largely reflect that the Rudebusch-Wu measure is more stable than the others and that the Cochrane-Piazzesi measure is more volatile.

The greater stability of the Rudebusch-Wu measure can be easily understood. Their underlying model attributes much of the variation in the 10-year Treasury yield to changes in the expected future path of short rates, reflecting, in their framework, variation in the perceived inflation target of the central bank. That assumption is supported by other research. For example, Gürkaynak, Sack, and Swanson (2005) found significant systematic variation in far-ahead for-

²² We extend the Kim-Wright measure back to 1961 by regressing the three Kim-Wright latent factors on the first three principal components of the yield curve and using these coefficients to estimate the Kim-Wright factors in prior years. Because the term premium in the model is a linear function of observed yields, and because the Kim-Wright model fits the yield-curve data very well, this exercise should come very close to deriving the same factors that would be implied if we extended their model back to 1961. Over the period where our proxy and the actual Kim-Wright term premium overlap, the correlation between the two measures is 0.998 and the average absolute difference between them is less than 4 basis points.

²³ These correlations are very high in comparison with, say, the zero correlations exhibited by various authors' measures of monetary policy shocks, as noted in Rudebusch (1998).

Figure 4**Five Measures of the 10-Year Term Premium****Table 1****Correlations Among the Five Measures of the Term Premium**

	BRS	RW	KW	CP	VAR
BRS	1.00				
RW	0.76	1.00			
KW	0.98	0.81	1.00		
CP	0.92	0.87	0.96	1.00	
VAR	0.96	0.68	0.94	0.88	1.00

ward nominal interest rates in response to macroeconomic news in a way that suggested changes in inflation expectations rather than changes in term premiums. Similarly, Kozicki and Tinsley (2001) found that statistical models that allow for a “moving endpoint” are able to fit interest rate and inflation time series much better than standard stationary or difference-stationary VARs. By attributing more of the movement in long rates to short-rate expectations, the Rudebusch-Wu

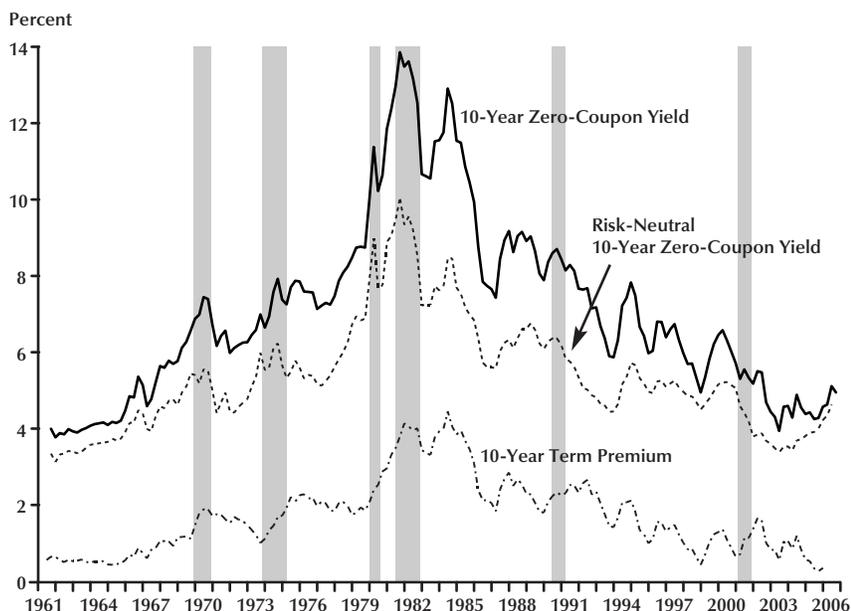
analysis does not need as much variation in the term premium to explain the observed variation in yields.²⁴

The behavior of the measure based on Cochrane and Piazzesi (2005) is harder to understand. This measure is well below the other meas-

²⁴ One could argue that a weakness of the other term-premium estimates is that they are based on models that assume that the long-run features of the economy, such as the steady-state real interest rate and rate of inflation, are completely anchored.

Figure 5

Kim-Wright Decomposition of the 10-Year Zero-Coupon Yield



NOTE: The shaded bars indicate recessions as dated by the National Bureau of Economic Research.

ures and is much more volatile. To a large extent, this behavior simply mimics the one-period expected excess returns computed by Cochrane and Piazzesi. Indeed, Sack (2006b) and Wright (2006) have pointed out that the implied one-period expected excess returns are surprisingly volatile and are currently very negative. This behavior partly shows through to the implied term-premium measure.

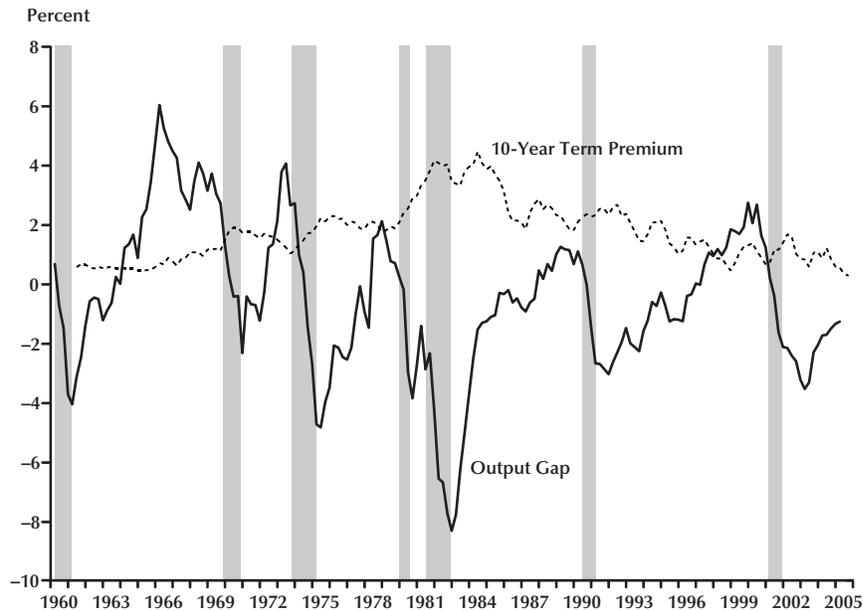
Overall, Figure 4 provides us with a menu of choices for the analysis that follows.²⁵ Even with the differences noted above, the five measures show considerable similarities in their variation over this sample. Indeed, the first principal component captures 95 percent of the variation in these five term-premium estimates. In the analysis in the next section, we focus our attention on the

Kim-Wright measure. This measure appears to be representative of the other measures considered. In fact, it is very highly correlated (0.99) with the first principal component of all five measures. Moreover, it has the advantage that it can be extended back to the early 1960s, allowing us to conduct our analysis over a longer sample.

The 10-year zero-coupon yield is shown in Figure 5 along with the two components based on the Kim-Wright term-premium estimate.²⁶ As can be seen, both short-rate expectations and the term premium contributed to the run-up in yields through the early 1980s and, since then, to the decline in yields. As noted by Kim and Wright (2005), the term premium recently has fallen to very low levels, a pattern consistent with the

²⁵ In contrast to the measures shown in Figure 4, Ludvigson and Ng (2006) provide one that has considerable high-frequency variation and little persistence or predictive power for economic activity. However, we have some reservations about their identification of the term premium and exclude it from our analysis.

²⁶ The yield data considered here are from the Gürkaynak, Sack, and Wright (2006) database. Those authors do not recommend using the 10-year Treasury yield before 1971, as there are very few maturities at that horizon for estimating the yield curve. However, their 10-year yield is highly correlated with the Treasury constant-maturity 10-year yield over that period, which justified its use. All results that follow are robust to beginning the sample in 1971.

Figure 6**Kim-Wright Term Premium and the CBO Output Gap**

NOTE: The shaded bars indicate recessions as dated by the National Bureau of Economic Research.

conundrum discussed by former Chairman Alan Greenspan.

Figure 6 plots this term-premium measure along with the Congressional Budget Office (CBO) output gap and provides the first hint of a negative relationship between the two. It is this relationship that we now explore in more detail.

New Evidence on the Implications of the Term Premium

We begin by estimating the standard relationship between the slope of the yield curve and subsequent GDP growth, using the specification in equation (19). The long rate is a 10-year zero-coupon Treasury yield, taken from the Gürkaynak, Sack, and Wright (2006) database. The short rate is the 3-month Treasury bill rate from the Federal Reserve's H.15 data release. All data are quarterly averages, and the sample ranges from 1961:Q3 to 2005:Q4. We examine both this full sample and a shorter subsample beginning in 1984, which

arguably has a more consistent monetary policy regime (e.g., Rudebusch and Wu, 2007).

Results are presented in the first column of Table 2. Over the full sample, we find that the coefficient for the yield-curve slope is highly statistically significant and has a positive sign. This estimate implies that a flatter yield curve predicts slower GDP growth, the standard finding in the academic literature. Over the shorter sample, the estimated coefficient loses its significance, reflecting another fact that is well-appreciated among researchers—that the predictive power of the yield-curve slope for growth appears to have diminished in recent decades.

As discussed above, this approach is purely a reduced-form exercise that is not explicitly tied to a theoretical structure. However, a common motivation for using the yield-curve slope as a predictor is that it serves as a proxy for the stance of monetary policy relative to its neutral level. Given this motivation, one would prefer to measure the yield-curve slope based strictly on the

Table 2**Prediction Equations for GDP Growth****Dependent Variable: $y_{t+4} - y_t$**

	(1)	(2)	(3)	(4)
1962-2005 Sample				
$y_t - y_{t-4}$	0.15 (1.57)	0.12 (1.18)	0.32 (3.04)	0.38 (4.22)
$i_t^{(n)} - i_t$	0.64 (3.64)			
$exsp_t$		0.68 (4.03)	1.03 (5.64)	
$exsp_{t-4}$			-0.79 (-3.49)	
tp_t		0.30 (0.92)	-0.61 (-1.34)	
tp_{t-4}			0.54 (1.24)	
$exsp_t - exsp_{t-4}$				0.96 (5.62)
$tp_t - tp_{t-4}$				-0.77 (-1.95)
1985-2005 Sample				
$y_t - y_{t-4}$	0.26 (2.54)	0.32 (2.31)	0.36 (2.30)	0.36 (2.68)
$i_t^{(n)} - i_t$	0.28 (1.29)			
$exsp_t$		0.35 (1.59)	0.46 (1.92)	
$exsp_{t-4}$			-0.07 (-0.32)	
tp_t		0.07 (0.25)	-0.46 (-1.15)	
tp_{t-4}			0.61 (2.18)	
$exsp_t - exsp_{t-4}$				0.30 (1.37)
$tp_t - tp_{t-4}$				-0.59 (-1.93)

NOTE: Coefficient estimates are shown with their t -statistics in parentheses (t -statistics have been corrected for residual heteroskedasticity and autocorrelation). Each regression includes a constant that is not reported.

portion of the long-term interest rate associated with expectations of the short-term rate. In that context, we can also ask how the other component of the long rate—the term premium—affects growth. This consideration leads to specification (21) above, in which the two components of the yield-curve slope are allowed to have different predictive effects for subsequent GDP growth.

We can implement this approach using the term-premium measure described above.²⁷ The results are shown in column 2. For both samples, the expectations-based component of the yield slope has slightly stronger predictive power than the pure yield-curve slope (that is, the coefficient

on this component is slightly larger and more significant than the coefficient on the overall slope reported in column 1), and the coefficient on the term premium, β_3 , is not significantly different from zero. However, we are unable to reject the hypothesis that $\beta_2 = \beta_3$ at even the 10 percent level over either the post-1962 or post-1985 sample.

Our findings are similar in spirit to the existing empirical evidence that the term premium has a different effect on subsequent growth than the expectations-related component of the yield curve. Note that the only purpose of having a term-premium measure, according to these results, is to determine the expectations component of the yield slope more accurately. The term premium itself has no predictive power for future growth.

²⁷ In our analysis, we ignore any potential issues associated with generated regressors.

However, the specification of these regression equations seems somewhat at odds with the models we presented earlier. For example, the New Keynesian IS curve (2) could be used to motivate the use of the yield-curve slope, as it assumes that output is determined by the deviation of the real short-term interest rate from its equilibrium level. The expectations component of the yield-curve slope might capture this variable, but it should then be related to the *level* of the output gap. In contrast, the reduced-form specifications (19) and (21) relate the slope of the yield curve to the *growth rate* of output. Thus, this specification seems to differ from the more structural models by a derivative. Moreover, the term premium in Figures 5 and 6 appears to be nonstationary or nearly nonstationary, while GDP growth is much closer to being stationary. Thus, from a statistical point of view, specifications (19) and (21) are also highly suspect.

If we difference equation (2) to arrive at a specification in growth rates, it would suggest that it is *changes* in the stance of monetary policy that predict future GDP growth.²⁸ This suggests investigating whether GDP growth is tied to changes in the stance of policy and changes in the term premium, as opposed to the levels of those variables.

As an exploratory step in this direction, we re-estimate equation (21) with an additional one-year lag of the right-hand-side variables included in the regression. The results, shown in column 3 of Table 2, strongly hint that there is greater predictive power associated with the changes in these variables than with their levels. Indeed, one can reject the hypothesis that the coefficients on the lagged variables are zero (at the 1 percent significance level). Moreover, one cannot reject that the right-hand-side variables enter the regression only as changes. That is, the hypothesis that the coefficients on the lag of these components equal the negative of the coefficients on their current levels cannot be rejected even at the 10 percent

significance level. A similar (though less striking) pattern is found in the shorter sample.

Because both the theory and the hypothesis tests in the preceding paragraph suggest that only differences should matter, column (4) of the table presents results from estimating the baseline forecasting regression equation in differences, namely,

$$(22) \quad y_{t+4} - y_t = \beta_0 + \beta_1(y_t - y_{t-4}) + \beta_2(exsp_t - exsp_{t-4}) + \beta_3(tp_t - tp_{t-4}) + \varepsilon_t.$$

The full-sample results indicate that both components of the yield-curve slope matter for future growth. The coefficient on the risk-neutral expectations component of the yield-curve slope is now larger and more statistically significant than in any of the earlier specifications. We can also overwhelmingly reject the hypothesis that $\beta_2 = \beta_3$ (with p -values less than 10^{-4}). This finding indicates that GDP growth is expected to be higher not when the short-term interest rate is merely low relative to its long-run level, but when it has *fallen* relative to that level.

More importantly for this paper, we find that the estimated coefficient on the term premium is now negative and (marginally) statistically significant. According to these results, a decline in the term premium tends to be followed by *faster* GDP growth—the opposite sign of the relationship uncovered by previous empirical studies. (In the shorter sample, all of the coefficients are again less significant. However, we still reject the hypothesis that $\beta_2 = \beta_3$ [with a p -value of 0.0395] in column 4, and the coefficient on the change in the term premium is again negative and borderline statistically significant.²⁹)

Our findings line up with the intuition expressed by Chairman Bernanke when he suggested that the declining term premium signaled additional stimulus to the economy. Our results

²⁸ Some might argue that the dependent variable here should be the growth of the output gap rather than GDP. As discussed below, we obtained similar results using the change in the CBO output gap as the dependent variable.

²⁹ Furthermore, using the year-on-year change in the CBO output gap as the predicted variable rather than the year-on-year change in output itself gave similar results. Specifically, the coefficient on the term premium remains negative, with a p -value just less than 0.05. These results suggest that a decline in the term premium predicts a higher future value of the output gap and that policymakers might want to take that prediction into account when formulating the optimal policy response.

are the first piece of evidence (that we are aware of) to support this hypothesis, and they stand in sharp contrast to the previous empirical evidence presented by Hamilton and Kim (2002), Favero, Kaminska, and Söderström (2005), and Wright (2006).

CONCLUSIONS

Our results can be usefully summarized from the perspective of advising monetary policymakers. Specifically, policymakers may wonder how they should respond when confronted with a substantial change in the term premium, such as the recent decline that appears to have taken place during 2004 and 2005.

The first, and perhaps most important, conclusion from our analysis is that policymakers should always try to determine the source of the change in the term premium. If that source can be identified, then policymakers are advised to consider the repercussions of that underlying driving force more broadly rather than focusing exclusively on the change in the term premium. In this way, policymakers can take into account the macroeconomic implications of the structural shifts or disturbances that are affecting the term premium.

Of course, policymakers often may be uncertain about the reasons for changes in the term premium. Indeed, during the past few years, a variety of only tentative explanations have been offered for the seemingly low term premium. In such a situation, policymakers may find our reduced-form analysis of the implications of the term premium for future economic activity to be a useful baseline. Our results suggest that a decline in the term premium has typically been associated with higher future GDP growth, which appears consistent with the practitioner view. Indeed, according to our reduced-form analysis, the attention that Federal Reserve officials paid to the seemingly large decline in the term premium in 2004 and 2005 may have been justified.

Finally, our finding that changes in the term premium have a significant correlation with future GDP growth is not captured by many macroeco-

nomical models. Understanding and incorporating this correlation within the framework of a model would appear to be a useful addition to the research agenda. In this regard, we only speculate that our empirical findings may reflect a heterogeneous population in which a decline in the term premium makes financial market conditions more accommodative for certain classes of borrowers.

REFERENCES

- Adolfson, Malin; Laséen, Stefan; Lindé, Jesper and Villani, Mattias. "Evaluating an Estimated New Keynesian Small Open Economy Model." *Journal of Economic Dynamics and Control*, 2007 (forthcoming).
- Altig, David; Christiano, Lawrence J.; Eichenbaum, Martin and Lindé, Jesper. "Firm-Specific Capital, Nominal Rigidities, and the Business Cycle." Unpublished manuscript, Northwestern University, 2005.
- Alvarez, Fernando and Jermann, Urban. "Quantitative Asset Pricing Implications of Endogenous Solvency Constraints." *Review of Financial Studies*, Winter 2001, 14(4), pp. 1117-51.
- Ang, Andrew and Piazzesi, Monika. "No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables." *Journal of Monetary Economics*, 2003, 50(4), pp. 745-87.
- Ang, Andrew; Piazzesi, Monika and Wei, Min. "What Does the Yield Curve Tell Us about GDP Growth?" *Journal of Econometrics*, March/April 2006, 131(1/2), pp. 359-403.
- Backus, David K.; Gregory, Allan W. and Zin, Stanley E. "Risk Premiums in the Term Structure." *Journal of Monetary Economics*, November 1989, 24(3), pp. 371-99.
- Bekaert, Geert; Cho, Seonghoon and Moreno, Antonio. "New-Keynesian Macroeconomics and the Term Structure." Unpublished manuscript, Columbia Business School, 2005.
- Bernanke, Ben S. "Reflections on the Yield Curve and Monetary Policy." Presented at the Economic

- Club of New York, March 20, 2006; www.federalreserve.gov/boarddocs/speeches/2006/.
- Bernanke, Ben S.; Reinhart, Vincent R. and Sack, Brian P. "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment." *Brookings Papers on Economic Activity*, 2004, 2, pp. 1-100.
- Campbell, John Y. and Cochrane, John H. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, April 1999, 107(2), pp. 205-51.
- Chapman, David. "The Cyclical Properties of Consumption Growth and the Real Term Structure." *Journal of Monetary Economics*, 1997, 39(2), pp. 145-72.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, February 2005, 113(1), pp. 1-45.
- Cochrane, John H. *Asset Pricing*. Princeton: Princeton University Press, 2001.
- Cochrane, John and Piazzesi, Monika. "Bond Risk Premia." *American Economic Review*, March 2005, 95(1), pp. 138-60.
- Cochrane, John and Piazzesi, Monika. "Decomposing the Yield Curve." Unpublished manuscript, University of Chicago, 2006.
- Constantinides, George and Duffie, Darrell. "Asset Pricing with Heterogeneous Consumers." *Journal of Political Economy*, April 1996, 104(2), pp. 219-40.
- Dai, Qiang and Singleton, Kenneth J. "Specification Analysis of Affine Term Structure Models." *Journal of Finance*, October 2000, 55(5), pp. 1943-78.
- Den Haan, Wouter J. "The Term Structure of Interest Rates in Real and Monetary Economies." *Journal of Economic Dynamics and Control*, July-September 1995, 19(5-7), pp. 909-40.
- Dewachter, Hans and Lyrio, Marco. "Macro Factors and the Term Structure of Interest Rates." *Journal of Money, Credit, and Banking*, February 2006a, 38(1), pp. 119-40.
- Dewachter, Hans and Lyrio, Marco. "Learning, Macroeconomic Dynamics and the Term Structure of Interest Rates." Unpublished manuscript, 2006b.
- Diebold, Francis X.; Rudebusch, Glenn D. and Aruoba, S. Boragan. "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach." *Journal of Econometrics*, March/April 2006, 131(1/2), pp. 309-38.
- Diebold, Francis; Piazzesi, Monika and Rudebusch, Glenn D. "Modeling Bond Yields in Finance and Macroeconomics." *American Economic Review*, May 2005, 95(2), pp. 415-20.
- Donaldson, John B; Johnsen, Thore and Mehra, Rajnish. "On the Term Structure of Interest Rates." *Journal of Economic Dynamics and Control*, July-October 1990, 14(3/4), pp. 571-96.
- Duffie, Darrell and Kan, Rui. "A Yield-Factor Model of Interest Rates." *Mathematical Finance*, October 1996, 6(4), pp. 379-406.
- Epstein, Lawrence and Zin, Stanley. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 1989, 57(4), pp. 937-69.
- Erceg, Christopher J.; Guerrieri, Luca and Gust, Christopher. "SIGMA: A New Open Economy Model for Policy Analysis." *International Journal of Central Banking*, March 2006, 2(1), pp. 1-50.
- Evans, Charles and Marshall, David. "Economic Determinants of the Term Structure of Nominal Interest Rates." Unpublished manuscript, Federal Reserve Bank of Chicago, 2001.
- Favero, Carlo A.; Kaminska, Iryna and Söderström, Ulf. "The Predictive Power of the Yield Spread: Further Evidence and a Structural Interpretation." Unpublished manuscript, Università Bocconi, 2005.
- Fuhrer, Jeffrey C. and Rudebusch, Glenn D. "Estimating the Euler Equation for Output." *Journal of Monetary Economics*, September 2004, 51(6), pp. 1133-53.

Rudebusch, Sack, Swanson

- Gallmeyer, Michael F.; Hollifield, Burton and Zin, Stanley E. "Taylor Rules, McCallum Rules and the Term Structure of Interest Rates." *Journal of Monetary Economics*, July 2005, 52(5), pp. 921-50.
- Greenspan, Alan. Monetary Policy Report to the Congress. Board of Governors of the Federal Reserve System, February 16, 2005.
- Greenspan, Alan. Monetary Policy Report to the Congress. Board of Governors of the Federal Reserve System, July 20, 2005.
- Gürkaynak, Refet S.; Sack, Brian and Swanson, Eric. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review*, March 2005, 95(1), pp. 425-36.
- Gürkaynak, Refet S.; Sack, Brian and Wright, Jonathan. "The U.S. Treasury Yield Curve: 1961 to the Present." Finance and Economics Discussion Series 2006-28, Board of Governors of the Federal Reserve System, 2006.
- Hamilton, James D. and Kim, Dong Heon. "A Reexamination of the Predictability of Economic Activity Using the Yield Spread." *Journal of Money, Credit, and Banking*, May 2002, 34(2), pp. 340-60.
- Hördahl, Peter; Tristani, Oreste and Vestin, David. "A Joint Econometric Model of Macroeconomic and Term Structure Dynamics." *Journal of Econometrics*, March-April 2006a, 131(1/2), pp. 405-44.
- Hördahl, Peter; Tristani, Oreste and Vestin, David. "The Yield Curve and Macroeconomic Dynamics." Unpublished manuscript, European Central Bank, 2006b.
- Kim, Don H. and Wright, Jonathan H. "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates." Finance and Economics Discussion Series 2005-33, Board of Governors of the Federal Reserve System, 2005.
- Kohn, Donald L. "Monetary Policy Perspectives on Risk Premiums in Financial Markets." Presented at the Financial Markets Risk Premiums Conference, Board of Governors of the Federal Reserve System, Washington, DC, July 21, 2005; www.federalreserve.gov/boarddocs/speeches/2005/.
- Kozicki, Sharon and Tinsley, Peter A. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics*, January 2001, 47(3), pp. 613-52.
- Kozicki, Sharon and Tinsley, Peter A. "What Do You Expect? Imperfect Policy Credibility and Tests of the Expectations Hypothesis." *Journal of Monetary Economics*, March 2005, 52(2), pp. 421-47.
- Lettau, Martin and Uhlig, Harald. "Can Habit Formation Be Reconciled with Business Cycle Facts?" *Review of Economic Dynamics*, January 2000, 3(1), pp. 79-99.
- Ludvigson, Sydney and Ng, Serena. "Macro Factors in Bond Risk Premia." Unpublished manuscript, New York University, 2006.
- Macroeconomic Advisers. "The Rise in Long-Term Interest Rates: A Reason to Stop Sooner." *Monetary Policy Insights Fixed Income Focus* (newsletter), April 18, 2006.
- McGough, Bruce; Rudebusch, Glenn D. and Williams, John C. "Using a Long-Term Interest Rate as the Monetary Policy Instrument." *Journal of Monetary Economics*, July 2005, 52(5), pp. 855-79.
- Pesenti, Paolo. "The IMF Global Economy Model (GEM): Theoretical Framework." Unpublished manuscript, International Monetary Fund, 2002.
- Ravenna, Federico and Seppälä, Juha. "Monetary Policy and the Term Structure of Interest Rates." Unpublished manuscript, University of California, Santa Cruz, 2006.
- Rudebusch, Glenn D. "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review*, November 1998, 39(4), pp. 907-31.
- Rudebusch, Glenn D. "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia." *Journal of Monetary Economics*, September 2002, 49(6), pp. 1161-87.

- Rudebusch, Glenn D. "Monetary Policy Inertia: Fact or Fiction?" *International Journal of Central Banking*, December 2006, 2(4), pp. 85-135.
- Rudebusch, Glenn D.; Swanson, Eric and Wu, Tao. "The Bond Yield 'Conundrum' from a Macro-Finance Perspective." *Monetary and Economic Studies*, 2006, 24(S-1), pp. 83-109.
- Rudebusch, Glenn D. and Wu, Tao. "A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy." Working Paper 2003-17, Federal Reserve Bank of San Francisco, 2003 (forthcoming in *Economic Journal*).
- Rudebusch, Glenn D. and Wu, Tao. "Accounting for a Shift in Term Structure Behavior with No-Arbitrage and Macro-Finance Models." *Journal of Money, Credit, and Banking*, 2007, 39(2-3), pp. 395-422.
- Sack, Brian P. "The Cochrane-Piazzesi Return Forecasting Factor and the Term Premium." Unpublished manuscript, Macroeconomic Advisers, 2006a.
- Sack, Brian P. "Comment on 'Can Central Banks Target Bond Prices?'" in *Monetary Policy in an Environment of Low Inflation: Proceedings of the Bank of Korea International Conference 2006*. Seoul, Korea: Bank of Korea, 2006b.
- Smets, Frank and Wouters, Raf. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of European Economic Association*, September 2003, 1(5), pp. 1123-75.
- Swanson, Eric; Anderson, Gary and Levin, Andrew. "Higher-Order Perturbation Solutions to Dynamic, Discrete-Time Rational Expectations Models." Working Paper 2006-01, Federal Reserve Bank of San Francisco, 2006.
- Wachter, Jessica A. "A Consumption-Based Model of the Term Structure of Interest Rates." *Journal of Financial Economics*, February 2006, 79(2), pp. 365-99.
- Wright, Jonathan. "The Yield Curve and Predicting Recessions." Finance and Economics Discussion Series 2006-7, Board of Governors of the Federal Reserve System, 2006.
- Wu, Tao. "Macro Factors and the Affine Term Structure of Interest Rates." *Journal of Money, Credit, and Banking*, October 2006, 38(7), pp. 1847-75.

APPENDIX

Benchmark New Keynesian Model

To better understand the structural relationship between the term premium and the macroeconomy, we define a simple New Keynesian DSGE model to use as a benchmark. This appendix provides a detailed description of the model, the benchmark parameter values we used in computing the impulse responses in Figures 1 to 3, and our solution algorithm.

The economy contains a continuum of households with a total mass of unity. Households are representative and seek to maximize utility over consumption and labor streams given by

$$(23) \quad \max E_t \sum_{t=0}^{\infty} \beta^t \left(\frac{(c_t - bh_t)^{1-\gamma}}{1-\gamma} - \chi_0 \frac{l_t^{1+\chi}}{1+\chi} \right),$$

where β denotes the household's discount factor, c_t denotes consumption in period t , l_t denotes labor, h_t denotes a predetermined stock of consumption habits, and γ, χ, χ_0 , and b are parameters. We will set $h_t = C_{t-1}$, the level of aggregate consumption in the previous period, so that the habit stock is external to the household.³⁰ The household's stochastic discount factor from period t to $t+j$ thus satisfies

$$m_{t,t+j} \equiv \beta^j \frac{(c_{t+j} - bC_{t+j-1})^{-\gamma}}{(c_t - bC_{t-1})^{-\gamma}} \frac{P_t}{P_{t+j}}.$$

The economy also consists of a continuum of monopolistically competitive intermediate goods firms indexed by $f \in [0,1]$. Firms have Cobb-Douglas production functions:

$$(24) \quad y_t(f) = A_t \bar{k}^\alpha l_t(f)^{1-\alpha},$$

where \bar{k} is a fixed, firm-specific capital stock (identical across firms) and A_t denotes an aggregate technology shock that affects all firms. The level of aggregate technology follows an exogenous AR(1) process:

$$(25) \quad \log A_t = \rho_A \log A_{t-1} + \varepsilon_t^A,$$

where ε_t^A denotes an i.i.d. aggregate technology shock with mean zero and variance σ_A^2 . Intermediate goods are purchased by a perfectly competitive final goods sector that produces the final good with a constant elasticity of substitution production technology:

$$(26) \quad Y_t = \left[\int_0^1 y_t(f)^{1/(1+\theta)} df \right]^{1+\theta}.$$

Each intermediate goods firm f thus faces a downward-sloping demand curve for its product given by

$$(27) \quad y_t(f) = \left(\frac{p_t(f)}{P_t} \right)^{-(1+\theta)/\theta} Y_t,$$

³⁰ Campbell and Cochrane (1999) consider instead a habit stock, which is an infinite sum of past aggregate consumption with geometrically decaying weights, and a slightly different specification of the utility kernel. They argue that this specification fits asset prices better than the one-period habits used here. However, Lettau and Uhlig (2000) argue that the Campbell-Cochrane specification significantly worsens the model's ability to fit consumption and labor data.

where

$$(28) \quad P_t \equiv \left[\int_0^1 p_t(f)^{-1/\theta} df \right]^{-\theta}$$

is the constant elasticity of substitution aggregate price of a unit of the final good.

Each firm sets its price, $p_t(f)$, according to a Calvo contract that expires with probability $1 - \xi$ each period. There is no indexation, so the price, $p_t(f)$, is fixed over the life of the contract. When a contract expires, the firm is free to reset its price as it chooses. In each period t , firms must supply whatever output is demanded at the posted price, $p_t(f)$. Firms hire labor, $l_t(f)$, from households in a competitive labor market, paying the nominal market wage, w_t . Marginal cost for firm f at time t is thus given by

$$(29) \quad mc_t(f) = \frac{w_t l_t(f)}{(1 - \alpha) y_t(f)}.$$

Firms are collectively owned by households and distribute profits and losses back to the households. When a firm's price contract expires and it is able to set a new contract price, the firm maximizes the expected present discounted value of profits over the lifetime of the contract:

$$(30) \quad E_t \sum_{j=0}^{\infty} \xi^j m_{t,t+j} \left[p_t(f) y_{t+j}(f) - w_{t+j} l_{t+j}(f) \right],$$

where $m_{t,t+j}$ is the representative household's stochastic discount factor from period t to $t+j$. The firm's optimal contract price, $p_t^*(f)$, thus satisfies

$$(31) \quad p_t^*(f) = \frac{(1 + \theta) E_t \sum_{j=0}^{\infty} \xi^j m_{t,t+j} mc_{t+j}(f) y_{t+j}(f)}{E_t \sum_{j=0}^{\infty} \xi^j m_{t,t+j} y_{t+j}(f)}.$$

To aggregate up from firm-level variables to aggregate variables, it is useful to define the cross-sectional price dispersion, Δ_t :

$$(32) \quad \Delta_t^{1/(1-\alpha)} \equiv (1 - \xi) \sum_{j=0}^{\infty} \xi^j p_{t-j}^*(f)^{-(1+\theta)/(\theta(1-\alpha))},$$

where the exponent $1/(1 - \alpha)$ arises from the firm-specificity of capital.³¹ We can then write

$$(33) \quad Y_t = \Delta_t^{-1} A_t \bar{K}^\alpha L_t^{1-\alpha},$$

where $\bar{K} = \bar{k}$ and

$$(34) \quad L_t \equiv \int_0^1 l_t(f) df$$

and equilibrium in the labor market requires $L_t = l_t$, where l_t is the labor supplied by households.

Optimizing behavior by households gives rise to the intratemporal condition,

$$(35) \quad \frac{w_t}{P_t} = \frac{\chi_0 l_t^\chi}{(c_t - b c_{t-1})^{-\gamma}},$$

³¹ Allowing a competitive capital market with free mobility of capital across sectors or considering industry-specific labor markets as well as firm-specific capital does not alter our basic findings presented in the benchmark model section.

and the intertemporal Euler equation,

$$(36) \quad (c_t - bC_{t-1})^{-\gamma} = \beta \exp(i_t) E_t (c_{t+1} - bC_t)^{-\gamma} P_t / P_{t+1},$$

where i_t denotes the continuously compounded interest rate on the riskless one-period nominal bond. There is no investment in physical capital in the model.

There is a government in the economy, which levies lump-sum taxes, G_t , on households and destroys the resources it collects. The aggregate resource constraint implies that

$$(37) \quad Y_t = C_t + G_t,$$

where $C_t = c_t$, with c_t denoting the consumption of the representative household. Government consumption follows an exogenous AR(1) process:

$$(38) \quad \log G_t = \rho_G \log G_{t-1} + \varepsilon_t^G,$$

where ε_t^G denotes an i.i.d. government consumption shock with mean zero and variance σ_G^2 .

Finally, there is a monetary authority in the economy that sets the one-period nominal interest rate according to a Taylor-type policy rule:

$$(39) \quad i_t = \rho_i i_{t-1} + (1 - \rho_i) [i^* + g_y (Y_t - Y_{t-1}) + g_\pi \pi_t] + \varepsilon_t^i,$$

where i^* denotes the steady-state nominal interest rate, π_t denotes the inflation rate (equal to $P_t/P_{t-1} - 1$), ε_t^i denotes an i.i.d. stochastic monetary policy shock with mean zero and variance σ_i^2 , and ρ_i , g_y , and g_π are parameters. Of course, the steady-state inflation rate, π^* , in this economy must satisfy $1 + \pi^* = \beta \exp(i^*)$.

As noted above, households have access to a long-term default-free nominal consol that pays one dollar every period in perpetuity. The nominal consol's price, $p_t^{(\infty)}$, thus satisfies

$$(40) \quad p_t^{(\infty)} = 1 + E_t m_{t+1} p_{t+1}^{(\infty)},$$

where $m_{t+1} \equiv m_{t,t+1}$ is the representative household's stochastic discount factor. We define the risk-neutral consol price, $p_t^{(\infty)rn}$, to be

$$(41) \quad p_t^{(\infty)rn} = 1 + \exp(-i_t^{(1)}) E_t p_{t+1}^{(\infty)rn},$$

and the implied term premium is then given by

$$(42) \quad \log \left(\frac{p_t^{(\infty)}}{p_t^{(\infty)} - 1} \right) - \log \left(\frac{p_t^{(\infty)rn}}{p_t^{(\infty)rn} - 1} \right).$$

Note that under our baseline parameterization, the consol in our model has a duration of about 25 years.

This completes the specification of the benchmark model referred to in the text. In computing impulse response functions, we use the parameter values as specified in Table A1. A technical issue that arises in solving the model above is the relatively large number of state variables, eight in all: C_{t-1} , A_{t-1} , G_{t-1} , i_{t-1} , Δ_{t-1} , plus the three shocks ε_t^A , ε_t^G , ε_t^i .³² Because of dauntingly high dimensionality, value-function iteration-based methods, such as projection methods (or, even worse, discretization methods), are computationally intractable. We instead solve the model above using a standard macroeconomic

³² The number of state variables can be reduced a bit by noting that G_t and A_t are sufficient to incorporate all of the information from G_{t-1} , A_{t-1} , ε_t^G , and ε_t^A , but the basic point remains valid—namely, that the number of state variables in the model is large from a computational point of view.

technique that approximates the model's solution around the nonstochastic steady state—a so-called perturbation method.

As discussed in the text, a first-order approximation (i.e., a linearization or log-linearization) of the model around the steady state eliminates the term premium from the model entirely, because equations (40) and (41) are identical in the first-order approximation. A second-order approximation produces a nonzero but constant term premium (a sum of the variances σ_A^2 , σ_G^2 , and σ_i^2). Because our interest in this paper is not just in the level of the term premium but also in its variation over time, we must compute a third-order approximation to the solution of the model around the nonstochastic steady state. We do so using the n th-order perturbation AIM algorithm of Swanson, Anderson, and Levin (2006). This algorithm requires that the equations of the model be put into a recursive form, which for the model above is fairly standard. The most difficult equation is (31), which can be written in recursive form as

$$(43) \quad \left(\frac{P_t^*(f)}{P_t} \right)^{1 + \frac{\alpha}{1-\alpha} \frac{1+\theta}{\theta}} = \frac{Z_{n,t}}{Z_{d,t}}$$

$$(44) \quad z_{d,t} = Y_t (C_t - bC_{t-1})^{-\gamma} + \beta \xi E_t \pi_{t+1}^{1/\theta} z_{d,t+1}$$

$$(45) \quad z_{n,t} = (1 + \theta) \frac{\chi_0}{1 - \alpha} L_t^{1+\chi} \Delta_t^{-1/(1-\alpha)} + \beta \xi E_t \pi_{t+1}^{\frac{1+\theta}{\theta} \frac{1}{1-\alpha}} z_{n,t+1}$$

The computational time required to solve our model to the third order is minimal—no more than about 10 seconds on a standard laptop computer.

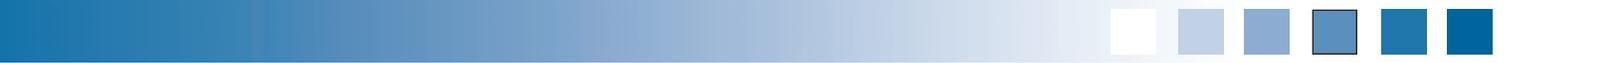
Computing impulse responses for this model is actually simpler than the use of a third-order approximation might suggest. We are interested in the responses of output and the term premium to an exogenous shock to ε_t^A , ε_t^G , or ε_t^i . For output, we plot the standard first-order (i.e., log-linear) responses of output to each shock. For small shocks, such as those of the size we are considering here (1 percent), these responses are highly accurate. For the term premium, of course, the first- and second-order responses of that variable to each shock would be identically zero, so we plot the third-order responses of that variable. These third-order terms are all of the form $\sigma_Z^2 X$, where $Z \in \{A, G, i\}$ and X is one of the state variables of the model,³³ so if we plug in the values of σ_A^2 , σ_G^2 , and σ_i^2 given in Table A1, these terms are linear as well, which makes them easy to plot.

Table A1

Benchmark Model Parameter Values

α	0.3	ρ_A	0.9
β	0.99	ρ_G	0.9
θ	0.2	ρ_i	0.7
ξ	0.75	σ_A^2	0.01 ²
γ	2	σ_G^2	0.004 ²
χ_0	$(1 - b)^{-\gamma}$	σ_i^2	0.004 ²
χ	1.5	\bar{K}	1
b	0.66	π^*	0
g_x	2		
g_y	0.5		

³³ In perturbation analysis, stochastic shocks of the model are given an auxiliary “scaling” parameter, so these shocks are third-order in a rigorous sense. See Swanson, Anderson, and Levin (2006) for details.



Commentary

John H. Cochrane

The paper by Glenn Rudebusch, Brian Sack, and Eric Swanson (2007) is an impressive survey of several literatures concerned with monetary economics and interest rates. It is well done, so I think the best thing for me to do is to highlight what I think are the central points and to give my views on those points.

WHAT “CONUNDRUM”?

Figure 1 presents the federal funds rate and 1- to 15-year forward rates through the past two recessions. This comparison lets us easily consider to what extent the recent behavior of long-term forward rates represents an unusual experience or not.

My first reaction to Figure 1 is that the patterns are strikingly similar. Short-term yields and forwards decline, spreads widen, and then yields and forwards recover as spreads tighten again. In both episodes there is a little blip on the way down in which long-term yields and forwards rise much more than short-term ones, despite no movement in the funds rate (late 1992 and 2002). In both episodes there is an event on the way up in which all yields and forwards increase sharply (late 1994 and 2004).

The main difference between the two episodes is that the rise in the federal funds rate in 2004-06 is much smoother and more predictable and long-term forward rates, in particular the 10-year rate,

falls while the funds rate is rising. Though long-term forwards decline overall in both recoveries (1994-96 and 2004-06), the earlier experience includes a blip up in all rates through 1995, which is later reversed. This experience is missing in the second period. This remaining difference is Greenspan’s “conundrum.”

The difference is already small. Furthermore, because the rise in the funds rate was much steadier in the later episode, the behavior of *market* rates relative to the funds rate (which reflects different behavior by the Fed) is even less different across the two episodes than the overall behavior of interest rates. If one regards long-term rates as dynamically driven by the federal funds rate, it’s not obvious that there is *any* difference in the behavior of markets.

Even if the later period is different, why is it puzzling? First, long forwards *should* fall when the Fed tightens. This is exactly how the world is *supposed* to work. Tighter policy now means lower inflation later, and thus *lower* nominal rates in 10 years. There is no model or estimate anywhere in which the Fed can raise real rates for 10 years without reducing inflation. Prices are not that sticky! In 1994, the opposite nearly one-for-one *rise* of long forwards with rises in the federal funds rate was viewed as a conundrum for just this reason. The main, somewhat convoluted, story used to explain the 1994 events is that an interest rate rise communicates bad news about inflation from the Fed to the markets—information

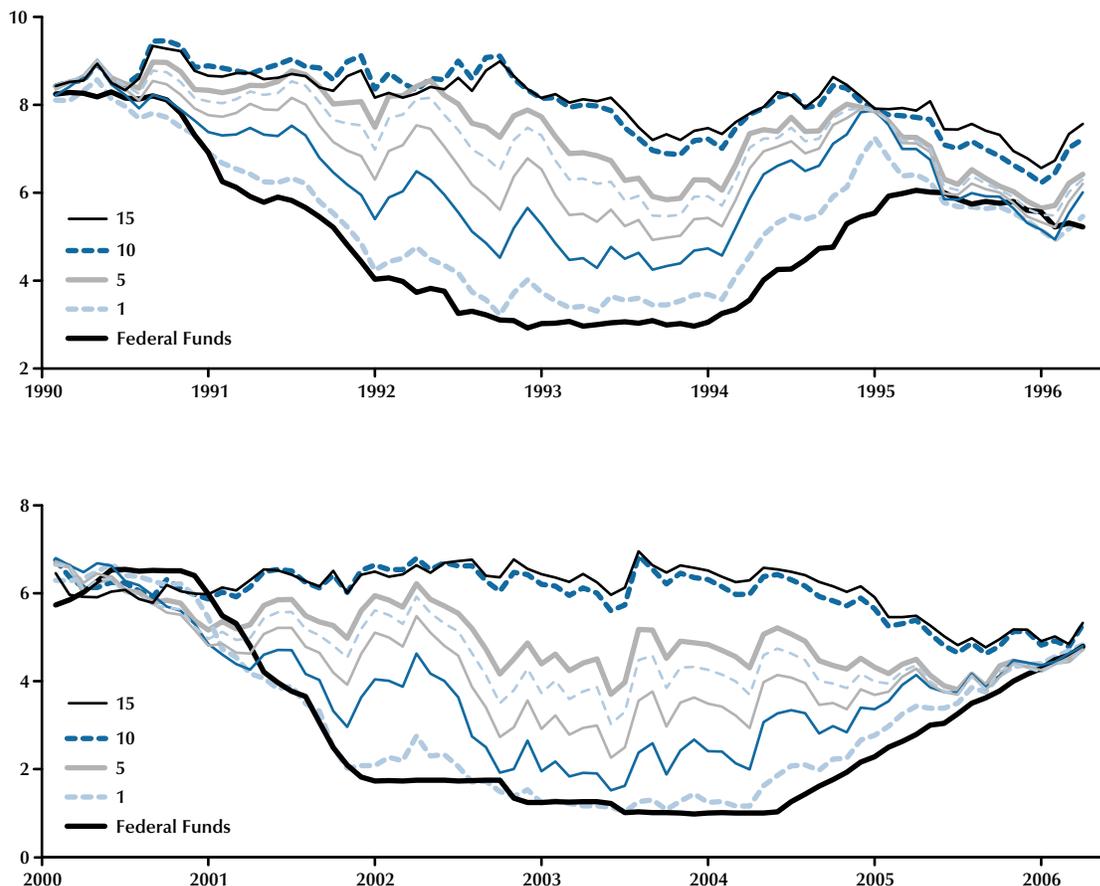
John H. Cochrane is a professor of finance at the Graduate School of Business at the University of Chicago and a research associate at the National Bureau of Economic Research. The author acknowledges research support from the Center for Research and Security Prices and from a National Science Foundation grant administered by the National Bureau of Economic Research. A draft of this paper with color graphics is available at <http://faculty.chicagosb.edu/john.cochrane/research/Papers/>.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 271-82.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Figure 1

Federal Funds Rate and Selected 1- to 15-Year Forward Rates Through Two Recessions



SOURCE: Forward-rate data are from Gürkaynak, Sack, and Wright (2006).

that for some reason the markets did not already have, of course. Greenspan himself echoed this view in 1994¹:

In early February, we thought long-term rates would move a little higher as we tightened. The sharp jump in [long] rates that occurred appeared to reflect the dramatic rise in market expectations of economic growth and associated concerns about possible inflation pressures.

Of course, this is a simplistic discussion. A tightening has to be unanticipated in order for it

to lower forward rates through this channel, and evidence from other sources, such as the Treasury inflation-protected securities mentioned by Chairman Greenspan or foreign interest rates, also bears on the issue. Still, where did anyone get the idea that monetary policy should control long-term rates and that it is puzzling if long-term rates do not “respond” positively to tightening? The natural benchmark predicts exactly the opposite, if any, effect.

Second, to the extent that the decline in forward rates represents a cyclical or secular decline in term premia, that decline also is perfectly natural. Term premia, like all risk premia, *should*

¹ I owe the quote to Gallmeyer et al. (2007).

decline as we come out of recessions, and have done so in every past recession. Even negative term premia are not a puzzle—they *should* be negative. In a world with stable inflation, interest rate variation comes from variation in *real* rates; and, in such a world, long-term bonds are safer investments for long-term investors. Rolling over *short-term* bonds runs the “reinvestment risk” that short-term (real) rates will change, so short-term bonds should bear the burden of any bond risk premia. We expect only a *positive* term premium in a world with unstable inflation and relatively constant real rates, such as the 1970s. Because short-term rates adapt quickly to inflation changes, rolling over short-term bonds has less risk to a long-term investor than does buying only long-term bonds in this environment.

In sum, were I a Fed Chairman testifying to Congress with the plots of Figure 1 in hand, I would be tempted to point out that, far from a “conundrum,” the world is finally behaving exactly the way it should—and so is the central bank. The increased transparency and predictability of operating procedures, seen in the steadiness of the rise in funds rates in 2004-06 versus the less predictable rise in 1994-96, has communicated to the markets the Fed’s steadfastness in controlling inflation. We are moving to the sensible world of negative risk premia, which is exactly what we should see once markets understand that inflation is vanquished forever. The conquest of inflation has removed an unnecessary risk premium for long-run investors and issuers of long-dated nominal bonds. I don’t necessarily believe all this, of course, but it would be awfully tempting to make this argument were I defending the Fed’s actions before a congressional committee. The “conundrum” is Greenspan: Why did he say anything else?

Finally, it is academics’ job to remind policy debaters of basic economics, so I think we should pounce anytime somebody says something like “[the] decline in the term premium...is financially stimulative and argues for greater monetary policy restraint.” Every price reflects both supply and demand. Low interest rates can reflect a lack of good investment projects as easily as they can reflect an abundance of savings. To take a local

example, low housing prices in East St. Louis do not seem to be particularly “stimulative.”

DECOMPOSING THE YIELD CURVE

My grumpy comments about “conundrum” and the “stimulative” effects of low prices notwithstanding, this episode *does* highlight the importance of splitting the yield curve into expected future rates and risk premia and of understanding the dynamic structure of risk premia and their macroeconomic underpinnings. Here Rudebusch, Sack, and Swanson provide a very nice summary of the state of the art.

I think the bottom line is that we know less than we think about this decomposition and far less than the pronouncements in policymakers’ quotes imply. The paper can be read as a comprehensive survey of one failure after another. Here, let me give two quick, and I hope memorable, points in this litany of ignorance.

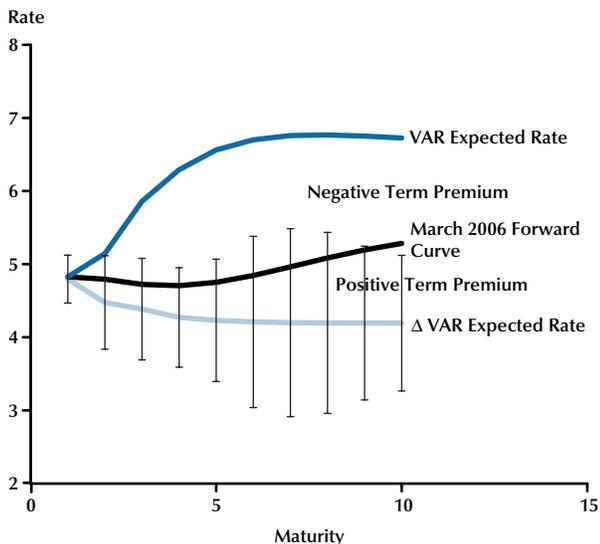
Levels, Differences, and Standard Errors

I learned two important lessons while Monika Piazzesi and I (2006) investigated this kind of decomposition. First, *how you specify trends, cointegration, etc.—which the data say very little about—is overwhelmingly the most important issue in driving the decomposition of the long-maturity end of the yield curve*. Second, *the standard errors are very large*. For these reasons alone, any statements decomposing the recent experience of forward rates into changes in expected interest rates versus declining term premia are subject to huge uncertainty.

To see this point, let’s try the simplest approach to decomposing the yield curve. I run a vector autoregression (VAR) of five forward rates on their lags (I use the Fama-Bliss data available from the Center for Research and Security Prices, and I use a three-month moving average of forward rates on the right-hand side, which Piazzesi and I (2005) find improves forecasts by mitigating measurement error):

Figure 2

March 2006 Forward Curve and Expected One-Year Rates from a VAR in Levels and from a Cointegrated VAR



NOTE: The one-standard-error bars are computed from a direct regression forecast, $X_t'cov(\hat{\beta})X_t$, using Hansen-Hodrick correction for serial correlation due to overlap.

$$\begin{bmatrix} y_{t+1}^{(1)} \\ f_{t+1}^{(2)} \\ \vdots \\ f_{t+1}^{(5)} \end{bmatrix} = A + B \begin{bmatrix} y_t^{(1)} \\ f_t^{(2)} \\ \vdots \\ f_t^{(5)} \end{bmatrix} + \varepsilon_{t+1}$$

where

$f_t^{(n)}$ = forward at time t for loans from $t + n - 1$ to $t + n$

$y_t^{(1)}$ = one-year rate at time t .

We can use this VAR to generate forecasts at each date of future one-year rates, leaving (“estimating”) the term premium as a residual,

$$f_t^{(n)} = E_t \left(y_{t+n}^{(1)} \right) + rp f_t^{(n)}.$$

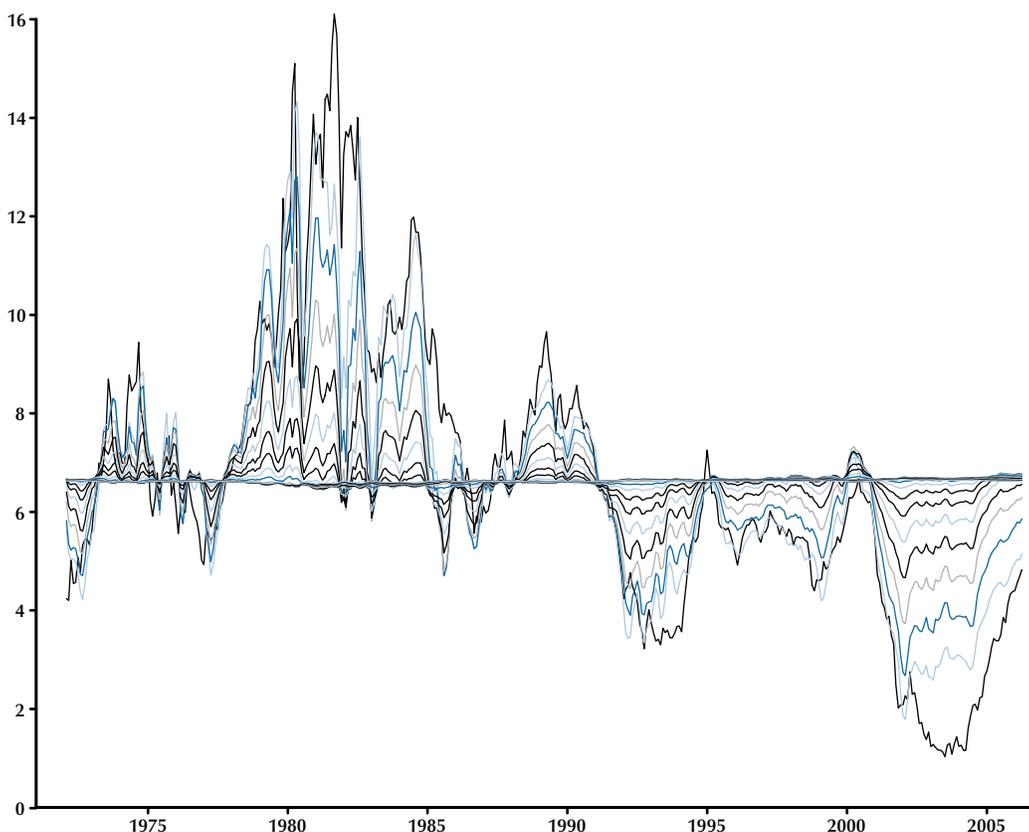
You don’t *have* to estimate fancy term-structure models to decompose the yield curve into expected interest rates and a risk premium.

Figure 2 presents the results, evaluated on March 2006 (the last point in my data sample), in the line labeled “VAR expected rate.” The line captures a lot of common opinion: It says that interest rates are expected to rise gently over the next few years, leaving a negative term premium, which is puzzling until you think through the economics of long-term bond investing in a low-inflation world. This kind of decomposition also says that much of the recent decline in forward rates comes from the term premium rather than changes in expected long-term rates.

This all seems very sensible. However, Figure 3 examines the same calculation over a longer time interval. The lines represent, at each date, expected one-year rates one, two, three, etc., years in the future: that is, for $E_t(y_{t+k}^{(1)})$ for $k = 1, 2, 3, \dots$ at each t . The graph dramatically makes the point that long-horizon expected one-year rates calculated by this method simply reflect reversion to the mean. The 6.25 percent asymptote in Figure 2 represented no specially sophisticated regression forecast; it was simply the sample interest rate.

There is nothing logically or econometrically wrong with this conclusion, but do we really believe it? For example, in 1980, this decomposition says that everyone knew interest rates would decline from 16 percent back to an unconditional mean of a bit over 6 percent, and rather rapidly, so the then-flat yield curves represented very large risk premia for holding long-term bonds. But did people really believe inflation would be tamed, or did perhaps the flat yield curves of the time really represent a good chance that inflation would re-emerge? Similarly, perhaps the sample mean is now too high an estimate. Our data come from inflation and its conquest. Perhaps it is sensible now to think a “structural shift” has happened, so the long-run mean should be a good deal less than 6.25 percent.

As an alternative, let us try a forecast that ignores this “level” information. On a statistical basis, forward rates are clearly best modeled by a single common trend that has a root that is near if not equal to 1 and stationary spreads around

Figure 3**Multiperiod One-Year-Rate Forecasts from a VAR in Levels**

NOTE: Current one-year rate and expectations of one-year rates one, two, three, etc., years in the future, calculated by a simple VAR.

that trend. I estimate a VAR imposing that restriction:

$$\begin{bmatrix} y_{t+1}^{(1)} - y_t^{(1)} \\ f_{t+1}^{(2)} - y_{t+1}^{(1)} \\ \vdots \\ f_{t+1}^{(5)} - y_{t+1}^{(1)} \end{bmatrix} = A + B \begin{bmatrix} f_t^{(2)} - y_t^{(1)} \\ \vdots \\ f_t^{(5)} - y_t^{(1)} \end{bmatrix} + \varepsilon_{t+1}.$$

This is equivalent to simply running forecasting regressions that set to zero a coefficient on the level of interest rates:

$$\text{Before : } f_{t+1}^{(n)} = a^{(n)} + B_{n,1} \times y_t^{(1)} + B_{n,2} [f_t^{(2)} - y_t^{(1)}] + B_{n,3} [f_t^{(3)} - y_t^{(1)}] + \dots$$

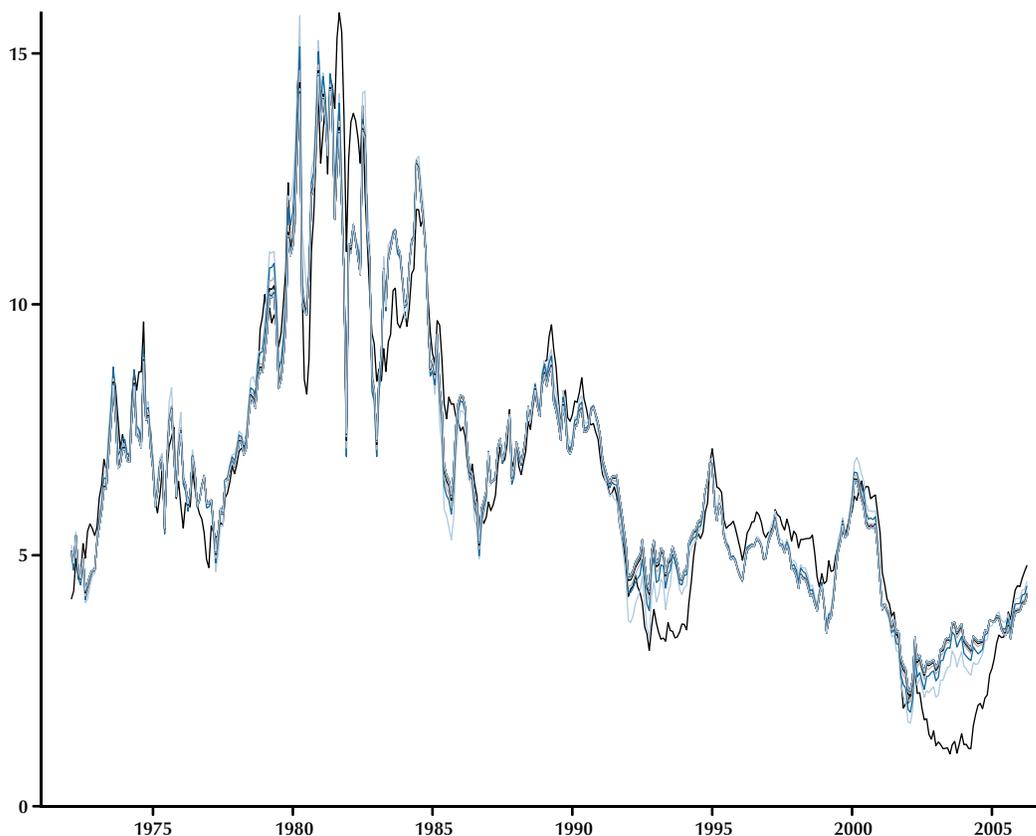
$$\text{now : } f_{t+1}^{(n)} = a^{(n)} + 0 \times y_t^{(1)} + B_{n,2} [f_t^{(2)} - y_t^{(1)}] + B_{n,3} [f_t^{(3)} - y_t^{(1)}] + \dots$$

Intuitively, we still allow information such as “the yield curve is upward sloping” to forecast interest rate changes. We ignore information such as “interest rates are low” to tell us interest rates will rise.

Figure 4 presents expected one-year rates over time by this method. You can see the *huge* difference. One-year rates are certainly not being forecast to revert to a constant unconditional mean! In particular, the flat yield curves of the 1980s are not now interpreted to reveal huge risk premia plus expected declines in interest rates. This is

Figure 4

Multiperiod One-Year-Rate Forecasts from a VAR in Levels



NOTE: Current one-year rate and expectations of one-year rates one, two, three, etc., years in the future, calculated by a VAR that imposes a single common trend in forward rates.

not a pure random-walk model, and there is still some forecastability left. For example, the steeply upward-sloping yield curves of 2003-04 do forecast substantial rises in short rates.

Figure 2 includes the March 2006 one-year rate forecast from this method, in the line labeled “ Δ VAR Expected Rate.” This is also a sensible forecast. Because we no longer use the information that the current one-year rate is slightly below its sample mean, we are left only with slope information. The unusually flat slope of the forward curve means, in this forecast, that interest rates will *decline* somewhat, so that the term premium is still somewhat positive. However, Figure 2 shows that long-term interest rate forecasts by

this method have been rising in recent years; so, the decline in forward rates since 2004 is attributed even more to declining term premia by this method than by the VAR in levels.

Can statistics help us? Alas, no. Testing for unit roots, cointegration, etc., and imposing the resulting structure on the analysis is not fruitful. One naturally wants to think about “structural shifts,” changing means, and so forth, and these will be even more imprecisely estimated in now-shorter samples. It is certainly true that the dominant root of a persistent set of variables is estimated with downward bias, so the actual reversion to the mean is slower than the VAR in levels indicates, but whether that mean makes

any sense in the first place is not something statistics can really help us with.

Can fancier models help us? In particular, most of the term-structure literature does not look at simple VAR forecasts. Instead, it estimates the parameters of “affine models.” To think about what these can do, it’s useful to have a specific example in front of us, so here is the Ang and Piazzesi (2003) model that we use in Cochrane and Piazzesi (2005 and 2006). A vector of state variables X_t follows an AR(1) process; the stochastic discount factor is an exponential function of the state variables, with “market prices of risk” (loadings of M on shocks to X) that also depend on the state variables,

$$(1) \quad X_t = \mu + \phi X_{t-1} + \Sigma \varepsilon_t$$

$$(2) \quad M_t = \exp\left(-\delta'_0 X_t - \frac{1}{2} \lambda'_t \lambda_t - \lambda'_t \varepsilon_t\right)$$

$$\lambda_t = \lambda_0 + \lambda'_1 X_t.$$

Assuming the shocks ε are i.i.d. normal with unit variance, we can find this model’s prediction for bond prices,

$$(3) \quad P_t^{(n)} = E_t(M_{t+1} M_{t+2} \dots M_{t+n})$$

$$(4) \quad p_t^{(n)} = A_n + B_n X_t,$$

and then yields and forward rates. Inverting (4), we can reveal the “state variables” from bond prices, yields, or forward rates. Thus, this model becomes a structured factor model in which a large collection of prices, yields, or forward rates are described in terms of a few linear combinations of those same prices, yields, or forward rates.

But, underlying the whole thing, we see a VAR(1) in yields, prices, or forward rates—just as we have been estimating all along! Thus the *only* way the affine model can give us *any* different answers from those of the ordinary least squares—estimated VARs above is if the structure of market prices of risk means that we use information in the cross-section of bond prices to infer something about the dynamics. In general, this is not the case. In Cochrane and Piazzesi (2005) we show how to *construct* market prices of risk, λ , from a

given discrete-time VAR(1) to turn it into an affine model. Thus, in general, the affine model lives on top of a VAR estimate of long-term forward rates and adds *nothing* to it. (There remains the possibility that by *restricting* or *modeling* market prices of risk, λ , in sensible ways, one obtains information about the VAR (1), and this is the point of our 2006 paper. But this is not [yet] a common idea, and its success lies in the believability of a priori restrictions on λ .)

In addition, once we have settled on a specification, we have to wonder how much sampling uncertainty in estimating the parameters translates into uncertainty about the forecasts. To address this question in a simple and transparent way, I run direct forecasting regressions,

$$y_{t+k}^{(1)} - y_t^{(1)} =$$

$$\left[1 \quad f_t^{(2)} - y_t^{(1)} \quad f_t^{(3)} - y_t^{(1)} \quad \dots \quad f_t^{(5)} - y_t^{(1)}\right] \beta_k + \varepsilon_{t+k},$$

$$y_{t+k}^{(1)} - y_t^{(1)} = X_t' \beta_k + \varepsilon_{t+k},$$

where the second equation defines notation. I find the covariance matrix of $\hat{\beta}_k$, including a Hansen-Hodrick correction for serial correlation due to overlapping data, and then I calculate the error as

$$\sigma_t^2 \left[\hat{E}_t \left(y_{t+k}^{(1)} - y_t^{(1)} \right) \right] = X_t' \text{cov}(\hat{\beta}_k, \hat{\beta}_k') X_t.$$

This is the error in the measurement of expected interest rates due to sampling uncertainty in the coefficients that comprise the regression forecast. It is not the forecast error—that is, it is not a measure of how large $\sigma^2(\varepsilon_{t+k})$ is. The one-standard-error bars in Figure 2 present this calculation. The term premium is not statistically significant, and the large difference between the two specifications is barely two standard errors. The Hansen-Hodrick correction for serial correlation is undoubtedly optimistic—at the right-hand end of the graph we’re forecasting interest rates 10 years ahead in 45 years of data—so the true sampling uncertainty is undoubtedly a good deal larger.

Now, understanding that large roots and common trends, which often must be specified a priori, are crucial to long-term forecasts and that long-run forecasts are subject to enormous sampling uncertainty is not news. However, as I read it, this sensitivity is not at all considered by the

literature that uses affine models to compute long-term yield-curve decompositions. We are usually treated only to one estimate based on one a priori specification, usually in levels, and usually with no measure of the huge sampling uncertainty. Needless to say, the usual habit of estimating 10-year interest rate forecasts by extrapolating models fit to *weekly* or monthly data is no help, and possibly a hindrance. The 520th power of a matrix is a difficult object to estimate.

In sum, when a policymaker says something that sounds definite, such as “long-run forward rates have declined, while interest rate expectations have remained constant, so risk premia have declined,” he is really guessing, and we really have no idea whether this is a fact.

Measuring Risk Premia

We also know a good deal less about long-term risk premia than we think we do. Quotes such as those at the beginning of the paper suggest that risk premia are well *measured* if perhaps poorly understood. Nothing of the sort is true. We may have a decent handle on *one-year* risk premia, as surveyed in the paper and the subject of my next set of comments, but the 10-year forward-rate premium reflects not only *this* year’s expected excess bond returns, but this year’s expectations of *next* year’s expected returns, and so on and so forth. If you like equations, an easy one in which to see this point is

$$(5) \quad y_t^{(n)} = \frac{1}{n} \left[y_t^{(1)} + E_t(y_{t+1}^{(1)}) + E_t(y_{t+2}^{(1)}) + \dots + E_t(y_{t+n-1}^{(1)}) \right] + \left[E_t(rx_{t+1}^{(n)}) + E_t(rx_{t+2}^{(n-1)}) + \dots + E_t(rx_{t+n-1}^{(2)}) \right],$$

where $y_t^{(1)}$ = one-year yield and $rx_{t+1}^{(n)} = r_{t+1}^{(n)} - y_t^{(1)}$ = excess returns. The first term is the expectations hypothesis. The second term is the risk premium, and you see that the risk premium depends on *future* expected excess returns, not just on *current* expected excess returns.

Now, if expected excess returns lived off in their own space, moving away in response to shocks and then recovering without relation to the rest of the yield curve, then, yes, there would

be one “risk premium” that accounts for expected excess returns, as well as long-horizon forward and yield-curve risk premia. Alas, this is not the case. Today’s level, slope, and curvature have strong power to forecast next year’s expected excess returns. (Characterizing these dynamics is a major point of Cochrane and Piazzesi, 2006.) We can easily be in a situation that this year’s expected excess return, $E_t rx_{t+1}$, is large and positive, while future expected excess returns, $E_t(rx_{t+k})$, are strongly negative, so the risk premium in the yield curve can be negative as well. The one-year expected excess return can be positive while the 10-year forward rate is below its corresponding expected one-year rate. It is precisely by such differences in expected *future* risk premia that the two decompositions shown in Figure 2 can produce forward rate premia of different signs, despite the same initial return risk premium.

In sum, there is no single “risk premium.” There is a full-term structure of return risk premia, which moves over time in interesting and still poorly measured ways. Sure statements that risk premia have moved down over time do not reflect any solid and independent measurement.

FORECASTING, TERM PREMIA, AND MACROECONOMICS

One of the major contributions of the Rudebusch, Sack, and Swanson paper is the empirical work linking bond risk premia and macroeconomics. By restating the points in my own way and slightly disagreeing with some conclusions, I think I can usefully highlight this important part of the paper.

Naturally, I like the Cochrane and Piazzesi (2005) measurement of the risk premium, so I’ll focus my comments on that paper. Briefly, we noticed that regressions of excess returns on ex ante forward rates follow a nearly exact one-factor structure: That is, that regressions

$$rx_{t+1}^{(n)} = \alpha_n + \beta_{n,1} y_t^{(1)} + \beta_{n,2} f_t^{(2)} + \dots + \beta_{n,5} f_t^{(5)} + \varepsilon_{t+1}^{(n)}$$

almost exactly follow

$$\begin{aligned} rX_{t+1}^{(n)} &= b_n (\gamma_0 + \gamma_1 y_t^{(1)} + \gamma_2 f_t^{(2)} + \dots + \gamma_5 f_t^{(5)}) + \varepsilon_{t+1}^{(n)} \\ &= b_n (\gamma' f_t) + \varepsilon_{t+1}^{(n)}, \end{aligned}$$

where the last equality defines notation. A single “return-forecasting factor,” $\gamma' f_t$, describes expected excess returns of all maturities. Longer-maturity bonds’ expected excess returns move more, and shorter-maturity bonds’ expected excess returns move less, but they all move in lockstep. Thus, we estimate the common “return-forecasting factor” by running a single regression of average (across maturity) returns on all forward rates:

$$\begin{aligned} \bar{rX}_{t+1} &= \frac{1}{4} \sum_{n=2}^5 rX_{t+1}^{(n)} \\ &= \gamma_0 + \gamma_1 y_t^{(1)} + \gamma_2 f_t^{(2)} + \dots + \gamma_5 f_t^{(5)} + \varepsilon_{t+1}, \end{aligned}$$

where the first equality defines notation. Sensitive to “levels” issues, we obtain nearly identical results by ruling out a level effect:

$$\begin{aligned} \bar{rX}_{t+1}^n &= \gamma_0 + \gamma_2 (f_t^{(2)} - y_t^{(1)}) \\ &\quad + \gamma_3 (f_t^{(3)} - y_t^{(1)}) + \dots + \gamma_5 (f_t^{(5)} - y_t^{(1)}) + \varepsilon_{t+1}. \end{aligned}$$

The coefficients γ have a pretty tent shape. This measure of bond risk premia values *curvature* in the forward curve, not *slope* in the forward curve.

Figure 5 shows how this works and the connection between macroeconomics and bond risk premia: In January 2002 (shown by the first vertical lines in panels A, B, and D), the recession and interest rates have just finished their stage of steep decline, as seen also in the unemployment rate (panel D). The forward curve is upward sloping, but it is also very curved (panel C). The *curved* forward rate, interacting with the tent-shaped γ , is the sign of risk premia. This means (statistically) that the upward slope will not be soon matched by rises in interest rates, so the greater yields on long-term bonds are (risky) profit for investors. The risk premium (panel B) is very high. In fact, this prediction is borne out: Interest rates do *not* rise for several years, so investors who bought long-term bonds in January 2002 profited handsomely for a few years.

By contrast, consider January 2004. Now, the forward curve still *slopes up* substantially (panel

C), but it is no longer particularly *curved*, so the tent-shaped γ coefficients no longer predict much of a risk premium (panel B). Now, the upward-sloping yield curve *does* signal rises in interest rates; the expectations hypothesis is working; returns on long-term bonds will be no higher (on average) than those on short-term bonds. Again, this prediction is borne out. This time, interest rates do rise. This is a repeated statistical pattern, working the same way in many previous recessions.

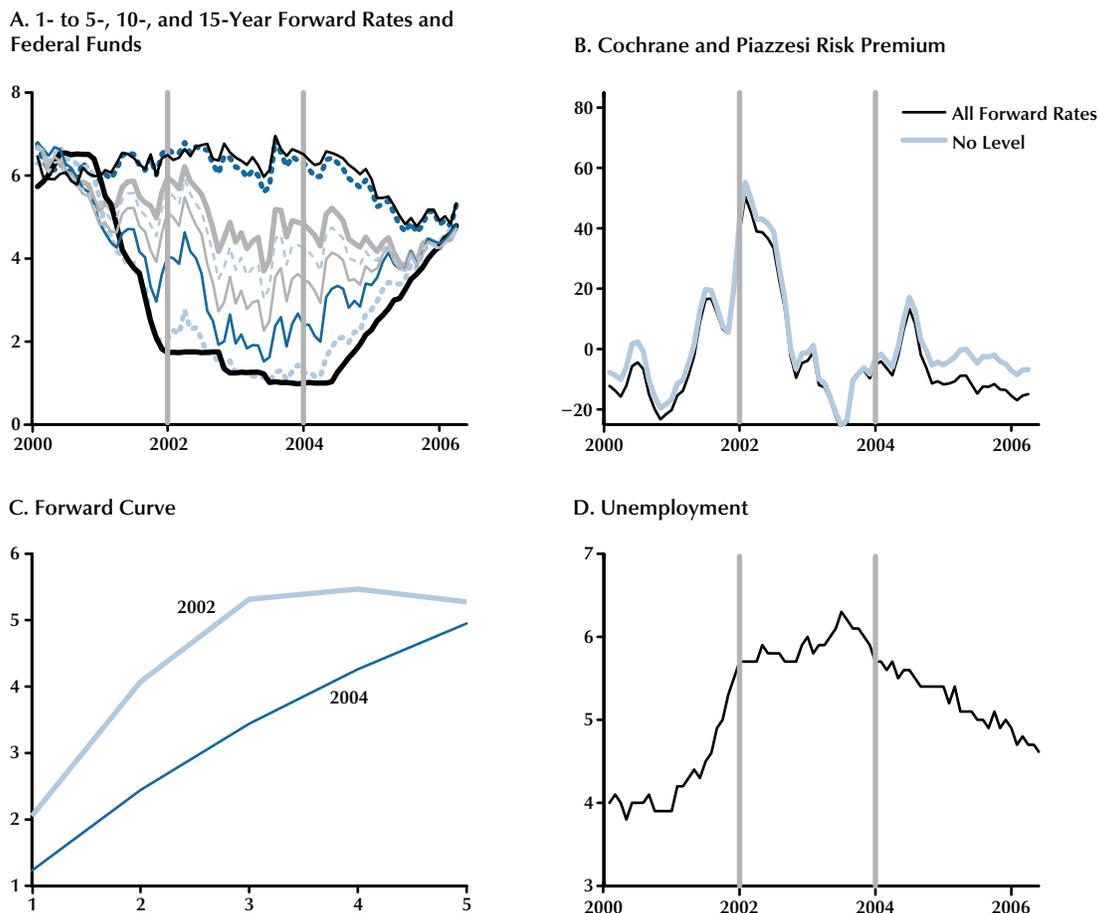
Having digested what term-premium forecasts are and how they work, we see that the graphs show several patterns seen in more-formal regressions. First, *the term premium* ($\gamma' f$ here, as well as other measures in the paper) *drives out slope variables for forecasting bond excess returns*.

Previously, Fama and Bliss (1987), Campbell and Shiller (1991), and others found that measures of the term-structure slope forecast excess returns. Yes, we see the slope is high in 2002, when long-term bond holders turn out to make money. But it is also high in 2004 when they don’t. When you put the slope and the curvature of the forward rate together in a multiple regression, the curvature measured by γ wins out. The slope seemed to forecast bond returns because it was correlated with the curve measure. (See, for example, Table A3 of the appendix to Cochrane and Piazzesi, 2005.)

Second, *the term premium is high in the depths of a recession*. In Figure 5, this is measured by the association of the term premium (panel B) with unemployment (panel D). The association is even stronger in previous recessions. In macroeconomic terms (that’s why we’re here), this is natural. The risk premium is high at the early stage of a recession, a time in which investors don’t want to hold risk of any kind. Stock prices are low, predicting higher-than-average stock returns; interest rates are low relative to foreign interest rates, predicting high returns for holding exchange-rate risk. By January 2004, however, the recession is over, the period of growth and rising interest rates has set in, and everybody knows it. It’s not a surprise that the premium for holding risk during recessions has vanished.

Figure 5

Macroeconomics and the Yield Curve



NOTE: Vertical lines mark interesting dates. Panel A: federal funds rate and 1- to 15-year forward rates through the previous recession. Panel B: Cochrane and Piazzesi (2005) measures of the bond risk premium. Panel C: forward curves on the two indicated dates. Panel D: unemployment rate.

SOURCE: Data for yields and forward rates past a five-year maturity are from Gürkaynak, Sack, and Wright (2006).

Third, and the major point of the authors' paper (as I see it), the *slope* of the yield curve drives out the *risk premium* for forecasting 1-year gross domestic product (GDP) growth. We see this in panel B of Figure 5: The *risk premium* is high in 2002, when GDP is not about to grow. The risk premium is low in 2004, when GDP is about to grow. The slope is high in both times. Thus, the *slope* carries GDP forecast power, and the slope, purged of its correlation with the risk premium, forecasts GDP even better.

This point is made in their paper in the regression of Table 2, last column, which I take to be the central result:

$$y_{t+4} - y_t = 0.38(4.22) + 0.96(5.62)(exsp_t - exsp_{t-4}) - 0.59(-1.93)(tp_t - tp_{t-4}) + \epsilon_{t+4},$$

where y = GDP; $exsp$ is the expectations-hypothesis component of the 10-year rate; tp is the term premium component of the 10-year rate, as in (5); t -statistics are in parentheses; and the sample is from 1962 to 2005.

The authors (p. 261) say of this regression, “The coefficient on the risk-neutral expectations component of the yield curve slope [0.96] is now larger and more statistically significant than in any of the earlier specifications,” which is the same interpretation I gave in discussing the figure.

The authors also say (p. 261), “More importantly for this paper, we find that the estimated coefficient on the term premium is now negative and (marginally) statistically significant. According to these results, a decline in the term premium tends to be followed by *faster* GDP growth—the opposite sign of the relationship uncovered by previous empirical studies.” I read the evidence differently: Rather than accept a marginally significant coefficient with the wrong sign, it seems to me the right lesson is that the second coefficient is zero. The *slope* of the yield curve forecasts GDP growth, but not risk premia. The *curvature* of the forward curve measures risk premia, but not GDP growth. Risk premia are high precisely when we are not sure whether the recession is over.

Table 8 of Ang, Piazzesi, and Wei (2006) runs the same sort of regression. At a four-quarter horizon, they find

$$y_{t+4} - y_t = a + 1.15(5.00)EH_t - 0.47(0.30)RP_t + \varepsilon_{t+4},$$

where EH = expectations hypothesis; RP = risk premium in the 20-quarter term spread (i.e., the terms in (5), estimated from a macro-affine model); and t -statistics are in parentheses. In this slightly different specification, they confirm the huge significance of the expectations-hypothesis term, but find an insignificant contribution due to the risk-premium term.

This view dovetails with the other side of risk premia that Monika Piazzesi and I (2006) have recently started investigating. From the basic asset-pricing relation, $1 = E_t(M_{t+1}R_{t+1})$, and (2), we can write

$$(6) \quad E_t(rx_{t+1}^{(n)}) + \frac{1}{2}\sigma_t^2(rx_{t+1}^{(n)}) = cov(rx_{t+1}^{(n)}, \varepsilon'_{t+1})\lambda_t.$$

Note the absence of an n index on λ . The point of this equation is that expected excess returns on each bond must be earned in compensation for, and in proportion to, the covariance of that bond's return with macroeconomic shocks, ε . So

far, we have been talking about the left-hand side: What models or state variables drive variation over time in expected excess returns? Now, it's time to start working on the right-hand side: What are the shocks? Piazzesi and I find that the term premium is almost exactly earned entirely in compensation for shocks to the level of the term structure. The prices of risk, λ , corresponding to other term-structure shocks are essentially zero.

In particular, expected returns seem *not* to be earned in compensation for “slope” shocks. This finding lets us start to think about macroeconomics. Whatever the macroeconomic determinants of bond risk premia, they must be variables with “level” effects on the term structure. This observation quickly rules out monetary policy, whose shocks typically raise short rates while lowering or leaving unchanged long rates—a “slope” shock.

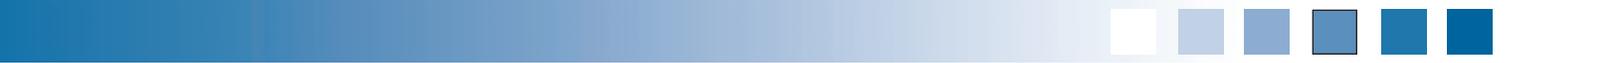
CONCLUDING COMMENT

In sum, I think we are headed to a view that slope movements in the yield curve, which are related to monetary policy, are also related to expected GDP growth, as seen in the usual impulse-response functions. But slope movements do not signal risk premia (left-hand side of (6)), nor does covariance with monetary policy shocks generate real risk premia (right-hand side of (6)). Term premia are large in the early phases of recessions, when it's not clear how long the recession will last; they are revealed by the curvature of the forward rate, and they are earned in compensation for macroeconomic risks that correspond to shocks in the level of the yield curve.

Of course, we have no economic models of any of these fascinating statistical regularities. This point is made clear in the brilliant survey of total failures that occupies a large part of the paper. First, what are the *macroeconomic* state variables that drive variation in expected returns? What exactly are the times and states of nature in which expected returns are high? Second, expected returns are generated by the covariance of returns with *macroeconomic* shocks. What are these macroeconomic shocks? These questions are, as ever, the Holy Grail of macro-finance.

REFERENCES

- Ang, Andrew and Piazzesi, Monika. "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables." *Journal of Monetary Economics*, July 2003, 50(5), pp. 745-87.
- Ang, Andrew; Piazzesi, Monika and Wei, Min. "What Does the Yield Curve Tell Us about GDP Growth?" *Journal of Econometrics*, March/April 2006, 131(1-2), pp. 359-403.
- Campbell, John Y. and Shiller, Robert J. "Yield Spreads and Interest Rate Movements: A Bird's Eye View." *Review of Economic Studies*, May 1991, 58(3), pp. 495-514.
- Cochrane, John H. and Piazzesi, Monika. "The Fed and Interest Rates—A High Frequency Identification." *American Economic Review*, May 2002, 92(2), pp. 90-95.
- Cochrane, John H. and Piazzesi, Monika. "Bond Risk Premia." *American Economic Review*, March 2005, 95(1), pp. 138-60.
- Cochrane, John H. and Piazzesi, Monika. "Decomposing the Yield Curve." Unpublished manuscript, University of Chicago, 2006.
- Fama, Eugene F. and Bliss, Robert R. "The Information in Long-Maturity Forward Rates." *American Economic Review*, September 1987, 77(4), pp. 680-92.
- Gallmeyer, Michael F.; Hollifield, Burton; Palomino, Francisco and Zin, Stanley F. "Arbitrage-Free Pricing with Dynamic Macroeconomics Models." *Federal Reserve Bank of St. Louis Review*, July/August 2007, 89(4), pp. 305-26.
- Greenspan, Alan. "Statement by Alan Greenspan, Chairman, Board of Governors of the Federal Reserve System, before the Committee on Banking, Housing, and Urban Affairs, May 27, 1994." *Federal Reserve Bulletin*, July 1994.
- Gürkaynak, Refet S.; Sack, Brian and Wright, Jonathan H. "The U.S. Treasury Yield Curve: 1961 to the Present." Unpublished manuscript, Board of Governors of the Federal Reserve System, 2006.
- Rudebusch, Glenn D.; Sack, Brian P. and Swanson, Eric T. "Macroeconomic Implications of Changes in the Term Premium." *Federal Reserve Bank of St. Louis Review*, July/August 2007, 89(4), pp. 241-69.



Long-Run Risks and Financial Markets

Ravi Bansal

The recently developed long-run risks asset pricing model shows that concerns about long-run expected growth and time-varying uncertainty (i.e., volatility) about future economic prospects drive asset prices. These two channels of economic risks can account for the risk premia and asset price fluctuations. In addition, the model can empirically account for the cross-sectional differences in asset returns. Hence, the long-run risks model provides a coherent and systematic framework for analyzing financial markets. (JEL G0, G00, G1, G10, G12)

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 283-99.

From the perspective of theoretical models, several key features of asset markets are puzzling. Among others, these puzzling features include the level of equity premium (see Mehra and Prescott, 1985), asset price volatility (see Shiller, 1981), and the large cross-sectional differences in average returns across equity portfolios such as value and growth portfolios. In bond and foreign exchange markets, the violations of the expectations hypothesis (see Fama and Bliss, 1987; Fama, 1984) and the ensuing return predictability is quantitatively difficult to explain. What risks and investor concerns can provide a unified explanation for these asset market facts? One potential explanation of all these anomalies is the long-run risks model developed in Bansal and Yaron (2004) (henceforth BY). In this model, fluctuations in the long-run growth prospects of the economy and the time-varying level of economic uncertainty (consumption or output volatility) drive financial markets. Recent work indicates that many of the asset prices anomalies are a natural outcome of these channels developed in BY. In this article I explain the key mechanisms in the BY model that enable it to account for the asset market puzzles.

In the BY model, the first economic channel relates to expected growth: Consumption and dividend growth rates contain a small long-run component in the mean. That is, current shocks to expected growth alter expectations about future economic growth not only for short horizons but also for the very long run. The second channel pertains to varying economic uncertainty: Conditional volatility of consumption is time varying. Fluctuations in consumption volatility lead to time variation in risk premia. Agents fear adverse movements in the long-run growth and volatility components because they lower equilibrium consumption, wealth, and asset prices. This makes holding equity quite risky, leading to high risk compensation in equity markets.

The preferences developed in Epstein and Zin (1989) play an important role in the long-run risks model. These preferences allow for a separation between risk aversion and intertemporal elasticity of substitution (IES) of investors: The magnitude of the risk aversion relative to the reciprocal of the IES determines whether agents prefer early or late resolution of uncertainty regarding the consumption path. In the BY model, agents prefer early resolution of uncertainty; that is, risk aver-

Ravi Bansal is a professor of finance at the Fuqua School of Business, Duke University. He thanks Dana Kiku, Tom Sargent, Ivan Shaliastovich, and Amir Yaron for comments.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Bansal

sion is larger than the reciprocal of the IES. This ensures that the compensation for long-run expected growth risk is positive. The resulting model has three distinct sources of risks that determine the risk premia: short-run, long-run, and consumption volatility risks. In the traditional power utility model, only the first risk source carries a distinct risk price and the other two risks have zero risk compensation. Separate risk compensation for shocks to consumption volatility and expected consumption growth is a novel feature of the BY model relative to earlier asset pricing models.

To derive model implications for asset prices, the preference parameters are calibrated. The calibrated magnitude of the risk aversion and the IES is an empirical issue. Hansen and Singleton (1982), Attanasio and Weber (1989), and Vissing-Jorgensen and Attanasio (2003) estimate the IES to be well in excess of 1. Hall (1988) and Campbell (1999), on the other hand, estimate its value to be well below 1. BY show that even if the population value of the IES is larger than 1, the estimation methods used by Hall would measure the IES to be close to zero. That is, in the presence of time-varying consumption volatility, there is a severe downward bias in the point estimates of the IES. Using data from financial markets, Bansal, Khatchatrian, and Yaron (2005) and Bansal and Shaliastovich (2007) provide further evidence on the magnitude of the IES.

Different techniques are employed to provide empirical and theoretical support for the existence of long-run components in consumption and dividends. Whereas BY calibrate parameters to match the annual moments of consumption and dividend growth rates, Bansal, Gallant, and Tauchen (2007) and Bansal, Kiku, and Yaron (2006) formally test the model using the efficient and generalized method of moments, respectively. Using multivariate analysis, Hansen, Heaton, and Li (2005) and Bansal, Kiku, and Yaron (2006) present evidence for long-run components in consumption growth. Colacito and Croce (2006) also provide statistical support for the long-run components in consumption data for the United States and other developed economies. Lochstoer and Kaltenbrunner (2006) provide a production-

based motivation for long-run risks in consumption. They show that in a standard production economy, where consumption is endogenous, the consumption growth process contains a predictable long-run component similar to that in the BY model. There is considerable support for the volatility channel as well. Bansal, Khatchatrian, and Yaron (2005) show that consumption volatility is time varying and that its current level predicts future asset valuations (price-dividend ratio) with a significantly negative projection coefficient; this implies that asset markets dislike economic uncertainty. Exploiting the BY uncertainty channel, Lettau, Ludvigson, and Wachter (2007) provide interesting market premium implications of the low-frequency decline in consumption volatility.

BY show that their long-run risks model can explain the risk-free rate, the level of the equity premium, asset price volatility, and many of the return and dividend growth predictability dimensions that have been characterized in earlier work. The time-varying volatility in consumption is important to capture some of the economic outcomes that relate to time-varying risk premia.

The arguments presented in their work also have immediate implications for the cross-sectional differences in mean returns across assets. Bansal, Dittmar, and Lundblad (2002 and 2005) show that the systematic risks across firms should be related to the systematic long-run risks in firms' cash flows that investors receive. Firms whose expected cash-flow (profits) growth rates move with the economy are more exposed to long-run risks and hence should carry higher risk compensation. These authors develop methods to measure the exposure of cash flows to long-run risks, and show that these cash flow betas can account for the differences in risk premia across assets. They show that the high book-to-market portfolio (i.e., value portfolio) has a larger long-run risks beta relative to the low book-to-market portfolio (i.e., growth portfolio). Hence, the high mean return of value firms relative to growth firms is not puzzling. The Bansal, Dittmar, and Lundblad (2002 and 2005) evidence supports a long-run risks explanation for the cross-sectional differences in mean returns.

Several recent papers use the BY long-run risks model to address a rich array of asset market questions; among others, these include Kiku (2006), Colacito and Croce (2006), Lochstoer and Kaltenbrunner (2006), Chen, Collin-Dufresne, and Goldstein (2006), Chen (2006), Eraker (2006), Piazzesi and Schneider (2005), and Bansal and Shaliastovich (2007). Kiku (2006) shows that the long-run risks model can account for the violations of the capital asset pricing model (CAPM) and consumption CAPM (C-CAPM) in explaining the cross-sectional differences in mean returns. Further, the model can capture the entire transition density of value or growth returns, which underscores the importance of long-run risks in accounting for equity markets' behavior. Eraker (2006) and Piazzesi and Schneider (2005) consider the implications of the model for the risk premia on U.S. Treasury bonds and show how to account for some of the average premium puzzles in the term structure literature. Colacito and Croce (2006) extend the long-run risks model to a two-country setup and explain the issues about international risk sharing and exchange rate volatility. Bansal and Shaliastovich (2007) show that the long-run risks model can simultaneously account for the behavior of equity markets, yields, and foreign exchange and explain the nature of predictability and violations of the expectations hypothesis in foreign exchange and Treasury markets. Chen, Collin-Dufresne, and Goldstein (2006) and Chen (2006) analyze the ability of the long-run risks model to explain the credit spread and leverage puzzles of the corporate sector.

Hansen, Heaton, and Li (2005) consider a long-run risks model with a unit IES specification. Using different methods to measure long-run risks exposures of portfolios sorted by book-to-market ratio, they find, as in Bansal, Dittmar, and Lundblad (2005), that these alternative long-run risk measures do line-up in the cross-section with the average returns. They further show that the measurement of long-run risks can be sensitive to the econometric methods used. Hansen and Sargent (2006) highlight the interesting implications of robust decisionmaking for risks in financial markets when the representative agent entertains the long-run risks model as a baseline description of the economy.

The above results indicate that the long-run risks model can go a long way toward providing an explanation for many of the key features of asset markets.

The remainder of the article has three sections. The next section reviews the long-run risks model of Bansal and Yaron (2004). The third section discusses the empirical evidence of the model and, in particular, its implications for the equity, bond, and currency markets. The final section concludes.

LONG-RUN RISKS MODEL

Preferences and the Environment

Consider a representative agent with the following Epstein and Zin (1989) recursive preferences,

$$U_t = \left\{ (1-\delta)C_t^{\frac{1-\gamma}{\theta}} + \delta \left(E_t \left[U_{t+1}^{1-\gamma} \right] \right)^{1/\theta} \right\}^{\frac{\theta}{1-\gamma}},$$

where the rate of time preference is determined by δ , with $0 < \delta < 1$. The parameter θ is determined by the risk aversion and the IES—specifically,

$$\theta \equiv \frac{1-\gamma}{1-\frac{1}{\psi}},$$

where $\gamma \geq 0$ is the risk aversion parameter and $\psi \geq 0$ is the IES. The sign of θ is determined by the magnitudes of the risk aversion and the elasticity of substitution. In particular, if $\psi > 1$ and $\gamma > 1$, then θ will be negative. Note that, when $\theta = 1$ (that is $\gamma = 1/\psi$), one obtains the standard case of expected utility.

As is pointed out in Epstein and Zin (1989), when risk aversion equals the reciprocal of IES (expected utility), the agent is indifferent to the timing of the resolution of the uncertainty of the consumption path. When risk aversion exceeds (is less than) the reciprocal of IES, the agent prefers early (late) resolution of uncertainty of consumption path. In the long-run risks model, agents prefer early resolution of the uncertainty of the consumption path.

The period budget constraint for the agent with wealth, W_t , and consumption, C_t , at date t is

$$(1) \quad W_{t+1} = (W_t - C_t)R_{a,t+1}.$$

$$R_{a,t+1} = \frac{W_{t+1}}{(W_t - C_t)}$$

is the return on the aggregate portfolio held by the agent. As in Lucas (1978), we normalize the supply of all equity claims to be 1 and the risk-free asset to be in zero net supply. In equilibrium, aggregate dividends in the economy (which also include any claims to labor income) equals aggregate consumption of the representative agent. Given a process for aggregate consumption, the return on the aggregate portfolio corresponds to the return on an asset that delivers aggregate consumption as its dividends each time period.

The logarithm of the intertemporal marginal rate of substitution (IMRS), m_{t+1} , for these preferences (Epstein and Zin, 1989) is

$$(2) \quad m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1)r_{a,t+1},$$

and the asset pricing restriction on any continuous return, $r_{i,t+1}$, is

$$(3) \quad E_t \left[\exp \left(\theta \log \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1)r_{a,t+1} + r_{i,t+1} \right) \right] = 1,$$

where g_{t+1} equals $\log(C_{t+1}/C_t)$ —the log growth rate of aggregate consumption. The return, $r_{a,t+1}$, is the log of the return (i.e., continuous return) on an asset that delivers aggregate consumption as its dividends each time period.

The return on the aggregate consumption claim, $r_{a,t+1}$, is not observed in the data, whereas the return on the dividend claim corresponds to the observed return on the market portfolio, $r_{m,t+1}$. The levels of market dividends and consumption are not equal; aggregate consumption is much larger than aggregate dividends. The difference is financed by labor income. In the BY model, aggregate consumption and aggregate dividends are treated as two separate processes and the difference between them defines the agent's labor income process.

The key ideas of the model are developed, and the intuition is provided by means of approximate analytical solutions. However, for the key qualitative results, the model is solved numerically. To derive the approximate solutions for the model, we use the standard Campbell and Shiller (1988) return approximation,

$$(4) \quad r_{a,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + g_{t+1},$$

where lowercase letters refer to variables in logs, in particular, $r_{a,t+1} = \log(R_{a,t+1})$ is the continuous return on the consumption claim and the price-to-consumption ratio is $z_t = \log(P_t/C_t)$. Analogously, $r_{m,t+1}$ and $z_{m,t}$ correspond to the continuous return on the dividend claim and the log of the price-to-dividend ratio. The approximating constants, κ_0 and κ_1 , are specific to the asset under consideration and depend only on the average level of z_t , as shown in Campbell and Shiller (1988). It is important to keep in mind that the average value of z_t for any asset is endogenous to the model and depends on all its parameters and the dynamics of the asset's dividends.

From equation (2), it follows that the innovation in IMRS, m_{t+1} , is driven by the innovations in g_{t+1} and $r_{a,t+1}$. Covariation with the innovation in m_{t+1} determines the risk premium for any asset. We characterize the nature of risk sources and their compensation in the next section.

Long-Run Growth and Economic Uncertainty Risks

The agent's IMRS depends on the endogenous consumption return, $r_{a,t+1}$. The risk compensation on all assets depends on this return, which itself is determined by the process for consumption growth. The dividend process is needed for determining the return on the market portfolio. To capture long-run risks, consumption and dividend growth rates, g_{t+1} and $g_{d,t+1}$, respectively, are modeled to contain a small persistent predictable component, x_t , while fluctuating economic uncertainty is introduced through the time-varying volatility of the cash flows:

$$\begin{aligned}
x_{t+1} &= \rho x_t + \phi_e \sigma_t e_{t+1} \\
g_{t+1} &= \mu + x_t + \sigma_t \eta_{t+1} \\
(5) \quad g_{d,t+1} &= \mu_d + \phi x_t + \phi_d \sigma_t u_{t+1} \\
\sigma_{t+1}^2 &= \sigma^2 + v_1 (\sigma_t^2 - \sigma^2) + \sigma_w w_{t+1} \\
e_{t+1}, u_{t+1}, \eta_{t+1}, w_{t+1} &\sim N.i.i.d.(0, 1),
\end{aligned}$$

with the shocks e_{t+1} , u_{t+1} , η_{t+1} , and w_{t+1} assumed to be mutually independent. The parameter ρ determines the persistence of the expected growth rate process. First, note that when $\phi_e = 0$, the processes g_t and $g_{d,t+1}$ have zero autocorrelation. Second, if $e_{t+1} = \eta_{t+1}$, the process for consumption is ARMA(1,1) used in Campbell (1999), Cecchetti, Lam, and Mark (1993), and Bansal and Lundblad (2002). If in addition $\phi_e = \rho$, then consumption growth corresponds to an AR(1) process used in Mehra and Prescott (1985). The variable σ_{t+1} represents the time-varying volatility of consumption and captures the intuition that there are fluctuations in the level of uncertainty in the economy. The unconditional volatility of consumption is σ^2 . To maintain parsimony, it is assumed that the shocks are uncorrelated and that there is only one source of time-varying economic uncertainty that affects consumption and dividends.

Two parameters, $\phi > 1$ and $\phi_d > 1$, calibrate the overall volatility of dividends and its correlation with consumption. The parameter ϕ is larger than 1 because corporate profits are more sensitive to changing economic conditions relative to consumption. Note that consumption and dividends are not cointegrated in the above specification; Bansal, Gallant, and Tauchen (2007) develop a specification that does allow for cointegration between consumption and dividends.

To better understand the role of long-run risks, consider the scaled long-run variance (or variance ratio) of consumption for horizon J ,

$$\sigma_{c,J}^2 = \frac{\text{Var} \left[\sum_{j=1}^J g_{t+j} \right]}{J \text{Var} [g_t]}.$$

The magnitude of this consumption growth volatility is the same for all J if consumption is uncorrelated across time. This scaled variance increases with the horizon when the expected growth is persistent. Hence, agents face larger

aggregate consumption volatility at longer horizons. As the persistence in x and/or its variance increases, the magnitude of long-run volatility will rise. In equilibrium, this increase in magnitude of aggregate consumption volatility will require a sizeable compensation if the agents prefer early resolution of uncertainty about the consumption path.

Using multivariate statistical analysis, Hansen, Heaton, and Li (2005) and Bansal, Kiku, and Yaron (2006) provide evidence on the existence of the long-run component in observed consumption and dividends. Using simulation methods, Bansal and Yaron (2005) document the presence of the long-run component in U.S. consumption data, whereas Colacito and Croce (2006) estimate this component in consumption for many developed economies. Note that there can be considerable persistence in the time-varying consumption volatility as well; hence, the long-run variance of the conditional volatility of consumption can be very large as well.

To see the importance of the small low-frequency movements for asset prices, consider the quantity

$$E_t \left[\sum_{j=1}^{\infty} \kappa_1^j g_{t+j} \right].$$

With $\kappa_1 < 1$, this expectation equals

$$\frac{\kappa_1 x_t}{1 - \kappa_1 \rho}.$$

Even though the variance of x is tiny, while ρ is fairly high, shocks to x_t (the expected growth rate component) can still alter growth rate expectations for the long run, leading to volatile asset prices. Hence, investor concerns about expected long-run growth rates can alter asset prices quite significantly.

Equilibrium and Asset Prices

The consumption and dividend growth rate processes are exogenous in this endowment economy. Further, the IMRS depends on an endogenous return, $r_{a,t+1}$. To characterize the IMRS and the behavior of asset returns, a solution for the log price-consumption ratio, z_t , and

the log price-dividend ratio, $z_{m,t}$, is needed. The relevant state variables for z_t and $z_{m,t}$ are the expected growth rate of consumption, x_t , and the conditional consumption volatility, σ_t^2 .

Exploiting the Euler equation (3), the approximate solution for the log price-consumption ratio, z_t , has the form $z_t = A_0 + A_1x_t + A_2\sigma_t^2$. The solution for A_1 is

$$(6) \quad A_1 = \frac{1 - \frac{1}{\psi}}{1 - \kappa_1\rho}$$

This coefficient is positive if the IES, ψ , is greater than 1. In this case, the intertemporal substitution effect dominates the wealth effect. In response to higher expected growth, agents buy more assets and consequently the wealth-to-consumption ratio rises. In the standard power utility model with risk aversion larger than 1, the IES is less than 1 and therefore A_1 is negative—a rise in expected growth potentially lowers asset valuations. That is, the wealth effect dominates the substitution effect.¹

Corporate payouts (i.e., dividends), with $\phi > 1$, are more sensitive to long-run risks and changes in the expected growth rate lead to a larger reaction in the price of the dividend claim than in the price of the consumption claim. Hence, for the dividend asset,

$$A_{1,m} = \frac{\phi - \frac{1}{\psi}}{1 - \kappa_{1,m}\rho}$$

and is larger in absolute value than the consumption asset.

The solution coefficient, A_2 , for measuring the sensitivity of the price-consumption ratio to volatility fluctuations is

$$(7) \quad A_2 = \frac{0.5 \left[\left(\theta - \frac{\theta}{\psi} \right)^2 + (\theta A_1 \kappa_1 \phi_e)^2 \right]}{\theta(1 - \kappa_1 v_1)}$$

¹ An alternative interpretation with the power utility model is that higher expected growth rates increase the risk-free rate to an extent that discounting dominates the effects of higher expected growth rates. This leads to a fall in asset prices.

An analogous coefficient for the market price-dividend ratio, $A_{2,m}$, is provided in BY.

The expression for A_2 provides two valuable insights. First, if the IES and risk aversion are larger than 1, then θ and consequently A_2 are negative. In this case, a rise in consumption volatility lowers asset valuations and increases the risk premia on all assets. Second, an increase in the permanence of volatility shocks—that is, an increase in v_1 —magnifies the effects of volatility shocks on valuation ratios as investors perceive changes in economic uncertainty as very long lasting.

Pricing of Short-Run, Long-Run, and Volatility Risks

Substituting the solutions for the price-consumption ratio, z_t , into the expression for equilibrium return for $r_{a,t+1}$ in equation (4), one can now characterize the solution for the IMRS that can be used to value all assets. The log of the IMRS m_{t+1} can always be stated as the sum of its conditional mean and its one-step-ahead innovation. The conditional mean is affine in expected mean and conditional variance of consumption growth and can be expressed as

$$(8) \quad E_t(m_{t+1}) = m_0 - \frac{1}{\psi}x_t + \frac{\left(\frac{1}{\psi} - \gamma\right)(\gamma - 1)}{2} \left[1 + \left(\frac{\kappa_1 \phi_e}{1 - \kappa_1 \rho} \right)^2 \right] \sigma_t^2,$$

where m_0 is a constant determined by the preference and consumption dynamics parameters.

The innovation in the IMRS is very important for thinking about risk compensation (risk premia) in various markets. Specifically, it is equal to

$$(9) \quad m_{t+1} - E_t(m_{t+1}) = -\lambda_{m,\eta}\sigma_t\eta_{t+1} - \lambda_{m,e}\sigma_t e_{t+1} - \lambda_{m,w}\sigma_w w_{t+1},$$

where $\lambda_{m,\eta}$, $\lambda_{m,e}$, and $\lambda_{m,w}$ are the market prices for the short-run, long-run, and volatility risks, respectively. The market prices of systematic risks, including the compensation for stochastic volatility risk in consumption, can be expressed in terms of the underlying preferences parameters that govern the evolution of consumption growth:

$$\begin{aligned}
 \lambda_{m,\eta} &= \gamma \\
 \lambda_{m,e} &= \left(\gamma - \frac{1}{\psi} \right) \left[\frac{\kappa_1 \varphi_e}{1 - \kappa_1 \rho} \right] \\
 \lambda_{m,w} &= \left(\gamma - \frac{1}{\psi} \right) (1 - \gamma) \left[\frac{\kappa_1 \left(1 + \left(\frac{\kappa_1 \varphi_e}{1 - \kappa_1 \rho} \right)^2 \right)}{2(1 - \kappa_1 v_1)} \right].
 \end{aligned}
 \tag{10}$$

The risk compensation for the η_{t+1} shocks is very standard, and $\lambda_{m,\eta}$ equals the risk aversion parameter, γ . In the special case of power utility, $\gamma = 1/\psi$, the risk compensation parameters $\lambda_{m,e}$ and $\lambda_{m,w}$ are zero. Long-run risks and volatility are priced only when the agent is not indifferent to the timing of the uncertainty resolution for the consumption path—that is, when risk aversion is different from the reciprocal of the IES. For this to be the case, γ should be larger than $1/\psi$. The market prices of long-run and volatility risks are sensitive to the magnitude of the permanence parameter, ρ , as well. The risk compensation for long-run risks and volatility risks rises as the permanence parameter, ρ , rises.

The equity premium in the presence of time-varying economic uncertainty is

$$\begin{aligned}
 E_t(r_{m,t+1} - r_{f,t}) &= \\
 \beta_{m,\eta} \lambda_{m,\eta} \sigma_t^2 + \beta_{m,e} \lambda_{m,e} \sigma_t^2 + \beta_{m,w} \lambda_{m,w} \sigma_w^2 - 0.5 \text{Var}_t(r_{m,t+1}).
 \end{aligned}
 \tag{11}$$

The first beta corresponds to the exposure to short-run risks and the second to long-run risks. The third beta (that is, $\beta_{m,w}$) captures the return's exposure to volatility risks. It is important to note that all the betas in this general equilibrium framework are endogenous. They are completely pinned down by the dynamics of the asset's dividends and the preferences parameters of the agent. The quantitative magnitude of the betas and the risk premium for the consumption claim is discussed below.

The risk premium on the market portfolio is time varying as σ_t fluctuates. The ratio of the conditional risk premium to the conditional volatility of the market portfolio fluctuates with

σ_t and therefore the Sharpe ratio is time varying. The maximal Sharpe ratio in this model, which approximately equals the conditional volatility of the log IMRS, also varies with σ_t . During periods of high economic uncertainty (i.e., consumption volatility), all risk premia rise.

The first-order effects on the level of the risk-free rate, as discussed in Bansal and Yaron (2005), are the rate of time preference and the average consumption growth rate divided by the IES. Increasing the IES keeps the level low. The variance of the risk-free rate is determined by the volatility of the expected consumption growth rate and the IES. Increasing the IES lowers the volatility of the risk-free rate. In addition, incorporating economic uncertainty leads to an interesting channel for interpreting fluctuations in the real risk-free rate. In addition, this has serious implications for the measurement of the IES in the data, which heavily relies on the link between the risk-free rate and expected consumption growth. In the presence of varying volatility, the estimates of the IES based on the projections considered in Hall (1988) and Campbell (1999) are seriously biased downward.

Hansen, Heaton, and Li (2005) also consider a long-run risks model specification where the IES is pinned at 1. This specific case affords considerable simplicity, as the wealth-to-consumption ratio is constant. To solve the model at values of the IES that differ from 1, the authors provide approximations around the case where the IES is 1. Bansal, Kiku, and Yaron (2006) provide an alternative approximate solution that relies on equation (4); they show how to derive the return $r_{a,t}$ along with the endogenous approximating constants, κ_1 and κ_0 , for any configuration of preferences parameters.

DATA AND MODEL IMPLICATIONS

Data and Growth Rate Dynamics

BY calibrate the model described in (5) at the monthly frequency. From this monthly model they derive time-aggregated annual growth rates of consumption and dividends to match key aspects of annual aggregate consumption and

Table 1
Time-Series Properties of Data

Variable	Estimate	Standard error
$\sigma(g)$	2.93	(0.69)
AC(1)	0.49	(0.14)
AC(2)	0.15	(0.22)
AC(5)	-0.08	(0.10)
AC(10)	0.05	(0.09)
VR(2)	1.61	(0.34)
VR(5)	2.01	(1.23)
VR(10)	1.57	(2.07)
$\sigma(g_d)$	11.49	(1.98)
AC(1)	0.21	(0.13)
$\text{corr}(g, g_d)$	0.55	(0.34)

NOTE: This table displays the time-series properties of aggregate consumption and dividend growth rates: g and g_d , respectively. The statistics are based on annual observations from 1929 to 1998. Consumption is real per capita consumption of non-durables and services; dividends are the sum of real dividends across all CRSP firms. $AC(j)$ is the j th autocorrelation, $VR(j)$ is the j th variance ratio, σ is the volatility, and corr denotes the correlation. Standard errors are Newey and West (1987)-corrected using 10 lags.

dividends data. For consumption, Bureau of Economic Analysis data on real per capita annual consumption growth of non-durables and services for the period 1929-98 is used. Dividends and the value-weighted market return data are taken from the Center for Research in Security Prices (CRSP). All nominal quantities are deflated using the consumer price index.

The mean annual real per capita consumption growth rate is 1.8 percent, and its standard deviation is about 2.9 percent. Table 1, adapted from BY, shows that, in the data, consumption growth has a large first-order autocorrelation coefficient and a small second-order coefficient. The standard errors in the data for these autocorrelations are sizeable. An alternative way to view the long-horizon property of the consumption and dividend growth rates is to use variance ratios, which are themselves determined by the autocorrelations (Cochrane, 1988). In the data, the variance ratios

first rise significantly and at a horizon of about seven years start to decline. The standard errors on these variance ratios, not surprisingly, are quite substantial.

In terms of the specific parameters for the consumption dynamics, BY calibrate ρ at 0.979, which determines the persistence in the long-run component in growth rates. Their choice of φ_e and σ ensures that the model matches the unconditional variance and the autocorrelation function of annual consumption growth. The standard deviation of the innovation in consumption equals 0.0078. This parameter configuration implies that the predictable variation in monthly consumption growth is very small—the implied R^2 is only 4.4 percent. The exposure of the corporate sector to long-run risks is governed by ϕ , and its magnitude is similar to that in Abel (1999). The standard deviation of the monthly innovation in dividends, $\varphi_d\sigma$, is 0.0351. The parameters of the volatility process are chosen to capture the persistence in consumption volatility. Based on the evidence of slow decay in volatility shocks, BY calibrate v_1 , the parameter governing the persistence of conditional volatility, at 0.987. The shocks to the volatility process have very small volatility; σ_w is calibrated at 0.23×10^{-5} . At the calibrated parameters, the modeled consumption and dividend growth rates very closely match the key consumption and dividends data features reported in Table 1. Bansal, Gallant, and Tauchen (2007) provide simulation-based estimation evidence that supports this configuration as well.

Table 2 presents the targeted asset market data for 1929-98. The equity risk premium is 6.33 percent per annum, and the real risk-free rate is 0.86 percent. The annual market return volatility is 19.42 percent, and that of the real risk-free rate is quite small, about 1 percent per annum. The volatility of the price-dividend ratio is quite high, and it is a very persistent series. In addition to these data dimensions, BY also evaluate the ability of the model to capture predictability of returns. Bansal, Khatchatrian, and Yaron (2005) show that, consistent with the implications of the BY model, price-dividend ratios are negatively correlated with consumption volatility at long leads and lags.

It is often argued that, in the data, consumption and dividend growth are close to being i.i.d. BY show that their model of consumption and dividends is also consistent with the observed data on consumption and dividends growth rates. However, although the financial market data are hard to interpret from the perspective of the i.i.d. growth rate dynamics, BY show that it is interpretable from the perspective of the growth rate dynamics that incorporate long-run risks. Given these difficulties in discrimination across these two models, Hansen and Sargent (2006) use features of the long-run model for developing economic models where agents update their model beliefs in a manner that incorporates robustness against model misspecification.

Preference Parameters

The preference parameters take account of economic considerations. The time preference parameter, $\delta < 1$, and the risk aversion parameter, γ , is either 7.5 or 10. Mehra and Prescott (1985) argue that a reasonable upper bound for risk aversion is around 10. The IES is set at 1.5: An IES value that is not less than 1 is important for the quantitative results.

There is considerable debate about the magnitude of the IES. Hansen and Singleton (1982) and Attanasio and Weber (1989) estimate the IES to be well in excess of 1. More recently, Guvenen (2001) and Vissing-Jorgensen and Attanasio (2003) also estimate the IES over 1; they show that their estimates are close to that used in BY. However, Hall (1988) and Campbell (1999) estimate the IES to be well below 1. BY argue that the low IES estimates of Hall and Campbell are based on a model without time-varying volatility. They show that ignoring the effects of time-varying consumption volatility leads to a serious downward bias in the estimates of the IES. If the population value of the IES in the BY model is 1.5, then the estimated value of the IES using Hall estimation methods will be less than 0.3. With fluctuating consumption volatility, the projection of consumption growth on the level of the risk-free rate does not equal the IES, leading to the downward bias. This suggests that Hall's and Campbell's estimates may not be a robust guide for calibrating the IES.

Table 2

Asset Market Data

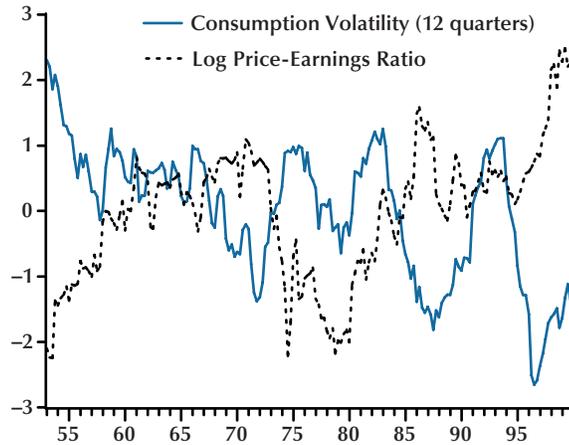
Variable	Estimate	Standard error
Returns		
$E(r_m - r_f)$	6.33	(2.15)
$E(r_f)$	0.86	(0.42)
$\sigma(r_m)$	19.42	(3.07)
$\sigma(r_f)$	0.97	(0.28)
Price-dividend ratio		
$E(\exp(p - d))$	26.56	(2.53)
$\sigma(p - d)$	0.29	(0.04)
$AC1(p - d)$	0.81	(0.09)
$AC2(p - d)$	0.64	(0.15)

NOTE: This table presents descriptive statistics of asset market data. The moments are calculated using annual observations from 1929 through 1998. $E(r_m - r_f)$ and $E(r_f)$ are, respectively, the annualized equity premium and mean risk-free rate; $\sigma(r_m)$, $\sigma(r_f)$, and $\sigma(p - d)$ are, respectively, the annualized volatilities of the market return, risk-free rate, and log price-dividend ratio; AC1 and AC2 denote the first and second autocorrelations. Standard errors are Newey and West (1987)-corrected using 10 lags.

In addition to the above arguments, the empirical evidence in Bansal, Khatchatrian, and Yaron (2005) shows that a rise in consumption volatility sharply lowers asset prices at long leads and lags, and high current asset valuations forecast higher future corporate earnings growth. Figures 1 through 4 use data from the United States, United Kingdom, Germany, and Japan to evaluate the volatility channel. The asset valuation measure is the price-to-earnings ratio, and the consumption volatility measure is constructed by averaging eight lags of the absolute value of consumption residuals. It is evident from the figures that a rise in consumption volatility lowers asset valuations for all countries under consideration; this highlights the volatility channel and motivates the specification of an IES larger than 1. In a two-country extension of the model, Bansal and Shaliastovich (2007) show that dollar prices of foreign currency and forward premia co-move negatively with the consumption volatility differential, whereas the ex ante currency returns

Figure 1

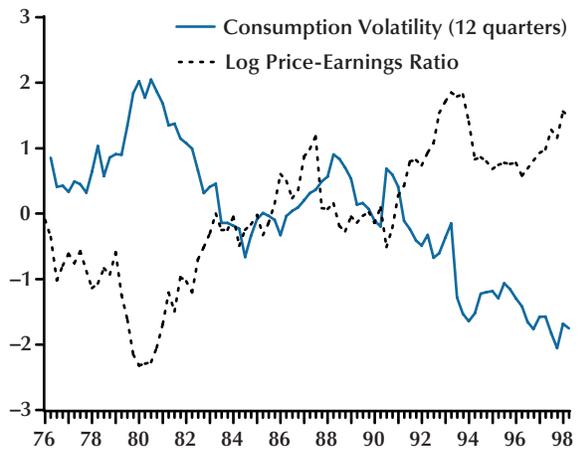
**P/E Ratio and Consumption Volatility:
United States**



NOTE: This figure plots consumption volatility along with the logarithm of the price-earnings ratio for the United States. Both series are standardized.

Figure 2

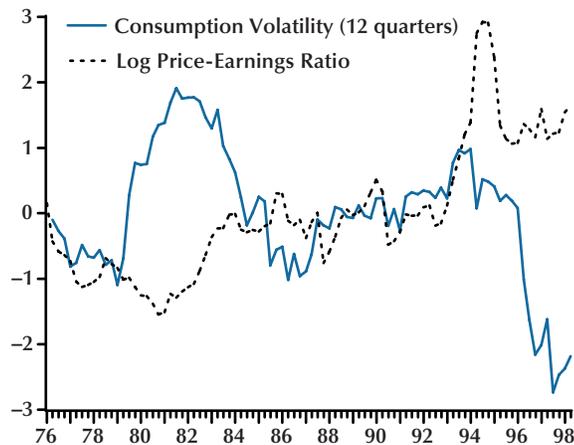
**P/E Ratio and Consumption Volatility:
United Kingdom**



NOTE: This figure plots consumption volatility along with the logarithm of the price-earnings ratio for the United Kingdom. Both series are standardized.

Figure 3

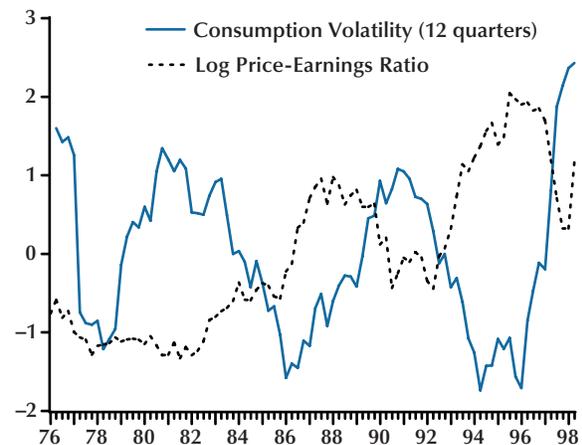
**P/E Ratio and Consumption Volatility:
Germany**



NOTE: This figure plots consumption volatility along with the logarithm of the price-earnings ratio for Germany. Both series are standardized.

Figure 4

**P/E Ratio and Consumption Volatility:
Japan**



NOTE: This figure plots consumption volatility along with the logarithm of the price-earnings ratio for Japan. Both series are standardized.

Table 3
Risk Components and Risk Compensation

	$\psi = 0.1$	$\psi = 0.5$	$\psi = 1.5$
mpr_{η}	93.60	93.60	93.60
mpr_e	0.00	137.23	160.05
mpr_w	0.00	-27.05	-31.56
β_{η}	1.00	1.00	1.00
β_e	-16.49	-1.83	0.61
β_w	11,026.45	1,225.16	-408.39
prm_{η}	0.73	0.73	0.73
prm_e	0.00	-1.96	0.76
prm_w	0.00	-0.08	0.03

NOTE: This table presents model-implied components of the risk premium on the consumption asset for different values of the intertemporal elasticity of substitution parameter, ψ . All entries are based on $\gamma = 10$. The parameters that govern the dynamics of the consumption process in equation (5) are identical to those in Bansal and Yaron (2004): $\rho = 0.979$, $\sigma = 0.0078$, $\phi_e = 0.044$, $v_1 = 0.987$, $\sigma_w = 0.23 \times 10^{-5}$, and $\kappa_1 = 0.997$. The first three rows report the annualized percentage prices of risk for innovations in consumption, the expected growth risks, and the consumption volatility risks— mpr_{η} , mpr_e , and mpr_w , respectively. These prices of risks correspond to annualized percentage values for $\lambda_{m,\eta}\sigma$, $\lambda_{m,e}\sigma$, $\lambda_{m,w}\sigma_w$ in equation (9). The exposures of the consumption asset to the three systematic risks, β_{η} , β_e , and β_w , are presented in the middle part of the table. Total risk compensation in annual percentage terms for each risk is reported as prm_* and equals the product of the price of risk, the standard deviation of the shock, and the beta for the specific risk, respectively.

have positive correlations with it. This provides further empirical support for a magnitude of the IES. In terms of growth rate predictability, Ang and Bekaert (2007) and Bansal, Khatchatrian, and Yaron (2005) report a positive relation between asset valuations and expected earnings growth. These data features, as discussed in the theory sections above, again require an IES larger than 1.

Asset Pricing Implications

To underscore the importance of two key aspects of the model, preferences and long-run risks, first consider the genesis of the risk premium on $r_{a,t+1}$ —the return on the asset that delivers aggregate consumption as its dividends. The determination of risk premia for other dividend claims follows the same logic.

Table 3 shows the market price of risk and the breakdown of the risk premium from various risk sources. Column 1 considers the case of power utility where the IES equals the reciprocal of the risk aversion parameter. As discussed ear-

lier, the prices of long-run and volatility risks are zero. In the power utility case, the main risk is the short-run risk and the risk premium on the consumption asset equals $\gamma\sigma^2$, which is 0.7 percent per annum.

Column 2 of Table 3 considers the case with an IES less than 1 (set at 0.5). For long-run growth rate risks, the price of risk is positive; for volatility risks, the price of risk is negative, as γ is larger than the reciprocal of the IES. However, the consumption asset's beta for long-run risks (beta with regard to the innovations in x_{t+1}) is negative. This, as discussed earlier, is because A_1 is negative (see equation (6)), implying that a rise in expected growth lowers the wealth-to-consumption ratio. Consequently, long-run risks in this case contribute a negative risk premium of -1.96 percent per annum. The market price of volatility risks is negative and small; however, the asset's beta for this risk source is large and positive, reflecting the fact that asset prices rise when economic uncertainty rises (see equation (7)). In all, when the IES is less than 1, the risk premium on the

consumption asset is negative, which is highly counterintuitive, and highlights the implausibility of this parameter configuration.

Column 3 of Table 3 shows that when the IES is larger than 1, the price of long-run growth risks rises. More importantly, the asset's beta for long-run growth risks is positive and that for volatility risks is negative. Both these risk sources contribute toward a positive risk premium. The risk premium from long-run growth is 0.76 percent and that for the short-run consumption shock is 0.73 percent. The overall risk premia for this consumption asset is 1.52 percent. This evidence shows that an IES larger than 1 is required for the long-run and volatility risks to carry a positive risk premium.

It is clear from Table 3 that the price of risk is highest for the long-run risks (see columns 2 and 3) and smallest for the volatility risks. A comparison of columns 2 and 3 also shows that raising the IES increases the prices of long-run and volatility risks in absolute value. The magnitudes reported in Table 3 are with $\rho = 0.979$ —lowering this persistence parameter also lowers the prices of long-run and volatility risks (in absolute value). Increasing the risk aversion parameter increases the prices of all consumption risks. Hansen and Jagannathan (1991) document the importance of the maximal Sharpe ratio, determined by the volatility of the IMRS, in assessing asset pricing models. Incorporating long-run risks increases the maximal Sharpe ratio in the model, which easily satisfies the non-parametric bounds of Hansen and Jagannathan (1991).

The risk premium on the market portfolio (i.e., the dividend asset) is also affected by the presence of long-run risks. To underscore their importance, assume that consumption and dividend growth rates are i.i.d. This shuts off the long-run risks channel. The market risk premium in this case is

$$(12) \quad E_t(r_{m,t+1} - r_{f,t}) = \gamma \text{cov}(g_{t+1}, g_{d,t+1}) - 0.5 \text{Var}(g_{d,t+1}),$$

and market return volatility equals the dividend growth rate volatility. If shocks to consumption and dividends are uncorrelated, then the geo-

metric risk premium is negative and equals $-0.5 \text{Var}(g_{d,t+1})$. If the correlation between monthly consumption and dividend growth is 0.25, then the equity premium is 0.08 percent per annum, which is similar to the evidence documented in Mehra and Prescott (1985) and Weil (1989).

BY show that their model, which incorporates long-run growth rate risks and fluctuating economic uncertainty, provides a very close match to the asset market data reported in Table 2. That is, the model can account for the *low risk-free rate*, *high equity premium*, *high asset price volatility*, and *low risk-free rate volatility*. The BY model matches additional data features, such as (i) predictability of returns at short and long horizons using dividend yield as a predictive variable, (ii) time-varying and persistent market return volatility, (iii) negative correlation between market return and volatility shocks, i.e., the volatility feedback effect, and the (iv) negative relation between consumption volatility and asset prices at long leads and lags. (Also see Bansal, Khatchatrian, and Yaron, 2005.)

Cross-Sectional Implications

Table 4, shows that there are sizable differences in mean real returns across portfolios sorted by book-to-market ratio, size, and momentum for quarterly data from 1967 to 2001. The size and book-to-market sorts place firms into different deciles once per year, and the subsequent return on these portfolios is used for empirical work. For momentum assets, CRSP-covered New York Stock Exchange and American Stock and Options Exchange stocks are sorted on the basis of their cumulative return over months $t-12$ through $t-1$. The loser portfolio (M1) includes firms with the worst performance over the past year, and the winner portfolio (M10) includes firms with the best performance. The data show that subsequent returns on these portfolios have a large spread (i.e., M10 return – M1 return), about 4.62 percent per quarter: This is the *momentum spread puzzle*. Similarly, the highest book-to-market firms (B10) earn average real quarterly returns of 3.27 percent, whereas the lowest book-to-market (B1) firms average 1.54 percent per quarter. The value spread (return on B10 – return on B1) is about 2 percent

Table 4**Portfolio Returns**

	Mean	Standard deviation		Mean	Standard deviation		Mean	Standard deviation
S1	0.0230	0.1370	B1	0.0154	0.1058	M1	-0.0104	0.1541
S2	0.0231	0.1265	B2	0.0199	0.0956	M2	0.0070	0.1192
S3	0.0233	0.1200	B3	0.0211	0.0921	M3	0.0122	0.1089
S4	0.0233	0.1174	B4	0.0218	0.0915	M4	0.0197	0.0943
S5	0.0242	0.1112	B5	0.0200	0.0798	M5	0.0135	0.0869
S6	0.0207	0.1050	B6	0.0234	0.0813	M6	0.0160	0.0876
S7	0.0224	0.1041	B7	0.0237	0.0839	M7	0.0200	0.0886
S8	0.0219	0.1001	B8	0.0259	0.0837	M8	0.0237	0.0825
S9	0.0207	0.0913	B9	0.0273	0.0892	M9	0.0283	0.0931
S10	0.0181	0.0827	B10	0.0327	0.1034	M10	0.0358	0.1139

NOTE: This table presents descriptive statistics for the returns on the 30 characteristic-sorted decile portfolios. Value-weighted returns are presented for portfolios formed on market capitalization (S), book-to-market ratio (B), and momentum (M). M1 represents the lowest momentum (loser) decile, S1 the lowest size (small firms) decile, and B1 the lowest book-to-market decile. Data are converted to real values using the personal consumption expenditure deflator, are sampled at the quarterly frequency, and cover 1967:Q1–2001:Q4.

per quarter: This is the *value spread puzzle*. What explains these big differences in mean returns across portfolios?

Bansal, Dittmar, and Lundblad (2002 and 2005) connect systematic risks to cash-flow risks. They show that an asset's risk measure (i.e., its beta) is determined by its cash-flow properties. In particular, their paper shows that cross-sectional differences in asset betas mostly reflect differences in systematic risks in cash flows. Hence, systematic risks in cash flows ought to explain differences in mean returns across assets. They develop two ways to measure the long-run risks in cash flows. First they model dividend and consumption growth rates as a VAR and measure the discounted impulse response of the dividend growth rates to consumption innovations. This is one measure of risks in cash flows. Their second measure is based on stochastic cointegration, which is estimated by regressing the log level of dividends for each portfolio on a time trend and the log level of consumption. Specifically, consider the projection

$$d_t = \tau(0) + \tau(1)t + \tau(2)c_t + \zeta_t,$$

where the projection coefficient, $\tau(2)$, measures the long-run consumption risk in the asset's dividends. The coefficient $\tau(2)$ will be different for different assets.

Bansal, Dittmar, and Lundblad (2002 and 2005) show that the exposure of dividend growth rates to the long-run component in consumption has considerable cross-sectional explanatory power. That is, dividends' exposure to long-run consumption risks is an important explanatory variable in accounting for differences in mean returns across portfolios. Portfolios with high mean returns also have higher dividend exposure to consumption risks. The cointegration-based measure of risk, τ_2 , also provides very valuable information about mean returns on assets. The cross-sectional R^2 from regressing the mean returns on the dividend-based risk measures is well over 65 percent. In contrast, other approaches find it quite hard to explain the differences in mean returns for the 30-asset menu used in Bansal, Dittmar, and Lundblad (2005). The standard consumption betas (i.e., C-CAPM) and the market-based CAPM asset betas have close to zero explanatory power. The R^2 for the C-CAPM

Bansal

is 2.7 percent, and that for the market CAPM is 6.5 percent, with an implausible negative slope coefficient. The Fama and French three-factor empirical specification also generates point estimates with negative, and difficult to interpret, prices of risk for the market and size factors; the cross-sectional R^2 is about 36 percent. Compared with all these models, the cash-flow risks model of Bansal, Dittmar, and Lundblad (2005) is able to capture a significant portion of the differences in risk premia across assets. Hansen, Heaton, and Li (2005) inquire about the robustness of the stochastic cointegration-based risk measures considered in Bansal, Dittmar, and Lundblad (2002). They argue that the dividend-based consumption betas—particularly, the cointegration-based risk measures—are imprecisely estimated in the time series. Interestingly, across the different estimation procedures, the cash-flow beta risk measures across portfolios line-up closely with the average returns across assets. That is, in the cross-section of assets (as opposed to the time series), the price of risk associated with the long-run risks measures is reliably significant.

Bansal, Dittmar, and Kiku (2006) derive new results that link this cointegration parameter to consumption betas by investment horizon and evaluate the ability of their model to explain differences in mean returns for different horizons. They provide new evidence regarding the robustness of the stochastic cointegration-based measures of permanent risks in equity markets. Parker and Julliard (2005) evaluate whether long-run risks in aggregate consumption can account for the cross-section of expected returns. Malloy, Moskowitz, and Vissing-Jorgensen (2005) evaluate whether long-run risks in stockholders' consumption relative to aggregate consumption has greater ability to explain the cross-section of equity returns, relative to aggregate consumption measures.

Term Structure and Currency Markets

Colacito and Croce (2006) consider a two-country version of the BY model. They show that this model can account for the low correlation in consumption growth across countries but high

correlation in marginal utilities across countries (high risk sharing despite a low measured cross-country consumption correlation). This feature of international data is highlighted in Brandt, Cochrane, and Santa-Clara (2006). The key idea that Colacito and Croce pursue is that the long-run risks component is very similar across countries, but in the short-run consumption growth can be very different. That is, countries share very similar long-run prospects, but in the short-run they can look very different. This dimension, they show, is sufficient to induce high correlation in marginal utilities across countries. It also accounts for high real exchange volatility.

BY derive implications for the real term structure of interest rates for the long-run risks model. More recent papers by Eraker (2006) and Piazzesi and Schneider (2005) also consider the quantitative implications for the nominal term structure using the long-run risks model. Bansal and Shaliastovich (2007) show that the BY model can simultaneously account for the upward-sloping terms structure, the violations of the expectations hypothesis in the bond markets, the violations in the foreign currency markets, and the equity returns. This evidence indicates that the long-run risks model provides a solid baseline model for understanding financial markets. With simple modifications the model can be used to analyze the impact of changing short-term interest rates on financial markets; that is, it can help in designing policy.

CONCLUSION

The work of Bansal and Lundblad (2002), Bansal and Yaron (2004), and Bansal, Dittmar, and Lundblad (2005) shows that the long-run risks model can help interpret several features of financial markets. These papers argue that investors care about the long-run growth prospects and the uncertainty (time-varying consumption volatility) surrounding the growth rate. Risks associated with changing long-run growth prospects and varying economic uncertainty drive the level of equity returns, asset price volatility, risk premia across assets, and predictability of returns in financial markets.

Recent papers indicate that the channel in this model can account for nominal yield curve features, such as the violation the expectations hypothesis and the average upward-sloping nominal yield curve. Evidence presented in Colacito and Croce (2006) and Bansal and Shaliastovich (2007) shows that the model also accounts for key aspects of foreign exchange markets.

Growing evidence suggests that the long-run risks model can explain a rich array of financial market facts. This suggests that the model can be used to analyze the impact of economic policy on financial markets.

REFERENCES

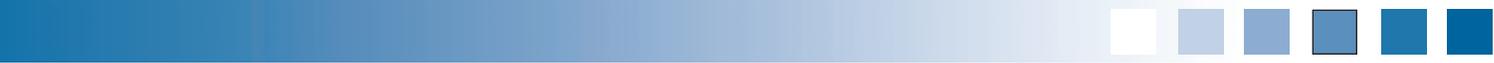
- Abel, Andrew. "Risk Premia and Term Premia in General Equilibrium." *Journal of Monetary Economics*, 1999, 43(1), pp. 3-33.
- Ang, Andrew and Bekaert, Geert. "Stock Return Predictability: Is It There?" *Review of Financial Studies*, May 2007, 20(3), pp. 651-707.
- Attanasio, Orazio P. and Weber, Guglielmo. "Intertemporal Substitution, Risk Aversion, and the Euler Equation for Consumption." *Economic Journal*, March 1989, 99(Suppl 395), pp. 59-73.
- Bansal, Ravi; Dittmar, Robert F. and Kiku, Dana. "Cointegration and Consumption Risks in Asset Returns." Forthcoming in *Review of Financial Studies*, 2007.
- Bansal, Ravi; Dittmar, Robert F. and Lundblad, Christian T. "Interpreting Risk Premia Across Size, Value and Industry Portfolios." Working paper, Duke University, 2002.
- Bansal, Ravi; Dittmar, Robert F. and Lundblad, Christian T. "Consumption, Dividends, and the Cross-Section of Equity Returns." *Journal of Finance*, August 2005, 60(4), pp. 1639-72.
- Bansal, Ravi; Gallant, A. Ronald and Tauchen, George. "Rational Pessimism, Rational Exuberance, and Markets for Macro Risks." Forthcoming in *Review of Economic Studies*, 2007.
- Bansal, Ravi; Khatchatrian, Varoujan and Yaron, Amir. "Interpretable Asset Markets." *European Economic Review*, April 2005, 49(3), pp. 531-60.
- Bansal, Ravi; Kiku, Dana and Yaron, Amir. "Risks for the Long-Run: Estimation and Inference." Working paper, Duke University, 2006.
- Bansal, Ravi and Lundblad, Christian T. "Market Efficiency, Asset Returns, and the Size of the Risk Premium in Global Equity Markets." *Journal of Econometrics*, August 2002, 109(2), pp. 195-237.
- Bansal, Ravi and Shaliastovich, Ivan. "Risk and Return in Bond, Currency and Equity Markets: A Unified Approach." Working paper, Duke University, 2007.
- Bansal, Ravi and Yaron, Amir. "Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, August 2004, 59(4), pp. 1481-509.
- Bansal, Ravi and Yaron, Amir. "The Asset Pricing Macro Nexus and Return Cash Flow Predictability." Working paper, Duke University, 2005.
- Brandt, Michael W.; Cochrane, John H. and Santa-Clara, Pedro. "International Risk Sharing Is Better Than You Think (or Exchange Rates Are Much Too Smooth)." *Journal of Monetary Economics*, May 2006, 53(4), pp 671-98.
- Campbell, John Y. "Asset Prices, Consumption and the Business Cycle," in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*. Volume 1. Amsterdam: Elsevier Science, North Holland, 1999.
- Campbell, John Y. and Shiller, Robert J. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Review of Financial Studies*, 1988, 1(3), pp. 195-228.
- Cecchetti, Stephen G.; Lam, Pok-Sang and Mark, Nelson C. "The Equity Premium and the Risk-Free Rate: Matching the Moments." *Journal of Monetary Economics*, February 1993, 31(1), pp. 21-46.

Bansal

- Chen, Hui. "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure." Working paper, University of Chicago, 2006.
- Chen, Long; Collin-Dufresne, Pierre and Goldstein, Robert S. "On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle." Working paper, University of California–Berkeley, 2006.
- Cochrane, John H. "How Big Is the Random Walk in GNP?" *Journal of Political Economy*, October 1988, 96(5), pp. 893-920.
- Colacito, Riccardo and Croce, Mariano M. "Risks for the Long-Run and the Real Exchange Rate." Working paper, University of North Carolina, 2006.
- Epstein, Larry G. and Zin, Stanley E. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, July 1989, 57(4), pp. 937-69.
- Eraker, Bjorn. "Affine General Equilibrium Models." Working paper, Duke University, 2006.
- Fama, Eugene F. "Forward and Spot Exchange Rates." *Journal of Monetary Economics*, November 1984, 14(3), pp. 319-38.
- Fama, Eugene F. and Bliss, Robert R. "The Information in Long-Maturity Forward Rates." *American Economic Review*, September 1987, 77(4), pp. 680-92.
- Guvenen, Fatih. "Mismeasurement of the Elasticity of Intertemporal Substitution: The Role of Limited Stock Market Participation." Working paper, University of Rochester, 2001.
- Hall, Robert E. "Intertemporal Substitution in Consumption." *Journal of Political Economy*, April 1988, 96(2), pp. 339-57.
- Hansen, Lars P.; Heaton, John C. and Li, Nan. "Consumption Strikes Back? Measuring Long-Run Risk." Working paper, University of Chicago, 2005.
- Hansen, Lars P. and Jagannathan, Ravi. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, April 1991, 99(2), pp. 225-62.
- Hansen, Lars P. and Sargent, Thomas. "Fragile Beliefs and the Price of Model Uncertainty." Working paper, University of Chicago, 2006.
- Hansen, Lars P. and Singleton, Kenneth. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica*, September 1982, 50(5), pp. 1269-86.
- Kiku, Dana. "Is the Value Premium a Puzzle?" Working paper, University of Pennsylvania, 2006.
- Lettau, Martin; Ludvigson, Sydney and Wachter, Jessica. "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" Forthcoming in *Review of Financial Studies*, 2007.
- Lochstoer, Lars A. and Kaltenbrunner, Georg. "Long-Run Risk through Consumption Smoothing." Working paper, London Business School, 2006.
- Lucas, Robert E. Jr. "Asset Prices in an Exchange Economy." *Econometrica*, November 1978, 46(6), pp. 1429-46.
- Malloy, Christopher J.; Moskowitz, Tobias J. and Vissing-Jorgensen, Annette. "Long-Run Stockholder Consumption Risk and Asset Returns." Working paper, University of Chicago, 2005.
- Mehra, Rajnish and Prescott, Edward C. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, March 1985, 15(2), pp. 145-61.
- Parker, Jonathan and Julliard, Christian. "Consumption Risk and the Cross-Section of Asset Returns." *Journal of Political Economy*, February 2005, 113(1), pp. 185-222.
- Piazzesi, Monika and Schneider, Martin. "Equilibrium Yield Curves." NBER Working Paper No. 12609, National Bureau of Economic Research, 2005.
- Shiller, Robert J. "Do Stock Prices Move Too Much To Be Justified by Subsequent Changes in Dividends?" *American Economic Review*, July 1981, 71(3), pp. 421-36.

Vissing-Jorgensen, Annette and Attanasio, Orazio P.
“Stock Market Participation, Intertemporal
Substitution, and Risk-Aversion.” *American
Economic Review (Papers and Proceedings)*, May
2003, 93(2), pp. 383-91.

Weil, Philippe. “The Equity Premium Puzzle and the
Risk-Free Rate Puzzle.” *Journal of Monetary
Economics*, November 1989, 24(3), pp. 401-21.



Commentary

Thomas J. Sargent

In several recent papers—including the paper from this conference, Bansal (2007)—Ravi Bansal and his coauthors have constructed an interpretation of some asset pricing puzzles that I think macroeconomists should pay attention to. Why? Because a representative agent’s consumption Euler equation that links a one-period real interest rate to the consumption growth rate is the “IS curve” that is central to the policy transmission mechanism in today’s New Keynesian models. A long list of empirical failures called puzzles come from applying the stochastic discount factor implied by that Euler equation. Until we succeed in getting a consumption-based asset pricing model that works well, the New Keynesian IS curve is built on sand.

In several exciting papers, Bansal and his coauthors propose a way to explain some of those asset pricing puzzles by (i) specifying the intertemporal structure of risks to put long-run risks into consumption and assets’ cash flows and (ii) altering preferences to make the representative consumer care more about those long-run risks.

LONG-RUN RISK

Let c_t be the logarithm of aggregate consumption. A workhorse model that does a good job of fitting per capita U.S. consumption of nondurables and services makes c_t a random walk with constant drift and i.i.d. Gaussian innovations. Bansal

and his coauthors begin from the observation that it is difficult to distinguish that specification from an alternative one in which the drift in log consumption growth is itself a highly persistent covariance stationary process with low conditional volatility but high unconditional volatility. Thus, the drift itself is almost but not quite a random walk. The high unconditional volatility of the drift confronts the representative consumer with what Bansal and his coauthors call long-run risk because the conditional mean of consumption growth is not constant but wanders.

Bansal also posits that cash flows on particular portfolios differ in the extent to which they are subject to long-run risks that are more or less correlated with the long-run risk in aggregate consumption. For example, Bansal and coauthors as well as Hansen, Heaton, and Li (2006) have offered evidence that the cash flows from Fama and French’s high book-to-market portfolios have long-run components that are more highly correlated with long-run components of consumption than are the cash flows from low book-to-market portfolios. Can the need to compensate the representative consumer for that higher long-run correlation with consumption explain why those high book-to-market portfolios have higher returns?

PREFERENCES

The answer is no if the representative consumer’s preferences are usual ones assumed by

Thomas J. Sargent is a professor of economics and business at New York University, a senior fellow at the Hoover Institution, and the president of the American Economic Association.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 301-03.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Sargent

macroeconomists—for example, time-separable logarithmic preferences with discount factor $\beta \in (0, 1)$. Why? Those preferences lead to the usual stochastic discount factor whose logarithm is

$$s_{t+1,t} = \log \beta - \Delta c_{t+1},$$

which has the property that the representative consumer just doesn't care enough about those long-run risks to pump up the returns on those high book-to-market portfolios. Therefore, Bansal and his coauthors adopt a preference specification of Epstein and Zin (1989) that separates the intertemporal elasticity of substitution (IES) from the reciprocal of the coefficient of relative risk aversion, both of which are unity for the log preference specification mentioned above. With the IES held fixed at 1, those preferences imply a stochastic discount factor whose logarithm is

$$s_{t+1,t} = \log \beta - \Delta c_{t+1} - (\gamma - 1) \sum_{j=0}^{\infty} (E_{t+1} - E_t) \Delta c_{t+j+1} - \frac{1}{2} (\gamma - 1)^2 \text{var}_t \left(\sum_{j=0}^{\infty} \beta^j (E_{t+1} - E_t) \Delta c_{t+j+1} \right),$$

where γ is the coefficient of relative risk aversion and E_t denotes mathematical expectation conditioned on time- t information. Setting $\gamma > 1$ adds forward-looking terms to the stochastic discount factor that make the representative consumer care today about rates of log consumption growth far in the future. For a γ high enough, the representative consumer does have to be compensated for long-run cash flow risk correlated with long-run consumption growth risk in amounts that are big enough to explain how the market prices those Fama and French portfolios. Furthermore, Tallarini (2000), Bansal and Yaron (2004) and others have shown that with a high enough γ this kind of preference specification can provide a neat explanation for both the risk-free rate and the equity premium puzzles.

ASSESSMENT

It is possible to reinterpret the above stochastic discount factor with IES = 1 and atemporal

risk aversion $\gamma > 1$ in terms of a representative consumer who has IES and risk aversion both equal to 1, but where now $\gamma > 1$ measures his doubts about the stochastic specification of his model for consumption growth and cash flows. This reinterpretation is achieved by noting that the continuation value in Epstein and Zin's formulas equals the indirect utility function for a robust valuation problem in which a malevolent nature helps the decisionmaker construct valuations that are robust to misspecification by choosing a worst-case model from a set of models surrounding the decisionmaker's approximation model. Now γ acquires the interpretation of a penalty parameter on the relative entropy between the approximating and the distorted model. The additional terms in the log stochastic discount factor that appear when $\gamma > 1$ encode the likelihood ratio of the worst case to the approximating model.¹ This reinterpretation is of special interest for the work of Bansal and his coauthors, who reason as follows:

1. Statistically, it is difficult to distinguish a stochastic specification with long run-risk from one without it.
2. Therefore, without attributing wacky ideas to their representative consumer, Bansal and coauthors can assume that the representative consumer assigns probability 1 to the long-run risk model and takes our original i.i.d. log consumption growth model off the table.
3. Besides, by using the rational expectations cross-equation restrictions associated with the consumption Euler equation, we can infer that the representative consumer has to believe the long-run risk specification to explain the asset pricing data.

There is more to say here. Although long-run risks are difficult to detect (assertion 1), Epstein-Zin preferences or concerns about robustness make it vital for the representative consumer to

¹ Barillas, Hansen, and Sargent (2007) show that by interpreting γ as measuring fear of model misspecification rather than risk-aversion, a moderate fear of model misspecification can do most of the job of the large risk-aversion parameters of Tallarini (2000) and Bansal and Yaron (2004) in explaining the equity premium.

care about them. The tenuous part of this argument is how the representative consumer can come to be sure about the presence of these long-run risks when they are so difficult to detect statistically. Assertion 3 differs from arguments in the least-squares learning literature that typically have an agent learn about a forcing process by way of a least-squares learning algorithm on that process itself, not by using prices and the rational expectations cross-equation restrictions to reverse engineer what that process must be.

The robustness interpretation can help with these learning issues. Hansen and Sargent (2007) address these in the context of a model with a representative consumer who responds to assertion 1 by leaving both the i.i.d. and long-run risk models for log consumption growth on the table, attaching a prior initialized at the equal ignorance value of 0.5 to the long-run risk model, then updating by Bayes' law. We show that a consumer who distrusts both submodels and the posterior over submodels that emerges from Bayes' law will behave in a way that supports much of what Bansal and his coauthors do. In particular, because the long-run risk model is worse for the representative consumer, his worst-case probabilities become slanted toward that model and possibly put almost all of the mass on that model. This is nice because it provides an alternative defense of Bansal's assumption that the representative consumer acts as if he puts probability 1 on the long-run risk model.

But the structure of Hansen and Sargent (2007) yields other interesting outcomes, too. Even when the robust investor slants his worst-case probability to put probability 1 on the long-run risk model, the gap between the ordinary Bayesian probability and this worst-case probability contributes a potential source of time-varying countercyclical risk premia.

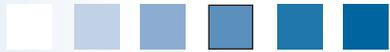
CONCLUSION

Bansal and Yaron's idea of stressing long-run risks that are difficult to detect but, with Epstein-Zin preferences or fear of model misspecification,

easy to care about is worth taking seriously. When many macroeconomists are now busy attaching loosely interpreted shocks or wedges to agents' first-order condition to make our dynamic stochastic general equilibrium (DSGE) models fit better, I welcome Bansal's new approach.

REFERENCES

- Bansal, Ravi. "Long-Run Risks and Financial Markets." Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 283-99.
- Bansal, Ravi and Yaron, Amir. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, August 2004, 59(4), pp. 1481-509.
- Barillas, Francisco; Hansen, Lars Peter and Sargent, Thomas J. "Doubts or Variability?" Working paper, University of Chicago and New York University, 2007.
- Epstein, Larry G. and Zin, Stanley E. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, July 1989, 57(4), pp. 937-69.
- Hansen, Lars Peter; Heaton, John and Li, Nan. "Consumption Strikes Back: Measuring Long Run Risk." Working paper, University of Chicago, 2006.
- Hansen, Lars Peter and Sargent, Thomas J. "Fragile Beliefs and the Price of Model Uncertainty." Unpublished manuscript, 2007.
- Tallarini, Thomas D. Jr. "Risk-Sensitive Real Business Cycles." *Journal of Monetary Economics*, June 2000, 45(3), pp. 507-32.



Arbitrage-Free Bond Pricing with Dynamic Macroeconomic Models

Michael F. Gallmeyer, Burton Hollifield, Francisco J. Palomino, and Stanley E. Zin

The authors examine the relationship between changes in short-term interest rates induced by monetary policy and the yields on long-maturity default-free bonds. The volatility of the long end of the term structure and its relationship with monetary policy are puzzling from the perspective of simple structural macroeconomic models. The authors explore whether richer models of risk premiums, specifically stochastic volatility models combined with Epstein-Zin recursive utility, can account for such patterns. They study the properties of the yield curve when inflation is an exogenous process and compare this with the yield curve when inflation is endogenous and determined through an interest rate (Taylor) rule. When inflation is exogenous, it is difficult to match the shape of the historical average yield curve. Capturing its upward slope is especially difficult because the nominal pricing kernel with exogenous inflation does not exhibit any negative autocorrelation—a necessary condition for an upward-sloping yield curve, as shown in Backus and Zin. Endogenizing inflation provides a substantially better fit of the historical yield curve because the Taylor rule provides additional flexibility in introducing negative autocorrelation into the nominal pricing kernel. Additionally, endogenous inflation provides for a flatter term structure of yield volatilities, which better fits historical bond data. (JEL G0, G1, E4)

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 305-26.

The response of long-term interest rates to changes in short-term interest rates is a feature of the economy that often puzzles policymakers. For example, in remarks made on May 27, 1994, Alan Greenspan (1994, p. 5) expressed concern that long rates moved too much in response to an increase in short rates:

In early February, we had thought long-term rates would move a little higher temporarily as we tightened...The sharp jump in [long] rates that occurred appeared to reflect the dramatic rise in market expectations of economic growth and associated concerns about possible inflation pressures.

Then in his February 16, 2005, testimony, Chairman Greenspan (2005) expressed a completely different concern about long rates:

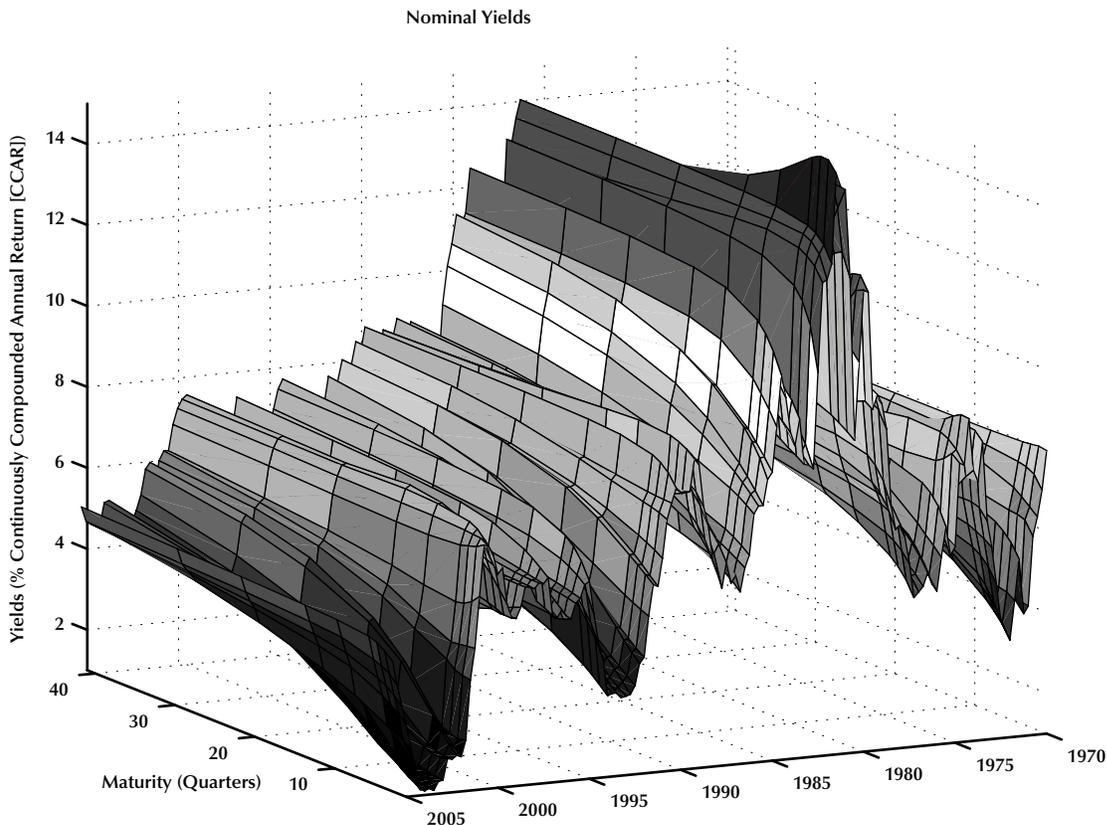
[L]ong-term interest rates have trended lower in recent months even as the Federal Reserve has raised the level of the target federal funds rate by 150 basis points...Historically, though, even these distant forward rates have tended to rise in association with monetary policy tightening...For the moment, the broadly unanticipated behavior of world bond markets remains a conundrum.

Chairman Greenspan's comments reflect the fact that we do not yet have a satisfactory under-

Michael F. Gallmeyer is an assistant professor of finance at the Mays Business School, Texas A&M University; Burton Hollifield is an associate professor of financial economics at the David A. Tepper School of Business, Carnegie Mellon University; Francisco J. Palomino is a PhD candidate at the David A. Tepper School of Business, Carnegie Mellon University; and Stanley E. Zin is a professor of economics at the David A. Tepper School of Business, Carnegie Mellon University, and a research fellow at the National Bureau of Economic Research. The authors thank David Backus and Pamela Labadie for valuable comments and suggestions, Bill Gavin for suggesting the topic, Monika Piazzesi and Chris Telmer for providing data, and Zhanhui Chen for excellent research assistance.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Figure 1
Time-Series Properties of the Yield Curve, 1970:Q1–2005:Q4



standing of how the yield curve is related to the structural features of the macroeconomy, such as investors’ preferences, the fundamental sources of economic risk, and monetary policy.

Figure 1 plots the nominal yield curve for a variety of maturities from 1 quarter—which we refer to as the short rate—up to 40 quarters for U.S. Treasuries, starting in the first quarter of 1970 and ending in the last quarter of 2005.¹ Figure 2 plots the average yield curve for the entire sample and for two subsamples. Figure 3 plots the standard deviation of yields against their maturities. Two basic patterns of yields are clear from these figures: (i) On average, the yield curve is upward

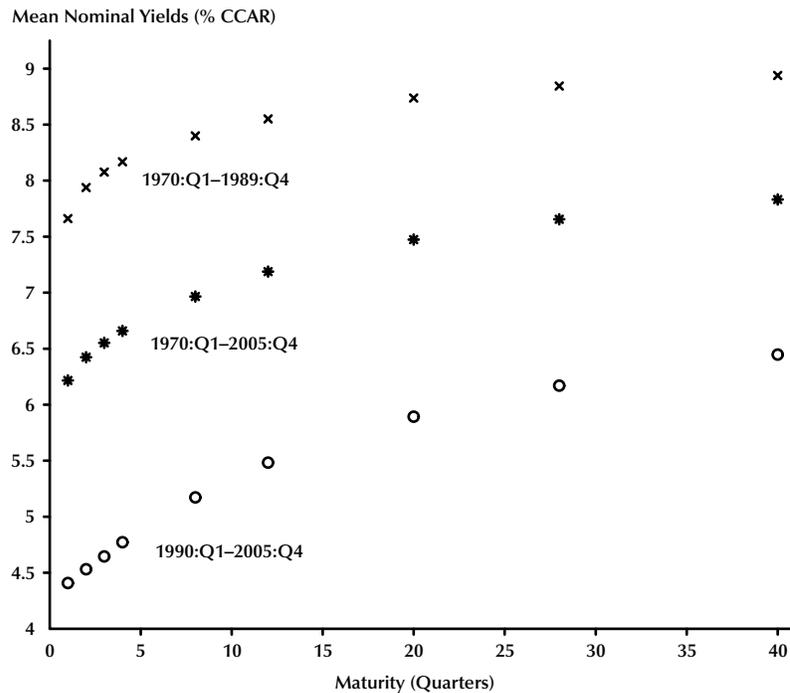
sloping and (ii) there is substantial volatility in yields at all maturities. Chairman Greenspan’s comments, therefore, must be framed by the fact that long yields are almost as volatile as short rates. The issue, however, is the relationship of the volatility at the long end to the volatility at the short end, and the correlation between changes in short-term interest rates and changes in long-term interest rates.

We can decompose forward interest rates into expectations of future short-term interest rates and interest rate risk premia. Because long-term interest rates are averages of forward rates, long-run interest rates depend on expectations of future short-term interest rates and interest rate risk premiums. A significant component of long rates is

¹ Yields up to 1991 are from McCulloch and Kwon (1993) then Datastream from 1991 to 2005.

Figure 2

Average Yield-Curve Behavior



the risk premium, and there is now a great deal of empirical evidence documenting that the risk premiums are time-varying and stochastic. Movements in long rates can therefore be attributed to movements in expectations of future nominal short rates, movements in risk premiums, or some combination of movements in both.

Moreover, if monetary policy is implemented using a short-term interest rate feedback rule—for example, a Taylor rule—then inflation rates must adjust so that the bond market clears. The resulting endogenous equilibrium inflation rate will then depend on the same risk factors that drive risk premiums in long rates. Monetary policy itself, therefore, could be a source of fluctuations in the yield curve in equilibrium.

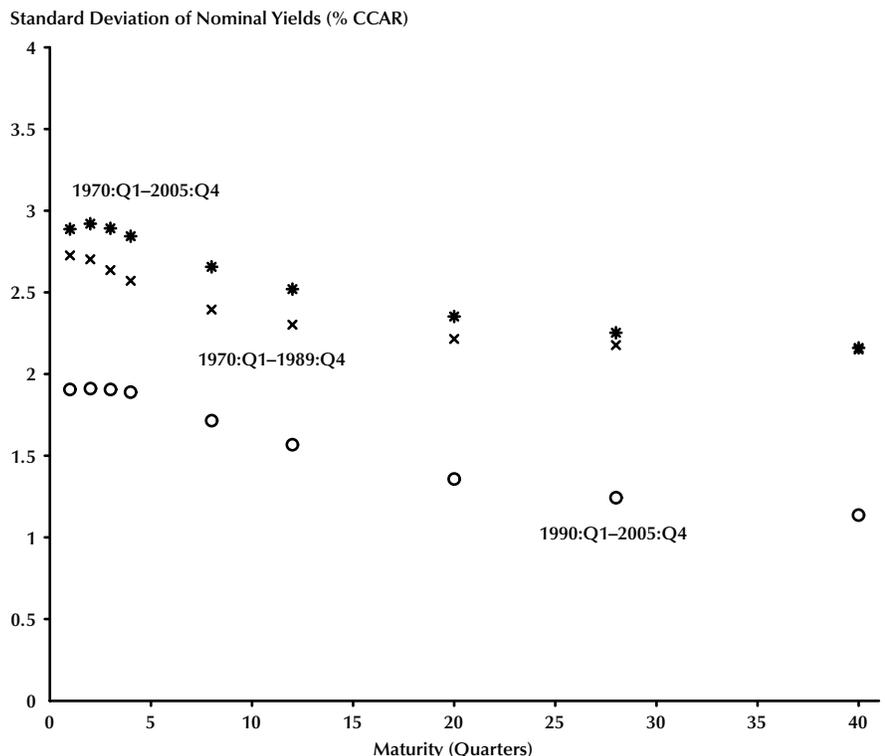
We explore such possibilities in a model of time-varying risk premiums generated by the recursive utility model of Epstein and Zin (1989) combined with stochastic volatility of endowment

growth. We show how the model can be easily solved using now-standard affine term-structure methods. Affine term-structure models have the convenient property that yields are maturity-dependent linear functions of state variables. We examine some general properties of multi-period default-free bonds in our model, assuming first that inflation is an exogenous process and by allowing inflation to be endogenous and determined by an interest rate feedback rule. We show that the interest rate feedback rule—the form of monetary policy—can have significant impacts on properties of the term structure of interest rates.

THE DUFFIE-KAN AFFINE TERM-STRUCTURE MODEL

The Duffie and Kan (1996) class of affine term-structure models, translated into discrete time

Figure 3
Volatility of Yields of Various Maturities



by Backus, Foresi, and Telmer (2001), is based on a k -dimensional vector of state variables z that follows a “square root” model:

$$z_{t+1} = (I - \Phi)\theta + \Phi z_t + \Sigma(z_t)^{1/2} \varepsilon_{t+1},$$

where $\{\varepsilon_t\} \sim \text{NID}(0, 1)$, $\Sigma(z)$ is a diagonal matrix with a typical element given by $\sigma_i(z) = a_i + b_i'z$, where b_i has nonnegative elements, and Φ is stable with positive diagonal elements. The process for z requires that the volatility functions, $\sigma_i(z)$, be positive, which places additional restrictions on the parameters.

The asset-pricing implications of the model are given by the pricing kernel, m_{t+1} , a positive random variable that prices all financial assets. That is, if a security has a random payoff, h_{t+1} , at date $t + 1$, then its date- t price is $E_t[m_{t+1}h_{t+1}]$. The pricing kernel in the affine model takes the form

$$-\log m_{t+1} = \delta + \gamma'z_t + \lambda'\Sigma(z_t)^{1/2} \varepsilon_{t+1},$$

where the $k \times 1$ vector γ is referred to as the “factor loadings” for the pricing kernel, the $k \times 1$ vector λ is referred to as the “price of risk” vector because it controls the size of the conditional correlation of the pricing kernel and the underlying sources of risk, and the $k \times k$ matrix $\Sigma(z_t)$ is the stochastic variance-covariance matrix of the unforecastable shock.

Let $b_t^{(n)}$ be the price at date t of a default-free pure-discount bond that pays 1 at date $t + n$, with $b_t^{(0)} = 1$. Multi-period default-free discount bond prices are built up using the arbitrage-free pricing restriction,

$$(1) \quad b_t^{(n)} = E_t \left[m_{t+1} b_{t+1}^{(n-1)} \right].$$

Bond prices of all maturities are log-linear functions of the state:

$$-\log b_t^{(n)} = A^{(n)} + B^{(n)}z_t,$$

where $A^{(n)}$ is a scalar and $B^{(n)}$ is a $1 \times k$ row vector.

The intercept and slope parameters, which we often refer to as “yield-factor loadings,” of these bond prices can be found recursively according to

(2)

$$A^{(n+1)} = A^{(n)} + \delta + B^{(n)}(I - \Phi)\theta - \frac{1}{2} \sum_{j=1}^k (\lambda_j + B_j^{(n)})^2 a_j,$$

$$B^{(n+1)} = (\gamma' + B^{(n)}\Phi) - \frac{1}{2} \sum_{j=1}^k (\lambda_j + B_j^{(n)})^2 b_j',$$

where $B_j^{(n)}$ is the j th element of the vector $B^{(n)}$. Because $b^{(0)} = 1$, we can start these recursions using $A^{(0)} = 0$ and $B_j^{(0)} = 0, j = 1, 2, \dots, k$.

Continuously compounded yields, $y_t^{(n)}$, are defined by $b_t^{(n)} = \exp(-ny_t^{(n)})$, which implies $y_t^{(n)} = -(\log b_t^{(n)})/n$. We refer to the short rate, i_t , as the one-period yield: $i_t \equiv y_t^{(1)}$.

This is an arbitrage-free model of bond pricing because it satisfies equation (1) for a given pricing kernel m_t . It is not yet a structural equilibrium model, because the mapping of the parameters of the pricing model to deeper structural parameters of investors’ preferences and opportunities has not yet been specified. The equilibrium structural models we consider will all lie within this general class, hence, can be easily solved using these pricing equations.

A TWO-FACTOR MODEL WITH EPSTEIN-ZIN PREFERENCES

We begin our analysis of structural models of the yield curve by solving for equilibrium real yields in a representative-agent exchange economy. Following Backus and Zin (2006), we consider a representative agent who chooses consumption to maximize the recursive utility function given in Epstein and Zin (1989). Given a sequence of consumption, $\{c_t, c_{t+1}, c_{t+2}, \dots\}$, where future consumptions can be random outcomes,

the intertemporal utility function, U_t , is the solution to the recursive equation,

$$(3) \quad U_t = \left[(1 - \beta)c_t^\rho + \beta\mu_t(U_{t+1})^\rho \right]^{1/\rho},$$

where $0 < \beta < 1$ characterizes impatience (the marginal rate of time preference is $1 - 1/\beta$), $\rho \leq 1$ measures the preference for intertemporal substitution (the elasticity of intertemporal substitution for deterministic consumption paths is $1/(1 - \rho)$), and the certainty equivalent of random future utility is

$$(4) \quad \mu_t(U_{t+1}) \equiv E_t \left[U_{t+1}^\alpha \right]^{1/\alpha},$$

where $\alpha \leq 1$ measures static risk aversion (the coefficient of relative risk aversion for static gambles is $1 - \alpha$). The marginal rate of intertemporal substitution, m_{t+1} , is

$$m_{t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{\rho-1} \left(\frac{U_{t+1}}{\mu_t(U_{t+1})} \right)^{\alpha-\rho}.$$

Time-additive expected utility corresponds to the parameter restriction $\rho = \alpha$.

In equilibrium, the representative agent consumes the stochastic endowment, e_t , so that, $\log(c_{t+1}/c_t) = \log(e_{t+1}/e_t) = x_{t+1}$, where x_{t+1} is the log of the ratio of endowments in $t + 1$ relative to t . The log of the equilibrium marginal rate of substitution, referred to as the real pricing kernel, is therefore given by

$$(5) \quad \begin{aligned} \log m_{t+1} &= \log \beta + (\rho - 1)x_{t+1} \\ &+ (\alpha - \rho)[\log W_{t+1} - \log \mu_t(W_{t+1})], \end{aligned}$$

where W_t is the value of utility in equilibrium.

The first two terms in the marginal rate of substitution are standard expected utility terms: the pure time preference parameter, β , and a consumption growth term times the inverse of the negative of the intertemporal elasticity of substitution. The third term in the pricing kernel is a new term coming from the Epstein-Zin preferences.

The endowment-growth process evolves stochastically according to

$$x_{t+1} = (1 - \phi_x)\theta_x + \phi_x x_t + v_t^{1/2} \varepsilon_{t+1}^x,$$

where

$$v_{t+1} = (1 - \phi_v)\theta_v + \phi_v v_t + \sigma_v \varepsilon_{t+1}^v$$

is the process for the conditional volatility of endowment growth. We will refer to v_t as stochastic volatility. The innovations ε_t^x and ε_t^v are distributed NID(0,1).

Note that the state vector in this model conforms with the setup of the Duffie-Kan model above. Define the state vector $z_t \equiv [x_t \ v_t]'$, which implies parameters for the Duffie-Kan model:

$$\theta = [\theta_x \ \theta_v]'$$

$$\Phi = \text{diag}\{\phi_x, \phi_v\}$$

$$\Sigma(z_t) = \text{diag}\{a_1 + b_1'z_t, a_2 + b_2'z_t\}$$

$$a_1 = 0, b_1 = [0 \ 1]', a_2 = \sigma_v^2, b_2 = [0 \ 0]'$$

Following the analysis in Hansen, Heaton, and Li (2005), we will work with the logarithm of the value function scaled by the endowment:

$$(6) \quad \begin{aligned} W_t / e_t &= \left[(1 - \beta) + \beta (\mu_t(W_{t+1}) / e_t)^\rho \right]^{1/\rho} \\ &= \left[(1 - \beta) + \beta \left(\mu_t \left(\frac{W_{t+1}}{e_{t+1}} \times \frac{e_{t+1}}{e_t} \right)^\rho \right) \right]^{1/\rho}, \end{aligned}$$

where we have used the linear homogeneity of μ_t (see equation (4)). Take logarithms of (6) to obtain

$$w_t = \rho^{-1} \log \left[(1 - \beta) + \beta \exp(\rho u_t) \right],$$

where $w_t \equiv \log(W_t/e_t)$ and $u_t \equiv \log(\mu_t(\exp(w_{t+1} + x_{t+1})))$. Consider a linear approximation of the right-hand side of this equation as a function of u_t around the point \bar{m} :

$$\begin{aligned} w_t &\approx \rho^{-1} \log \left[(1 - \beta) + \beta \exp(\rho \bar{m}) \right] \\ &\quad + \left[\frac{\beta \exp(\rho \bar{m})}{1 - \beta + \beta \exp(\rho \bar{m})} \right] (u_t - \bar{m}) \\ &\equiv \bar{\kappa} + \kappa u_t, \end{aligned}$$

where $\kappa < 1$. For the special case with $\rho = 0$, that is, a log time aggregator, the linear approximation is exact, implying $\bar{\kappa} = 1 - \beta$ and $\kappa = \beta$ (see Hansen, Heaton, and Li, 2005). Similarly, approximating around $\bar{m} = 0$, results in $\bar{\kappa} = 0$ and $\kappa = \beta$.

Given the state variables and the log-linear structure of the model, we conjecture a solution for the log value function of the form

$$w_t = \bar{\omega} + \omega_x x_t + \omega_v v_t,$$

where $\bar{\omega}$, ω_x , and ω_v are constants to be determined. By substituting,

$$w_{t+1} + x_{t+1} = \bar{\omega} + (\omega_x + 1)x_{t+1} + \omega_v v_{t+1}.$$

Because x_{t+1} and v_{t+1} are jointly normally distributed, the properties of normal random variables can be used to solve for u_t :

$$\begin{aligned} u_t &\equiv \log \left(\mu_t \left(\exp(w_{t+1} + x_{t+1}) \right) \right) \\ &= \log \left(E_t \left[\exp(w_{t+1} + x_{t+1})^\alpha \right]^{\frac{1}{\alpha}} \right) \\ &= E_t [w_{t+1} + x_{t+1}] + \frac{\alpha}{2} \text{Var}_t [w_{t+1} + x_{t+1}] \\ &= \bar{\omega} + (\omega_x + 1)(1 - \phi_x)\theta_x + \omega_v (1 - \phi_v)\theta_v \\ &\quad + (\omega_x + 1)\phi_x x_t + \omega_v \phi_v v_t \\ &\quad + \frac{\alpha}{2} (\omega_x + 1)^2 v_t + \frac{\alpha}{2} \omega_v^2 \sigma_v^2. \end{aligned}$$

We can use the above expression to solve for the value-function parameters and verify its log-linear solution:

$$\begin{aligned} \omega_x &= \kappa (\omega_x + 1) \phi_x \\ \Rightarrow \omega_x &= \left(\frac{\kappa}{1 - \kappa \phi_x} \right) \phi_x \\ \omega_v &= \kappa \left[\omega_v \phi_v + \frac{\alpha}{2} (\omega_x + 1)^2 \right] \\ \Rightarrow \omega_v &= \left(\frac{\kappa}{1 - \kappa \phi_v} \right) \left[\frac{\alpha}{2} \left(\frac{1}{1 - \kappa \phi_x} \right)^2 \right] \\ \bar{\omega} &= \frac{\bar{\kappa}}{1 - \kappa} + \frac{\kappa}{1 - \kappa} \left[(\omega_x + 1)(1 - \phi_x)\theta_x \right. \\ &\quad \left. + \omega_v (1 - \phi_v)\theta_v + \frac{\alpha}{2} \omega_v^2 \sigma_v^2 \right]. \end{aligned}$$

The solution allows us to simplify the term $[\log W_{t+1} - \log \mu_t(W_{t+1})]$ in the real pricing kernel in equation (5):

$$\begin{aligned} \log W_{t+1} - \log \mu_t(W_{t+1}) &= w_{t+1} + x_{t+1} - \log \mu_t(w_{t+1} + x_{t+1}) \\ &= (\omega_x + 1)[x_{t+1} - E_t x_{t+1}] + \omega_v[v_{t+1} - E_t v_{t+1}] \\ &\quad - \frac{\alpha}{2}(\omega_x + 1)^2 \text{Var}_t[x_{t+1}] - \frac{\alpha}{2}\omega_v^2 \text{Var}_t[v_{t+1}] \\ &= (\omega_x + 1)v_t^{1/2}\varepsilon_{t+1}^x + \omega_v\sigma_v\varepsilon_{t+1}^v - \frac{\alpha}{2}(\omega_x + 1)^2 v_t - \frac{\alpha}{2}\omega_v^2\sigma_v^2. \end{aligned}$$

The real pricing kernel, therefore, is a member of the Duffie-Kan class with two factors and parameters:

$$\begin{aligned} \delta &= -\log(\beta) + (1 - \rho)(1 - \phi_x)\theta_x + \frac{\alpha}{2}(\alpha - \rho)\omega_v^2\sigma_v^2 \\ \gamma &= [\gamma_x \quad \gamma_v]' \\ &= \left[(1 - \rho)\phi_x \quad \frac{\alpha}{2}(\alpha - \rho)\left(\frac{1}{1 - \kappa\phi_x}\right)^2 \right]' \\ (7) \quad \lambda &= [\lambda_x \quad \lambda_v]' \\ &= \left[(1 - \alpha) - (\alpha - \rho)\left(\frac{\kappa\phi_x}{1 - \kappa\phi_x}\right) \right. \\ &\quad \left. - \left(\frac{\alpha}{2}\right)\left(\frac{\kappa(\alpha - \rho)}{1 - \kappa\phi_v}\right)\left(\frac{1}{1 - \kappa\phi_x}\right)^2 \right]. \end{aligned}$$

We can now use the recursive formulas in equation (2) to solve for real discount bond prices and the real yield curve.

Note how the factor loadings and prices of risk depend on the deeper structural parameters and the greatly reduced dimensionality of the parameter space relative to the general affine model. Also, for the time-additive expected utility special case, $\alpha = \rho$, the volatility factor does not enter the conditional mean of the pricing kernel because $\gamma_v = 0$; and the price of risk for the volatility factor is zero because $\lambda_v = 0$. Finally, we can see from the expressions for bond prices that the two key preference parameters, ρ and α , provide freedom in controlling both the factor loadings and the prices of risk in the real pricing kernel.

NOMINAL BOND PRICING

To understand the price of nominal bonds, we need a nominal pricing kernel. If we assume

that there is a frictionless conversion of money for goods in this economy, the nominal kernel is given by

$$(8) \quad \log(m_{t+1}^s) = \log(m_{t+1}) - p_{t+1},$$

where p_{t+1} is the log of the money price of goods at time $t+1$ relative to the money price of goods at time t , that is, the inflation rate between t and $t+1$. Clearly then, the source of inflation, its random properties, and its relationship to the real pricing kernel is of central interest for nominal bond pricing. We next consider two different specifications for equilibrium inflation.

Exogenous Inflation

If we expand the state space to include an exogenous inflation process, p_t , the state vector becomes $z_t = [x_t \ v_t \ p_t]'$. The stochastic process for exogenous inflation is

$$p_{t+1} = (1 - \phi_p)\theta_p + \phi_p p_t + \sigma_p \varepsilon_{t+1}^p,$$

where ε_{t+1}^p is also normally distributed independently of the other two shocks. In this case, the parameters for the affine nominal pricing kernel are

$$\begin{aligned} \delta^s &= \delta + (1 - \phi_p)\theta_p \\ \gamma^s &= [\gamma_x \quad \gamma_v \quad \phi_p]' \\ \lambda^s &= [\lambda_x \quad \lambda_v \quad 1]'. \end{aligned}$$

In the exogenous inflation model, the price of inflation risk is always exactly 1 and does not change with the values of any of the other structural parameters in the model. In addition, the factor loadings and prices of risk for output growth and stochastic volatility are the same as in the real pricing kernel. We will refer to this nominal pricing kernel specification as the exogenous inflation economy.

Monetary Policy and Endogenous Inflation

We begin by assuming that monetary policy follows a simple nominal interest rate rule. We will abuse conventional terminology and often

refer to the interest rate rule as a Taylor rule. Although there are a variety of ways to specify a Taylor rule—see Ang, Dong, and Piazzesi (2004)—we will consider a rule in which the short-term interest rate depends on contemporaneous output, inflation, and a policy shock:

$$(9) \quad i_t = \bar{\tau} + \tau_x x_t + \tau_p p_t + s_t,$$

where the monetary policy shock satisfies

$$s_t = \phi_s s_{t-1} + \sigma_s \varepsilon_t^s$$

and where $\varepsilon_t^s \sim \text{NID}(0,1)$ is independent of the other two real shocks.

Because this nominal interest rate rule must also be consistent with equilibrium in the bond market, that is, it must be consistent with the nominal pricing kernel in equation (8) as well as equation (9), we can use these two equations to find the equilibrium process for inflation. Conjecture a log-linear solution for p_t ,

$$(10) \quad p_t = \bar{\pi} + \pi_x x_t + \pi_v v_t + \pi_s s_t,$$

with $\bar{\pi}$, π_x , and π_s constants to be solved.

To solve for a rational expectations solution to the model, we substitute the guess for the inflation rate into both the Taylor rule and the nominal pricing kernel and solve for the parameters $\bar{\pi}$, π_x , π_v , and π_s that equate the short rate determined by the pricing kernel with the short rate determined by the Taylor rule.

From the dynamics of x_{t+1} , v_{t+1} , and s_{t+1} , inflation, p_{t+1} , is given by

$$\begin{aligned} p_{t+1} &= \bar{\pi} + \pi_x x_{t+1} + \pi_v v_{t+1} + \pi_s s_{t+1} \\ &= \bar{\pi} + \pi_x (1 - \phi_x) \theta_x + \pi_v (1 - \phi_v) \theta_v \\ &\quad + \pi_x \phi_x x_t + \pi_v \phi_v v_t + \pi_s \phi_s s_t \\ &\quad + \pi_x v_t^{1/2} \varepsilon_{t+1}^x + \pi_v \sigma_v \varepsilon_{t+1}^v + \pi_s \sigma_s \varepsilon_{t+1}^s. \end{aligned}$$

Substituting into the nominal pricing kernel,

$$\begin{aligned} -\log(m_{t+1}^s) &= -\log(m_{t+1}) + p_{t+1} \\ &= \delta + \gamma_x x_t + \gamma_v v_t + \lambda_x v_t^{1/2} \varepsilon_{t+1}^x + \lambda_v \sigma_v \varepsilon_{t+1}^v + p_{t+1} \\ &= \delta + \bar{\pi} + \pi_x (1 - \phi_x) \theta_x + \pi_v (1 - \phi_v) \theta_v \\ &\quad + (\gamma_x + \pi_x \phi_x) x_t + (\gamma_v + \pi_v \phi_v) v_t + \pi_s \phi_s s_t \\ &\quad + (\lambda_x + \pi_x) v_t^{1/2} \varepsilon_{t+1}^x + (\lambda_v + \pi_v) \sigma_v \varepsilon_{t+1}^v + \pi_s \sigma_s \varepsilon_{t+1}^s. \end{aligned}$$

From these dynamics, the nominal one-period interest rate, $i_t = -\log(E_t[m_{t+1}^s])$, is

$$\begin{aligned} i_t &= \delta + \bar{\pi} + \pi_x (1 - \phi_x) \theta_x + \pi_v (1 - \phi_v) \theta_v \\ &\quad + (\gamma_x + \pi_x \phi_x) x_t + (\gamma_v + \pi_v \phi_v) v_t + \pi_s \phi_s s_t \\ &\quad - \frac{1}{2} (\lambda_x + \pi_x)^2 v_t - \frac{1}{2} (\lambda_v + \pi_v)^2 \sigma_v^2 - \frac{1}{2} \pi_s^2 \sigma_s^2. \end{aligned}$$

Comparing this with the interest rate rule,

$$i_t = \bar{\tau} + \tau_x x + \tau_p (\bar{\pi} + \pi_x x_t + \pi_v v_t + \pi_s s_t) + s_t,$$

gives the parameter restrictions consistent with equilibrium:

$$(11) \quad \begin{aligned} \pi_x &= \frac{\gamma_x - \tau_x}{\tau_p - \phi_x} \\ \pi_v &= \frac{\gamma_v - \frac{1}{2} (\lambda_x + \pi_x)^2}{\tau_p - \phi_v} \\ \pi_s &= -\frac{1}{\tau_p - \phi_s} \\ \bar{\pi} &= \frac{1}{\tau_p - 1} \left(\delta - \bar{\tau} + \pi_x (1 - \phi_x) \theta_x + \pi_v (1 - \phi_v) \theta_v \right. \\ &\quad \left. - \frac{1}{2} (\lambda_v + \pi_v)^2 \sigma_v^2 - \frac{1}{2} \pi_s^2 \sigma_s^2 \right). \end{aligned}$$

These expressions form a recursive system we use to solve for the equilibrium parameters of the inflation process. See Cochrane (2006) for a more detailed account of this rational expectations solution method.

It is clear from these expressions that the equilibrium inflation process will depend on the preference parameters of the household generally and attitudes toward risk specifically.

In a similar fashion, we can extend the analysis to any Taylor-type rule that is linear in the state variables, including (i) lagged short rates, (ii) other contemporaneous yields at any maturity, and (iii) forward-looking rules, such as those in Clarida, Galí, and Gertler (2000). Such extensions are possible because, in the affine framework, interest rates are all simply linear functions of the current state variables. See Ang, Dong, and Piazzesi (2004) and Gallmeyer, Hollifield, and Zin (2005) for some concrete examples.

A Monetary Policy–Consistent Pricing Kernel

Substituting the equilibrium inflation process from equations (10) and (11) into the nominal pricing kernel, we obtain an equilibrium three-factor affine term-structure model that is consistent with the nominal interest rate rule. The state space is

$$\begin{aligned} z_t &\equiv [x_t \ v_t \ s_t]' \\ \Phi &= \text{diag}\{\phi_x, \phi_v, \phi_s\} \\ \theta &= [\theta_x \ \theta_v \ 0]' \\ \Sigma(z_t) &= \text{diag}\{a_1 + b_1'z_t, a_2 + b_2'z_t, a_3 + b_3'z_t\} \\ a_1 &= 0, b_1 = [0 \ 1 \ 0]' \\ a_2 &= \sigma_v^2, b_2 = [0 \ 0 \ 0]' \\ a_3 &= \sigma_s^2, b_3 = [0 \ 0 \ 0]', \end{aligned}$$

and the parameters of the pricing kernel are

$$\begin{aligned} \delta^s &= \delta + \bar{\pi} + \pi_x(1 - \phi_x)\theta_x + \pi_v(1 - \phi_v)\theta_v \\ \gamma^s &= [\gamma_x + \phi_x\pi_x \ \gamma_v + \phi_v\pi_v \ \phi_s\pi_s]' \\ \lambda^s &= [\lambda_x + \pi_x \ \lambda_v + \pi_v \ \pi_s]'. \end{aligned}$$

We will often refer to this nominal pricing kernel specification as the endogenous inflation economy.

The Taylor rule parameters, through their determination of the equilibrium inflation process, affect both the factor loadings on the real factors as well as their prices of risk. Monetary policy through its effects on endogenous inflation, therefore, can result in risk premiums in the term structure that are significantly different from those in the exogenous inflation model. We explore such a possibility through numerical examples.

QUANTITATIVE EXERCISES

We calibrate the exogenous processes in our model to quarterly postwar U.S. data as follows:

1. Endowment growth: $\phi_x = 0.36$, $\theta_x = 0.006$, $\sigma_x = 0.0048(1 - \phi_x^2)^{1/2}$

2. Inflation: $\phi_p = 0.8471$, $\theta_p = 0.0093$, $\sigma_p = 0.0063(1 - \phi_p^2)^{1/2}$
3. Stochastic volatility: $\phi_v = 0.973$, $\theta_v = 0.0001825$, $\sigma_v = 0.9884 \times 10^{-5}$
4. Policy shock: $\phi_s = 0.922$, $\sigma_s = (0.023 \times 10^{-4})^{1/2}$

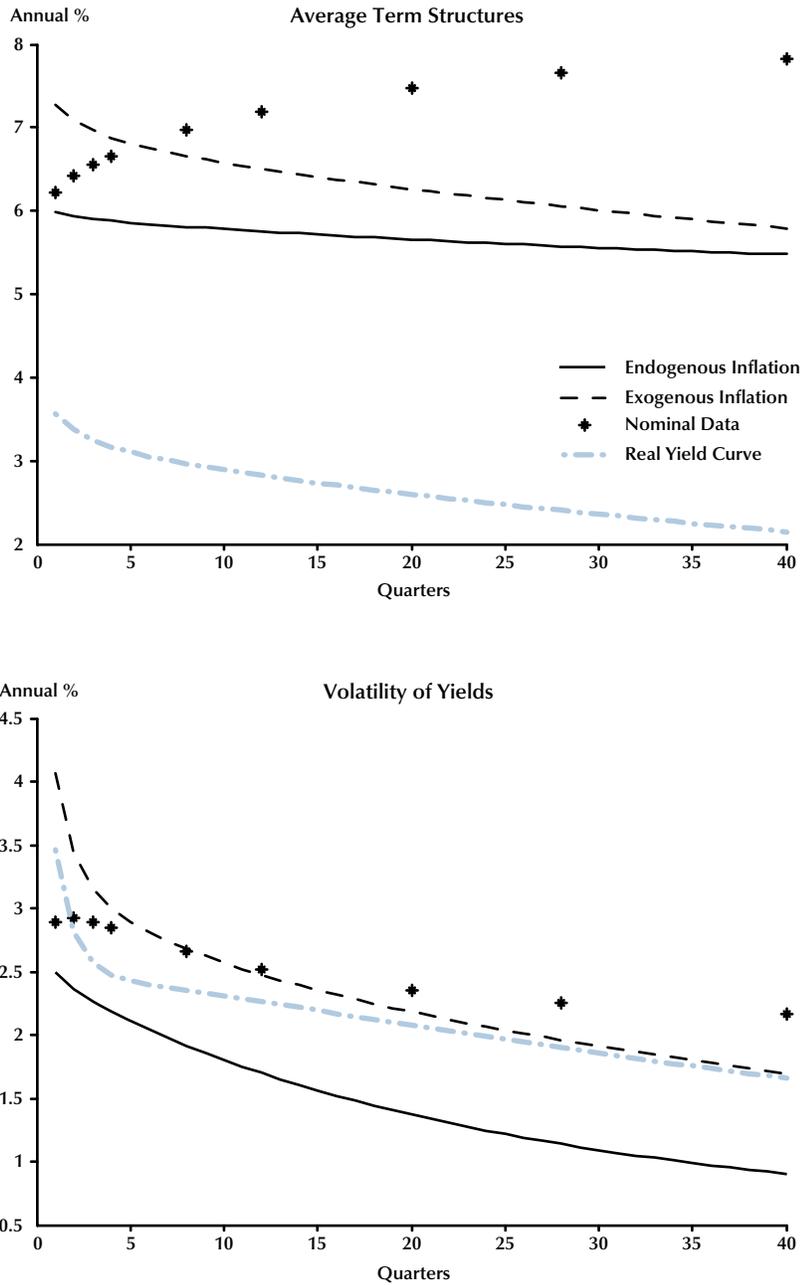
The endowment growth process is calibrated to quarterly per capita consumption of durable goods and services, and inflation is calibrated to the nondurables and services deflator, similarly to Piazzesi and Schneider (2007). The volatility process is taken from Bansal and Yaron (2004), who calibrate their model to monthly data. We adjust their parameters to deal with quarterly time-aggregation. We take the parameters for the policy shock from Ang, Dong, and Piazzesi (2004), who estimate a Taylor rule using an affine term-structure model with macroeconomic factors and an unobserved policy shock.

Figures 4 through 7 depict the average yield curves and yield volatilities for different preference parameters for the exogenous and endogenous inflation models. In the top panel of each figure, asterisks denote the empirical average nominal yield curve, a blue dashed-dotted line denotes the average real yield curve common across both inflation models, a dashed line denotes the average nominal yield curve in the exogenous inflation economy, and a solid line denotes the average nominal yield curve in the endogenous inflation economy. The bottom panel depicts yield volatilities for the same cases as the average yield curve in the top panel. (Asterisks in Figures 4 through 9 are the moments—means and standard deviations—of the data in Figure 1.)

Each figure is computed using a different set of preference parameters. We fix a level of the intertemporal elasticity parameter, ρ , for each panel and pick the remaining preference parameters—the risk aversion coefficient, α , and the rate of time preference, β —to minimize the distance between the average nominal yields and yield volatilities in the data and those implied by the exogenous inflation economy. We pick the Taylor rule parameters to minimize the distance between the average nominal yields and yield volatilities in the data and those implied by the endogenous inflation economy. Table 1 reports

Figure 4

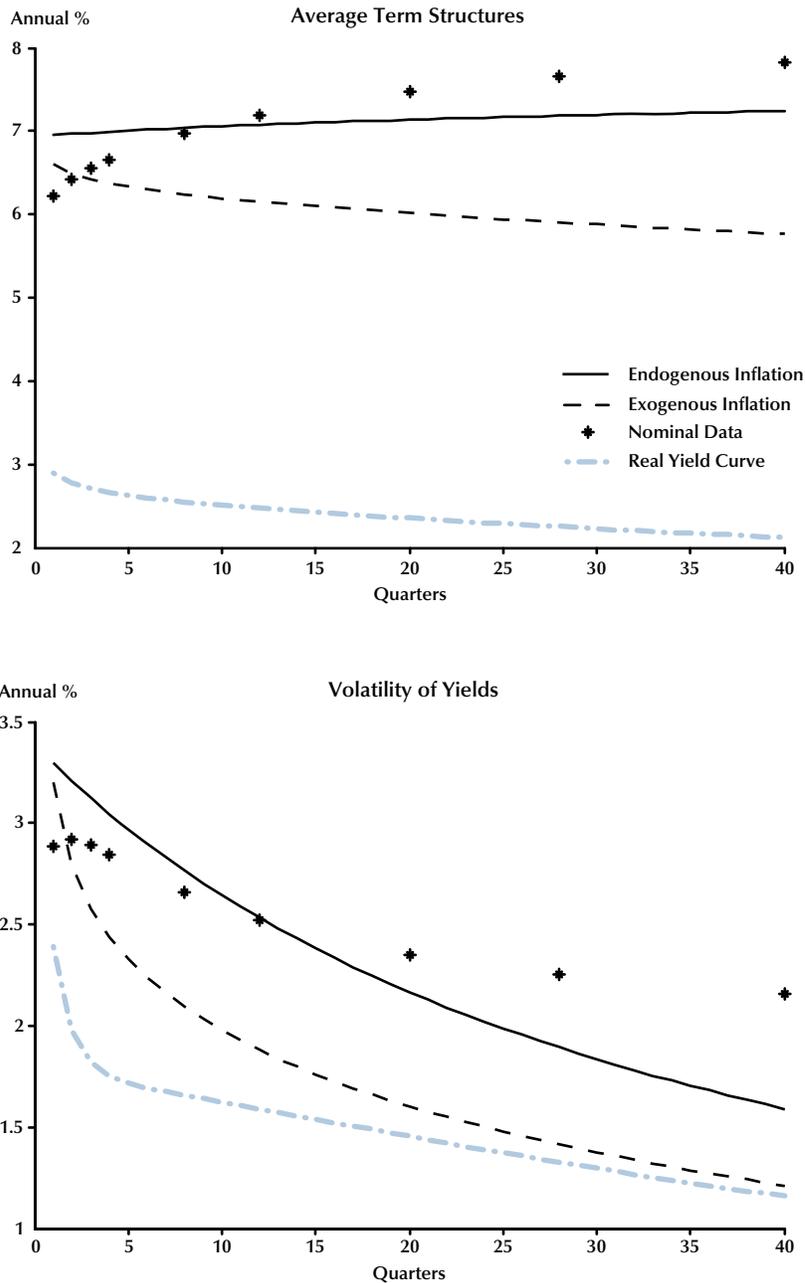
Average Yield Curve and Volatilities for the Epstein-Zin Model with Stochastic Volatility



NOTE: The parameters are $\rho = -0.5$, $\alpha = -4.835$, $\beta = 0.999$, $\bar{\tau} = 0.003$, $\tau_x = 1.2475$, and $\tau_p = 1.000$. The empirical moments for the full sample (1970:Q1–2005:Q4) are plotted with asterisks, properties of the real yield curve are plotted with a dashed-dotted blue line, properties of the yield curve in the exogenous inflation economy are plotted with a dashed black line, and properties of the yield curve in the economy with endogenous inflation are plotted with a solid black line.

Figure 5

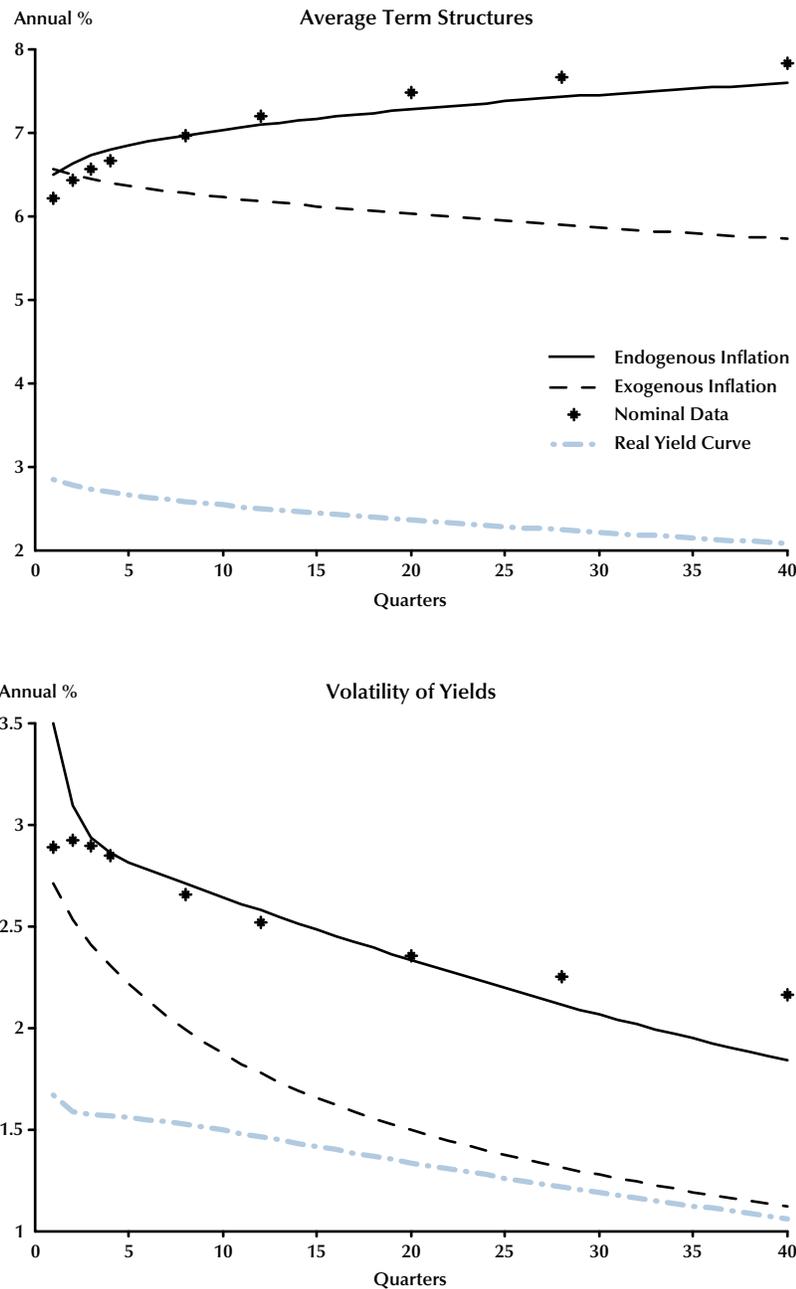
Average Yield Curve and Volatilities for the Epstein-Zin Model with Stochastic Volatility



NOTE: The parameters are $\rho = 0.0$, $\alpha = -4.061$, $\beta = 0.998$, $\bar{\tau} = 0.003$, $\tau_x = 0.973$, and $\tau_p = 0.973$. The empirical moments for the full sample (1970:Q1–2005:Q4) are plotted with asterisks, properties of the real yield curve are plotted with a dashed-dotted blue line, properties of the yield curve in the exogenous inflation economy are plotted with a dashed black line, and properties of the yield curve in the economy with endogenous inflation are plotted with a solid black line.

Figure 6

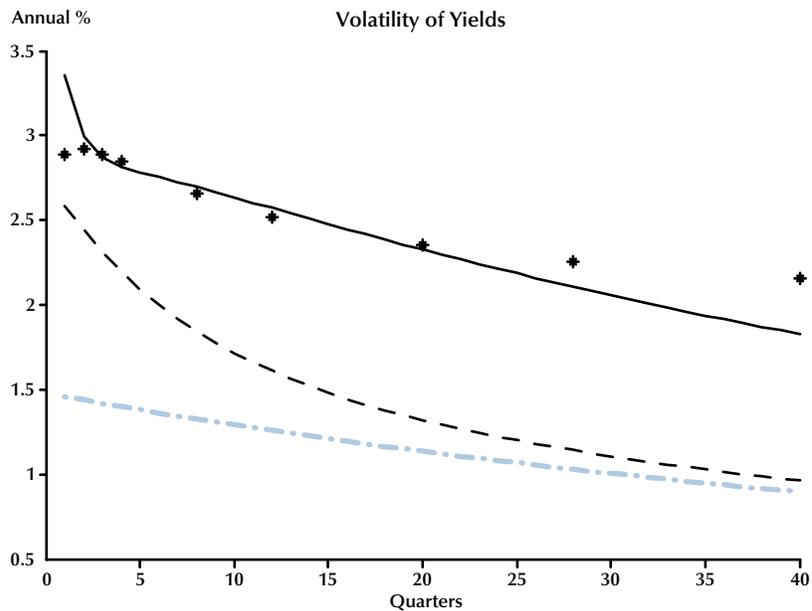
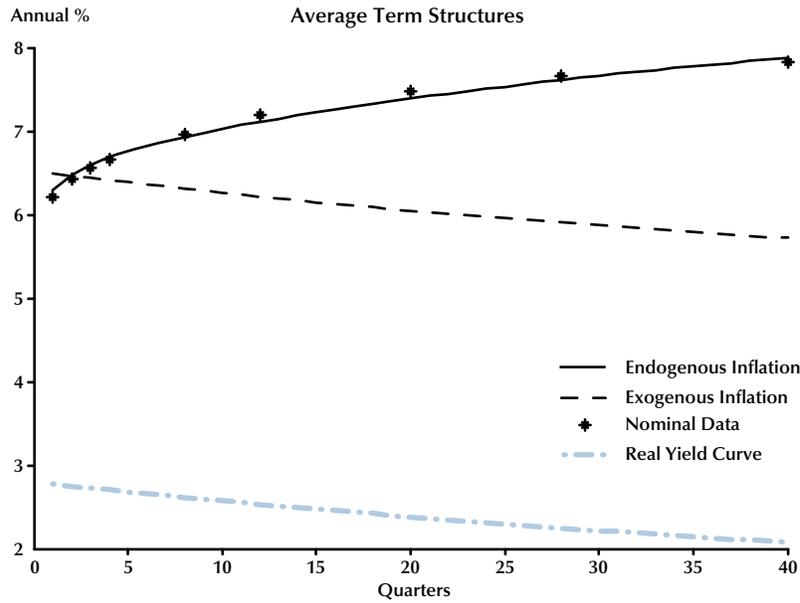
Average Yield Curve and Volatilities for the Epstein-Zin Model with Stochastic Volatility



NOTE: The parameters are $\rho = 0.5$, $\alpha = -4.911$, $\beta = 0.994$, $\bar{r} = -0.015$, $\tau_x = 3.064$, and $\tau_p = 2.006$. The empirical moments for the full sample (1970:Q1–2005:Q4) are plotted with asterisks, properties of the real yield curve are plotted with a dashed-dotted blue line, properties of the yield curve in the exogenous inflation economy are plotted with a dashed black line, and properties of the yield curve in the economy with endogenous inflation are plotted with a solid black line.

Figure 7

Average Yield Curve and Volatilities for the Epstein-Zin Model with Stochastic Volatility



NOTE: The parameters are $\rho = 1.0$, $\alpha = -6.079$, $\beta = 0.990$, $\bar{\tau} = -0.004$, $\tau_x = 1.534$, and $\tau_p = 1.607$. The empirical moments for the full sample (1970:Q1–2005:Q4) are plotted with asterisks, properties of the real yield curve are plotted with a dashed-dotted blue line, properties of the yield curve in the exogenous inflation economy are plotted with a dashed black line, and properties of the yield curve in the economy with endogenous inflation are plotted with a solid black line.

Table 1**Factor Loadings and Prices of Risk**

	Constant	Factor loadings (γ s)				Prices of risk (λ s)			
		x_t	v_t	p_t	s_t	ε_{t+1}^x	ε_{t+1}^v	ε_{t+1}^p	ε_{t+1}^s
A. $\rho = -0.5$, $\alpha = -4.835$, $\beta = 0.999$, $\bar{\tau} = 0.003$, $\tau_x = 1.2475$, $\tau_p = 1.000$									
Real kernel	0.01	0.54	25.45	—	—	8.25	-902.98	—	—
Exogenous inflation	0.01	0.54	25.45	0.85	—	8.25	-902.98	1.00	—
Endogenous inflation	0.02	0.14	21.63	—	-1.44	7.15	-906.90	—	-1.56
B. $\rho = 0.0$, $\alpha = -4.061$, $\beta = 0.998$, $\bar{\tau} = 0.003$, $\tau_x = 0.973$, $\tau_p = 0.973$									
Real kernel	0.01	0.36	20.07	—	—	7.34	-677.11	—	—
Exogenous inflation	0.01	0.36	20.07	0.85	—	7.34	-677.11	1.00	—
Endogenous inflation	0.02	0.00	33.56	—	-1.51	6.34	-663.24	—	-1.63
C. $\rho = 0.5$, $\alpha = -4.911$, $\beta = 0.994$, $\bar{\tau} = -0.015$, $\tau_x = 3.064$, $\tau_p = 2.006$									
Real kernel	0.01	0.18	32.23	—	—	8.93	-972.61	—	—
Exogenous inflation	0.01	0.18	32.23	0.85	—	8.93	-972.61	1.00	—
Endogenous inflation	0.02	-0.45	38.34	—	-0.56	7.18	-966.33	—	-0.61
D. $\rho = 1.0$, $\alpha = -6.079$, $\beta = 0.990$, $\bar{\tau} = 0.004$, $\tau_x = 1.534$, $\tau_p = 1.607$									
Real kernel	0.02	0.00	51.82	—	—	10.99	-1,398.00	—	—
Exogenous inflation	0.02	0.00	51.82	0.85	—	10.99	-1,398.00	1.00	—
Endogenous inflation	0.03	-0.44	58.30	—	-0.74	9.76	-1,391.30	—	-0.80

NOTE: The table reports the affine term-structure parameters for the real term structure, the nominal term structure in the exogenous inflation economy, and the nominal term structure in the endogenous inflation economy. The parameters in each panel are computed using a different set of preference parameters. We fix a level of the intertemporal elasticity parameter, ρ , and choose the remaining preference parameters—the risk aversion coefficient, α , and the rate-of-time preference, β , to minimize the distance between the average nominal yields and yield volatilities in the data and those implied by the exogenous inflation economy. We pick the Taylor rule parameters to minimize the distance between the average nominal yields and yield volatilities in the data and the those implied by the endogenous inflation economy.

the factor loadings and the prices of risk for each economy corresponding to the figures. Table 2 reports the coefficients on the equilibrium inflation rate and properties of the equilibrium inflation rate in the endogenous inflation economy.

Figure 4 reports the results with $\rho = -0.5$; here the representative agent has a low intertemporal elasticity of substitution. The remaining preference parameters are $\alpha = -4.835$ and $\beta = 0.999$. With this choice of parameters, the average real term structure is slightly downward sloping.

Backus and Zin (1994) show that a necessary condition for the average yield curve to be upward sloping is negative autocorrelation in the pricing

kernel.² Consider an affine model with independent factors $z_t^1, z_t^2, \dots, z_t^k$ with an innovation ε_t^j on the j th factor, a factor loading γ_j on the j th factor, and a price of risk λ_j on the j th factor. In such a model, the j th factor contributes

$$(12) \quad \gamma_j^2 \text{Autocov}(z_t^j) + \gamma_j \lambda_j \text{Cov}(z_t^j, \varepsilon_t^j)$$

to the autocovariance of the pricing kernel.

In our calibration, the exogenous factors in

² Piazzesi and Schneider (2007) argue that an upward-sloping nominal yield curve can be generated if inflation is bad news for consumption growth. Such a structure leads to negative autocorrelation in the nominal pricing kernel.

Table 2
Properties of p_t in the Endogenous Inflation Economy

	$\bar{\pi}$	π_x	π_v	π_s	$E(p_t)$	$\sigma(p_t)$	AR(1)
A. $\rho = -0.5, \alpha = -4.835, \beta = 0.999, \bar{\tau} = 0.003, \tau_x = 1.2475, \tau_p = 1.000$	0.01	-1.11	-3.92	-1.56	0.01	0.02	0.37
B. $\rho = 0.0, \alpha = -4.061, \beta = 0.998, \bar{\tau} = 0.003, \tau_x = 0.973, \tau_p = 0.973$	0.01	-1.00	13.87	-1.63	0.01	0.02	0.44
C. $\rho = 0.5, \alpha = -4.911, \beta = 0.994, \bar{\tau} = -0.015, \tau_x = 3.064, \tau_p = 2.006$	0.02	-1.75	6.28	-0.61	0.01	0.03	0.37
D. $\rho = 1.0, \alpha = -6.079, \beta = 0.990, \bar{\tau} = 0.004, \tau_x = 1.534, \tau_p = 1.607$	0.01	-1.23	6.66	-0.80	0.01	0.02	0.37

NOTE: The first four columns are coefficients of the inflation rate: the equilibrium inflation rate coefficients on a constant, output, stochastic volatility, and the monetary policy shock, respectively. The last two columns are properties of inflation: the unconditional standard deviation and the first-order autocorrelation of inflation, respectively.

the real economy—output growth and stochastic volatility—all have positive autocorrelations and the factor innovations have positive covariances to the factor levels. This implies that $\gamma_j^2 \text{Autocov}(z_t^j)$ and $\text{Cov}(z_t^j, \varepsilon_t^j)$ are both positive. For a factor to contribute negatively to the autocorrelation of the pricing kernel, the factor loading γ_j and the price of risk λ_j must have opposite signs. Additionally, the price of risk λ_j must be large enough relative to the factor loading γ_j to counteract the positive autocovariance term $\gamma_j^2 \text{Autocov}(z_t^j)$.

Output growth has a lower autocorrelation coefficient than stochastic volatility in our calibration, but because output growth has a much higher unconditional volatility, it has a much higher autocovariance than stochastic volatility. In the real economy, the factor loading γ_x on the level of output growth is equal to $(1 - \rho)\phi_x$, which is nonnegative for all $\rho \leq 1$. Also, the price of risk for output growth, λ_x , is positive at the parameter values used in Figure 4 because a sufficient condition for it to be positive is $\alpha \leq 0$ and $|\rho| \leq |\alpha|$. From (12), output growth contributes positively to the autocovariance of the pricing kernel.

From the real pricing kernel parameters given in (7), the price of risk for volatility is related to the factor loading on the level of volatility by

$$\lambda_v = -\frac{\beta}{1 - \beta\phi_v} \gamma_v.$$

Because $1 - \beta\phi_v > 0$, the volatility price of risk, λ_v , and the volatility factor loading, γ_v , have opposite signs, implying that the volatility factor can contribute a negative autocovariance to the pricing kernel. But output growth has the strongest effect on the autocovariance of the pricing kernel, leading to positive autocovariance in the pricing kernel. As a consequence, the average real yield curve is downward sloping. The numerical values for the real pricing kernel's factor loadings and prices of risk from Figure 4 are reported in panel A of Table 1.

In the exogenous inflation economy, shocks to inflation are uncorrelated to output growth and stochastic volatility—the factor loadings and prices of risk on output growth and stochastic volatility in the nominal pricing kernel are the same as in the real pricing kernel. Average nominal yields in the exogenous inflation economy are equal to the real yields plus expected inflation and inflation volatility with an adjustment for properties of the inflation process. The inflation shocks are positively autocorrelated, with a factor

loading and a price of risk that are both positive. The average nominal yield curve has approximately the same shape as the real yield curve—it is downward sloping.

In the endogenous inflation economy, inflation is a linear combination of output growth, stochastic volatility, and the monetary policy shock. As shown in panel A of Table 2, endogenous inflation's loading on output, π_x , is negative. This implies that the nominal pricing kernel's output-growth factor loading and price of risk are lower than in the exogenous inflation economy. As a consequence, output growth contributes much less to the autocovariance of the pricing kernel with endogenous inflation. The factor loading and price of risk for stochastic volatility are also lower in the endogenous inflation economy. The policy shocks are positively autocorrelated, but the factor loading and the price of risk for the policy shock are of opposite sign. The average nominal yield curve in the endogenous inflation economy is therefore flatter than both the real yield curve and the nominal yield curve with exogenous inflation.

Turning to the volatilities in the bottom panel of Figure 4, the exogenous inflation economy exhibits more volatility in short rates and less volatility in long rates than found in the data. This is a fairly standard finding for term-structure models with stationary dynamics (see Backus and Zin, 1994). The volatility of long rates is mainly driven by the loading on the factor with the largest innovation variance and that factor's autocorrelation. The closer that autocorrelation is to zero, the faster that yield volatility decreases as bond maturity increases. In our calibration, output growth has the largest innovation variance and a fast rate of mean reversion, equal to 0.36. Yield volatility drops quite quickly as bond maturity increases. In general, the lower the loading on output growth, the slower that yield volatility drops as bond maturity increases. Because endogenous inflation is negatively related to output growth, the factor loading on output growth is lower. Yield volatility drops at a slower rate (relative to maturity) in the endogenous inflation economy than in the exogenous inflation economy.

Figure 5, panel B of Table 1, and panel B of Table 2 report yield-curve properties with a higher intertemporal elasticity of substitution ($\rho = 0$) or a log time aggregator. Piazzesi and Schneider (2007) study a model with the same preferences, but without stochastic volatility. The factor loading on output growth in the real economy is higher than in the economy with $\rho = -0.5$ reported in Figure 4 (compare panel A with panel B of Table 1). The average real yield curve and the average nominal yield curve with exogenous inflation are less downward sloping when $\rho = 0$ than when $\rho = -0.5$. Similarly, increasing ρ further to 0.5 (see Figure 6) or 1.0 (see Figure 7) leads to a real yield curve that is less downward-sloping. Because increasing ρ decreases the factor loading on output growth, it also decreases the volatility of real yields: See the bottom panels in Figures 4 through 7.

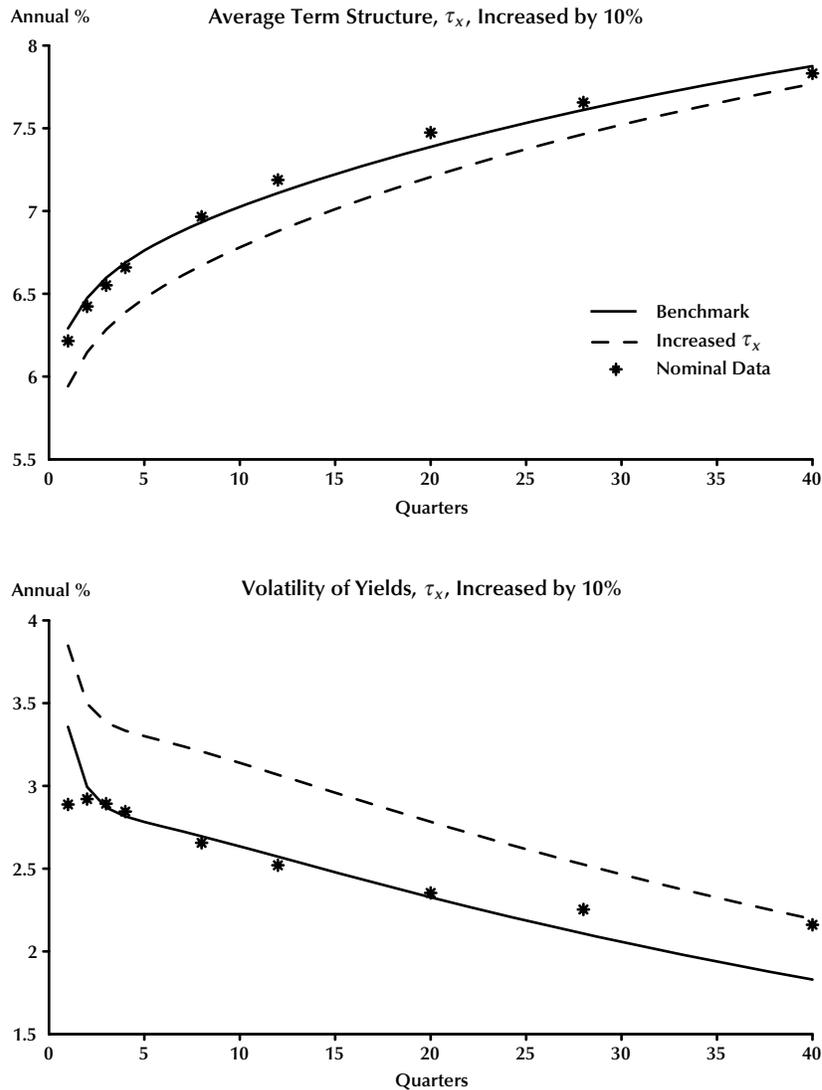
As ρ increases, the representative agent's intertemporal elasticity of substitution increases, implying less demand for smoothing consumption over time. Increasing ρ decreases the representative agent's demand for long-term bonds for the purpose of intertemporal consumption smoothing and leads to lower equilibrium prices and higher yields for real long-term bonds. The average real yield curve therefore is less downward-sloping as ρ increases. Increasing ρ also reduces the sensitivity of long-term real yields to output growth, leading to less volatile long-term yields: See the bottom panels in Figures 4 through 7.

Nominal yields in the economies with exogenous inflation are approximately equal to the real yields plus a maturity-independent constant. But in the economies with endogenous inflation, inflation and output growth have a negative covariance, leading to a decrease in the factor loading on output growth: See panels C and D of Tables 1 and 2. For $\rho \geq 0.5$ (see Figures 6 and 7), the average nominal yield curve is upward sloping and the shape of the volatility term structure decays similarly to that observed in the data.

The final three columns of Table 2 report unconditional moments of inflation in the economy with endogenous inflation. There are a few notable features. First, the unconditional moments

Figure 8

The Effects of Increasing τ_x



NOTE: The baseline parameters are $\rho = 1.0$, $\alpha = -6.079$, $\beta = 0.990$, $\bar{\tau} = -0.004$, $\tau_x = 1.534$, and $\tau_p = 1.607$. Empirical moments for the full sample (1970:Q1–2005:Q4) are plotted with asterisks, results from the baseline parameters are plotted with a solid black line, and results when the feedback from output growth to short-term interest rates is increased by 10 percent are plotted with a dashed black line.

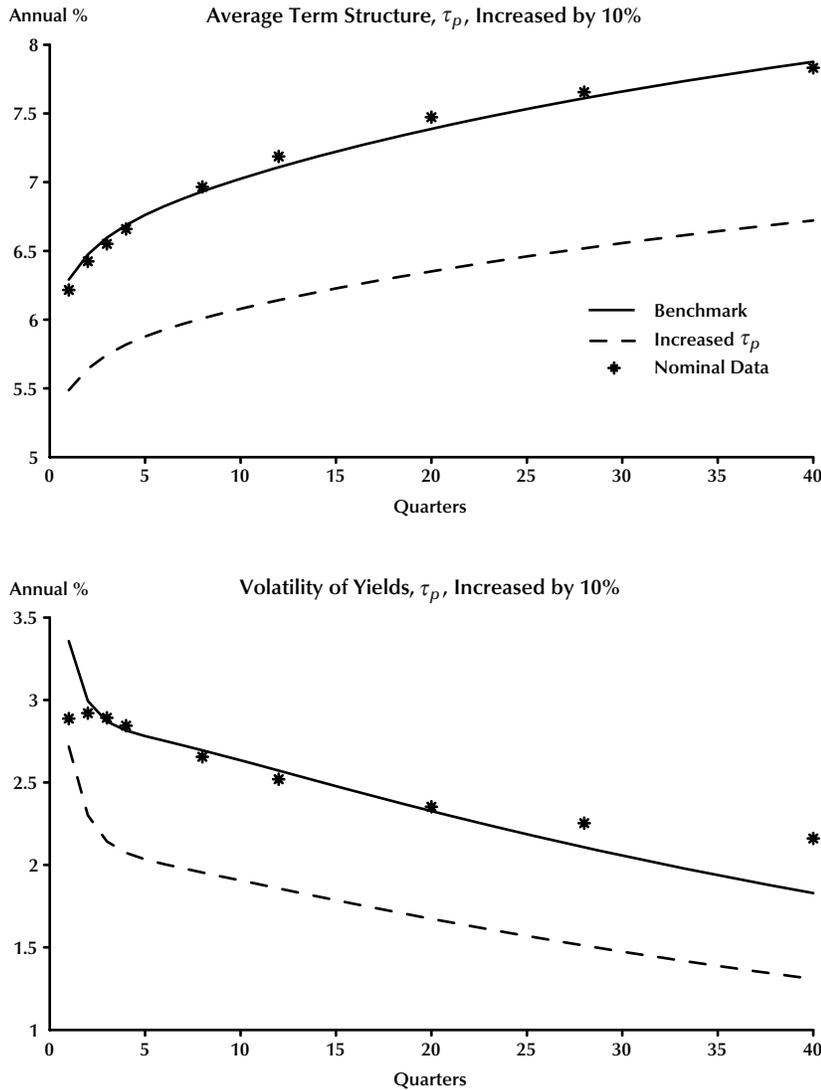
are not particularly sensitive to the intertemporal elasticity of substitution. Second, the unconditional variance of inflation in the calibrated economy is an order of magnitude higher than that in the data: 0.0033 in empirical data and about 0.02 in these economies. Finally, inflation is much

more autocorrelated in the data—the AR(1) coefficient is 0.85 in the data and about 0.4 in the model economies.

Figure 8, Figure 9, and Table 3 show results from changing the Taylor rule parameters. We keep the remaining parameters fixed at the values

Figure 9

The Effects of Increasing τ_p



NOTE: The baseline parameters are $\rho = 1.0$, $\alpha = -6.079$, $\beta = 0.990$, $\bar{r} = -0.004$, $\tau_x = 1.534$, and $\tau_p = 1.607$. Empirical moments for the full sample (1970:Q1–2005:Q4) are plotted with asterisks, results from the baseline parameters are plotted with a solid black line, and results when the feedback from inflation to short-term interest rates is increased by 10 percent are plotted with a dashed black line.

used to generate Figure 7. Figure 8 shows that increasing τ_x , the interest rate’s responsiveness to output growth shocks, leads to a reduction in average nominal yields and a steepening in the average yield curve (top panel), as well as an increase in yield volatility (bottom panel).

Panel A of Table 3 shows that increasing τ_x decreases the constant in the nominal pricing kernel, decreases the factor loading on output growth, decreases the price of risk for output growth, and also increases the factor loading on stochastic volatility. The loading on output growth

Table 3
Comparative Statics for the Taylor Rule Parameters

	Nominal pricing kernel										
	Constant	Factor loadings (γ s)			Prices of risk (λ s)			Equilibrium inflation loadings			
		x_t	v_t	s_t	ε_{t+1}^x	ε_{t+1}^v	ε_{t+1}^s	$\bar{\pi}$	π_x	π_v	π_s
A. τ_x increased by 10%, from 1.53 to 1.69											
Baseline	0.03	-0.44	58.30	-0.74	9.76	-1,391.30	-0.80	0.01	-1.23	6.66	-0.80
Increased τ_x	0.02	-0.49	60.13	-0.74	9.63	-1,389.40	-0.80	0.01	-1.35	8.55	-0.80
B. τ_p increased by 10%, from 1.61 to 1.77											
Baseline	0.03	-0.44	58.30	-0.74	9.76	-1,391.30	-0.80	0.01	-1.23	6.66	-0.80
Increased τ_p	0.02	-0.39	55.30	-0.66	9.90	-1,394.40	-0.71	0.01	-1.09	3.58	-0.71

NOTE: The table reports the effect of changing the Taylor rule parameter τ_x or τ_p on the affine term-structure parameters as well as properties of p_t in the endogenous inflation economy. The equilibrium inflation rate coefficients on output, stochastic volatility, and the monetary policy shock are reported. The baseline parameters are $\rho = 1.0$, $\alpha = -6.08$, $\beta = 0.990$, $\bar{r} = -0.004$, $\tau_x = 1.53$, and $\tau_p = 1.61$.

in the pricing kernel drops because the sensitivity of the inflation rate to output growth drops; in turn, the sensitivity of inflation to stochastic volatility increases by a large amount—from 6.66 to 8.55.

Figure 9 shows that increasing τ_p , the interest rate responsiveness to inflation, leads to a reduction in average nominal yields and a flattening in the average yield curve (top panel) and a decrease in yield volatility (bottom panel).

Panel B of Table 3 shows that increasing τ_p decreases the constant in the nominal pricing kernel, increases the factor loading on output growth, increases the price of risk for output growth, decreases the factor loading on stochastic volatility, and also drops the factor loading on the monetary policy shock. The constant in the pricing kernel drops because the constant in the inflation rate drops, the factor loading on output growth increases because the sensitivity of the inflation rate to output growth increases; in turn, the sensitivity of inflation to stochastic volatility decreases by a large amount—from 6.66 to 3.58.

Overall, the experiments reported in Figure 8 and Figure 9 show that properties of the term structure depend on the form of the monetary authorities' interest rate feedback rule. In particular, the factor loading on stochastic volatility is

quite sensitive to the interest rate rule. In this economy, because stochastic volatility is driving time-variation in interest rate risk premiums, monetary policy can have large impacts on interest rate risk premiums.

RELATED RESEARCH

The model we develop is similar to a version of Bansal and Yaron's (2004), which includes stochastic volatility; however, our simple autoregressive state-variable process does not capture their richer ARMA specification. Our work is also related to Piazzesi and Schneider (2007), who emphasize that, for a structural model to generate an upward-sloping nominal yield curve, it requires joint assumptions on preferences and the distribution of fundamentals. Our work highlights how an upward-sloping yield curve can also be generated through the monetary authority's interest rate feedback rule.

Our paper adds to a large and growing literature combining structural macroeconomic models that include Taylor rules with arbitrage-free term-structure models. Ang and Piazzesi (2003), following work by Piazzesi (2005), have shown that a factor model of the term structure that imposes arbitrage-free conditions can provide a better

empirical model of the term structure than a model based on unobserved factors or latent variables alone. Estrella and Mishkin (1997), Evans and Marshall (1998 and 2001), Hördahl, Tristani, and Vestin (2004), Bekaert, Cho, and Moreno (2005), and Ravenna and Seppala (2006) also provide evidence of the benefits of building arbitrage-free term-structure models with macroeconomic fundamentals. Rudebusch and Wu (2004) and Ang, Dong, and Piazzesi (2004) investigate the empirical consequences of imposing a Taylor rule on the performance of arbitrage-free term-structure models.

For an alternative linkage between short- and long-maturity bond yields, see Vayanos and Vila (2006), who show how the shape of the term structure is determined in the presence of risk-averse arbitrageurs, investor clienteles for specific bond maturities, and an exogenous short rate that could be driven by the central bank's monetary policy.

CONCLUSIONS

We demonstrate that an endogenous monetary policy that involves an interest rate feedback rule can contribute to the riskiness of multi-period bonds by creating an endogenous inflation process that exhibits significant covariance risk with the pricing kernel. We explore this through a recursive utility model with stochastic volatility that generates sizable average risk premiums. Our results point to a number of additional questions. First, the Taylor rule that we work with is arbitrary, so how would the predictions of the model change with alternative specification of the rule? In particular, how would adding monetary non-neutralities along the lines of a New Keynesian Phillips curve as in Clarida, Galí, and Gertler (2000) and Gallmeyer, Hollifield, and Zin (2005) alter the monetary policy-consistent pricing kernel? Second, what Taylor rule would implement an optimal monetary policy in this context? Because preferences have changed relative to the models in the literature, this is a nontrivial theoretical question.

In addition, the simple calibration exercise in this paper is not a very good substitute for a more

serious econometric exercise. Further research will explore the trade-offs between shock specifications, preference parameters, and monetary policy rules for empirical yield-curve models that more closely match historical evidence.

Finally, it would be instructive to compare and contrast the recursive utility model with stochastic volatility with other preference specifications that are capable of generating realistic risk premiums. The leading candidate on this dimension is the external habits models of Campbell and Cochrane (1999). We are currently pursuing an extension of the external habits model in Gallmeyer, Hollifield, and Zin (2005) to include an endogenous, Taylor rule-driven inflation process.

REFERENCES

- Ang, Andrew; Dong, Sen and Piazzesi, Monika. "No-Arbitrage Taylor Rules." Working paper, Columbia University, 2004.
- Ang, Andrew and Piazzesi, Monika. "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables." *Journal of Monetary Economics*, May 2003, 50(4), pp. 745-87.
- Backus, David K.; Foresi, Silverio and Telmer, Chris I. "Affine Term Structure Models and the Forward Premium Anomaly." *Journal of Finance*, February 2001, 56(1), pp. 279-304.
- Backus, David K. and Zin, Stanley E. "Reverse Engineering the Yield Curve." NBER Working Paper No. 4676, National Bureau of Economic Research, 1994.
- Backus, David K. and Zin, Stanley E. "Bond Pricing with Recursive Utility and Stochastic Volatility." Working paper, New York University, 2006.
- Bansal, Ravi and Yaron, Amir. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *Journal of Finance*, August 2004, 59(4), pp. 1481-510.
- Bekaert, Geert; Cho, Seonghoon and Moreno, Antonio. "New-Keynesian Macroeconomics and the Term

- Structure.” Working paper, Columbia University, 2005.
- Campbell, John Y. and Cochrane, John H. “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior.” *Journal of Political Economy*, April 1999, 107(2), pp. 205-51.
- Clarida, Richard; Galí, Jordi and Gertler, Mark. “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory.” *Quarterly Journal of Economics*, February 2000, 115(1), pp. 147-80.
- Cochrane, John H. “Identification and Price Determination with Taylor Rules: A Critical Review.” Working paper, University of Chicago, 2006.
- Duffie, Darrell and Kan, Rui. “A Yield-Factor Model of Interest Rates.” *Mathematical Finance*, October 1996, 6(4), pp. 379-406.
- Epstein, Larry G. and Zin, Stanley E. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework.” *Econometrica*, July 1989, 57(4), pp. 937-69.
- Estrella, Arturo and Mishkin, Frederic S. “The Predictive Power of the Term Structure of Interest Rates in Europe and the United States: Implications for the European Central Bank.” *European Economic Review*, July 1997, 41(7), pp. 1375-401.
- Evans, Charles L. and Marshall, David A. “Monetary Policy and the Term Structure of Nominal Interest Rates: Evidence and Theory.” *Carnegie-Rochester Conference Series on Public Policy*, December 1998, 49, pp. 53-111.
- Evans, Charles L. and Marshall, David A. “Economic Determinants of the Term Structure of Nominal Interest Rates.” Working Paper WP-01-16, Federal Reserve Bank of Chicago, 2001.
- Gallmeyer, Michael F.; Hollifield, Burton and Zin, Stanley E. “Taylor Rules, McCallum Rules, and the Term Structure of Interest Rates.” *Journal of Monetary Economics*, July 2005, 52(5), pp. 921-50.
- Greenspan, Alan. Testimony before the U.S. Senate Committee on Banking, Housing, and Urban Affairs. Board of Governors of the Federal Reserve System, May 27, 1994; http://fraser.stlouisfed.org/historicaldocs/ag94/download/27982/Greenspan_19940527.pdf.
- Greenspan, Alan. Testimony before the U.S. Senate Committee on Banking, Housing, and Urban Affairs. Board of Governors of the Federal Reserve System, February 16, 2005; www.federalreserve.gov/boarddocs/hh/2005/february/testimony.htm.
- Hansen, Lars Peter; Heaton, John and Li, Nan. “Consumption Strikes Back? Measuring Long-Run Risk.” NBER Working Paper No. 11476, National Bureau of Economic Research, 2005.
- Hördahl, Peter; Tristani, Oreste and Vestin, David. “A Joint Econometric Model of Macroeconomic and Term Structure Dynamics.” Working Paper Series No. 405, European Central Bank, 2004.
- McCulloch, J. Huston and Kwon, Heon-Chul. “U.S. Term Structure Data, 1947-1991.” Working Paper No. 93-6, The Ohio State University, 1993.
- Piazzesi, Monika. “Bond Yields and the Federal Reserve.” *Journal of Political Economy*, April 2005, 113(2), pp. 311-44.
- Piazzesi, Monika and Schneider, Martin. “Equilibrium Yield Curves,” in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., *NBER Macroeconomics Annual 2006*. Cambridge, MA: MIT Press, 2007.
- Ravenna, Federico and Seppala, Juha. “Monetary Policy and Rejections of the Expectations Hypothesis.” Working paper, University of Illinois, 2006.
- Rudebusch, Glenn D. and Wu, Tao. “A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy.” Presented at the Federal Reserve Bank of San Francisco conference *Interest Rates and Monetary Policy*, March 19-20, 2004.

Gallmeyer, Hollifield, Palomino, Zin

Vayanos, Dimitri and Vila, Jean-Luc. "A Preferred-Habitat Model of the Term Structure of Interest Rates." Working paper, London School of Economics, 2006.

Pamela A. Labadie

The questions addressed in the Gallmeyer et al. (2007) paper are important ones: How does monetary policy affect long-term interest rates? How can we explain the volatility of the long end of the yield curve and its relationship with monetary policy? A successful model is one that answers these questions and is useful to policymakers. To be useful to policymakers, the model should yield quantitative answers to questions such as, What will happen to long-term rates if the central bank raises the federal funds rate by 25 basis points? Since money is not introduced in a way that is essential,¹ this model is not designed to answer questions about the mechanism by which monetary policy affects the yield curve. This is not necessarily a deficiency, because models where money is essential have yet to prove useful in policy discussions and empirical results.² The strength of this model is that it provides quantitative answers to the questions. Hence, it is reasonable to set a standard of success where the model-generated time series must, in some sense, “look like” actual time series.

As in any paper where hard modeling decisions are made, there are both strengths and weaknesses in the choices made by the authors. To describe these choices concisely, I’ll provide a brief overview of the model.

¹ Wallace (2001) says that money is “essential” in an economy if it permits allocations that would otherwise not be achieved.

² See Kocherlakota (2002).

OVERVIEW OF THE MODEL

There is a single, exogenous, stochastic, and perishable endowment good. The endowment grows at the rate x_t , with stochastic volatility v_t :

$$\begin{aligned} x_{t+1} &= (1 - \phi_x)\theta_x + \phi_x x_t + v_t^2 \varepsilon_{t+1}^x \\ v_{t+1} &= (1 - \phi_v)\theta_v + \phi_v v_t + \sigma_v \varepsilon_{t+1}^v \end{aligned}$$

Hence, the endowment process is characterized by a parameter vector $(\phi_x, \phi_v, \theta_x, \theta_v, \sigma_v)$.

There is a representative agent with Epstein-Zin preferences. The advantages of this preference structure, with its property of separating relative risk aversion from the elasticity of intertemporal substitution, are well known. The preference parameter vector is (β, ρ, α) . The pricing kernel is the intertemporal marginal rate of substitution in consumption, denoted $\log(m_{t+1})$. The price of default-free discount bonds can be determined recursively through an arbitrage-free restriction of the form

$$b_t^{(n)} = E_t m_{t+1} b_{t+1}^{(n-1)}.$$

With this structure for the endowment process and the preferences, the model can be mapped into the Duffie-Kan affine term-structure model with two factors. In particular, the authors guess a form of the value function and then verify that this guess is a solution. The value function then implies the form of the real pricing kernel. The process by which the authors relate the deeper preference and endowment parameters to the

Pamela A. Labadie is a professor of economics at The George Washington University.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 327-29.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

discrete-time affine term-structure model is interesting and elegant.

To introduce money and inflation into the model, the authors add a stochastic process to the real pricing kernel:

$$\log(m_{t+1}^s) = \log(m_{t+1}) - p_{t+1},$$

where p_{t+1} is inflation. Hence, money, prices, and inflation in the model are merely noise, creating a wedge between the real and nominal pricing kernels. Two specifications of the inflation process are studied: exogenous and endogenous.

Exogenous Inflation

Inflation is conjectured to take the form

$$p_{t+1} = (1 - \phi_p)\theta_p + \phi_p p_t + \sigma_p \varepsilon_{t+1}^p.$$

This specification does not link inflation to the endowment process or to any money-growth process. The state space is expanded to x_t, v_t, p_t . They calibrate the model to data and set the parameters values at

$$\phi_p = 0.8471$$

$$\theta_p = 0.0093$$

$$\sigma_p = 0.0063(1 - \phi_p^2)^{\frac{1}{2}}.$$

The preference parameters are set at $\rho = 0$ and $\alpha = -2.91$, where ρ is the elasticity of intertemporal substitution and α is relative risk aversion,³ and there is little discussion of this choice. With this set of parameter values, the model is used to derive a yield curve. The result is a model-generated yield curve that matches the shape of the historical yield curve, but exhibits less volatility in long rates. Since an explicit goal of the paper is to explain volatility at the long end of the yield curve, this answer is not satisfactory. Three comments arise: The first is why are these values for the preference parameters chosen? How does varying the preference parameters change the results? Should we just aimlessly search the parameter space for a better fit? This

leads to my second comment. The model, as specified, can be estimated using formal econometric techniques. The endowment process and inflation processes, along with the real and nominal pricing kernels, form a system of equations. Using data on nominal interest rates, one can estimate inflation and consumption growth, the parameters of the model, and in particular the preference parameters. The preference-parameter estimates could then be used to generate a real pricing kernel. What are the estimated preference parameters and does economic intuition suggest that they are sensible? How does the real pricing kernel behave? It seems a missed opportunity.

Finally, the model, as posed, severely restricts the price of inflation risk by fixing the price of inflation risk at unity. A very simple cash-in-advance model with fixed velocity has the property that inflation is a function of both money growth and output growth, so there is a state-varying inflation risk premium. Even with a zero mean, the inflation-risk premium may be an additional source of variability and may help to remedy the lack of volatility at the long end of the yield curve in the model.

Endogenous Inflation

To make inflation endogenous, the authors assume that monetary policy follows a nominal interest rate rule (Taylor rule) of the form

$$i_t = \bar{\tau} + \tau_x x_t + \tau_p p_t + s_t.$$

This rule raises short-term rates aggressively in response to inflation. There are many other specifications that could be considered, and it would be helpful to have a discussion on why this rule is chosen over other specifications. Is this the type of rule the authors believe monetary authorities are using? Is this the rule that gives the best results in the sense that the model-generated yield curve matches the data? Is it chosen for tractability?

This process must be consistent with the other equations in the model, which requires the derivation of an inflation process consistent with the interest rate rule and other equations. To link the rule to the nominal pricing kernel and bond-

³ These are the values chosen in the version of the paper presented at the conference.

market equilibrium, they use a guess-and-verify method to derive a consistent inflation process of the form

$$p_t = \bar{\pi} + \pi_x x_t + \pi_v v_t + \pi_s s_t.$$

The state space is now $z_t = (x_t, v_t, s_t)$, and there are additional restrictions on the means and conditional variances. Once again they calibrate the model, choosing parameter values taken from the data. The calibrated model fits the average yield curve and has a volatility pattern closer to the data, especially at the long end.

This is a major goal of the paper and, in that sense, it is successful; but the question arises as to whether the endogenous inflation process in any way resembles the inflation process in the data. My conjecture is that it does not—and in some important ways—and the differences need to be made explicit. How do the endogenous and exogenous inflation processes compare? If the calibrated model fits the average yield curve and closely matches the volatility, but is based on an inflation process that differs significantly from the actual inflation process, how useful is this to policymakers? How useful is a model that fits the yield curve and its volatility but is greatly at odds with the actual inflation process? Common sense would suggest it is of limited usefulness.

CONCLUSION

The authors are to be commended for devising a model with such a rich potential for explaining yield curves and their volatility. Linking the Epstein-Zin preferences to a discrete-time affine

term-structure model is no easy task, although they seem to do so effortlessly. The model is devised to be estimated with standard econometric methods, using data on bond prices, consumption growth, and inflation. Such an exercise would provide useful insights into the real pricing kernel and the parameters of the Epstein-Zin preferences. The exogenous inflation specification is too restrictive and should permit a variable inflation-risk premium. Finally, alternative interest rate rules should be examined with the explicit goal of generating an inflation process using the model that matches the actual inflation process, according to some explicit criterion. While much of this may sound negative, I want to emphasize that this model has the potential to be very useful to policymakers and the steps needed to make it so are straightforward ones to take.

REFERENCES

- Gallmeyer, Michael F.; and Hollifield, Burton; Palomino, Francisco and Zin, Stanley. "Arbitrage-Free Bond Pricing with Dynamic Macroeconomic Models." *Federal Reserve Bank of St. Louis Review*, July/August 2007, 89(4), pp. 305-26.
- Kocherlakota, Narayana. "Money: What's the Question and Why Should We Care about the Answer?" *American Economic Review*, May 2002, 92(2), pp. 58-61.
- Wallace, Neil. "Lawrence R. Klein Lecture 2000: Whither Monetary Economics?" *International Economic Review*, November 2001, 42(4), pp. 847-69.



Monetary Policy as Equilibrium Selection

Gaetano Antinolfi, Costas Azariadis, and James B. Bullard

Can monetary policy guide expectations toward desirable outcomes when equilibrium and welfare are sensitive to alternative, commonly held rational beliefs? This paper studies this question in an exchange economy with endogenous debt limits in which dynamic complementarities between dated debt limits support two Pareto-ranked steady states: a suboptimal, locally stable autarkic state and a constrained optimal, locally unstable trading state. The authors identify feedback policies that reverse the stability properties of the two steady states and ensure rapid convergence to the constrained optimal state. (JEL E31, E42, E58)

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 331-41.

INTRODUCTION

Overview

Indeterminacy, or non-uniqueness, of rational expectations equilibrium has been a prominent feature of monetary policy analysis since Sargent and Wallace (1975) found that passive interest rate policies cause indeterminacy in an IS-LM framework with rational expectations. Generally speaking, policy choices influence equilibrium outcomes, and passive choices can support multiple equilibria. This situation has been viewed as one to be avoided if at all possible; the prospect of the economy coordinating on the “wrong” set of self-confirming beliefs is unnecessary at best and detrimental to welfare at worst. In the standard New Keynesian model, for example, the monetary policymaker must follow a sufficiently active policy to avoid indeterminacy. “Active” means that the policy instrument cannot be held fixed, or allowed to fluctuate randomly, but instead must adjust to the state of the economy according to a specific, widely understood rule. A policy that is too

passive—say, too close to a nominal interest rate peg—allows indeterminacy.

Results with this flavor depend critically on the expectations of the private sector regarding future monetary policy actions, and this has led many to describe the problem of monetary policy as one of managing or shaping expectations to rule out private sector beliefs that may send the economy toward a suboptimal course. How can policy be designed to stop this process? Can policy somehow strengthen rational beliefs in the desired inflation target and in moderate inflationary expectations?

This paper considers indeterminacy and monetary policy from a dynamic general equilibrium perspective in order to study the robustness of activist monetary policy advice, like that coming from the large literature on Taylor-type rules.¹ We find these results to be quite robust. In fact, multiple Pareto-ranked dynamic equilibria turn out to occur whenever the monetary instrument

¹ See the discussion in Woodford (2003) and Bullard and Mitra (2002). For a discussion of the Taylor principle, see Woodford (2001).

Gaetano Antinolfi is an associate professor of economics and Costas Azariadis is a professor of economics at Washington University in St. Louis. Azariadis is a visiting scholar and James B. Bullard is a vice president and deputy director of monetary analysis at the Federal Reserve Bank of St. Louis. The authors thank Peter Ireland for extensive comments on an earlier draft.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

is used passively, without regard to the state of the economy. In contrast, some types of informed active policies eliminate most of the indeterminacy and force the economy onto a constrained optimal path. Our framework also begins to address interesting questions concerning the nature of the interaction between monetary policy and the smooth operation of credit markets.

What We Do

We study a dynamic general equilibrium model of pure exchange that is a simplified version of Azariadis and Kaas (2007). Following Eaton and Gersovitz (1981) and Kehoe and Levine (1993), endogenous debt limits deter default by households that cannot be forced to repay debts. These households live forever and have variable incomes. To keep the analysis tractable, we focus on an economy with just two types of agents who share a constant flow of total income. Income shares fluctuate between low and high levels in alternating periods. To smooth consumption perfectly, high-income agents could in principle lend a large enough amount to low-income agents each period to ensure that every household's share of total consumption remains constant. We show that, under certain reasonable assumptions, this first-best outcome *cannot* be achieved as an equilibrium with endogenous debt limits. Instead, there are two steady states: a constrained efficient outcome at a high interest rate in which credit markets work as well as possible and an inefficient autarkic outcome at a low interest rate in which credit markets break down and agents are unable to smooth consumption at all.² A continuum of dynamical equilibria indexed by initial conditions all tend toward the suboptimal steady state.

We introduce policy into this environment. We discuss the possibility of fiscal tax-transfer schemes that would in principle work well, but which also require the policymaker to use detailed information concerning household incomes to make the correct resource realloca-

tions. Passive monetary policy, which we think of as a constant rate of growth of the money stock, is always associated with indeterminacy and particularly poor dynamics. We then turn to active monetary policy, in which the policy instrument is adjusted in a particular way in reaction to the current state of the economy. We show that credible commitment to a certain active policy can converge to the constrained efficient outcome immediately if the policymaker reacts to the entire state of the economy and gradually if the policy rule responds only to prices. We regard this as a version of the policy advice coming from related literature on monetary policy in the face of important frictions in the economy, even though the friction in this paper is quite different. We also think this result suggests that good monetary policy is partly responsible for the smooth functioning of credit markets, a sentiment that is often stated in monetary policy circles.

Recent Related Literature

It is a typical result from the literature that models with a role for fiat money tend to have a nonmonetary steady state and an associated indeterminacy. This is true in models of overlapping generations; but the demand for money depends on beliefs in the search-theoretic monetary literature as well.³ The model here is more closely related to Bewley-type economies.⁴

In the New Keynesian literature, such as Woodford (2003), credit markets are complete and work perfectly, even though there are other frictions in that model. We also have complete markets, but the friction in our setting directly affects the incentives of households to lend appropriately. Thus, monetary policy in our framework improves the operation of credit markets, whereas in Woodford (2003) it has no particular effect on the operation of these markets.

² In the model, credit markets break down completely, but we think of this as representing poorly functioning credit markets in which the volume of borrowing and lending is less than it could be.

³ See, for example, the discussion in Wright (2005). Other examples of indeterminacy include older Keynesian models with rational expectations, and dynamic general equilibrium structures with bubbles, complementarities, and increasing returns, such as those reviewed in Boldrin and Woodford (1990), Cooper (1999), and Benhabib and Farmer (1999).

⁴ See Bewley (1980) and Townsend (1980).

Our results have a certain global flavor. In part that is because the model is simple enough that we can characterize the entire set of equilibria in a fairly straightforward way. Other authors have focused on a global perspective in models of monetary policy, perhaps most prominently Benhabib, Schmitt-Grohé, and Uribe (2001) and Benhabib and Eusepi (2005). They emphasize that active policies may be associated with local determinacy but global indeterminacy. In Benhabib, Schmitt-Grohé, and Uribe (2001), the second steady state (the one not associated with the inflation target of the monetary authorities) is close to the Friedman rule, whereas the second steady state in our framework is associated with high inflation. Benhabib and his collaborators emphasize how the design of policy may or may not be able to avoid too low an inflation rate relative to the target, whereas we stress how the design of monetary policy can avoid inflation rates that exceed any reasonable target. In particular, Benhabib and Eusepi (2005) show that, in a model with sticky prices, a feedback rule can eliminate global indeterminacy if the monetary instrument responds to the output gap.

A NONMONETARY MODEL

The economy we have in mind, but do not analyze here, consists of a large number of agents, possibly a continuum, with a common utility function and a large variety of income processes. Aggregate income can be thought of as constant so that we may focus on fluctuations in the *distribution* of income among households and on the asset trades they will conduct as they attempt to smooth consumption. Individual consumption shares will be constant if asset markets are perfect, but will necessarily fluctuate if endogenous debt limits constrain household borrowing.

To simplify matters and maintain tractability, we analyze an economy with deterministic individual incomes populated by two agents indexed by $i = 0, 1$. Time is discrete and denoted by $t = 0, 1, 2, \dots$. Each agent i has preferences given by

$$(1) \quad \sum_{t=0}^{\infty} \beta^t u(c_t^i),$$

with $0 < \beta < 1$. The aggregate endowment is constant at two units, but its distribution over agents changes deterministically over time. In particular, individual endowments are periodic⁵; that is,

$$(2) \quad (\omega_t^0, \omega_t^1) = \begin{cases} (1 + \alpha, 1 - \alpha) & \text{if } t = 0, 2, \dots \\ (1 - \alpha, 1 + \alpha) & \text{if } t = 1, 3, \dots \end{cases}$$

with $\alpha \in (0, 1)$. In addition, agent zero owes an initial debt, $B = \alpha / (1 + \beta)$, to agent one. This debt makes the initial wealth of the two agents identical when incomes are discounted at the common rate of time preference. In a more complicated economy, agents would be indexed by $\alpha \in (0, 1)$; some individual incomes would fluctuate only a little, others would fluctuate quite a bit.

Perfect Enforcement

To fix ideas and notation, we start with a standard dynamic general equilibrium model with perfect enforcement of loan contracts. In this setting, an equilibrium is an infinite sequence (c_t^H, c_t^L, R_t) that describes for each period t consumption for the high- and low-income agents and the gross yield on loans. This sequence satisfies consumption Euler equations for each person, two intertemporal budget constraints, and market clearing. Based on our assumptions concerning the initial distribution of wealth, it is obvious that the unique equilibrium is $(c_t^H, c_t^L, R_t) = (1, 1, 1/\beta)$ for all t , and it is Pareto optimal. Individual consumption is a *constant fraction* of aggregate consumption at all times.

Commitment to repay debts is essential in achieving this allocation of resources. If borrowers can in principle default on their loan obligations at the cost of perpetual exclusion from both sides of the asset market, as suggested by Kehoe and Levine (1993), then the Pareto-optimal allocation cannot be decentralized as a competitive equilibrium with limited enforcement unless it is weakly

⁵ In a growing economy, individual incomes need not be negatively correlated but income *shares* must be. This simple deterministic endowment process is the degenerate case of a stochastic economy with two Markovian states with a zero probability of remaining in the same state. Markovian endowments with two states are a straightforward extension. The assumption of two states or dates has obvious geometric advantages, but it is not innocuous where policy is concerned. We discuss this point further in the conclusion.

preferred to autarky by all agents at all times. It is easy to check that the current autarky payoff is

$$(3) \quad \frac{u(1+\alpha) + \beta u(1-\alpha)}{1-\beta^2}$$

for a high-income agent and

$$(4) \quad \frac{u(1-\alpha) + \beta u(1+\alpha)}{1-\beta^2}$$

for a low-income one. These are dominated by market participation and perpetual consumption of one unit if, and only if,

$$(5) \quad u(1+\alpha) + \beta u(1-\alpha) \leq (1+\beta)u(1).$$

This inequality holds under conditions similar to those enumerated in Alvarez and Jermann (2000, Proposition 4.9), which require that *all* individuals have a strong need for consumption smoothing. In particular, inequality (5) holds if all individuals have a low intertemporal elasticity of substitution, or a low rate of time preference, or are subject to large individual income shocks. Reasonable as they might seem for an economy with two agents, these conditions are difficult to achieve in an environment with a large variety of agent types, some of whom will necessarily experience small income shocks. In what follows we assume that inequality (5) fails⁶ and that autarky is a state with a low implied rate of interest. Specifically, we assume

$$(6) \quad u(1+\alpha) + \beta u(1-\alpha) > (1+\beta)u(1)$$

and

$$(7) \quad u'(1+\alpha) < \beta u'(1-\alpha).$$

In a more complicated model with a continuum of agents indexed by α , inequality (7) would have to hold for some interval of α , in particular for the highest values of α .⁷

⁶ If inequality (5) fails, then it is straightforward to show that high-income agents will prefer autarky to the perfect enforcement allocation for *any* initial distribution of debt, not just for the distribution assumed in this paper.

⁷ If the utility function were logarithmic, inequality (7) would require that the maximal value of α should exceed $(1-\beta)/(1+\beta)$, which implies that the maximal annual fluctuation in individual

These relations are shown in Figure 1, where the first-best allocation is on the diagonal and point *A* represents autarky. An implied interest factor of unity corresponds to point *M*.

Limited Enforcement

In environments where loan contracts are enforced by perpetual exclusion of defaulters from asset markets, equilibria are defined somewhat differently from standard models. In particular, an equilibrium is an infinite sequence, (c_t^H, c_t^L, R_t, b_t) , where b_t is the debt limit assigned to the low-income person at t . Agents maximize taking R_t and b_t as given, markets clear, and b_t is the largest possible debt limit that will keep borrowers at t from defaulting at date $t+1$. These limits must be binding by inequality (6), which states that the first-best allocation $(c_t^H, c_t^L) = (1, 1) \forall t$ is ruled out by debt limits. In particular, (i) the consumption Euler equation holds for the high-income agent and fails for the low-income agent; that is,

$$(8) \quad \beta R_t = \frac{u'(c_t^H)}{u'(c_{t+1}^L)} < \frac{u'(c_t^L)}{u'(c_{t+1}^H)}.$$

(ii) Budget constraints apply, with the low-income agent borrowing at the debt limit from the high-income agent; that is,

$$(9) \quad c_t^H = 1 + \alpha - R_{t-1}b_{t-1} - b_t$$

and

$$(10) \quad c_t^L = 1 - \alpha + R_{t-1}b_{t-1} + b_t.$$

(iii) Markets clear; that is,

$$(11) \quad c_t^H + c_t^L = 2.$$

And (iv) debt limits equate the autarkic and market payoffs for a high-income consumer who is about to repay last period's debt; specifically,

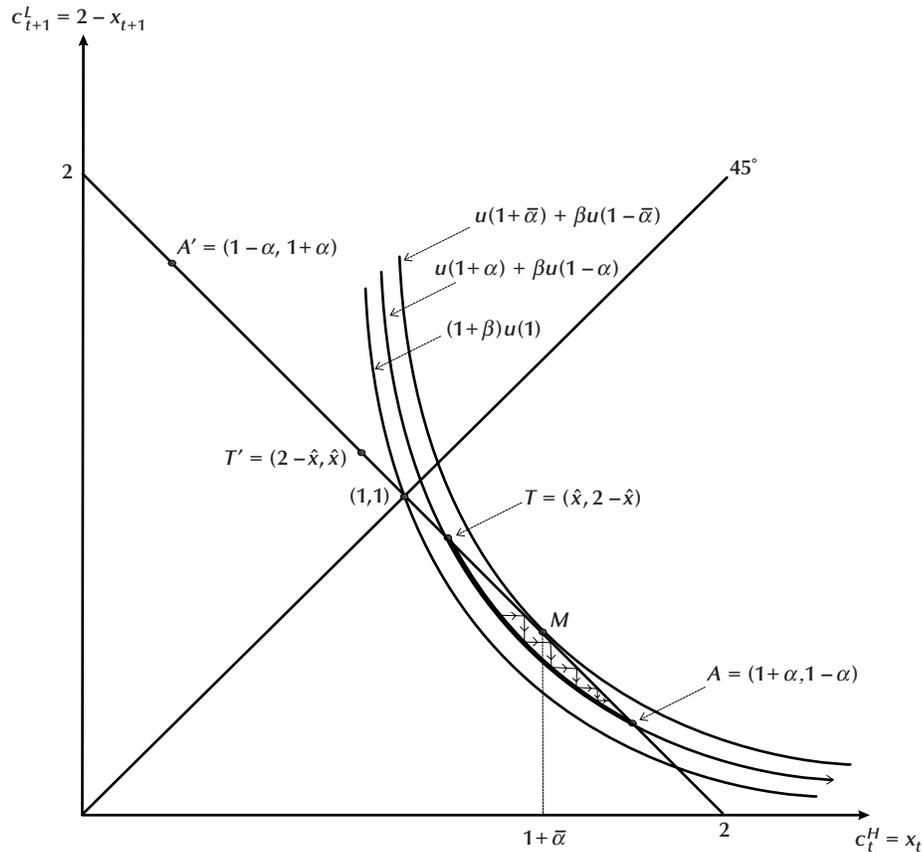
$$(12) \quad u(c_t^H) + \beta u(c_{t+1}^L) = u(1+\alpha) + \beta u(1-\alpha)$$

for all t .

income should be no less than approximately 2 percent. Hence, it seems quite plausible that the first-best allocation will be prevented by endogenous debt limits.

Figure 1

The Fundamental Diagram



If we define $c_t^H = x_t \in [1, 1 + \alpha]$, then it is clear that equilibria are solution sequences to equation (12), that is, to

$$(13) \quad u(x_t) + \beta u(2 - x_{t+1}) = u(1 + \alpha) + \beta u(1 - \alpha).$$

These sequences are shown in Figure 1.

Real Indeterminacy

If inequalities (6) and (7) hold, Figure 1 shows that there are two steady states. This first is a stable autarkic state, $(c_t^H, c_t^L, R_t, b_t) = (1 + \alpha, 1 - \alpha, \bar{R}, 0)$ for all t , where

$$(14) \quad \bar{R} = \frac{u'(1 + \alpha)}{\beta u'(1 - \alpha)} < 1.$$

This state corresponds to point A in Figure 1. The loan market is shut down in this state. The second is an unstable trade state, $(c_t^H, c_t^L, R_t, b_t) = (\hat{x}, 2 - \hat{x}, \hat{R}, \hat{b})$ for all t , where $\hat{x} \in (1, 1 + \alpha)$ is the unique solution to

$$(15) \quad u(x) + \beta u(2 - x) = u(1 + \alpha) + \beta u(1 - \alpha),$$

$$(16) \quad \hat{R} = \frac{u'(\hat{x})}{\beta u'(2 - \hat{x})},$$

and

$$(17) \quad \hat{b} = \frac{1 + \alpha - \hat{x}}{1 + \hat{R}}.$$

This state corresponds to point T in Figure 1. The loan market is active in this state. Because T lies between point M and the diagonal, we have

$$(18) \quad \hat{R} \in (1, 1/\beta).$$

Because autarky is associated with an interest factor below 1, and the trading state with an interest factor above 1, it follows from Alvarez and Jermann (2000, Proposition 4.6) that the trading state is constrained optimal⁸ and the autarkic state is not. Individual consumption shares fluctuate less in the constrained optimal state than they do in the autarkic state.

In addition to the two steady states, there is a continuum of equilibrium sequences (x_t) indexed on $x_0 \in (\hat{x}, 1 + \alpha)$, which converge to autarky. See again Figure 1. All of these sequences can be Pareto ranked by the initial consumption, x_0 .

Equilibrium outcomes are indeterminate in this nonmonetary economy for reasons that have nothing to do with the intertemporal elasticity of substitution in consumption or the lack of gross substitutes as commonly understood. Instead, indeterminacy in this environment comes from *dynamic complementarities* between current and expected future debt limits. In particular, low *future* debt limits reduce gains from future asset trading and lower the *current* payoff to solvency. This, in turn, raises the incentive to default, which must be deterred by tighter debt limits now.

We conclude that the constrained optimal allocation of consumption $(\hat{x}, 2 - \hat{x})$ can be achieved only if all future debt limits are expected to stay *exactly* at \hat{b} . Any other expectations will lead inevitably to autarky or to the nonexistence of equilibrium. In the remainder of the paper, we will explore whether, and how, policies can guide individual expectations in a manner that leads away from autarky and, perhaps, toward the constrained optimal allocation.

⁸ An allocation is constrained optimal if it satisfies the usual resource constraints, is weakly preferred to autarky by all agents at all times, and cannot be dominated by another feasible allocation.

PASSIVE FISCAL AND MONETARY POLICIES

Fiscal Policy with Zero Debt Limits

We explore here the possibility of achieving the constrained efficient allocation by a passive fiscal or monetary policy, that is, by choosing policy instruments that are invariant to the history of economic events. We start with a constant lump-sum tax, τ , on the high-income agent and an equal subsidy to the low-income agent. Any tax we choose must support an equilibrium allocation that is *weakly preferred to autarky by all agents at all times*. This feasibility requirement excludes tax and transfer schemes that would equalize post-transfer endowments in all periods, thus implementing the first-best allocation $x_t = 1 \forall t$ as an autarkic equilibrium. However, the policy $\tau = 1 + \alpha - \hat{x}$ shifts the endowment point from point A to point T in Figure 1 and implements the constrained optimal allocation as a *unique* post-transfer autarkic equilibrium at the high-interest yield, $\hat{R} \in (1, 1/\beta)$. All agents weakly prefer this outcome to the pre-transfer autarkic equilibrium at the low-interest yield, $\bar{R} < 1$.

The only problem with this policy is that it relies on precise information about individual incomes, especially if there were a large variety of income types. Policy in this setting must be able to tailor individual transfers to individual incomes. Are there simpler ways to achieve desirable outcomes with a blunter policy instrument that requires less information—that is, that does not discriminate between individuals?

Monetary Policy with Zero Debt Limits

One completely anonymous instrument is fiat money printed to pay an equal lump-sum transfer to all agents. Positive lump-sum transfers flatten the distribution of current resources among households, and negative transfers skew that distribution in favor of high-income persons. This, in turn, enables monetary policy to control the real yield on money,⁹ which is the reciprocal

⁹ Identical outcomes can be achieved by changes in the stock of public debt because money and debt are perfect substitutes in our economy. We use the term “monetary policy” advisedly here

of the inflation factor, along any equilibrium path. To see this, we let M_t be the stock of money per agent, μ_t be the gross rate of money growth, p_t be the price level, τ_t be the real value of the transfer, and $m_t = M_t/p_t$. Policymakers choose the sequence (μ_t) under the restriction that the resulting monetary equilibrium is weakly preferred to autarky by all agents at all times. We assume that agents have the option of rejecting monetary transfers and taxes in favor of autarky.

Assuming for the moment that debt limits are zero (we relax this assumption in the next subsection), budget constraints are

$$(19) \quad c_t^H = 1 + \alpha + \tau_t - m_t^d,$$

where m_t^d is the demand for money by high-income agents and

$$(20) \quad c_{t+1}^L = 1 - \alpha + \tau_{t+1} + R_t m_t^d,$$

where $R_t = p_t/p_{t+1}$ is the real rate of return on money. Low-income agents are assumed to be rationed and to spend their entire money balances to raise current consumption.

Equilibrium in this economy satisfies the consumption Euler equation (8) for the high-income agent, rewritten here as

$$(21) \quad u'(c_t^H) = \beta R_t u'(c_{t+1}^L),$$

as well as equilibrium in the goods and money markets; that is,

$$(22) \quad c_t^H + c_t^L = 2,$$

$$(23) \quad m_t^d = 2m_t.$$

In addition, individual budget constraints apply; that is,

$$(24) \quad c_t^H \equiv x(m_t, \mu_t) = 1 + \alpha - \left(1 + \frac{1}{\mu_t}\right)m_t.$$

The real return on money is

$$(25) \quad R_t = \frac{p_t}{p_{t+1}} = \frac{m_{t+1}}{\mu_{t+1} m_t}.$$

We conclude that, for a given policy sequence (μ_t) , equilibria are bounded non-negative solution sequences (m_t) to the nonautonomous equation,

$$(26) \quad u'[x(m_t, \mu_t)] = \beta \frac{m_{t+1}}{\mu_{t+1} m_t} u'[2 - x(m_{t+1}, \mu_{t+1})].$$

Equilibria that converge to autarky are driven by self-confirming inflationary expectations that reduce the demand for real money balances and diminish trading between high- and low-income agents. One important drawback of passive policies is that they are unable to connect future returns on money with the current state of the economy and therefore cannot counter inflationary expectations with tighter monetary policy, that is, by lowering μ_t .

Figure 2 shows the qualitative properties of solutions that correspond to a *passive* monetary policy, $\mu_t = \mu \in [1/\hat{R}, 1/\bar{R}] \forall t$, for an economy in which dated consumption goods are gross substitutes. Each policy is associated with two steady states: a stable autarkic state with $m = 0$ and an unstable trading state with $m^*(\mu) > 0$. In general, higher values of μ correspond to lower steady-state returns on money and to a lower demand for money, $m^*(\mu)$. For example, $\mu = 1/\hat{R} \in (\beta, 1)$ supports the constrained optimal trading state, $c_t^H = \hat{x}$, for all t by raising the value of real balances to

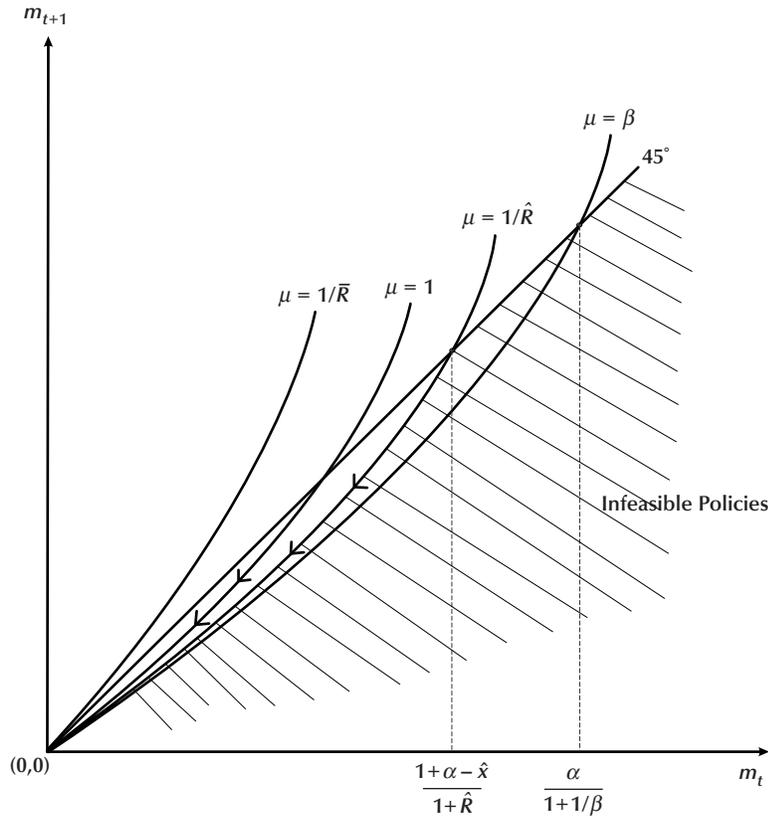
$$(27) \quad m^*(1/\hat{R}) = \frac{1 + \alpha - \hat{x}}{1 + \hat{R}}.$$

This value of μ , which involves a mild deflation at the steady state of an economy with zero income growth, is the lowest feasible rate of growth consistent with all agents preferring monetary equilibrium to autarky at all times. The Friedman rule, $\mu = \beta$, which would support the first-best allocation $c_t^H = 1$ for all t is simply not feasible: It imposes too large a tax on high-income agents, causing them to choose autarky over the use of money. As we raise the value of μ above $1/\hat{R}$, the amount of trading between the two groups of agents shrinks, vanishing at $\mu = 1/\hat{R}$; here the steady-state return on money is equal to the autarkic rate of return.

Figure 2 reveals that feasible passive policies cannot overcome the indeterminacy problem of

because we expect that our results carry over to economies where debt dominates money in rate of return.

Figure 2
Passive Policies



self-fulfilling inflationary expectations that leads to reduced trade among households. To solve this problem, policy should connect the current state of the economy with future returns on money, that is, with expectations of future inflation. We explore active feedback policies in the “Active Monetary Policy” section.

Monetary Policy with Positive Debt Limits

Positive debt constraints add nothing essential to the equilibria described in the previous subsection because holdings of private debt are a perfect substitute for balances of fiat money. A

high-income individual has the same payoff when trading loans as he would if he held only money. Loan default does not hurt the trading opportunities of any individual.

If debt limits are a positive sequence (b_t), the budget constraint for a high-income individual,

$$\begin{aligned}
 (28) \quad c_t^H &= 1 + \alpha - \left(1 + \frac{1}{\mu_t}\right) m_t - R_{t-1} b_{t-1} - b_t \\
 &= 1 + \alpha - (m_t + b_t) - R_{t-1} (m_{t-1} + b_{t-1}),
 \end{aligned}$$

indicates that only the size of the asset portfolio matters for individual plans, not its division into fiat money and debt. Money displaces private loans at a one-to-one rate.

ACTIVE MONETARY POLICY

Suppose next that monetary policy can control directly the real yield on money balances in a manner that depends on the entire state of the economy. We index that state by (x_t, R_t) , where x_t is twice the consumption share of the high-income agent and R_t is the real yield on money and debt. Equivalently, we may index the state by (m_t, R_t) , where m_t is real balances per capita. Then we write the policy function as

$$(29) \quad R_{t+1} = f(x_t, R_t).$$

The arbitrary function f maps the product space $[1, 1 + \alpha] \times [\bar{R}, \hat{R}]$ into $[\bar{R}, \hat{R}]$. It should be consistent with autarky (because autarky is an equilibrium whenever households refuse to accept fiat money in exchange for goods) and with the constrained optimal state (\hat{x}, \hat{R}) (because this state is a reasonable target for a benevolent policymaker). This requires that

$$(30) \quad \bar{R} = f(1 + \alpha, \bar{R}),$$

$$(31) \quad \hat{R} = f(\hat{x}, \hat{R}).$$

An example of this type of policy is

$$(32) \quad f(x, R) = \begin{cases} \bar{R} & \text{if } (x, R) = (1 + \alpha, \bar{R}) \\ \frac{2 - \hat{x}}{\beta(2 - \beta R x)} & \text{otherwise.} \end{cases}$$

Given the policy f , an equilibrium is a sequence (x_t, R_t) that satisfies (29), plus the consumption Euler equation of the high-income agent, rewritten here as

$$(33) \quad u'(x_t) = \beta R_t u'(2 - x_{t+1}).$$

By construction, the dynamical system (29) and (33) has two steady states, $(x, R) = (1 + \alpha, \bar{R})$ and (\hat{x}, \hat{R}) . A sensible policy $f(\cdot)$ ensures that equilibria starting at any point (x_0, R_0) move away from the suboptimal autarkic state $(1 + \alpha, \bar{R})$ and converge rapidly to the constrained optimal state (\hat{x}, \hat{R}) .

We do not attempt a global characterization of policies that achieve this objective for an arbitrary utility function $u(\cdot)$. We instead confine ourselves to exploring the properties of such

policies for a logarithmic utility function in the neighborhood of each steady state. In this class of economies, equilibria satisfy

$$(34) \quad x_{t+1} = 2 - \beta R_t x_t,$$

$$(35) \quad R_{t+1} = f(x_t, R_t).$$

All we need do is study the characteristic polynomial $p(\lambda)$ in the neighborhood of any steady state (x, R) . That polynomial is

$$(36) \quad g(\lambda) = \lambda^2 + \beta R(\beta R - \varepsilon_R)\lambda + \beta R(\varepsilon_x - \varepsilon_R),$$

where ε_x and ε_R are partial elasticities of the policy function f with respect to x and R , and $\beta R < 1$ at each steady state.

Desirable policies, as we have described them, should turn the constrained optimal steady state into an attractor, or sink, and the suboptimal state into a source. The eigenvalues, or roots of the polynomial $g(\lambda)$, should be inside the unit circle at $(x, R) = (\hat{x}, \hat{R})$ and outside the unit circle at (\bar{x}, \bar{R}) . One way to choose is to focus on functions f that raise future real yields whenever households with currently high incomes consume “too much” (relative to the efficient outcome, \hat{x}) and demand “too little” money.

It is easy to check whether the policy function in equation (32) has exactly this property, which would furthermore guarantee immediate convergence to the constrained efficient state (\hat{x}, \hat{R}) from any initial condition other than autarky. Under this policy, the dynamical system consisting of equations (34) and (35) has a double real eigenvalue with modulus zero at (\hat{x}, \hat{R}) .

Exactly what does the monetary authority have to do to control the real rate of interest in the manner specified by the policy function in (32)? This can be answered with a logarithmic utility function: Combine the budget constraint (24) with the Euler equation (26) to obtain an expression that connects monetary policy at date $t+1$ with monetary policy at t and with the state of the economy at t . In particular, we find that

$$(37) \quad 1 + \mu_{t+1} = \frac{1 + \mu_t \left(\frac{\beta x_t - \frac{1 - \alpha}{R_t}}{1 + \alpha - x_t} \right)}{\mu_t}.$$

We conclude that monetary policy tightens (μ_{t+1} falls) subsequent to a rise in the rate of inflation and a drop in the rate of return on money. Equation (37) says that the tightening appears to be substantial. For example, if α is small relative to unity, β is about 0.95, and the ratio of money to income is about 1:7; then, each additional 1 percentage point of inflation near the constrained efficient steady state causes the money growth rate to drop by about 7 percentage points. To see this, we calculate the response of the money growth rate to changes to the past inflation rate from equations (37) and (24) and obtain

$$\frac{d\mu_{t+1}}{d(1/R_t)} = \frac{1-\alpha}{m_t}.$$

Then, we set $\alpha = 0$ and $m_t = 1/7$.

This is the sort of strongly reactive policy that guides inflationary expectations to just the level needed to support the constrained efficient outcome. In a similar but not quite as effective way, outcomes can be achieved if the policy rule simply maps current inflation into future inflation, ignoring current quantities such as x_t and m_t . Specifically, if we employ the rule

$$(38) \quad R_{t+1} = \phi(R_t),$$

where ϕ maps the interval $[\bar{R}, \hat{R}]$ into itself, then the dynamical system consisting of equations (34) and (38) has two real eigenvalues, $-\beta R$ and $\phi'(R)$, at any steady state (x, R) . Recall that $\beta R \in (0, 1)$ at both \bar{R} and \hat{R} .

Therefore, any policy rule such that $|\phi'(\bar{R})| > 1$ and $|\phi'(\hat{R})| < 1$ will convert the autarkic state A into a saddle and the trading state T into a sink. For most initial conditions (x_0, R_0) , monetary policy leads the economy to converge asymptotically, but not immediately, to the constrained efficient state.

CONCLUSIONS AND EXTENSIONS

This paper provides general equilibrium examples of how active monetary policy can be used to select a desirable outcome in economies where passive policies are associated with many Pareto-ranked dynamic equilibria.

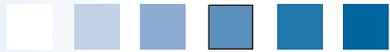
In our setting, monetary policy works directly on rational beliefs about future values of the inflation rate, debt limits, and other financial variables. It does so by *committing* to a feedback rule that connects current financial conditions with future values of the policy instrument and, in particular, to a shared belief that asset returns will improve substantially when the volume of asset trading falls below what is consistent with an efficient allocation of resources.

When viewed as an exercise in equilibrium selection, monetary policy is an attempt to foster expectations that lead to socially desirable states of the economy as rapidly as possible. This attempt is completely successful in our simple setting where money is a perfect substitute for private debt, that is, a store of value for two agents trading in complete, albeit imperfect, asset markets. Efficiency in this setting is achieved when the volume of loans is as large as capital-market imperfections will allow. If that volume is less than it should be, properly valued money can act as a substitute for private loans. The job of the central bank is to defend the correct value of money by connecting expectations of future inflation with current economic conditions and intervening aggressively to pin inflation expectations to the right value.

We are not sure that a blunt policy instrument such as anonymous monetary policy will be as successful in selecting constrained optimal outcomes in a richer environment with many agents and uncertainty. In particular, if money and debt are imperfect substitutes because the former has a liquidity advantage over the latter, then monetary policy has implications for debt limits and for the participation of households in financial markets. In addition, policy choices may not be conditioned on the entire state of the economy if that state includes detailed information about individual incomes and trading plans. In that case, the policymaker may have to settle for something less than constrained efficiency, as in Benhabib and Eusepi (2005). These implications need to be carefully explored before we can design monetary rules with any degree of confidence.

REFERENCES

- Alvarez, Fernando and Jermann, Urban J. "Efficiency, Equilibrium, and Asset Pricing with Risk of Default." *Econometrica*, July 2000, 68(4), pp. 775-97.
- Azariadis, Costas and Kaas, Leo. "Asset Price Fluctuations Without Aggregate Shocks." *Journal of Economic Theory*, 2007 (forthcoming).
- Benhabib, Jess and Eusepi, Stefano. "The Design of Monetary and Fiscal Policy: A Global Perspective." *Journal of Economic Theory*, July 2005, 123(1), pp. 40-73.
- Benhabib, Jess and Farmer, Roger E.A. "Indeterminacy and Sunspots in Macroeconomics," in J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*. Volume 1A. Amsterdam: North-Holland, 1999, pp. 387-448.
- Benhabib, Jess; Schmitt-Grohé, Stephanie and Uribe, Martin. "The Perils of Taylor Rules." *Journal of Economic Theory*, January/February 2001, 96(1-2), pp. 40-69.
- Bewley, Truman. "The Optimum Quantity of Money," in J. Kareken and N. Wallace, eds., *Models of Monetary Economies*. Federal Reserve Bank of Minneapolis, 1980.
- Boldrin, Michele and Woodford, Michael. "Equilibrium Models Displaying Endogenous Fluctuations and Chaos: A Survey." *Journal of Monetary Economics*, March 1990, 25(2), pp. 189-222.
- Bullard, James B. and Mitra, Kaushik. "Learning about Monetary Policy Rules." *Journal of Monetary Economics*, September 2002, 49(6), pp. 1105-29.
- Cooper, Russel W. *Coordination Games: Complementarities and Macroeconomics*. Cambridge, MA: MIT Press, 1999.
- Eaton, Jonathan and Gersovitz, Mark. "Debt with Potential Repudiation: Theoretical and Empirical Analysis." *Review of Economic Studies*, April 1981, 48(2), pp. 289-309.
- Kehoe, Timothy J. and Levine, David K. "Debt-Constrained Asset Markets." *Review of Economic Studies*, October 1993, 60(4), pp. 865-88.
- Sargent, Thomas J. and Wallace, Neil. "Rational Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule." *Journal of Political Economy*, April 1975, 83(2), pp. 241-54.
- Townsend, Robert M. "Models of Money with Spatially Separated Agents," in J. Kareken and N. Wallace, eds., *Models of Monetary Economies*. Federal Reserve Bank of Minneapolis, 1980.
- Woodford, Michael. "The Taylor Rule and Optimal Monetary Policy." *American Economic Review*, May 2001, 91(2), pp. 232-37.
- Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press, 2003.
- Wright, Randall. "Introduction to 'Models of Monetary Economies II: The Next Generation.'" *International Economic Review*, May 2005, 46(2), pp. 305-16.



Commentary

Peter N. Ireland

In their conference paper, Gaetano Antinolfi, Costas Azariadis, and James Bullard (2007) develop and analyze a macroeconomic model with heterogeneous agents in which individual incomes fluctuate, aggregate income remains constant, and frictions that inhibit the strict enforcement of private contracts place endogenous limits on agents' ability to borrow and lend and hence to engage in intertemporal trade. Further, because the model features only a single good, intertemporal trade is the only trade that potentially takes place in equilibrium. Bewley (1980), Townsend (1980), Kehoe and Levine (1993), Kocherlakota (1996), and Alvarez and Jermann (2000) previously and famously considered similar models. Here, however, Antinolfi, Azariadis, and Bullard go beyond all of this previous work by highlighting that these models typically feature multiple equilibria.

Here, in fact, the authors' model has two steady-state equilibria under *laissez-faire*. In one, no trade takes place, so that equilibrium allocations are autarkic; in the other, agents trade actively. Hence, the two steady states can be Pareto-ranked: All agents prefer the good equilibrium with trade to the bad equilibrium without. The authors' policy problem then arises, because the bad steady state is stable and the good steady state is unstable, implying that even if the economy begins arbitrarily close to but not exactly in the good steady state, it will converge over time to the bad steady state. In Antinolfi, Azariadis, and Bullard's analysis, the government's stabiliza-

tion policy aims at keeping the economy at or near the good steady state.

Stabilization policy in this analysis therefore plays an important but somewhat unfamiliar role. Typically, in mainstream macroeconomic models, stabilization policy calls for the monetary and fiscal authorities to adjust their policy instruments in response to shocks that buffet the economy around a given steady state. In Antinolfi, Azariadis, and Bullard's model, by contrast, stabilization policy works on a more fundamental level, to actually pick out the steady state toward which the economy gravitates. Hence their paper's title, "Monetary Policy as Equilibrium Selection."

Here, monetary policy helps achieve this stabilization goal by reversing the properties of the two steady states, rendering the good steady state stable and the bad steady state unstable.

Specifically, the authors show that active policies that call for the government to adjust its policy instruments vigorously in response to changes in the underlying state of the economy succeed in achieving this goal. By contrast, passive policies—including constant money growth rate rules—that call for little or no policy response to changes in the economy fail by leaving the bad steady state as the economy's most likely destination.

Antinolfi, Azariadis, and Bullard's analysis, results, and conclusions combine to make their paper quite interesting and useful. The paper is novel in its focus on active versus passive policy rules in models of the type used in Bewley (1980) and Townsend (1980). Ljungqvist and Sargent

Peter N. Ireland is a professor of economics at Boston College and a research associate at the National Bureau of Economic Research.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 343-48.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

Ireland

(2000), for example, take a far more limited approach to policy analysis in a version of Townsend's (1980) turnpike model. Specifically, Ljungqvist and Sargent simply assume that the economy starts in its good steady state and then ask what the optimal constant rate of money growth in that good steady state is. Antinolfi, Azariadis, and Bullard qualify and extend these earlier results in an important way by making clear that Ljungqvist and Sargent's preferred constant money growth rate rule does not prevent the economy from leaving a neighborhood of its good steady state and converging to the bad steady state instead.

By highlighting the importance of this active-versus-passive distinction for the design of welfare-enhancing monetary policy, Antinolfi, Azariadis, and Bullard's paper also becomes quite useful, as it draws previously unnoticed links between the branch of the literature that works with Bewley-Townsend-type models and another branch of the literature in monetary economics that works with a very different class of models. In particular, recent work with New Keynesian models featuring monopolistic competition and staggered nominal price setting in goods markets establishes what Woodford (2003) and others call the "Taylor principle." This Taylor principle indicates that the central bank can stabilize the inflation rate around a desired target value through the use of an interest rate rule for monetary policy of the kind proposed by Taylor (1993), provided that rule is sufficiently active, calling for a vigorous adjustment of the short-term nominal interest rate instrument in response to shocks that push the inflation rate away from target. Antinolfi, Azariadis, and Bullard's results favor the use of active monetary policy rules as well, helping to establish the generality and robustness of these findings across two otherwise divergent branches of inquiry.

This new paper by Antinolfi, Azariadis, and Bullard thereby contributes importantly to the literature. It extends, as the other conference papers do, the "Frontiers in Monetary Policy Research." What's more, like many other papers that extend the frontiers of research—particularly

in monetary economics, it seems—this new paper raises a host of additional questions at the same time that it provides answers to existing ones. The remainder of my discussion focuses on some of these additional questions, pointing as Antinolfi, Azariadis, and Bullard's paper itself does to promising avenues for future research.

IS PUBLIC POLICY REALLY NECESSARY?

This first and most basic question asks whether public policy is really crucial in an economic environment like the one described by Antinolfi, Azariadis, and Bullard's model. In their paper, the authors themselves provide a partial response to this question by indicating that the answer is "no" if private credit markets work well to begin with. In particular, the authors show that, when contracts can be perfectly enforced, trading in private credit markets supports an equilibrium allocation that is Pareto optimal. In this special case, government policy cannot help; *laissez-faire* works best.

Yet one might go a step further, as I am tempted to do, and note that even with limited contractual enforcement, the scope for welfare-enhancing public policy, though present, will necessarily be limited to the extent that the autarkic equilibrium is really not so bad. I raise this possibility with a specific concern in mind. The point is that all of these terms—"bad equilibrium," "autarkic allocations," "unstable steady states," and so on—have very specific meanings when used in the context of a formal study in macroeconomic theory like Antinolfi, Azariadis, and Bullard's. Of course, the authors very carefully and properly use these terms in their paper. However, the risk remains that, when presented to a broader audience of nonspecialists and policymakers, these words will unintentionally conjure up images of disastrous outcomes under *laissez-faire*; in fact, though, a full, quantitative assessment of the welfare properties of equilibrium outcomes with and without government intervention—perhaps along the same lines as that presented by Krueger and Perri (2005) but applied

to the specific environment studied here—remains a task for future research.

In his famous essay “The Role of Monetary Policy,” Milton Friedman (1968, p. 14) cautions against the tendency toward overconfidence in economists offering policy advice: “[I]n this area particularly,” he warns, “the best is likely to be the enemy of the good.” It is almost surely true that, in reality, as in Antinolfi, Azariadis, and Bullard’s model, frictions prevent private markets—especially private credit markets—from operating with total efficiency so as to bring equilibrium allocations in line with Pareto-optimal outcomes. Yet, as Friedman emphasizes, it seems equally true that, in reality, even the most carefully designed government policies introduced into environments in which outcomes under *laissez-faire* are clearly suboptimal have often made matters much worse instead of much better. The inefficiencies in private credit markets are usefully highlighted in Antinolfi, Azariadis, and Bullard’s model. But, before we lean too heavily on those inefficiencies as the basis for justifying activist government intervention in those same segments of the U.S. economy, future research must more forcefully establish that those inefficiencies are severe enough, quantitatively, to also justify the risk that a well-designed public policy will be poorly implemented or will otherwise have unintended and detrimental consequences. Many sad lessons from history teach us that “reversion to autarky,” in the vernacular as opposed to the language for formal economic theory, most frequently occurs precisely because of excessive government involvement in private markets.

IS MONETARY POLICY REALLY NECESSARY?

Although fiscal policy, in the form of a carefully designed system of income taxes and transfers, might seem to be the most direct and effective way of helping private agents in Antinolfi, Azariadis, and Bullard’s model stabilize their consumptions in the face of their fluctuating income streams, the authors point out that the successful implementation of such a policy requires the

government to obtain and exploit detailed information about individual agents’ economic circumstances. On these grounds, they advocate the search for monetary policy rules that help accomplish the same goal of income redistribution.

Along the same lines, however, one might also note that the monetary policy rule that the authors propose later, shown in their equation (37), requires the central bank to adjust the rate of money growth in response not just to movements in the aggregate variable R , which measures the real return to money (or the inverse of the inflation rate), but also to the variable x , which measures not aggregate income or consumption but rather the share of aggregate consumption enjoyed by high-income agents. In a more complicated model with richer forms of heterogeneity, the analog to the variable x would be a statistic or set of statistics summarizing the cross-sectional distribution of consumption. Successful implementation of this preferred monetary policy, therefore, also requires the government to collect and process much of the same individual-specific data needed to run an optimal tax-and-transfer fiscal scheme.

For this reason, an alternative policy rule that takes the form of the authors’ equation (38) and therefore calls for a monetary response to changes in the aggregate variable R alone may represent a more appealing and realistic alternative to pure *laissez-faire* or to a passive constant money growth rate rule. In any case, working out the implications of private information and the incentives that the government can offer agents to truthfully reveal that private information in settings like that described by Antinolfi, Azariadis, and Bullard’s model remains another important task for future research; those implications may draw sharper and more reliable distinctions between fiscal and monetary policies as effective tools for income redistribution.

IS TIME CONSISTENCY A PROBLEM?

In Antinolfi, Azariadis, and Bullard’s model, activist policy works to stabilize the economy

Ireland

around its good steady state by influencing private expectations of future inflation under various contingencies that arise both in and out of equilibrium. As Kydland and Prescott (1977) emphasize, however, public policymakers who make announcements in an attempt to shape private expectations often fall victim to the time-consistency problem. Once private expectations based on a policymaker's announcements have been built into private decisions, that same policymaker may have an incentive to deviate from his or her promised action. The problem then arises because private agents recognize that the policymaker has this incentive to renege on any initial promise. In equilibrium, a policymaker without the ability to commit strongly to a preannounced policy may be unable to influence expectations in the desired way.

All of Antinolfi, Azariadis, and Bullard's analysis proceeds under the assumption that the central bank has this ability to commit. At the same time, however, their model builds directly and importantly on the idea that private agents' inability to precommit to their own future actions is precisely what provides room for welfare-enhancing public policy in the first place. What justifies this assumption that the government faces no similar commitment problem? And if the optimal activist monetary policy rules shown in equations (37) and (38) turn out to be time inconsistent, how do optimal policies under discretion compare with these counterparts under commitment, both in terms of their implications for the behavior of the money stock and inflation and in terms of their ability to stabilize the economy around the good steady state? These questions, too, remain to be answered in future research.

IS CREDIBILITY A PROBLEM? WHICH EQUILIBRIA ARE EXPECTATIONALLY STABLE?

In addition to the time-consistency problem described above, a second potential difficulty may arise when the central bank tries to use the optimal activist policies described by equations

(37) and (38) to stabilize the economy around its good steady state: Once the economy reaches the good steady state—immediately under (37) and eventually under (38)—these activist policies call for constant money growth and inflation rates and may therefore appear to private agents as being observationally equivalent to passive policies, such as a constant money growth rate rule. Hence, once the economy reaches the good steady state, either of these activist policies retains its power to stabilize the economy only through the effects that the central bank's commitment to the policy rule has on private expectations of what would happen, out of equilibrium, if the economy begins to slip away from that good steady state.

Given the potential tenuousness of the expectational forces keeping the economy in the good steady state, even under an active monetary policy rule, one might reasonably ask, What would happen if, instead of forming their expectations based on how they believe the government would behave out of equilibrium, private agents formed their expectations based on how they actually observe the government to behave in equilibrium? Would the central bank have to act, periodically at least, to maintain the credibility of its commitment to the optimal rule?

Often, in the literature following Kydland and Prescott (1977), "credibility" is used synonymously with "time consistency." In this case, however, the term as I use it refers to ideas that are closer in spirit to the concepts of "expectational stability" and "learnability" that, in previous work, Bullard (2006) uses to characterize the government's ability to keep the economy in or around a desired steady state when private agents form their expectations adaptively, based on historical data as opposed to full knowledge of the economy's true structure.

Examining the need and scope for activist monetary policy to stabilize the economy described by Antinolfi, Azariadis, and Bullard's model around the good steady state when expectations are formed through adaptive learning also remains an important and useful task for future research.

ARE ACTIVE POLICIES ROBUSTLY OPTIMAL?

This last question comes full circle, back to Milton Friedman's (1968) caveats about activist public policymaking. Antinolfi, Azariadis, and Bullard's main result, concerning the optimality of activist policy rules, seems quite sensible: If the central bank wants to stabilize the economy around a desirable steady state, then it certainly stands to reason that its monetary policy ought to react strongly whenever the economy begins to deviate from that steady state. The authors' main result shares the same powerful, intuitive appeal as the Taylor principle from the literature on New Keynesian economics.

However, their statement about robustness—that, looking across many different macroeconomic models, optimal policy rules are all activist—remains logically distinct from (and therefore does not imply) another statement about robustness: that any given activist policy rule, fine-tuned to fit the special features of any given model, will continue to work well across many different macroeconomic models. Barnett and He (2002) make this point quite forcefully, using methods and arguments that are quite similar to Antinolfi, Azariadis, and Bullard's.

In this earlier paper, Barnett and He focus on a macroeconomic model that is quite different from the one studied here by Antinolfi, Azariadis, and Bullard; specifically, Barnett and He work with an older-style, medium-scale macroeconomic model developed originally by Bergstrom, Nowman, and Wymer (1992). Nevertheless, Barnett and He begin their analysis just as Antinolfi, Azariadis, and Bullard do, by demonstrating that, although the Bergstrom-Nowman-Wymer model has an unstable steady state under *laissez-faire*, it can be stabilized by an appropriately designed activist fiscal policy rule. At the same time, however, Barnett and He also show that this activist fiscal policy rule, when properly calibrated to stabilize the economy under a given configuration of the model's nonpolicy parameters, works counterproductively to destabilize the economy still further when improperly calibrated to a slightly different set of nonpolicy parameters.

Barnett and He's results thereby echo Friedman's caveat about the best being the enemy of the good by confirming that an activist policy that is fine-tuned to work well within one particular model may perform quite poorly when applied to a very similar, but still slightly different, economic environment. Barnett and He's results clearly indicate that additional careful and rigorous analyses like Antinolfi, Azariadis, and Bullard's are needed to establish the robustness of optimal activist fiscal and monetary policies.

CONCLUSION

Antinolfi, Azariadis, and Bullard's conference paper contributes to scientific knowledge in several ways. It stands as the first paper to consider the important distinction between active and passive policy rules in a heterogeneous-agent model with endogenously incomplete markets that builds on Bewley's (1980) and Townsend's (1980) early formulations. By considering this distinction and by highlighting the stabilizing powers of activist monetary policy rules, it also draws useful and previously unnoticed links between the branch of the literature in monetary economics that studies the properties and implications of Bewley-Townsend-type models and the until-now completely distinct branch of the literature that studies New Keynesian models of monopolistic competition and nominal price rigidity. Finally, Antinolfi, Azariadis, and Bullard's paper contributes to scientific knowledge by raising a host of questions for future researchers who share these authors' technical sophistication, fine attention to detail, and intellectual rigor.

Before closing, let me ask some of these questions again, phrasing them in a slightly different way than they appear in my discussion above. The optimal activist monetary policy characterized by Antinolfi, Azariadis, and Bullard's equation (37) calls, in the authors' own words (p. 340), for the "money growth rate to drop by about 7 percentage points" in response to "each additional 1 percentage point of inflation." Is this optimal policy time consistent? Is this optimal policy credible or expectationally stable? Is this optimal

Ireland

policy robust to various changes in the economy environment? And is this optimal policy really necessary? All of these questions await the same type of careful and rigorous analysis contained in Antinolfi, Azariadis, and Bullard's very fine conference paper.

REFERENCES

- Alvarez, Fernando and Jermann, Urban J. "Efficiency, Equilibrium, and Asset Pricing with Risk of Default." *Econometrica*, July 2000, 68(4), pp. 775-97.
- Antinolfi, Gaetano; Azariadis, Costas and Bullard, James B. "Monetary Policy as Equilibrium Selection." Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 331-41.
- Barnett, William A. and He, Yijun. "Stabilization Policy as Bifurcation Selection: Would Stabilization Policy Work if the Economy Really Were Unstable?" *Macroeconomic Dynamics*, November 2002, 6(5), pp. 713-47.
- Bergstrom, A.R.; Nowman, K.B. and Wymer, C.R. "Gaussian Estimation of a Second Order Continuous Time Macroeconometric Model of the UK." *Economic Modelling*, October 1992, 9(4), pp. 313-51.
- Bewley, Truman. "The Optimum Quantity of Money," in John H. Kareken and Neil Wallace, eds., *Models of Monetary Economies*. Federal Reserve Bank of Minneapolis, 1980.
- Bullard, James B. "The Learnability Criterion and Monetary Policy." Federal Reserve Bank of St. Louis *Review*, May/June 2006, 88(3), pp. 203-17.
- Friedman, Milton. "The Role of Monetary Policy." *American Economic Review*, March 1968, 58(1), pp. 1-17.
- Kehoe, Timothy J. and Levine, David K. "Debt-Constrained Asset Markets." *Review of Economic Studies*, October 1993, 60(205), pp. 865-88.
- Kocherlakota, Narayana R. "Implications of Efficient Risk Sharing without Commitment." *Review of Economic Studies*, October 1996, 63(217), pp. 595-609.
- Krueger, Dirk and Perri, Fabrizio. "Public versus Private Risk Sharing." Unpublished manuscript. Department of Economics of Goethe University, Frankfurt, December 2005.
- Kydland, Finn E. and Prescott, Edward C. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, June 1977, 85(3), pp. 473-92.
- Ljungqvist, Lars and Sargent, Thomas J. *Recursive Macroeconomic Theory*. Cambridge: MIT Press, 2000.
- Taylor, John B. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, December 1993, 39, pp. 195-214.
- Townsend, Robert M. "Models of Money with Spatially Separated Agents," in John H. Kareken and Neil Wallace, eds., *Models of Monetary Economies*. Federal Reserve Bank of Minneapolis, 1980.
- Woodford, Michael. *Interest and Prices; Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press, 2003.

Model Fit and Model Selection

Narayana R. Kocherlakota

This paper uses an example to show that a model that fits the available data perfectly may provide worse answers to policy questions than an alternative, imperfectly fitting model. The author argues that, in the context of Bayesian estimation, this result can be interpreted as being due to the use of an inappropriate prior over the parameters of shock processes. He urges the use of priors that are obtained from explicit auxiliary information, not from the desire to obtain identification. (JEL C11, E40, E60)

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 349-60.

In an influential recent paper, Smets and Wouters (2003) construct a dynamic stochastic general equilibrium (DSGE) model with a large number of real and nominal frictions and estimate the unknown parameters of the model using sophisticated Bayesian techniques. They document that the estimated model has out-of-sample forecasting performance superior to that of an unrestricted vector autoregression. They write of their findings (p. 1125), “This suggests that the current generation of SDGE [stochastic dynamic general equilibrium] models with sticky prices and wages is sufficiently rich to capture the stochastics and the dynamics in the data, as long as a sufficient number of structural shocks is considered. *These models can therefore provide a useful tool for monetary policy analysis*” (italics added for emphasis). The European Central Bank (ECB) agrees. They are planning to begin using models with explicit micro-foundations for the first time in their analyses of monetary policy. In doing so, they are explicitly motivated by the Smets and Wouters (2003) analysis.¹

Smets and Wouters and the ECB are adherents to what one might call the *principle of fit*. According to this principle, models that fit the available data well should be used for policy analysis; models that do not fit the data well should not be. The principle underlies much of applied economic analysis. It is certainly not special to sophisticated users of econometrics: Even calibrators who use little or no econometrics in their analyses believe in the principle of fit. Indeed, there are literally dozens of calibration papers concerned with figuring out what perturbation in a given model will lead it to fit one or two more extra moments (like the correlation between hours and output or the equity premium).

In this paper, I demonstrate that the principle of fit does not always work. I construct a simple example economy that I treat as if it were the true world. In this economy, I consider an investigator who wants to answer a policy question of interest and estimates two models to do so. I show that

¹ See www.ecb.int/home/html/researcher_swm.en.html for details.

Narayana R. Kocherlakota is a professor of economics at the University of Minnesota, a consultant at the Federal Reserve Bank of Minneapolis, and a research associate at the National Bureau of Economic Research. The author thanks Ricardo DiCecio, Lee Ohanian, Tom Sargent, Adam Slawski, Hakki Yazici, Stan Zin, and especially Barbara McCutcheon for conversations about this paper. He learned much of what is in this paper from joint work that he did in the 1990s with Beth Ingram and Gene Savin when they were colleagues at the University of Iowa. The views expressed herein are the author's and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

model 1, which has a perfect fit to the available data, may actually provide worse answers than model 2, which has an imperfect fit.

The intuition behind this result is quite simple. The policy question of interest concerns how labor responds to a change in the tax rate. The answer depends on the elasticity of the labor supply. In both models, the estimate of this parameter hinges on a particular non-testable assumption about how stochastic shocks to the labor-supply curve covary with tax rates. When model 2's identification restriction is closer to being correct than model 1's, model 2 provides a better answer to the policy question, even though its fit is always worse.

In the second part of the paper, I consider a potential fix. I enrich the class of possible models by discarding the non-testable assumption mentioned above. The resultant class of models is, by construction, only *partially identified*; there is a continuum of possible parameter estimates that are consistent with the observed data. I argue that, from the Bayesian perspective, a user of model 1 essentially has an incorrect prior over the set of parameters of this richer third model. As a solution, I suggest using a prior that is carefully motivated from auxiliary information, so that it does not assign zero probability to positive-probability events.

In general, there is much prior information available about behavioral parameters, such as those governing preferences and technology. However, there is much less prior information about the parameters governing shock processes. One possible response to this problem is to be what I will term *agnostic*—that is, to be fully flexible about the specification of the prior concerning the shock-process parameters. I argue in the context of the example that if one takes such an agnostic approach, the data themselves reveal no information about the behavioral parameters. I interpret this result as an indirect argument for the procedure commonly called *calibration*, in which an investigator picks a plausible range for technology and preference parameters based only on prior auxiliary information.

In the final part of the paper, I return to Smets and Wouters (2003). Using the above analysis, I

offer a critique of their approach to estimation and model evaluation. I suggest how one might change their estimation and evaluation approach to ensure more reliable policy analyses.

The second part of the paper is about Bayesian estimation of models given limited a priori information, which is, as far as I know, novel. However, the first part is not new: It is well-known that there are potential problems with the principle of fit. In early contributions, Marschak (1950) and Hurwicz (1950) emphasize that multiple structures (mappings between interventions and outcomes) may be consistent with a given reduced form (probabilistic description of available data). Liu (1960) argues that this potential problem is, in fact, endemic: The available data never serve to identify the true structure uniquely. In perhaps the most related work, Sims (1980) argues explicitly that large-scale models may fit the data well and yet provide misleading answers to (some) questions because their estimates are based on incredible identification restrictions.

Though it lacks novelty, my discussion about the principle of fit serves three purposes. First, the principle remains a dominant one among policymakers and others (as my opening paragraphs indicate). Given the recent excitement about Smets and Wouters (2003) and other related papers, it is worthwhile (I believe) to remind everyone of the principle's limitations.

Second, I want to make absolutely clear that we cannot resolve the problem by using structural models. Most macroeconomists are highly cognizant of the Lucas critique (1976). It correctly emphasizes that to assess a policy intervention a model's parameters should be structural—that is, invariant to the intervention. In response, most macroeconomists now use structural models to analyze policy interventions. My paper demonstrates that this response is not a panacea. In particular, I show that even if it is *structural* and *well-fitting*, a model may provide misleading answers to policy questions.

Finally, my argument is not just that better fit *can* lead to worse answers but that we should expect that obtaining better fit *will* lead to worse answers. Archetypal macroeconomic models are usually under-shocked relative to the data under

consideration. Hence, to get macroeconomic models that fit well, we need to add shocks. But we generally know little about these shocks. It should not be at all surprising if adding them were to create new, possibly substantial, sources of error.

AN ARTIFICIAL WORLD AND POLICY INTERVENTIONS

The basic structure of this paper is akin to a Monte Carlo study. I first set up an artificial world over which I have complete control. I introduce an investigator (econometrician) into this artificial world who does not know the structure of the artificial world but is instead limited to using one of two possible (classes of) models. Both are false because the artificial world is not a special case of either model; however, the investigator does not know that they are false. Based on data from the artificial world, the investigator uses a variety of possible methods to determine which model has superior fit.

In this section, I describe the artificial world and a class of policy interventions under consideration in that world. In the artificial world, agents decide how much to work at each date. Their decisions are influenced by shocks to labor productivity, taxes, and preferences. The means of these random variables are hit by observable shocks in each period. Preference shocks and tax rates covary; it is this covariance that makes estimation of the parameters of the model challenging.

The Artificial World

Time is discrete and continues forever. There is a unit measure of agents who live forever, and preferences are given by

$$\sum_{t=1}^{\infty} \delta^{t-1} \left[\ln c_t - \exp(\psi_t) n_t^* / \gamma^* \right], 0 < \delta < 1,$$

where c_t is consumption in period t and n_t is labor in period t . Technology is given by

$$y_t = \exp(A_t) n_t,$$

where y_t is the amount of consumption produced

in period t . Agents are taxed at rate τ_t , where τ_t is governed by the policy rule

$$\begin{aligned} \ln(1 - \tau_t) &= -\beta \psi_t + \varepsilon_t \\ \beta &> 0. \end{aligned}$$

The proceeds of the taxes are handed back lump-sum to the agents.

The random variables $(A_t, \psi_t, \varepsilon_t)$ are i.i.d., over time. There is another random variable λ_t , which is equally likely to be 0 or 1. Conditional on $\lambda_t = i$, the random variables (A, ψ, ε) are all Gaussian and mutually independent, with means $(\mu_A(i), \mu_\psi(i), \mu_\varepsilon(i))_{i=0}^1$ and positive variances $(\sigma_A^2, \sigma_\psi^2, \sigma_\varepsilon^2)$. Note that the means depend on i , but the variances do not.

It is easy to prove that, in this economy, there is a unique equilibrium of the form

$$\begin{aligned} \ln n_t &= (\ln(1 - \tau_t) - \psi_t) / \gamma^* \\ \ln y_t &= A_t + \ln n_t. \end{aligned}$$

Interventions

Consider the following class of *interventions*, indexed by the real variable Δ . With intervention Δ , the tax rate follows the rule

$$\ln(1 - \tau_t(\Delta)) = \Delta + \ln(1 - \tau_t).$$

The *policy question* is this: How much does average logged output change in response to a change in the tax rate? Mathematically, let $y_t(\Delta)$ denote per capita output under intervention Δ . What is $E(\ln(y_t(\Delta^*))) - E(\ln(y_t))$, where Δ^* is a given intervention? The true answer to this question is Δ^* / γ^* .

TWO (IDENTIFIED) MODELS

There is an investigator who wants to know the answer to the given policy question. The investigator does not know the structure of the artificial world, but does observe the following data:

$$(\ln y_t, \ln n_t, \ln(1 - \tau_t), \lambda_t)_{t=1}^{\infty}.$$

The investigator has two possible models to use to answer the question. The basic economic ele-

ments of the models are the same as that of the artificial world itself. In each model, there is a unit measure of identical agents who work to produce output. The agents face a linear tax on output, and the proceeds of this tax are handed out lump-sum. However, the shock-generation processes in the two models are different from each other and from the artificial world.

Model 1

In model 1, preferences are of the form

$$\sum_{t=1}^{\infty} \delta^{t-1} [\ln c_t - \exp(\psi_{1t}) n_t^{\gamma_1} / \gamma_1];$$

technology is given by $y_t = \exp(A_{1t}) n_t$; and agents are taxed at rate τ_t , where $\ln(1 - \tau_t) = \varepsilon_{1t}$. The random variables $(A_{1t}, \psi_{1t}, \varepsilon_{1t})$ are i.i.d. over time and mutually independent. The random variable λ_t has support $\{0, 1\}$; the probability that λ_t equals 1 is given by p_1 . Conditional on $\lambda_t = i$, the random variables $(A_{1t}, \psi_{1t}, \varepsilon_{1t})$ are Gaussian, with means $(\mu_{1A}(i), \mu_{1\psi}(i), \mu_{1\varepsilon}(i))_{i=0}^1$ and variances $(\sigma_{1A}^2, \sigma_{1\psi}^2, \sigma_{1\varepsilon}^2)$. The investigator does not know these means and variances; they will have to be estimated in some fashion from the data. Put another way, this is actually a class of models indexed by the 11 parameters $(\gamma_1, (\mu_{1A}(i), \mu_{1\psi}(i), \mu_{1\varepsilon}(i))_{i=0}^1, \sigma_{1\varepsilon}^2, \sigma_{1A}^2, \sigma_{1\psi}^2, p_1)$.

Model 1 implies that in equilibrium

$$\begin{aligned} \ln(n_t) &= [\ln(1 - \tau_t) - \psi_{1t}] / \gamma_1 \\ \ln(y_t) &= A_{1t} + \ln(n_t). \end{aligned}$$

How does model 1 differ from the artificial world? It is alike in all respects except one: In model 1, the parameter β has been set to zero. As we shall see, this additional restriction allows the investigator to estimate γ_1 from the available data.

Model 2

In model 2, preferences are given by

$$\sum_{t=1}^{\infty} \delta^{t-1} [\ln c_t - \exp(\psi_2) n_t^{\gamma_2} / \gamma_2]$$

and technology is given by $y_t = \exp[A_2(1)\lambda_t + A_2(0)(1 - \lambda_t)] n_t$. Here, ψ_2 , $A_2(1)$, and $A_2(0)$ are all

constants; λ_t is a random variable; and agents are taxed at rate τ_t , where

$$\ln(1 - \tau_t) = \varepsilon_2(1)\lambda_t + \varepsilon_2(0)(1 - \lambda_t).$$

The random variable λ_t is i.i.d. over time, with support $\{0, 1\}$, and the probability that λ_t equals 1 is given by p_2 . The parameters $\varepsilon_2(1)$ and $\varepsilon_2(0)$ are both constants. Hence, in model 2, there are seven unknown parameters, $(\gamma_2, \psi_2, A_2(1), A_2(0), \varepsilon_2(1), \varepsilon_2(0), p_2)$. The model implies that

$$\begin{aligned} \ln(n_t) &= [\ln(1 - \tau_t) - \psi_2] / \gamma_2 \\ \ln(y_t) &= A_2(1)\lambda_t + A_2(0)(1 - \lambda_t) + \ln(n_t). \end{aligned}$$

How does model 2 differ from the artificial world? In model 2, tastes do not vary at all. As well, the variances of the other shocks around their means are both set to zero. Like many modern macroeconomic models, model 2 has relatively few sources of uncertainty compared with what is true of the (artificial) world.

THE FALLACY OF FIT

The investigator has two models available. He wants to use his infinitely long sample to decide which model to use in order to answer the policy question. The sample $(\ln(y_t), \ln(n_t), \ln(1 - \tau_t))_{t=1}^{\infty}$ is jointly Gaussian conditional on $\lambda_t = i$, for $i = 0, 1$. The means of the conditional distributions depend on λ_t ; the conditional distributions have the same variance-covariance matrix. Hence, the sample can be fully summarized by 13 moments: the probability p that λ_t equals 1, the means $(\mu_y(i), \mu_n(i), \mu_{\tau}(i))$ of $(\ln(y), \ln(n), \ln(1 - \tau))$ conditional on $\lambda_t = i$, and the variance-covariance matrix Σ of $(\ln(y), \ln(n), \ln(1 - \tau))$ conditional on λ .

Note that in these data, there are two distinct kinds of variation. The first kind is because of λ . Movements in λ generate changes in $(\mu_y(i), \mu_n(i), \mu_{\tau}(i))$; these changes can provide information about the unknown parameters of the two models. At the same time, $(\ln(y), \ln(n), \ln(1 - \tau))$ vary around these fluctuating means. This information is summarized by the six moments of Σ . The goal

of the investigator is to use these two sources of variation to estimate the unknown parameter γ .

Given this information, the investigator has available three methods of estimating/evaluating the models.

Method 1: Maximum Likelihood

In this subsection, I suppose that the investigator estimates the unknown parameters of each model by maximum likelihood, and then compares the models' abilities to fit the 13 population moments.

Model 2 implies that, conditional on $\lambda_t = i$, the data is deterministic. In other words, according to model 2, the conditional variance-covariance matrix of $(\ln(y_t), \ln(n_t), \ln(1-\tau_t))$ contains only zeros. It follows that the likelihood of the data, conditional on any specification of model 2, is zero.²

For model 1, the maximum-likelihood estimates of the 11 unknown parameters are given by

$$\begin{aligned}
 \hat{p}_1 &= 1/2 \\
 \hat{\gamma}_1 &= \Sigma_{\tau\tau} / \Sigma_{n\tau} \\
 \hat{\sigma}_{1A}^2 &= \Sigma_{yy} - \Sigma_{nn} \\
 \hat{\sigma}_{1\epsilon}^2 &= \Sigma_{\tau\tau} \\
 \hat{\sigma}_{1\psi}^2 &= (\hat{\gamma}_1)^2 \Sigma_{nn} - \Sigma_{\tau\tau} \\
 \hat{\mu}_{1A}(i) &= \mu_y(i) - \mu_n(i), \quad i = 0, 1 \\
 \hat{\mu}_{1\epsilon}(i) &= \mu_\tau(i) \\
 \hat{\mu}_{1\psi}(i) &= \mu_\tau(i) - \mu_n(i) \hat{\gamma}_1, \quad i = 0, 1.
 \end{aligned}
 \tag{1}$$

Given the infinitely long sample, these estimates are very precise; the likelihood of the sample is 1 given this parameter setting and zero given all others. Note that under this parameter setting, the model fits all 13 moments of the data exactly.

Hence, according to maximum likelihood, only model 1 should be used to answer the policy question (with the parameter estimates (1)); no specification of model 2 should be used. The lack of fit is because model 2 is “under shocked” rela-

tive to the data. The world has four distinct shocks generating the data, but model 2 has only one. Maximum likelihood punishes this kind of discrepancy severely; from a statistical point of view, it is the most readily detectable form of misspecification.

Method 2: Bayesian Estimation

In this subsection, I suppose that the investigator applies Bayesian estimation methods to the available data from the artificial world. Obviously, models 1 and 2 are nested—if model 2 is true, model 1 is also true. Consider an econometrician who has a prior over the 11 unknown parameters of model 1. The prior is such that it puts probability q on the parameters being consistent with model 2 and puts probability $(1 - q)$ on the parameters being inconsistent with model 2.

Now suppose the econometrician observes an infinite sequence of data from the artificial world. The data bleaches out the effect of the initial prior; the econometrician's posterior will be concentrated on the parameter estimates (1). A Bayesian econometrician with an infinitely long sample will reach the same policy conclusions as does a classical econometrician using maximum likelihood.

Method 3: Method of Moments

As we have seen, maximum likelihood and Bayesian estimation simply discard all under-shocked models. Now, we consider a less severe measure of fit: method of moments, by which I mean the following. Consider the 13 population moments that characterize the sample. Pick 13 positive weights that sum to 1. Estimate the unknown parameters in each model by minimizing the weighted sum of squared deviations between model-generated moments and sample moments. Then compare model 1 and model 2 by the value of the minimized objective.

Note again that the models are nested. Because we are minimizing the same objective for each model, model 1 must do at least as well as model 2. By setting the parameters in model 1 according to (1), the objective is set equal to zero. Because model 2 (incorrectly) generates a non-invertible variance-covariance matrix for any parameter

² It is worth noting that model 2 implies that the sample variance-covariance matrix, conditional on $\lambda_t = i$, is noninvertible in any finite sample. Hence, the likelihood of any finite sample, conditional on model 2, is zero.

Kocherlakota

setting, the objective must be strictly larger than 0. Model 1 must fit the data better than model 2, according to this measure of fit.

However, using method of moments, we can now actually estimate parameters for model 2, as opposed to simply discarding it as maximum likelihood does. For model 2, regardless of the weights, the estimated seven parameters are

$$\begin{aligned}\hat{p}_2 &= 1/2 \\ \hat{A}_2(i) &= [\mu_y(i) - \mu_n(i)], i = 0, 1 \\ \hat{\epsilon}_2(i) &= \mu_\tau(i), i = 0, 1 \\ \hat{\gamma}_2 &= \frac{\mu_\tau(1) - \mu_\tau(0)}{\mu_n(1) - \mu_n(0)} \\ \hat{\psi}_2 &= [\mu_\tau(1) - \hat{\gamma}_2 \mu_n(1)] = [\mu_\tau(0) - \hat{\gamma}_2 \mu_n(0)].\end{aligned}$$

The seven parameters are set so that the model generates the values in the data for the moments $(p, \mu_y(1), \mu_y(0), \mu_\tau(1), \mu_\tau(0), \mu_n(1), \mu_n(0))$. Model 2 predicts that the other moments are zero for any choice of parameters, so that part of the minimization problem is irrelevant for parameter estimation.

Using the Estimated Models to Answer the Policy Question

Recall that the policy question is this: What is the value of

$$E(\ln(y_t)(\Delta^*)) - E(\ln(y_t))$$

when taxes are changed so that $\ln(1 - \tau_t(\Delta^*)) - \ln(1 - \tau_t) = \Delta^*$? The true answer to this question is given by Δ^*/γ^* .

Here is what the two models deliver. Under model 1, the answer is $\Delta^*/\hat{\gamma}_1 = \Delta^* \Sigma_{n\tau} / \Sigma_{\tau\tau}$, where $\Sigma_{n\tau}$ is the population covariance of $\ln(n_t)$ and $\ln(1 - \tau_t)$ in the artificial world, and $\Sigma_{\tau\tau}$ is the population variance of $\ln(1 - \tau_t)$ in the true world. (This is Δ^* multiplied by the population regression coefficient.) We can calculate these population moments to find that that answer in model 1 is given by

$$ANS_1 = \Delta^* \left(1/\gamma^* + \beta \gamma^{*-1} \sigma_\psi^2 / \{ \beta^2 \sigma_\psi^2 + \sigma_\epsilon^2 \} \right),$$

which is too large in absolute value relative to the true answer of Δ^*/γ^* .

Under model 2, the answer is given by

$$\begin{aligned}ANS_2 &= \frac{[\mu_n(1) - \mu_n(0)]}{[\mu_\tau(1) - \mu_\tau(0)]} \Delta^* \\ &= \Delta^* \frac{[\mu_\epsilon(1) - \mu_\epsilon(0)]/\gamma^* - \beta[\mu_\psi(1) - \mu_\psi(0)]/\gamma^* - [\mu_\psi(1) - \mu_\psi(0)]/\gamma^*}{\mu_\epsilon(1) - \mu_\epsilon(0) - \beta[\mu_\psi(1) - \mu_\psi(0)]} \\ &= \Delta^* / \gamma^* - (\Delta^* / \gamma^*) \frac{[\mu_\psi(1) - \mu_\psi(0)]}{\mu_\epsilon(1) - \mu_\epsilon(0) - \beta[\mu_\psi(1) - \mu_\psi(0)]}.\end{aligned}$$

Note that if $(\mu_\psi(1) - \mu_\psi(0))$ is sufficiently close to zero in absolute value, ANS_2 is nearer to $1/\gamma^*$ than is ANS_1 . Even though model 2's fit is worse than that of model 1, model 2 may still deliver a superior answer to the policy question.

Why Doesn't Fit Work?

In the above discussion, we have seen that the model that fits better—indeed, the model that fits the available data perfectly—may well deliver a worse answer to the policy question. What is going on here? The policy question is this: What happens to hours worked if we increase the tax rate on labor? The answer is wholly governed by the elasticity of labor, which is equal to $1/\gamma^*$ in the artificial world. To answer the question the investigator has to estimate γ^* well, but there is a traditional difficulty associated with doing so. If there are no shifts in the labor supply, then the comovement in hours and tax rates will pin down the elasticity of the labor supply. However, if the labor supply shifts (that is, movements in ψ) are correlated with the variation in tax rates, then the investigator will achieve biased estimates of $1/\gamma^*$.

How do the two models estimate γ^* ? In the artificial world, there are two sorts of variation in the data. The first is that the means of the distributions of $(A, \psi, \ln(1 - \tau))$ fluctuate over time. The second is that the realizations of the random variables fluctuate around their means. The good news is that, because λ is observable, the two kinds of variations are distinct. The bad news is that both kinds of fluctuations feature potential comovement between ψ_t and τ_t —comovement that makes our task of estimating γ more difficult.

The two models differ in their estimates of γ because each one relies on a different type of

fluctuation to pin down γ . Model 1 assumes (incorrectly) that the fluctuations of $\ln(1-\tau_t)$ and ψ_t around their means are independent. It then exploits the fluctuations in tax rates and hours around their means to estimate γ . Model 2 assumes (incorrectly) that the mean of ψ_t does not fluctuate at all. It then uses the shifts in the means of hours and tax rates over time to estimate γ . Which one works better depends on which incorrect assumption is a better approximation to reality. Nothing in the data answers this question.

The key point is that the relative fit of the models does not tell us which of these assumptions is closer to being right. More generally, parameter estimation of any kind always relies on two sources of information: the data and nontestable identification assumptions. The fit of a model tells us nothing about the reliability of the latter.³ Yet their reliability is essential if one is to obtain accurate parameter estimates.

PRIOR CARE

The problem with model 1 is that it includes a false restriction, which is included solely to identify the unknown parameter. In this section, I consider a richer model than model 1, in which I dispense with the false identification restriction. By construction, this model is only partially identified. I argue that one way to interpret the problem with model 1 is that the investigator is using an incorrect prior over the larger parameter space of this richer model. I suggest a simple fix to these problems: estimate the larger, partially identified model in a Bayesian fashion while being meticulous in building the prior explicitly from auxiliary information.

As before, assume that there is an investigator who has an infinite sample $(\ln(y_t), \ln(n_t), \ln(1-\tau_t), \lambda_t)_{t=1}^\infty$ from the artificial world described

in the first section. The investigator does not use model 1 or model 2 though. Instead, the investigator uses a new model, model 3.

Model 3

In model 3, preferences are of the form

$$\sum_{t=1}^\infty \delta^{t-1} \left[\ln c_t - \exp(\psi_{3t}) n_t^{\gamma_3} / \gamma_3 \right], \gamma_3 \geq 0;$$

technology is given by $y_t = \exp(A_{3t}) n_t$; and agents are taxed at rate τ_t , where

$$\ln(1 - \tau_t) = -\beta_3 \psi_{3t} + \varepsilon_{3t}, \beta_3 \in R.$$

The random variables $(A_{3t}, \psi_{3t}, \varepsilon_{3t})$ are i.i.d. over time and mutually independent. The collection of random variables $\{\lambda_t\}_{t=1}^\infty$ are i.i.d., with support $\{0, 1\}$; the probability that λ_t equals 1 is given by p_3 . Conditional on $\lambda_t = i$, the random variables $(A_{3t}, \psi_{3t}, \varepsilon_{3t})$ are Gaussian, with means $(\mu_{3A}(i), \mu_{3\psi}(i), \mu_{3\varepsilon}(i))_{i=0}^1$ and variances $(\sigma_{3A}^2, \sigma_{3\psi}^2, \sigma_{3\varepsilon}^2)$. In this class of models, there are 12 unknown parameters $(\gamma_3, \beta_3, (\mu_{3A}(i), \mu_{3\psi}(i), \mu_{3\varepsilon}(i))_{i=0}^1, \sigma_{3A}^2, \sigma_{3\psi}^2, \sigma_{3\varepsilon}^2, p_3)$.

Model 3 implies that in equilibrium

$$\begin{aligned} \ln(n_t) &= \left[\ln(1 - \tau_t) - \psi_{3t} \right] / \gamma_1 \\ \ln(y_t) &= A_{3t} + \ln(n_t). \end{aligned}$$

Model 3 is exactly the same as model 1, except that, in model 3, tax rates may be correlated with the preference shock ψ_3 . This change means that model 3 is sufficiently rich—to nest both model 1 and the artificial world.

Model 3 is only partially identified. Suppose $\gamma_3 = \hat{\gamma}_3$. Then, there is a unique specification of the other 11 parameters so that model 3 fits the available data exactly. In particular, let

$$\begin{aligned} \hat{p}_3 &= 1/2 \\ \hat{\sigma}_{3\psi}^2 &= (0|\hat{\gamma}_3|-1)' \Sigma(0|\hat{\gamma}_3|-1) \\ \hat{\beta}_3 &= \frac{\hat{\gamma}_3 \Sigma_{n\tau} - \Sigma_{\tau\tau}}{\hat{\sigma}_{3\psi}^2} \\ \hat{\sigma}_{3\varepsilon}^2 &= \frac{(\hat{\gamma}_3)^2 \left[\Sigma_{nn} \Sigma_{\tau\tau} - \Sigma_{n\tau}^2 \right]}{\hat{\sigma}_{3\psi}^2} \\ \hat{\sigma}_{3A}^2 &= \Sigma_{yy} - \Sigma_{nn} \end{aligned} \tag{2}$$

³ Model 1 is a just-identified model; the number of identifying restrictions is equal to the number of estimated parameters. More generally, there may be more identifying restrictions than unknown parameters. It is commonplace to construct tests of the overidentifying restrictions in such models. However, it is important to keep in mind that these are tests only of some linear combinations of the restrictions. The other linear combinations are being used to estimate the parameters and are, as in the just-identified case, nontestable.

$$\begin{aligned} \hat{\mu}_{3A}(i) &= \mu_y(i) - \mu_n(i), \quad i = 0, 1 \\ \hat{\mu}_{3\psi}(i) &= \mu_\tau(i) - \mu_n(i)\hat{\gamma}_3 \\ \hat{\mu}_{3\varepsilon}(i) &= \mu_\tau(i) + \hat{\beta}_3\hat{\mu}_{3\psi}(i), \end{aligned}$$

and then the 13 moments generated by model 3 correspond to the moments of the sample. (Note that all parameter estimates that are supposed to be non-negative [that is, variances] are in fact non-negative.) Hence, for each specification of the parameter $\hat{\gamma}_3$, there exists a specification of the other 11 parameters so that the model exactly fits the data.

Recently, there has been a great deal of work on classical methods to estimate partially identified models (see Manski, forthcoming, for a useful survey). However, I believe it is most useful to consider the estimation of model 3 from a Bayesian perspective.⁴ Specifically, let $\theta = (\beta_3, (\mu_{3A}(i), \mu_{3\psi}(i), \mu_{3\varepsilon}(i))_{i=0}^1, \sigma_{3A}^2, \sigma_{3\psi}^2, \sigma_{3\varepsilon}^2, p_3)$ represent the parameters of the model other than γ_3 . Suppose that the parameter space for (γ_3, θ_3) is given by $R_+ \times \Theta$, where $\Theta = R^7 \times R_+^3 \times [0, 1]$. This parameter space is a 12-dimensional manifold. I assume that the investigator has a prior density over this manifold such that γ_3 is stochastically independent from θ . I will let the marginal prior density over γ_3 be denoted by f and the prior density over θ be denoted by g .

A basic intuition in Bayesian estimation/learning is that the prior is essentially irrelevant if one has a large amount of data. Intuitively, the impact of the data is sufficiently large to bleach out the initial information in the prior. However, this intuition applies only when the model is identified. As we shall see, when one uses the partially identified model 3, the prior over (γ_3, θ) affects the posterior distribution over γ_3 , even though the investigator has access to an infinite sample.

⁴ Lubik and Schorfheide (2004) use a Bayesian procedure to estimate a partially identified model. As Schorfheide (forthcoming) emphasizes, identification problems—that is, the presence of ridges or multiple peaks in the likelihood—do not create any problems for Bayesian estimation: “Regardless, the posterior provides a coherent summary of pre-sample and sample information and can be used for inference and decision making.”

A Mistaken Prior: The Case of Model 1

Suppose g is such that the prior puts probability 1 on the event $\beta_3 = 0$. In this case, after seeing the available data, the investigator’s posterior is concentrated on the vector (1). With this kind of prior, model 3 is equivalent to model 1.

We have seen that using model 1 gives misleading answers to the policy question. A prior like g implicitly contains a great deal of information, because no amount of data can shift the investigator from his belief that $\beta_3 = 0$. It should not be used unless the investigator actually has this information about the world.

An Arbitrary Prior

Suppose instead that g and f are such that the support of the investigator’s prior is the entire parameter space. Let $h(\hat{\gamma}_3; Data)$ represent the parameter estimates (2) when $\gamma_3 = \hat{\gamma}_3$. Then, after seeing the infinite sample, the investigator’s posterior is concentrated on a one-dimensional manifold $(\hat{\gamma}_3, h(\hat{\gamma}_3; Data))$ indexed by $\hat{\gamma}_3 \in [\gamma_L, \gamma_H]$. His posterior over this one-dimensional manifold is proportional to $\phi(\hat{\gamma}_3)$, where

$$\phi(\hat{\gamma}_3) = f(\hat{\gamma}_3)g(h(\hat{\gamma}_3; Data)) \prod_{n=1}^{11} \frac{\partial h_n(\hat{\gamma}_3; Data)}{\partial \gamma_3}.$$

(Here, h_n represents the n th component of the function h .) Given this posterior uncertainty, the investigator’s answer to the policy question is no longer a single number. Instead, the investigator’s answer is now a random variable, with support equal to the interval $[\Delta^*/\gamma_H, \Delta^*/\gamma_L]$ and density proportional to

$$\phi(1/x) / x^2,$$

where x represents the answer to the policy question.

Because the model is only partially identified, the investigator’s posterior over the answer to the policy question is influenced by his marginal prior f over the preference parameter γ_3 and his prior g over the other parameters. This dependence exists even though he sees an infinite sample.⁵

⁵ Note that in (2), the estimates $(\hat{\sigma}_{3A}^2, \hat{\mu}_{3A}^0, \hat{\mu}_{3A}^1, \hat{\beta}_3)$ depend only on the

This means that the investigator cannot count on the available data to eventually correct all misinformation in his initial prior. Instead, he must be sure that his prior truly represents information about these parameters derived from auxiliary sources.

No Prior Information About Shock Processes

It is easy to see how to construct a prior f over the preference parameters (or over technology parameters in a more general context). We can derive information about such behavioral parameters from other data sources, from introspection, or from experiments. However, it is more difficult to obtain this information about the joint shock process (A, τ, ψ) . In at least some, and perhaps most, cases, there will be no auxiliary information available about these processes. What should be done?⁶

In this subsection, I assume that the investigator has information that leads to a prior f with support $[\gamma_L, \gamma_H]$, where $\gamma_L > 0$. The investigator has no auxiliary information about the shock processes.

The Bayesian Approach. One possible response to this no-information situation is to formulate a purely subjective prior belief over the 11 parameters of the shock process and then proceed in a standard Bayesian fashion. In doing so, it is important to keep two issues in mind. First, as we have seen above, when the model is partially identified, the prior impacts the answer to the policy question regardless of how large the sample is. The subjective beliefs always matter.

Second, every prior—regardless of how neutral it may seem—has some information embedded in it. To appreciate this last point, suppose there is a parameter α . All that an investigator

truly knows about α is that α lies in $[0, 1]$; he wants his prior over α to be neutral over its location within that interval. It is tempting to conclude that we can capture this neutrality by using a uniform distribution over $[0, 1]$. But now consider $y = \alpha^{1000}$. What does the investigator know about y ? Presumably, all that the investigator knows about y is that it lies in $[0, 1]$ —if he knew more, he would have known more about α . However, if the investigator has a uniform prior over α , then the investigator's prior over y is proportional to $y^{-999/1000}$. This density is far from uniform over $[0, 1]$; it places a lot more weight on low values of y than on high values of y . The uniform density over α actually does smuggle information about α into the analysis.

An Agnostic Approach. The Bayesian approach weds the investigator to a single prior g . As I suggest above, this prior contains information that the investigator does not literally have. One response to this problem is to use what I would call an *agnostic* approach: Be flexible about the choice of g and compute a posterior density for each possible prior g over γ_3 . By doing so, the investigator's answer to the policy question is no longer a single number, or even a single posterior, but rather a collection of posteriors generated by varying g . All of these posteriors have support $[\Delta^*/\gamma_H, \Delta^*/\gamma_L]$.

The resulting collection of posteriors is large. In particular, let p be any continuous probability density function over $[\Delta^*/\gamma_H, \Delta^*/\gamma_L]$. Let g_p be a continuous function mapping Θ into R_+ such that

$$g_p(h(1/x; Data)) = \frac{x^2 p(\Delta^{*-1}x)}{f(1/x) \prod_{n=1}^N h_n(1/x; Data)}$$

for all x in $[1/\gamma_H, 1/\gamma_L]$. (This pins down the behavior of g_p only on a given line in Θ .) If the investigator's prior over Θ is given by g_p , then his posterior over $[\Delta^*/\gamma_H, \Delta^*/\gamma_L]$ is given by p . Thus, the agnostic approach imposes no discipline on the question of interest beyond the upper or lower bounds on γ_3 imposed by the prior f .

This kind of agnostic analysis is reminiscent of calibration. Under calibration, an investigator uses information from auxiliary sources to pin

data and not on $\hat{\gamma}_3$. Hence, the posterior over these four parameters is concentrated on a single vector after the investigator sees an infinite sample.

⁶ Recently, del Negro and Schorfheide (2006) have suggested using prior beliefs about endogenous variables (such as output and inflation) as a way to construct legitimate priors about exogenous shocks. This approach is potentially interesting. One concern is that usually our prior beliefs about endogenous variables come from the macroeconomic data that will, in fact, be used for estimation.

down a range of possible values for behavioral parameters. He then reports answers to the policy question for all of the parameter settings in this range. It is exactly this information that the investigator ends up reporting under the agnostic approach: a range of possible values for the policy question given the range of possible values for the behavioral parameters.

It is important to emphasize that this conclusion does not mean that the data is useless under the agnostic approach. Estimation collapses the support of the original joint prior from a 12-dimensional manifold to the 1-dimensional support of the posterior. Hence, the prior information about γ_3 , combined with the data, does help the investigator learn a great deal about the nature of the shocks hitting the economy. It is true that this information is irrelevant given the policy question posed. For other potential questions, though, this information may well be valuable.

The Agnostic Approach and Decisionmaking.

Of course, more generally, the investigator may have some information about the underlying shocks that restricts the possible specifications of g . Then, the agnostic approach is not equivalent to calibration. In this general case, the agnostic approach implies that each policy intervention (each Δ) leads to a set of posterior probability distributions over outcomes.

It is interesting to consider the problem of choosing Δ in this setting. Suppose there is a social welfare $W(p)$ associated with a given posterior p over the set of outcomes. Let $\Pi(\Delta)$ represent the set of possible posteriors implied by a given Δ . Then, choosing Δ is akin to optimizing under Knightian uncertainty, as opposed to risk. It is standard in such settings to use a maximin approach, under which the choice of Δ solves the problem:

$$\max_{\Delta} \min_{p \in \Pi(\Delta)} W(p).$$

Hurwicz (1950, p. 257) provides a similar resolution to the problem of decisionmaking with partially identified models.⁷

⁷ See Gilboa and Schmeidler (1989) for an axiomatization of this approach to uncertainty.

RELATIONSHIP TO SMETS AND WOUTERS (2003)

As reported in the introduction, Smets and Wouters (2003) estimate a DSGE monetary model. They note that their model is highly similar to that of Christiano, Eichenbaum, and Evans (2005). The big difference between the two specifications is in the number of shocks: Smets and Wouters allow for 10 different shock processes. None of these shock processes represent measurement error. Instead, they all play a substantive economic role.

Smets and Wouters use a Bayesian procedure to estimate their model. As argued above, the prior plays an important role in this kind of estimation. Smets and Wouters correctly spend a great deal of time in their paper discussing the specification of the prior over the preference and technology parameters. They motivate this part of the prior thoroughly using explicit auxiliary information.

The motivation for their choice of prior over the 10 shock processes is quite different. They write (p. 1140), "Identification is achieved by assuming that each of the structural shocks [is] uncorrelated and that four of the ten shocks... follow a white noise process." In other words, they choose the prior over the 10 shock processes in order to achieve identification, not because of auxiliary information. The first example makes clear the problems with this approach. Like Smets and Wouters, the user of model 1 chooses the prior over the shock processes to achieve identification. Because this prior does not truly reflect auxiliary information, the resulting estimates are severely biased, even though the model fits the data exactly. Smets and Wouters give us no reason to believe why the same should not be true of the estimates of their model.⁸

The second part of the current paper suggests an alternative approach. The investigator should not pick a prior that is designed to achieve iden-

⁸ In their recent discussion of identification of DSGE models, Canova and Sala (2006, p. 40) write that "resisting the temptation to arbitrarily induce identifiability is the only way to make DSGE models verifiable and knowledge about them accumulate on solid ground." I agree.

tification. Instead, the prior—or collection of priors—over the shock processes should reflect the investigator’s actual beliefs about those processes. The resulting set of posteriors will naturally contain less information—but also be more reliable. The key property of model 3 is that it is sufficiently rich to include as a special case the artificial world that is actually generating the data. It is certainly difficult to build such a class of models in the real world. Nonetheless, Bayesian estimation techniques (or any other for that matter) are only reliable if one does so.⁹

CONCLUSIONS

A model-based analysis of a policy intervention has two steps. The first is to figure out the key parameters that shape the quantitative impact of the intervention. The second is to gather information about these parameters. This information can come in two forms: prior information and information derived from estimating the model using a particular data set. The first part of this paper argues that the fit of a model tells us little about the quality of information coming from either source. The second part of the paper argues that the latter source of information (estimation) is not useful unless the investigator has reliable prior information about shock processes.

There is an important lesson for the analysis of monetary policy. Simply adding shocks to models in order to make them fit the data better should not improve our confidence in those models’ predictions for the impact of policy changes. Instead, we need to find ways to improve our information about the models’ key parameters (for example, the costs and the frequency of price adjustments). It is possible that this improved information may come from estimation of model parameters using macroeconomic data. However, as we have seen, this kind of estimation is only useful if we have reliable a priori evidence about the shock processes. My own belief is that this kind of a priori evidence is unlikely to be avail-

able. Then, auxiliary data sources, such as the microeconomic evidence set forth by Bils and Klenow (2004), will serve as our best source of reliable information about the key parameters in monetary models.

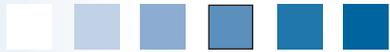
REFERENCES

- Bils, Mark and Klenow, Peter J. “Some Evidence on the Importance of Sticky Prices.” *Journal of Political Economy*, October 2004, 112(5), pp. 947-85.
- Canova, Fabio and Sala, Luca. “Back to Square One: Identification Issues in DSGE Models.” Working Paper 583, European Central Bank, 2006.
- Christiano, Lawrence J; Eichenbaum, Martin and Evans, Charles. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy*, February 2005, 113(1), pp. 1-45.
- del Negro, Marco and Schorfheide, Frank. “Forming Priors for DSGE Models and How It Affects the Assessment of Nominal Rigidities.” Working paper, University of Pennsylvania, 2006.
- Gilboa, Itzhak and Schmeidler, David. “Maxmin Expected Utility with a Non-unique Prior.” *Journal of Mathematical Economics*, 1989, 18(2), pp. 141-53.
- Hurwicz, Leonid. “Generalization of the Concept of Identification,” in Tjalling Koopmans, ed., *Statistical Inference in Dynamic Economic Models*. New York: John Wiley and Sons, 1950.
- Liu, Ta-Chung. “Underidentification, Structural Estimation, and Forecasting.” *Econometrica*, October 1960, 28(4), pp. 855-65.
- Lubik, Thomas A. and Schorfheide, Frank. “Testing for Indeterminacy: An Application to U.S. Monetary Policy.” *American Economic Review*, March 2004, 94(1), pp. 190-217.
- Lucas, Robert E. “Econometric Policy Evaluation: A Critique,” in K. Brunner and A. Meltzer, eds., *The Phillips Curve and Labor Markets*. Amsterdam: North-Holland, 1976.

⁹ See Schorfheide (2000) for a discussion of how to augment Bayesian techniques to allow for the possibility that no model under consideration is true.

Kocherlakota

- Manski, Charles F. "Partial Identification in Econometrics," in Stephen Durlauf and Lawrence Blume, eds., *The New Palgrave Dictionary of Economics*. Second Edition (forthcoming).
- Marschak, Jacob. "Statistical Inference in Economics," in T. Koopmans, ed., *Statistical Inference in Dynamic Economic Models*. New York: John Wiley and Sons, 1950.
- Schorfheide, Frank. "Loss Function Based Evaluation of DSGE Models." *Journal of Applied Econometrics*, November-December 2000, 15(6), pp. 645-70.
- Schorfheide, Frank. "Bayesian Methods in Macroeconometrics," in Stephen Durlauf and Lawrence Blume eds., *The New Palgrave Dictionary of Economics*. Second Edition (forthcoming).
- Sims, Christopher A. "Macroeconomics and Reality." *Econometrica*, January 1980, 48(1), pp. 1-48.
- Smets, Frank and Wouters, Raf. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, September 2003, 1(5), pp. 1123-75.



Commentary

Lee E. Ohanian

In “Model Fit and Model Selection,” Narayana Kocherlakota (2007) warns econometricians and the users of econometric analyses that macroeconomic models that fit the data well—as measured by a high R^2 and/or low residual variance—may not be very useful for policy advice because key parameters may not be identified. As an alternative, Kocherlakota provides a Bayesian approach that recognizes the significant challenge of identifying all parameters in a fully specified general equilibrium model and that also treats uncertainty about parameter values in a consistent fashion.

Kocherlakota’s agnostic approach means that the range of uncertainty associated with the conditional forecasts (policy advice) generated by the macroeconomic models used by central banks and other policymaking agencies is probably much larger than recognized by macroeconomic practitioners. This also suggests that policymakers, who use the forecasts from these models as an input into policymaking, should also modify their priors and recognize the considerable uncertainty in conditional forecasts.

There is substantial evidence that supports Kocherlakota’s recommendation, and there is also substantial evidence that this practice—or other practices that explicitly recognize the degree of uncertainty in modeling the economy—is not followed by macroeconomic model builders. Nor is agnosticism followed by policymakers regarding the structure of the economy. Instead, current model-building practice focuses largely

on models that feature a Phillips curve, and recent monetary policy decisions also appear to focus on the Phillips curve. I first discuss Kocherlakota’s analysis of fit and identification and then discuss the broader issues of agnosticism in choosing among alternative theoretical frameworks for evaluating policy and the role of agnosticism in making monetary policy.

AGNOSTICISM IN MODELING

Kocherlakota reminds macroeconomic modelers and the users of these models that a model with a good fit may not be a useful tool for conditional forecasting and policy analysis. To summarize practitioners’ current view about fit and its implications for policy analysis, Kocherlakota quotes from recent influential work by Smets and Wouters (2003), who suggest that a useful model for policy analysis is one that includes enough shocks to fit the data well. Smets and Wouters’s view is quite representative of macroeconomic modeling strategies used today.

To illustrate why a model that fits well may be not be useful for conditional forecasting, Kocherlakota constructs a model economy in which a model fits the data perfectly. He then shows that despite the perfect fit, the model cannot accurately forecast the impact of a tax cut on the economy because the elasticity of labor supply is unidentified. The reason that the perfect-fitting model provides a poor conditional forecast is

Lee E. Ohanian is a professor of economics at the University of California at Los Angeles, a visiting scholar at the Federal Reserve Bank of Minneapolis, and a research associate of the National Bureau of Economic Research.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 361-69.

© 2007, The Federal Reserve Bank of St. Louis. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

because it includes a non-testable identifying restriction that is false.

Kocherlakota suggests an alternative procedure that completely discards the non-testable identifying restriction. As an alternative, Kocherlakota recommends using auxiliary information to construct a range of values that a parameter can take. This is not standard Bayesian analysis, however, in that the standard approach commits the investigator to a single prior. Instead, Kocherlakota's procedure considers many priors over the parameter of interest. This delivers a collection of posteriors, which are restricted only in that they have support between the minimum and maximum values specified for the parameter. In principle, this approach is very sensible, as it explicitly and systematically allows the researcher to conduct a sensitivity analysis.

From a practical perspective, however, this procedure may be difficult to apply. To see this, note that in Kocherlakota's example of the impact of a tax cut, there is a single parameter. In this case, the application of the agnostic procedure yields a collection of posteriors over this single parameter, with support over a minimum and maximum value. This one-dimensional case is fairly simple to implement and to investigate. However, in a high-dimensional setting, the investigator must specify *many* priors over *several* parameters. Specifying multidimensional priors can be difficult; in practice, specifications are often chosen for computational ease, but in this case the prior is particularly important because the effect of the prior does not wash out as the sample size becomes arbitrarily large, as in standard analysis with full identification. Understanding how various multidimensional priors affect the analysis is very much an open, and difficult, question. Moreover, understanding how to distill and interpret the information from a collection of posteriors is also an open question. Making progress on these fronts seems necessary to successfully apply the agnostic, Bayesian approach in any rich model that includes many parameters and/or shocks.

Kocherlakota's agnostic approach presumes that there is an inherent identification problem in macroeconomics that is not easy to resolve. Is

the identification problem in macroeconomics as difficult as suggested by Kocherlakota? Unfortunately, there is no easy answer to this question; the profession has wrestled with identification in aggregate economics since the work of Tinbergen (1937), Haavelmo (1944), the Cowles Commission (Koopmans, 1950), Liu (1960), Sims (1980), and it continues today (Canova and Sala, 2006).

The identification debate in macroeconomics has suggested many different resolutions to the problem. One must understand the differences: Sims (1980) viewed the identification challenge in macroeconomics a sufficiently tall order to fill as to recommend relatively unrestricted vector autoregression (VAR) models that achieved identification by imposing a sufficient number of lags in the VAR to generate white noise innovations and then impose a Wold causal ordering on the innovation covariance matrix. In contrast, Kocherlakota recommends a very different approach, in which the behavioral equations of the model are tightly restricted by theory, but only minimal restrictions are imposed on the structure of the shock processes. Regarding the relative merits of these two different approaches, identification achieved through restrictions on shock processes are often difficult to justify because economic theory typically does not shed much light on the correct stochastic specification of shocks. Moreover, evaluating the identification of shocks is difficult, as identifying shocks almost always requires strong non-testable restrictions.

Some economists argue that shocks should be uncorrelated, and that this apparently innocuous assumption can go a long ways toward achieving identification. But we have several observations that shocks can be correlated. For example, there were several scientific, productivity breakthroughs in World War II that were largely the consequence of the large wartime government spending shock. Similarly, World War II monetary shocks were due to fiscal spending requirements that induce the need for seignorage finance. The deregulation of financial markets over the past 30 years has led to significant technological change in financial intermediation. The Great Depression led to enormous changes in economic regulation and government management of the economy. These

are just a few examples that indicate that achieving identification through orthogonality assumptions can be at variance with the data.

Identification will always be a difficult issue in macroeconomics, as maintained identifying restrictions are by definition not testable and are almost always open to debate. But what researchers can do is not make the identification problem any more difficult than it needs to be, and here Kocherlakota's recommendations are particularly valuable. The focus on fit, as exemplified by Smets and Wouters, tends to increase the difficulty of the identification problem. This is because increasing the richness of the model—by including more shocks—makes identification harder by requiring more restrictions to be placed on the shock process. From this perspective, relatively simple models may be easier to identify than densely parameterized models.

It is puzzling that the profession needs to be reminded of the “fallacy of fit” (Kocherlakota's words). The profession learned this dictum the hard way in the 1970s, when the apparently well-fitting large-scale macroeconomic models broke down, particularly the Phillips curve (the inflation-unemployment relationship), which was a central component of these models. Specifically, the 1970s witnessed both unemployment and inflation rising to levels far outside their fitted historical relationship. At the same time Charles Nelson (1972) showed that atheoretic, low-order univariate ARMA models of macroeconomic time series—that typically were characterized by a worse fit than the large-scale models—dominated the large-scale models in forecasting competitions. Further improvements in forecasting were generated by pseudo-Bayesian VARs, which imposed random-walk priors on time series to reduce the problem of overparameterization that is inherent in VARS. Bayesian VARs are used for forecasting at several research agencies and commercial banks, including the Minneapolis Fed and the Richmond Fed. All of these events led to traditional large-scale econometric models playing a much smaller role in central bank research and in policymaking.

So why did central bank researchers and policymakers return to macroeconomic models

after these failures? One reason stems from the fact that current models feature a much deeper structure than their large-scale predecessors. Today's models include dynamically maximizing households, maximizing firms, an internally consistent set of expectations, and a precise definition of equilibrium. All of these advances were absent from earlier models, and it is believed by many that these improvements would allow macroeconomic models to successfully confront the Lucas critique. Nevertheless, it is critical to distinguish between a model with a deep structure that in principle can be used for conditional forecasting and a model in which the parameters are reasonably identified. The first feature doesn't imply the second; ironically, specifying rich, fully articulated general equilibrium models will tend to make the identification problem even more difficult.

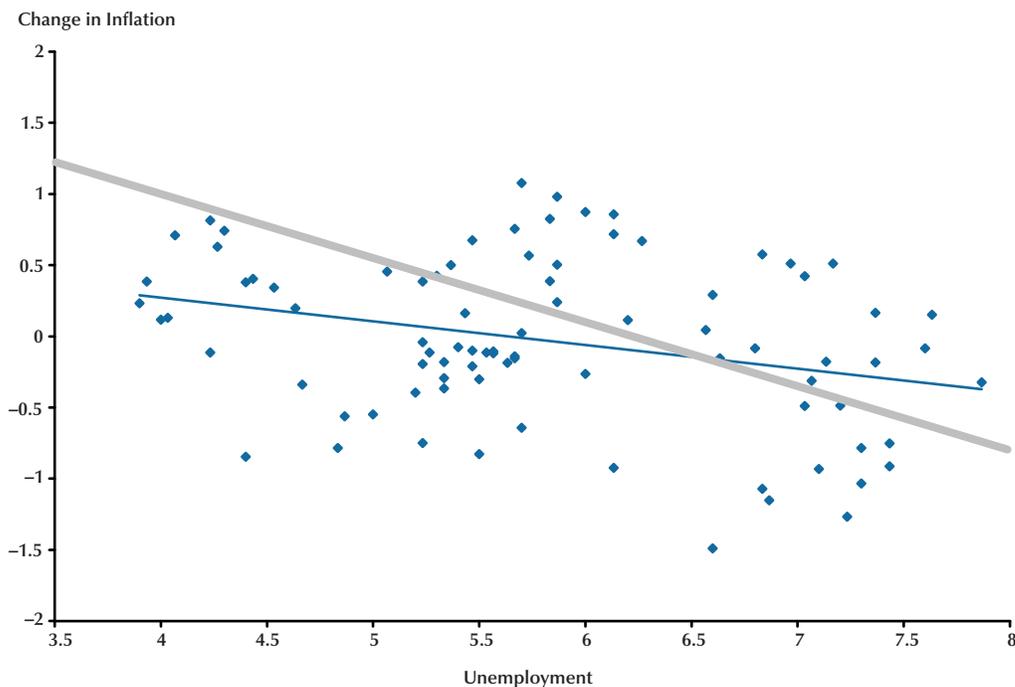
Applied economists face a difficult trade-off in specifying macroeconomic models. Simple models may in principle be easier to identify, conditional on correct specification, but simple models will tend to be false, and thus may not be useful for parameter estimation, at least from a classical perspective. Richer models may be more difficult to identify but, conditional on identification, may fare better in terms of parameter estimation. This trade-off is one reason why *calibration*, which sidesteps these difficult issues, has been so popular among applied macroeconomists. Kocherlakota's approach is another proposal in a research program that has attempted to place calibration into either an explicit classical or Bayesian framework (see Watson, 1993; Diebold, Ohanian, and Berkowitz, 1998; Schorfheide, 2000; and Fernández-Villaverde and Rubio-Ramírez, 2004).

THE PHILLIPS CURVE PRIOR IN MACROECONOMETRIC MODELING

Central bank research staff have strong priors on the class of models that are used in monetary policy analysis. The dominant class of models are those with a *Phillips curve*. Here, I define the

Figure 1

Is the NAIRU Phillips Curve Still There? Unemployment vs. Future Changes in Inflation, 1984-2006



Phillips curve as the view that during periods of slack economic capacity—such as a period of relatively high unemployment—rapid economic growth will not raise inflation very much; but during periods of low unemployment, economic growth can raise inflation considerably. The Phillips curve view has implications for monetary policy. Specifically, it implies that monetary stimulus during periods of high unemployment will not raise inflation very much and that there is scope for the Fed to moderate recessions (which are periods of slack capacity) through expansionary monetary policy. It also implies that, as capacity becomes tight, the Fed controls inflation by attempting to reduce the growth rate of the real economy through monetary contraction.

It may be reasonable for model builders and policymakers to narrowly focus on this class of models if there is strong empirical support for the Phillips curve. In contrast, if there is limited

support for this class of models, then there is scope to consider alternative theoretical channels for the determination of inflation. Here, I present U.S. time-series evidence that shows little support for the view that inflationary risks are significantly higher during periods of rapid growth and tight capacity, relative to rapid growth and slack capacity.

Atkeson and Ohanian (2001), Stock and Watson (2006), and others have recently analyzed the Phillips curve in U.S. time series. Under the Phillips curve view, low unemployment should be associated with rising inflation. This is often referred to as the NAIRU (non-accelerating inflation rate of unemployment) Phillips curve. Figure 1 shows the NAIRU Phillips curve by presenting the change in inflation and the unemployment rate for the period 1960-2006. The figure updates my earlier study with Atkeson to include data from 2000-06. The heavy gray line is an ordi-

nary least squares (OLS) regression line, which shows a modest negative relationship between the change in inflation and unemployment for 1960-83. The blue line is the OLS regression line between these variables for 1984-2006. The slope coefficient is very close to zero and is also statistically insignificantly different from zero. This latter result indicates that there has been no systematic relationship between the change in the inflation rate and unemployment since 1984. In other words, there has been no simple NAIRU Phillips curve in U.S. data for more than 20 years.

In Atkeson and Ohanian (2001), we extended our analysis of the Phillips curve by examining whether changes in unemployment, or other measures of slack capacity, help forecast future inflation relative to a naive forecasting model that simply extrapolates the current inflation rate into the future. Surprisingly, we found that inflation was not forecasted well by measures of slack capacity. In particular, the root mean-squared forecast error (RMSE) for the core consumer price index, which is a measure of inflation that excludes volatile food and energy prices and a key indicator of inflation for both financial markets and central banks, is as much as 94 percent higher compared with the forecast from the naive model that extrapolates the current inflation rate into the future. More sophisticated forecasting models did not fare much better: for example, inflation forecasts from Stock and Watson's macroeconomic activity index model, which forecasts inflation from a much larger information set than just unemployment. This model had an RMSE of between 33 and 81 percent higher than the naive forecast. These results indicate that there is no significant, predictable relationship between cyclical fluctuations in the real economy and future inflation. Paradoxically, forecasts from sophisticated models, which clearly fit much better in sample, are deficient to those from very simple models that do not fit so well in sample, such as the naive model we used in Atkeson and Ohanian (2001).

Tables 1 through 3 show the results of other tests of the Phillips curve. These tests evaluate whether economic growth generates inflation differentially when unemployment is low (tight

Table 1

Testing the Phillips Curve Correlation Between GDP Growth and Inflation, Controlling for Unemployment

Period	Deflator	Core CPI
1957:Q1–2006:Q2	–0.23	–0.20
1960s	–0.20	–0.23
1970s	–0.43	–0.40
1980s	–0.28	–0.11
1990s	–0.45	–0.44
2000:Q1–2006:Q2	–0.05	–0.22

capacity) relative to when unemployment is high (slack capacity). Table 1 shows the correlation between quarterly gross domestic product (GDP) growth and inflation, measured by the core CPI and by the GDP deflator for the period 1957-2006, and also for each decade during the past 50 years. The data are conditioned on the unemployment rate as a measure of capacity. The most striking finding is that the relationship between growth and inflation is negative, not positive as suggested by the Phillips curve. Tables 2 and 3 show correlations between GDP growth and inflation, conditioned on unemployment, for various leads and lags up to 1 year (4 quarters). The correlations are primarily negative or close to zero.

These tests indicate that there is not sufficient evidence to support the strong prior for the Phillips curve that characterizes current econometric practice in central banks. Instead, these results suggest that models with alternative inflation mechanisms should be analyzed.

THE PHILLIPS CURVE PRIOR AND MONETARY POLICY

The strong emphasis of the Phillips curve in macroeconomic research at central banks suggests that this view is also prominent among the policy-makers who are the primary users of central bank economic research. The importance of the Phillips curve in policymaking appears to vary over time. There does not appear to be a substantial Phillips

Table 2**Testing the Phillips Curve Correlation Between GDP Growth and Inflation (Core CPI): Leads and Lags, Controlling for Unemployment**

1957:Q1–2006:Q2	<i>i</i> = 4	<i>i</i> = 3	<i>i</i> = 2	<i>i</i> = 1
$\Delta \ln (y_{t+i}), \Delta \ln (p_t)$	-0.23	-0.45	-0.41	-0.28
$\Delta \ln (y_t), \Delta \ln (p_{t+i})$	-0.02	0.04	-0.10	-0.12

Table 3**Testing the Phillips Curve Correlation Between GDP Growth and Inflation (Deflator): Leads and Lags, Controlling for Unemployment**

1957:Q1–2006:Q2	<i>i</i> = 4	<i>i</i> = 3	<i>i</i> = 2	<i>i</i> = 1
$\Delta \ln (y_{t+i}), \Delta \ln (p_t)$	-0.27	-0.29	-0.26	-0.22
$\Delta \ln (y_t), \Delta \ln (p_{t+i})$	0.04	-0.04	-0.09	-0.16

curve policy bias during the Volcker-Greenspan disinflation of 1982-95. This period is certainly one of the great triumphs of central banking. After engineering the largest peacetime inflation in the history of the United States, with inflation rising from about 1 percent in the early 1960s to more than 13 percent by 1980, the Fed lost credibility with financial markets. By 1980, long-term interest rates rose to 13 percent. A standard Fisher equation decomposition, which relates nominal interest rates to expected inflation over the horizon of the security, clearly indicates that financial markets were systematically expecting permanent high inflation. Beginning at this time, however, Paul Volcker initiated a low-inflation monetary policy, and inflation declined to less than 3 percent by the mid-1990s.

To analyze the potential impact of the Phillips curve on monetary policy, I examine the relationship between the federal funds rate and the unemployment rate. If there is a strong influence of the Phillips curve on policy, then we should observe a systematic inverse relationship between the funds rate and the unemployment rate. Figure 2 shows monthly data on these variables between 1981 and 1995. Note that there is little systematic relationship between the federal funds rate and the unemployment rate (the correlation is about 0.3), suggesting that policy was not par-

ticularly focused on the Phillips curve. This is not surprising, as there is little disagreement among economists or financial market participants that monetary policy during this period was unconditionally committed to reducing inflation, without much reference to the business cycle. The policy was indeed effective, as inflation fell and long-term interest rates fell.

But the nature of policy seemed to change considerably after inflation declined. Figure 3 shows the funds rate and the unemployment rate between 1996 and 2006. This figure shows a distinct and systematic inverse relationship between unemployment and the funds rate, with a correlation of -0.91 . As the unemployment rate declined to 4 percent in 1999 and 2000, Fed officials worried about tight labor markets and inflationary pressures and raised the funds rate. Then, as the unemployment rate rose from 4 percent to more than 6 percent, the Fed pursued a more expansionary policy, driving the funds rate down from 6.5 percent in late 2000 to just 1 percent in late 2003. The policy record since 1995 is consistent with a strong Phillips curve prior and is reminiscent of the “fine tuning” that policymakers pursued in the 1960s and 1970s. U.S. time-series data provides little support that a fine-tuning policy based on Phillips curves will be successful.

Figure 2

Monetary Policy Without a Phillips Curve Focus

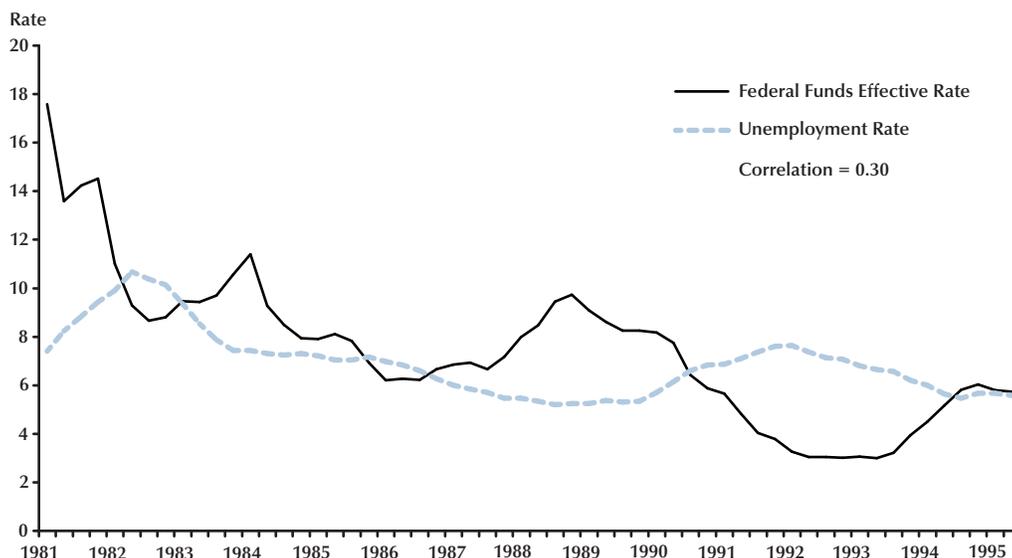
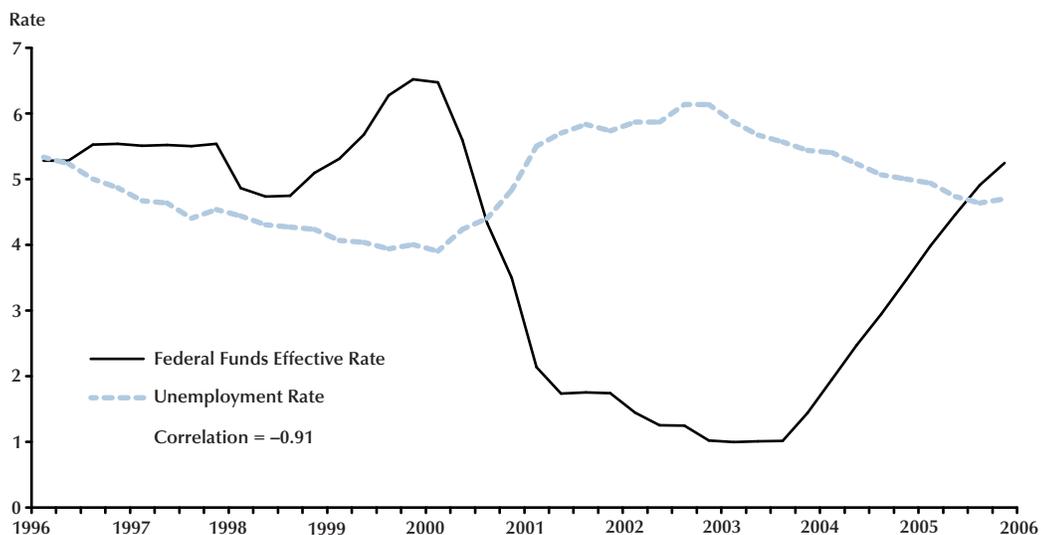


Figure 3

Unemployment Rate–Driven Monetary Policy? The Return of Phillips Curve Policymaking



CONCLUSION

During the 1960s, policymakers believed they understood the U.S. economy so well that they could achieve virtually any desired result with the appropriate mix of fiscal and monetary stimulus or contraction. Part of this belief stemmed from the close fit macroeconomic model builders were able to achieve with large-scale macroeconomic models. This belief ended abruptly in the stagflation of the 1970s, in which the perceived stable and systematic trade-off between unemployment and inflation broke down and both unemployment and inflation rose to unprecedented postwar levels. The belief that tight-fitting models could generate accurate conditional forecasts also broke down and formed the basis of Robert Lucas's famous critique of econometric models.

Macroeconomics and economic modeling have advanced enormously since the large-scale models of the 1960s, and these advances are largely responsible for the return of macroeconomic model building to the forefront of central bank research and policymaking. But as Kocherlakota points out, identification of all parameters in these models is tenuous, particularly in models with many shock processes. Good macroeconomic practice almost by necessity requires sensitivity analysis that provides a systematic treatment of the uncertainty underlying model parameters. And when there is considerable uncertainty in conditional forecasts, policymakers should recognize this uncertainty as well.

Current policy and the current menu of models analyzed appear to be too responsive to the Phillips curve, more so than is warranted by the data. Model fit is a seductive property; it is hard for model builders to resist modifying model equations to achieve a better fit, even when the modifications do not have strong theoretical underpinnings. Fromm and Klein (1965) show how model builders of the 1960s focused on fit and modified models to achieve low mean square error, despite the fact there were few, if any, economic foundations for these modifications.

Econometric practice today is in some ways reminiscent of 1960s practice. Shocks are being

added to various equations to achieve a close fit to the data, but without necessarily understanding deeply what the frictions or market imperfections underlying these shocks are. And current policymaking seems far too responsive to a perceived Phillips curve that is not present in the data. We know all too well the outcome of the fitting exercises of the 1970s and the reliance on the Phillips curve. Perhaps the best way to avoid the monetary policy mistakes of the past is to remember that these mistakes were partly the consequence of relying too much on an empirical relationship that does not have strong theoretical underpinnings and that is not a robust feature of U.S. data. Agnostic approaches to modeling, as suggested by Kocherlakota, can significantly aid in the process of quantifying macroeconomic uncertainty and understanding its implications for monetary policy.

REFERENCES

- Atkeson, Andrew and Ohanian, Lee E. "Are Phillips Curves Useful for Forecasting Inflation?" Federal Reserve Bank of Minneapolis *Quarterly Review*, Winter 2001, 25(1), pp. 2-11.
- Canova, Fabio and Sala, Luca. "Back to Square One: Identification Issues in DSGE Models." Working Paper 583, European Central Bank, 2006.
- Diebold, Francis X.; Ohanian, Lee E. and Berkowitz, Jeremy. "Dynamic Equilibrium Economies: A Framework for Comparing Models and Data." *Review of Economic Studies*, July 1998, 65(3), pp. 433-51.
- Fromm, Gary and Klein, Laurence R. *The Brookings Quarterly Econometric Model of the United States*. Chicago: Rand-McNally, 1965.
- Haavelmo, Trygve. "The Probability Approach in Econometrics." *Econometrica*, 1944, 12(Suppl), pp. 1-118.
- Kocherlakota, Narayana. "Model Fit And Model Selection." Federal Reserve Bank of St. Louis *Review*, July/August 2007, 89(4), pp. 349-60.

- Koopmans, Tjalling, ed. *Statistical Inference in Dynamic Economic Models*. Cowles Commission Monograph Series No. 10. New York: John Wiley, 1950.
- Liu, Ta-Chung. "Underidentification, Structural Estimation, and Forecasting." *Econometrica*, October 1960, 28(4), pp. 855-65
- Nelson, Charles R. "The Prediction Performance of the FRB-MIT-PENN Model of the U.S. Economy." *American Economic Review*, December 1972, 62(5), pp. 902-17,
- Fernández-Villaverde, Jesús and Rubio-Ramírez, Juan. "Comparing Dynamic Equilibrium Economies to Data: A Bayesian Approach." *Journal of Econometrics*, November 2004, 123(1), pp. 153-87.
- Schorfheide, Frank. "Loss Function-Based Evaluation of DSGE Models." *Journal of Applied Econometrics*, November-December 2000, 15(6), pp. 645-70.
- Sims, Christopher. "Macroeconomics and Reality." *Econometrica*, January 1980, 48(1), pp. 1-48.
- Smets, Frank and Wouters, Raf. "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, September 2003, 1(5), pp. 1123-75.
- Stock, James and Watson, Mark. "Has Inflation Become Harder to Forecast?" Discussion paper, Harvard University, 2006.
- Tinbergen, Jan. *An Econometric Approach to Business Cycles*. Paris: Fr. Hermann Publishing, 1937.
- Watson, Mark W. "Measures of Fit for Calibrated Models." *Journal of Political Economy*, December 1993, 101(6), pp. 1011-41.

