Learning-by-Doing or Habit Formation?*

Hafedh Bouakez† Takashi Kano‡

This Version: May 11, 2005

Abstract

In a recent paper, Chang, Gomes, and Schorfheide (2002) extend the standard real business cycle (RBC) model to allow for a learning-by-doing (LBD) mechanism whereby current labour supply affects future productivity. They show that this feature magnifies the propagation of shocks and improves the matching performance of the standard RBC model. In this paper, the authors show that the LBD model is nearly observationally equivalent to an RBC model with habit formation in labour (or, equivalently, in leisure). Under the same calibration of the parameters, the two models share the same equilibrium paths of output, consumption, and investment, but have different implications for hours worked. Using Bayesian techniques, the authors investigate which of the LBD and habit models fits the U.S. data best. Their results suggest that the habit specification is more strongly supported by the data.

JEL classification: C52, E32, J22
Bank classification: Business Cycle Fluctuation and Cycles, Labour Markets, Economic Models, Econometric and Statistical Methods

*We thank Mick Devereux, Thomas Lubik, Jim Nason, Frank Schorfheide, and seminar participants at the Bank of Canada, CIREQ, Concordia University, and HEC Montréal for helpful comments and suggestions. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

†Department of Economics and CIREQ, HEC Montréal, 3000 chemin de la Côte-Sainte-Catherine, Montréal, Québec, Canada H3T 2A7. Tel.: 1-514-340-7003; Fax: 1-514-340-6469; E-mail: hafedh.bouakez@hec.ca.

‡International Department, Bank of Canada, 234 Wellington St., Ottawa, Canada K1A OG9. Tel.: 1-613-782-7391; Fax: 1-613-782-7658; E-mail: tkano@bank-banque-canada.ca.

1. Introduction

Owing to their elegance and parsimony, real business cycle (RBC) models have always been popular in macroeconomics. However, as emphasized by Rouwenhorst (1991) and Cogley and Nason (1995), among others, the standard RBC model is hampered by the weakness of its internal propagation mechanism. This deficiency manifests itself in several ways. In particular, the standard RBC model fails to generate a persistent and hump-shaped response of output to a transitory shock, as is typically found in the empirical literature (see, for example, Blanchard and Quah 1989). It also fails to replicate the autocorrelation function of output growth: while the data show that output growth is positively autocorrelated over short horizons, the autocorrelations predicted by the standard RBC model are essentially zero over all horizons.[1]

During the past decade, several studies have attempted to improve the matching performance of the standard RBC model by extending it along various dimensions. Extensions made by earlier papers include labour-adjustment costs (Cogley and Nason 1995), factor hoarding (Burnside and Eichenbaum 1996), labour-market search and matching (Andolfatto 1996 and den Haan, Ramey, and Watson 2000) and the combination of habit formation in leisure and increasing returns to scale (Wen 1998). These mechanisms magnify the propagation of shocks in the economy, thus enabling the model to replicate the dynamic pattern of output and other salient features of the business cycle.

More recently, Chang, Gomes, and Schorfheide (2002, hereafter CGS) have extended the standard RBC model to allow for a learning-by-doing (LBD) mechanism whereby hours worked in a given period increase workers' skill, which in turn increases their labour productivity in subsequent periods.[2] CGS demonstrate that this mechanism improves the ability of the standard RBC model to generate empirically plausible output fluctuations. This suggests that the LBD-augmented model embodies a quantitatively important propagation mechanism.

In this paper, we show that the LBD model developed by CGS is nearly observationally equiva-

[1] More generally, the spectral density function of output growth generated by the standard RBC model is flat. The actual spectral density, in contrast, exhibits spectral peaks over business cycle frequencies.

[2] Cooper and Johri (2002) introduce an alternative version of LBD, in which the production technology depends on the stock of organizational capital, which is accumulated through past production activities.

lent to an RBC model with habit formation in labour (or, equivalently, in leisure), henceforth called the habit model. Under a coefficient restriction, the LBD and habit models share the same equilibrium paths of output, consumption, and investment, but differ in implications regarding hours worked. Thus, data on output, consumption, and investment provide no useful information in distinguishing between the two models. Only when a measure of hours worked (or labour productivity) is included in the information set of an econometrician are the two models distinguishable.

Habit formation in labour implies that current utility depends on current labour supply relative to a reference level determined by past hours worked. This specification implies that habit-forming agents dislike large swings in their hours of work (or leisure time), and are more willing to smooth the path of their labour supply than agents with time-separable preferences. As a result, habit formation in labour will typically lead to a sluggish adjustment of hours worked in response to shocks. Although both learning-by-doing and habit formation make the optimal choice of labour supply non-time-separable, the two mechanisms are distinctly different in nature. Learning-by-doing is akin to a production externality that is internalized by workers. It implies that past hours of work affect both labour demand and labour supply. Habit formation, on the other hand, is a preference feature that affects only the labour supply schedule.

The purpose of this paper, therefore, is to determine which of the LBD and habit models is best supported by the data. Addressing this question is important because, as stated above, the two specifications have different implications for the dynamics of hours worked in the economy. But, more importantly, our results will provide some insights into the fundamental question of whether preferences are important in explaining business cycle fluctuations.

We start by showing that, under the same calibration of the structural parameters, the habit model generates hump-shaped responses of hours worked to permanent and transitory shocks, while the LBD model fails to do so. This suggests that habit formation is a stronger propagation mechanism than learning-by-doing, because it leads to richer dynamics in the labour market. To assess the empirical plausibility of the habit model and whether it fits the data better than the LBD model, we estimate both models using U.S. data on output growth and hours worked, and perform formal statistical comparisons. As in CGS, our econometric analysis is based on Bayesian

techniques. Our results indicate that, in almost every dimension (overall statistical fit, impulse-response functions, and autocorrelations), the habit model explains the data better than does the LBD model.

The rest of this paper is organized as follows. In section 2, we show the near-observational equivalence between the LBD and habit models. In section 3, we use impulse-response analysis to compare the theoretical predictions of habit formation and learning-by-doing. Section 4 describes the econometric methodology and discusses the results. Section 5 performs a robustness analysis. Section 6 concludes.

2. Learning-by-Doing and Habit Formation: Near-Observational Equivalence

CGS extend the standard one-sector RBC model to allow for an LBD mechanism associated with labour effort. In their model, households internalize the LBD process when choosing their labour supply. Under this assumption, the decentralized equilibrium is equivalent to the command optimum. Hence, without loss of generality, we rewrite the decentralized-economy model of CGS as one in which a benevolent central planner solves the following problem:

$$\max_{\{C_i\}_{i=t}^{\infty},\, \{H_i\}_{i=t}^{\infty},\, \{X_{i+1}\}_{i=t}^{\infty},\, \{K_{i+1}\}_{i=t}^{\infty}} E_t \sum_{i=0}^{\infty} \beta^i \left[\ln C_{t+i} - B_{t+i} \frac{H_{t+i}^{1+1/v}}{1+1/v} \right], \tag{1}$$

subject to

$$\begin{aligned}
Y_t &= C_t + I_t, && (2)\\
Y_t &= K_t^{1-\alpha}(A_t X_t H_t)^{\alpha}, && (3)\\
K_{t+1} &= (1-\delta)K_t + I_t, && (4)\\
\ln(X_t/X) &= \phi \ln(X_{t-1}/X) + \mu \ln(H_{t-1}/H), \qquad 0 \le \phi < 1, \quad \mu \ge 0, && (5)\\
\ln A_t &= \gamma + \ln A_{t-1} + \epsilon_t^a, \qquad \epsilon_t^a \sim \mathcal{N}(0, \sigma_a^2), && (6)\\
\ln B_t &= (1-\rho)\ln B + \rho \ln B_{t-1} + \epsilon_t^b, \qquad 0 \le \rho < 1, \quad \epsilon_t^b \sim \mathcal{N}(0, \sigma_b^2), && (7)
\end{aligned}$$

and the initial conditions $K_t \ge 0$, $X_t > 0$, $A_{t-1} > 0$, and $B_{t-1} > 0$.

The utility of the "representative household" is defined over consumption, C_t, and hours of work, H_t. A preference shock, B_t, affects the marginal disutility of labour. Output, which is either consumed or invested, is produced using a Cobb-Douglas technology described by (3), where K_t is the stock of capital, X_t is the skill level, and A_t is an exogenous technology shock. Investment increases the stock of capital according to (4), where δ is the depreciation rate of capital. The skill level evolves according to (5). This process implies that the current stock of skill is a moving average of past hours worked. It also indicates that the effect of past hours of work on the skill level is persistent but not permanent, with the persistence measured by the parameter ϕ. The technology shock, A_t, follows a random walk with drift, while the preference shock, B_t, follows a stationary AR(1) process.

First-order necessary conditions for the planner's problem are:

$$1 = \beta E_t \left(C_t/C_{t+1}\right)\left[\left(1-\alpha\right)\left(Y_{t+1}/K_{t+1}\right)+1-\delta\right], \tag{8}$$

$$B_t H_t^{1+1/v} = \alpha\left(Y_t/C_t\right)+\beta E_t\left[\phi B_{t+1} H_{t+1}^{1+1/v} - \alpha(\phi-\mu)\left(Y_{t+1}/C_{t+1}\right)\right]. \tag{9}$$

Equations (8) and (9) are, respectively, the standard Euler equation for consumption and the equilibrium condition for hours worked. Unlike the RBC model, where the (static) first-order condition for the optimal choice of hours equates the marginal disutility of work to the marginal utility of consuming the share of output accruing to labour in a given period, the corresponding condition in the LBD model (i.e., equation (9)) is dynamic. To understand the intuition behind this equation, consider the marginal costs and benefits of working an additional unit of time in period t. The immediate marginal cost is a utility loss equal to $B_t H_t^{1/v}$, while the immediate marginal benefit is a utility gain from consuming additional output in period t equal to $\alpha Y_t/C_t H_t$. In the standard RBC model, the intratemporal optimality condition simply equates these two terms. But in the LBD model, the additional unit of time worked in period t raises the skill level in period $t+1$ by $\mu X_{t+1}/H_t$. This implies that less hours in period $t+1$ are required to achieve the same skill level, X_{t+2}. In particular, hours worked in $t+1$ decrease by $\phi H_{t+1}/H_t$. The higher skill level and the lower labour supply in period $t+1$ yield additional marginal costs and benefits beyond those of the current period. On the one hand, utility increases from consuming the additional output

generated by the improvement in productivity ($\alpha\beta\mu Y_{t+1}/C_{t+1}H_t$), and from enjoying more leisure in period $t+1$ ($\phi\beta B_{t+1}H_{t+1}^{1+1/v}/H_t$). On the other hand, working less hours decreases output and the household's utility in period $t+1$ ($-\alpha\beta\phi Y_{t+1}/C_{t+1}H_t$). The Euler equation (9) states that, for the path of hours worked to be optimal, the current and future (expected) marginal costs and benefits of working an additional unit of time in period t must be equal. Note that it is the dependence of the skill level on past hours of work that makes the first-order condition (9) dynamic and allows the LBD model to endogenously generate persistence.[3]

In what follows, we show that under the coefficient restriction $0 \leq \phi - \mu < 1$, the LBD model, (1)–(7), can be rewritten as an RBC model with (internal) habit formation in labour. Indeed, by defining $N_t \equiv X_t H_t$, the central planner's problem becomes

$$\max_{\{C_i\}_{i=t}^{\infty},\, \{N_i\}_{i=t}^{\infty},\, \{X_{i+1}\}_{i=t}^{\infty},\, \{K_{i+1}\}_{i=t}^{\infty}} E_t \sum_{i=0}^{\infty} \beta^i \left[\ln C_{t+i} - B_{t+i} \frac{(N_{t+i}/X_{t+i})^{1+1/v}}{1+1/v} \right], \tag{10}$$

subject to

$$\begin{aligned}
Y_t &= C_t + I_t, && (11)\\
Y_t &= K_t^{1-\alpha}(A_t N_t)^{\alpha}, && (12)\\
K_{t+1} &= (1-\delta)K_t + I_t, && (13)\\
\ln(X_t/X) &= (\phi-\mu)\ln(X_{t-1}/X) + \mu \ln(N_{t-1}/N), \qquad 0 \leq \phi - \mu < 1, \quad \mu \geq 0, && (14)\\
\ln A_t &= \gamma + \ln A_{t-1} + \epsilon_t^a, \qquad \epsilon_t^a \sim \mathcal{N}(0, \sigma_a^2), && (15)\\
\ln B_t &= (1-\rho)\ln B + \rho \ln B_{t-1} + \epsilon_t^b, \qquad 0 \leq \rho < 1, \quad \epsilon_t^b \sim \mathcal{N}(0, \sigma_b^2), && (16)
\end{aligned}$$

and the initial conditions $K_t \geq 0$, $X_t > 0$, $A_{t-1} > 0$, $B_{t-1} > 0$.

In this model, current utility depends not only on current hours worked, now denoted by N_t, but also on a habit stock, X_t, determined by past hours worked (or, equivalently, by past leisure).[4] Specifically, the habit stock evolves according to the equation of motion (14), where $\phi - \mu$ measures

[3] By constructing and analyzing residuals from first-order conditions for labour, Johri and Letendre (2004) show that static first-order conditions are inconsistent with U.S. data, and that dynamic terms are needed to explain the data successfully.

[4] Earlier papers that allow preferences to depend on past labour/leisure include those by Kydland and Prescott (1982), Eichenbaum, Hansen, and Singleton (1988), Hotz, Kydland, and Sedlacek (1988), Yun (1996), Wen (1998), and Lettau and Uhlig (2000).

the persistence of habits. When $\phi = \mu$, only hours worked in the previous period affect the habit stock, whereas with $\phi > \mu$, the latter is a moving average of hours worked in the past. Notice that past hours of work increase current utility, and only hours worked over and above the habit stock effectively decrease utility. That is, past and current labour supplies are complements. This specification implies that habit-forming agents dislike large swings in their hours worked, and are more willing to smooth the path of their labour supply than agents with time-separable preferences. As a result, habits in labour will typically lead to a sluggish adjustment of labour in response to shocks. Thus, habit formation constitutes an alternative propagation mechanism that is different in nature from the LBD mechanism.

To see this, note that the first-order condition for the optimal choice of hours in the habit model is

$$B_t \left(N_t/X_t\right)^{1+1/v} - \beta E_t \left[\phi B_{t+1} \left(N_{t+1}/X_{t+1}\right)^{1+1/v} - \alpha(\phi - \mu) \left(Y_{t+1}/C_{t+1}\right)\right] = \alpha \left(Y_t/C_t\right). \qquad (17)$$

This condition equates the marginal disutility of work to the marginal utility of consuming the labour share of national income. Since past hours worked affect current utility through the habit stock, the marginal disutility of labour is no longer static, as it is in the RBC model. Unlike the corresponding equation in the LBD model, which involves only current and expected future hours of work, equation (17) depends in addition on past hours worked. Hence, the equilibrium path of hours worked will be different across the LBD and habit models. But, more importantly, equation (17) shows that the habit specification leads to richer labour dynamics and is likely to generate more persistence.

It is straightforward to show that, for given values of the structural parameters, the LBD and habit models imply the same equilibrium paths of output, consumption, capital, and investment. As discussed above, however, the equilibrium path of hours will be different across the two models. This implies that stochastic variations in output, consumption, or investment provide no useful information in distinguishing between the two models. From the viewpoint of classical statistics, the two models are observationally equivalent with respect to these variables. That is, they yield the same likelihood when evaluated using output, consumption, or investment data. Only when a

measure of hours worked (or labour productivity) is included in the information set of an econometrician are the two models distinguishable.

3. Counterfactual Experiments

To gain some insights into how the behaviour of hours worked differs across the habit and LBD models, it is useful to compare the impulse-response functions generated by the two models under the same calibration of the structural parameters. As CGS show, a model augmented with LBD captures output dynamics significantly better than does the standard RBC model. The LBD mechanism, however, only marginally improves the performance of the standard RBC model in matching the dynamic pattern of hours worked. In particular, the LBD model fails to replicate the hump-shaped response of hours worked to a transitory or a permanent shock, as is typically found in the vector autoregression (VAR) literature. The purpose of this section is to investigate whether, for a given parameterization, the habit model improves upon the LBD model in replicating the actual responses of hours worked.

To generate impulse-response functions from the LBD and habit models, we calibrate the structural parameters according to the posterior means reported by CGS (Table 2, p. 1507). We also compute the impulse-response functions implied by the standard RBC model (by setting $\phi = \mu = 0$), and by a benchmark bivariate VAR.[5] The permanent and the transitory shocks in the VAR are identified using the method developed by Blanchard and Quah (1989).

The top panels of Figure 1 depict the impulse responses of output and hours worked to a one-standard-deviation transitory (preference) shock. The figure shows that a positive preference shock (a decrease in B_t) triggers an increase in hours worked. Intuitively, a decrease in B_t lowers the weight of leisure in the utility function, thus inducing the representative household to work more, *ceteris paribus*. Because capital is predetermined, the rise in hours worked translates into a rise in output on impact. Unlike the standard RBC model, where output rises initially but decays monotonically in subsequent periods, the LBD and habit models generate a hump-shaped output response similar to that obtained from the VAR.[6] As discussed above, however, the LBD model,

[5] As in CGS, we use a fourth-order VAR with "Minnesota prior" as a benchmark.

[6] The LBD and habit models generate identical output responses when they are similarly calibrated because they

as well as the standard RBC model, fails to reproduce the hump-shaped response of hours worked predicted by the VAR. In contrast, the habit model is capable of generating an empirically plausible response that exhibits a hump-shaped pattern.

The bottom panels of Figure 1 show the impulse responses to a one-standard-deviation permanent shock. Unlike the VAR-based response, whereby output "overshoots," rising more in the short run than in the long run, the standard RBC, LBD, and habit models predict that output converges monotonically to its new steady-state level. The three models, however, diverge regarding the response of hours worked: while the standard RBC and LBD models generate a monotonic response, the habit model predicts a hump-shaped response that better matches the VAR impulse-response function.

These results raise the obvious question: Why, for given values of the parameters, does the habit model generate hump-shaped responses of hours worked, and the LBD model fail to do so? To answer this question, it is useful to remember that hours of work in the habit and LBD models are related through the identity $N_t = X_t H_t$. When a shock hits the economy, H_t and N_t increase by the same amount on impact. The variable X_t, on the other hand, is predetermined and therefore remains unaffected. In the next period, however, X_t increases by μ times the initial rise in H_t (or N_t). Thus, the response of N_t in the period following the shock will be higher than the corresponding response of H_t. With the specific calibration of the parameters ϕ and μ used in this exercise, H_t decreases monotonically in subsequent periods. But, for a few periods following the shock, the decay in H_t is actually slower than the rise in X_t, which leads to a hump-shaped response of N_t.

An intuitive explanation of why habit formation generates more persistence in hours worked than learning-by-doing is as follows: Consider a transitory shock that hits the economy at time t and causes hours to increase in current and subsequent periods before returning to their steady-state level. In the LBD model, such a shock has two conflicting effects on hours worked in period $t+1$. The first is a substitution effect that causes hours to increase (and leisure to decrease) to benefit from the favourable shock. The second is a wealth effect that arises from the higher labour

are observationally equivalent in this case.

productivity induced by the higher skill level resulting from the increase in hours worked in period t. This effect tends to decrease hours in period $t+1$. In the habit model, on the other hand, the shock increases hours worked in period t and thus the stock of habit in period $t+1$. Since habit formation implies that past and current labour supplies are complements, hours worked in period $t+1$ will increase more than in the RBC or LBD model. In addition, this substitution effect is not partially offset by any wealth effect, as it is in the LBD model.[7]

Because the above analysis is based on a counterfactual calibration of the habit model, one cannot conclude that the latter fits the data better than the LBD model. It does, nonetheless, suggest that habit formation is a stronger propagation mechanism than learning-by-doing. To assess the empirical plausibility of the habit model and whether it is more strongly supported by the data than the LBD model, we estimate both models using the same data and perform formal statistical comparisons.

4. Econometric Analysis

In this section, we estimate and statistically evaluate the habit and LBD models following the approach described in CGS. Denote these two models by $\mathcal{M}_i$, $i = 1, 2$, and let θ_i be the vector of their structural parameters. The estimation procedure consists in the following steps. From the state-space representation of each model, $\mathcal{M}_i$, and given the time-series data, $\mathbf{Y}$, we construct the likelihood function $p(\mathbf{Y}|\theta_i, \mathcal{M}_i)$ recursively using the Kalman filter.[8] For each model, the likelihood function is combined with a prior distribution, $p(\theta_i)$, to obtain a posterior distribution of the structural parameters, $p(\theta_i|\mathbf{Y}, \mathcal{M}_i)$. Bayesian estimates (i.e., posterior means of the structural parameters) are computed by generating random draws from the posterior distribution using a random-walk Metropolis-Hastings algorithm.

CGS estimate their LBD model using U.S. data on output growth and hours worked. To obtain priors for the parameters ϕ and μ, they resort to evidence from micro-level panel data. More

[7] In other words, in the context of a decentralized economy, the skill level shifts both labour supply and labour demand in the LBD model, whereas the habit stock affects only labour supply in the habit model.

[8] We obtain the state-space representation of each model by taking a log-linear approximation of the stochastically detrended equilibrium path around a deterministic steady state.

specifically, they use Panel Study of Income Dynamics (PSID) data to estimate a wage equation in which the current wage depends on the market wage rate for the efficiency unit of labour, past hours of work, and other control variables, such as age and schooling. For the remaining parameters, the priors reflect standard values used in the RBC literature. To make our results comparable with those reported by CGS, we estimate the LBD and habit models using the same data and prior distributions, which we summarize in Table 1.[9]

Table 2 reports the posterior means and standard deviations of the structural parameters for the LBD and habit models. Posterior estimates are very similar across the two models. The only exception is the autocorrelation of the preference shock, which is found to be slightly higher under the LBD specification. This could be viewed as an additional indication that habit formation is a stronger propagation mechanism than learning-by-doing, since the habit model requires less-persistent exogenous shocks to fit the data.

In what follows, we formally assess the ability of the LBD and habit models to explain the data. As in CGS, we use a bivariate VAR as a reference model, which is denoted by $\mathcal{M}_0$ with parameters θ_0. We start by evaluating the overall statistical fit of each model. If the fit of the LBD and habit models is poor (due to a potential misspecification), the VAR can be used as a benchmark to derive posterior estimates of the moments that we are interested in. The predicted moments obtained from the two competing models can then be compared with their VAR-based counterparts.

4.1 Overall statistical fit

We evaluate the overall statistical fit of model $\mathcal{M}_i$ by computing its posterior model probability $\pi_{i,T}$ given by

$$\pi_{i,T} = \frac{\pi_{i,0} p(\mathbf{Y}|\mathcal{M}_i)}{\sum_{i=0}^{2} \pi_{i,0} p(\mathbf{Y}|\mathcal{M}_i)},$$

where $\pi_{i,0}$ and $p(\mathbf{Y}|\mathcal{M}_i) = \int p(\mathbf{Y}|\theta_i, \mathcal{M}_i) p(\theta_i) d\theta_i$ are, respectively, the prior probability and the marginal data density of model $\mathcal{M}_i$. We then calculate the posterior odds ratio of model $\mathcal{M}_i$ versus the LBD model, $\pi_{i,T}/\pi_{LBD,T}$. A value larger than 1 indicates that model $\mathcal{M}_i$ matches the data

[9] We thank Yongsung Chang, Joao Gomes, and Frank Schorfheide for making their data and Gauss codes available at <http://www.econ.upenn.edu/~schorf/research.htm>. Their paper provides a detailed description of the data and a discussion of their priors.

better than the LBD model, and vice versa.

Table 3 reports the marginal data densities and posterior model probabilities of the LBD, habit, and VAR models. The marginal data densities of the LBD and habit models are computed using Geweke's (1999) modified harmonic-mean estimator, whereas the marginal data density of the VAR is computed by Monte Carlo approximation of one-step-ahead predictive densities (see CGS for details). Table 3 shows that the marginal data density of the habit model is higher than that of the LBD model. In computing the posterior model probabilities, equal prior probabilities (of 1/3) were assigned to the LBD, habit, and VAR models. Posterior odds ratios indicate that the habit model is favoured by a factor of 45,510 to 1, suggesting strong evidence in favour of the habit specification. The fit of both models, however, is much worse than that of the VAR, as reflected by the posterior probability of the latter. Thus, as discussed above, the VAR will serve as a benchmark to which the predictions of the LBD and habit models will be compared. We focus on two sets of predicted moments: impulse-response functions and autocorrelations.

4.2 Impulse-response functions

The top panels of Figure 2 depict the posterior means of the impulse responses to a one-standard-deviation transitory shock, generated by the LBD, habit, and VAR models. These responses are computed by averaging over random draws from the posterior distributions of the impulse-response functions. To allow statistical comparisons across models, 75 per cent Bayesian highest-posterior-density confidence bands from the VAR are also plotted. Figure 2 shows that both the LBD and habit models generate hump-shaped output responses similar to that obtained from the VAR. It is obvious, however, that the output response predicted by the habit model is (statistically) closer to that predicted by the VAR. On the other hand, only the habit model is capable of replicating the hump-shaped response of hours worked predicted by the VAR. In fact, the response obtained from the habit model is not statistically different from the VAR impulse-response function, since it lies mostly within the VAR's confidence band.

Impulse responses to a permanent shock are shown in the bottom panels of Figure 2. Both the LBD and habit models fail to explain the output overshooting observed in the VAR. The habit

model, however, is more successful in matching the hump-shaped VAR response of hours worked. In contrast, the LBD model yields a monotonic decay of hours worked.

This analysis clearly shows that the habit model is more successful in replicating the VAR impulse-response functions than the LBD model. To the extent that the VAR provides a reasonable characterization of the time-series properties of the data, our results suggest that the habit model fits the data better than the LBD model.

4.3 Autocorrelations

Another important dimension along which the standard RBC model fails to match the data is the autocorrelation of output growth. While this variable is found to be positively autocorrelated over short horizons in the data, Cogley and Nason (1995) show that the standard RBC model predicts zero autocorrelations at all horizons. CGS show that the LBD model improves upon the RBC model in generating positive serial correlations of output growth. This result is illustrated in the top panel of Table 4, which shows the autocorrelations of output growth predicted by the LBD, habit, and VAR models up to 4 lags. Table 4 shows that the habit model is also capable of generating positive autocorrelations of output growth, though to a lesser extent than the LBD model.

A formal evaluation of the ability of the two models to match the autocorrelation function of the data is based on the posterior expected loss (risk). Let φ denote the population autocorrelations (i.e., those obtained from the benchmark VAR), and $\hat{\varphi}_i$ the predictions of model $\mathcal{M}_i$, $i = 1$, 2. Given the posterior distribution of φ conditional on the VAR, $p(\varphi|\mathbf{Y}, \mathcal{M}_0)$, the posterior risk associated with model $\mathcal{M}_i$ is given by $R(\hat{\varphi}_i|\mathbf{Y}, \mathcal{M}_0) = \int L(\varphi, \hat{\varphi}_i)p(\varphi|\mathbf{Y}, \mathcal{M}_0)d\varphi$, where $L(\varphi, \hat{\varphi}_i)$ is a loss function. Following CGS, our results are based on two loss functions, L_q and L_{χ^2}.[10]

Both measures of loss reported in Table 4 confirm that the LBD model does marginally better than the habit model in explaining output-growth autocorrelations.

The bottom panel of Table 4 reports the autocorrelations of hours worked predicted by the three models. A lag-by-lag comparison indicates that both the LBD and habit models are successful in

[10]See Schorfheide (2000) for a detailed discussion of these loss functions and their interpretations.

replicating the sample autocorrelations, but the loss statistics suggest that the latter performs significantly better than the former.

5. Robustness Analysis

The results discussed in the previous section are conditional on using tight priors for the parameters ϕ and μ. As stated earlier, to obtain these priors, CGS estimate a micro-level wage equation that links the current wage rate to the market wage rate for the efficiency unit of labour, and to past hours of work. This equation is the micro counterpart of the wage equation in the LBD model, where the skill level raises the marginal product of labour and therefore the wage rate. More precisely, in the LBD model, a worker with skill X_t earns a competitive wage rate of

$$W_t(X_t) = W_t^* X_t,$$

where W_t^* is the market wage rate for the efficiency unit of labour. In the habit model, the habit stock, X_t, does not affect directly the marginal product of labour (the wage rate). That is, the competitive wage rate is always equal to the market wage rate for the efficiency unit of labour.[11] For this reason, the priors obtained from the micro-level data, albeit sensible for the LBD model, might not be appropriate for the habit model. To address this issue, we repeat the analysis carried out in section 4 using non-informative priors for the parameters ϕ and μ. Under non-informative priors, the estimation procedure ignores the microeconomic evidence and consists in selecting the parameters ϕ and μ that maximize the likelihood function of the aggregate models. To obtain non-informative priors for ϕ and μ, we scale their covariance matrix by a factor of 10^6.[12]

Table 5 reports posterior estimates under the non-informative priors. The posterior means of ϕ and μ are now quite different across the LBD and habit models. In both cases, however, the estimate of ϕ is lower and that of μ is higher than what we obtain under the tight priors (Table 2). The posterior odds ratio, reported in Table 6, indicates that the habit model is favoured by a

[11] This, of course, does not imply that past hours of work do not affect the competitive wage rate in equilibrium. It simply means that, for any strictly positive value of X_t, the competitive wage rate will be lower in the habit model than in the LBD model, *ceteris paribus*.

[12] Alternatively, one could consider flat priors by assuming a uniform prior distribution. The parameter μ, however, is defined on the positive real axis, implying that the flat prior would be improper.

factor of 3.6×10^5 to 1 in this case. Thus, the overall statistical fit of the habit model is still much better than that of the LBD model under non-informative priors for the parameters ϕ and μ.

Impulse-response functions to a transitory shock, illustrated in the top panels of Figure 3, show that the LBD model generates a small one-period hump in the response of hours worked. The habit model, on the other hand, produces a more pronounced hump that reaches its peak two periods after the shock, exactly as predicted by the VAR. In addition, the output response generated by the habit model tracks its VAR-based counterpart much more closely than that implied by the LBD model. In response to a permanent shock, the bottom panels of Figure 3 show that, under the non-informative priors, the habit model predicts that output overshoots, rising slightly more in the short run than in the long run, whereas the LBD model still implies a monotonic convergence of output towards its new steady-state level. The two models, however, fail to match the shape and the magnitude of the VAR response, and no clear conclusion can be drawn as to which model performs better in this dimension. The habit model does perform better in replicating the VAR impulse response of hours worked to a permanent shock. Interestingly, Figure 3 indicates that the matching performance of the habit model is better under non-informative priors for the parameters ϕ and μ than under tight priors. Thus, we can conclude that results based on impulse-response analysis are robust, and, if anything, better when we consider non-informative priors for the parameters ϕ and μ.

The final robustness check is related to the autocorrelation functions. The results reported in Table 7 indicate that the LBD and habit models are even more successful in replicating the autocorrelation functions of output and hours worked under the non-informative priors. In comparison to the case with tight priors, a similar conclusion is reached regarding the relative performance of the two models: the LBD model does marginally better in matching the autocorrelations of output growth, but the habit model performs significantly better in matching those of hours worked.

In summary, our result that the habit model fits the data better than the LBD model appears to be robust to the use of non-informative priors for the parameters ϕ and μ.[13] It should be emphasized, however, that all our findings are conditional on the information set used in the analysis. To identify

[13] We reach the same conclusion when we use diffuse rather than non-informative priors for ϕ and μ.

the LBD and habit models, we have used data on aggregate hours worked in the United States. Alternatively, one could use aggregate data on wages to discriminate between the two models, since they have different implications for the equilibrium wage rate. We leave this for future research.

6. Conclusion

In an important paper, CGS augment the standard RBC model with a learning-by-doing mechanism and show that this feature provides an important propagation mechanism that can help the standard RBC model explain the persistence of aggregate U.S. output.

In this paper, we have shown that the LBD model is nearly observationally equivalent to an RBC model with habit formation in labour (or, equivalently, in leisure). Under a coefficient restriction, the LBD and habit models deliver identical predictions for output, consumption, and investment, and are therefore indistinguishable using data on these quantities. The two models differ in implications, however, regarding hours worked. We have exploited this difference to discriminate between the two models. Using Bayesian techniques, we have found that habit formation is a stronger propagation mechanism and that the habit model fits aggregate U.S. data better than the LBD model proposed by CGS.

These findings suggest that time non-separable preferences are important in accounting for business cycle fluctuations, and that more effort should be devoted to refining the modelling assumptions of consumers' behaviour. Beyond the specific issue addressed in this paper, our study draws attention to the fact that dynamic stochastic general-equilibrium models may not be identifiable, and that an econometrician may draw a mistaken inference about the theory underlying a given model if their information set does not allow them to discriminate between observationally equivalent models. In this regard, our paper is related to recent work by Beyer and Farmer (2004), who discuss the lack of identification in a class of linear rational-expectation models, namely new-Keynesian monetary models.

References

Andolfatto, D. 1996. "Business Cycles and Labour Market Search." *American Economic Review* 86(1): 112–32.

Beyer, A. and R. Farmer. 2004. "On the Indeterminacy of New-Keynesian Economics." European Central Bank Working Paper No. 323.

Blanchard, O.J. and D. Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79(4): 655–73.

Burnside, C. and M.S. Eichenbaum. 1996. "Factor Hoarding and the Propagation of Business-Cycle Shocks." *American Economic Review* 86(5): 1154–74.

Chang, Y., J.F. Gomes, and F. Schorfheide. 2002. "Learning-By-Doing as a Propagation Mechanism." *American Economic Review* 92(5): 1498–1520.

Cogley, T. and J.M. Nason. 1995. "Output Dynamics in Real-Business-Cycle Models." *American Economic Review* 85(3): 492–511.

Cooper, R. and A. Johri. 2002. "Learning-By-Doing and Aggregate Fluctuations." *Journal of Monetary Economics* 49(8): 1539–66.

den Haan, W., G. Ramey, and J. Watson. 2000. "Job Destruction and Propagation of Shocks." *American Economic Review* 90(3): 482–98.

Eichenbaum, M.S., L.P. Hansen, and K.J. Singleton. 1988. "A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty." *Quarterly Journal of Economics* 103(1): 51–78.

Geweke, J.F. 1999. "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication." *Econometric Reviews* 18: 1–126.

Hotz, V.J., F.E. Kydland, and G.L. Sedlacek. 1988. "Intertemporal Preferences and Labour Supply." *Econometrica* 56(2): 335–60.

Johri, A. and M.-A. Letendre. 2004. "What Do Residuals from First-Order Conditions Reveal about DGE Models?" McMaster University. Photocopy.

Kydland, F.E. and E.C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50(6): 1345–71.

Lettau, M. and H. Uhlig. 2000. "Can Habit Formation Be Reconciled with Business Cycle Facts?" *Review of Economic Dynamics* 3(1): 79–99.

Rouwenhorst, K.G. 1991. "Time to Build and Aggregate Fluctuations: A Reconsideration." *American Economic Review* 27(2): 241–54.

Schorfheide, F. 2000. "Loss Function-Based Evaluation of DSGE Models." *Journal of Applied Econometrics* 15(6): 645–70.

Wen, Y. 1998. "Can a Real Business Cycle Model Pass the Watson Test?" *Journal of Monetary Economics* 42(1): 185–203.

Yun, T. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37(2-3): 345–70.

Table 1: Prior Distributions for the Structural Parameters

Parameter	Range	Density	Mean	S.D.
μ	R^+	Truncated Bivariate Normal	0.111	0.004
ϕ	$[0,1]$		0.798	0.012
α	$[0,1]$	Beta	0.660	0.020
β	$[0,1]$	Beta	0.993	0.002
γ	R	Normal	0.005	0.005
δ	$[0,1]$	Beta	0.025	0.005
ν	R^+	Gamma	2.000	0.500
ρ	$[0,1]$	Beta	0.800	0.100
σ_a	R^+	Inverse gamma	N/A	N/A
σ_b	R^+	Inverse gamma	N/A	N/A

Notes: S.D. is standard deviation. N/A = not available.

Source: Chang, Gomes, and Schorfheide (2002, p. 1507).

Table 2: Estimation Results

	LBD		Habit	
Parameter	Mean	S.D.	Mean	S.D.
μ	0.1107	0.0039	0.1115	0.0038
ϕ	0.7968	0.0120	0.7976	0.0116
α	0.6524	0.0200	0.6590	0.0198
β	0.9936	0.0018	0.9933	0.0019
γ	0.0040	0.0009	0.0041	0.0009
δ	0.0224	0.0045	0.0231	0.0045
ν	1.5293	0.3583	1.5986	0.3739
ρ	0.9474	0.0314	0.8687	0.0437
σ_a	0.0118	0.0008	0.0116	0.0008
σ_b	0.0087	0.0013	0.0075	0.0011

Note: S.D. is standard deviation.

Table 3: Goodness of Fit

Statistic	LBD	Habit	VAR(4)
Prior probability, $\pi_{i,0}$	1/3	1/3	1/3
Marginal data density, $\ln p(\mathbf{Y}\|\mathcal{M}_i)$	1055.41	1066.14	1082.69
Posterior probability, $\pi_{i,T}$	0.00	0.00	1.00
Posterior odds ratio, $\pi_{i,T}/\pi_{LBD,T}$	1.00	45510.72	7.02×10^{11}

Table 4: Autocorrelations

Statistic	Lag	LBD	Habit	VAR (mean)	VAR (band)
			Output growth		
Autocorrelation	1	0.0634	0.0480	0.3097	[0.1553, 0.4663]
	2	0.0526	0.0382	0.1578	[0.0153, 0.3071]
	3	0.0435	0.0302	0.0297	[−0.1068, 0.1723]
	4	0.0358	0.0237	−0.0269	[−0.1563, 0.1051]
L_q risk	$1-4$	0.0758	0.0854		
L_{χ^2} risk	$1-4$	0.9686	0.9758		
			(Log) Hours worked		
Autocorrelation	1	0.9558	0.9533	0.9575	[0.9319, 0.9833]
	2	0.9113	0.9003	0.8864	[0.8200, 0.9548]
	3	0.8671	0.8438	0.8008	[0.6855, 0.9195]
	4	0.8237	0.7857	0.7104	[0.5462, 0.8831]
L_q risk	$1-4$	0.0179	0.0077		
L_{χ^2} risk	$1-4$	0.9865	0.9684		

Table 5: Estimation Results (Non-Informative Priors)

	LBD		Habit	
Parameter	Mean	S.D.	Mean	S.D.
μ	0.2921	0.1130	0.3797	0.0847
ϕ	0.2335	0.1669	0.6975	0.0896
α	0.6514	0.0200	0.6597	0.0195
β	0.9938	0.0017	0.9933	0.0019
γ	0.0040	0.0009	0.0040	0.0009
δ	0.0220	0.0050	0.0249	0.0048
ν	1.3952	0.3993	1.2492	0.3501
ρ	0.9412	0.0259	0.6852	0.0896
σ_a	0.0117	0.0008	0.0116	0.0008
σ_b	0.0085	0.0018	0.0079	0.0013

Note: S.D. is standard deviation.

Table 6: Goodness of Fit (Non-Informative Priors)

Statistic	LBD	Habit	VAR(4)
Prior probability, $\pi_{i,0}$	1/3	1/3	1/3
Marginal data density, $\ln p(\mathbf{Y}\|\mathcal{M}_i)$	1052.02	1064.83	1082.69
Posterior probability, $\pi_{i,T}$	0.00	0.00	1.00
Posterior odds ratio, $\pi_{i,T}/\pi_{LBD,T}$	1.00	3.62×10^5	2.08×10^{13}

Table 7: Autocorrelations (Non-Informative Priors)

Statistic	Lag	LBD	Habit	VAR (mean)	VAR (band)
			Output growth		
Autocorrelation	1	0.1382	0.1283	0.3097	[0.1553, 0.4663]
	2	0.0467	0.0858	0.1578	[0.0153, 0.3071]
	3	0.0176	0.0574	0.0297	[−0.1068, 0.1723]
	4	0.0068	0.0380	−0.0269	[−0.1563, 0.1051]
L_q risk	$1-4$	0.0430	0.0431		
L_{χ^2} risk	$1-4$	0.7987	0.8674		
			(Log) Hours worked		
Autocorrelation	1	0.9486	0.9615	0.9575	[0.9319, 0.9833]
	2	0.8895	0.9007	0.8864	[0.8200, 0.9548]
	3	0.8307	0.8296	0.8008	[0.6855, 0.9195]
	4	0.7747	0.7554	0.7104	[0.5462, 0.8831]
L_q risk	$1-4$	0.0051	0.0031		
L_{χ^2} risk	$1-4$	0.9759	0.0351		

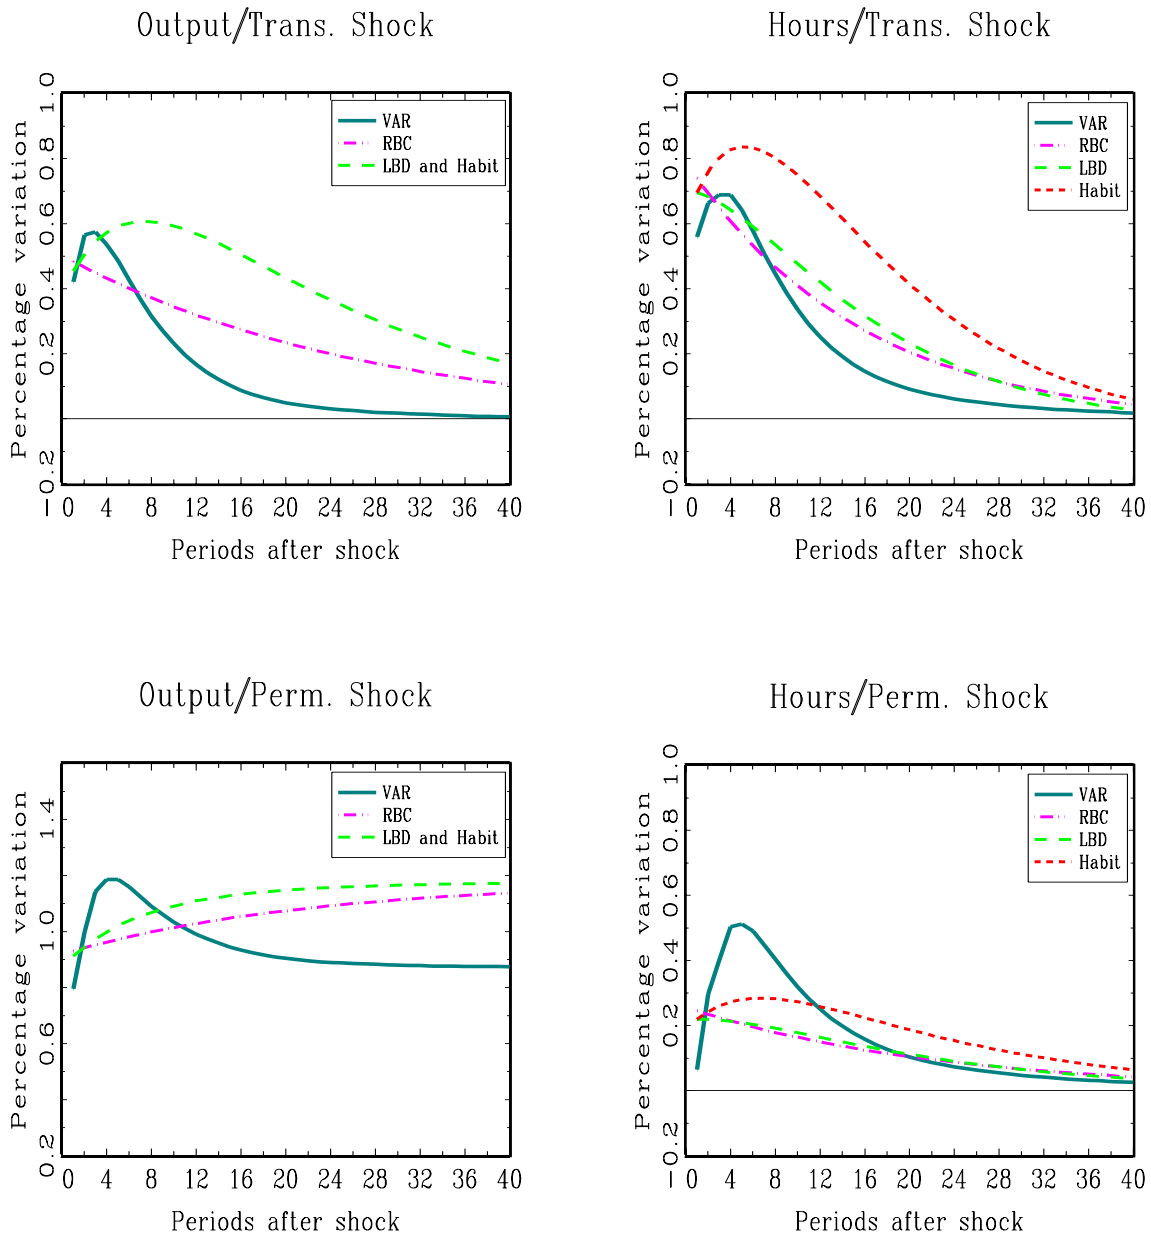

Figure 1: Impulse-Response Functions: Counterfactual Calibration

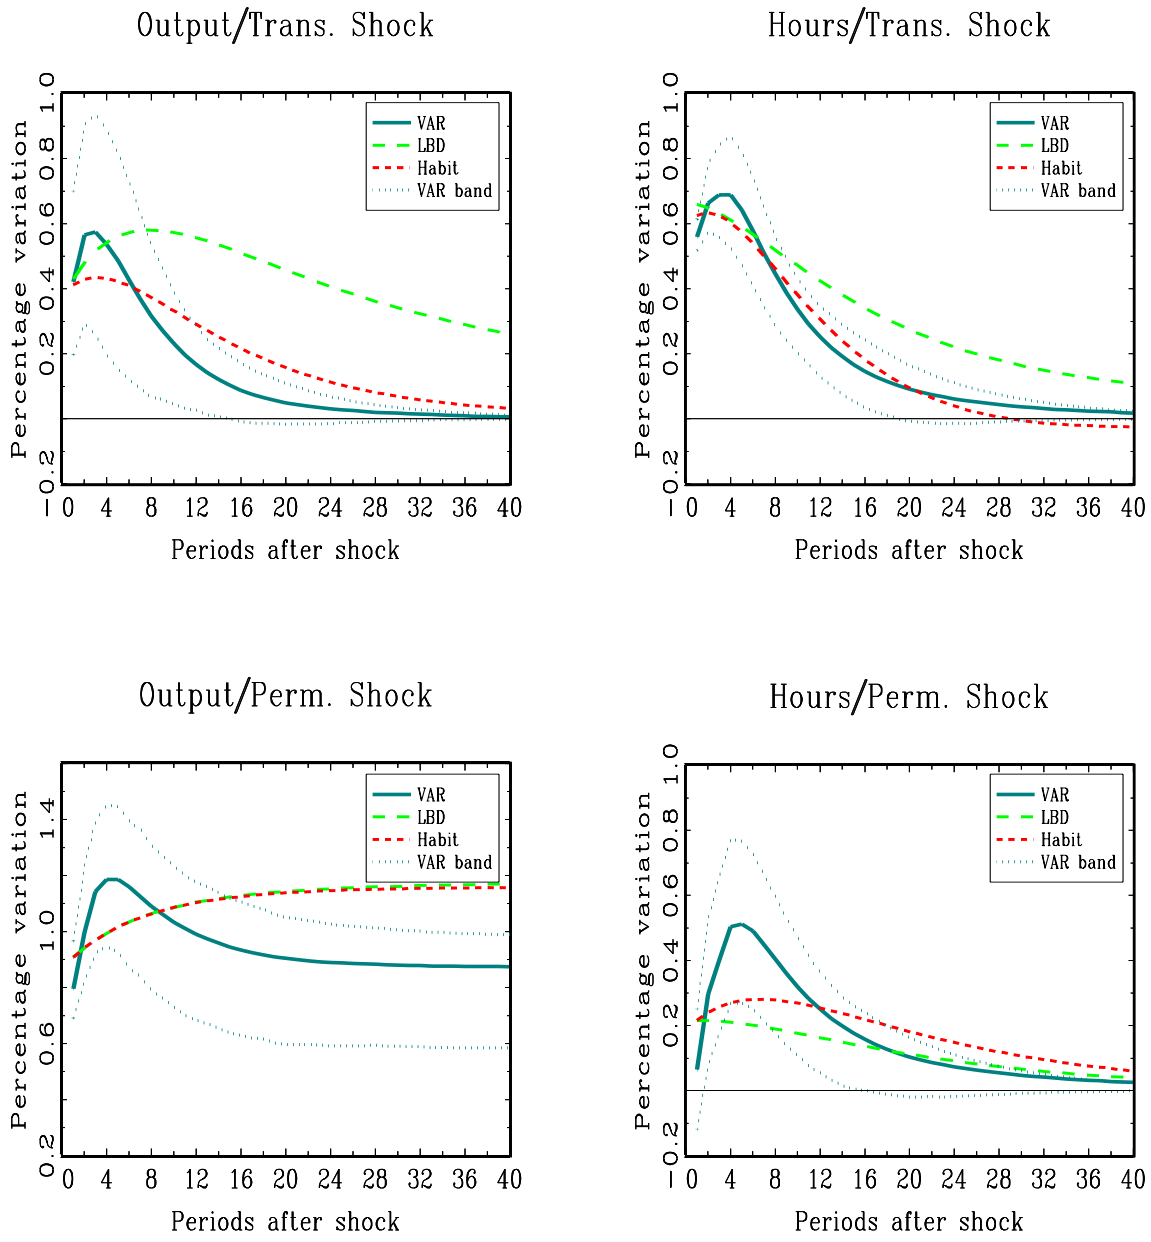

Figure 2: Impulse-Response Functions: Posterior Means

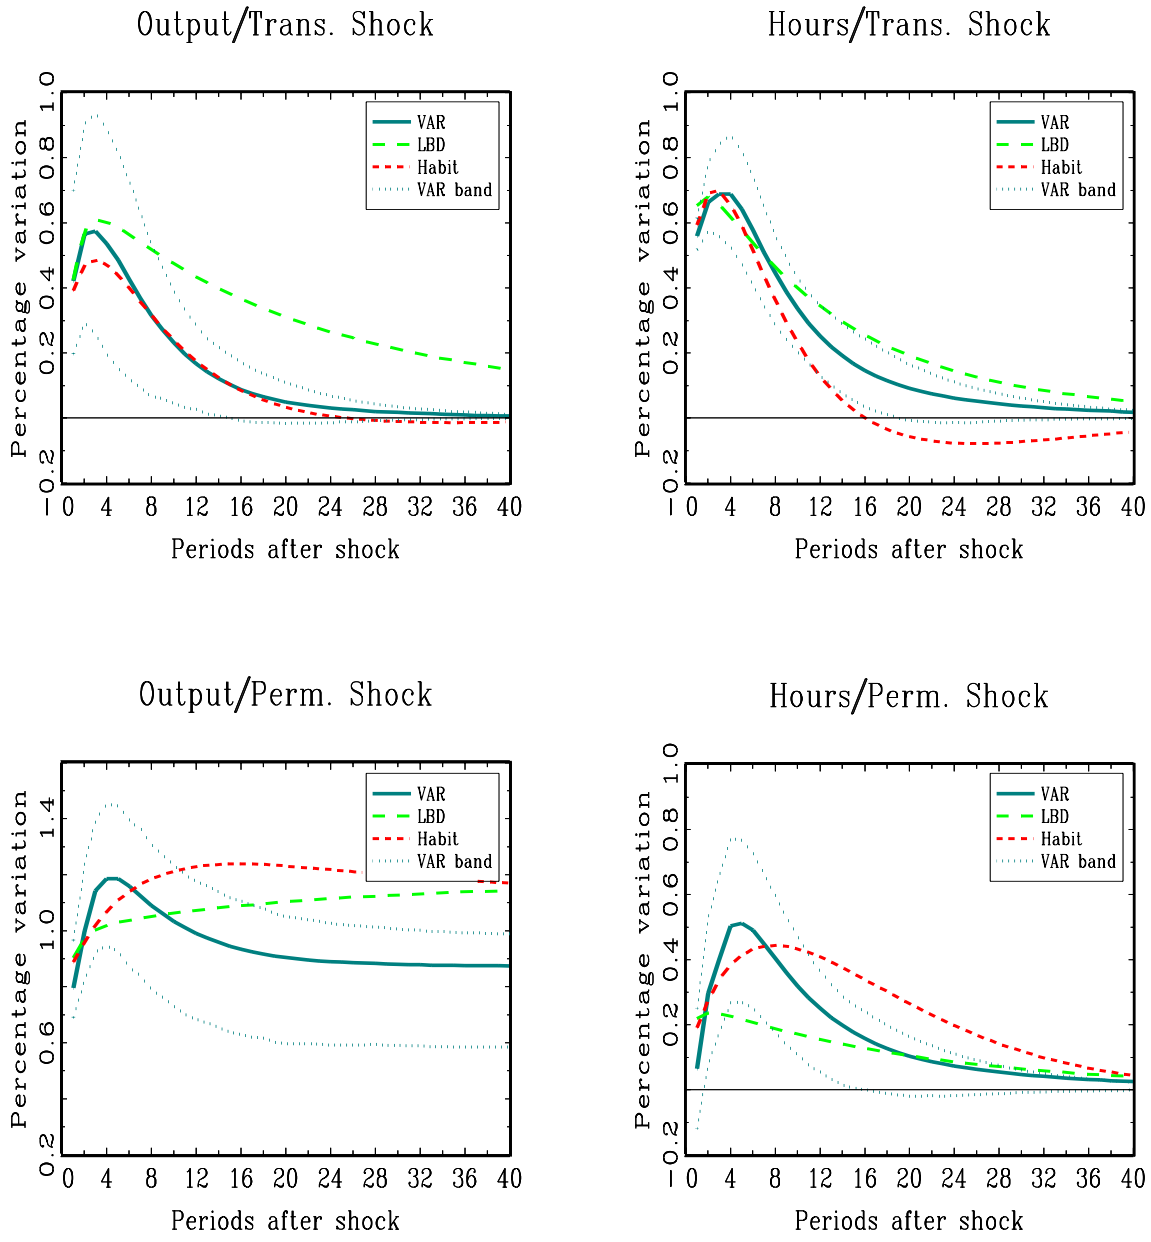

Figure 3: Impulse-Response Functions: Posterior Means (non-informative priors)