Keynes vs. Prescott and Solow: Identifying Sources of Business Cycle Fluctuations

David N. DeJong*, Beth F. Ingram**, and Charles H. Whiteman**

March 1995

Abstract

Who was closer to the source of business cycle fluctuations--Keynes or Prescott and Solow? Two types of business-cycle impulses which have been associated with their names -- marginal efficiency of investment shocks (Keynes) and technology shocks (Prescott and Solow) -- are studied here in a neoclassical model which builds on the Greenwood, Hercowitz, and Huffman (1988) variable-utilization framework. The important parameters of the model are estimated using a Bayesian procedure which accommodates prior uncertainty about their magnitudes; from these estimates, posterior distributions of the two shocks are obtained. The postwar U.S. experience suggests that *both* shocks are important in understanding fluctuations, but that investment shocks are primarily responsible for beginning and ending recessions.

* The University of Pittsburgh
** The University of Iowa

We thank participants in the Iowa Macroeconomics Workshop, especially Narayana Kocherlakota, for helpful comments during the formative stages of this project, and Yi Wen for providing valuable research assistance. We also gratefully acknowledge the support provided for this research by the NSF under grants SBR 9422828 (to the University of Pittsburgh) and SBR 9422873 (to the University of Iowa). The source of errors is us.

... the trade cycle and, especially, the regularity of time-sequence and of duration which justifies us in calling it a cycle, *is mainly due to the way in which the marginal efficiency of capital fluctuates.*

John Maynard Keynes

... the observed increase in output per head is a consequence of "technical progress" rather than of increased capital per head.

Robert M. Solow

Economic fluctuations are optimal responses to uncertainty in the rate of technological change.

Edward C. Prescott

I. Introduction

According to Prescott (1986) the neoclassical growth model with technology shocks dramatically and surprisingly accounts for much of the observed variation in output known as the business cycle. Though there are some dimensions (several of which were pinpointed by Prescott) along which the model describes the data poorly, a substantial literature has developed in an attempt to explain economic fluctuations by building on the real business cycle models of Kydland and Prescott (1982) and Long and Plosser (1983).

The importance assigned in the modern literature to the role that technology shocks play in generating business cycle activity stands in contrast with Keynes' (1936) view that shocks to the marginal efficiency of capital are of primary importance in generating cycles. Recently, Greenwood, Hercowitz, and Huffman (GHH 1988) investigated the extent to which shocks of this sort can account for business cycle activity by incorporating this type of shock in a real business cycle model which features variable capital utilization and depreciation. To isolate the ability of the marginal efficiency of investment shock to produce fluctuations, they did not include a technology shock in their model. Yet after calibrating their model using long-run averages for the U.S., they were able to match observed fluctuations about as well as the growth *cum* technology shock model described in Prescott (1986).

The purpose of this paper is to investigate empirically the relative importance of investment and technology shocks in driving U.S. business-cycle activity. To pursue this investigation, we adopt a modified version of the GHH model which is augmented to include both investment and technology shocks. We use this model to identify probabalistically the joint time-series behavior of these shocks that is consistent with observed fluctuations in postwar U.S. output and investment.

Shock identification is facilitated using a formal Bayesian inferential procedure. The specification of the theoretical model yields a likelihood function for the observed data; combining this function with a prior distribution specified for the parameters of the model, we obtain a posterior distribution for these parameters. The time series of shocks induced by a particular parameterization of the model (and the observed data) is computable using a Kalman filter algorithm, thus moments of posterior distributions of the shocks can be computed by integrating over the posterior distribution of the parameters of the model.

Mindful that the inferences generated by our procedure are conditional on the specification of the theoretical model we consider, the prior distribution over its parameters we adopt, and the observed data, our analysis provides several insights concerning the pattern of shock behavior which underlies U.S. business-cycle activity. Despite the specification of prior diffusion over the parameters of the theoretical model, the posterior distributions of the shocks we obtain are rather tightly distributed, suggesting that the model and data are informative about the time-series behavior of the shocks. Just as output and investment exhibit recurrent patterns of behavior over the stages of the business cycle, so too do the shocks. Indeed, the specification of a simple logit model which employs current and lagged realizations of the shocks as explanatory variables yields reasonably accurate predictions of turning points identified by the NBER. While the patterns of shock behavior are somewhat irregular, large positive technology shocks, coupled with negative marginal efficiency of investment shocks, tend to precede recessions; large positive marginal efficiency of investment shocks, coupled with negative technology shocks, tend to precede recoveries. The rather surprising reaction of output to technology shocks emphasizes the importance of studying investment and technology shocks simultaneously: the shocks exhibit important interactions at key stages of the business cycle, and both are important in accounting for observed patterns of business-cycle activity.

II. A Real Business Cycle Model with Investment Shocks and Variable Capital Utilization

With one exception, we adopt the representative agent real business cycle framework described in Greenwood, Hercowitz and Huffman (GHH). This model departs from a standard framework in three respects: economic fluctuations are attributable to a single shock to investment; agents may vary the rate of capital utilization in response to this shock; and higher rates of capital utilization induce higher rates of depreciation of the existing stock of capital. We add to the model a second source of fluctuations: a standard total factor productivity (technology) shock.

The agent has available a constant-returns-to-scale production technology in which capital (k_t) and labor (λ_t) combine to produce output (y_t):

(1) $$y_t = A_t (k_t h_t) \quad l_t^{1-} \quad ,$$

where A_t is a technology shock which alters the marginal product of both factors of production, and $\alpha \in (0, 1)$ represents capital's share of output. The innovation in the model of GHH is the inclusion of h_t, which is an index of period-t capital utilization: in choosing h_t, the agent has the ability to employ the given capital stock more or less intensely during a period in response to shocks. Of course, absent any other constraint on utilization, the agent has no incentive to set h_t to a value less than one. However, GHH also assume that higher utilization rates induce higher rates of capital depreciation. Specifically, depreciation is given by $\ (h_t) = h_t \ / \ $, with $\omega > 1$; note that $\delta'(h) > 0$ and $\delta''(h) > 0$, so that the marginal cost of utilization of the capital stock is increasing in the rate of utilization.

Output is divided between two uses, consumption (c_t) and investment (i_t):

(2) $$y_t = c_t + i_t.$$

Investment increases the stock of capital according to the following accumulation relationship:

(3) $$k_{t+1} = k_t (1- \ (h_t)) + i_t \quad _t .$$

The shock ε_t alters the marginal efficiency of investment: a unit of output which is diverted from consumption to investment generates a bigger boost in the capital stock when ε_t exceeds unity than when it does not. In other words, ε_t determines the rate at which consumption goods can be substituted for capital goods which become productive tomorrow.

The representative agent chooses c_t and λ_t to maximiz*e the expected value of lifetime utility:

(4) $$E_t \sum_{t=0}^{\infty} \quad ^{t} \frac{1}{1-} \left[\left(c_t - \frac{l_t^{1+}}{1+} \right)^{1-} \quad -1 \right].$$

The parameter $\beta \in (0,1)$ represents the rate at which the agent discounts the future, $\gamma > 0$ represents the degree of risk aversion exhibited by the agent, and $1/\theta > 0$ measures the elasticity of substitution in labor supply. As pointed out in GHH, one feature of the period utility function is that the marginal rate of substitution between consumption and labor depends only on labor: income effects do not influence labor decisions. Hence the supply of labor hours does not move in response to expected future shocks.

The first-order conditions which result from maximizing (4) subject to (1) - (3) can be reduced to the following expressions (in addition to (1) - (3)):

(5) $$l_t^{\quad +1} = (1- \quad)y_t$$

(6) $$h_t \; = \quad {}_t y_t \,/\, k_t$$

(7) $$\frac{U_{ct}}{\quad_t} = \quad E_t\left\{U_{ct+1}\left(\frac{1}{\quad_{t+1}} + \quad \frac{y_{t+1}}{k_{t+1}}(1-\frac{1}{\quad})\right)\right\}.$$

Note that equations (1), (5) and (6) can be used to derive expressions for capital utilization and the supply of labor hours as a function of the beginning-of-period capital stock and the current shocks to investment and technology. Hence neither variable responds to expected future shocks. As an implication, the level of output is also unresponsive to expected future shocks.

In the current period, positive technology shocks directly raise labor supply, utilization and output; the increase in output is divided between consumption and investment, as agents attempt to smooth their lifetime consumption profile. Positive *investment* shocks increase current-period utilization, thus increasing labor supply and output. In addition, investment shocks reduce the relative price of capital, and hence induce agents to augment the future capital stock through additional investment. In the periods following the realization of a positive shock of either type, responses of output and consumption remain positive, while the response of utilization and investment become negative (i.e., utilization and investment fall below their steady state values) due to the buildup in the capital stock undertaken in the period in which the shock was realized.

To complete the specification of the model, we assume that the logs of the two shocks follow independent AR(1) processes:

(8) $$\begin{aligned} \ln A_t &= \quad_A \ln A_{t-1} + \quad_{At}, \qquad {}_{At} \sim N(0, \quad_A^2), \\ \ln \quad_t &= \quad \ln \quad_{t-1} + \quad_t, \qquad {}_t \sim N(0, \quad^2). \end{aligned}$$

To characterize approximate solutions to this model, we use a standard log-linearization procedure which produces a first-order difference equation in the state of the system, $s_t = [k_t\ A_t\ \varepsilon_t]'$:

$$\begin{bmatrix} \ln k_t \\ \ln A_t \\ \ln \quad_t \end{bmatrix} = \mathrm{H}_{-1}\begin{bmatrix} \ln k_{t-1} \\ \ln A_{t-1} \\ \ln \quad_{t-1} \end{bmatrix} + \begin{bmatrix} \quad_k & \quad_A & \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}\begin{bmatrix} 0 \\ \quad_{At} \\ \quad_t \end{bmatrix}.$$

Letting $\nu_t = [0\ \nu_{At}\ \nu_{\varepsilon t}]'$, we can write this more succinctly as

(9) $$s_t = \mathrm{H}_{-1} s_{t-1} + \mathrm{H} \quad_t \,.$$

Each choice variable at date t is a linear function of the state at time t:

$$\begin{bmatrix} \ln c_t \\ \ln l_t \\ \ln k_{t+1} \\ \ln h_t \end{bmatrix} = \Lambda \begin{bmatrix} \ln k_t \\ \ln A_t \\ \ln \ {}_t \end{bmatrix}.$$

In addition, output and investment can be written as linear functions of capital and the two shocks:

(10) $$\begin{bmatrix} \ln y_t \\ \ln i_t \end{bmatrix} = \quad \begin{bmatrix} \ln k_t \\ \ln A_t \\ \ln \ {}_t \end{bmatrix}.$$

The entries in the matrices are functions of the underlying parameters of the model, denoted by the 5x1 vector $\mu=[\alpha\ \beta\ \gamma\ \theta\ \omega]'$ and the 4x1 vector $\zeta = [\rho_A\ \rho_\varepsilon\ \sigma_A\ \sigma_\varepsilon]'$. Given a particular setting for μ and ζ, we can calculate H_{-1}, H_ν, Λ, and λ numerically. Our interest in this model lies in its ability to explain fluctuations in output and investment, hence we concentrate on the state transition equation (9) and the mapping from the state to output and investment given in (10).

III. Estimation

We measure investment using nonresidential fixed investment, and output using the sum of investment and consumption of nondurables and services. The series are measured in 1987 dollars, and were converted to *per capita* terms by dividing by the noninstitutionalized population (including armed forces) over 16 years of age. The sample runs from 1950:I to 1993:IV, and was obtained from CITIBASE. We denote the data set as $\{X_t = [\ln y_t\ \ \ln i_t]', t=1,...,T\}$.

The vector μ contains the parameters of the model which have clear economic interpretations. We have well-defined prior views concerning plausible specifications of these parameters, and are interested in whether the values favored by the data accord with these views. Alternatively, the parameters in the vector ζ merely govern the behavior of the exogenous shocks in the model; rather than estimating these parameters, we instead calibrated them. We did this to reduce the computational demands of our procedure, and to focus attention on the parameters of economic interest in the analysis. We calibrated ζ so that under the prior means specified for the parameters in μ (described below), the first-order serial correlation coefficient and standard deviation of output implied by the model matched the values observed in the data.[1]

---

[1] The lag coefficients for $\ln(A_t)$ and $\ln(\varepsilon_t)$ were 0.973 and 0.948; associated innovation standard deviations were 0.0068 and 0.00288. In calibrating these parameters, α, γ, and ω were set to the mean of the prior

Given equations (8), (9) and (10), the likelihood function for μ given the data $\{X_t\}$ can be constructed and evaluated using a Kalman filter algorithm. Here, equations (9) and (10) represent the state and observation equations. A classical approach to the problem of estimating μ involves searching over values of μ to find the particular value which maximizes the likelihood function. However, in this case we have clear prior views concerning reasonable values for the elements of μ, and wish to incorporate these views in our empirical analysis. (For example, small values for β imply unreasonably large yearly interest rates (e.g., 52 percent for β=0.9); we wish to rule out such values *a priori*.) We do so by adopting a Bayesian procedure which enables us to incorporate our prior views formally.

Aside from two sources of complication, our estimation procedure is a straightforward application of Bayesian analysis which involves the evaluation of posterior distributions of functions of interest, which we denote by $g(\mu)$. Letting $\pi(\mu)$ and $L(\mu|X)$ denote a prior distribution and likelihood function specified for μ, the posterior distribution of μ, by Bayes' Rule, is given by

$$P(\mu|X) \propto \pi(\mu)\, L(\mu|X).$$

Moreover, the expectation of $g(\mu)$ under the posterior distribution is given by

$$E[g(\mu)] = \int g(\mu)P(\mu|X)\, d\mu \,/\, \int P(\mu|X)\, d\mu. \quad (11)$$

Hence, our estimation problem involves evaluating (11) for suitably defined functions $g(\mu)$.

In specifying $\pi(\mu)$, we assumed that the components of μ were independently and normally distributed. With the exception of γ, the means we specified for $\pi(\mu)$ corresponded to the values employed by GHH in their quantitative analysis, which are consistent with the values generally used in the real business cycle literature. The mean specified for β is 0.99, which corresponds to a yearly interest rate of 4 percent. Since there seems to be general agreement about the value of this parameter, we assumed a relatively small standard deviation, 0.00125; the two-standard deviation band around β implies a yearly interest rate which varies between 3 and 5 percent. The mean specified for capital's share of output (α) is 0.29, with standard deviation 0.025, implying a two-standard-deviation band of [0.24, 0.34]. The parameter ω measures the curvature of the depreciation function and, in conjunction with β, determines the steady state rate of depreciation. The mean specified for ω is 1.6, with standard deviation 0.15; given the prior specified for β, this corresponds to an annual rate of depreciation in the range of 1 to 10 percent.

---

distribution described below. Since there were two moments and four parameters to calibrate, presumably there are other values for the lag coefficients and innovation variances which would work equally well.

We do not have sharply defined views concerning appropriate values for θ, which represents the inverse of the elasticity of labor supply. Hence we center the distribution of θ at the value specified by GHH -- 0.6 -- and specify a relatively large standard deviation of 0.1. The parameter γ governs the degree of risk aversion of the agent; as GHH noted, the value of this parameter is somewhat controversial. Mindful of this, they specified two values for γ: 1 and 2. Here, we compromise, setting the prior mean of γ to 1.5 and adopting a relatively large standard deviation of 0.25.

In general, equation (11) cannot be evaluated analytically, but must be approximated using numerical integration techniques. Given the ability to obtain drawings of μ directly from $P(\mu | y)$, equation (11) is well-approximated by

(12) $$\bar{g}_n = \sum_{i=1}^{n} g(\mu_i)$$

for large n.[2] However, the first complication we face in this application is that it is not possible to obtain drawings of μ directly from $P(\mu | X)$. Instead, we employ the importance sampling technique described by Geweke (1989). This technique involves obtaining drawings of μ from an importance density $I(\mu)$, and calculating

(13) $$\bar{g}_n = \sum_{i=1}^{n} g(\mu_i)\, w(\mu_i),$$

where the weight function $I(\mu_i)$ appears in the denominator of $w(\mu_i)$ to offset the direct influence that $I(\mu)$ has in obtaining the particular drawing μ_i. Given that the support of $I(\mu)$ includes that of $P(\mu | X)$, Geweke shows that $\bar{g}_n$ converges almost surely to $E[g(\mu)]$, so long as $E[g(\mu)]$ exists and is finite.

The practical implication of the first complication concerns the rate of convergence of $\bar{g}_n$ to $E[g(\mu)]$: the more closely $I(\mu)$ mimics $P(\mu | X)$, the more rapid the rate of convergence, and hence the fewer drawings of μ that are required in the numerical analysis to achieve small numerical standard errors. (Indeed, note that if $I(\mu) = P(\mu | X)$, (12) and (13) are identical.) In this application, the problem of tailoring $I(\mu)$ to $P(\mu | X)$ is exacerbated by the second complication we face in this application: since the state (containing the capital stock and the two shocks) is unobservable, we

[2] For example, suppose $g(\mu)$ denotes a Bernoulli trial with success probability q, which is estimated by the number of successes divided by the number of trials. Then the numerical standard error of $\bar{g}_n$ is $[q(1 - q)/n]^{1/2}$, which is no greater than 0.005 when n = 10,000.

cannot express L(μ|X) analytically, hence the properties of P(μ|X) are difficult to ascertain. Evaluation of L(μ|X) for a specific value of μ requires the use of a Kalman filter algorithm on the system (9) and (10). For every μ, we must linearize the model to obtain the mapping from μ to the matrices in (9) and (10); given H_v, H_{-1} and λ, we then loop through the data set to obtain L(μ|X). Hence, it is somewhat computationally expensive to explore the behavior of P(μ|X).

We specified a multivariate t distribution for I(μ) to insure that its support included that of P(μ|X). The mean and covariance matrix of I(μ) we ultimately employed resulted from a sequence of preliminary runs. Initially, we assigned the prior mean and covariance matrix to I(μ), and over 50,000 drawings computed first-pass approximations of the posterior mean and covariance matrix of μ using (13). However, very few of the drawings obtained in this manner were assigned appreciable weight by the posterior distribution, so we relocated I(μ) at our first-pass approximations and obtained second-pass approximations using 50,000 more drawings. After several rounds our moment calculations converged (subject to numerical sampling error) to those used in deriving the results presented below, which are based on 75,000 drawings. Of these drawings, that which was assigned the greatest weight received only 0.35 percent of the total assigned weight, hence we are confident that our results closely approximate the actual posterior calculations we seek.

The posterior distributions of interest in our empirical analysis include the distributions of the elements of μ, and impulse response functions which map the response of output and investment to the investment and technology shocks. However, of primary interest are posterior distributions of the shocks themselves. These are obtained as a byproduct of the Kalman filter calculations used in evaluating the likelihood function L(μ|X). Conditional on information available through time t-1 and a particular drawing of μ, the mean of the state vector at time t is given by $\bar{s}_{t|t-1} = E(s_t | s_{t-j}, j > 0, \ )$. Also, the smoothed estimate of the state at time t is given by the mean of the state conditional on the entire data set: $\bar{s}_{t|T} = E(s_t | s_j, 0 < j \leq T, \ )$. Using smoothed estimates of the state in period t, we obtain estimates of the corresponding shocks:

(14) $$\exists_{t|T} = H^{-1}(\bar{s}_{t|T} - H_{-1}\bar{s}_{t-1|T}).$$

Hence for a given drawing μ, we can obtain the entire time series of shocks implied by the model and the actual data.[3] Integrating over μ yields period-by-period posterior distributions of these

[3] Since the smoothed estimates of the state vector have nonzero conditional variances, so too do the smoothed estimates of the shock vector. Here, we used two observed series to recover two shocks, hence for a given

shocks: conditional on the model and the observed data, these distributions permit probability statements concerning the sequence of shocks that have driven U.S. business cycle fluctuations.

Each of the posterior distributions we analyze is approximated using histograms constructed over the 75,000 drawings of μ mentioned above. For example, the posterior distribution of the technology shock at time t was constructed by calculating the time-t shock computed for each drawing μ_i, and assigning the weight $w(\mu_i)$ to the relevant bin of the histogram established for the shock at that date.[4] The posteriors we obtained in this manner are described below.

IV. Results

Figure 1 presents posterior distributions of the parameters of the model. The solid line in each panel represents the posterior distribution of the relevant parameter, the long dashed line represents the prior distribution, and the short dashed line represents the flat-prior posterior distribution.[5] In addition, Table 1 provides the posterior moments and correlation matrix of the parameters. The parameter α is fairly tightly estimated around 0.29, the value most commonly used in the real business cycle literature; notice that values exceeding 0.33 receive almost no posterior weight. The posterior distribution of β lies to the right of the prior distribution, hence agents are more patient than is suggested by our prior. Agents are also more risk averse, as the posterior distribution of γ also lies to the right of our prior: values of γ less than 1.2 receive virtually no posterior weight. The posterior distribution of the depreciation parameter ω is centered on top of our prior, and is far less diffuse, assigning negligible weight outside the range [1.4, 1.8]. Finally, the posterior distribution for θ is centered near our prior but is relatively diffuse. This result is perhaps to be expected: since θ governs the behavior of intratemporal substitution between leisure and consumption, we would not expect data on output and investment to contain as much information about this parameter as it contains for others such as ω.

---

drawing of μ the only impediment to uncovering the precise realization of the implied shock is that the initial capital stock k_0 is unobserved. Since the system is stationary, the effect of k_0 on the system decreases over time, hence the precision of our estimates of the shock vector improves as we move away from the endpoints of the sample. In practice, the variance of the shock estimates, conditional on μ and the entire data set, was zero except for the first few and final few periods in the sample.

[4] There are 176 observations in our data set and one is lost because of the lagged state variable, hence we construct 175 histograms for both the investment and technology shocks.

[5] This latter distribution was obtained by rescaling the weights in each bin of the model-based posterior histogram by a factor proportional to the inverse of the corresponding prior weights.

The most interesting entry in the correlation table is that for the correlation between β and ω, which is -0.859. We believe this strong negative correlation arises because the data have a strong predilection for certain values of the rate of depreciation. In this model, the steady state rate of depreciation is given by

$$\overset{-}{\ } = \frac{1/\quad -1}{\quad -1}.$$

Steady state depreciation, then, is a decreasing function of both β and ω: to keep steady state depreciation constant, a rise in β must be accompanied by a decline in ω. Here, the data prefer a steady state rate of depreciation of 1.7 percent per quarter (about 7 percent annually). This rate is somewhat lower that that typically used in the business cycle literature, though it is similar to the magnitude reported by Ingram, Kocherlakota and Savin (1994).

The results in Figure 1 and Table 1 indicate that the typical real business cycle literature values for the parameters of interest are not wildly at variance with the data. Our prior distributions were independent, and located in accord with these values; the data did little to move us from away from those views. But with the possible exception of ω, the data also did not refine the precision in our views much: the posterior distributions are quite diffuse relative to the amount of diffusion (none) permitted in the usual real business cycle model calibration exercise.

Figure 2 provides an illustration of how the model behaves when beset by shocks. Using the posterior-mean values of the parameters given in Table 1, we generated impulse response functions of output and investment to the technology and marginal efficiency of investment shocks, and obtained two notable results. First, unit-for-unit, technology shocks have much larger effects on output and investment than do investment shocks--by a factor of about five. Thus in our model, a typical technology shock has a greater economic impact than does a typical investment shock: Round I on the Keynes vs. Prescott-Solow card goes to Prescott-Solow. Second, the effect on output of an investment shock is essentially permanent. In contrast, technology shocks die out slowly, having a half-life of about eight years; the effect on investment of an investment shock dissipates even faster. But despite the fact (not evident in the figure) that the effect on output of an investment shock eventually dies out, the dissipation is negligible after eight years: only in the very long run is the marginal efficiency of investment shock dead. Round II: Keynes.

These two features of the impulse responses begin to tell the story of how the right sort of coincident movements in the two shocks can begin or end recessions. To explore this further, it is helpful to examine posterior distributions of the shocks themselves.

Figure 3 displays time series of output and investment growth (bottom panel) along with posterior median values of the two shocks (top panel). It also highlights recessions, as identified by the NBER, with shaded bars. There are eight such recessions in our sample. The onset of five of these recessions is attributable exclusively to negative marginal productivity of investment shocks: the 1953, 1970, 1974, 1980, and 1982 recessions coincide with the realization of large negative investment shocks coupled with positive productivity shocks. The 1957 and 1960 recessions were preceded by the realization of negative shocks of each type, while the 1990 recession is attributable to a negative productivity shock. Round III: Keynes.

Figure 3 also indicates that the onset of six of the eight recoveries in our sample is attributable exclusively to positive investment shocks: the 1953, 1958, 1960, 1974, 1982 and 1990 recoveries coincide with the realization of large positive investment shocks coupled with large negative productivity shocks. The 1970 and 1980 recoveries coincide with the realization of positive shocks of each type. Round IV: Keynes.

In considering these patterns of shock behavior, it is important to recognize two regularities in the behavior of output and investment growth prior to and during recessions. At the onset of recessions, both series exhibit falling if not negative growth; at the onset of recoveries, output growth turns positive, while investment growth often remains negative. Indeed, of the seven most dramatic realizations of negative investment growth in our sample, six coincide with recoveries. In the context of the model, the key to understanding how upturns in output can coincide with negative investment growth lies in the behavior of capital growth, k_{t+1} - k_t. Even with negative investment growth, capital growth can be positive, for two reasons: (i) utilization is low during recessions, hence so too is depreciation; (ii) negative investment growth in the face of a positive marginal efficiency of investment shock can leave the gross addition to the capital stock positive. In the context of the onset of a recovery, these two effects together imply that a positive (and highly persistent) investment shock can cause output to rise despite a coincident negative technology shock: the technology shock induces negative investment growth but does not offset the positive and virtually permanent effect on output induced by the investment shock.

Figure 4 highlights the behavior of the two shocks during the 1974 and 1982 recessions by displaying 2nd, 50th and 98th percent quantiles of their posterior distributions during these periods. This figure highlights the patterns of behavior described above, and also illustrates that the posterior distributions of the shocks are rather tightly distributed. This latter characteristic indicates that the model and data are relatively informative concerning the patterns of shock behavior that are consistent with business-cycle fluctuations: in any given period, the distributions of the investment and technology shocks rarely overlap, enabling us to clearly distinguish between their relative magnitudes. The distributions also display little overlap with distributions in adjoining periods, enabling us to clearly identify movements in the shocks over time.

To highlight the recurrent patterns of behavior exhibited by the shocks over the stages of the business cycle, we specified three logit models which employed functions of the shocks as explanatory variables, and assessed the success of these models in predicting NBER turning points. The first two models examined the predictive power of current and four lagged shocks in isolation, and the last examined their joint predictive power. Figure 5 presents the predictions generated by these models.

The striking feature of Figure 5 is clear improvement in predictive power that is gained by incorporating the joint behavior of the shocks in the logit model. Using a fitted value in excess of 0.5 as the trigger for the prediction of a recession, the model which employs only the investment shocks predicts four false recessions and misses three of the recessions which actually occurred, and the model which employs only the technology shock predicts five false recessions and misses four. In contrast, the joint model misses only the 1970 and 1990 recessions, and generates only one false prediction. These results serve to highlight the importance of studying these shocks simultaneously: the shocks exhibit recurrent and significant interactions over the stages of the business cycle, and both are important in accounting for observed patterns of business-cycle activity. Round V: draw.

V. Conclusion

It may well be that growth is due primarily to “technical progress” rather than increases in capital per head. We also recognize the important role that technological change plays in driving business-cycle fluctuations. However, our view of the relative importance of this source of cyclical behavior has been altered dramatically by the results of this analysis. Studying the joint behavior of shocks to both technology and the marginal efficiency of investment over the stages of the business cycle, we are struck by the propensity for investment shocks to begin and end recessions. At the risk of being labeled Keynsians, or more eponymously, Bayesian Keynsians, we conclude that the data support Keynes’ view that shocks to the marginal efficiency of investment are the predominant source of business-cycle fluctuations.

References

Geweke, John (1989), “Bayesian Inference in Econometric Models Using Monte Carlo Integration,” *Econometrica*, 57: 1317-1339.

Greenwood, Jeremy, Zvi Hercowitz and Gregory Huffman (1988), "Investment, capacity utilization, and the real business cycle," *American Economic Review*, 78:402-417.

Ingram, B., Kocherlakota, N., and N.E. Savin (1994), “Explaining Business Cycles: A Multiple Shock Approach,” *Journal of Monetary Economics* 34: 415-428.

____ (1993b), “Letting the Macroeconomic Data Tell Us About the Great Productivity Slowdown,” University of Iowa manuscript.

Kydland, Finn and Edward Prescott (1982), "Time to build and aggregate fluctuations," *Econometrica*, 50:1345-1370.

Long, J. B., and C. I. Plosser (1983), "Real Business Cycles," *Journal of Political Economy* 91 (February):39-69.

Prescott, E. (1986), "Theory Ahead of Business Cycle Measurement," Federal Reserve Bank of Minneapolis *Quarterly Review* 10:9-22.

Table 1

	α	β	γ	ω	θ
Prior Mean	0.290	0.990	1.500	1.600	0.600
Post. Mean	0.289	0.991	1.836	1.589	0.602
Prior Stan. Dev.	0.025	0.00125	0.250	0.150	0.100
Stan. Dev.	0.017	0.001	0.195	0.057	0.086

Table 1a. Prior and posterior moments of parameter vector.

	α	β	γ	ω	θ
α	1.00				
β	-0.197	1.00			
γ	-0.239	-0.022	1.00		
ω	-0.197	-0.850	0.193	1.00	
θ	-0.293	-0.103	-0.191	-0.137	1.00

Table 1b. Posterior correlations among parameters.

(Prior cross-correlations are zero.)

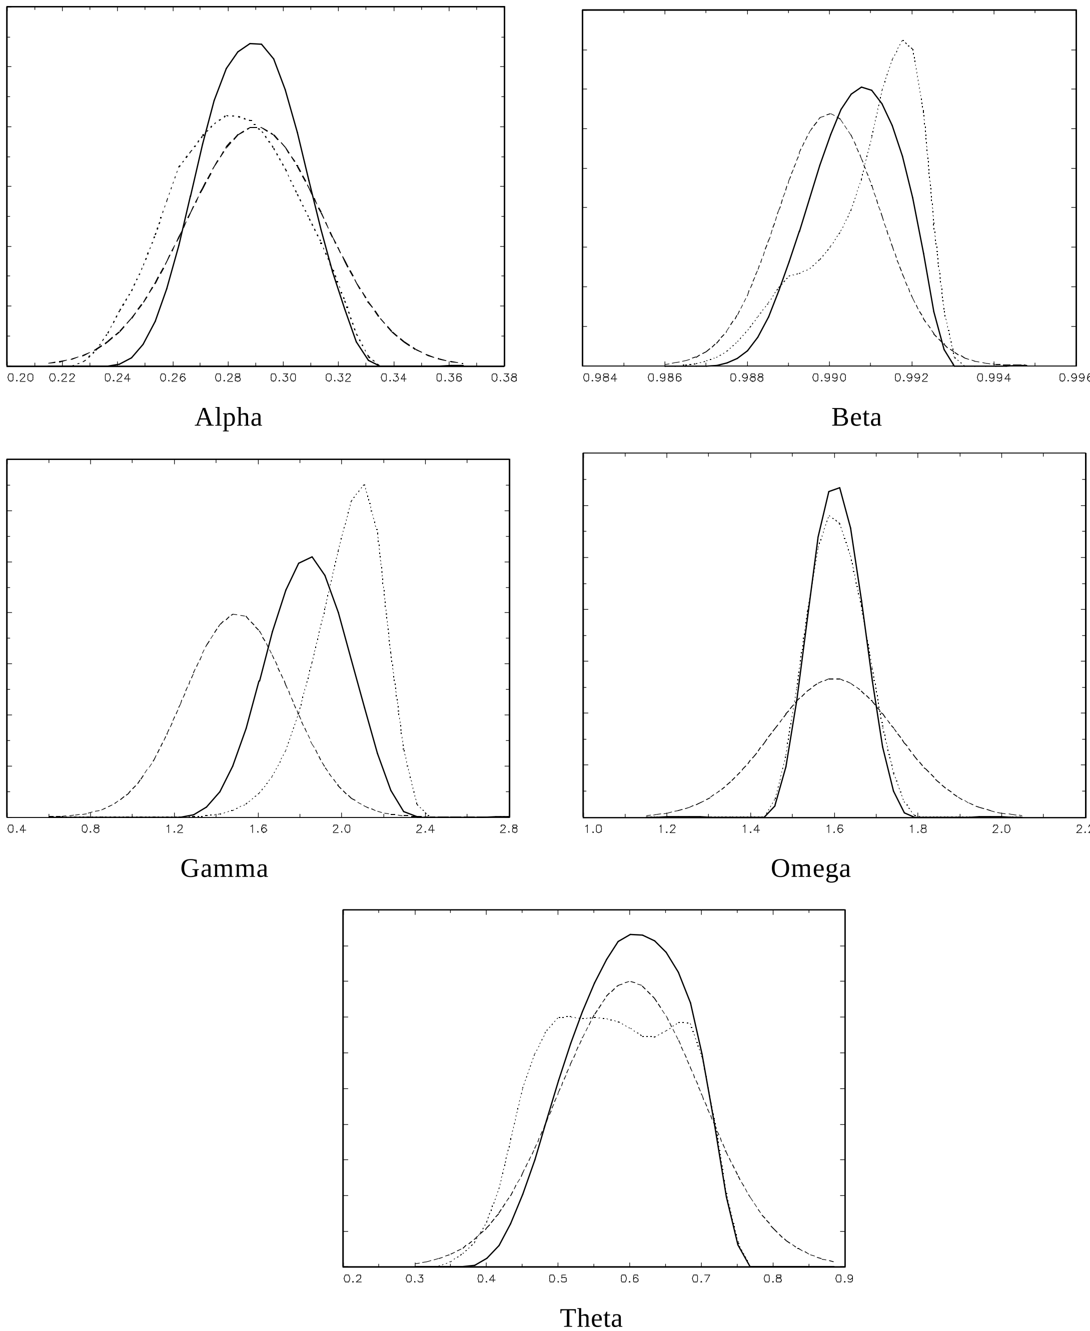

Figure 1. Parameter Histograms. Prior distribution (long dashes), flat-prior posterior distribution (short dashes), and posterior distribution under model prior (solid line) for the five parameters of interest.

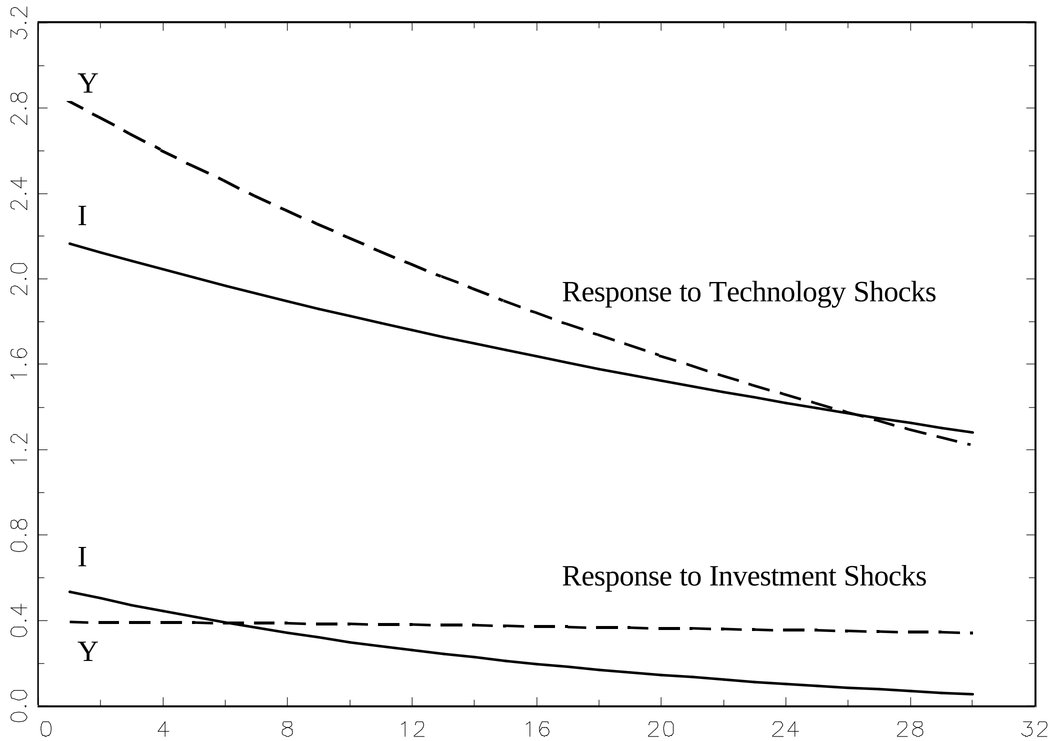

Figure 2. Impulse Response Functions. Impulse response functions to a one-standard-deviation technology shock (upper lines) and investment shock (lower lines). Results obtained using posterior means of parameters. Dashed lines represent responses of output; solid lines represent responses of investment. The vertical scale respresents own-standard-deviation units.

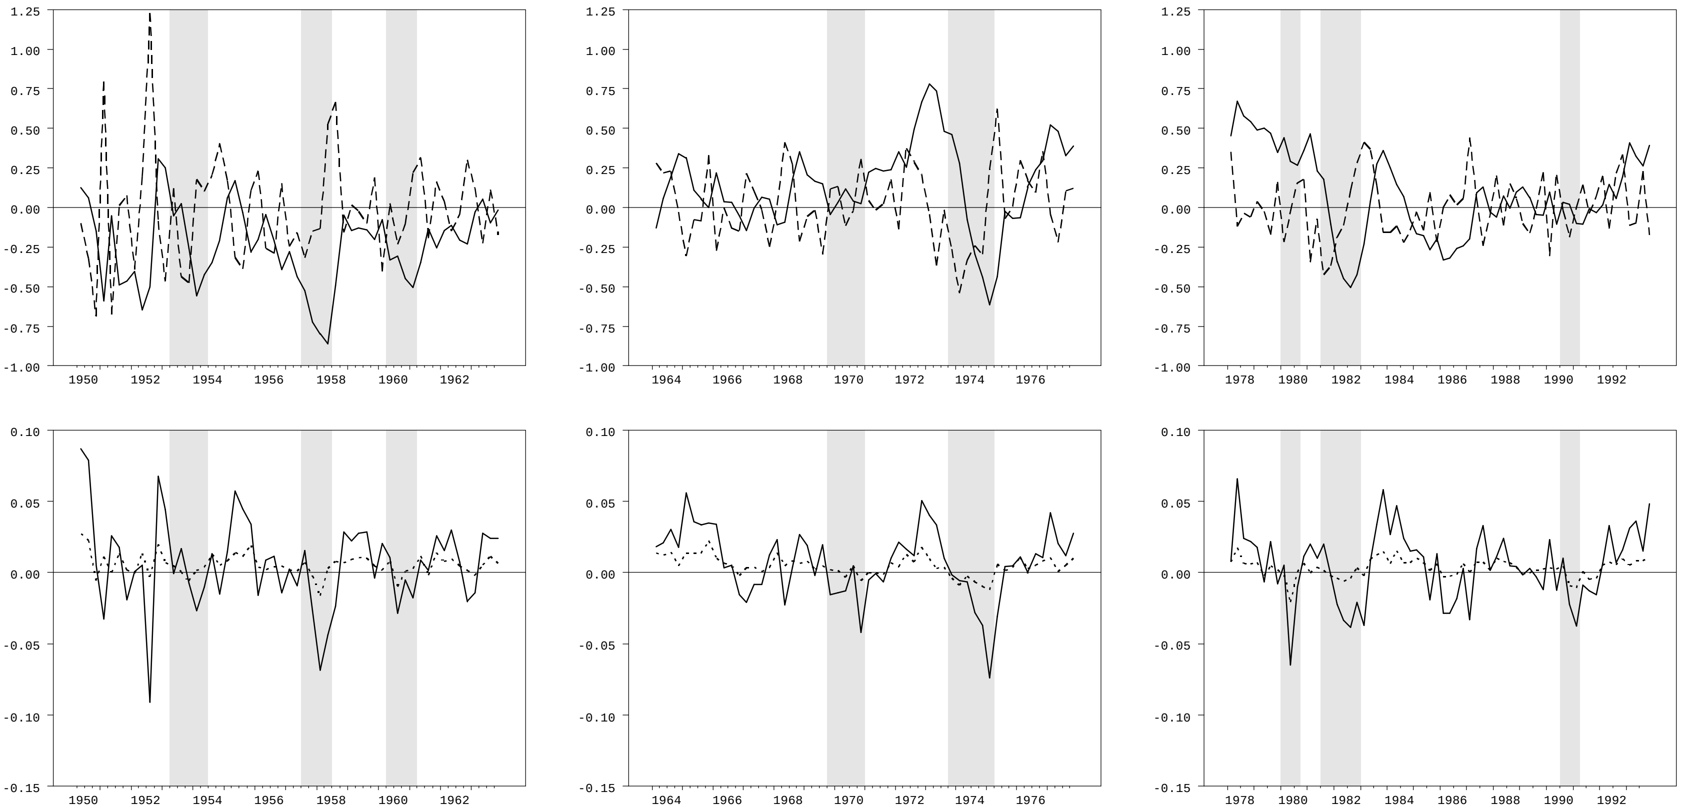

Figure 3. Time Series of Data and Shocks. In the top panel, the dotted line represents the 50th quantiles of the posterior distributions of the investment shocks; the solid line is analagous for the technology shocks. Each set of shocks is measured in own-standard-deviation units. In the bottom panel, the dotted line represents the growth rate in output, and the solid line represents the growth rate in investment. Shaded bars denote recessions (NBER dating scheme).

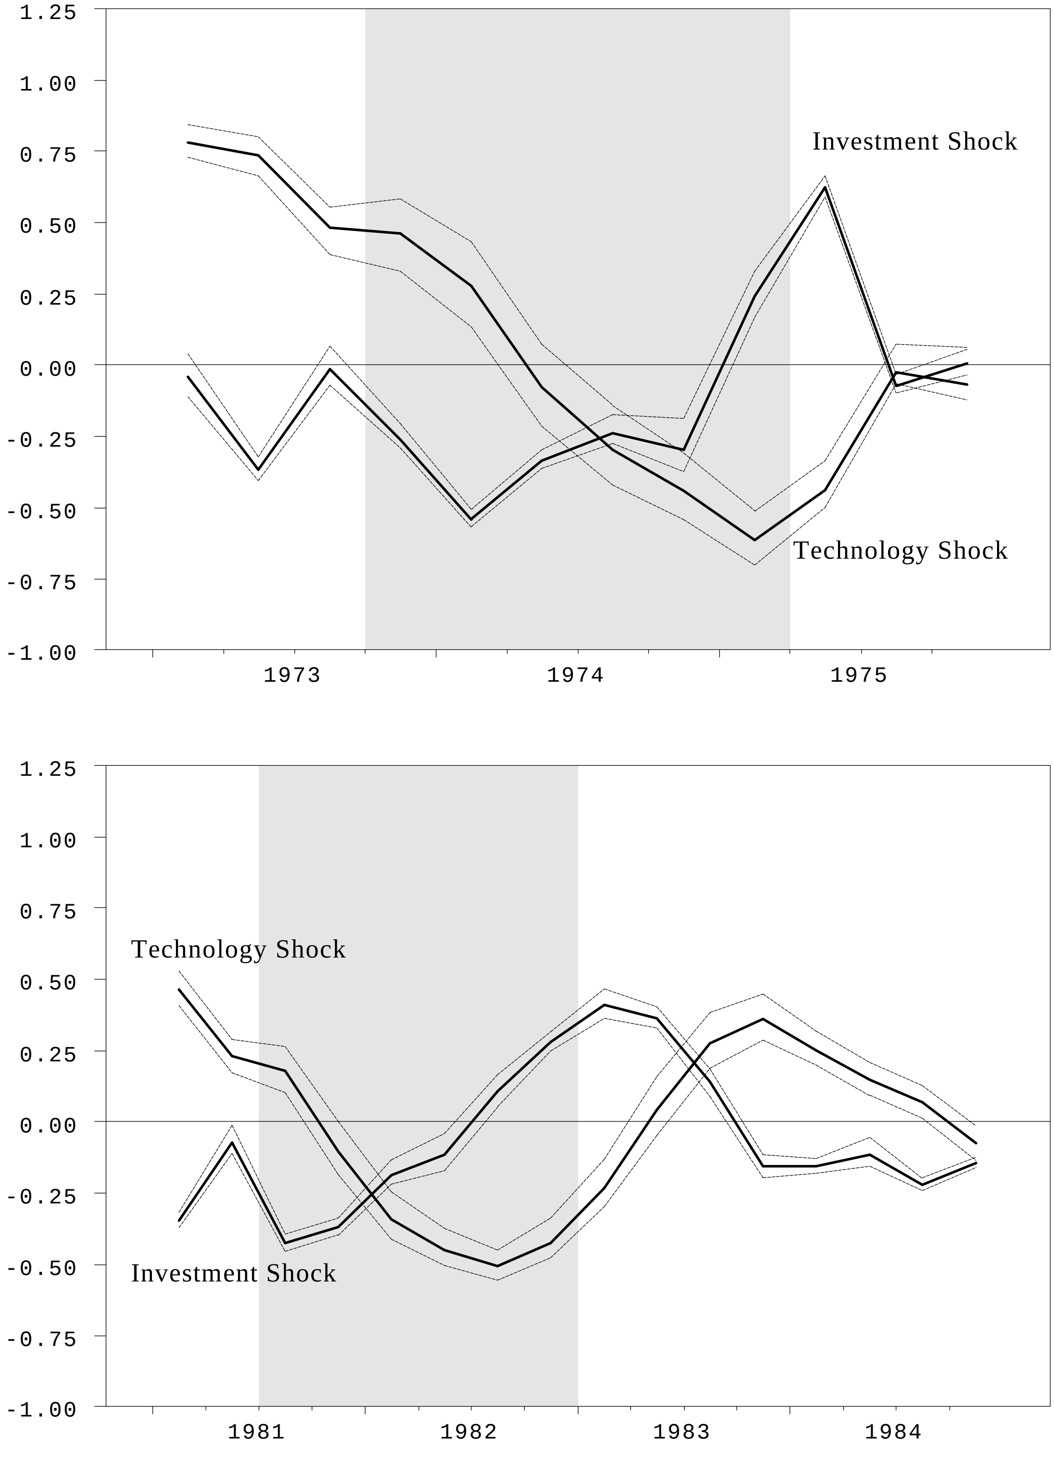

Figure 4. Close-ups of 1974 and 1982 Recessions. 2nd, 50th and 98th percent quantiles of posterior distributions of shocks during the 1975 and 1982 recessions.

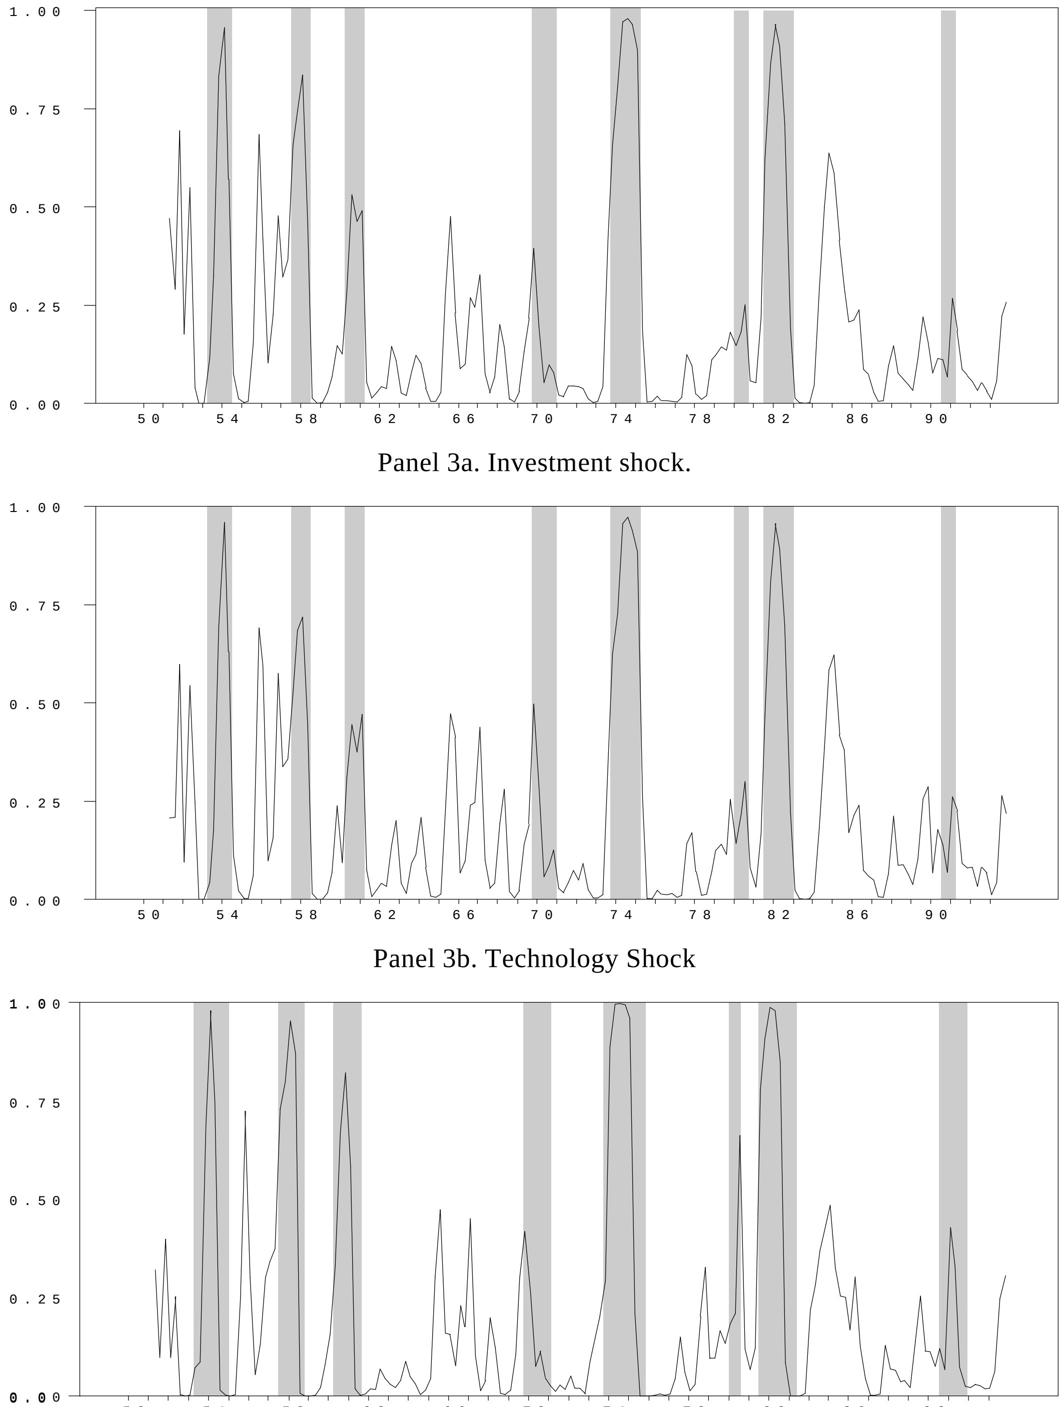

Panel 3a. Investment shock.

Panel 3b. Technology Shock

Panel 3c. Combined Investment and Technology Shocks

Figure 5. Probabilities of NBER Recessions. Probabilities generated by logit models which employed current and four lagged realizations of the indicated shocks. Shaded bars denote recessions (NBER dating scheme).