An exploration into Pigou's theory of cycles [☆]

Paul Beaudry[a,b,c], Franck Portier[d,e,f],[*]

[a]University of British Columbia, Vancouver, B.C, Canada V6T 1Z2
[b]CIAR, Toronto, Ont., Canada, M5G 1ZB
[c]National Bureau of Economic Research, Cambridge, MA 02138, USA
[d]GREMAQ-IDEI-LEERNA, Université de Toulouse, 31000 Toulouse, France
[e]Institut Universitaire de France, 75005 Paris, France
[f]CEPR, London EC1V 7RR, UK

Abstract

This paper explores a theory of business cycles in which recessions and booms arise due to difficulties encountered by agents in properly forecasting the economy's future needs in terms of capital. The idea has a long history in the macroeconomic literature, as reflected by the work of Pigou (Industrial Fluctuation, MacMillan, London, 1926). The contribution of this paper is twofold. First, we illustrate the type of general equilibrium structure that can give rise to such phenomena. Second, we examine the extent to which such a model can explain the observed pattern of U.S. recessions (frequency, depth) without relying on technological regress. We argue that such a model offer a framework for understanding elements of both the recent U.S. recession and of the Asia downturns of the late 1990s.
© 2004 Elsevier B.V. All rights reserved.

1. Introduction

Equilibrium business cycle theory is often criticized on the ground that it does not provide a convincing theory of recession. In particular, it is well known that standard real business cycles models have difficulties explaining recessions[1]—at least of the size observed in Post War US data—without invoking technological regress.[2] This criticism is particularly relevant when examining the recession of 2001. Very few economic commentators regard the recession of 2001 as resulting from a negative technology shock. A more common view among economists is that the collapse of investment observed in 2001 resulted from some combination of changes in expectations about the profitability of new investments as well as a possible feedback from a period of very high investment in the late nineties and early in 2000.

The view that recession and booms may arise as the result of investment swings generated by agents' difficulties to properly forecast the economy's need in terms of capital has a long tradition in economics.[3] For example, this difficulty was seen by Pigou as being an inherent feature of any economy with technological progress. As emphasized in Pigou (1926), when agents are optimistic about the future and decide to build up capital in expectation of future demand then, in the case where their expectations are not met, there will be a period of retrenched investment which is likely to cause a recession. The object of this paper is to offer a formalization of this idea, which we call Pigou cycles, and to explore its quantitative plausibility as a theory of recessions. A key aspect of this paper is to explore the extent to which such a mechanism can explain the depth and frequency of recessions within an equilibrium framework where technological regress never occurs.

At first glance, the idea of a business cycle model where optimism and pessimism play a dominant role may appear counter to the notion of rational expectations. However, this will not be the case in our model. In effect, we consider an environment where agents get imperfect signals about future productivity growth and use these signals to make decisions about investment; knowing that the received signals are imperfect. The notion of optimism simply refers to a state where agents receive an above average signal. In this environment, periodic recessions are most likely to arise when agents signals about the future are precise. In effect, in our framework, occasional recessions are a sign of a well-functioning economy since they reflect the availability of good quality information upon which people act.

---

[1]Similarly, nominal-real confusion models (see Lucas, 1972) generate persistent downturns only if agents' ability to access price and money supply information is severely limited.

[2]See Kydland and Prescott (1982) or King et al. (1988). A notable exception is King and Rebelo (1999), where it is shown that a one sector business cycle model can explain business cycles with "a low probability of technological regress", provided that we are in a "high substitution economy", characterized by large elasticity of labor supply and elastic capacity utilization.

[3]This view have been recently surveyed by De Long (1991), and advocated by Black (1995) and Greenwood and Yorukoglu (1997).

The analysis conducted in this paper can be viewed as being complementary to the literature emphasizing how rational herding and information cascades may be important for understanding macroeconomic phenomena (see for example Banerjee, 1992; Bickhchandani et al., 1992; Chamley and Gale, 1994; Caplin and Leahy, 1993; Zeira, 1994). In particular, this strand of literature has emphasized how information may occasionally be aggregated improperly thereby leading to significant forecast errors that are shared by a large fraction of the population. The current paper adds to this research program by examining whether (rational/non-systematic) aggregate forecast errors can explain the observed pattern of recessions within a fully specified dynamic general equilibrium model. It should also be noted that the mechanisms at work in this paper are very close to those discussed in Phelps (1999). In this sense, this paper can be seen as offering a particular formalization to Phelps idea of structural booms and structural slumps.

The remaining sections of the paper are structured as follows. In Section 2 we illustrate why standard equilibrium models used in the macroeconomic literature do not produce typical business cycle co-movements as the result of forecast errors. In Section 3 we present an alternative model, which can be seen as either a three-sectors model or a particular two-sector model. We use the model to illustrate how Pigou cycles can arise in this dynamic general equilibrium framework as the result of forecast errors. In Section 4 we evaluate the quantitative properties of this model using standard numerical techniques. The main question addressed here is whether such a model can explain the observed depth and frequency of recession without invoking technological regress. Section 5 examines some of the price implications of the Pigou cycle story and Section 6 offers concluding comments.

2. Expectation driven fluctuations in standard equilibrium models

In this section, we want to highlight the type of fluctuations generated by expectations and forecast errors in standard one- and two-sector equilibrium models. Our main claim is that such models are incapable of generating Pigou cycles, that is they are incapable of generating equilibrium paths in which: (i) a forecast of future technological improvement first leads to a boom defined as an increase in aggregate output, employment, investment and consumption, and (ii) the realization that a forecast is too optimistic leads to a recession defined as a fall in all the same aggregate quantities.

We successively examine a baseline one-sector model, a one-sector model augmented with capital adjustment costs and a typical two-sector model. To most clearly present the properties of these models, we temporarily adopt a non-stochastic framework. This approach has the advantage that the dynamics of the models can be illustrated using impulse response functions to the announcement of a future permanent shocks that is eventually not realized. However, this framework has the disadvantage that forecast errors must be modelled as complete surprises. This drawback will be remedied in the following section where we embed our model in a stochastic setting where rational agents receive signals, make and revise forecasts,

and take decisions, knowing that received signals may be wrong. Our approach is to present the properties of each model using numerical examples. However, as we outline in Appendix A, the properties emphasized in this section are not sensitive to the particular parameter values used.

2.1. Expectations errors in a one-sector model

The first model we consider is a simple one-sector optimal growth model as used in the RBC literature. The equilibrium dynamics for this model is the solution of the following social planner's program:

$$\max_{C_t, l_t, K_t} \quad \mathrm{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ \frac{C_t^{1-\sigma}}{1-\sigma} + v_0(\bar{l} - l_t)^\eta \right\} \right]$$
$$\text{s.t.} \quad C_t + I_t = \theta_t l_t^\alpha K_t^{1-\alpha},$$
$$K_{t+1} = (1-\delta)K_t + I_t,$$

where C stands for consumptions, I for investment, K for capital, l for worked hours and $\bar{l}$ for total time endowment. In this model, as well as in the two following ones, the discount rate β is set to .98, the depreciation rate δ is set to .05 and the intertemporal labor supply elasticity is assumed to be infinite ($\eta = 1$). The share of labor income in total income is assumed to be 2/3, so that $\alpha = 2/3$. We consider two values for the intertemporal elasticity of substitution in consumption ($1/\sigma$), which are $\sigma = 1$ and 1/4. The reason for considering two values for σ is that they give rise to different impact responses for investment and consumption.

The exogenous process we feed into the model aims at capturing the type of expectational phenomena mentioned in the introduction. To this end, at period 1 the following announcement is made: θ will permanently increase by 1% in period $t + 4$. However, in $t + 4$ this technological improvement is not realized, so that the path of technology through the experiment is actually constant as individuals make a forecast error about growth.

Fig. 1 graphs the response of the one-sector economy to a technological announcement (news) that is eventually not realized. Let us first consider the case $\sigma = 1$. The good news instantaneously increases consumption through a wealth effect. As leisure is also a normal good, worked hours decrease, so does output. The only way consumption can be increased while hours are decreased is by decreasing investment: a good news creates an opposite movement in consumption and investment, and causes an output recession that lasts until period 4, when agents realize they were wrong. At that time, this bad news causes a drop in consumption, but a boom in investment, hours worked and output. Obviously, these dynamics do not correspond to what we are calling Pigou cycles. It is nevertheless possible to reverse the initial response of investment by assuming a high intertemporal elasticity of substitution in consumption, which is the case with $\sigma = 1/4$. However, it is still the case that consumption and investment move in opposite directions: a good news now
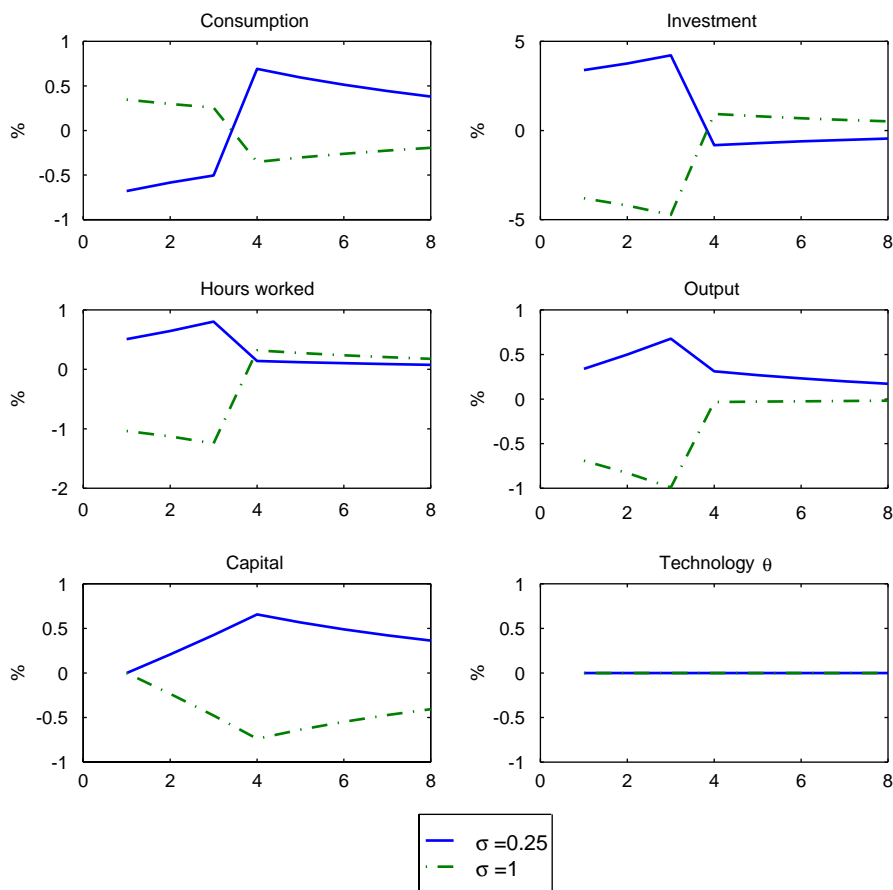
Fig. 1. Response of the one-sector economy without capital adjustment costs to an announcement at time 1 of future positive shock on the technology and no realization of that shock at time 4.

creates a consumption recession, while the bad news at the time of the revision creates a consumption boom. The induced negative co-movement between consumption and investment can be shown to hold for any set of parameters in such an economy (see Appendix A for a proof), and therefore this class of models cannot display Pigou cycles.

2.2. A one-sector model with capital adjustment costs

A natural extension of the above model is to introduce capital adjustment costs in hope of restoring a positive co-movement between consumption and investment. Let us explore this idea by considering a Tobin's q type model, whose equilibrium

dynamics is the solution of the following social planner's program:

$$\max_{C_t, l_t, K_t} \quad \mathrm{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ \frac{C_t^{1-\sigma}}{1-\sigma} + v_0(\bar{l} - l_t)^\eta \right\} \right]$$

$$\text{s.t.} \quad C_t + I_t = \theta_t l_t^\alpha K_t^{1-\alpha},$$

$$K_{t+1} = (1-\delta)K_t + \Phi\left(\frac{I_t}{K_t}\right) K_t.$$

We assume that the elasticity of Φ is .5, and that utility is log in consumption ($\sigma = 1$).

Fig. 2 shows that this extension does not allow the model to generate Pigou cycles. In effect, investment and hours worked do increase on impact, but consumption falls, while the reverse occurs at the time of the expectations revision. Again, it can be shown that this class of economies cannot generate expectational driven booms and bust with a positive co-movement between investment and consumption (see Appendix A).
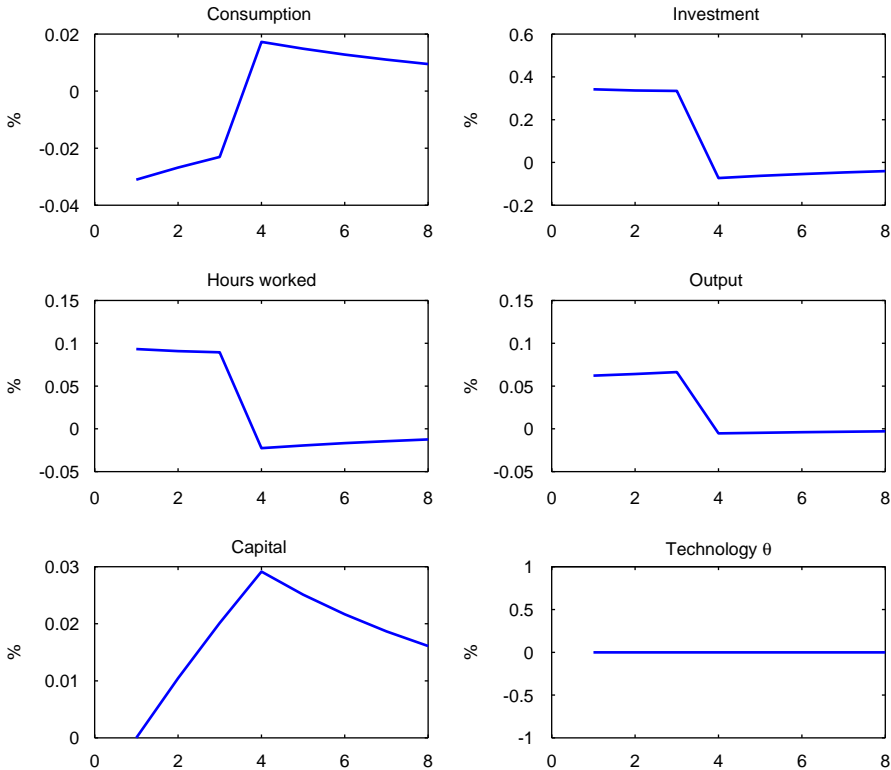

Fig. 2. Response of the one-sector economy with capital adjustment costs to an announcement at time 1 of future positive shock on the technology and no realization of that shock at time 4.

2.3. A typical two-sector model

We now consider now a simple two sectors optimal growth model. Its equilibrium dynamics is the solution of the following social planner's program:

$$\max_{C_t, l_{c,t}, l_{k,t}, K_{c,t}, K_{k,t}, K_t} \quad \mathrm{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ \frac{C_t^{1-\sigma}}{1-\sigma} + v_0 (\bar{l} - l_{c,t} - l_{k,t})^\eta \right\} \right]$$

$$\text{s.t.} \quad C_t = \theta_{c,t} l_{c,t}^{\alpha_c} K_{c,t}^{1-\alpha_c},$$

$$I_t = \theta_{k,t} l_{k,t}^{\alpha_k} K_{k,t}^{1-\alpha_k},$$

$$K_{t+1} = (1-\delta) K_t + I_t,$$

$$l_{c,t} + l_{k,t} \leqslant \bar{l},$$

$$K_{c,t} + K_{k,t} \leqslant K_t.$$

The subscript c stands for the consumption good sector and the subscript k for the capital good sector. The values of the discount rate β and the discount factor δ are the same as before. We take as a benchmark the case of log utility of consumption ($\sigma = 1$) and linear disutility of labor ($\eta = 1$). We assume $\alpha_c = .59$ and $\alpha_k = .9$ so that, in conjunction with $\theta_c = \theta_x = 1$, the labor share is 2/3 and consumption's share in total output is 75% at the steady state. Note that we are imposing very little decreasing returns to labor in the capital good sector as to potentially allow this sector to react substantially to news.

We successively consider shocks to θ_k and θ_c. The response of the economy to these two shocks is presented in Fig. 3. Two observations can be made. First, the announcement of a future increase of productivity in the investment sector (an expected increase in θ_k) moves consumption up, but investment, hours and output down from period 1 to 3. The opposite is true later on. Again, a good news creates a recession, while the bad one causes an expansion: the model does not display Pigou cycles. Second, the economy does not respond at all to an expected increase in productivity for the consumption sector (θ_c). This later result comes from the fact that with log preferences, wealth and substitution effect exactly offset following this announcement.[4]

In order to move away from the knife edge case in the presence of an expected change in θ_c, we allow for more or less curvature in the utility of consumption ($\sigma = 2$ or 1/2). As shown in Fig. 4, we once again do not obtain from the model a joint increase of consumption, investment, output and hours following a good news, and an aggregate recession following a downward revision of expectations. In effect, we find such a result to be maintained for all plausible configuration of parameters.

In summary, this section has illustrated why standard equilibrium models do not capture the type of dynamics discussed by Pigou. However, this does not imply

---

[4]The result is similar to what we would get in the case of a log-utility, Cobb–Douglas-technology and full-depreciation one-sector model. In this configuration, one can analytically solve the model, and prove that labor and the ratio C/I are constant along any equilibrium path. Therefore, the economy does not respond to a technological news.

Fig. 3. Response of the two-sectors economy to an announcement at time 1 of future positive shock on θ_k or θ_c and no realization of that shock at time 4.

that it is an impossibility inherent to all general equilibrium models. In effect, in the next section we will present a simple equilibrium model that can support Pigou cycles.

3. The model

The object of this section is to present what we believe is the simplest dynamic general equilibrium model in which Pigou cycles can arise. The model can be thought as either a three-sector model or as a particular two-sector model. We will present it as a three-sector model.
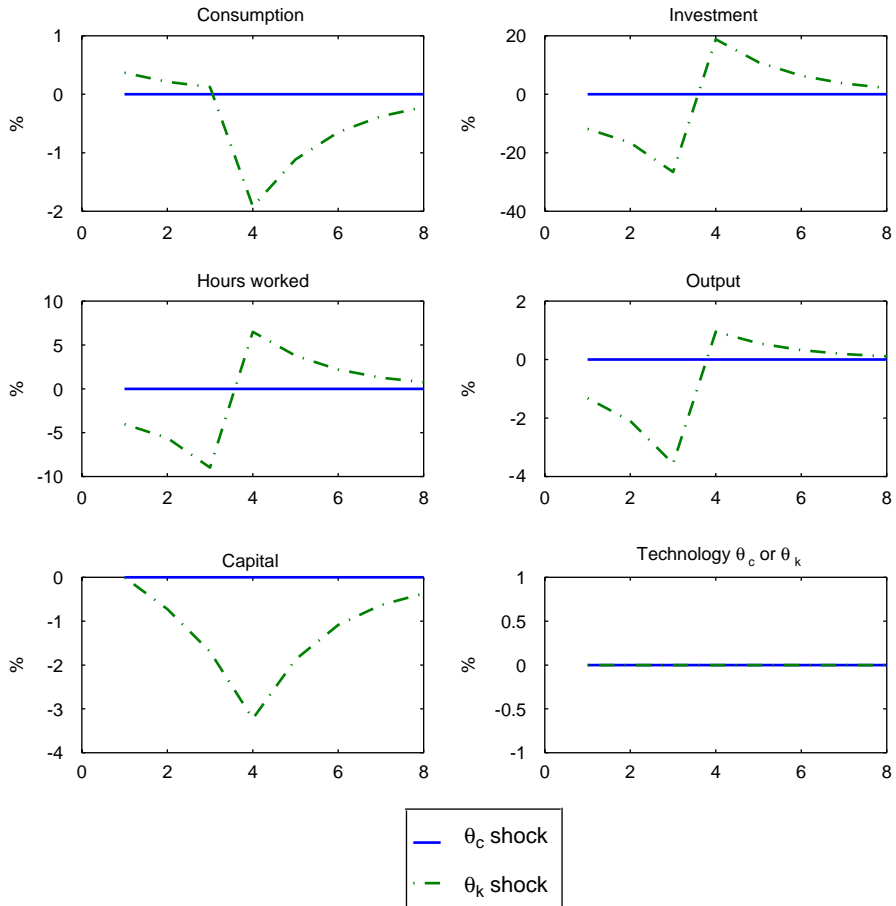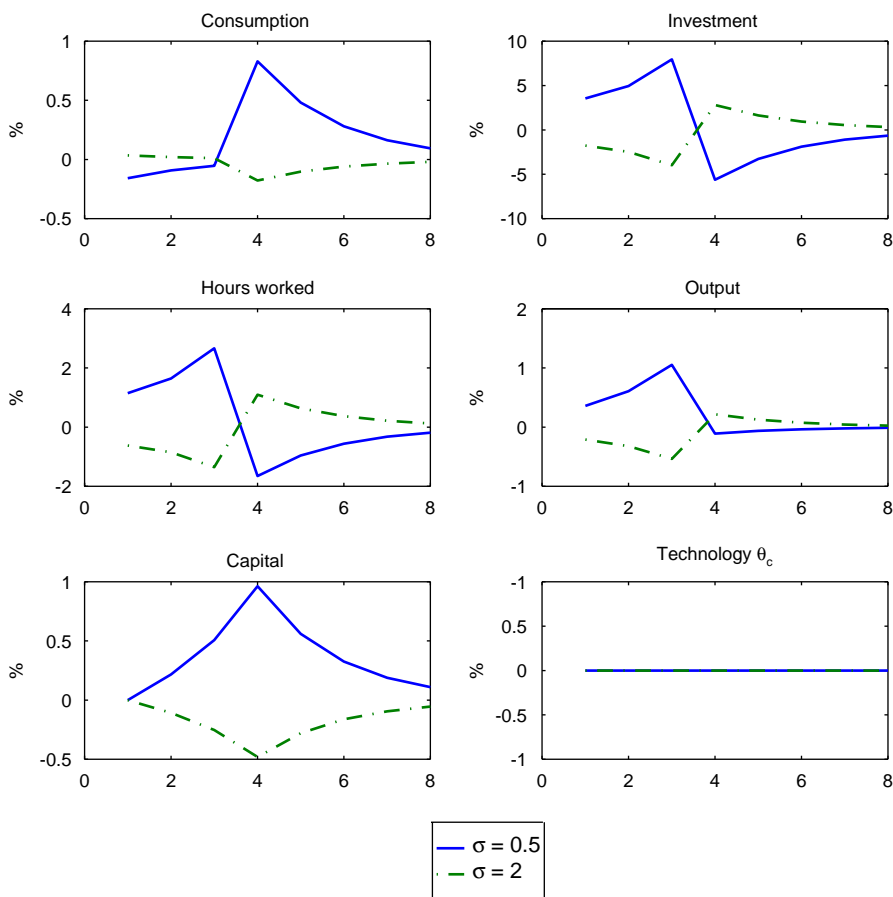
Fig. 4. Response of the two-sectors economy to an announcement at time 1 of future positive shock on θ_c and no realization of that shock at time 4.

3.1. The production sector

The economy is composed of three sectors: a final consumption goods sector, a non-durable goods (or intermediate good) sector and a durable goods sector. The durable good sector is best thought as the construction industry with the stock of the durable good representing plant and housing infrastructure. The final good, denoted C_t, is produced as CES composite of the nondurable good (or service) X_t and the stock of infrastructure K_t:

$$C_t = (aX_t^v + (1-a)K_t^v)^{1/v}, \quad v \leqslant 0.$$

The final good C_t is a flow of consumption services, which could be modelled as being either produced inside the household (by households purchasing X_t and K_t) or

in the market. For the sake of concreteness, we choose to treat C_t as being produced in the market.

The non-durable good X_t is produced using labor according to:

$$X_t = \theta_{x,t} l_{x,t}^{\alpha_x} \; \tilde{l}_x^{(1-\alpha_x)}, \quad 0 < \alpha_x \leqslant 1,$$

where $\theta_{x,t}$ is the state of technology in the non-durable goods sector and $l_{x,t}$ is the level of employment in this sector. $\tilde{l}_x$ represents a fixed factor that is required in production. The introduction of the fixed factor assures that overall returns to scale are constant, but forces returns to scale in the variable factor to be decreasing.

The capital good accumulates according to

$$K_{t+1} = (1 - \delta)K_t + I_t,$$

where δ is the rate of depreciation and I_t is investment which is provided by the construction sector. Production in the construction sector depends on the state of technology in this sector, $\theta_{k,t}$, the levels of employment $l_{k,t}$ and a fixed factor $\tilde{l}_k$.

$$I_t = \theta_{k,t} l_{k,t}^{\alpha_k} \; \tilde{l}_k^{(1-\alpha_k)}, \quad 0 < \alpha_k \leqslant 1.$$

We will restrict attention to cases where the elasticity of substitution between K_t and X_t in the final goods sector is no greater that one (which seems reasonable given our interpretation of K_t as infrastructure). Obviously, both the intermediate good sector and the construction sector should have production technologies which use both physical capital (machines) and labor. However, in order to make our model concise we exclude this possibility and instead introduce fixed factors. This simplifies exposition greatly since it allows us to remain in the family of models with only one capital stock.

3.2. The household sector

The representative household has preferences defined over consumption of the final good and over the labor supplied in each of the two sectors. The household's objective is to maximize:

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t \{ \log (C_t) + v_0(\bar{l} - l_{x,t} - l_{k,t}) \} \right],$$

where C_t is the level of consumption of the final good, $\bar{l}$ is the endowment of labor available in each period, β is the discount factor and v_0 is a positive constant. Note that household preferences are assumed to be separable in consumption and in leisure and that, following Hansen (1985) and Rogerson (1988), we assume that preferences are linear with respect to labor at the representative agent level. The household's within period budget constraint is

$$C_t + p_t I_t = w_{x,t}(l_{x,t}) + w_{k,t}(l_{k,t}) + r_t K_t + \Pi_{x,t} + \Pi_{k,t},$$

where the final good C_t is the numéraire, p_t is the price of capital, r_t is the rental rate of capital, $w_{x,t}$ and $\Pi_{x,t}$ are, respectively, the wage rate and returns to the fixed factor

in the intermediate goods sector, and finally $w_{k,t}$ and $\Pi_{k,t}$ are the wage rate and returns to the fixed factor in the construction sector.

3.3. Equilibrium

A Walrasian Equilibrium for this economy is a set of time paths for K, l_x, l_k, C, r, p, w_x and w_k such that (1) allocations are optimal given prices (that is, consumers maximize utility and firms maximize profits) and (2) markets clear. Given an initial capital stock K_0 and processes for $\theta_{x,t}$ and $\theta_{k,t}$, equilibrium allocations for this economy can be found by solving the following social planner's problem:

$$\max_{C_t, l_{x,t}, l_{k,t}, K_t} \quad E_0 \left[\sum_{t=0}^{\infty} \beta^t \{ \log(C_t) + v_0(\bar{l} - l_{x,t} - l_{k,t}) \} \right]$$

$$\text{s.t.} \quad C_t = (a(\theta_{x,t} l_{x,t}^{\alpha_x} \tilde{l}_x^{(1-\alpha_x)})^v + (1-a) K_t^v)^{1/v},$$

$$K_{t+1} = (1-\delta)K_t + \theta_{k,t} l_{k,t}^{\alpha_k} \tilde{l}_k^{1-\alpha_k},$$

$$l_{x,t} + l_{k,t} \leqslant \bar{l}.$$

3.4. Processes for technology

Our objective here is to examine whether the above model, once embedded in a stochastic setting with imperfect signals and rational expectations, can display Pigou cycles and mimic some of the quantitative features of recessions; in particular, their frequency and depth. As a simplifying assumption, we will assume that technology grows stochastically only in the non-durable goods sector, while it grows in a deterministic way in the construction sector. This choice appears reasonable given that expectations about technological improvements in the construction sector do not stand out as an important driving force behind business cycles. Moreover, technological improvements in the non-durable goods sector can be interpreted (and formalized) as the arrival of new differentiated goods in an economy with tastes for variety. In this interpretation, it is the expected arrival of new goods and its associated infrastructure requirements which would lead a Pigou cycle.

We want to examine an economy where (i) technology only improves (never regresses), (ii) the economy exhibits balanced growth in the long-run and (iii) technological progress is stochastic only in the non-durable good sector. To this end, we assume technology in the construction sector grows deterministically according to

$$\log \theta_{k,t} = g_{0,k} + g_1 t,$$

while technology in the intermediate goods sector evolves stochastically according to:

$$\log \theta_{x,t} = g_{0,x} + g_1 t + \log \hat{\theta}_{x,t},$$

$$\log \hat{\theta}_{x,t} = \lambda \log \hat{\theta}_{x,t-1} + \varepsilon_t, \quad 0 < \lambda < 1,$$

where ε is a zero mean i.i.d. random variable. We assume that ε_t can either take on a high or a low level. The low level corresponds to the absence of technological progress, which in the case where $\lambda = 1$ corresponds to $\varepsilon = -g_1$. The respective probabilities of the high and low states are $1 - p$ and p. With the restriction that ε has a zero mean, p fully characterizes the distribution of technological innovations. In effect, these restrictions imply that ε_t takes on the value $g_1 \times p/(1 - p)$ in the growth state. When λ is very close but smaller than 1, which will be the case in our simulations ($\lambda = .999$), this parameterization of ε guarantees that technological regressions almost never regresses. In effect, we have checked that technology indeed never regresses in our simulations.

3.5. Information structure

Whereas we adopt a process for technology that is fairly standard, we want to depart from the RBC literature by allowing for a richer informational structure. To this end, in every period we allow agents to observe, in addition to ε_t, an i.i.d. zero-mean signal v_t. The signal v_t is assumed to bring new information on the j-periods ahead technological innovation ε_{t+j}, for $0 < j \leqslant n$. That is, the signal brings information on the growth of technology between time t and $t + n$. Without loss of generality, the information content of the signals can be summarized by a matrix of correlation between the ε_{t+j} and v_{t-k}, where $0 < j \leqslant n$ and $0 < k \leqslant n$. The (j, k) element of the matrix M, denoted m_{jk} represents the correlation between the current signal v_{t-k} and the future technological innovation ε_{t+j}.

In every period t, given the information set $\{\theta_x(t - 1), \varepsilon_t, v_t, \ldots, v_{t-n+1}\}$ and the correlation matrix M, agents form rational expectations on the sequence of future productivity levels $\hat{\theta}_x(t + j)$, $j > 0$. In the following period, given the new information in contemporaneous ε and v, expectations on $\hat{\theta}_x(t + 1 + j)$, $j > 0$ are revised. Note that the standard information structure used in most of the literature is a special case of this structure where v_t is uninformative but nevertheless, given the information revealed in period t by ε_t, the expectation on the sequence of $\hat{\theta}_x(t + j)$, $j \geqslant 1$ is updated every period. In our model, we simply allow the agents to have more information each period than that contained in $\{\theta_x(t - 1), \varepsilon_t\}$.

A key issue is how best to specify the joint distribution of v_{t-k} and ε_{t+j}. Our approach is to favor extreme parsimony in that we would like to summarize the matrix M by one parameter. We choose to restrict the number of parameters this way so that the fit of the model cannot be attributed to simple over-parameterization. Accordingly, we assume that v_t only brings information on the n-periods ahead innovation. In effect, we assume that the signal can take only the values $-g_1$ and $g_1 \times p/(1 - p)$, with q being the probability that signal is right about the n-step ahead technological innovation ($v_t = \varepsilon_{t+n}$) and, $1 - q$ being the probability that the signal is uninformative in the sense of being an i.i.d. draw on the two states. This process nicely captures the idea that signals may sometimes be entirely void of information, as suggested in the herding literature (e.g. Banerjee, 1992). The attractive feature of this process is that it allows the matrix M to take the following

simple form:

$$
\begin{array}{c}
\\
\varepsilon_{t+n} \\
\\
\varepsilon_{t+n-1} \\
\\
\vdots \\
\\
\varepsilon_{t+n-j} \\
\\
\vdots \\
\\
\varepsilon_{t+1} \\
\\
\varepsilon_{t}
\end{array}
\begin{array}{ccccccc}
v_t & v_{t-1} & \cdots & v_{t-j} & \cdots & v_{t-n+1} & v_{t-n} \\
\left(\begin{array}{ccccccc}
q & 0 & \cdots & 0 & \cdots & 0 & 0 \\
0 & q & \cdots & 0 & \cdots & 0 & 0 \\
\vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & \cdots & q & \cdots & 0 & 0 \\
\vdots & \vdots & \cdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & \cdots & 0 & \cdots & q & 0 \\
0 & 0 & \cdots & 0 & \cdots & 0 & q
\end{array}\right)
\end{array}.
$$

Given the two states process that we have assumed for ε, the economy can therefore go through one of the following four realizations of signal and subsequent growth: a growth signal at time t which is validated by technological growth at time $t + n$ (probability $q(1 - p) + (1 - q)(1 - p)^2$); a growth signal at t but no realized growth at time $t + n$ (probability $(1 - q)p(1 - p)$); a no-growth signal at time t but a growth realization at $t + n$ (probability $(1 - p)(1 - q)p$); a no-growth signal at time t and a no-growth realization at time $t + n$ (probability $pq + (1 - q)p^2$). Note that under this information process, the expected levels of technology are weekly increasing over time even though the expected innovations are not necessarily monotonic.

4. A quantitative evaluation of the model

4.1. Calibration

In our calibration exercise, our goal is not to suggest that a three-sector model is a fully adequate description of the economy. In effect, we believe that the above model is an extreme simplification of reality and that it omits many important elements (for example: adjustment costs, variable rates of factor utilization, inventories, additional capital stocks). Nonetheless, we believe that a calibration exercise is useful for evaluating whether the theoretical mechanism by which this model produces booms and recessions (in the absence of technological regress) can be considered quantitatively relevant. Hence, we will examine whether a reasonably calibrated version of the model can reproduce the certain observed pattern of recessions—especially frequency and depth—while simultaneously capturing the variances and co-movements emphasized in much of the modern business cycle literature. Throughout this exercise, we will interpret a time period as representing 6-months. The reason for adopting a semester as our unit of time is that it allows a decline in output in the model to be referred to as a recession.

There are several parameters in our model, some of which do not have immediate counterparts in the literature. Therefore we approach this calibration exercise by first setting parameters (as is most commonly done) based on known estimates or based on matching certain steady-state properties. We estimate the remaining parameters using a simulated method of moments technique. In particular, the discount factor β is set equal to .98, the depreciation rate δ is set to .05. Total disposable time $\bar{l}$ is normalized to 2, and the disutility of labor scale parameter v_0 is set to 1, so that one third of total time is devoted to work in the steady state. The average growth factor of productivity is set to its observed level in our sample period (see below for a description of the data set). The ratio $\theta_{0,x}/\theta_{0,k}$ and the relative weight of K and X in the CES production function, that is the parameter a, are set so that, in conjunction with the other parameters, the labor share is 66% and consumption's share in total output is 75%.[5] We also need to set values for the short-run returns-to-labor parameters α_x and α_k. The literature on scale parameters suggests that the short-run returns to labor are close to the labor share in output.[6] However, the literature on the construction industry arrives at a somewhat different conclusion. Allen (1985), for example, estimates the short-run return to labor in construction to be very close to one. In order to reflect these two considerations, we set $\alpha_x = .6$ and $\alpha_k = .97$, that is we impose very little decreasing returns to labor in the construction industry.[7]

There remain four parameters that we cannot infer from previous studies, namely the two parameters governing the technology and information processes (p, q), the technological parameter v,[8] and the number of periods n between the arrival of a signal and the related realization of ε. Therefore we choose to estimate these four parameters by Simulated Method of Moments.[9] We implement this procedure by finding, for different values of n, a vector $\pi = (p, q, v)$ that provides the best match for the following six moments: the volatilities of output, consumption and investment calculated for both the Hodrick–Prescott (*HP* filter) cyclical

---

[5]It should be noted that total output (GDP) in our model is calculated as the sum of the flow of consumption services C_t plus the value of investment $p_t \times I_t$. The production of the non-durable good is treated as an intermediate input.

[6]See for example Burnside et al. (1995).

[7]These values for α_x and α_k can also be justified based on some of the results found in Burnside et al. (1995). In particular, when focusing on industries for which there are good direct measures of output, Burnside, Eichenbaum and Rebelo estimate the short run return to labor in durable manufacturing to be .98, while the counterpart for non-durable manufacturing is estimated to be .61 (see their last columns of Table 10). Clearly, these are estimates offer an alternative justification for the returns to scale parameters that we use in our calibration.

[8]Note that the parameter representing the elasticity of substitution between capital and non-durable goods is for the final goods production function. Given that this production function describes the process of aggregating goods and services into a final flow of consumption goods, it does not seem appropriate to set it based on estimates derived from industry studies.

[9]Roughly speaking, simulated method of moments consists of choosing those model parameters values that produce the best match between a set of empirical and simulated moments, where the distance between those moments is evaluated using the inverse of a consistent estimate of the moments estimators asymptotic variance matrix. See Duffie and Singleton (1993) for an exposition and Hairault et al. (1997) for an application.

components [10] and growth rate ($1 - L$ filter). We denote this set of six moment by $M^\mathrm{o} = (\sigma_y, \sigma_c, \sigma_i)_{HP,1-L}$. The HP moments are chosen for reasons of comparability with previous studies, while our interest in "classical cycles" (cycles in terms of growth rate) suggests the use of the $1 - L$ filter. Let us denote by Ω the variance-covariance matrix of these estimators. For a given vector of parameters π and n, we simulate the model N times for T periods ($N = 20$ and $T = 77$ [11]) and compute a vector of simulated moments $M^\mathrm{s}(\pi)$. We performed the simulations using a log-linearized approximation of the model (around its (locally) unique steady state). The estimate of π is then

$$\hat{\pi} = \operatorname{Arg\,Min}_\pi \quad J = \frac{NT}{NT+1}(M^\mathrm{s}(\pi) - M^\mathrm{o})\Omega(M^\mathrm{s}(\pi) - M^\mathrm{o})'.$$

We estimated a π vector for each of the ten cases where n was allowed to vary between 1 and 10. We then chose n and the corresponding π vector based on the lowest value for the J statistic.

The data we use are US National Income and Product Account data covering the period 1959 to the end of 1997. We build the relevant empirical counterparts to our theoretical constructs in the following manner. Durable goods and inventories are considered investment, and net exports are split into consumption and investment according to the relative share of consumption and investment. More precisely, the three series are constructed as follows: Investment (I) = Fixed investment + Durable goods + Change in business inventories + Net export of good and service $\times(i/y)$, Consumption (C) = Non-durable goods + Services + Net export of good and service $\times(c/y)$, Output (Y) = Consumption + Investment. Variables are then expressed in per capita terms. Estimation results for the case where $n = 2$ (which corresponds to a minimum for the J statistic) [12] are given in Table 1, and the model's predictions relative to the targeted moments are given in Table 2.

The results from the estimation using simulated method of moments implies an economy where (i) infrastructure K and other goods X are strong complements (elasticity of substitution close to .2), (ii) agents receive rather informative signals, that is, signals are right 82% of the time, (iii) technological growth is quite sporadic with 71% of semesters registering no technological progress and 29% percent of semesters registering growth of 4.17%, and (iv) the delay between signals and realizations is one year (2 periods or semesters). It is interesting to note that, under the null that the model is the Data Generating Process and that $n = 2$, the J statistics (which conditional on n would follow a $\chi^2(3)$) cannot be rejected at a 5% level ($\chi^2(3)$ at 95% is 7.8). Although this is not an appropriate test since we are choosing n to minimize J (in which case it is more appropriate to compare J with a $\chi^2(2)$ distribution), it nevertheless suggests that this simple three-sector model can fit these data surprisingly well.

---

[10] In calculating HP filtered moments we set $\lambda = 800$ since this appeared to give reasonable cyclical components to other semestrial data.

[11] $T = 77$ corresponds to the length of our sample.

[12] We did not find any significant differences for J calculated using $n = 1$ or 2. However, the value of J does increase substantially for $n > 2$.

Table 1
SMM estimators of p, q and v (standard-deviations in parenthesis)

p	.71 (.04)
q	.82 (.31)
v	-3.78 (1.21)
n	2
J	3.30
$\chi^2(2)$ at 95%	5.99

Table 2
Targeted and simulated moments

	U.S. data	Model simulation
σ_c (HP)	1.055 (.087)	1.060
σ_c $(1-L)$	0.687 (.059)	.714
σ_y (HP)	2.162 (.222)	1.825
σ_y $(1-L)$	1.438 (.125)	1.477
σ_i (HP)	6.872 (.669)	5.742
σ_i $(1-L)$	4.996 (.488)	5.100

4.2. A first look at the model's quantitative properties

We illustrate here key properties of the model using impulse response functions derived from our quantified version of the model. However, it should be noted that most of the key properties of the equilibrium dynamics can be derived analytically, and are not dependent on the particular values derived above.[13]

Transitional dynamics : In Fig. 5, we plot the transitional dynamics of our model when the capital stock begins at 1% below its steady state. Fig. 5 shows that employment in the construction sector begins above its steady state and gradually converges to it, which allows capital to be built up. During this transition, employment in the non-durable goods sector is below its steady-state level, as some labor resources are shifted to the construction sector.

There are two more aspects to notice from Fig. 5. First, along the transition path, the aggregate level of employment (as defined by $l_x + l_k$) could be either above or below its steady state level, depending on the values of the parameters. Second, these dynamics are qualitatively similar to those derived for the one-sector model generally used in real business cycle models. However, the dynamics of this model will differ from those of the more standard model when anticipated technological change is introduced.

---

[13]See the working paper version of this work (Beaudry and Portier, 1999) for more details.
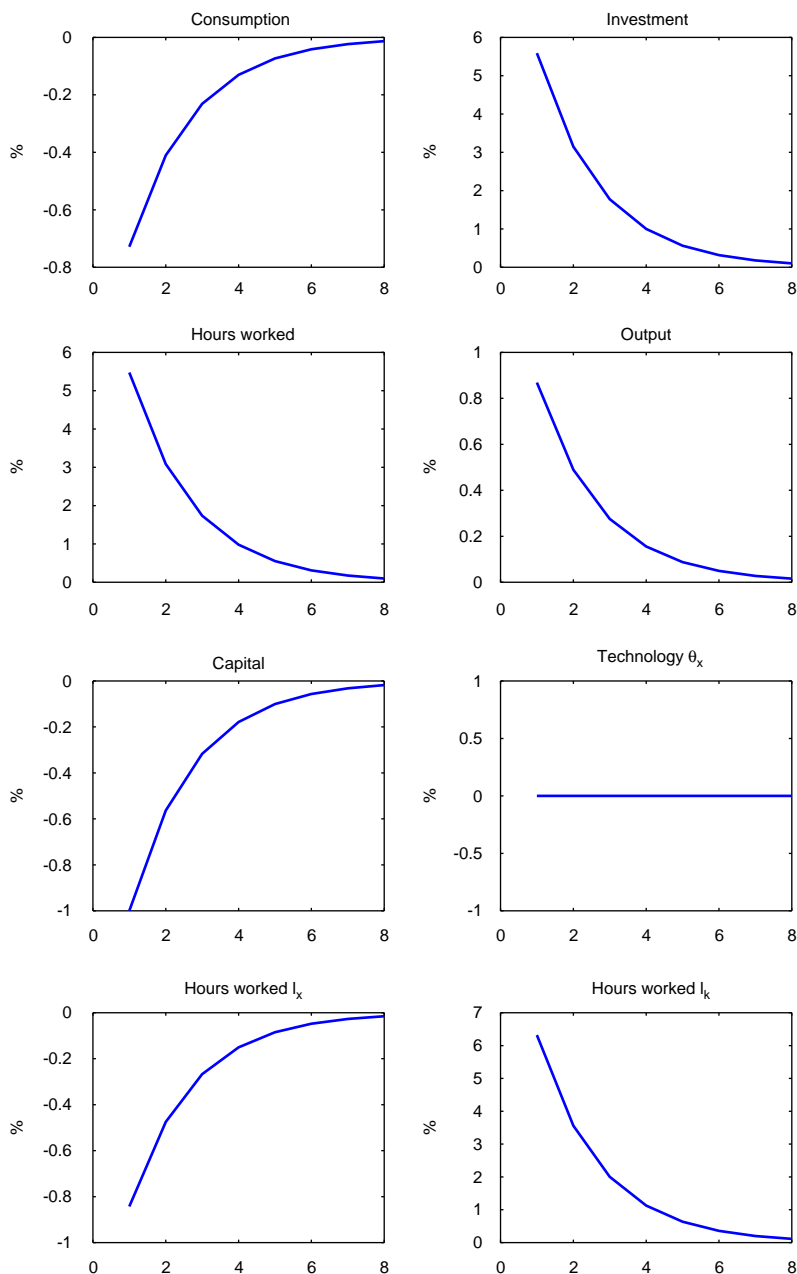
Fig. 5. Transitional dynamics with K_0 below steady state (all variables are measured in relative deviations from their respective steady-state level).

Expectationally driven booms and busts : Let us now turn to examining how such an economy responds to an anticipated increase in technology in the consumption good sector, as induced by a positive realization of the signal v_t. As a result of the signal, a technological improvement is expected to arise at period 4.

Fig. 6 reports the dynamics associated with the case where the expected increase is actually realized. As can be seen from the figure, employment in the construction sector immediately jumps, then continues to increase until period 4. At this point technology improves and employment in the construction sector jumps on to the new saddle path, and then gradually decreases to its steady-state level while the capital stock continuously increases. These dynamics are such that anticipated technological improvement can be said to cause an expectation lead boom, that is, from period 1 to 3, employment in both sectors, total output (defined as $C_t + p_t I_t$), investment and consumption are all increasing even though technology has not yet improved.

Let us now consider what happens if, at period 4, instead of technology improving as anticipated, individuals learn that their forecast is incorrect and that technology does not actually change (it remained at its initial level). In this case, there is a fall in output and employment in the construction sector at period 4, as individuals realize that they previously over-accumulated. Following this drop, employment gradually returns to its previous steady state as the capital stock returns to its initial level. Fig. 7 graphs the time paths of all the main variables through this entire sequence of anticipation and realization.

Note from Fig. 7 that the economy first experiences a boom and then a recession without ever having experienced an actual change in technology. In particular, at period 4 aggregate output, investment and employment all fall, while consumption falls with a lag. It is interesting to note the pattern of investment prices (*p*) that decentralizes this behavior (shown in Fig. 8). During the first phase, from period 1 to period 4, individuals invest in infrastructure in anticipation of realizing capital gains. Throughout this phase, the price of infrastructure increases thereby fulfilling these expectations of capital gains. At period 4, however, the price of infrastructure falls in recognition of an over-supply. The new low price for infrastructure makes investment unprofitable and therefore employment in the construction industry collapses. As the oversupply of capital slowly diminishes due to depreciation, incentives for new investment reemerge and thereby pulling the economy out of the recession.

Discussion and policy implications: In our view, these dynamics capture the idea, suggested by Pigou and others, that forecast errors may be key in understanding recessions. In effect, in this model, a boom and a recession can arise as the result of overly optimistic expectations about future technological growth. Two questions arise immediately. Which properties of this model allows it to generate Pigou cycles, and are these properties reasonable? As we will now make clear, the conditions under which Pigou cycles can arise are rather stringent, and this is the reason why our model has several particular features.

Two conditions must be met if a model is to generate Pigou cycles. These are (a) agents must receive signals about future productivity which lead them to increase their current demand for investment and (b) this increase in investment demand must
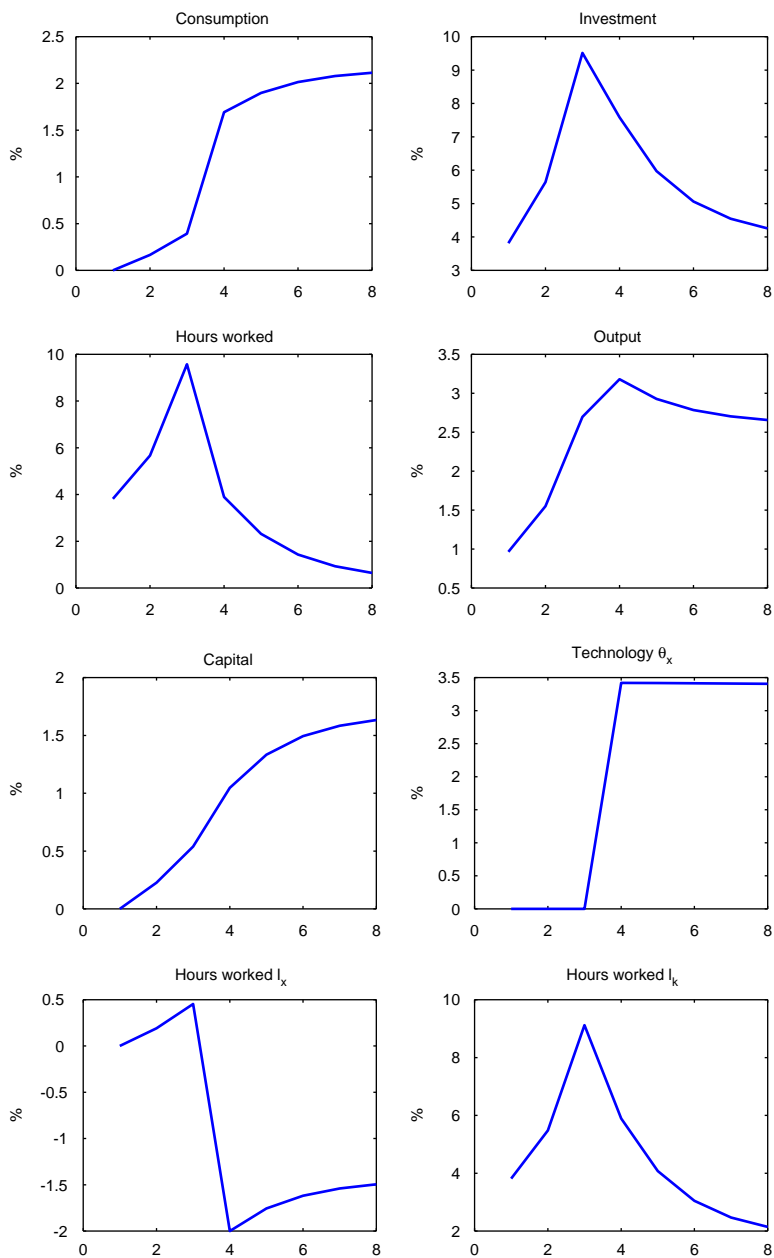
Fig. 6. Response of the three-sectors economy to an announcement at time 1 of future positive shock on the technology and a realization of that shock at time 4.
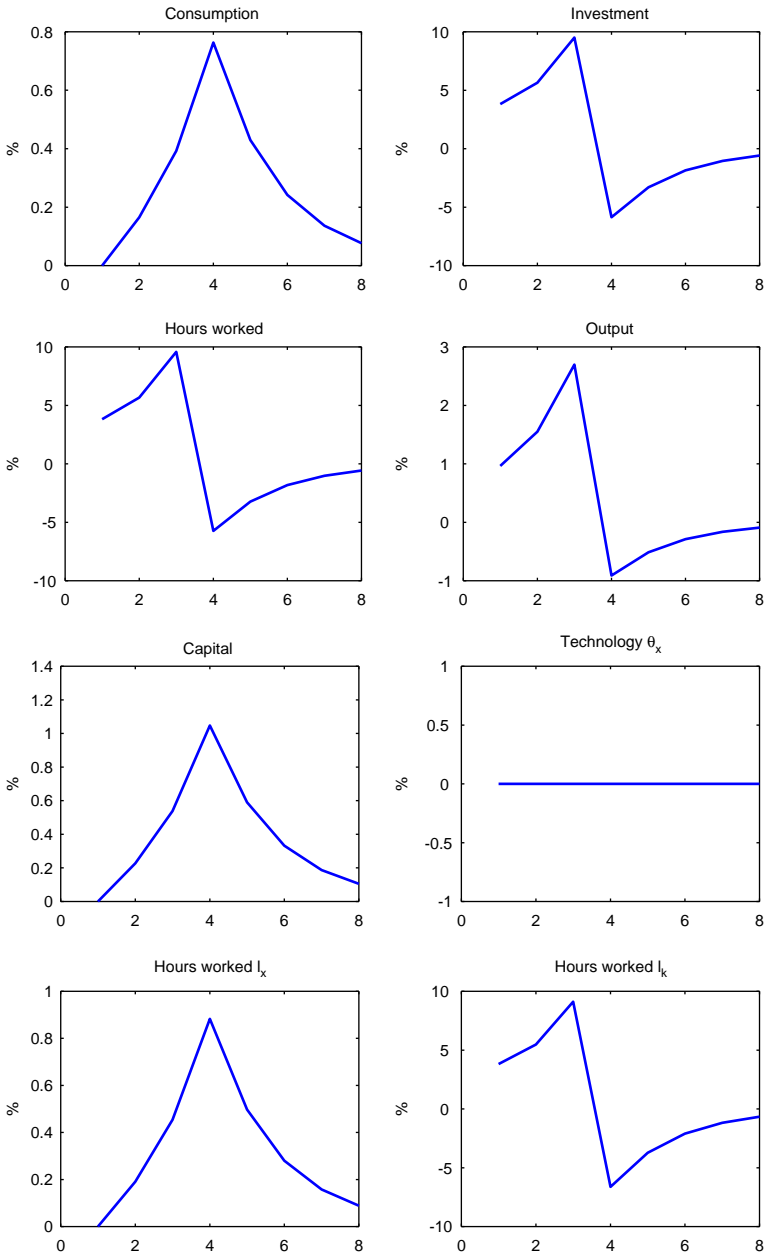
Fig. 7. Response of the three-sectors economy to an announcement at time 1 of future positive shock on the technology and no realization of that shock at time 4.
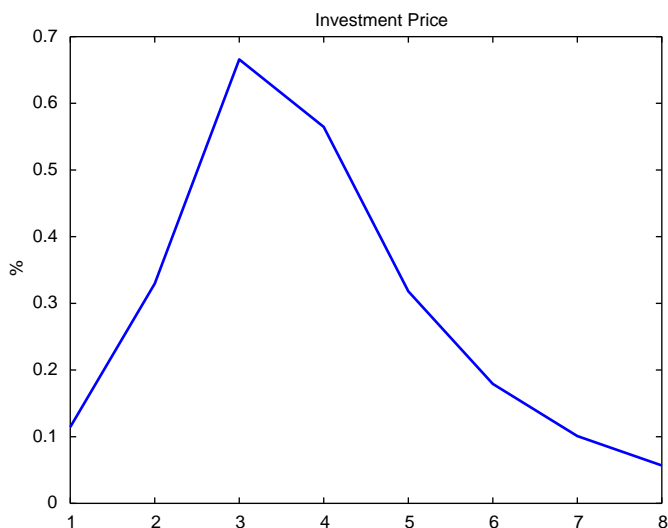
Fig. 8. The behavior of investment price in a Pigou cycle.

be met by an increase in employment, not a decrease in consumption. Point (a) is easy to understand and can arise in many standard models. For example, in a one-sector model, adjustment costs to capital are sufficient to generate increased investment demand following a positive signal about future productivity growth. However, in such a model, an increase in investment demand will be met by a decrease in consumption and hence will not produce a Pigou cycle. In two-sector models or in our three-sector model, an increase in investment demand will arise following a signal only if the signal received does not relate only to productivity improvement in the capital good sector. Otherwise, if the signal relates only to an improvement in the capital good sector, it is always optimal for agents to wait for the realization of the shock before increasing investment. In the case of news about improvement in the capital good sector, good news would lead to a recession not a boom. Fig. 9 shows how the economy responds to such a shock: macroeconomic aggregates move together, but downwards. Furthermore, in a multi-sector model like ours, expected improvement in the non-durable good is not in itself sufficient to guaranty increased demand for current investment. In effect, in order to guaranty that investment increases following news about future productivity in the non-durable good sector, the production structure must exhibit enough complementarity between capital and the non-durable good. If instead capital and the non durable good were highly substitutable, agents would decide to wait until the arrival of the shock before increasing investment, then reallocate labor from the non durable sector to the capital good sector, and substitute non durable good for capital to maintain their flow of consumption. To confirm that this mechanism is indeed at work in our model, Fig. 10 reports the response of the economy to an expected improvement in the consumption good sector when production in that sector is of
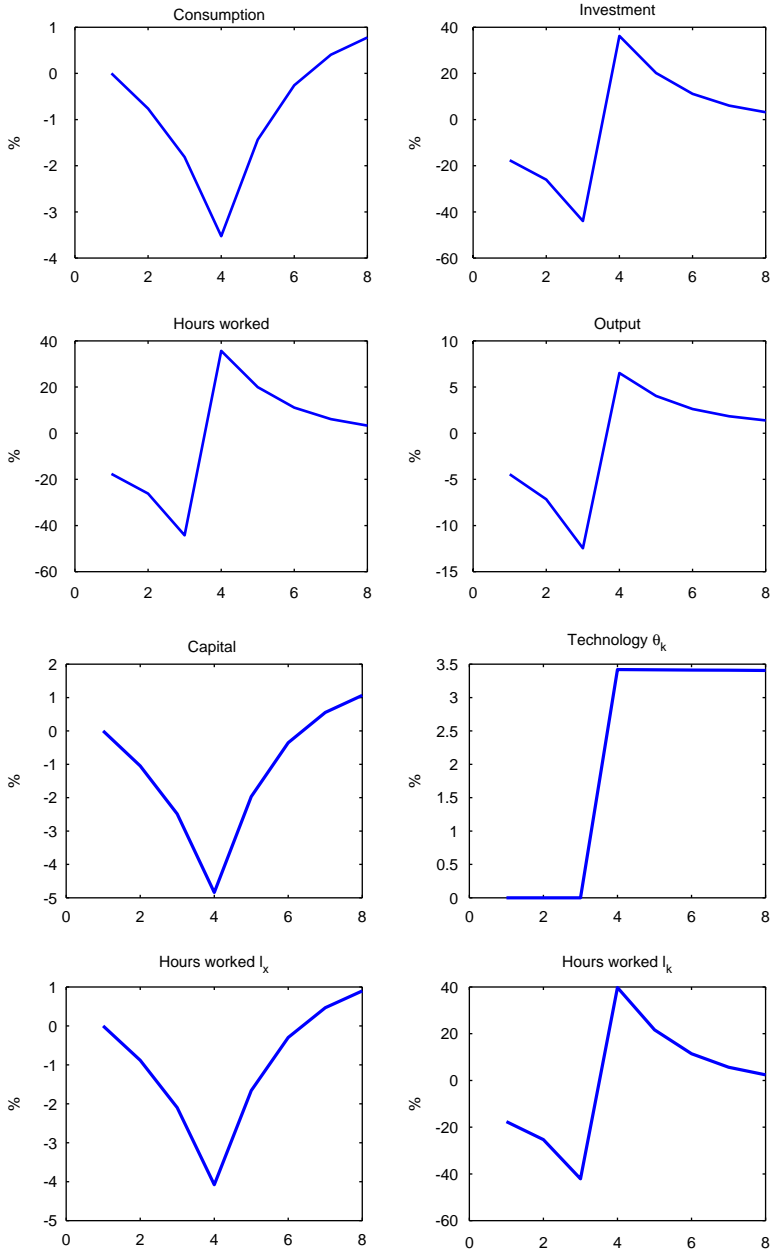
Fig. 9. Response of the three-sectors economy to an announcement at time 1 of future positive shock on the technology and a realization of that shock at time 4. Model with technological shock in the capital good sector (θ_k).
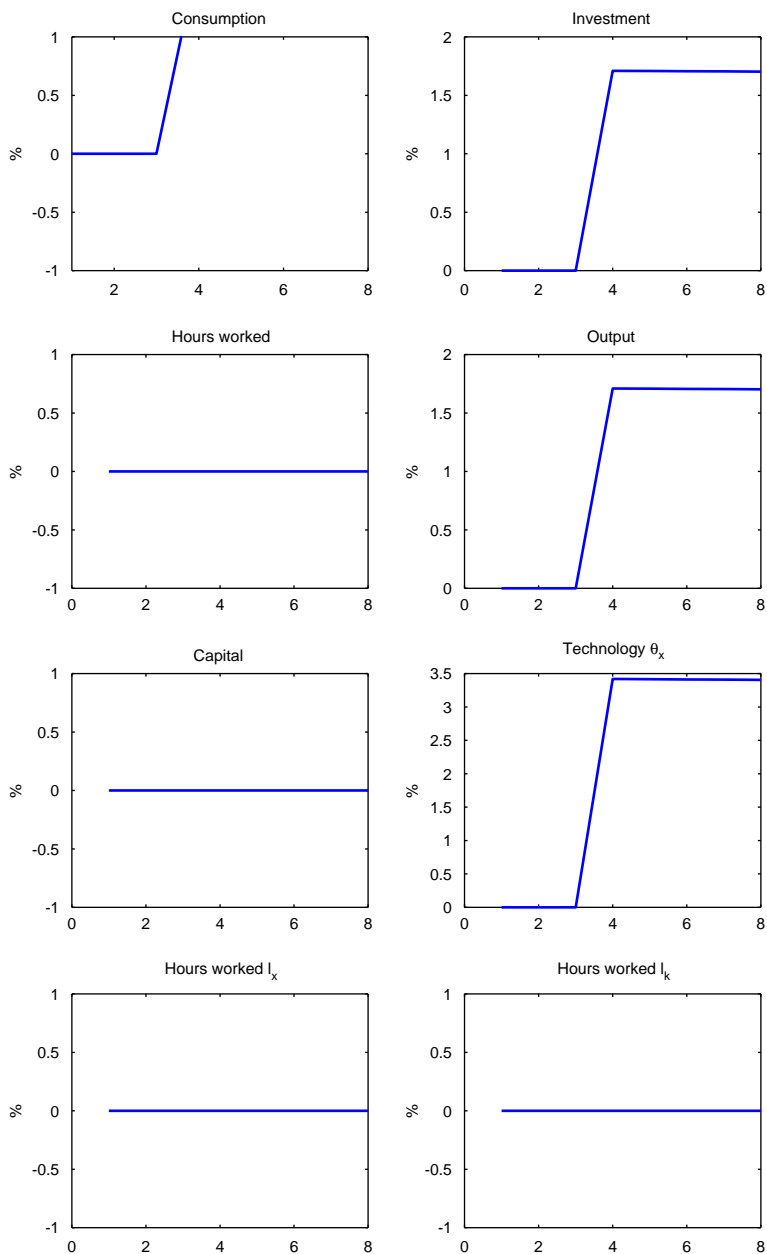
Fig. 10. Response of the three-sectors economy to an announcement at time 1 of future positive shock on the technology and a realization of that shock at time 4. Model with Cobb–Douglas production function ($C_t = X_t^a K_t^{1-a}$).

the Cobb–Douglas form $C_t = X_t^a K_t^{1-a}$, that is, the case with a unitary elasticity of substitution. As can be seen, the economy does not respond to news in this case: all responses are flat until the arrival of the technological improvement. For there to be a positive response of investment following news, the elasticity of substitution between the non-durable good and capital in the production of the consumption good must be less than one.

The above discussion clarifies that Pigou cycles can arise only if good news favors current capital accumulation. This happens in our model due to both our assumptions that the news is about improvements in the non-durable good sector and the fact that non-durable goods are complementary to capital in the production of consumption goods. As we have mentioned, an increase in investment demand following news can be achieved by many other means, and hence this is not a feature that is particular to our model. In effect, what is most novel about our model is its capacity to satisfy point (b) above, that is, to allow increased investment demand to be met by an increase in employment instead of a fall in consumption. This latter property arises in our model because current consumption decisions are essentially decoupled from current investment decisions. In other words, agents in the economy can determine how much investment in infrastructure to undertake without this decision having a direct feedback on how much the economy can currently consume. In particular, the shadow price of investment in our three-sector model is not directly related to foregone consumption but only relates to the cost of reduced leisure for workers in the construction industry. We believe that this feature, whereby increased investment is not directly reducing consumption possibilities, is a property that may be a sensible description of short-term substitutability constraints in a modern economy. For example, if an economy has an oversupply of buildings, it seems a reasonable simplification to exclude—at least in the short run—the possibility of immediately transforming the output of construction sector workers into others goods. It is this type of property which most differentiates the above three-sector model from more standard macromodels and thereby allows for Pigou cycles to emerge.

Let us finish by stressing that having high labor supply elasticity is a feature that helps the decoupling of investment and consumption decisions in the model and thereby favors the emergence of Pigou cycles. If instead agents valued smooth leisure sufficiently, they would be less willing to work harder in both sector following a news and this would limit the possibility of Pigou cycles. Nonetheless, it is worth noting that our model does not require an infinite elasticity of labor supply to have expectations generate substantial fluctuations. For example, when we consider a version of our model with a unitary intertemporal elasticity of labor supply (i.e., preferences of the form $\log(C_t) + v_0 \log(\bar{l} - l_{x,t} - l_{k,t})$), the effects of an anticipated change in productivity are still substantial. To see this, in Fig. 11 we report the economy's response to an anticipated change in productivity when labor supply has a unitary elasticity. The shock underlying Fig. 11 is of the same size as that used in Fig. 6 and hence the size of the responses can be directly compared. As can be seen from the figures, the response of investment and consumption in the anticipation phase is about a third less when labor supply has a unitary elasticity relative than
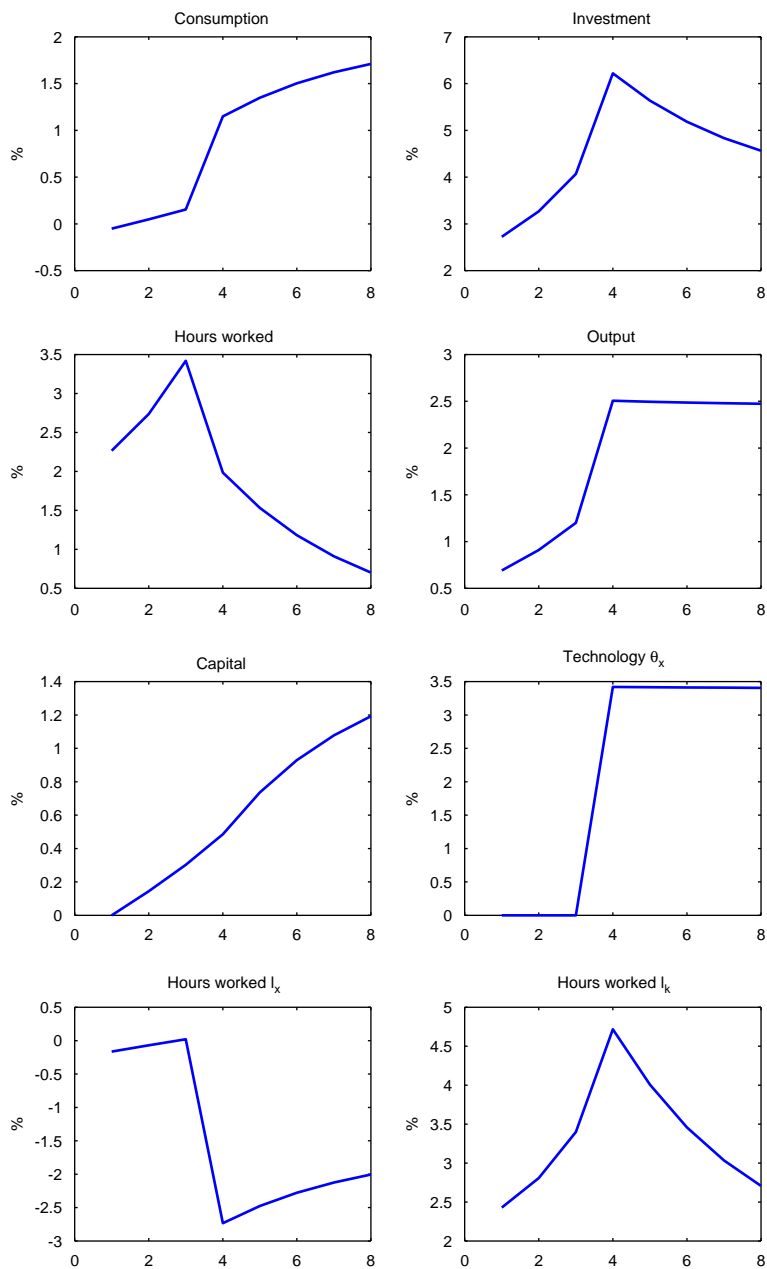
Fig. 11. Response of the three-sectors economy to an announcement at time 1 of future positive shock on the technology and a realization of that shock at time 4. Model with Log disutility of labor $(\log(C_t) + v_0 \log(\bar{l} - l_{x,t} - l_{k,t}))$.

when this elasticity is infinite. Although this is a sizeable reduction, it nicely illustrates how expectations can cause substantial fluctuations even in the presence of a unitary elasticity of labor supply.[14]

It is interesting to conjecture how individuals may perceive a downturn generated by a Pigou cycle and how this may lead them to choose inappropriate policies. For example, at the onset of a recession, individuals in our model are likely to perceive the cause of the recession as being a fall in aggregate demand. In response, they may be tempted to favor policies, that would stimulate investment demand; such as temporary tax breaks or investment tax credits. However, such policies would be misplaced in this model since it is precisely an excess of investment that caused the recession. Policies which stimulate investment may even appear to individuals as a cure to downturns—since they would temporarily increase employment and output—when in fact such policies would at best be a postponement of needed adjustment.

4.3. Does the model produce sizable recessions?

In this section we explore the extent to which our model economy is capable to generate recessions. To this end, summary business cycle statistics and a set of recession statistics for the U.S. economy are given in Tables 3 and 4. Our construction and reporting of recession statistics, in addition to standard business cycle statistics, reflect our desire to evaluate the capacity of our model to explain this particular phase of the cycle.

In these two tables, the figures in parenthesis are standard deviations of estimators. In Table 3, the variables of the form $\overline{\Delta^- x}$ represents the average growth rate of x conditional on Δx being negative, $F(\Delta^- x)$ represents the percentage of semesters for which Δx is negative, $\min(\Delta^- x)$ represents the largest recession (percentage decrease) of x, and finally $\overline{\Delta x}|_{\Delta y < 0}$ is the average growth rate of x conditional on Δy (growth in aggregate output) being negative. Note that recessions are not rare events: almost one fifth of semesters experienced output drops and on average the falls are 1.1%. For investment, recessions happen one third of the time, and the average fall in a semester is almost 4%. In contrast, recession for consumptions happen rarely and when they do happen they are on average very shallow at .1%.

Let us now turn to the statistics generated by the model. We evaluate the model's ability to match the data in the following way. We generate 1000 simulations of length 77 (the number of observations in our sample), and compute the mean and the standard deviations of the moments of interest. We then ask the question: "Is the data at odds with the statistics generated by the model?". We ask this question for

---

[14]At close inspection, one can see in Fig. 11 that a news shock leads to an increase in investment and l_k on impact, but a slight decrease in l_x and therefore a small decrease on impact of consumption. However, one can still reasonably say that the economy exhibit an expansion following the news shock in this case since consumption starts to increase and passes above its steady-state level after only one period following the news, although no technological improvement is yet observed.

Table 3
Statistics on recessions (U.S. NIPA, 59:1–97:2, semi-annual)

$\overline{\Delta^- y}$	$F(\Delta^- y)$	$\overline{\Delta^- c}$	$F(\Delta^- c)$	$\overline{\Delta^- i}$	$F(\Delta^- i)$
−1.1	19.5	−.15	7.8	−3.87	33.8
(.2)	(5.3)	(.03)	(3.7)	(.69)	(6.5)
min$(\Delta^- y)$	min$(\Delta^- i)$	min$(\Delta^- c)$	$\overline{\Delta c}|_{\Delta y < 0}$	$\overline{\Delta i}|_{\Delta y < 0}$	
−2.56	−13.56	−.29	.51	−5.8	
—	—	—	(.17)	(.79)	

Table 4
Statistics on business cycle (U.S. NIPA, 59:1–97:2, semi-annual)

HP filter data

σ_y		σ_c		σ_i
2.16		1.06		6.85
(.22)		(.09)		(.67)
ρ_y		ρ_c		ρ_i
.79		.84		.74
(.08)		(.06)		(.09)
cor(y, c)		cor(y, i)		
.64		.95		
(.08)		(.01)		

Table 5
Statistics on recessions (Model)

$\overline{\Delta^- y}$	$F(\Delta^- y)$	$\overline{\Delta^- c}$	$F(\Delta^- c)$	$\overline{\Delta^- i}$	$F(\Delta^- i)$
−.95★	13.99★	−.00	.03	−3.17★	40.27★
(.35)	(4.2)	(.02)	(.39)	(.99)	(6.3)
min$(\Delta^- y)$	min$(\Delta^- i)$	min$(\Delta^- c)$	$\overline{\Delta c}|_{\Delta y < 0}$	$\overline{\Delta i}|_{\Delta y < 0}$	
−3.22★	−16.25★	.27	.97	−6.61★	
(.84)	(3.42)	(.14)	(.10)	(1.78)	

several different moments, focusing on one moment at a time. Tables 5 and 6 report statistics generated by the model, with standard deviations given in parenthesis. A ★ on a statistic indicates that the empirical moment lies within a interval of ±2 standard deviations around the mean of the model simulations (★★ for ±3 standard deviations). We interpret this as follows: as far as this particular moment is concerned, we cannot reject that the data could have been generated by our model.

Table 6
Statistics on business cycle (Model)

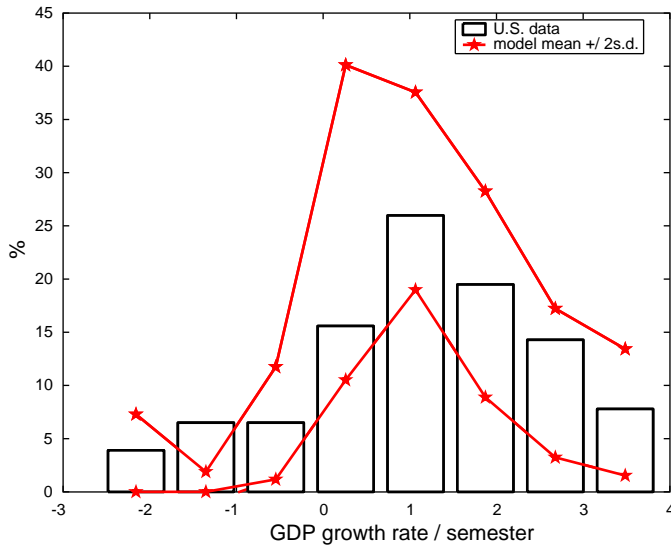| HP filter data		
σ_y	σ_c	σ_i
1.85★	1.07★	5.84★
(.27)	(.18)	(.80)
ρ_y	ρ_c	ρ_i
.67★	.81★	.58★
(.09)	(.05)	(.10)
cor(y, c)	cor(y, i)	
.63★	.91★★	
(.06)	(.01)	

Fig. 12. Output growth histogram.

We first comment on the ability of the model to reproduce standard business cycle statistics as reported in Table 6. Recall that the model has been calibrated to give a good fit for the standard deviations of *HP*-filtered output, consumption and investment. However, the model was not calibrated to match the other statistics in Table 6. In particular, it is interesting to note that the serial correlations and cross-correlation of the *HP* filtered data are well reproduced by the model. We interpret these results as suggesting that the model does a good job at matching the moments most often discussed in the *RBC* literature. Let us now look at recession statistics

(Table 5). Again recall that the calibration has been done without targeting these statistics. As far as output and investment are concerned, the model does a very good job at reproducing the recession statistics. Average and maximum depth of recessions, as well as frequency of recessions, are all matched by the model even though there is never technological regress. The only major failure of the model is that consumption is too smooth: in effect, the model does not produce significant recessions in per capita consumption, it only produces significant slowdowns. However, we do not interpret this failure of the model to be a fatal drawback given that drops in consumption are also rare and small in the data (see Table 3).

In order to get an additional view of the model's ability to reproduce observed output growth, Fig. 12 plots the histogram of output growth. As can be seen from the figure, the empirical histogram lies almost entirely within the 2 standard deviations bands associated with the histogram generated by the model. This figure nicely illustrates how a simple three-sector equilibrium model without any technological regress can reproduce important patterns of output growth.

5. Asset prices implications

Our model has sharp implications for the price of capital. In effect, if one believes that fluctuations are largely driven by expectations of technological improvements then, as illustrated previously in Fig. 8, the relative price of capital should co-vary positively with changes in output that are induced by such expectations. In this section we want to explore the relevance of our model along this dimension. To do so, we proceed in two steps. First, we examine some un-conditional correlations between output and the price of capital. Second we report some conditional correlations and impulse responses inferred from a VAR using long-run restrictions. The advantage of the second approach is that it allows us to examine more closely properties of the data that can be tied directly to changes in technological opportunities. However, before going to the data, we need to identify what is meant by the price of capital in our model. In accordance with the spirit of the model, we will focus on three price indices for capital. The first two measures derive naturally from the idea that the capital stock in our model is the infrastructure built by the construction industry. For this reason we will use both the price of residential investment and the price of non-residential buildings as measures of the price of capital. However, in order to allow for a larger definition of capital, we will also examine the behavior of a stock market index as given by the Standard & Poors 500.[15]

---

[15] Our sample runs quarterly from 1948:1 to 2000:4. The variables that we use here are obtained from the U.S. BEA. The implicit price deflator (hereafter IPD) of non-residential buildings and the IDP of residential investment are divided by the IPD of nondurable goods to obtain relative prices. Output is measured as the Non Farm Private Business Sector Gross Domestic Product. When computing correlations, series are filtered using Hodrick–Prescott filter with $\lambda = 1600$.

5.1. Unconditional correlations

When we consider the Standard & Poors 500 as our measure of the price of capital, we obtain a positive correlation of .33 between H-P filtered output and price data for the U.S. postwar period. When we consider the price of residential investment as our measure of the price of capital, we obtain a correlation of .33. Finally, when we consider the price of non-residential buildings as the relevant price, we obtain a correlation of .29.[16] Although these correlations are not exceedingly high, they do support the predictions of the model. One reason for why these correlations may not be as high as would be predicted by the model is that fluctuations in the data are likely induced by many sources of shocks, not just the ones considered in the model. In order to explore this possibility, we now turn to examining some conditional moments.

5.2. Conditioning

If we interpret our model broadly, it actually should be viewed as delivering predictions for conditional correlation, as it predicts that the relative price of capital increases when the economy is hit by expected but permanent shocks to technology. To get an estimate of this conditional correlation from the data, we estimate bivariate VARs $(\Delta y, \Delta p)$ where y is the log of output and p the log of a relative price of capital.[17] The series used are the same as those used to compute unconditional correlations. From the estimated VAR, we compute the responses to the permanent shock to output.[18] The VAR analysis is repeated for the price of residential investment, the price of non-residential buildings and the stock price index. Impulse response functions to a permanent shock are displayed in Fig. 13, together with 90% confidence bands. The impact responses of residential investment, non-residential buildings prices and stock market index are all significantly positive. When we compute the correlation between output and price along the response paths we obtain a conditional correlation is .89 for residential investment price, .79 for non-residential buildings and .75 for the stock price. These correlations bring substantial support for the type of mechanism we have illustrated.

---

[16]Non-residential buildings constitutes about two thirds of the broad category referred to structures. However, we should note that the price of non-residential buildings behaves substantially different from the other components in structures which are mining exploration, shafts, wells and utilities. In particular, these last three components have a slight negative covariance with output.

[17]All VARs are estimated with 5 lags, on the largest available sample. The bi-variate system for output and the S&P 500 index is estimated in VECM form instead of as a VAR in differences since cointegration between these two variables is highly plausible.

[18]Formally, in the spirit of Blanchard and Quah (1989), we identify two orthogonal shocks to the bivariate system. One is restricted to have zero long-run impact on the level of output. The other one is our permanent shock. It is indeed a permanent shock, as the series are all non stationary.
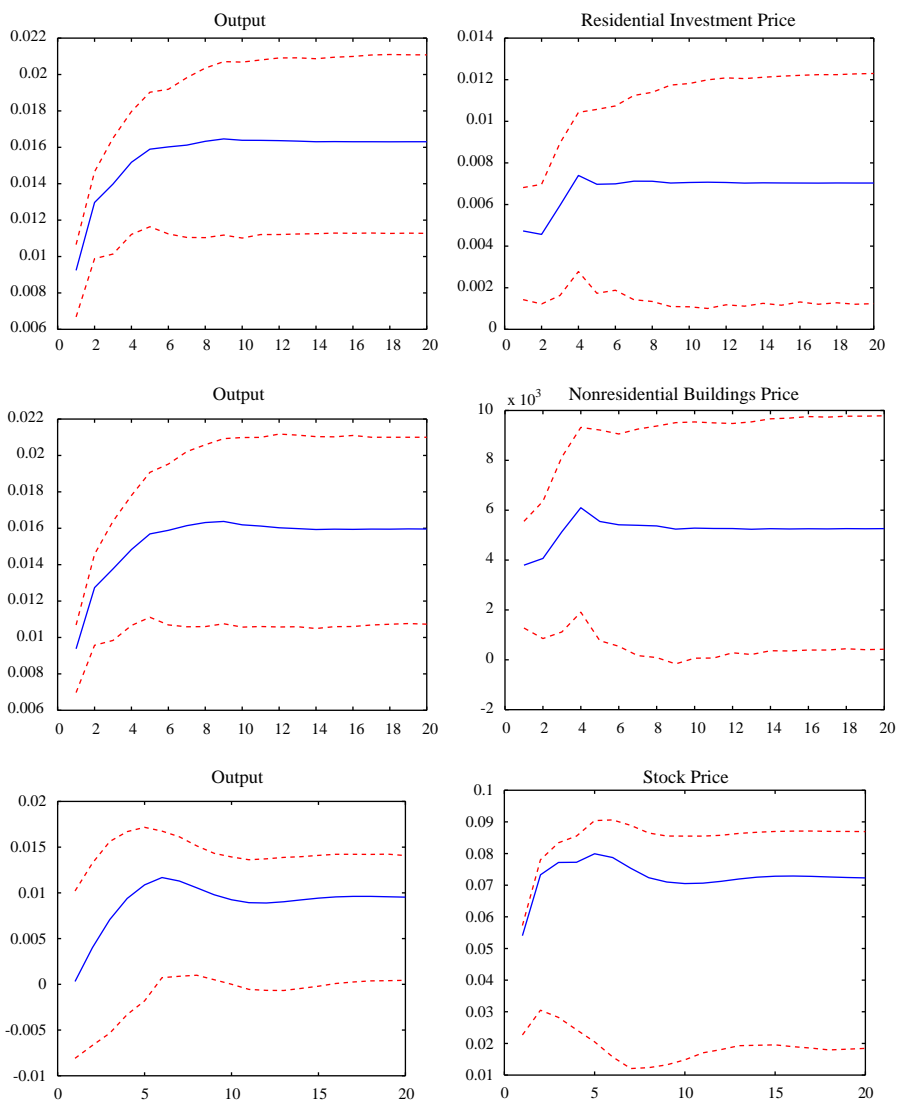
Fig. 13. Asset prices responses to a permanent shock to output.

6. Conclusion

In this paper we have illustrated an equilibrium business cycle model where anticipations and realizations of technological growth were qualitatively and quantitatively able to explain several patterns associated with business cycles and recessions. We think that the mechanism of this model—the importance of forecasts

and forecast errors in explaining aggregate movements of activity—may help understand certain episodes of cyclical downturns in industrialized economies. In particular, this type of model may provide a useful framework for understanding the recent downturns in South-East Asia since it has been argued that revisions of expected growth were central in generating the crises observed in these economies.

We also believe that the U.S. recession of 2001 may reflects a downturn of a Pigou cycle. Although the period of the late nineties was one where there were significant *realized* technological improvements in the capital goods sector, and there was an anticipation that such a trend would continue, this is not where we believe anticipation where most relevant nor overly optimistic. In effect, if the only important source of anticipation over this period was expected future decrease in the relative price of capital goods, this should not have had lead to an investment boom, it should have lead to postponement. Instead, what we believe to be the most important source of anticipation over this period, which eventually was revised downward in 2001, regarded the anticipated speed of discovery and development of new uses for IT goods. According, we believe that it was the expectation of such development that drove the investment process, whether it be in telecommunications, computers or structures; and it was the revision of this expectation that lead to a recession. Hence in our view, as in the model, it is expected developments of sectors that use capital good as inputs which is key to understanding the period.

However, we leave a detailed exploration of these particular episodes for future research.

Appendix A. Proof of the impossibility of Pigou cycles in standard one-sector models

In this appendix we prove that, in a one-sector economy with or without adjustment cost to capital, Pigou cycles cannot occur, i.e. we cannot have positive co-movements of consumption, investment and hours induced by an expectational change without any change to current fundamentals.

Let us first describe the environment. The economy is competitive, populated with one representative firm and one representative household. Preferences are ordered by

$$E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, Lt), \tag{A.1}$$

where U is a twice continuously differentiable utility function and $\beta \in]0, 1[$. In particular, we assume that $U_C > 0$ and $U_{CC}, U_{LL}, U_{CL} < 0$, which implies that leisure is a normal good. In order to provide a general characterization, it is helpful to specify the production set as follows:

$$C_t = G(K_t, X_t L_t, I_t), \tag{A.2}$$

where L_t is labor input, K_t is the capital stock, X_t is an index of technology, I_t is investment and where G is a twice continuously differentiable constant return to scale function with $G_I < 0$. For our two cases of interest, the function G take the following forms.

In the case corresponding to the standard one-sector model:

$$G(K_t, X_t L_t, I_t) = F(K_t, X_t L_t) - I_t$$

with $F_1, F_2, F_{12} > 0$ and $F_{22}, F_{11} < 0$. The second case corresponds to a model with adjustment costs:

$$G(K_t, X_t L_t, I_t) = F(K_t, X_t L_t) - \Psi\left(\frac{I_t}{K_t}\right) K_t$$

with again $F_1, F_2, F_{12} > 0$ and $F_{22}, F_{11} < 0$, and where Ψ is a convex function. The important aspect to remember about these two cases is that G_{LI} is equal to zero.

Capital accumulates over time according to the following law of motion:

$$K_{t+1} \leqslant (1 - \delta) K_t + I_t \tag{A.3}$$

with $0 < \delta < 1$. As we consider convex economies with perfect competition, the two theorems of welfare hold, and we can obtain equilibrium allocations from the following social planner's program:

$$\max_{\{C_{t+j}, K_{t+j+1}, L_{t+j}\}} \quad \mathrm{E}_t \sum_{j=0}^{\infty} \beta^t U(C_{t+j}, L_{t+j})$$
$$\text{s.t.} \quad C_{t+j} = G(K_{t+j}, X_{t+j} L_{t+j}, I_{t+j}),$$
$$K_{t+j+1} = (1 - \delta) K_{t+j} + I_{t+j},$$
$$K_{t+j} quad \text{given.}$$

At any given period t, the first-order conditions of the problem are given by

$$U_C(t) X_t G_L(t) = -U_L(t), \tag{A.4}$$

$$C_t = G(K_t, X_t L_t, I_t), \tag{A.5}$$

$$-G_I(t) U_C(t) = \beta E_t [U_C(t+1)(G_K(t+1) - (1 - \delta) G_I(t+1))]. \tag{A.6}$$

We now want to explore the conditions under which a change in expectations—without a current change in X_t—necessarily leads to a negative co-movement between C_t and I_t. To derive such a condition, all that is needed is to examine the determinants of $\mathrm{d}C/\mathrm{d}I$ implicitly defined by Eqs. (A.4) and (A.5), since these two equations define the set of temporary equilibria.

This derivative is given by

$$\frac{\mathrm{d}C}{\mathrm{d}I} = \frac{G_I + G_L \kappa_3}{1 - G_L \kappa_1 \kappa_2}, \tag{A.7}$$

where

$$\kappa_1 = -(G_L U_{CL} + U_C G_{LL} + U_{LL})^{-1} > 0,$$
$$\kappa_2 = G_L U_{CC} + U_{LC} < 0,$$
$$\kappa_3 = U_C G_{LI} \lesseqgtr 0.$$

From this derivative, we can immediately note that if $G_{LI} \leqslant 0$ then $\mathrm{d}C/\mathrm{d}I < 0$. Hence, it is necessarily the case that, within the standard one-sector model with or

without adjustment costs to capital, an expected change in future technological opportunities leads to a negative co-movement between C_t and I_t since both these models have the property that $G_{LI} \leqslant 0$. Note that it is easy to verify that for the three-sector model we present in the paper, the implied G function has $G_{LI} > 0$.

References

Allen, S., 1985. Why construction industry productivity is declining? Review of Economics and Statistics 67, 661–669.

Banerjee, A., 1992. A sample model of herd behavior. Quarterly Journal of Economics 107, 797–817.

Beaudry, P., Portier, F., 1999. An exploration into Pigou's theory of cycles. Mimeo, CEPR Discussion Paper 2996, University of British Columbia and Université de Toulouse, October 2001.

Bickhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy 100, 992–1026.

Black, F., 1995. Exploring General Equilibrium. MIT Press, Cambridge, MA.

Blanchard, O., Quah, D., 1989. The dynamic effects of aggregate demand and supply disturbances. American Economic Review 79 (4), 655–673.

Burnside, C., Eichenbaum, M., Rebelo, S., 1995. Capital utilization and return to scale. NBER Macroeconomics Annual 1995, pp. 67–109.

Caplin, A., Leahy, J., 1993. Sectoral shocks, learning and aggregate fluctuations. Review of Economic Studies 60, 1065–1085.

Chamley, C., Gale, D., 1994. Information revelation and strategic delay in a model of investment. Econometrica 62 (5), 1065–1085.

De Long, B., 1991. Liquidation cycles and the great depression. Mimeo, Harvard University.

Duffie, D., Singleton, K., 1993. Simulated moments estimators of Markov models of asset prices. Econometrica 61 (4), 929–952.

Greenwood, J., Yorukoglu, M., 1997. 1974, Carnegie–Rochester Conference Series on Public Policy 46 (1), 49–95.

Hairault, J., Langot, F., Portier, F., 1997. Time to implement and aggregate fluctuations. Journal of Economics Dynamics and Control 22, 109–121.

Hansen, G., 1985. Indivisible labor and the business cycles. Journal of Monetary Economics 16 (3), 309–327.

King, R., Rebelo, S., 1999. Resuscitating real business cycles. In: Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 1C. North-Holland, Amsterdam.

King, R., Plosser, C., Rebelo, S., 1988. Production, growth and business cycles I. Journal of Monetary Economics 21 (2/3), 196–232.

Kydland, F., Prescott, E., 1982. Time to build and aggregate fluctuations. Econometrica 50 (6), 1345–1370.

Lucas, R., 1972. Expectations and the neutrality of money. Journal of Economic Theory 4 (2), 103–124.

Phelps, E., 1999. Behind this structural boom: the role of asset valuations. American Economic Review 89 (8), 63–68.

Pigou, A., 1926. Industrial Fluctuations. MacMillan, London.

Rogerson, R., 1988. Indivisible labor, lotteries and equilibrium. Journal of Monetary Economics 21 (1), 3–16.

Zeira, J., 1994. Informational cycles. Review of Economic Studies 61 (1), 31–44.