Gold rush fever in business cycles

Paul Beaudry [a], Fabrice Collard [b], Franck Portier [c,*]

[a] University of British Columbia and NBER, Canada
[b] The University of Adelaide, Australia
[c] Toulouse School of Economics Aile Jean-Jacques Laffont, Manufacture des Tabacs, Université Toulouse 1 Capitole, 21 allée de Brienne, 31000 Toulouse, France

A R T I C L E   I N F O

A B S T R A C T

A flexible price model of the business cycle is proposed, in which fluctuations are driven primarily by inefficient movements in investment around a stochastic trend. A boom in the model arises when investors rush to exploit new market opportunities even though the resulting investments simply crowd out the value of previous investments. A metaphor for such profit driven fluctuations are gold rushes, as they are periods of economic boom associated with expenditures aimed at securing claims near new found veins of gold. An attractive feature of the model is its capacity to provide a simple structural interpretation to the properties of a standard consumption and output Vector Autoregression.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

There is a large literature aimed at decomposing business cycles into temporary and permanent components. A common finding in this literature is that there is a significant temporary component in business cycle fluctuations; that is, an important fraction of business cycles appears to be driven by impulses that have no long run impact. While technology shocks have arisen as a leading candidate explanation to the permanent component (whether these shocks be surprise increases in technological capacities, or news about future possibilities), there remains substantial debate regarding the driving forces behind the temporary component of macroeconomic fluctuations. Several potential explanations to the temporary component have been advanced and explored in the literature; the most notable being monetary shocks and government spending shocks. While such disturbances can create temporary business cycle movements, quantitative evaluation of their effects have generally found that they account for a very small fraction of macroeconomic fluctuations.[1] Hence, the puzzle regarding the driving force behind temporary fluctuations persists. Since the most obvious – and most easily measured – candidates have not been convincingly shown to adequately explain temporary fluctuations, part of the literature has turned to explore the potential role of shocks that are conceptually more difficult to measure. A prominent example of this alternative line of research is the literature related to sunspot shocks.

---

* Corresponding author. Tel.: +33 5 61 12 88 40; fax: +33 5 61 12 86 37.
  E-mail address: fportier@cict.fr (F. Portier).

[1] Following a Bayesian likelihood approach to estimate a dynamic stochastic general equilibrium model for the US economy using seven macro-economic time series (Smets and Wouters, 2007) found that no more than 10% of the variance of output can be explained by monetary shocks between one quarter and one year, while an "exogenous spending" shock can explain 35% of the variance of output at a one quarter horizon, but only 15% after one year. In an estimated new neoclassical synthesis model of the U.S. economy, Justiniano et al. (2010) find that government spending shocks explain no more than 2% of output variance at business cycle frequencies (from 6 to 32 quarters). Using a Vector Autoregression, Uhlig (2005) finds that monetary policy shocks account for 10% of the variations in real GDP at all horizons.

While several papers have argued that sunspot shocks offer a good explanation to temporary business cycle fluctuations (see Benhabib and Farmer, 1999 for a survey), much of the profession has remained skeptical.[2] The present research proposes and evaluates a theory of temporary business cycle fluctuations which has some similarities with sunspot shocks, in that expectation changes are the initial driving force. However, our approach is fundamentally different since it does not rely on indeterminacy of equilibrium nor on increasing returns to scale. Instead our model builds on the intuition derived from gold rushes, where expectations play an important role but are nevertheless based on fundamentals. Furthermore, like in a gold rush, the individual level gains from investment are clear while the social gains may be small or nil.

To help motivate our approach, let us briefly discuss the properties of a gold rush. For example, consider the case of Sutter's Mill near Coloma, California. On January 24, 1848, James W. Marshall, a carpenter from New Jersey, found a gold nugget in a sawmill ditch. This was the starting point of one of the most famous Gold Rushes in history, the California Gold Rush of 1848–1858. More than 90,000 people made their way to California in the two years following Marshall's discovery, and more than 300,000 by 1854—or one of about every 90 people then living in the United States. The population of San Francisco exploded from a mere 1000 in 1848 to 20,000 full-time residents by 1850. More than a century later, the San Francisco 49ers NFL team is still named for the prospectors of the California Gold Rush. Another famous episode, which inspired Charlie Chaplin's movie "The Gold Rush" and Jack London's book the "Call of the Wild", is the Klondike Gold Rush of 1896–1904. Gold prospecting took place along the Klondike River near Dawson City in the Yukon Territory, Canada. An estimated 100,000 people participated in the gold rush and about 30,000 made it to Dawson City in 1898. By 1910, when the first census was taken, the population had declined to 9000. As these examples make clear, gold rushes are periods of economic boom, generally associated with large increases in expenditures aimed at securing claims near new found veins of gold. We are aware that gold rush episodes do not occur at business cycle frequency, but they will serve here as a useful metaphorical example.

This paper explores whether business cycle fluctuations may sometimes be driven by a phenomenon akin to a gold rush. In particular, an analytic dynamic general equilibrium model is constructed, in which the opening of new market opportunities causes an economic expansion by favoring competition for market share. Those episodes are called market rushes. To capture the idea of a market rush, the model is an expanding varieties one, in which agents compete to secure monopoly positions in new markets, as often done in the growth literature (see for example Romer, 1987, 1990) and in some business cycle models (see for example Devereux et al., 1993), although the growth in the potential set of varieties is technologically driven and exogenous. In this setting, when agents perceive an increase in the set of technologically feasible products, they invest to set up a prototype firm (or product) with the hope of securing a monopoly position in the new market. It is therefore the perception of these new market opportunities that causes the onset of a market rush and the associated economic expansion. After the initial rush, there is a shake out period where one of the prototypes secures the dominant position in the market.[3] The long term effect of such a market rush depends on whether the expansion in variety has an external effect on productivity. In the case where it does not have an external effect, the induced cycle is socially wasteful as it only contributes to the redistribution of market rents. In contrast, when the expansion of variety does exert positive external effects, the induced cycle can have social value but will generally induce output fluctuations that are excessively large.[4] In the case where the market expansion has no external effect, the model is capable of explaining the salient qualitative features obtained from a permanent–temporary decomposition of a consumption–output vector autoregression (VAR).

Section 2 presents a set of properties of the data that models of fluctuations should aim to explain. Several of these features are well known and extensively discussed in Cochrane (1994). In a bivariate output–consumption vector error correction model (VECM) of the U.S. postwar economy, consumption is, at all horizons, almost solely accounted for by a permanent shock recovered using a long run restriction. In contrast, the associated temporary shock of the system is found to explain an important part of the short run volatility of output—i.e. the business cycle. This temporary shock also explains much of the fluctuations in hours worked and investment. These robust features of the data are quite challenging for business cycle models since even temporary shocks generally imply some reaction of consumption. Furthermore, the literature remains divided as to a structural interpretation for the temporary shock. As we think that a market rush is a potential candidate, Section 3 builds a model[5] which can be solved analytically and whose properties can therefore be clearly stated. In this model, the current economic activity depends positively on the expectation of next period's activity and on the perceived opening of new markets. Hence, when agents believe that the economy is starting a prolonged period of market expansion, this induces an immediate increase in investment and an associated economic expansion. In contrast,

---

[2] There are at least two reasons for why the profession has remained skeptical about the importance of sunspot shocks in business cycles. First, the empirical evidence has not provided great support for the theoretical features of the economy needed to allow for sunspot shocks. Second, the coordination of beliefs implicit in the underlying mechanism is hard to understand.

[3] The assumption that all markets are monopolistically competitive is made for analytical convenience. A richer model would make the degree of competition on the market a function of the number of startups. Such a model is presented in the online technical appendix to this paper.

[4] A potential example of such a process is the "dot com" frenzy of the late 1990s,where large investments were made by firms trying to secure a position in the expanding internet market. At the end of this process, there was a large shake out as many firms went bankrupt and only a small percentage survived and obtained a substantial market position. The long run productivity gains and social value associated with this process are still debated.

[5] The model presented belongs to the class of models in which nominal rigidities play no role. Our interpretation of such models is that they can correspond to models with sticky prices in which monetary authorities follow rules that implement the flexible price outcomes.

when there are no newly perceived market opportunities, the economy experiences a slump. Section 4 highlights the properties of this simple model in relation to the empirical properties of a consumption–output VECM. In particular, our market rush model is shown to display several of the qualitative properties of consumption–output VECM: consumption does not respond at all to the temporary (but persistent) shock, while this shock contributes to the short run dynamics of output, investment and worked hours. These patterns are often interpreted as providing evidence in favor of the permanent income hypothesis. However, it must be emphasized that these properties are aggregate properties and not partial equilibrium ones, which implies that a coherent explanation to these patterns requires a general equilibrium model that gives rise to permanent–temporary decomposition with no temporary component in consumption. As shown in Section 4, such patterns are not consistent with the standard analytical RBC model.

2. A target set of observations

The set of observations presented here provides a rich, though concise, description of fluctuations in output, consumption, investment and hours worked. Some of these observations are well-known, and some are not. The set of observations presented is meant to capture important features of fluctuations that business cycle theory should aim at explaining. These observations will be used to evaluate the potential role of market rushes in explaining macroeconomic fluctuations.

2.1. An output–consumption VECM and two identifications

Let us begin by reviewing properties of the bi-variate process for consumption and output in a VAR with one cointegrating relation. The main properties of this system were originally discussed in Cochrane (1994). As in this paper, two schemes are used to orthogonalize the innovations of the process: a long run orthogonalization scheme *à la* Blanchard and Quah (1989), and a short run or impact scheme *à la* Sims (1980). At this point, these two schemes should be viewed as devices for presenting properties of the data. There is no claim that these schemes identify structural shocks, nor that these data should be explained by a model with only two shocks.

Our empirical analysis is based on quarterly data for the U.S. economy. The sample spans the period 1947Q1 to 2004Q4. Consumption, C, is defined as real personal consumption expenditures on nondurable goods and services and output, Y, is real gross domestic product. Both series are first deflated by the 15–64 U.S. population and expressed in logarithms.[6] Standard Dickey–Fuller likelihood ratio and cointegration tests indicate that C and Y are $I(1)$ processes and do cointegrate. The joint behavior of those variables is therefore modeled with a VECM, where the cointegrating relation coefficients are $[1; -1]$ (meaning that the (log) consumption to output ratio is stationary).[7] Likelihood ratio tests suggest that the VECM should include three lags. Omitting constants, the joint behavior of (C,Y) admits the following Wold representation:

$$\begin{pmatrix} \Delta C_t \\ \Delta Y_t \end{pmatrix} = A(L) \begin{pmatrix} \mu_{1,t} \\ \mu_{2,t} \end{pmatrix}, \tag{1}$$

where L is the lag operator, $A(L) = I + \sum_{i=1}^{\infty} A_i L^i$, and where the covariance matrix of μ is given by Ω. As the system possesses one common stochastic trend, $A(1)$ is not full rank. Given $A(1)$, it is possible to derive a representation of the data in terms of permanent and transitory components of the form

$$\begin{pmatrix} \Delta C_t \\ \Delta Y_t \end{pmatrix} = \Gamma(L) \begin{pmatrix} \varepsilon_t^{\mathrm{P}} \\ \varepsilon_t^{\mathrm{T}} \end{pmatrix}, \tag{2}$$

where the covariance matrix of $(\varepsilon^{\mathrm{P}}, \varepsilon^{\mathrm{T}})$ is the identity matrix and $\Gamma(L) = \sum_{i=0}^{\infty} \Gamma_i L^i$. The Γ matrices solve

$$\begin{cases} \Gamma_0 \Gamma'_0 = \Omega, \\ \Gamma_i = A_i \Gamma_0 \quad \text{for } i > 0. \end{cases} \tag{3}$$

Note that once Γ_0 is known, all Γ_i are pinned down by the second set of relations. But, due to the symmetry of the covariance matrix Ω, the first part of the system only pins down three parameters of Γ_0. One remains to be set. This is achieved by imposing an additional restriction. The [1,2] element of the long run matrix $\Gamma(1) = \sum_{i=0}^{\infty} \Gamma_i$ is set to zero, meaning that the orthogonalization chosen is such that the disturbance ε^{T} has no long run impact on C and Y (the use of

this type of orthogonalization was first proposed by Blanchard and Quah, 1989). Hence, ε^T is labeled as a temporary shock, while ε^P is a permanent one. This orthogonalization is called the "long run" one.

Let us now consider an alternative orthogonalization that uses short run restrictions:

$$\begin{pmatrix} \Delta C_t \\ \Delta Y_t \end{pmatrix} = \tilde{\Gamma}(L) \begin{pmatrix} \varepsilon_t^C \\ \varepsilon_t^Y \end{pmatrix}, \tag{4}$$

where $\tilde{\Gamma}(L) = \sum_{i=0}^{\infty} \tilde{\Gamma}_i L^i$ and the covariance matrix of $(\varepsilon^C, \varepsilon^Y)$ is the identity matrix. The $\tilde{\Gamma}$ matrices are solution to a system of equations similar to (3). The system however departs from (3) and imposes that the 1,2 element of $\tilde{\Gamma}_0$ is equal to zero. Therefore, ε^Y can be called an output innovation, and by construction the contemporaneous response of C to ε^Y is zero. This orthogonalization is called the "short run" one.

2.2. Results

Consider first the long run identification. Fig. 1 graphs the impulse response functions of C and Y to both shocks as well as their associated 95% confidence bands, obtained by bootstrapping the VECM. Table 1 reports the corresponding variance decomposition of the process.

These results provide an interesting decomposition of macroeconomic fluctuations. The lower left panel of Fig. 1 clearly shows that consumption virtually does not respond to the transitory shock. This is confirmed by Table 1 which shows that the transitory shock accounts for less than 4% of consumption volatility at any horizon. Conversely, consumption is very responsive to the permanent shock and most of the adjustment dynamics take place in less than one year. In other words, consumption is almost a pure random walk that responds only to permanent shocks and has very little dynamics. On the contrary, short run fluctuations in output are mainly associated with the temporary shocks, which explain more than 60% of output volatility on impact. These patterns are often interpreted as simply reflecting the permanent income hypothesis. If the data corresponded to the consumption and investment decision of an individual facing a fixed interest rate, such interpretation would be correct. However, it must be emphasized that these properties are aggregate properties and not individual level properties, which implies that a coherent explanation to these patterns requires a general equilibrium model that exhibits a permanent–temporary decomposition with no temporary component in consumption. For example, in a standard real business cycle model, a temporary change in technology that generates a persistent increase in investment will also generate – because of general equilibrium constraints – a temporary rise in consumption.
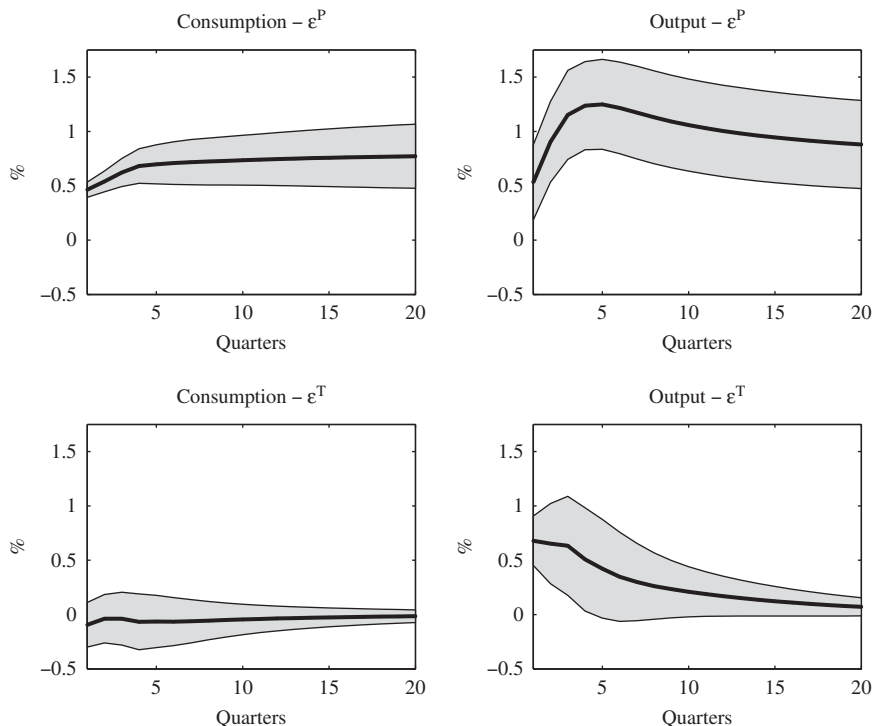
Fig. 1. Responses of output and consumption to ε^P and ε^T. This figure shows the responses of consumption and output to temporary ε^T and permanent ε^P one percent shocks. These impulse response functions are computed from a VECM (C,Y) estimated with one cointegrating relation $[1; -1]$, three lags, using quarterly per capita U.S. data over the period 1947Q1–2004Q4. The shaded area depicts the 95% confidence intervals obtained from 1000 bootstraps of the VECM.

Table 1
The contribution of the shocks to the volatility of output and consumption.

Horizon	Output		Consumption	
	ε^T	ε^Y	ε^T	ε^Y
1	62	80	4	0
4	28	46	1	1
8	17	33	1	1
20	10	22	0	2
∞	0	4	0	4

This table shows the k-period ahead share (in percentage points) of the forecast error variance of consumption and output that is attributable to the temporary shock ε^T in the long run orthogonalization and to the output shock ε^Y in the short run one, for $k=1, 4, 8, 20$ quarters and for $k \longrightarrow \infty$. Those shares are computed from a VECM (C,Y) estimated with one cointegrating relation $[1; -1]$, three lags, using quarterly per capita U.S. data over the period 1947Q1–2004Q4.

Fig. 2. Responses of output and consumption to ε^C and ε^Y. This figure shows the responses of consumption and output to consumption ε^C and output ε^Y one percent shocks obtained from a short run orthogonalization scheme. Those impulse response functions are computed from a VECM (C,Y) estimated with one cointegrating relation $[1; -1]$, three lags, using quarterly per capi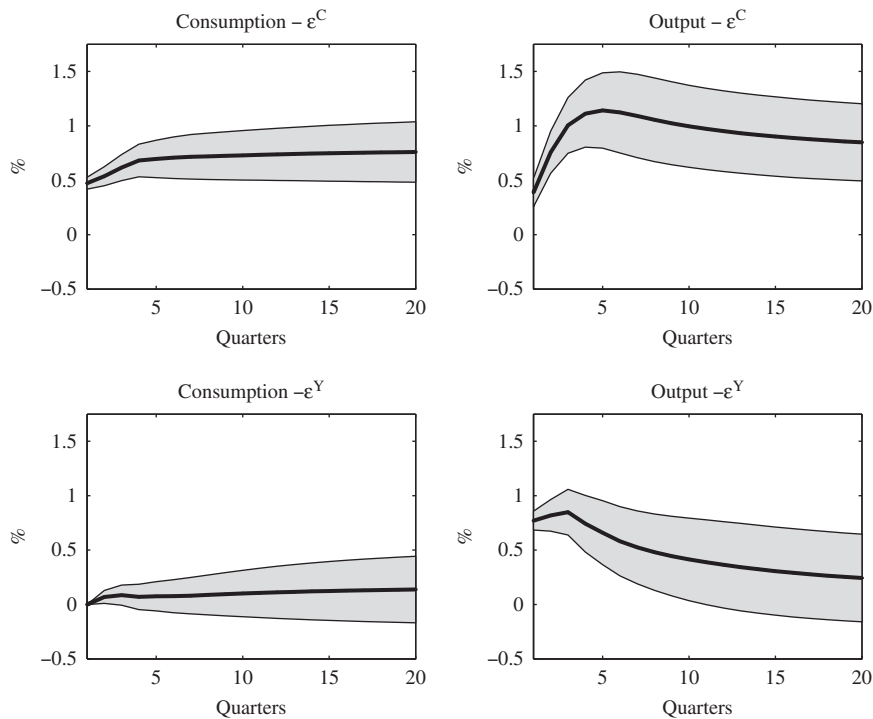ta U.S. data over the period 1947Q1–2004Q4. The shaded area depicts the 95% confidence intervals obtained from 1000 bootstraps of the VECM.

Fig. 2 graphs the impulse responses of C and Y associated with the second orthogonalization scheme. The associated variance decompositions are displayed in Table 1. The striking result from these estimations is that the consumption shock ε^C is almost identical to the permanent shock to consumption (ε^P in the long run orthogonalization scheme), so that the responses and variance decompositions are very similar to those obtained using the long run orthogonalization scheme. This observation is further confirmed by Fig. 3, which plots ε^P against ε^C and ε^T against ε^Y. It is striking to observe that both shocks align along the 45° line, indicating that the consumption innovation is essentially identical to the permanent component.

2.3. The movements of investment and hours worked

Let us now link the behavior of investment and hours worked to the above description of output and consumption. In particular, how much of the variance of those variables is associated with the temporary shock (or quasi-equivalently the output shock) versus the permanent shock recovered from the consumption–output VECM? To answer this question, the following approach is taken. Once the innovations ε^P and ε^T are recovered from the bivariate C–Y VECM, investment in

Fig. 3. Plots of ε^C against ε^P and ε^Y against ε^T. The left panel plots the estimated permanent innovation ε^P (from the long run orthogonalization scheme) against the consumption innovation ε^C (from the short run orthogonalization scheme). The right panel plots the estimated temporary innovation ε^T (from the long run orthogonalization scheme) against the output innovation ε^Y (from the short run orthogonalization scheme). In both panels, the straight line is the 45° line. These shocks are computed from a VECM (C,Y) estimated with one cointegrating relation $[1; -1]$, three lags, using quarterly per capita U.S. data over the period 1947Q1–2004Q4.

Table 2
Variance decomposition of investment and hours worked.

Horizon	Investment			Hours in level			Hours in difference		
	ε^P	ε^T	ε^I	ε^P	ε^T	ε^H	ε^P	ε^T	ε^H
1	1	97	2	19	75	6	21	74	5
4	37	62	1	37	56	7	46	52	2
8	50	48	2	61	32	7	66	32	2
20	44	49	7	60	21	19	69	28	3
40	23	60	17	54	20	26	57	38	5

This table shows the k-period ahead share (in percentage points) of the forecast error variance of hours and investment that is attributable to the temporary and permanent shocks ε^T and ε^P and to the residual shock, for $k=1, 4, 8, 20$ and 40 quarters. Those shares are computed from the estimation of (5). The shocks ε^T and ε^P are obtained in a first stage from the VECM (C,Y) estimated with one cointegrating relation $[1; -1]$, three lags, using quarterly per capita U.S. data over the period 1947Q1–2004Q4.

difference and hours worked (in levels or differences) are regressed on current and lagged values of these two shocks plus a moving-average error term denoted ε^I or ε^H, which is called an investment or hours specific shock.[8] An attractive feature of this approach (compared to estimating a tri-variate VAR) is that it delivers results that are robust to the specification of hours worked (level or difference).[9] More precisely, the regression estimated is

$$x_t = c + \sum_{k=0}^{K} (\alpha_k \varepsilon_{t-k}^P + \beta_k \varepsilon_{t-k}^T + \gamma_k \varepsilon_{t-k}^X), \qquad (5)$$

where x_t denotes either the (log) hours per capita in levels or the (log) difference of hours and investment. This model is estimated by maximum likelihood, choosing an arbitrarily large number of lags ($K=40$). For each horizon k is computed the share of the overall volatility of investment or hours worked accounted for by ε^P, ε^T, and by the specific shock ε^I or ε^H. Results are reported in Table 2.

The numbers reported in the table clearly indicate that investment and hours worked are primarily explained by the transitory component for 1–4 quarter horizons. This transitory component still explains one-half of the variance of investment and one-third of the variance of hours at a 8-quarter horizon. This is also illustrated in Fig. 4 that displays the estimated impulse response function of investment and hours worked to temporary and permanent shocks, as estimated from Eq. (5). The method we use to estimate the response of investment has the disadvantage of working with investment

---

[8] Such a two step strategy amounts to the estimation of the following restricted tri-variate moving-average process:

$$\begin{pmatrix} C_t \\ Y_t \\ X_t \end{pmatrix} = \begin{pmatrix} R(L) & 0_{2,1} \\ S(L) & T(L) \end{pmatrix} \begin{pmatrix} \varepsilon_t^P \\ \varepsilon_t^T \\ \varepsilon_t^X \end{pmatrix},$$

where $R(L)$ is a 2×2 polynomial matrix, $0_{2,1}$ is a 2×1 vector of zeros, $S(L)$ is a 1×2 polynomial matrix and $T(L)$ is a polynomial in lag operator. $R(L)$, ε^P and ε^T are recovered from the first step bivariate VECM, while $S(L)$, $T(L)$, and ε^H are estimated using a truncated approximation of the third line of the above MA process (which is Eq. (5)). In the case of an estimation in difference, X has to be replaced by $(1-L)X$.

[9] It is well known (see for instance the discussions in Gali, 1999; Gali and Rabanal, 2004; Chari et al., 2004; Christiano et al., 2004) that specification choice (levels versus first differences) matters a lot for VARs with hours worked. Results show that our procedure is robust to this specification choice.
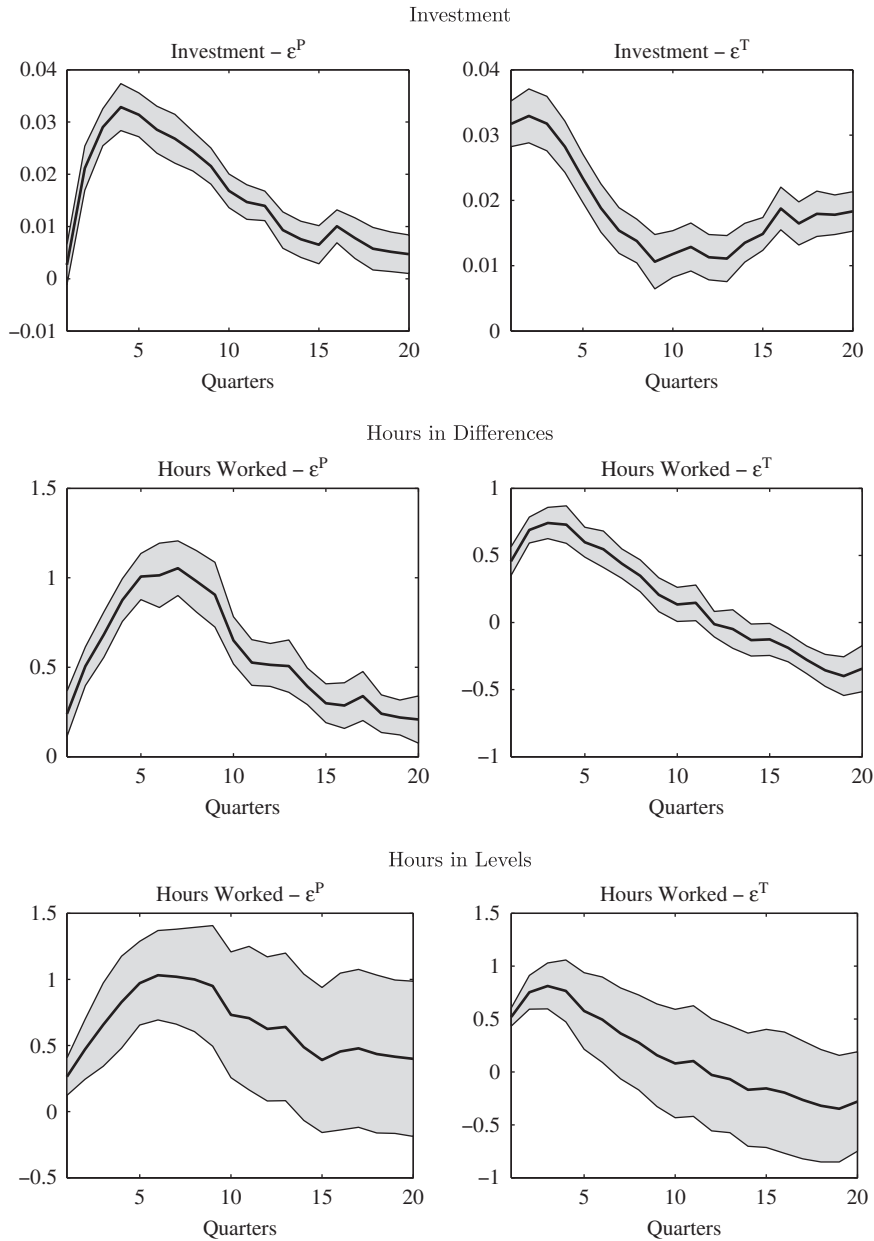
Investment

Hours in Differences

Hours in Levels

Fig. 4. Responses of investment and hours to ε^P and ε^T. This figure shows the response of investment and hours worked to the temporary ε^T and permanent ε^P shocks. Those impulse responses are computed using a two-step procedure. First ε^T and ε^P are derived from the estimation of a VECM (C,Y) with one cointegrating relation $[1;-1]$, three lags, using quarterly per capita U.S. data over the period 1947Q1–2004Q4. Then investment in difference or hours worked (in levels or difference depending on the specification) are projected on current and past values of those innovations plus a moving-average term in ε^I or ε^H. Confidence bands are obtained by a delta method.

first differences, and therefore not taking care of long run relations between investment, consumption and output. As a consequence, the temporary shock ε^T happens to still explain quite a large share of investment after 10 years (60%). An alternative method amounts to estimate a trivariate (Y,C,I) VECM with cointegrating relations between I, C and Y, and impose that the long run impact of the temporary shock is zero. This method is presented in the online technical appendix, and is shown to give very similar short run responses of investment to the temporary shock.

To summarize, four properties of the data are worth highlighting: (i) the permanent shock ε^P, as recovered from a long run restriction in a consumption–output VECM, is essentially the same shock as that corresponding to a consumption shock ε^C, as obtained from an impact restriction, (ii) the response of consumption to a temporary shock is extremely close to zero at all horizons, and there are almost no dynamics in the response of consumption to a permanent shock, as it jumps almost instantaneously to its long run level, (iii) the temporary shock (or the output shock in the short run orthogonalization) is

responsible for a significant share of output volatility at business cycle frequencies and (iv) investment and hours are largely explained by the transitory shock at business cycle frequencies. These facts emphasize that a substantial fraction of the business cycle action seems to be related to changes in investment and hours worked, without any short or long run implications for consumption. It is shown in the online technical appendix that these findings are robust both against changes in the specification of the VECM – by estimating rather than imposing the cointegration relation, adding additional lags, or estimating the VECM in levels – and against the data used to estimate the VECM – taking total consumption rather than the consumption of nondurables and services, measuring output as consumption plus investment only. In all these cases, no major changes in patterns are found. Since some emphasis has been put on the quasi-equivalence between the shocks recovered using a long run restriction, and shocks recovered using an impact restrictions, a formal test[10] for the equality between ε^Y and ε^T is conducted. At a 5% significance level, the hypothesis that the consumption shock is identical to the permanent shock cannot be rejected.

3. An analytical model of market rushes

In this section we present a simple analytical model of market rushes. The main element of the model is that agents receive, each period, information about potential new varieties of goods that could become profitable to produce. In response to these expectations of profits, agents invest in putting on the market a prototype of the new good. Since many agents may invest in such startups, they engage in a winner-takes-all competition for securing the market of a newly created variety. The winning firm becomes a monopolist on the market, but may randomly lose this position at an exogenous rate. Expansion in variety may or may not have a long run impact on productivity, so that the market rush is not forced *a priori* to satisfy the gold rush analogy.

3.1. The model

Firms: There exists a raw final good, denoted Q_t, produced by a representative firm using labor h_t and a set of intermediate goods $X_t(j)$ with mass N_t. The constant returns to scale technology is represented by the production function

$$Q_t = (\Theta_t h_t)^\alpha X_t^{1-\alpha}, \tag{6}$$

where $\alpha \in (0,1)$. Θ_t is an index of disembodied exogenous technological progress and X_t is an aggregate of intermediate goods:

$$X_t = N_t^\xi \left(\int_0^{N_t} X_t(j)^\chi \, dj \right)^{1/\chi}, \tag{7}$$

where $\chi \leq 1$ determines the elasticity of substitution between intermediate goods and ξ is a parameter that determines the long run effect of variety expansion. Since this final good will also serve to produce intermediate goods, Q_t will be referred to as the gross amount of final good. Also note that the raw final good will serve as the numéraire. The representative firm is price taker on the markets.

Each existing intermediate good is produced by a monopolist. It is assumed that the production of one unit of intermediate good requires the use of one unit of the raw final good as input. Since the final good serves as a numéraire, this leads to a situation where the price of each intermediate good is given by $P_t(j) = 1/\chi$. Therefore, the quantity of intermediate good j, $X_t(j)$, produced in equilibrium, is given by

$$X_t(j) = (\chi(1-\alpha))^{1/\alpha} \Theta_t N_t^{\phi-1} h_t, \tag{8}$$

where $\phi = ((1-\alpha)/\alpha)(\xi + (1-\chi)/\chi)$. The profits, $\Pi_t(j)$, generated by intermediate firm j are given by

$$\Pi_t(j) = \pi_0 \Theta_t N_t^{\phi-1} h_t, \tag{9}$$

where $\pi_0 = ((1-\chi)/\chi)(\chi(1-\alpha))^{1/\alpha}$. Equalization of the real wage with marginal product of labor implies

$$W_t = A\Theta_t N_t^\phi, \tag{10}$$

where $A = \alpha(\chi(1-\alpha))^{(1-\alpha)/\alpha}$.

Value added, Y_t, is then given by the quantity of raw final good, Q_t, net of that quantity used to produce the intermediate goods, $X_t(j)$. Substituting out for $X_t(j)$, and taking away the amount of Q_t used in the production of $X_t(t)$, one obtains

$$Y_t \equiv Q_t - \int_0^{N_t} X_t(j) \, dj = B\Theta_t N_t^\phi h_t, \tag{11}$$

where $B = (1 - \chi(1-\alpha))(\chi(1-\alpha))^{(1-\alpha)/\alpha}$.

---

[10] The online technical appendix shows that such a test amounts to testing the nullity of a_{12}, the [1,2] element of the long run matrix of the Wold decomposition. The confidence intervals for the estimate of a_{12} are obtained from 1000 bootstraps of the long run matrix. The coefficient $\hat{a}_{12}$ takes an average value of 0.2024 with a 95% confidence interval $[-0.2, 0.8]$.

The net amount of raw final good Y_t can be used for consumption C_t and startup expenditures S_t:

$$Y_t = C_t + S_t. \tag{12}$$

Variety dynamics: Let $\mathcal{N}_t$ denote the number of potential varieties in period t, and N_t denote the number of active varieties—*i.e.* those which are effectively produced, with $N_t \leq \mathcal{N}_t$. In each period, new potential varieties are created at the stochastic growth rate η_t. The $\mathcal{N}_t$ existing potential varieties of the period become obsolete at an exogenous rate $\mu \in (0,1)$. Therefore, the dynamics for the number of potential products is given by

$$\mathcal{N}_{t+1} = (1 + \eta_t - \mu)\mathcal{N}_t. \tag{13}$$

Note that η brings information about future potentially profitable varieties but does not immediately affect the production function. In the following, there is no drift in $\mathcal{N}$ as we assume $E(\eta_t) = \mu$.

The law of motion of the number of effectively produced goods is driven by an endogenous adoption decision. Any entrepreneur, who desires to produce a potential new variety, has to pay a fixed cost of $\kappa_t \equiv \kappa \Theta_t N_t^{\phi} > 0$ units of the final good to setup the startup. She does so if the expected discounted sum of profits of a startup exceeds κ_t. Let $N_{S,t}$ denote the number of startups and $S_t \equiv \kappa_t N_{S,t}$ denote total expenditures on setup costs. A time $t+1$, a startup will become a functioning new firm with a product monopoly with an endogenous probability ρ_t, and existing monopolies will disappear at rate μ. Therefore, the dynamics for the number of effectively produced goods is given by

$$N_{t+1} = (1-\mu)N_t + \rho_t N_{S,t}. \tag{14}$$

The $N_{S,t}$ startups of period t compete to secure the $\eta_t \mathcal{N}_t$ new monopoly positions. The successful startups are uniformly drawn among the $N_{S,t}$ existing ones. Therefore, the probability that a startup at time t will become a functioning firm at $t+1$ is given by $\rho_t = \min(1, \eta_t \mathcal{N}_t / N_{S,t})$, and the number of new goods created will be $\min(N_{S,t}, \eta_t \mathcal{N}_t)$. If it turns out that startups are not profitable enough, so that $N_{S,t} < \eta_t \mathcal{N}_t$, not all existing varieties will be exploited and therefore $N_t < \mathcal{N}_t$. In order to obtain a tractable solution, parameters are chosen to rule out this case of partial adoption. Allocations will have the property that it is always optimal for entrepreneurs to exploit the whole range of intermediate goods.[11] In other words, it amounts to assuming that the adoption cost κ_t is sufficiently small. This implies that there will be no difference in the model between the potential and the actual number of varieties in equilibrium, so that $N_t = \mathcal{N}_t$ $\forall t$.

Households: The preferences of the representative household are represented by the utility function

$$\mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^{\tau}(\log(C_{t+\tau}) + \psi(\overline{h} - h_{t+\tau})), \tag{15}$$

where $0 < \beta < 1$ is a constant discount factor, C_t denotes consumption in period t and h_t is the quantity of labor the household supplies. The household chooses how much to consume, supply labor, and hold equities in existing firms ($\mathcal{E}_t$) and in startups ($\mathcal{E}_t^S$) by maximizing (15) subject to the following sequence of budget constraints:

$$C_t + P_t^M \mathcal{E}_t + P_t^S \mathcal{E}_t^S = W_t h_t + \mathcal{E}_t \Pi_t + (1-\mu)P_t^M \mathcal{E}_{t-1} + \rho_{t-1} P_t^M \mathcal{E}_{t-1}^S, \tag{16}$$

where P_t^M is the beginning of period (prior to dividend payments Π_t) price of an existing monopoly equity, P_t^S is the price of startups and W_t is the wage rate.

3.2. Equilibrium allocations

The decision to invest in a startup is obtained by combining the first order conditions associated with the household's program and is given by

$$P_t^S = \beta \rho_t \mathbb{E}_t \sum_{\tau=0}^{\infty} \left[\beta^{\tau} \frac{C_t}{C_{t+\tau+1}} (1-\mu)^{\tau} \Pi_{t+\tau+1} \right]. \tag{17}$$

This condition states that the price of a startup is equal to the expected discounted sum of future profits. Free entry of startups drives to zero the expected discounted sum of profits (the right hand side of Eq. (17)) net of the setup cost. Therefore, one has in equilibrium $P_t^S = \kappa_t$. Using this last equation, the labor demand condition (10), the profit equation (9), the resource constraint (12), and the startup equity market equilibrium condition $\mathcal{E}_t^S = N_t^S$, the asset pricing equation (17) becomes

$$(h_t - \psi^{-1}) = \beta \delta_t \frac{\pi_0}{A} \mathbb{E}_t h_{t+1} + \beta \delta_t \mathbb{E}_t \left[\left(\frac{1}{\delta_{t+1}} - 1 \right)(h_{t+1} - \psi^{-1}) \right], \tag{18}$$

where $\delta_t = \eta_t / (1 - \mu + \eta_t)$ is an increasing function of the fraction of newly opened markets η_t.

Eq. (18) is a key equation of the model. It shows that current employment h_t depends on h_{t+1}, δ_t and δ_{t+1}, and therefore indirectly depends on all the future expected δ. As δ_t brings information about the future, employment is purely forward looking. The reason why future employment favors current employment is that higher future employment reflects higher

---

[11] Such an assumption would be definitively not appealing in a growth perspective, or to account for cross-country income differences (see for Comin and Hobijn, 2004), but seems to us acceptable from a business cycle perspective.

expected profits, which therefore stimulates new entries today. Note that the model exhibits certain salient neutrality properties, as the determination of employment does not depend on either current or future changes in disembodied technological change Θ_t.[12]

Iterating forward, the above equation can be written as a function of current and future values of δ only. Given the nonlinearity of Eq. (18), it is useful to compute a log-linear approximation around the deterministic steady-state value of hours worked, h_t[13]

$$\hat{h}_t = \gamma \mathbb{E}_t \hat{h}_{t+1} + \left(\frac{h-\psi^{-1}}{h}\right) \mathbb{E}_t[\hat{\delta}_t - \beta \hat{\delta}_{t+1}],$$

where $\hat{h}_t$ and $\hat{\delta}_t$ now represent relative deviations from the steady state, $h = \psi^{-1}(1-\beta(1-\delta))/(1-\beta\delta\pi_0/A - \beta(1-\delta))$ and $\gamma = \beta\delta(\pi_0/A) + \beta(1-\delta)$ with $\gamma \in (0,1)$.[14] Solving forward, this can be written as

$$\hat{h}_t = \left(\frac{h-\psi^{-1}}{h}\right)\left(\hat{\delta}_t - \beta\delta\left(\frac{A-\pi_0}{A}\right)\mathbb{E}_t\left[\sum_{i=0}^{\infty} \gamma^i \hat{\delta}_{t+1+i}\right]\right). \tag{19}$$

Note that, as $\gamma \in (0,1)$, the model possesses a unique determinate equilibrium path. Eq. (19) reveals that a positive $\hat{\delta}_t$ – *i.e.* an acceleration of variety expansion – causes an instantaneous increase in hours worked, output and investment in startups S. This boom arises as the result of the prospects of future profits derived from securing those new monopoly positions. This occurs irrespective of any current change in the technology or in the number of varieties. Such an expansion is therefore akin to a "demand driven" or "investment driven" boom.

Once the equilibrium path of h_t is computed, output is directly obtained from Eq. (11). Finally, combining labor demand (10) and the household's labor supply decision, one obtains an expression for aggregate consumption:

$$C_t = \frac{A}{\psi}\Theta_t N_t^{\phi}. \tag{20}$$

4. Equilibrium allocation properties

This section first derives the VECM representation of the model solution, and shows the similarity between some orthogonalized representations of the model and of the data. The optimal properties of equilibrium allocations are then discussed. Finally, our results are contrasted with the ones obtained from a baseline RBC model, and we discuss the empirical counterpart to our new markets metaphor.

4.1. A VECM representation of the model solution

As mentioned in the Introduction, it is attractive to represent macroeconomic fluctuations as responses to permanent and transitory shocks in a consumption and output autoregressive vector. We therefore begin by deriving a consumption–output VECM representation of the model solution. This representation will then be compared to the estimated VECM. It is assumed that disembodied technical change Θ_t follows (in log) a random walk without drift: $\log\Theta_t = \log\Theta_{t-1} + \sigma_\Theta \varepsilon_t^\Theta$, where ε_t^Θ are *i.i.d.* with zero mean and unit variance. The variety expansion shock η_t follows an AR(1) process of the form $\log(\eta_t) = \rho_\eta \log(\eta_{t-1}) + (1-\rho_\eta)\log(\tilde{\mu}) + \sigma_N \varepsilon_t^N$, where ε_t^N are *i.i.d.* with zero mean and unit variance, and with $0 < \rho_\eta < 1$.[15] The solution for hours worked is given by $\hat{h}_t = \omega\hat{\eta}_t$, with $\omega \equiv ((h-\psi^{-1})/h)(1-\delta)(1-\beta\rho)/(1-\gamma\rho)$. The logs of consumption and output are therefore given by

$$\log(Y_t) = k_y + \log(\Theta_t) + \phi\log(N_t) + \log(h_t), \tag{21}$$

$$\log(C_t) = k_c + \log(\Theta_t) + \phi\log(N_t), \tag{22}$$

where k_c and k_y are constant terms. Using Eq. (18) to replace h_t with its approximate solution, it is straightforward to derive the following $MA(\infty)$ representation of the system:

$$\begin{pmatrix} \Delta\log(C_t) \\ \Delta\log(Y_t) \end{pmatrix} = \begin{pmatrix} \sigma_\Theta & \frac{\phi L}{1-\rho L}\sigma_N \\ \sigma_\Theta & \frac{(\omega(1-L)+\phi L)}{1-\rho L}\sigma_N \end{pmatrix}\begin{pmatrix} \varepsilon_t^\Theta \\ \varepsilon_t^N \end{pmatrix} = C(L)\begin{pmatrix} \varepsilon_t^\Theta \\ \varepsilon_t^N \end{pmatrix}. \tag{23}$$

---

[12] This result is due to the functional forms chosen for preferences and technology. It is related to (i) the separability between consumption and hours in the utility function; (ii) logarithmic preferences for consumption and (iii) Cobb–Douglas production function.

[13] In the online technical appendix, an exact analytical solution to the model is derived in the case of i.i.d. shocks.

[14] This follows from the restriction $\beta\delta(1-\chi)(1-\alpha)/\alpha + \beta(1-\delta) < 1$ imposed on parameters to guarantee positive hours worked in the non-stochastic steady state.

[15] Note that $\tilde{\mu}$ takes the value $\mu 2(1-\rho_\eta^2)/\sigma_N^2$ so that $E[\eta] = \mu$.

4.2. Orthogonalized representations of equilibrium allocations

If the model of Section 3 is a data generating process, what would it imply for the orthogonalizations performed in Section 2? One way to answer this question would be to simulate data using the model, and then to estimate and orthogonalize a VECM on those simulated data. As our simple model has a tractable analytical solution, it is possible to derive exactly the VECM representation of equilibrium allocations. The impact matrix, $C(0)$, and long run matrix, $C(1)$, can be obtained from the system (23) as

$$C(0) = \begin{pmatrix} \sigma_\Theta & 0 \\ \sigma_\Theta & \omega\sigma_N \end{pmatrix} \quad \text{and} \quad C(1) = \begin{pmatrix} \sigma_\Theta & \frac{\phi}{1-\rho}\sigma_N \\ \sigma_\Theta & \frac{\phi}{1-\rho}\sigma_N \end{pmatrix}.$$

The VECM permanent and transitory shocks are then given by

$$\begin{cases} \varepsilon_t^P = \left(\sigma_\Theta^2 + \left(\frac{\phi}{1-\rho}\right)^2 \sigma_N^2\right)^{-1/2} \left(\sigma_\Theta \varepsilon_t^\Theta + \frac{\phi}{1-\rho}\sigma_N \varepsilon_t^N\right), \\ \varepsilon_t^T = \left(\sigma_\Theta^2 + \left(\frac{\phi}{1-\rho}\right)^2 \sigma_N^2\right)^{-1/2} \left(-\frac{\phi}{1-\rho}\sigma_N \varepsilon_t^\Theta + \sigma_\Theta \varepsilon_t^N\right). \end{cases} \tag{24}$$

Similarly, short run orthogonalization yields

$$\begin{cases} \varepsilon_t^Y = \varepsilon_t^\Theta, \\ \varepsilon_t^C = \varepsilon_t^N. \end{cases} \tag{25}$$

This simple model shares the importance of dynamic properties with the data when the parameter ϕ is set to zero. This corresponds to the case where $\xi = (\chi-1)/\chi$, meaning that an expansion in variety exerts no effect on labor productivity.

First of all the system (23) clearly shows that consumption and output do cointegrate ($C(1)$ is not full rank) with cointegrating vector $[1; -1]$. Second, it shows that consumption is a random walk, that is only affected – in the short run as well as in the long run – by technology shocks, ε^Θ. Output is also affected by the temporary shock, ε^N, in the short run. Hence, computing sequentially our short-run and long-run orthogonalization with this model would imply $\varepsilon^P = \varepsilon^C = \varepsilon^\Theta$ and $\varepsilon^T = \varepsilon^Y = \varepsilon^N$, as it can been seen from (24) and (25) in the case $\phi = 0$. Finally, it is the temporary shock ε^T (which is indeed ε^N) that explains all of hours worked volatility at any horizon, as $\hat{h}_t = \omega\hat{\eta}_t$. Such a model, therefore, allows for a structural interpretation of the results obtained in Section 2. Permanent shocks to C and Y are now interpretable as technology shocks. Consumption does not respond to variety expansion shocks, which however account for a lot of output fluctuations and all the fluctuations in hours worked. Variety expansion shocks create market rushes that are indeed gold rushes, generating inefficient business cycles as the social planner would choose not to respond to them (as shown below). In effect, these shocks only trigger rent seeking activities, as startups are means of appropriating a part of the economy pure profits.

Although simple, this model illustrates how the market mechanism we have put forward has the potential to account for some intriguing properties of the data; in particular, the equivalence of the short and long run identification schemes, and the complete absence of a temporary component in consumption.

4.3. Comparison between equilibrium and optimal allocations

Optimality properties of those allocations are worth discussing, and it is useful to compute the socially optimal allocations as a benchmark. The social planner problem is given by

$$\max \quad \mathbb{E}_t \sum_{i=0}^\infty [\log C_{t+i} + \psi(\overline{h} - h_{t+i})] \quad \text{s.t.} \quad \begin{cases} C_t \leq \hat{A}\Theta_t N_t^\phi h_t - \kappa_t \eta_t N_{S,t}, \\ \mathcal{N}_{t+1} = (1 + \eta_t - \mu)\mathcal{N}_t, \\ N_{t+1} = (1-\mu)N_t + \rho_t N_{S,t}, \\ N_t \leq \mathcal{N}_t, \end{cases}$$

with $\hat{A} = \alpha(1-\alpha)^{\alpha/(1-\alpha)}$ and where one has already solved for the optimal use in intermediate goods. Note that parameters are again assumed to be such that it is always socially optimal to invest in a new variety, so that $N_t = \mathcal{N}_t$. One necessary condition for full adoption to be socially optimal is that the long run effect of variety expansion is positive—*i.e* $\phi > 0 \Longleftrightarrow \xi > -(1-\alpha)(1-\chi)/\chi$. The first order condition of the social planner program is given by

$$\frac{\hat{A}\Theta_t N_t^{(\xi + (1-\alpha)(1/\chi-1))/\alpha}}{\hat{A}\Theta_t N_t^{(\xi + (1-\alpha)(1/\chi-1))/\alpha} h_t - \eta_t N_t \kappa_t} = \psi. \tag{26}$$

There are many sources of inefficiency in the decentralized allocations. One obvious source is the presence of imperfect competition: *ceteris paribus*, the social planner will produce more of each intermediate good. Another one is the congestion

effect associated with investment in startups, because only a fraction ρ_t of startups are successful. The social planner internalizes this congestion effect, and does not duplicate the fixed cost of startups, as the number of startups created is equal to the number of available slots for optimal allocations.[16] Because of these imperfections, the decentralized allocation differs from the optimal allocation along a balanced growth path.

The difference between the market and the socially optimal allocations that we want to highlight regards the response to expected future market shocks. It is remarkable that the socially optimal allocation decision for employment (26) is static, and only depends on η_t (positively). This stands in sharp contrast with the market outcome, as summarized by Eq. (18), in which all future values of η appear. To understand this difference, let us consider an increase in period t in the expected level of η_{t+1}. In the decentralized economy, larger η_{t+1} means more startup investment in $t+1$ and more firms in $t+2$. Those firms will affect other firms' profits from period $t+2$ onward. Therefore, a period t startup will face more competitors in $t+2$, which reduces its current value, and therefore decreases startup investment and output.[17] Such an expectation is not relevant for the social planner, which does not respond to changes in the future values of η. Therefore, in that simple analytical model, part of economic fluctuations are driven by investors (rational) forecasts about future profitability that are inefficient from a social point of view.[18] A stark result is obtained in the case when the returns to variety are nil, so that an expansion in the number for varieties has no long run impact on productivity. This case corresponds to $\phi = (\xi + (1-\alpha)(1-\chi)/\chi)/\alpha = 0$. In this particular case, investment in startups occurs in the decentralized equilibrium in response to market shocks, whereas the social planner would choose not to adopt any new good ($N_t = N_0 \forall t$), as implementing new goods costs κ_t and has no productive effect. In this very case, optimal allocations are invariant to market shocks η, while equilibrium allocations react suboptimally to those shocks. In particular, as hours are only affected by market shocks in equilibrium, all equilibrium fluctuations in hours are suboptimal. This case echoes with an interesting aspect of gold rushes. In effect, from a social point of view, part of the increased activity was wasteful since historically it mainly just contributed to the expansion of the stock of money.

4.4. Properties of an extended analytical RBC model

Let us now contrast the positive properties of our model with those obtained in analytical RBC model that is extended to have both TFP and investment specific shocks. In order to be fully analytical, logarithmic consumption utility, Cobb–Douglas technology and full depreciation is assumed. The representative household has the same preferences as in the preceding model

$$\mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^{\tau} [\log(C_{t+\tau}) + \psi(\bar{h} - h_{t+\tau})]. \tag{27}$$

The final good, Y, is produced according to

$$Y_t = K_t^{\alpha} (\Theta_t h_t)^{1-\alpha}, \tag{28}$$

where Θ_t is an exogenous TFP shock. Capital accumulates as

$$K_{t+1} = Q_t I_t, \tag{29}$$

where Q_t is an investment specific shock and I_t denotes investment.

Equilibrium allocations of such a model are given by[19]

$$h_t = h^{\star} = \frac{1-\alpha}{\psi(1-\alpha\beta)}, \tag{30}$$

$$Y_t = \Gamma_y (Q_{t-1} Y_{t-1})^{\alpha} \Theta_t^{1-\alpha}, \tag{31}$$

$$C_t = \Gamma_c (Q_{t-1} Y_{t-1})^{\alpha} \Theta_t^{1-\alpha}, \tag{32}$$

with $\Gamma_y \equiv (\alpha\beta)^{\alpha} h^{\star 1-\alpha}$ and $\Gamma_c = (1-\alpha\beta)\Gamma_y$.

Note that the saving rate is constant in this analytical model ($C_t = (1-\alpha\beta)Y_t$), so that any shock that does affect output proportionally affects consumption. As such, the model cannot replicate the facts, as the temporary shock increases consumption as much as investment (in percentage points). This rather extreme result is due to the very specific assumptions that were made in order to obtain an analytical solution, but we show in Beaudry et al. (2009) that the impossibility of such a model to replicate the VARs facts highlighted here extends to non-analytical models of that type.

---

[16] Note that it has been assumed here that parameter values are such that it is optimal to adopt all the new varieties. Another potential source of suboptimality would be an over or under adoption of new goods by the market. As shown in Benassy (1998) in a somewhat different setup with endogenous growth, the parameter ξ is then crucial in determining whether the decentralized allocations show too much or too little of new goods adoption.

[17] This is due to the typical "business stealing" effect found in the endogenous growth literature, for example in Aghion and Howitt (1992), and originally discussed in Spence (1976a, 1976b).

[18] The very result that it is socially optimal not to respond to such future shocks is of course not general, and depends on the utility and production function specification. The general result is not that it is socially optimal not to respond to shocks on future η, but that the decentralized allocations are inefficient in responding to those shocks.

[19] See the online technical appendix for a derivation of the model solution.

4.5. Discussion

An important question not yet discussed is the interpretation of "a new market" and the associated empirical observations with regard to its cyclical properties.[20] Our metaphor of new markets describes all new ways of introducing new products given existing technology or using new technologies.[21] Broadly speaking, a new market ranges from producing a newly invented product (say cellular phones) to producing old goods with newly developed uses (fiber-optic cable networks once the use of the internet has exploded) or new ways of designing old products (say producing shirts of a fashionable new color). Given this broad interpretation, it is difficult to obtain a comprehensive measure of our new market margin. In a very narrow sense, one could associate new markets with new firms, and therefore look at Net Business Formation. Net Business Formation is without ambiguity procyclical in the U.S., which is also one of our model predictions if one literally associates N with the number of firms. The problem is that the evidence suggests that smaller firms typically make up the majority of entries and exits, which is insufficient to account for a large share of hours worked and output variance at short horizons. A less restrictive interpretation is to look at variations in the number of establishments and franchises as an additional channel affecting the number of "operating units". The Business Employment Dynamics database documents job gains and job losses at the establishments level and quarterly frequency for the period between the third quarter of 1992 and the second quarter of 2005. Using these observations, Jaimovich (2004) finds that more than 20% of the cyclical fluctuations in job creation is accounted for by opening establishments, which is already a sizable number. Another dimension that could be associated to the new market margin is variation in the number of franchises. As Lafontaine and Blair (2005) show, numerous firms in a variety of industries have adopted franchising as a method of operation. Sales of goods and services through the franchising format amounted to more than 13% of real Gross Domestic Product in the 1980s and 34% of retail sales in 1986. Jaimovich (2004) documents that the variations in the number of franchises are procyclical at the business cycle frequency, which is again in line with the ideas put forward by the model. We take this empirical evidence as supporting the notion that agents' expectations about the possibility of new markets is potentially an important driving force of the business cycle.

5. Conclusion

This paper presented theory and evidence in support of the idea that expectations of new market openings may be a key element in explaining the temporary component in output fluctuations. In particular, we proposed a model where the opening of new market opportunities causes an economic expansion by favoring competition for market share. Such an episode was called a market rush in analogy to a gold rush. A simple analytical model of market rushes has been developed and it has been shown how it can replicate an important qualitative feature of the data, namely that the temporary component extracted from an output–consumption VECM is associated with virtually no movement in consumption at any frequency. It has been demonstrated that such a pattern arises in our model when most of the investment in new varieties is socially inefficient. While such an interpretation of business cycles is certainly controversial, it is worth noting that the properties of the consumption–output VECM suggest that the data can be generated by only two large classes of models. Either the data is generated by a model that does not admit a structural temporary–permanent decomposition, which would be the case if all shocks have permanent effects. Or, the data is generate by a model that does admit a structural temporary–permanent decomposition, in which case the induced temporary fluctuations should be explained in terms of socially inefficient investment as there are no associated gains in terms of consumption even though more work is exerted.[22] The contribution of this paper is to provide a candidate explanation to the second possibility.

A natural follow-up question to this paper is whether the market rush phenomenon can be quantitatively an important source of fluctuations? Such an exploration requires extending the model in several directions to make it more realistic. In a companion paper (Beaudry et al., 2009), we pursue this goal by introducing into the model capital accumulation, two types of intermediate goods and habit persistence in consumption. The extent to which the model is quantitatively capable of replicating the impulse responses presented here is investigated in this extended version. This ongoing work suggests that market rush phenomenon with social wasteful variety expansion may be a significant contributor to business cycle fluctuations.

---

[20] We have here benefited from comments and discussion with Nir Jaimovich.

[21] The new goods margin has been recently shown to be important in understanding the pattern of international trade (see Ghironi and Melitz, 2005; Kehoe and Ruhl, 2006).

[22] For example, the properties of the consumption–output VECM should be viewed as challenging to sticky price theories of the business cycles driven by one permanent shock and one temporary shock. In such models, the temporary shock induces temporary movements in consumption, while such predicted outcome is not apparent in the data.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.moneco.2011.01.001.

References

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60 (2), 323–351.

Beaudry, P., Collard, F., Portier, F., 2009. A theory of transitory macroeconomic shocks. Mimeo, Toulouse School of Economics.

Benassy, J., 1998. Is there always too little research in endogenous growth with expanding product variety? European Economic Review 42 (1), 61–69

Benhabib, J., Farmer, R., 1999. Indeterminacy and sunspots in macroeconomics. In: Talor, J., Woodford, M. (Eds.), Handbook of Macroeconomics, vol. 1A. North-Holland, New York, pp. 387–448.

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

Chari, V., Kehoe, P., McGrattan, E., 2004. A critique of structural vars using real business cycle theory. Working Paper 631, Federal Reserve Bank of Minneapolis.

Christiano, L., Eichenbaum, M., Vigfusson, R., 2004. What happens after a technology shock. Working Paper 9819, National Bureau of Economic Research.

Cochrane, J., 1994. Permanent and transitory components of gnp and stock prices. The Quarterly Journal of Economics 109 (1), 241–265.

Comin, D., Hobijn, B., 2004. Cross-country technology adoption: making the theories face the facts. Journal of Monetary Economics 51 (1), 39–83.

Devereux, M.B., Head, A.C., Lapham, B.J., 1993. Monopolistic competition, technology shocks, and aggregate fluctuations. Economics Letters 41 (1), 57–61.

Gali, J., 1999. Technology, employment and the business cycle: Do technology shocks explain aggregate fluctuations? The American Economic Review 89 (1), 249–271

Gali, J., Rabanal, P., 2004. Technology shocks and aggregate fluctuations: How well does the RBC model fit postwar U.S. data? In: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2004. The MIT Press, Cambridge, MA, pp. 225–288.

Ghironi, F., Melitz, M.J., 2005. International trade and macroeconomic dynamics with heterogeneous firms. The Quarterly Journal of Economics 120 (3), 865–915.

Jaimovich, N., 2004. Firm dynamics, markup variations and the business cycle. Mimeo, Stanford University.

Justiniano, A., Primiceri, G.E., Tambalotti, A., 2010. Investment shocks and business cycles. Journal of Monetary Economics 57 (2), 132–145.

Kehoe, T.J., Ruhl, K.J., 2006. How important is the new goods margin in international trade? 2006 Meeting Papers 733, Society for Economic Dynamics.

Lafontaine, F., Blair, R., 2005. The Economics of Franchising. Cambridge University Press, Cambridge, UK.

Romer, P.M., 1987. Growth based on increasing returns due to specialization. The American Economic Review 77 (2), 56–62.

Romer, P.M., 1990. Endogenous technological change. The Journal of Political Economy 98 (5), S71–S102.

Sims, C., 1980. Macroeconomics and reality. Econometrica 48 (1), 1–48.

Smets, F., Wouters, R., 2007. Shocks and frictions in us business cycles: a Bayesian DSGE approach. American Economic Review 97 (3), 586–606.

Spence, M., 1976a. Product differentiation and welfare. The American Economic Review Papers and Proceedings 76, 407–414.

Spence, M., 1976b. Product selection, fixed costs, and monopolistic competition. Review of Economic Studies 43, 217–235.

Uhlig, H., 2005. What are the effects of monetary policy on output? results from an agnostic identification procedure. Journal of Monetary Economics 52 (2), 381–419.

Whelan, K., 2003. A two-sector approach to modeling U.S. NIPA data. Journal of Money, Credit and Banking 35 (4), 627–656.