

Financial Crises and Total Factor Productivity

Felipe Meza

Erwan Quintin

Universidad Carlos III de Madrid

Federal Reserve Bank of Dallas*

March 22, 2005

Abstract

Total factor productivity (TFP) falls markedly during financial crises, as we document with recent evidence from Mexico and Asia. These falls are unusual in magnitude and present a difficult challenge for the standard small open economy neoclassical model. We show in the case of Mexico's 1994-95 crisis that the model predicts that inputs and output should have fallen much more than they did. Using models with endogenous factor utilization, we find that capital utilization and labor hoarding can account for a large fraction of the TFP fall during the crisis. However, these models also predict that output should fall significantly more than in the data. Given the behavior of TFP, the biggest challenge may not be explaining why output falls so much following financial crises, but rather why it falls so little.

Keywords: Financial crises; total factor productivity; output fluctuations

JEL classification: E32; F41; J24

*Corresponding author. Felipe Meza: Universidad Carlos III de Madrid, Calle Madrid 126, Getafe, Madrid 28903, Spain, +34 91 624 57 34 (phone), +34 91 624 98 75 (fax), fmeza@eco.uc3m.es. Erwan Quintin: Research Department, Federal Reserve Bank of Dallas, 2200 N. Pearl St. Dallas, TX 75201, +1 214 922 51 57 (phone), erwan.quintin@dal.frb.org. Felipe Meza thanks the Ministerio de Educación y Ciencia de España for financial support through project SEJ2004-00968. We thank seminar participants at the European Central Bank, Rice University, the University of Texas at Austin, ES European Meeting Madrid 2004, ASSET Meeting Barcelona 2004 and the Banco de México for their comments. The views expressed herein are those of the authors and may not reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 Introduction

Output falls drastically following financial crises, much more than what the behavior of capital and labor would lead one to expect. In the language of growth accounting, total factor productivity (TFP) falls markedly during financial crises, as we document with recent evidence from Mexico and East Asia. The magnitude of these falls is unusual: TFP falls by more than two standard deviations in all the cases we study. These declines are also puzzling because given their size, a standard small open economy model would predict that hours worked, capital and output should fall much more than they do in the data. Our goal in this paper is twofold: document the unusual behavior of TFP during crises, and describe the challenge that this behavior poses for standard small open economy neoclassical models.

Most of the existing literature focuses on what triggers a financial crisis in the first place. For instance, in the case of Mexico's 1994-95 "Tequila" crisis, Flood, Garber, and Kramer (1996) and Calvo and Mendoza (1996) study the role played by flow imbalances (liquid financial assets vs. broad monetary aggregates, and short-run debt vs. gross foreign reserves). Cole and Kehoe (1996) and Sachs, Tornell and Velasco (1996) conjecture that Mexico's large stock of short-term debt may have given rise to a self-fulfilling debt crisis. These and many related articles have shed light on what causes financial collapses in nations like Mexico, but they do not try to account for the behavior of output after the collapse. Like Calvo (2000) and despite some exceptions which we review below, our assessment is that there has been little emphasis on the deep consequences of crises on real activity. This paper contributes to filling this gap.

We study the real impact of financial crises in the context of various versions of the standard small open economy neoclassical model (as articulated for instance by Mendoza, 1991) and concentrate our attention on Mexico's Tequila crisis. While financial crises share many characteristics, a satisfactory quantitative study of the real impact of these episodes must incorporate country-specific features. In the case of Mexico's Tequila episode, a number of deep policy shocks accompanied the crisis. In particular, the Mexican government sharply raised energy prices and the Value Added Tax (VAT) rate in the first quarter of 1995. We model those shocks explicitly and find that changes in consumption taxes and the price of

energy in 1995 had an important effect on output. However, the impact of these fiscal shocks is markedly smaller than the impact of TFP. Our key result is that given the size of the fall in TFP Mexico experienced in 1995, hours worked, capital and output should have fallen twice as much as they did, in percentage terms.

The fact that the behavior of factor series diverges so much from the predictions of the standard model suggests that factor hoarding plays a large role during financial crises. Intuitively, one should expect large swings in capital utilization and effort during crises. For several quarters, interest rates are well above average, while TFP is well below average. This gives firms strong incentives to postpone the consumption of capital services (say, by leaving plants or machines temporarily idle) and economize on variable expenditures such as wear and tear until conditions improve. Similarly, if employment is costly to adjust, firms may use the effort margin to respond to the fall in the marginal product of labor. Standard models of factor hoarding (Greenwood, Hercowitz and Huffman, 1988, and Burnside, Eichenbaum and Rebelo, 1993) suggest that capital utilization and labor hoarding can account for a large fraction of the variance of TFP both during and outside the crisis period in Mexico. Nevertheless, we also find that modeling factor hoarding explicitly may not resolve the output puzzle. The models we consider continue to predict that, following the crisis, output should fall much more than in the data.

Our calculations complement some related investigations of the real impact of financial crises by stressing the importance of TFP. Most existing studies ignore the unusual behavior of TFP.¹ Among the few exceptions, Burstein, Eichenbaum and Rebelo (2002) study the behavior of inflation and output in Mexico and South Korea in a model with credit market frictions and four sectors: local, export, tradable and non-tradable goods. The model pro-

¹Burnside, Eichenbaum and Rebelo (2001), Corsetti, Pesenti and Roubini (1999) and Lahiri and Vegh (2002) provide qualitative explanations for the contraction of output. Cavallo, Kisselev, Perri and Roubini (2004) show that large falls in output are possible after crises in sticky-price models with a margin constraint. Similarly, Cook and Devereux (2005) simulate recent crises in Malaysia, South Korea and Thailand and show that output can drop sharply following shocks to a country's risk premium. All these papers assume that TFP is constant. Mendoza (2002) shows that a flexible-price model with a liquidity constraint can lead to sudden stops of capital flows and large output falls. He allows for TFP fluctuations, but only of average business cycle size: the standard deviation of TFP fluctuations coincides with that of output. Chari, Kehoe and McGrattan (2005) show that sudden stops of capital flows induce an output increase, not a fall, in a standard neoclassical model. They argue that more research is needed to find a "friction" that can overwhelm the positive effect of a sudden stop. Our results suggest that explicitly modeling and measuring TFP should suffice.

duces a yearly fall in output that slightly exceeds the actual fall in Mexico in 1995 provided TFP falls by 50% in the export sector. They do not provide any evidence that such a fall took place in 1995. Gertler, Gilchrist and Natalucci (2003) simulate the impact of shocks to a country's risk premium in a model with price-stickiness and endogenous capital utilization. Their model can predict a fall in output in South Korea of a magnitude similar to the one observed. They argue that measured TFP varies because capital utilization does. However, their quantitative analysis does not rely on measures of TFP that are consistent with the model: adjusted TFP remains constant during experiments. Our calculations suggest that in the Mexican case, capital utilization accounts for less than 30% of the fall in TFP. Similar calculations could be carried out for Asian nations.²

Given our findings, quantitative studies of the real impact of financial crises should take the behavior of TFP into account. Given the precipitous fall in TFP that follows crises episodes, the most puzzling aspect of financial crises may not be that output falls a lot, but rather that it falls too little.

2 Evidence

In this section we document the fact that financial crises are followed by unusually large falls in TFP and GDP using evidence from Mexico's 1994 crisis, and from the 1997 crisis in South Korea and Thailand. In addition, we find that these falls are persistent. Both GDP per capita and TFP remain below trend for several years after the crisis.

To measure TFP, we use the following specification of aggregate technological opportunities:

$$Y_t = A_t K_t^\alpha L_t^{1-\alpha},$$

where Y_t denotes GDP at date t , K_t is aggregate capital, L_t denotes aggregate hours worked and $\alpha \in (0, 1)$ measures the importance of capital in production. We assume that A_t ,

²With the data we present in the next section and using the capital utilization model of Greenwood, Hercowitz and Huffman (1988), one can calculate the fall in TFP adjusted for capital utilization in South Korea and Thailand in 1998. We set $\phi = 1.5$ which, assuming long-term interest rates of 4%, implies a steady state rate of depreciation of 8%. Given those parameters, capital utilization can account for only 36% and 17% of the fall of measured TFP in South Korea and Thailand, respectively.

aggregate TFP at date t , equals $z_t(1 + \gamma)^{t(1-\alpha)}$, where z_t is stationary and $\gamma \geq 0$ is an exogenous growth rate. Let y_t , k_t and l_t denote the per capita counterparts of Y_t , K_t and L_t , respectively. In the neoclassical growth model, per capita output and capital grow at constant rate γ along the balanced growth path, while per capita hours worked are constant. Letting \hat{y}_t and \hat{k}_t be detrended per capita output and capital, we have

$$\hat{y}_t = z_t \hat{k}_t^\alpha l_t^{1-\alpha}.$$

Measuring z_t requires empirical counterparts for \hat{y}_t , \hat{k}_t and l_t . We construct capital stock series using the perpetual inventory approach with geometric depreciation and yearly data from the International Financial Statistics (IFS) database (IMF, 2004). The IFS database reports nominal data on investment. Like Bergoeing, Kehoe, Kehoe and Soto (2002) we measure real investment as the ratio of nominal gross fixed capital formation to nominal GDP multiplied by real GDP. We assume that capital depreciates at a yearly depreciation rate of 8%. To set the initial capital stock, we follow Young (1995) and assume that the growth rate of investment in the first five years of the series is representative of the growth of investment in previous years. The resulting capital stock series begins in 1963 for Mexico and Thailand, and 1966 for South Korea. GDP series start in 1950 for Mexico and Thailand, and in 1953 for South Korea. For Mexico, we use total hours worked as measured in Bergoeing, Kehoe, Kehoe and Soto (2002). They report the product of total employment and average hours per worker in the manufacturing sector, as measured with data from a manufacturing sector survey. Calculations are similar for South Korea and Thailand except that an estimate of average hours worked is available for most sectors in those two countries.³ Labor series cover the periods 1980-2000, 1970-2002 and 1989-1999 for Mexico, South Korea and Thailand, respectively.

³For South Korea, we use data on total employment and average hours worked per week, as reported by the South Korean National Statistical Office. Total employment corresponds to employed individuals of age 15 and higher in all sectors. Average hours worked correspond to all industries, excluding agricultural activities. Data were downloaded from <http://www.nso.go.kr>. For Thailand, total employment corresponds to employed individuals of age 13 and higher in all sectors, as reported by the International Labour Office (ILO) and the Thai National Statistical Office. Average hours worked correspond to all industries, excluding agricultural activities and public administration, as reported by the ILO. Data were downloaded from <http://www.nso.go.th> and <http://laborsta.ilo.org>.

We calculate y_t , k_t , and l_t by dividing Y_t , K_t and L_t by the number of adults between ages 15 and 64.⁴ Population series for the three countries start in 1960. Detrended variables \hat{y}_t and \hat{k}_t are y_t and k_t divided by the average geometric growth factor of y_t in the period before the crisis episode. This factor is 1.7% for Mexico between 1960 and 1994, 5.3% for South Korea between 1960 and 1997, and 4.4% for Thailand between 1960 and 1997. Finally, we set $\alpha = 0.3$. Gollin (2002) finds that after distributing the income of the self-employed to capital and labor income, labor income shares do not vary much across countries and time, and take values around 70%.⁵

Figure 1 shows the resulting series for Mexico, Thailand and South Korea with vertical lines marking the onset of each crisis. Output falls by over 10% in all countries during the year following the crisis, and by as much as 15.6% in Thailand. Capital, on the other hand, remains practically constant after the crisis, and hours fall less than output in all cases. In fact, in Mexico and Thailand, hours worked fall by less than 2%. In South Korea, hours worked fall by a larger 8.6% in 1998. Since capital and labor fall relatively little during crises, TFP has to fall by a large amount to account for the fall in output: 15.1% in Thailand, 7.1% in South Korea and 8.6% in Mexico. The magnitude of these falls is very unusual for all countries. Yearly falls in GDP and TFP exceed two standard deviations. Also notice that the falls in output and TFP triggered by crises are quite persistent. They remain below trend in all cases for several years. Like us, Cook and Devereux (2005) find that output remains persistently below trend in Asia after 1997, using a different detrending procedure. They report that in Malaysia, South Korea and Thailand, output remains below trend for at least 8 quarters.

Naturally, these results could be sensitive to some of our measurement assumptions. Young (1995) argues for instance that data on changes in inventories are of poor quality in East Asia. Excluding changes in inventories from our calculations had negligible consequences on our results.⁶ Changing $(1 + \gamma)$ to 1.02 for all countries, the value Kehoe and

⁴We use population data for Mexico as reported by Bergoeing et al. (2002). For South Korea, data were downloaded from <http://www.nso.go.kr>. For Thailand, data were obtained from the World Bank Development Indicators database (World Bank, 2004).

⁵Young (1995) finds a value of $1 - \alpha = 0.703$ for South Korea with data from the 1966-1990 time period.

⁶Our TFP findings for South Korea can be compared to results in Young (1995). He reports that the average logarithmic annual growth rate of A_t in South Korea was 1.7% between 1966 and 1990. The main difference between his calculations and ours is that he takes into account changes in the quality of labor

Prescott (2002) propose,⁷ also has little effect on results, with one minor exception. In the case of South Korea, the effect of the 1997 crisis becomes less persistent as \hat{y}_t and z_t surpass their 1997 levels by 2000. Next, one can carry out all calculations using national sources of data for \hat{y}_t and \hat{k}_t instead of IMF data.⁸ National Income and Product Account series are much shorter because countries modify their systems of national accounts every now and again, which makes results more sensitive to the choice of initial capital. On the other hand, IMF data include only the most basic national accounts variables. Data from national sources are richer and allow one to construct more accurate empirical counterparts of theoretical variables, as we do in the next section. In particular, one can subtract indirect business taxes from GDP, impute the returns of government capital and of the stock of durable goods, and include public investment and durable goods purchases in investment (with different rates of depreciation for each type of capital.) After making those corrections, the behavior of detrended series changes little. It is still the case that the falls in \hat{y}_t and z_t after financial crises are unusually large.⁹

In summary, recent financial crises triggered unusually large falls in detrended GDP per capita and TFP in Mexico and East Asia. There is also some evidence that these falls are persistent. These facts beg several interesting questions. In the remainder of the paper, we study whether given the behavior of TFP a standard small open economy neoclassical model can account for the behavior of output and inputs in the Mexican case.

and capital. After excluding inventory changes as he does, we calculate that the average logarithmic annual growth rate of A_t for South Korea for the period 1970-1990 is 2.6%. The difference is large and is due to Young's adjustment for quality. Assuming a labor income share of 70% and using Young's data on raw inputs, we find that A_t in South Korea grew at an average rate of 2.7% between 1970 and 1990. In other words, our measurement of TFP leads to the same growth rate as the one found in Young (1995) if no adjustment for input quality is made.

⁷This is the US trend. They interpret productivity as the stock of knowledge useful in production and argue that knowledge is not country-specific.

⁸Mexican data were downloaded from <http://dgcnesyp.inegi.gob.mx>. South Korean data were downloaded from <http://www.nso.go.kr>. Thai data were downloaded from <http://www.nso.go.th>.

⁹Carrying out these adjustments can have important consequences for some countries where financial crises took place. For example, Indonesian GDP per capita fell by 13.1% between 1997 and 1998. If IBT are eliminated from GDP, then the fall in GDP is 8.0%. These adjustments make data consistent with variables in the simplest growth model.

3 The small open economy neoclassical model

This section evaluates the consistency of the small open economy neoclassical model with the behavior of output and inputs after the Mexican crisis of 1994. We model the crisis as exogenous shocks to TFP and interest rates. Feeding these shocks into the model yields paths for endogenous variables that we compare to data.

Because Mexico underwent deep fiscal changes in 1995 as part of the government's response to the crisis, we study a benchmark model where agents face distortionary taxes on consumption, capital income, and labor income. Also for fiscal reasons, Mexico's government significantly raised energy prices in 1995. To control for the impact of this shock, we model the role of energy in production. Incorporating these elements will enable us to measure the quantitative impact of fiscal shocks on the behavior of output in Mexico in 1995. But this complicates computations by preventing us from solving a standard planner's problem.

3.1 Benchmark model

Consider an economy in which time is discrete and infinite. The economy contains a continuum of mass one of identical households, and a continuum of mass one of identical firms. Households live forever. They order consumption and labor supply sequences $\{c_t, l_t\}_{t=0}^{+\infty}$ according to the following intertemporal utility function:

$$\sum_{t=0}^{+\infty} \beta^t \log \left(c_t - \frac{\rho}{\nu} l_t^\nu \right),$$

where $\beta \in (0, 1)$ is the discount factor, $\nu > 1$ determines the wage elasticity of labor supply and $\rho > 0$ measures the disutility from working. With these preferences, labor supply depends only on the current wage, w_t , and is independent of consumption or income. These preferences are commonly used in small open economy models (see e.g. Mendoza 1991, 2002, Correia, Neves and Rebelo, 1995 and Neumeyer and Perri, 2001). Correia, Neves and Rebelo (1995) argue that they improve the ability of small open economy models to replicate business cycle regularities.¹⁰

¹⁰Correia et al. (1995) and Neumeyer and Perri (2001) point out that these preferences are not consistent with a balanced growth path, unless the disutility of working increases with the rate of labor-augmenting

Households have access to an international capital market where one-period risk-free claims to a unit of the consumption good earn exogenous return r_t at the beginning of period t . We denote by a_t the risk-free asset holdings of households in period t . Households can also invest in physical capital, which they sell to firms at price $1 + r_t^k$. Let k_t be the quantity of capital held by households in period t . Adjusting capital across periods carries cost

$$\frac{\psi}{2} (k_{t+1} - k_t)^2,$$

where $\psi > 0$. As is well-known, adjustment costs are necessary in open economy models to prevent investment from being counterfactually volatile. Assuming that adjustment costs are borne by households rather than firms is immaterial. An equivalent decentralization would have firms make investment decisions and bear adjustment costs. The specification we use shortens the exposition by keeping the firm's problem static.

Households also face three types of taxes. In period t , consumption is taxed at rate τ_t^c , labor income is taxed at rate τ_t^l , and returns on physical capital and international assets are taxed at rate τ_t^k . In addition, households receive transfer T_t from the government. Therefore, households face the following budget constraint at date t :

$$c_t(1 + \tau_t^c) + k_{t+1} + a_{t+1} = l_t w_t (1 - \tau_t^l) + a_t(1 + r_t(1 - \tau_t^k)) + k_t(1 + r_t^k(1 - \tau_t^k)) - \frac{\psi}{2} (k_{t+1} - k_t)^2 + T_t.$$

We also assume that household borrowing is bounded below so as to rule out Ponzi schemes, and that the bound is low enough to never bind in equilibrium.

At date t , firms transform physical capital k_t^f , energy e_t and labor n_t into quantity $y_t = z_t \left(k_t^f\right)^{\alpha_k} n_t^{\alpha_n} e_t^{\alpha_e}$ of the consumption good, where z_t is TFP. Energy is available perfectly elastically at price p_t^e in date t . Fraction $\delta > 0$ of the physical capital firms purchase from households depreciates within each period. Therefore, at date t , firms choose (n_t, k_t^f, e_t) to maximize:

$$z_t \left(k_t^f\right)^{\alpha_k} n_t^{\alpha_n} e_t^{\alpha_e} + (1 - \delta)k_t^f - k_t^f (1 + r_t^k) - n_t w_t - e_t p_t^e.$$

technological change. In our model, there is no technological change. We follow Greenwood et al. (1988) and compare model predictions to data which have no trend. Mexican GDP per working age person displays no growth on average between 1980 and 2003.

The government collects tax revenues $\tau_t^c c_t + \tau_t^l l_t w_t + \tau_t^k (a_t r_t + k_t r_t^k)$ at date t . We assume that tax revenues and energy sales are rebated lump-sum to households. At date t , the government's budget constraint is:

$$\tau_t^c c_t + \tau_t^l l_t w_t + \tau_t^k (a_t r_t + k_t r_t^k) + e_t p_t^e = T_t. \quad (3.1)$$

We now define an equilibrium under the simplifying assumption that agents perfectly foresee the path of TFP, taxes and the exogenous interest rate. In the quantitative section, we consider other assumptions on expectations. Given an initial stock of capital and initial international assets (k_0, a_0) , an equilibrium in this environment is sequences of wages and prices of capital $\{w_t, r_t^k\}_{t=0}^{+\infty}$, consumption, labor supply and savings sequences $\{c_t, l_t, k_{t+1}, a_{t+1}\}_{t=0}^{+\infty}$, sequences of labor, capital and energy demands $\{n_t, k_t^f, e_t\}_{t=0}^{+\infty}$, and a sequence $\{T_t\}_{t=0}^{+\infty}$ of transfers such that, given prices, 1) $\{c_t, l_t, k_{t+1}, a_{t+1}\}_{t=0}^{+\infty}$ solves the household's problem, 2) $\{n_t, k_t^f, e_t\}_{t=0}^{+\infty}$ solves the firm's problem, 3) the market for physical capital clears ($k_t = k_t^f$ for all t), 4) the labor market clears ($n_t = l_t$ for all t) and 5) transfers satisfy (3.1). We will now ask whether this benchmark model can account for the behavior of output, labor, capital and energy after Mexico's Tequila Crisis.

3.2 Data and calibration

Computing the predictions of this benchmark model for Mexico requires paths for exogenous shocks $\{z_t, r_t, p_t^e, \tau_t^c, \tau_t^k, \tau_t^l\}_{t=0}^{+\infty}$ that are consistent with the theory. This requires a few adjustments to national income and product accounts data. Date t TFP in the benchmark model is:

$$z_t = \frac{y_t}{k_t^{\alpha_k} n_t^{\alpha_n} e_t^{\alpha_e}}. \quad (3.2)$$

Therefore, we need empirical counterparts for the theoretical variables y_t, k_t, n_t , and e_t . Appendix A describes the procedure we use in some detail. Our basic approach closely follows Atkeson and Kehoe (1999). We use quarterly data to construct the empirical counterparts of theoretical variables. National accounts data are from Mexico's national statistical institute (INEGI). There are four key conceptual differences between GDP as reported in the Mexican

national accounts (*measured GDP*) and output y_t in the model. First, y_t equals the sum of payments to labor, capital and energy, i.e. $y_t = w_t n_t + (r_t^k + \delta)k_t + p_t^e e_t$. GDP, on the other hand, treats energy as an intermediate output and thus corresponds to $y_t - p_t^e e_t = w_t n_t + (r_t^k + \delta)k_t$. Second, there is no energy-producing sector in the model, whereas measured GDP includes the value added by the energy sector. Third, measured GDP includes indirect business taxes (IBT), whereas output y_t does not. Finally, output includes the return to all types of capital in the model, whereas measured GDP does not. It excludes the return on government capital and the return plus depreciation of the stock of durable goods. We make the four corresponding adjustments to measured GDP to construct a measure of output consistent with y_t . We call this adjusted GDP *measured gross output*. We also construct capital, labor and energy series that are consistent with the model. In particular, we take into account the fact that in the model there is no energy-producing sector, and that in the model only firms use energy.

Besides empirical counterparts for y_t, k_t, n_t , and e_t , three technological parameters must be calibrated before measuring TFP. We assume that the share of labor income in GDP is 0.7. This assumption is supported by the work of Gollin (2002), who finds that after taking into account the income of the self-employed, labor income shares take values around 70%, across a large set of countries, and across time. We assume that the share of labor income in the energy-producing sector in Mexico is also 70%.¹¹ Given these assumptions, factor shares equal:

$$\begin{aligned}\alpha_n &= 0.7 \frac{GDP}{\text{Measured gross output}} = 0.664, \\ \alpha_k &= 0.3 \frac{GDP}{\text{Measured gross output}} = 0.285, \\ \alpha_e &= \frac{\text{Energy expenditure}}{\text{Measured gross output}} = 0.051,\end{aligned}$$

where all ratios are set to average yearly values. Having set those parameters, we generate a series for TFP using equation (3.2) in each period.

Turning now to exogenous shocks other than TFP, we calculate the interest rate r_t in

¹¹Verifying that this assumption is appropriate is difficult in the case of Mexico since Mexican national accounts do not provide compensation of employees data for oil and electricity companies.

period t as

$$r_t = \frac{(1 + Tbill\ rate_t)(1 + MX\ Brady\ spread_t)}{1 + US\ inflation_t} - 1,$$

where $Tbill\ rate_t$ is the interest rate on US Treasury bills, $MX\ Brady\ spread_t$ is the spread between the return paid by (dollar-denominated) Mexican Brady bonds and the interest rate paid by US Treasury bills, and $US\ inflation_t$ is the relative change in the US GDP deflator. In other words, our proxy for r_t is the real return paid by Mexican Brady bonds.¹² Our sample of Mexican Brady bond data starts in the last quarter of 1990 and ends in the first quarter of 2003. We measure the price of energy as a weighted average of the nominal price of natural gas, gasoline and electricity divided by Mexico's GDP deflator.¹³ Finally, we calculate taxes on consumption, labor income and returns from capital and international assets using the method of Mendoza, Razin and Tesar (1994).¹⁴ The calculated taxes are average effective tax rates, i.e the ratio of tax revenue to tax base.

Figure 2 plots the measured shocks. Most of these series underwent unusually large changes in 1995. In particular, and not surprisingly given the fact that capital and labor fall much less than output during 1995, TFP falls markedly during 1995. Measured gross output fell 10.1% between the last quarter of 1994 and the last quarter of 1995, while capital fell by 0.7%, labor fell by 2.5%, and energy fell by 24.4%. Given these data and our calculated technological shares, TFP must fall by 7.1% to account for the fall of measured gross output in 1995. Interest rates measured in annual terms rise from 8.7% on average during 1994 to 19.5% in the first quarter of 1995. The price of energy jumps by 43% between the last quarter of 1994 and the first quarter of 1995. The consumption tax rate rises from 10.4% to

¹²Neumeyer and Perri (2001) use a similar construct to study the relationship between business cycles and international interest rates in developing countries. We use end of quarter rates, using average rates does not alter our quantitative findings.

¹³We follow the method used by Atkeson and Kehoe (1999). Appendix A provides the details. We are constrained to use yearly data on prices and sales of energy to calculate the average price of energy. We assume that the nominal prices of different kinds of energy remain constant throughout each year. The Mexican government typically adjusts energy prices either at the end or at the beginning of each year.

¹⁴Only data on total income tax revenues are available in Mexico. We follow the estimate reported in Fernandez and Trigueros (2001) to split total income tax revenue into its components: individual and corporate. We use these components to measure the tax rate on labor income, and on capital and asset returns. Also, when measuring consumption taxes using OECD data, Mendoza et al. (1994) exclude the "Other taxes" item. Because this last item is large in magnitude in Mexico, we choose to include it. We are constrained to use yearly data to measure taxes. We assume in the numerical experiments that taxes remain constant throughout each year.

13.3% from the last quarter of 1994 to the first quarter of 1995. On the other hand, the tax rate on labor shows almost no change, falling from 12.5% to 12.2% between 1994 and 1995. The tax rate on capital income falls from 9.5% to 7.4%.

Overall, the Mexican economy underwent a number of severe negative shocks in 1995. We will now argue that given the magnitude of these shocks, the model predicts that output should have fallen much more than it did. We will also argue that the quantitative impact of changes in fiscal policy is small compared to the role of TFP. To make these points, we first need to calibrate preference and adjustment cost parameters. One way to calibrate the model would be to assume that at a given date Mexico was on a balanced growth path. However, we do not think that such an assumption is appropriate. Mexico underwent a series of deep crises in the 1980s after decades of brisk growth. Between 1980 and 2003, GDP per capita did not grow in Mexico, and we do not believe this to be a balanced growth path. Our calibration strategy consists of choosing parameter values to match certain statistical properties of inputs and investment before 1995.

Preference parameters ρ and ν determine the level and volatility of labor supply, respectively. We set ρ to match the average of our measure of hours worked per working age adult before 1995. As for ν , we begin by setting it to 1.5, which implies a wage elasticity of labor supply of 2, the value used in Mendoza (1991).¹⁵ It falls within the range mentioned by Greenwood, Hercowitz and Huffman (1988), who cite studies of labor supply in the United States. We were unable to find similar studies for Mexico. Setting $\nu = 1.5$ implies a volatility of hours worked that is near its pre-crisis counterpart in Mexico. The next section provides some sensitivity analysis on this key parameter. In this benchmark model the predicted path for input and output series is independent of β . We simply set it as in Correia, Neves and Rebelo (1995) to satisfy $\beta [1 + r (1 - \tau^k)] = 1$, where r and τ^k are the long run values of the international interest rate and the tax on the return on international assets. This assumption is necessary for an equilibrium with zero long run growth of consumption to exist. To obtain a long run value for the interest rate, we assume that the value it takes in the first quarter of 2003 (0.9% at a quarterly rate), the last date in our sample, will be Mexico's cost of international funds in the future. We also use the last value for τ_t^k in our sample

¹⁵The elasticity of labor supply is $\frac{1}{\nu-1}$.

(9.1%) as the long run value of the tax on capital income. Next, we choose ψ , the capital adjustment cost parameter, to match the observed standard deviation of the investment to measured gross output ratio before 1995. The resulting aggregate adjustment costs amount to a negligible fraction of GDP. Finally, in all quantitative experiments, we scale the relative price of energy to match the average energy level prior to 1995.

Having set all parameters, we calculate the path our model predicts for inputs and output under two assumptions on agents' expectations. In the first experiment (perfect foresight, PF) we assume that agents know the entire sequence of exogenous shocks shown in Figure 2 before making any decision. In the second experiment (perfect surprise, PS) we assume that agents foresee all shocks up to the last quarter of 1994. After 1994, agents expect all shocks other than the interest rate to permanently assume their average values before 1995. As for the interest rate, households expect it to be constant at $\frac{\beta^{-1}-1}{1-\tau_{average}^k}$, the only value compatible with zero long-run consumption growth, where $\tau_{average}^k$ is the average capital income tax rate before the crisis. In other words, under the PS scenario, agents do not expect a crisis to occur in 1995. When they observe the values of shocks in the first quarter of 1995, agents immediately revise their expectations of future shocks to the correct path. We view this experiment as an approximation to a situation where households assign a positive but very small probability to the possibility of a crisis in 1995. These simplifying assumptions on expectations enable us to use a simple non-linear solution method based on Euler equations. Specifically, first order conditions from the firm and the household problems' imply that the evolution of capital in this model is described by the following second-order difference equation for all t :

$$1 + r_{t+1}(1 - \tau_{t+1}^k) = \frac{1 + \left(\alpha_k \frac{y_{t+1}}{k_{t+1}} - \delta\right) (1 - \tau_{t+1}^k) + \psi (k_{t+2} - k_{t+1})}{1 + \psi (k_{t+1} - k_t)}. \quad (3.3)$$

Given the initial level of capital, we use a shooting algorithm to find the path of capital such that endogenous variables converge to steady state when exogenous variables stay at their level in the first quarter of 2003 forever. Appendix B provides details. The equilibrium path for other endogenous variables can then be calculated as a function of capital and exogenous shocks. Given the magnitude of shocks in 1995, using linear approximations around the

steady state would yield inaccurate results.¹⁶

3.3 Results

Figure 3 plots the predictions of the model for GDP, labor, the capital-GDP ratio, and energy, for both the PF and PS experiments, and compare them to data. Each time series is scaled by its respective value in the last quarter of 1994 to emphasize the impact of the crisis. The key result is that GDP, labor and capital fall more than twice as much in percentage terms as in the data, under both expectation scenarios. This is true, that is, whether or not agents saw the crisis coming. In both experiments, GDP falls by 19.5% between the last quarter of 1994 and the last quarter of 1995, compared with 9.8% in the data. The behavior of energy is also very similar across experiments. The yearly fall in predicted energy in 1995 is similar to the fall in the data.¹⁷

The two experiments make very different predictions for the behavior of output and inputs before the crisis. In the PF experiment, the capital-output ratio falls before 1995 as agents anticipate the crisis. This makes all variables fall in anticipation of the large changes in exogenous variables in 1995. Series predicted by the PS experiment track their empirical counterparts more closely as agents base their investment decisions on optimistic expectations.

To measure the relative role of each of the many shocks that hit the Mexican economy in 1995, we carried out PF experiments in which only one of the exogenous variables changes after the last quarter of 1994, while other variables remain constant at their values in the last quarter of 1994. We find that changes in the capital tax and the labor tax have little effect on the behavior of GDP in 1995. Shocks to the consumption tax, interest rates and the price of energy yield more pronounced falls in output: -2.9%, -2.2% and -1.4% during 1995 respectively.¹⁸ The impact of TFP outweighs that of all other shocks combined. Holding other exogenous variables at their end-of-1994 values, TFP alone would have caused GDP

¹⁶Dotsey and Mao (1992) find that the accuracy of linear approximation methods worsens as the variance of shocks rises.

¹⁷Energy falls much more on impact than in the data, but this owes in large part to our assumption that nominal prices remained constant in 1995 which we have to make for lack of quarterly data.

¹⁸We take into account the fact that keeping the interest rate fixed at its (high) value in the last quarter of 1994 induces a trend in endogenous variables. Results are reported net of this trend.

to fall by 15.4% in 1995 relative to 1994. The magnitude of the TFP shock accounts for the model's counterfactually large fall in output. The benchmark model's difficulties in matching the behavior of output and inputs during Mexico's 1995 crisis do not stem from fiscal shocks.

3.4 Sensitivity analysis

This section evaluates the robustness of the previous findings to our assumptions on preferences.

Elasticity of labor supply

Our findings are sensitive to the assumed elasticity of the labor supply on the part of households. In particular, a higher ν would render labor supply less elastic, which should reduce the predicted fall in hours worked, hence in output in 1995. In fact, it is clear that one can find a value for ν such that the model will predict the correct fall in hours worked during the crisis. Figure 4 shows that setting $\nu = 4.33$, which is at the upper bound of the estimates available for the United States, produces a fall in hours in 1995 that resembles quantitatively the fall in the data, in a PF experiment.¹⁹ The same is true for the behavior of GDP. Hours worked fall 2.5% in the data and 3.3% in the model. GDP falls 9.8% in the data and 11.7% in the model. Additionally, the model predicts a recovery in GDP very similar to the one observed, as can be seen on the graph. In fact, predicted and observed GDP meet by the first quarter of 2003.

However, such a value for ν predicts a counterfactually stable path for labor input outside the crisis. The standard deviation of predicted hours before the crisis and over the full length of our sample is quite smaller than in the data, as can be seen in the figure. The standard deviation of labor predicted by the model between 1990 and 2003 is 0.004, smaller than the actual value of 0.007. The volatility of predicted labor represents 57% of actual volatility. In short, it is not possible to find a value for ν such that the model yields a reasonable path for hours worked both during and outside of the crisis period.

¹⁹Greenwood et al. (1988) report estimates of the elasticity of labor supply. The value of $\nu = 4.33$ is implicit in the lowest estimate reported.

Standard preferences

Heretofore we have assumed preferences such that the wage elasticity of the labor supply is exogenous and invariant over time. As we have mentioned, these preferences are typically used in small open economy models because they improve the model's consistency with business cycle regularities. It is interesting nonetheless to consider the impact of giving households preferences that are more standard in closed economy exercises. Specifically, assume that households now order consumption and labor supply sequences $\{c_t, l_t\}_{t=0}^{+\infty}$ according to the following intertemporal utility function:

$$\sum_{t=0}^{+\infty} \beta^t \{ \log c_t + \rho \log(1 - l_t) \},$$

where $\rho > 0$ measures the weight of leisure in utility. Households face the same budget constraint as before. Solutions to the household problem must satisfy, for all t :

$$\frac{c_{t+1}}{c_t} = \frac{\beta(1 + \tau_t^c)}{1 + \tau_{t+1}^c} (1 + r_{t+1}(1 - \tau_{t+1}^k)) \quad (3.4)$$

$$\frac{\rho c_t}{1 - l_t} = \frac{w_t(1 - \tau_t^l)}{1 + \tau_t^c}. \quad (3.5)$$

Both conditions have the usual interpretation. The first says that the marginal rate of substitution between consumption in two consecutive periods must equal the return on savings (the marginal rate of transformation between date t and date $t + 1$ consumption). The second equates the marginal utility of leisure in each period to its opportunity cost, the net wage times the marginal utility of consumption. Using first order conditions for profit maximization by firms (those are unchanged), (3.5) can be rearranged to read:

$$l_t = \left(1 + \frac{(1 + \tau_t^c)\rho c_t}{(1 - \tau_t^l)\alpha_n y_t} \right)^{-1} \quad (3.6)$$

Condition (3.6) shows how standard preferences could help account for the behavior of hours worked in 1995. Hours worked are now a simple function of the consumption-output ratio. If the model predicts a fall in consumption comparable in relative size to the fall in output

in 1995, the model will also predict little change in hours, as in the data.²⁰

Computing the model requires solving for paths of consumption, hours worked, assets and capital that satisfy (3.4), (3.6), the household's budget constraint, and the same difference equation in capital as before. In implementing the algorithm described in appendix B, we set ρ to match the average level of hours worked before the crisis. We also choose the initial level of asset a_0 so that the model implies an approximate debt to GDP ratio of 35% for Mexico in 1994, as in the data.²¹ To match the fact that hours have no trend before 1995, we set the rate of time preference to the average net interest before the crisis, and set long term interest rates accordingly in our two expectation scenarios.

The model with standard preferences performs very poorly under perfect foresight, for obvious reasons. In this model, consumption rises at the after-tax rate of interest net of the rate of time preference. Since interest rates are very high in 1995, consumption rises throughout the year (see equation (3.5)) while TFP falls markedly. Correspondingly, the consumption-output ratio rises markedly and hours worked fall even more drastically than in the previous model. Those results are available upon request.

Under perfect surprise assumptions however, agents adjust consumption in the first quarter of 1995 after discovering the true path of exogenous series. In particular, consumption must be adjusted downward since agents realize that their future income will be much lower than expected. This could mitigate the negative impact on hours worked. In fact, as figure 5 shows, the consumption-output ratio actually falls, so that hours *rise* in the first quarter.²² But this effect is short-lived, as consumption then starts rising steeply due to high interest rates. Hours adjust downward and eventually bring output, hours and capital markedly below trend. In other words, once agents have adjusted to the crisis, hours and output and input series fall below their data counterpart.

The basic problem is that given persistently high interest rates after the crisis, the model predicts that consumption should rise faster than output, and that, correspondingly, hours

²⁰For consistency with the model, we measure consumption as the sum of private consumption, the returns and depreciation of durable goods, the returns to government capital, and government consumption minus Indirect Business Taxes.

²¹This is approximately the value reported in Lane and Milesi-Ferretti (2001).

²²We replace the energy panel of the figure with the consumption-output ratio path since this is the crucial statistic in this model.

should fall for several periods, which is at odds with the evidence. If one sets the rate of time preference to offset the high interest rates that prevail after the crisis, predicted hours trend steeply up before the crisis when interest rates are relatively low, which is also at odds with the evidence. This a manifestation of the difficulties standard preferences present for open economy models. Because interest rates are volatile and display large, persistent deviations from their average values in economies such as Mexico, predicted consumption and hours worked display counterfactually large fluctuations.

4 Factor utilization

The fact that the predictions for inputs deviate drastically from their empirical counterparts suggest that factor hoarding could play a big role during crises. In this section, we use standard models of factor utilization to quantify the importance of capital utilization and labor hoarding.

4.1 Capital utilization

Consider a small open economy model with variable capital utilization modeled as in Greenwood, Hercowitz and Huffman (1988). Household preferences are the same as in the benchmark model. Firms can now alter the rate at which they utilize capital. Raising utilization in a given period raises output, but it also raises the quantity of capital lost to depreciation.²³ Specifically, depreciation at date t depends on utilization u_t as follows:

$$\delta_t = \frac{u_t^\phi}{\phi},$$

where $\phi > 1$. Output at date t is now given by:

$$z_t^u \left(u_t k_t^f \right)^{\alpha_k} n_t^{\alpha_n} e_t^{\alpha_e}.$$

²³Gertler, Gilchrist and Natalucci (2003) also combine the framework in Greenwood et al. (1988) with a small open economy model.

Firms continue to take all prices as given and choose k_t^f , n_t , e_t and u_t each period to maximize:

$$z_t^u \left(u_t k_t^f \right)^{\alpha_k} n_t^{\alpha_n} e_t^{\alpha_e} + \left(1 - \frac{u_t^\phi}{\phi} \right) k_t^f - k_t^f (1 + r_t^k) - n_t w_t - e_t p_t^e.$$

This maximization problem yields the following condition for optimal utilization at date t :

$$u_t = \left(\alpha_k \frac{y_t}{k_t} \right)^{\frac{1}{\phi}}, \quad (4.1)$$

as in Greenwood, Hercowitz and Huffman (1988). That is, the capital-output ratio path implies a unique utilization path. Given measures of the capital to output ratio and a value for ϕ , TFP net of changes in capital utilization can then be computed as:

$$z_t^u = \frac{y_t}{(u_t k_t)^{\alpha_k} n_t^{\alpha_n} e_t^{\alpha_e}}.$$

While no further adjustment to national accounts data is needed to implement those calculations, the capital stock needs to be recalculated, because its evolution depends on utilization in each period. The capital stock and the utilization rate need to be calculated recursively. Using an initial capital stock and a value for ϕ we calculate utilization as defined by condition (4.1). Then next period's capital stock can be calculated using the following law of motion:

$$k_{t+1} = k_t \left(1 - \frac{u_t^\phi}{\phi} \right) + i_t,$$

where i_t is gross capital formation net of adjustment costs. Proceeding recursively yields a path of capital, utilization and therefore of TFP adjusted for utilization. Implementing this procedure requires a value for ϕ , the curvature of the depreciation schedule. Simple algebra shows that in this model the steady state depreciation rate is equal to $\frac{r}{\phi-1}$ where r is the steady state rate of interest. We choose $\phi = 1.44$ to imply a steady state yearly depreciation rate of 8% (the constant depreciation rate we assumed in the benchmark model), assuming that interest rates eventually become constant at their last value in our sample.²⁴ We also

²⁴This value for ϕ is close to values used in studies of the U.S. economy. Greenwood et al. (1988) use $\phi = 1.42$ to imply a steady state yearly depreciation rate of 10%. Burnside and Eichenbaum (1996) estimate

experimented with different values of ϕ , including a value such that the implied steady state depreciation rate is 5% on a yearly basis. The quantitative results are practically the same in all cases.

Because our measure of the capital-output ratio falls in 1995, utilization does as well. This makes intuitive sense. TFP in the first quarter of 1995 falls by a large amount while interest rates (the opportunity cost of capital) increase significantly. This gives firms an incentive to postpone the consumption of capital services. Specifically, we find that measured utilization fell 5.7% between the last quarter of 1994 and the last quarter of 1995. This implies that adjusted TFP falls less than unadjusted TFP (5.1% versus 7.1%), as shown in figure 6. Note however that it continues to fall by a large amount. In fact, relative to movements outside of the crisis period, the 1995 change in adjusted TFP is as much of an outlier as the change in unadjusted TFP.

The key question is whether making capital utilization endogenous improves the ability of the model to account for the behavior of output and inputs. To answer that question, we first recalibrate parameters to continue matching our calibration targets. Figure 7 plots the predictions of the model for GDP, labor, capital and utilization, under the PS scenario.²⁵

The results are quantitatively similar to those we obtained in the benchmark model. GDP, labor, and capital fall much more than in the data in 1995. In particular, GDP falls by 22.5%, even though adjusted TFP falls less than unadjusted TFP. The reason for this is simple. When confronted with exogenous shocks, firms can adjust labor and energy as before, but they can also vary capital utilization. This new margin of adjustment magnifies the economy's response to productivity shocks.²⁶ The predicted fall in utilization in 1995 (8.5% in the PS experiment) is higher than in the data as the model predicts a greater increase in the capital-output ratio than in the data. Finally, like in the benchmark model, TFP and the interest rate shock account for most of the fall in output. Experiments in which taxes and the price of energy are held constant at their pre-1995 level produce a fall

ϕ to be 1.56 in the U.S.

²⁵We do not show the results of the PF experiment to reduce clutter. Results are very similar to those we obtained with the benchmark model.

²⁶This is consistent with the findings of Burnside and Eichenbaum (1996). They find that the response of the economy to a given productivity shock is magnified once variable capital utilization is introduced in a real business cycle model.

in output that continues to markedly exceed its data counterpart. For instance, GDP falls by 19.9% between 1994 and 1995 in the PF case when fiscal variables are held constant.

In summary, including variable capital utilization helps account for some of the variance of TFP but does not improve the performance of the model during 1995. The model continues to predict that output and inputs should fall at least twice as much as observed in 1995.

4.2 Labor hoarding

Assume now that firms and households can use another margin of adjustment when confronted with exogenous shocks: effort. We model labor hoarding in the spirit of Burnside et al. (1993).²⁷ Time devoted to work by households is indivisible: employed households devote time $f > 0$ to work while unemployed households devote no time to work. As in Hansen (1985) and Rogerson (1988), we convexify the choice set of households by allowing them to randomize between employment and unemployment. Specifically, households choose a probability l_t of working in a given period, a level c_t^e of consumption when employed, a level c_t^u of consumption when unemployed, and a level ϵ_t of effort when employed. We further assume that working entails a fixed cost $\kappa > 0$.²⁸ Households maximize:

$$\sum_{t=0}^{+\infty} \beta^t \left[l_t \log \left(c_t^e - \kappa - \frac{1}{\nu} (f\epsilon_t)^\nu \right) + (1 - l_t) \log (c_t^u) \right].$$

With this utility function, effort is independent of consumption and income, as labor supply was in the benchmark model. This makes the model with labor hoarding comparable to the benchmark model in the sense that the short-run wage elasticity of aggregate labor supply is governed by exogenous parameter $\nu > 1$, and is independent of income and consumption. We further assume as Burnside et al. (1993) that adjusting labor between periods is costly, although our specification of the cost is different from theirs. Burnside et al. (1993) assume that it takes one period to adjust employment. This constraint would never bind in the perfect foresight case and would only bind in the first quarter of 1995 in the perfect surprise

²⁷Baxter and Farr (2001) construct a two-country model with utilization modeled as in this paper and with labor hoarding as in Bils and Cho (1994).

²⁸If there is no such cost, it is efficient for households to work in every period, because they can adjust their labor effort in response to exogenous shocks. See Burnside et al. (1993).

case. We assume instead that households who change their work probability from l_t to l_{t+1} in period $t + 1$ bear costs $\frac{\psi_l}{2}(l_{t+1} - l_t)^2$ where $\psi_l > 0$.²⁹ This specification is similar to the one used by Cogley and Nason (1995). As in the case of capital in the benchmark model, assuming that adjustment costs are borne by households rather than by firms is immaterial but simplifies the exposition by keeping the firm's problem static. Output in this model is given by

$$y_t = z_t^{u,e} \left(u_t k_t^f \right)^{\alpha_k} \left(n_t f \epsilon_t^f \right)^{\alpha_n} e_t^{\alpha_e},$$

where ϵ_t^f is the firm's effort choice, and n_t is the fraction of households that they employ. In equilibrium, $n_t = l_t$ and $\epsilon_t = \epsilon_t^f$ for all t . In appendix C, we show that optimal behavior on the part of firms and households in this environment imply the following condition for effort in equilibrium:

$$\epsilon_t = \left(\frac{\alpha_n(1 - \tau_t^l)y_t}{(1 + \tau_t^c)n_t f} \right)^{\frac{1}{\nu}}. \quad (4.2)$$

Note that effort depends negatively on both the tax on labor and the tax on consumption. Calibrating ν is difficult since independent evidence on this parameter is not available. We choose to experiment with various values centered around $\nu = 1.5$, the value we used for the curvature of the disutility of labor in the benchmark model. The fixed length of work f is set to 0.45. This number corresponds to average hours per worker before 1995, relative to approximate discretionary time available in a quarter, 1300 hours. These parameters are sufficient to infer effort from the observed path of hours worked $n_t f$ and output. Figure 6 shows the behavior of adjusted TFP when $\nu = 1.5$.³⁰ Combined, capital utilization and effort account for 90% of the fall in TFP in 1995. Note that factor hoarding accounts for much of the variance of TFP outside of the crisis as well.

We now ask whether including effort improves the consistency of the model with evidence. To that end, we need to assign values to a few more parameters. In all experiments we choose κ to match the initial level of employment in our sample. Following our previous calibration strategy, the natural way to choose a value for employment adjustment cost parameter ψ_l

²⁹Without adjustment costs, one easily shows that the optimal level of effort is constant across periods.

³⁰Results for a variety of other values of ν are available upon request. Reducing ν increases the elasticity of effort, and effort accounts for a greater share of the variance of TFP when ν is low. At the same time, for all values of ν considered, effort accounts for a large fraction of the volatility of TFP, both during and outside of the crisis.

is to match the standard deviation of employment before the crisis. However, we found that doing this led to unreasonably large fluctuations in employment after the crisis. To keep employment within reasonable bounds, we choose a value of ψ_l such that n_t remains between 40% and 60% throughout the simulation period in all experiments. One should bear in mind this calibration compromise when interpreting our results.

The predictions of the model in the perfect surprise case are shown in figure 8 for the case $\nu = 1.5$. Hours worked are very smooth, like capital, which is not surprising since their evolution is governed by a similar second order difference equation, and labor adjustment costs are set high.³¹ In particular, hours worked now fall less than observed in 1995: 2.0% versus 2.5%. Because hours fall very little, GDP falls much less than in previous experiments during the crisis. Predicted GDP falls slightly more than observed GDP: 10.0% versus 9.8%. The model predicts paths for the capital-output ratio and output-employment ratio very similar to the ones observed. Consequently, predicted utilization and effort track closely the fluctuations of observed counterparts. Energy falls significantly more than in the data in the first quarter of 1995, as in the previous experiments. But the relatively good behavior of the model during 1995 is short-lived. As labor slowly adjusts, GDP remains below its data counterpart and eventually diverges from it by magnitudes quite similar to what we obtained in the benchmark economy. The eventual real impact of the crisis is just as large as in the previous models. We also found that varying ν does not change results much.³²

Finally, we carried out a PS experiment in which taxes and the energy price remain constant at their pre-1995 average values, assuming $\nu = 1.5$. In sharp contrast to previous experiments, the real impact of the crisis becomes smaller than observed when only interest rates and TFP are allowed to fluctuate as in the data after 1994. Predicted GDP falls less than observed: 5.2% versus 9.8% between 1994 and 1995.

³¹Even though the resulting value of ψ_l is high, employment adjustment costs represent a small fraction of gross output, at most 0.5% in any given period.

³²Results are available upon request. The overall impact of the crisis generally rises with ν as the elasticity of effort falls.

5 Conclusion

In this paper, we document the fact that TFP falls by very unusual magnitudes after financial crises, and argue that these falls pose a serious challenge for the standard small open economy model. In the case of Mexico's Tequila crisis, the benchmark model predicts that inputs and output should have fallen at least twice as much in percentage terms as they did. Standard models of endogenous factor utilization suggest that factor hoarding can account for much of the fall of TFP during the crisis. However, we find that this does not solve the output puzzle. The factor hoarding models we study also predict that output should fall much more than it does following financial crises.

Our paper focuses for the most part on the Mexican case, but we provide some evidence that TFP fell by a very unusual amount during recent crises in Asia, a fact that is ignored by most quantitative studies of these episodes. Falling capital utilization may explain some of the real impact of these crises, as argued by Gertler et al. (2003), but like in the case of Mexico, this is unlikely to account for the behavior of output. More generally, we interpret our results as suggesting that a full understanding of the real impact of financial crises will require some modeling of the allocation of resources across sectors. For example, employment started growing briskly in the export sector in Mexico after the devaluation in 1994. The fall in output could reflect transitory losses in the quality of labor as employees devote time to learning new skills.

Finally, our findings have implications for the potential ability of the neoclassical growth model to account for depression episodes in Latin American nations. Kydland and Zarazaga (2002) use a closed economy version to account for Argentina's recent economic history. The most challenging time period for the model begins with Argentina's 1989 crisis, and extends past Argentina's 1995 crisis. The model predicts that investment should have fallen much more than it did in 1989, and recovered much faster than it did thereafter. Similarly, the results in Bergoeing et al. (2002) show that a closed economy model's predictions diverge from the evidence around Mexico's Tequila crisis. Output and especially labor fall much more than observed in 1995. Recent depression periods in Latin American are best described as series of financial crises. The closed economy neoclassical growth model's difficulties in

accounting for the behavior of output, the capital-output ratio and hours worked in those countries likely stems from the fact that it does not account well for the real impact of financial crises.

A Mexican data appendix

This appendix describes how we construct empirical counterparts for theoretical variables with Mexican data. We use quarterly data when available, and impute quarterly series from yearly series otherwise. Our sample goes from 1980.1 to 2003.1. Data from original sources are seasonally adjusted using the Census Bureau's X12-Arima procedure. All data are in 1993 prices. The following quarterly series are available from Mexico's Instituto Nacional de Estadística, Geografía e Informática (INEGI) and Mexico's Central Bank (detailed sources are available upon request):

1. Gross Domestic Product (GDP)
2. Gross Fixed Capital Formation, private (GFCFp)
3. Gross Fixed Capital Formation, public (GFCFg)
4. Change in inventories (CH)
5. Private Gross Capital Formation (GCFp)=GFCFp+CH
6. Gross Capital Formation=GCFp+GFCFg
7. Purchases of durable goods by households.

In our model there is no energy producing sector. A first step towards making the variables in the data consistent with theoretical ones is to eliminate the energy sector from the national accounts. There are no quarterly series for GDP or investment in the energy sector. INEGI provides GDP data at yearly frequency in 1993 prices for the oil and electricity sectors between 1988 and 2001. To construct quarterly data, we multiply quarterly GDP by the yearly ratio of the energy sector's GDP to overall GDP. For years before 1988 and after 2001 we use the ratio's average between 1988 and 2001. Yearly data for investment in the oil industry is available from INEGI for the 1980-2002 period, and after 1987 for the electricity sector. As in the case of GDP, we use yearly ratios of investment in the energy sector to overall investment to construct a quarterly series after 1987. For years before 1987 and after 2000 we assume that the ratio is equal to the value observed in 1987. We use the 1987 value rather than the post-1987 average because the ratio decreases from 1987 on.

In order to make GDP from national accounts consistent with output in our model, we need to construct a few more quarterly variables: Indirect business taxes (IBT), the returns of government capital, the returns and depreciation of the stock of durable goods, and energy expenditures by firms outside of the energy sector. IBT data are available from INEGI only at yearly frequency, and we use yearly ratios as above to impute quarterly data.

Using data on private and public investment, and purchases of durable goods, we construct three capital stock series using the perpetual inventory method.³³ We assume a yearly depreciation rate of 6% for private capital, 5% for government capital and 20% for durable goods. To construct the stock of total capital, we add up the three resulting stocks. The average yearly depreciation rate implied by the total stock of capital, total investment and the law of motion of capital in the benchmark model is 8%. To calculate gross returns to government capital and the stock of durables we assume a net yearly return of 4% and the same depreciation rates as above.

Quarterly energy consumption data come from INEGI.³⁴ Consumption numbers for the non-energy producing sector for *gas licuado* (LPG), *combustóleo* (fuel oil), diesel, and *gasolina* (gasoline) are based on internal sales (*ventas internas*) plus imports into Mexico. This quantity approximates consumption by all sectors other than the energy-producing sector. The residential and public sectors were removed using weights inferred from annual consumption data available from the *Secretaría de Energía* (SENER). Quarterly electricity data from INEGI include only the industrial sector, so we used annual industrial electricity consumption as a percentage of total business sector consumption from SENER to impute consumption by the rest of the business sector. All the series were converted to megajoules.

To obtain an energy price index, prices for different types of energy are necessary. Electricity prices are average prices charged by the public sector to the industrial sector (*Precios Promedio de Energía Eléctrica del Sector Eléctrico Paraestatal*). Oil related product prices are the prices charged in Mexico (*Precios Internos de los Principales Productos*) as reported by INEGI. After converting the prices per unit for the different types of energy into a common unit (pesos/megajoule), consumption numbers were used to calculate a weighted price index, which was then converted into real terms using the GDP deflator. We use this index to calculate expenditure on energy by firms outside of the energy sector.

After constructing these variables, we calculate the data counterpart of gross output in our model by subtracting from GDP indirect business taxes and the energy sector's GDP, and adding the imputed returns and depreciation of government capital and durable goods. Given that output in the model includes the expenditure on energy by firms, we also add this variable to GDP. To calculate per capita variables we use the yearly series for population of age 15 to 64 reported in Bergoing, Kehoe, Kehoe and Soto (2002). To construct quarterly working age population, we take yearly growth rates of population and calculate implicit quarterly rates.

We also need an empirical counterpart for labor input in the model. To that end, we first calculate (seasonally adjusted) average hours worked in the manufacturing sector from Mexico's Manufacturing Sector Survey available from INEGI. The survey produces monthly series for man-hours and for employment. There are two versions of the survey. The first one has data from 1987.01 to 1995.12. The second one has data starting in 1994.01. We splice the

³³We assume that gross capital formation data includes the empirical counterpart of theoretical adjustment costs. In our simulations, adjustment costs amount to at most 0.9% of GDP in the benchmark model in any given period.

³⁴In the process of producing these data, we discovered several errors in the electricity series published by INEGI, including the fact that they did not reflect the effects of major tariff changes in 1992. We are most grateful to Rafael del Villar from Banco de México and Jorge Garcia Peña from CFE for helping us construct the correct series. INEGI has now updated and corrected its series.

quarterly hours per employee of the two surveys. To calculate a measure of workers relative to total working age population, we multiply quarterly measures of the ratio of economically active population relative to population of age 12 and higher by the employment rate and by the fraction of employment in sectors other than the mining and electricity industries. These data are available from INEGI. The measure of the labor input per working age person consistent with the model is hours per employees times the ratio of employed persons to population. We scale the resulting series by 1300, an approximation of the total number of hours of discretionary time available in a quarter.

B Computational appendix

Benchmark model

Simple manipulations of first-order conditions for profit and utility maximization show that output can be reduced to a function of capital, so that equation (3.3) is a second order difference equation in capital only. We assume that after the first quarter of 2003 all exogenous variables stay forever at their level in the first quarter of 2003. Given k_0 , we look for the unique k_1 such that the economy eventually converges to steady state via a standard shooting algorithm. All endogenous variables can then be calculated as a function of the path of physical capital. In the perfect surprise (PS) experiment, the algorithm is re-started in the first quarter of 1995 using as initial value for capital the value agents would choose under the expectations assumed before 1995.

Standard preferences

Given parameter values and paths for exogenous shocks, the algorithm we use consists of the following steps:

1. Guess a_0 , the initial stock of risk-free bonds held by households.
2. Guess $\frac{c_0}{y_0}$ and calculate n_0 . Since k_0 is known and $e_0 = \alpha_e \frac{y_0}{p_0^e}$ from the firm's first order conditions, y_0 can be calculated using the definition of output. So then can c_0 .
3. Get c_1 from (3.4).
4. Guess k_1 and get y_1 and n_1 using (3.6) and the definition of output.
5. For $t \geq 0$ obtain c_{t+2} and k_{t+2} inductively by repeating steps 2 and 3.
6. Iterate on k_1 until path for capital is stable (i.e. variables converge to steady state values).
7. Update $\frac{c_0}{y_0}$ until path for assets is stable.
8. Update a_0 until the debt-GDP ratio predicted by the model for 1994 is 35%, as in the data.

Capital utilization

The model produces the same second order difference equation for capital as before, except that output and depreciation now depend on utilization. But utilization is a function of the capital-output ratio (see equation 4.1). Therefore, (3.3) can be written as a second order difference equation in capital only as in the benchmark model, and the same shooting algorithm can be used.

Labor hoarding

The algorithm is much more demanding in this case because one needs a simultaneous solution to two second-order difference equations: one for capital as before, and one for employment. Given initial values (k_0, n_0) for capital and labor, we carry out the following steps:

1. Guess a full path $\{n_t^{guess}\}_{t=0}^T$ for employment where T is a large number.
2. Choose k_1 so that, given the labor guess, the path predicted by the second order difference equation for capital is stable.
3. Find n_1 so that given the path for capital obtained in step 2, the path for labor predicted by (C.6) is stable.
4. Iterate until the paths for labor and capital are approximately invariant

In the perfect surprise experiment, the algorithm is re-started in the first quarter of 1995 using as initial value for capital and employment the values agents would choose under optimistic expectations before 1995.

C Labor hoarding model

Assume as in the text that households maximize:

$$\sum_{t=0}^{+\infty} \beta^t \left[l_t \log \left(c_t^e - \kappa - \frac{1}{\nu} (f \epsilon_t)^\nu \right) + (1 - l_t) \log (c_t^u) \right]$$

where ϵ_t is effort at date t , l_t is the probability that a household becomes employed, c_t^e is consumption if the household is employed, and c_t^u is consumption if the household is unemployed. Assuming quadratic labor adjustment costs and letting w_t be the price of labor services, households now face budget constraint:

$$\begin{aligned} & (l_t c_t^e + (1 - l_t) c_t^u) (1 + \tau_t^c) + k_{t+1} + a_{t+1} \\ &= l_t f \epsilon_t w_t (1 - \tau_t^l) + a_t (1 + r_t (1 - \tau_t^k)) + k_t (1 + r_t^k (1 - \tau_t^k)) - \frac{\psi}{2} (k_{t+1} - k_t)^2 - \frac{\psi_l}{2} (l_{t+1} - l_t)^2 + T_t. \end{aligned}$$

In Hansen (1985) or Rogerson (1988), it is optimal for agents to equate consumption across employment states. In our model this is not the case. It remains true that households equate

utility across employment states at all dates t :

$$c_t^e - \kappa - \frac{1}{\nu} (f\epsilon_t)^\nu = c_t^u. \quad (\text{C.1})$$

This implies that employed households consume more than unemployed households. Similarly, utility maximization by households implies that employment and effort solve:

$$\epsilon_t w_t (1 - \tau_t^l) + \psi_l (1 + r_{t+1} (1 - \tau_{t+1}^k)) (l_{t+2} - l_{t+1}) = (c_t^e - c_t^u) (1 + \tau_t^c) + \psi_l (l_{t+1} - l_t) \quad (\text{C.2})$$

$$(f\epsilon_t)^\nu = \epsilon_t \frac{w_t (1 - \tau_t^l)}{1 + \tau_t^c}. \quad (\text{C.3})$$

On the other hand, profit maximization by firms implies:

$$w_t = \frac{\alpha_n y_t}{n_t f \epsilon_t}. \quad (\text{C.4})$$

Equations (C.1-C.4), the fact that n_t must equal l_t in all periods, and some algebra imply the following condition for effort:

$$\epsilon_t = \left(\frac{\alpha_n (1 - \tau_t^l) y_t}{(1 + \tau_t^c) n_t f} \right)^{\frac{1}{\nu}}. \quad (\text{C.5})$$

Finally, first order conditions for profit maximization together with equations (C.1-C.4) yield the following second-order difference equation for employment:

$$n_{t+2} = n_{t+1} + \frac{1 - \tau_t^l}{\psi_n} \alpha_n \left(1 - \frac{1}{\nu} \right) \frac{y_{t+1}}{n_{t+1}} (n_{t+1} - n_t) (1 + r_{t+1} (1 - \tau_t^k)) + \frac{\kappa}{\psi_n} (1 + \tau_{t+1}^c). \quad (\text{C.6})$$

Bibliography

- Atkeson, A., and P. J. Kehoe (1999), “Models of Energy Use: Putty-Putty vs. Putty-Clay”, *American Economic Review*, Vol. 89, pp. 1028-1043.
- Baxter, M., and D.D. Farr (2001), “Variable Factor Utilization and International Business Cycles”, NBER Working Paper No. 8392. Forthcoming *Journal of International Economics*.
- Bergoing, R., P. J. Kehoe, T. J. Kehoe, and R. Soto (2002), “Decades Lost and Found: Mexico and Chile Since 1980”, *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 26, pp. 3-30.
- Bils, M., and J. Cho (1994), “Cyclical factor utilization”, *Journal of Monetary Economics*, Vol. 33, pp. 319-354.
- Burnside C., and M. Eichenbaum (1996), “Factor Hoarding and the Propagation of Business-Cycle Shocks”, *American Economic Review*, Vol. 86, pp. 1154-1174.
- Burnside, C., M. Eichenbaum, and S. Rebelo (1993), “Labor Hoarding and the Business Cycle”, *Journal of Political Economy*, Vol. 101, pp. 245-273.
- Burnside, C., M. Eichenbaum and S. Rebelo (2001), “Hedging and Financial Fragility in Fixed Exchange Rate Regimes”, *European Economic Review*, Vol. 45, pp. 1151-1193.
- Burstein, A., M. Eichenbaum, and S. Rebelo, “Why Is Inflation So Low After Large Devaluations?”, NBER Working Paper No. 8748, January 2002.
- Calvo, G.A., “Balance of Payments Crises in Emerging Markets: Large Capital Inflows and Sovereign Governments”, in Paul Krugman (editor) *Currency Crises*, Chicago, University of Chicago Press, 2000.
- Cavallo, M., K. Kisselev, F. Perri, and N. Roubini (2004), “Exchange rate overshooting and the costs of floating”, New York University.
- Chari, V.V., P.J. Kehoe, and E.R. McGrattan (2003), “Accounting for the Great Depression”, *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 27 pp. 2-8.
- Chari, V.V., P.J. Kehoe, and E.R. McGrattan (2005), “Sudden Stops and Output Drops”, *Staff Report 353*, Research Department, Federal Reserve Bank of Minneapolis.
- Cogley, T., and J. M. Nason (1995), “Output Dynamics in Real-Business-Cycle Models”, *American Economic Review*, Vol. 85, pp. 492-511.
- Cole, H.L., and T.J. Kehoe (1996), “A Self-Fulfilling Model of Mexico’s 1994-1995 Debt Crisis”, *Journal of International Economics*, Vol. 41, pp. 309-330.
- Cook, D., and M. B. Devereux (2005), “Accounting for the East Asian Crisis: A Quantitative Model of Capital Outflows in Small Open Economies”, Forthcoming, *Journal of Money, Credit and Banking*.

- Cooley, T. F., and E. C. Prescott (1995), “Economic Growth and Business Cycles”, in T. F. Cooley, editor, *Frontiers of Business Cycle Research*, Princeton, NJ: Princeton University Press, pp. 1-38.
- Correia, I., J. C. Neves, and S. Rebelo (1995), “Business Cycles in a Small Open Economy”, *European Economic Review*, Vol. 39, pp. 1089-1113.
- Corsetti, G, P. Pesenti, and N. Roubini (1999), “Paper Tigers? A Model of the Asian Crisis”, *European Economic Review*, Vol. 43 , pp. 1211-1236.
- Dotsey, M., and C.S. Mao (1992), “How well do linear approximation methods work?”, *Journal of Monetary Economics*, Vol. 29, pp. 25-58.
- Fernandez, A. M., and I. Trigueros (2001), “Análisis, Evaluación y Propuestas para una Reforma Tributaria”, *Gaceta de Economía*, Suplemento, pp. 69-144.
- Flood, R.P., P.M. Garber, and C. Kramer (1996), “Collapsing exchange rate regimes: another linear example.” *Journal of International Economics*, Vol. 41, pp. 223-34.
- Gertler, M. S. Gilchrist, and F. Natalucci (2003), “External Financial Constraints on Monetary Policy and the Financial Accelerator”, NBER Working Paper 10128.
- Gollin, D. (2002), “Getting Income Shares Right”, *Journal of Political Economy*, Vol. 110, pp. 458-474.
- Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988), “Investment, Capacity Utilization and the Real Business Cycle”, *American Economic Review*, Vol. 78, pp. 402-417.
- Hansen, G. D. (1985), “Indivisible Labor and the Business Cycle”, *Journal of Monetary Economics*, Vol. 16, pp. 309-327.
- IMF (2004), “International Financial Statistics”, CD database.
- Kehoe, T.J., and E.C. Prescott (2002) “Great Depressions of the 20th Century ”, *Review of Economic Dynamics*, Vol. 5, pp. 1-18.
- Kydland, F.E., and E.C. Prescott (1982), “Time to Build and Aggregate Fluctuations”, *Econometrica*, Vol. 50, pp. 1345-1371.
- Kydland, F.E., and Zarazaga, C.E.J.M. (2002), “Argentina’s lost decade”, *Review of Economic Dynamics*, Vol. 5, pp. 152-65.
- Lahiri, Amartya, and C.A. Vegh (2002), “Output Costs, BOP Crises, and Optimal Interest Rate Policy”, University of California at Los Angeles.
- Lane, P.R., and G.M. Milesi-Ferretti (2001), “The external wealth of nations: measures of foreign assets and liabilities for industrial and developing countries”, *Journal of International Economics*, Vol. 55, pp. 263-294.

- Mendoza, E. G. (1991), "Real Business Cycles in a Small Open Economy", *American Economic Review*, Vol. 81, pp. 797-818.
- Mendoza, E. G. (2002), "Credit, Prices, and Crashes: Business Cycles with a Sudden Stop", in S. Edwards and J. Frankel, editors, *Preventing Currency Crises in Emerging Markets*, NBER.
- Mendoza, E. G., A. Razin, and L. Tesar (1994), "Effective Tax Rates in Macroeconomics: Cross-country Estimates of Tax Rates on Factor Incomes and Consumption", *Journal of Monetary Economics*, Vol. 34, pp. 297-323.
- Neumeyer, P. A., and F. Perri (2005), "Business Cycles in Emerging Economies: The Role of Interest Rates", *Journal of Monetary Economics*, Vol. 52.
- Rogerson, R. (1988), "Indivisible Labor, Lotteries and Equilibrium", *Journal of Monetary Economics*, Vol. 21, pp. 3-16.
- Sachs, J., A. Tornell, and A. Velasco (1996), "The Mexican Peso Crisis: Sudden Death or Death Foretold?", *Journal of International Economics*, Vol. 41 (1996), pp. 265-283.
- Young, A. (1995), "The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience", *Quarterly Journal of Economics*, Vol. 110, pp. 641-680.
- World Bank (2004), *World Bank Development Indicators*, CD database.

Figure 1: Output, inputs and TFP during recent financial crises




Figure 2: Exogenous shocks during Mexico's Tequila crisis

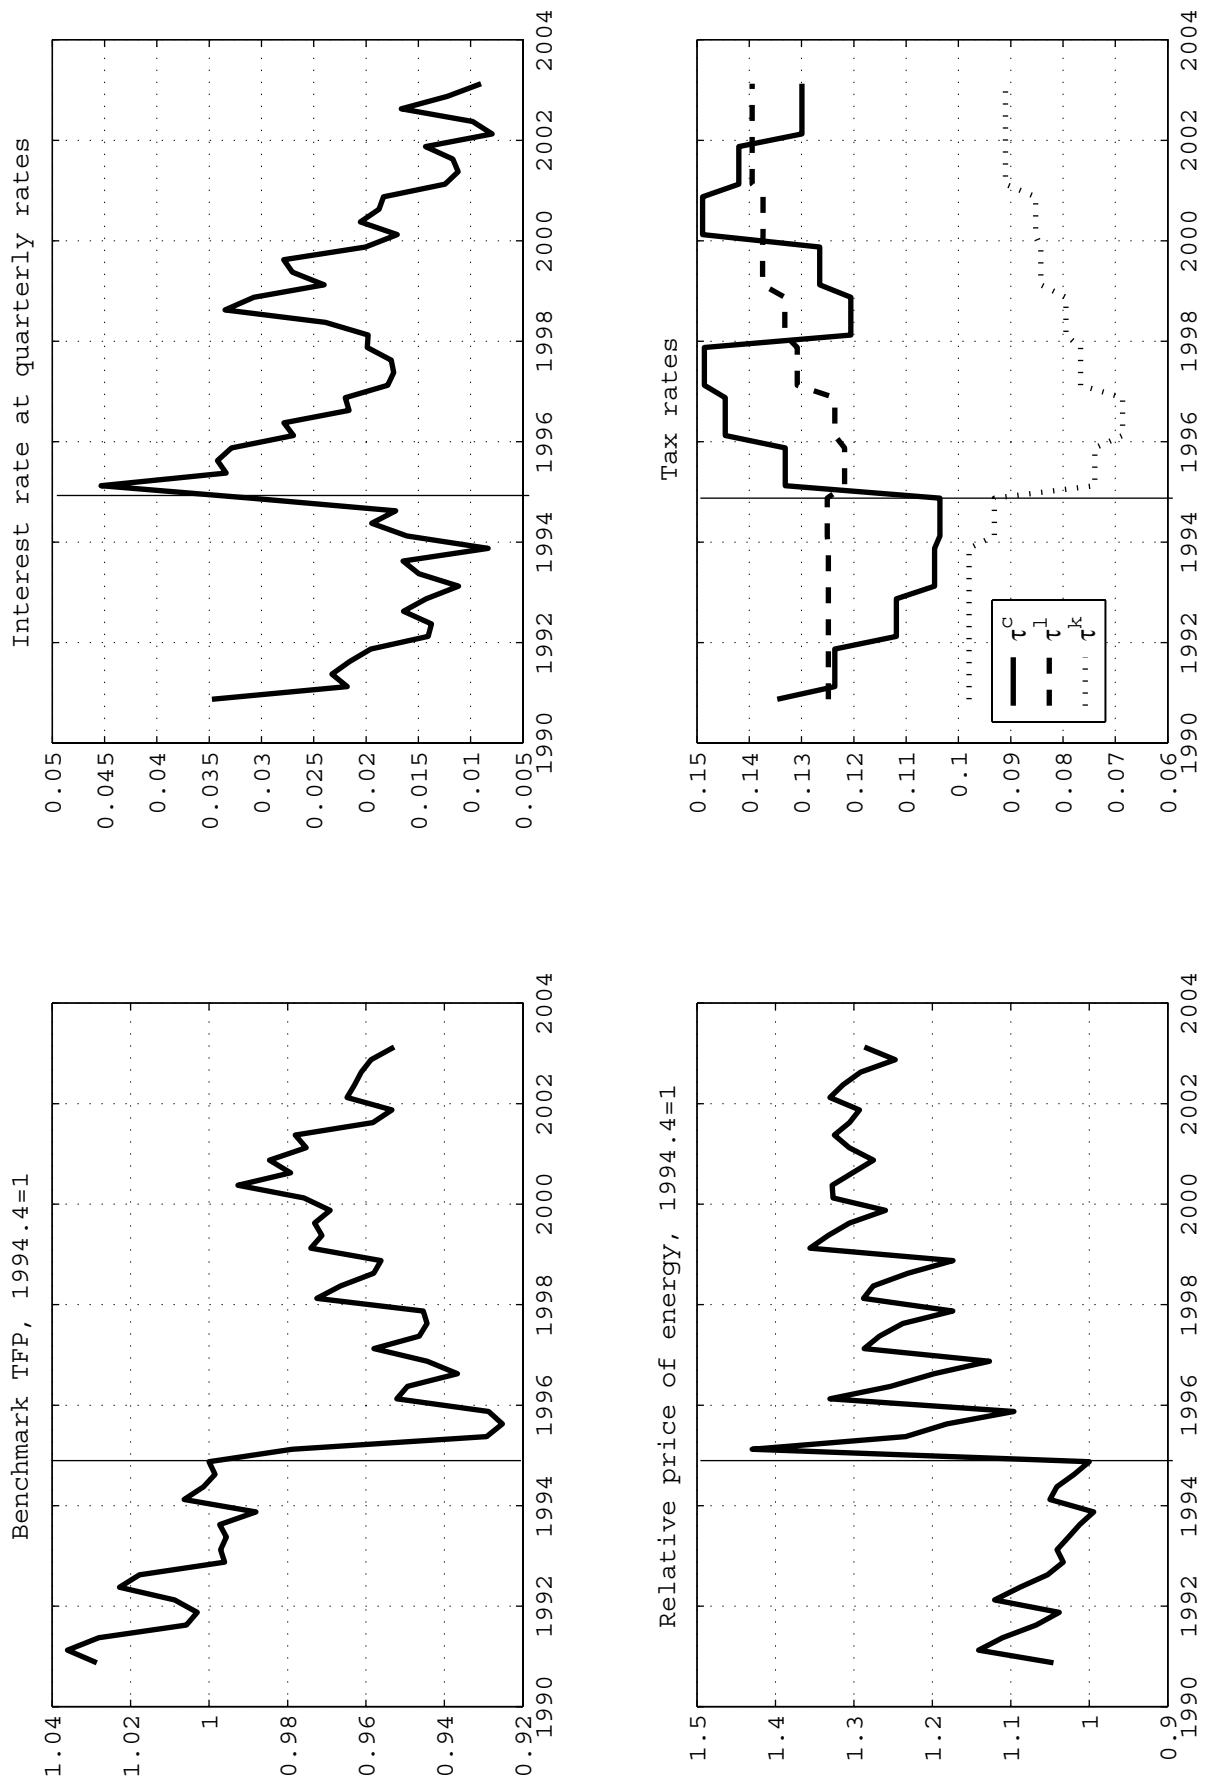


Figure 3: Predictions of the benchmark model

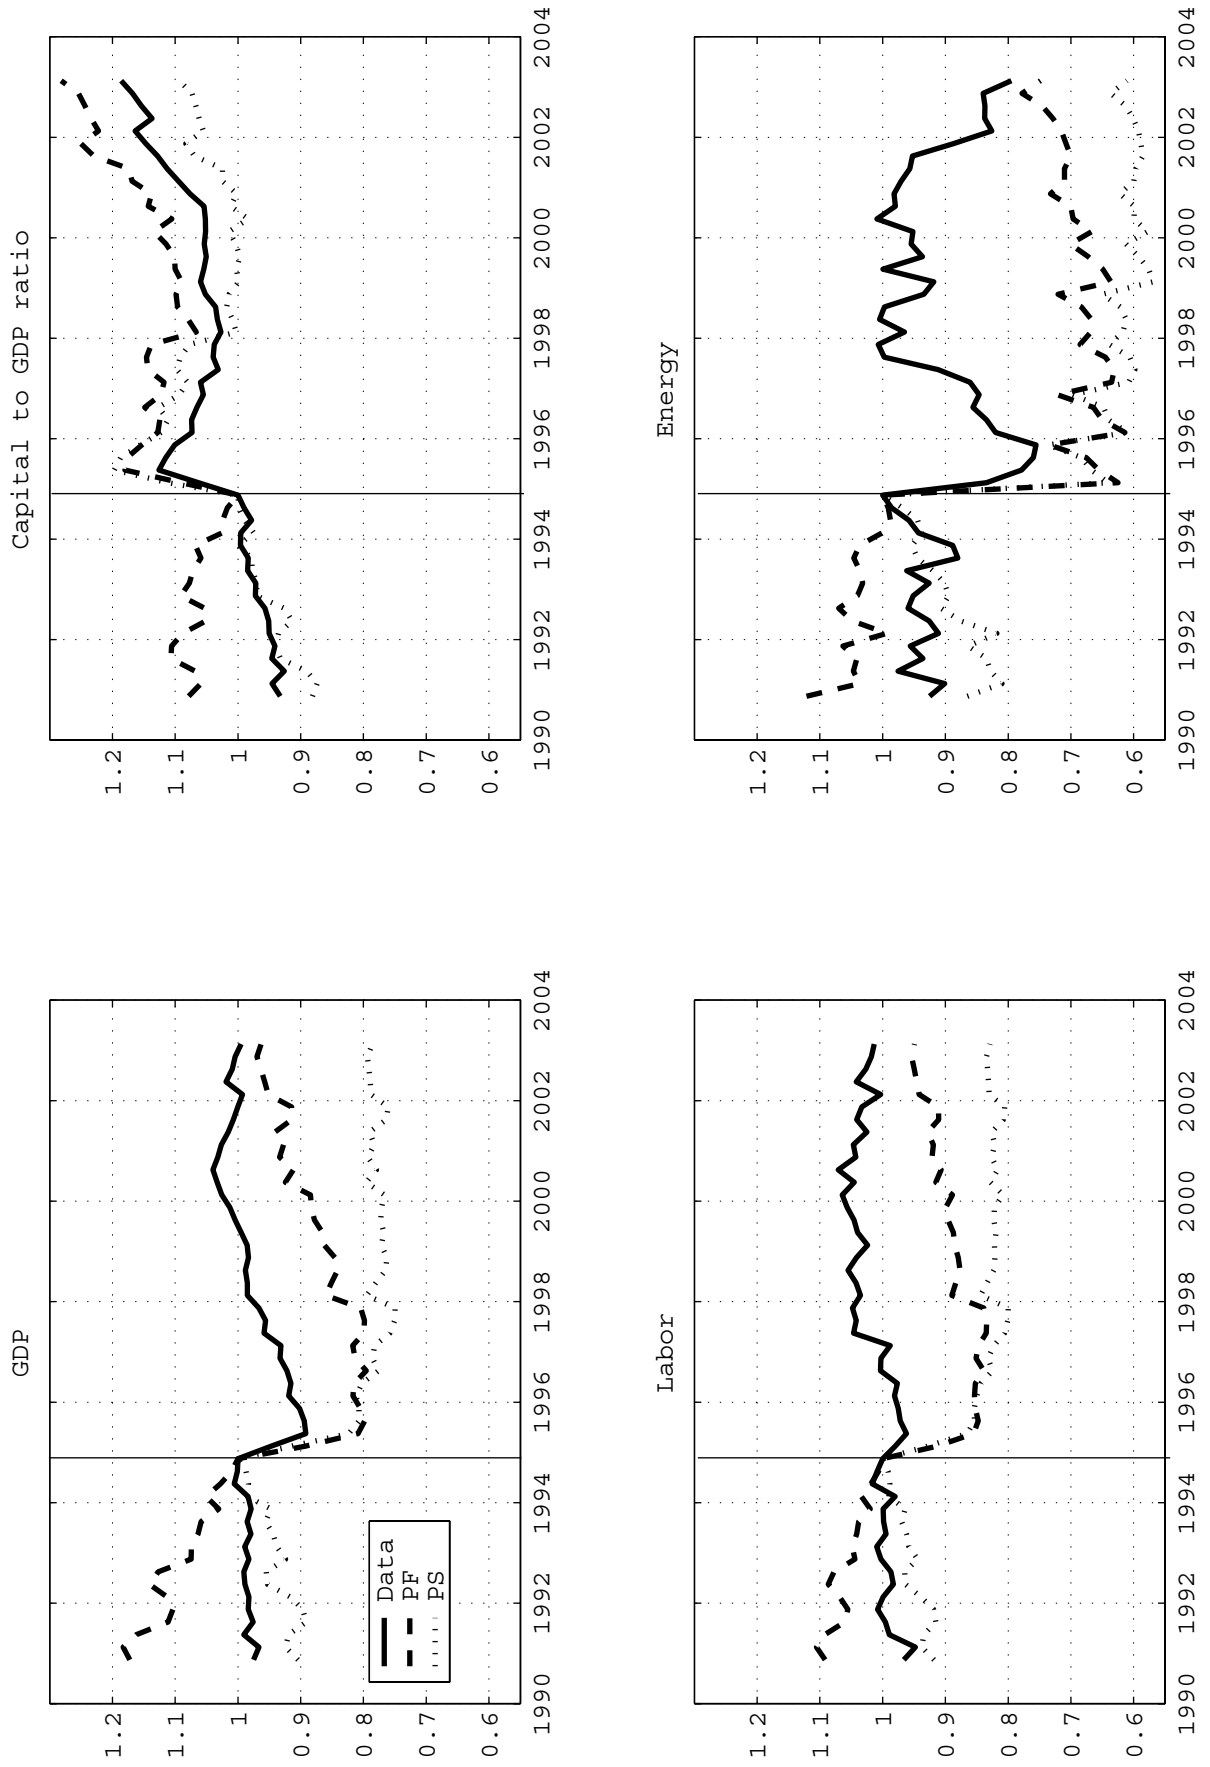


Figure 4: Benchmark model with low elasticity of labor supply

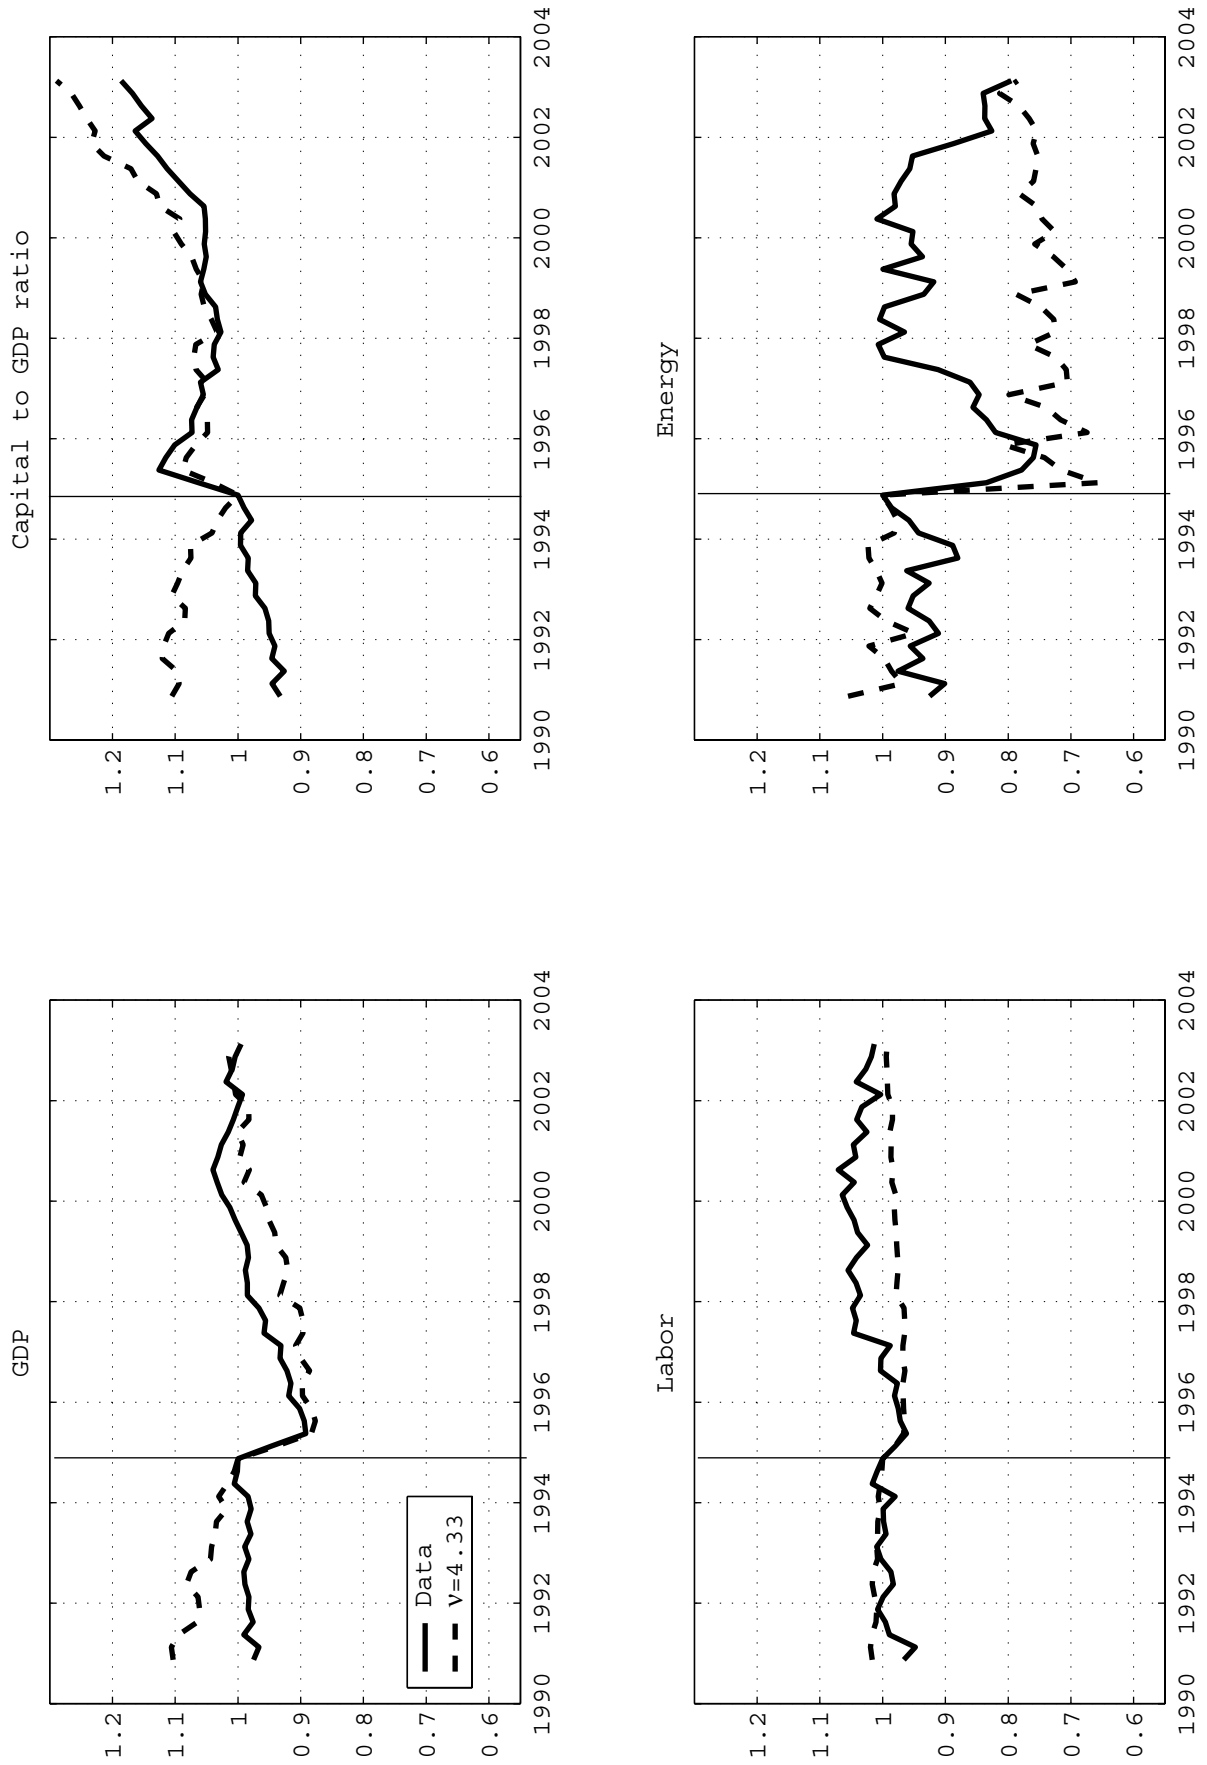


Figure 5: Logarithmic preferences




Figure 6: TFP adjusted for endogenous capital utilization and labor hoarding




Figure 7: Predictions of model with endogenous capital utilization

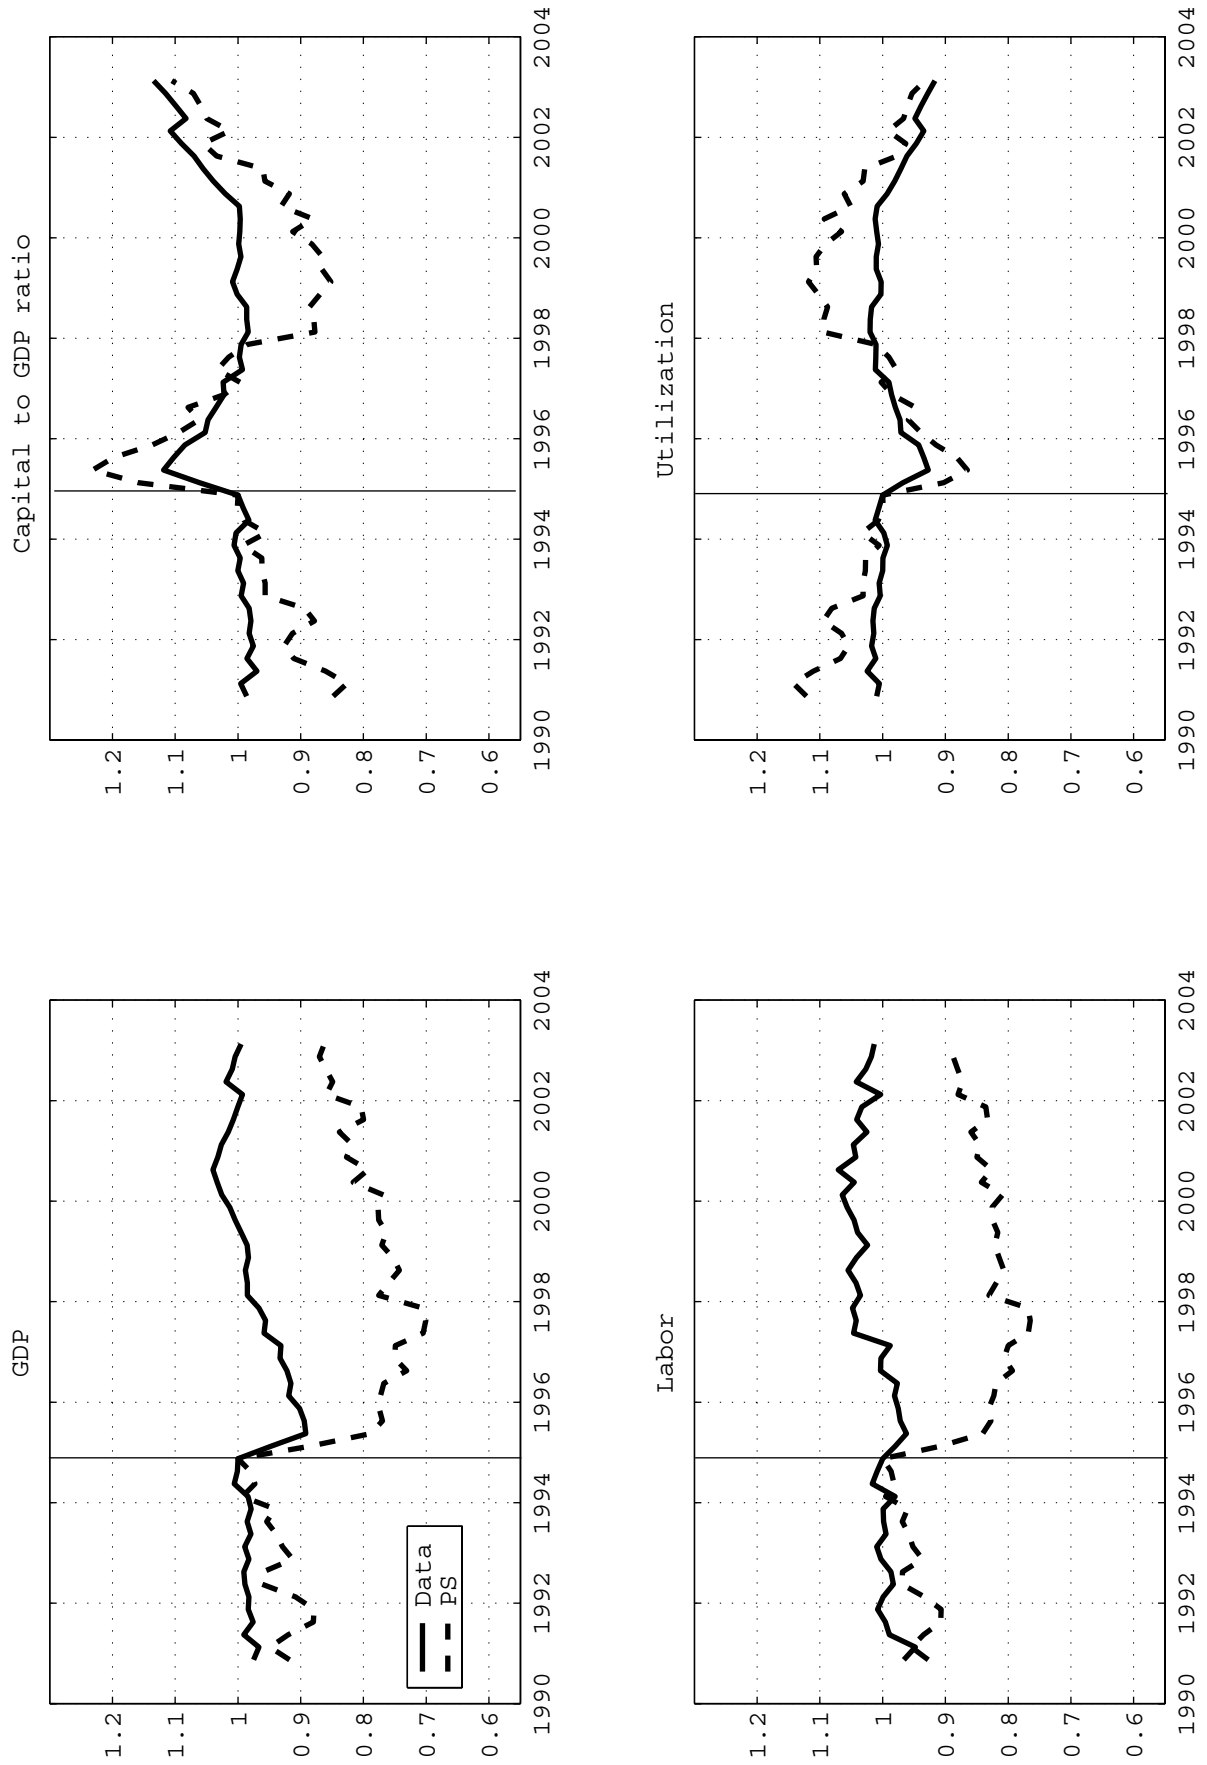


Figure 8: Predictions of model with labor hoarding

