Journal of Monetary Economics 60 (2013) 737–751

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Long run productivity risk and aggregate investment

Jack Favilukis [a,*], Xiaoji Lin [b]

[a] Department of Finance, London School of Economics and Political Science and FMG, Houghton Street, London WC2A 2AE, UK
[b] Fisher College of Business, Ohio State University, United States

ARTICLE INFO

Article history:
Received 25 November 2011
Received in revised form
10 May 2013
Accepted 14 May 2013
Available online 23 May 2013

ABSTRACT

Long-run productivity risk – shocks to the growth rate of productivity – offers an alternative to microfrictions explanations of aggregate investment non-linearities, in particular the heteroscedasticity of investment rate. Additionally, consistent with the data, these shocks imply that investment rate is history dependent (rising through expansions), its growth is positively autocorrelated, and it is positively correlated with output growth at various leads and lags. A standard model with shocks to the level of productivity either predicts opposite investment behavior or fails to quantitatively capture these features in the data.

1. Introduction

U.S. nonresidential private fixed investment displays nonlinearities and the causes of these non-linearities have been a source of debate for macroeconomists. This is because the behavior of aggregate investment can shed light on the importance of adjustment costs to firms, on the nature of the shocks affecting the economy, and on household preferences. In particular, Caballero et al. (1995), Caballero and Engel (1999), Cooper et al. (1999) show that in partial equilibrium non-convex adjustment costs can lead to investment nonlinearities (i.e. non-linear responses of investment to shocks, such as heteroscedasticity). Thomas (2002) and Khan and Thomas (2008) (henceforth KT) argue that general equilibrium effects on prices undo much of this while Bachmann et al. (2011) (henceforth BCE) show that non-convex frictions can give rise to conditional heteroscedasticity in a DSGE model providing a counterexample to KT.[1]

While non-convex costs may be important, we offer an alternative explanation for the behavior of aggregate investment. The key finding is that shocks to the growth rate, as opposed to the level, of total factor productivity (TFP) naturally imply that the aggregate investment rate is heteroscedastic if households have preference for smoothing consumption over time.[2] Moreover, beyond explaining the conditional heteroscedasticity in the aggregate investment rate, the model generates other interesting dynamics in aggregate investment that are difficult to explain by standard models. In particular, as in the data, the model implies that the investment rate is history dependent in that longer expansions are associated with larger

* Corresponding author. Tel.: +44 207 955 6948.
E-mail addresses: j.favilukis@lse.ac.uk, jack.favilukis@gmail.com (J. Favilukis).

[1] The heart of the debate lies in whether general equilibrium forces can cancel out aggregate investment demand implied by micro lumpy investment. Consider firms that face a non-convex (i.e. fixed) cost to invest; such firms will have a cut off rule in deciding whether to invest a large amount or none at all. Suppose many firms are just below the cutoff and not investing, then a small positive aggregate shock can drive a large number of firms over the hump, resulting in large swings of investment as everyone suddenly invests (extensive margin). Without fixed costs firms will only adjust the quantity of investment (intensive margin) and such large swings in response to small shocks would not occur. According to the intuition in KT, general equilibrium forces prevent a large number of firms from concentrating just below the cutoff. This is because investment is valuable and some firms would invest earlier in expectation of higher returns. However, BCE show that both adjustment costs and general equilibrium forces play a relevant role. In particular, when extensive margin is calibrated to have a more important role in shaping aggregate investment than general equilibrium constraints, non-convex frictions can have a consequential effect on aggregate quantities.

[2] The growth rate shocks are exactly the types of shocks necessary to produce high Sharpe Ratios in a Long Run Risk model such as Bansal and Yaron (2004).

0304-3932/$ - see front matter
http://dx.doi.org/10.1016/j.jmoneco.2013.05.002

increases in investment, that investment rate growth is positively autocorrelated, and that investment rate growth is positively correlated with output growth at various leads and lags. A standard model with shocks to the level of TFP cannot produce these features of aggregate investment. Finally, it is shown that if growth rate shocks are the drivers of business cycles, then matching the joint behavior of consumption, investment, and hours implies that the intertemporal elasticity of substitution (IES) should not be too low.

The importance of modeling growth rate (permanent) shocks and the interaction of such shocks with the IES has been a hotly discussed topic in finance (i.e. Bansal and Yaron, 2004; Alvarez and Jermann, 2005). Beyond providing a potential explanation of aggregate investment behavior, the findings offer additional confirmation for the importance of such shocks and for the likely range of the IES, even independently of asset pricing considerations.

The conditional heteroscedasticity in the aggregate investment rate refers to the conditional volatility of the investment rate being high in times of high past investment (see Fig. 1). As mentioned above, BCE show that this can be explained by non-convex adjustment costs. The model with growth rate shocks naturally implies that the investment rate is heteroscedastic, even without adjustment costs, as long as households prefer to smooth consumption over time. When there is no preference for smoothing consumption over time (this corresponds to an infinite intertemporal elasticity of substitution), capital adjusts to its optimal target capital (which is implied by the level of productivity and does not depend on past capital) immediately; therefore the investment rate is perfectly correlated with the realized growth rate of technology. If there is no heteroscedasticity in this growth rate, there will not be heteroscedasticity in the investment rate. When households have a preference for smoothing consumption, the realized investment rate will be positively related to both the realized growth rate of technology, and to the past investment rate; this is because in the past households were smoothing consumption and did not fully adjust capital to the long term trend. As a result, the high past investment rate amplifies shocks to the growth rate of technology, making the conditional volatility of investment rate higher when past investment rate is higher. The relationship between the investment rate, the growth rate of technology, and the past investment rate is true for both permanent (growth) shocks and transitory (level) shocks. However, because for level shocks the growth rate of technology and investment rate are negatively correlated, the amplification mechanism resulting in heteroscedasticity fails unless the transitory shocks are extremely persistent. Note that the growth rate of an AR(1) process is negatively related to its level (this weakens with high persistence). Because the past investment rate is high when the level of transitory productivity is high, which in turn is associated with a lower growth rate of future productivity, this works to dampen the mechanism resulting in near zero heteroscedasticity in a model with only transitory shocks.

History dependence is another important feature of aggregate investment. It refers to the investment rate rising through expansions, and falling through recessions (see Fig. 2). This feature of investment naturally arises in a model with growth rate shocks but *not* a model with level shocks. The intuition is as follows. When, as in standard models, shocks are to the level of productivity, firms have an optimal level of capital associated with each productivity level. When the productivity level increases due to a positive shock, so does optimal capital and firms choose the investment rate based on the distance to the optimum. Subsequent positive shocks are counterbalanced by mean reversion, resulting in little change to the currently optimal capital levels. The result is an initial jump in the aggregate investment rate, followed by a slow decline towards the long-run average, even as more positive shocks arrive. This is because firms are closer and closer to their optimal target capital. On the other hand, when the growth of productivity is persistent, a shock to productivity implies a permanent change in the level of productivity. Subsequent positive shocks are again counterbalanced by mean reversion, but this time it is the growth rate, rather than level of productivity that stays high. This results in further increases to productivity and to the optimal target capital, requiring even more investment. Thus the investment rate is history dependent, growing (falling) as the expansion (recession) gets longer.

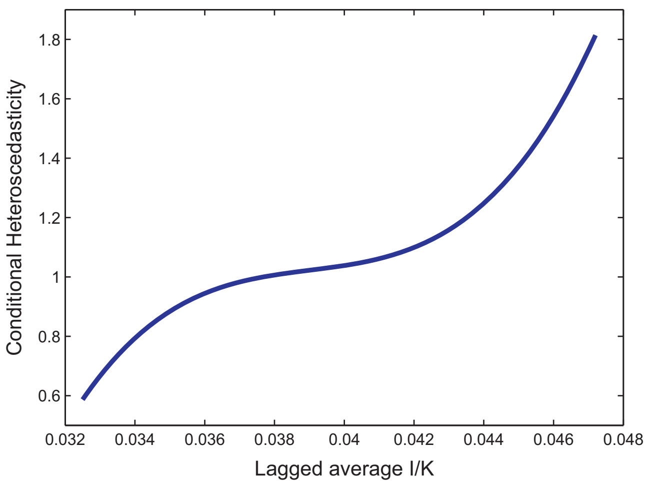

Fig. 1. Heteroscedasticity. This figure plots the heteroscedasticity range following the analysis in Bachmann et al. (2011). First I/K is regressed on its own lag, then residuals from this regression are squared. The mean of square residuals for each level of I/K lagged is plotted against I/K lagged.

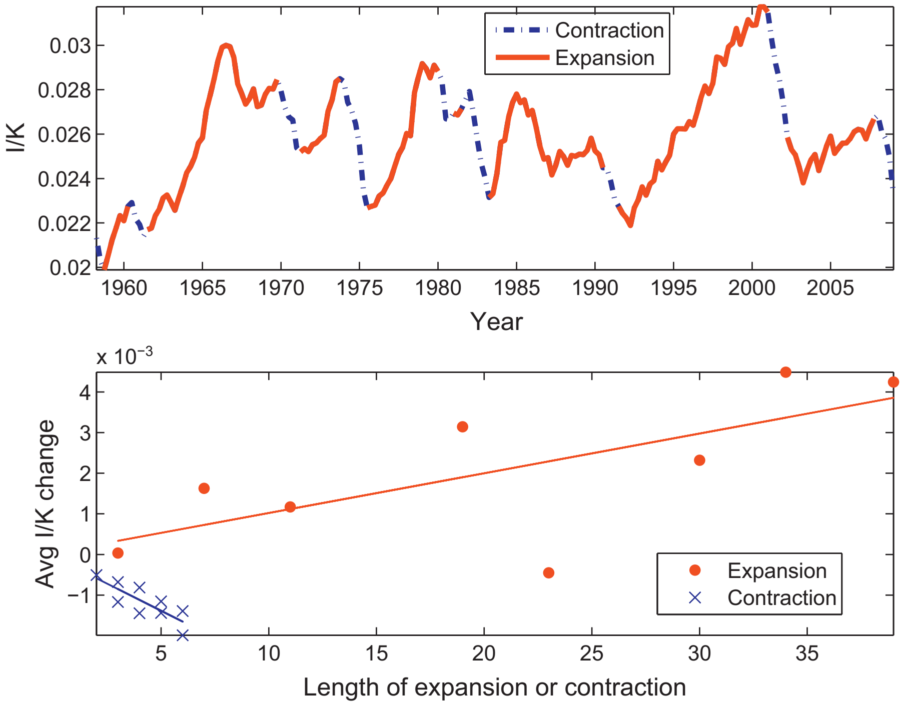

Fig. 2. History dependence. This figure displays the history dependence of aggregate investment rate. The change in investment rate is defined as in the text: $\overline{\Delta I/K}_j = (1/T_j)\sum_{t=0_j,T_j} I_t/K_t - I_0/K_0$ where T_j is the length of expansion or recession j. The top panel plots investment rate over time, with NBER contractions dashed. The bottom panel plots $\overline{\Delta I/K}$ against the length of the associated expansion (contraction).

In the data, investment rate growth is fairly persistent and positively correlated with output growth at various leads and lags. The reason that the growth shock model is able to match this persistence is that persistent shocks to the growth rate of TFP make the growth rate of output persistent and the growth rate of investment rate follows. Furthermore, because the investment rate growth and output growth cointegrate with the growth rate of TFP, these two series positively correlate at various leads and lags. On the other hand, a model with level shocks fails to match this behavior of the investment rate. This failure comes about because shocks to the level of TFP imply that TFP growth is negatively autocorrelated, which in turn implies negative autocorrelation of investment rate growth and low correlation between output and investment growth at various leads and lags.

Although the main focus is investment, this paper also explores implications for the cyclical behavior of employment and consumption. As in the data, in a model with growth rate shocks, a high intertemporal elasticity of substitution (strong substitution effect) leads to positive correlations between output, consumption, investment, and employment while a low IES (strong wealth effect) leads to negative correlations among some of these quantities. Recall that matching the heteroscedasticity of investment rate in the model required that households have preferences for smoothing consumption and leisure over time, that is the IES should not be too high. Taken together these suggest rough lower and upper bounds on the IES; this quantity is of great interest to economists, for example Hansen and Singleton (1982), Attanasio and Weber (1989), and Attanasio and Vissing-Jorgensen (2003) estimate the IES to be well above one, while Hall (1988), Campbell and Mankiw (1989), and Campbell (1999) cannot reject the IES being zero; Bansal et al. (2005) and Bansal and Shaliastovich (forthcoming) provide further evidence on the magnitude of the IES using data from financial markets.

The remainder of the paper is organized as follows. The next section briefly reviews the empirical facts. Section 3 presents a dynamic general equilibrium model. Section 4 discusses why the model is able to match the data. Section 5 concludes and is followed by several appendices.

2. Empirical facts

This section reviews several features of the data.[3] The following section presents a model which captures these features of the data (Table 1).

[3] The data source is standard, from the National Income and Product Accounts available from the Bureau of Economic Analysis (BEA). The sample period is 1958:I–2008:IV. A detailed construction of variables is in the online appendix which can be found at http://personal.lse.ac.uk/faviluki/research.html.

Table 1
Parameter values in the benchmark calibration. This table presents the values of parameters used in the benchmark calibration.

Notation	Value	Description
β	0.995	Subjective discount factor
ψ	1.5	Elasticity of intertemporal substitution
ϕ	0.35	Determines elasticity between consumption and leisure
γ	4	Risk aversion
α	0.36	Share of capital in output production
δ	0.025	Rate of depreciation for capital
g	0.005	Growth rate of log productivity
ρ_X	0.24	Persistence of the growth rate of log productivity
σ_X	0.0093	Conditional volatility of the growth rate of log productivity
ρ_L	0.96	Persistence of the level of log productivity
σ_L	0.0066	Conditional volatility of the level of log productivity

2.1. *Heteroscedasticity*

Bachmann et al. (2011) show that the volatility in the investment rate (I/K) is high when the past investment rate is high. To show this, BCE regress the aggregate investment rate on its lag and compute a time-series of residuals. They then regress the absolute values of these residuals (specification 1 in Table 3) or the squared residuals (specification 2 in Table 3) on the lagged investment rate and find that the slope of this regression is positive, indicating higher volatility of innovations to investment rate following high past investment rate.

To quantify this heteroscedasticity range they define the heteroscedasticity range as $\log(\sigma_{95}/\sigma_5)$ where σ_x is the fitted value of this regression at the x percentile of the investment rate distribution. Similar to Bachmann et al. (2011), the heteroscedasticity range is positive and significant, being equal to 0.16 and 0.22 for specifications 1 and 2, respectively. This indicates that the investment rate is more volatile when its lag is high.

The heteroscedasticity range is also computed for the output-to-capital and consumption-to-capital ratios, these are in Panels B and C of Table 3. Unlike the investment rate, these are insignificant and the point estimates are smaller in magnitude.

2.2. *History dependence*

The history dependence of investment rate is defined as the observation that the change in investment rate depends on the length of time the economy has been in a boom or bust. To summarize history dependence, the change in the investment rate is calculated for each recession and expansion. In particular, the quantity $\overline{\Delta I/K}_j = (1/T_j)\sum_{t=0_j,T_j} I_t/K_t - I_0/K_0$ is computed, where the expansion or recession j starts at 0_j and ends at T_j. This measure is plotted against the length of each recession and expansion in the lower panel of Fig. 2, where recessions are defined by NBER quarterly recession dates. Note that longer expansions are associated with larger rises in the investment rate while longer recessions are associated with larger falls.

Because NBER dates do not exist in the model, recessions are also defined in a statistical way: a recession is defined as any two consecutive quarters with negative output growth (this produces similar dates as NBER).[4] Next, $\overline{\Delta I/K}$ is regressed on the length of the associated recession; this is then separately done for expansions. The slope coefficients from these regressions are b^R and b^E, reported in Table 4. For both definitions of booms and busts, b^R is negative and significant and b^E is positive and significant.[5] A positive b^E implies that the investment rate tends to rise as the expansion gets longer, similarly a negative b^R implies that the investment rate tends to fall as a recession gets longer.

History dependence is also computed for the output-to-capital and consumption-to-capital ratios, these too, are in Table 4. These ratios follow a similar pattern as investment-to-capital: larger rises for longer recessions and larger falls for longer recessions.

2.3. *Higher order autocorrelations and cross-correlations*

Table 5 presents the higher order autocorrelations of the level of the investment rate, and of the growth rates of investment, output and consumption. The investment rate is highly persistent, ranging from 0.97 in the first order and slower declining to 0.56 at the sixth order. The growth rate of the investment rate is less persistent than the investment rate itself, but still quite persistent; for example the first order autocorrelation is 0.42 and remains positive for four quarters.

[4] As another alternative, the Bry and Boschan (1971) procedure was used to define recessions. Both procedures yield very similar results, therefore for brevity, the Bry–Boschan results are omitted from the text.

[5] Several other ways of defining recessions all deliver consistent results.

Table 2
Business cycle moments. This table presents the standard Business Cycle moments in the data and the simulated moments in the growth rate shock model and the level shock model. All lower case variables are in logs; c is consumption, y is output, i is investment, n is hours, and k is capital. All quantities except for Δc and I/K are HP filtered. This data comes from NIPA from 1958 to 2008. Hours come from BLS, the data starts in 1964. Bootstrapped standard errors for the data are in parentheses.

	Data			Growth shock			Level shock		
	$\sigma(x)$	AC(x)	$\rho(y,x)$	$\sigma(x)$	AC(x)	$\rho(y,x)$	$\sigma(x)$	AC(x)	$\rho(y,x)$
y	1.54	0.86	1.00	1.51	0.86	1.00	1.44	0.73	1.00
	(0.08)	(0.02)	(0.00)						
c	1.19	0.86	0.89	0.74	0.82	0.96	0.40	0.81	0.73
	(0.07)	(0.03)	(0.01)						
i	4.76	0.91	0.79	3.37	0.86	0.99	4.19	0.71	0.99
	(0.19)	(0.02)	(0.03)						
n	0.42	0.73	0.72	0.55	0.85	0.97	0.79	0.70	0.97
	(0.02)	(0.03)	(0.04)						
Δc	0.74	0.24	0.18	0.49	0.20	0.86	0.26	0.08	0.64
	(0.06)	(0.06)	(0.09)						
I/K	0.25	0.96	0.54	0.21	0.98	0.47	0.17	0.90	0.67
	(0.01)	(0.00)	(0.05)						

Table 3
Heteroscedasticity range of aggregate quantities. This table reports the conditional heteroscedasticity range for investment rate (I/K), the output-to-capital ratio (Y/K), and the consumption-to-output ratio (C/K). Specifications 1 and 2 are following BCE. An AR(1) process is estimated for I_t/K_t using OLS, then the residuals from this regression, denoted as ε_t, are used to estimate α and η via OLS from the two specifications below following BCE.

$$\text{Specification 1}: \ |\varepsilon_t| = \sqrt{\frac{2}{\pi}}(\alpha_1 + \eta(I_{t-1}/K_{t-1})) + \text{error}$$

$$\text{Specification 2}: \ \varepsilon_t^2 = \alpha_2 + \eta(I_{t-1}/K_{t-1}) + \text{error}.$$

t_η is the Newey–West t-statistics. σ_x is the fitted value of the x percentile of the regression. To compute t-statistics for simulated data, 25,000 quarters are simulated and broken into 122 intervals of 204 quarters each; the average t-statistic over these 122 intervals are reported.

Specification	I/K		Y/K		C/K	
	1	2	1	2	1	2
Panel A: Data						
$\eta*1000$	12.07	0.02	−5.71	−0.23	5.48	0.07
τ_η	2.01	3.27	−0.55	−0.93	0.56	0.49
$\log\left(\frac{\sigma Max}{\sigma Min}\right)$	0.23	0.33	−0.15	−0.24	0.15	0.13
$\log\left(\frac{\sigma 95}{\sigma 5}\right)$	0.16	0.22	−0.11	−0.18	0.12	0.10
Panel B: Growth shock						
$\eta*1000$	16.40	0.01	13.10	0.01	−15.37	0.00
τ_η	1.68	1.43	1.32	1.28	−1.11	−0.56
$\log\left(\frac{\sigma Max}{\sigma Min}\right)$	0.28	0.21	0.22	0.18	−0.21	−0.09
$\log\left(\frac{\sigma 95}{\sigma 5}\right)$	0.21	0.16	0.16	0.14	−0.15	−0.07
Panel C: Level shock						
$\eta*1000$	−3.06	−0.01	2.91	0.00	13.74	0.00
τ_η	−0.05	−0.11	0.13	0.01	0.73	0.67
$\log\left(\frac{\sigma Max}{\sigma Min}\right)$	−0.04	−0.06	0.05	0.00	0.26	0.25
$\log\left(\frac{\sigma 95}{\sigma 5}\right)$	−0.03	−0.04	0.03	−0.01	0.19	0.17

Output and consumption growth have a similar positive and declining pattern with consumption growth being somewhat more persistent at higher orders.

Panel A of Table 6 presents the correlations between HP-filtered output, consumption, investment, and hours. All four exhibit fairly high correlations with the lowest correlation being 0.59, between investment with hours. In Panel B this table reports the correlations at three leads and lags between investment rate growth and output growth. Investment growth is positively correlated at lags of two quarters and leads of three quarters with magnitude declining in both leads and lags.

Table 4
History dependence of aggregate quantities. This table reports the measure of history dependence for investment rate (I/K), the output-to-capital ratio (Y/K), and the consumption-to-output ratio (C/K). The measure is the slope of regressing the difference between the average I/K during expansions (recessions) and I/K at the start of expansions (recessions) on the length of expansions (recessions). b^E and b^R denote slopes of expansions and recessions, respectively; *t*-stat denotes the Newey and West *t*-statistics. To compute *t*-statistics for simulated data, 25,000 quarters are simulated and broken into 122 intervals of 204 quarters each; the average *t*-statistic over these 122 intervals are reported. Recessions are defined as two consecutive quarters of negative TFP growth.

	Data			Growth shock			Level shock		
	I/K	Y/K	C/K	I/K	Y/K	C/K	I/K	Y/K	C/K
b^E*1000	0.10	2.49	1.72	0.10	0.14	0.04	0.03	0.02	−0.01
t-Stat	21.48	12.80	11.40	7.19	6.84	6.09	2.37	1.60	−2.46
b^R*1000	−0.27	−12.03	−5.30	−0.19	−0.28	−0.09	−0.11	−0.11	0.00
t-Stat	−14.24	−17.46	−10.31	−10.70	−11.00	−10.33	−5.50	−4.68	−0.10

Table 5
Autocorrelations of aggregate quantities. This table reports the autocorrelations for investment rate (I/K), investment rate growth rate (Δ(I/K)), logged output growth rate (Δy), logged consumption growth rate (Δc). Data comes from NIPA from 1958 to 2008. Bootstrapped standard errors for the data are in parentheses.

Lag	1	2	3	4	5	6
Panel A data						
AC(I/K)	0.97	0.91	0.84	0.75	0.65	0.56
	(0.00)	(0.01)	(0.02)	(0.03)	(0.04)	(0.05)
AC(Δ(I/K))	0.42	0.27	0.13	0.17	−0.08	−0.17
	(0.05)	(0.06)	(0.07)	(0.07)	(0.07)	(0.06)
AC(Δc)	0.41	0.43	0.37	0.32	0.33	0.37
	(0.06)	(0.07)	(0.06)	(0.07)	(0.06)	(0.06)
AC(Δy)	0.37	0.31	0.20	0.22	0.14	0.17
	(0.07)	(0.08)	(0.07)	(0.08)	(0.08)	(0.06)
Panel B growth rate shock						
AC(I/K)	0.97	0.93	0.88	0.84	0.80	0.76
AC(Δ(I/K))	0.35	0.05	−0.04	−0.06	−0.06	−0.08
AC(Δc)	0.20	0.15	0.11	0.10	0.14	0.07
AC(Δy)	0.34	0.09	0.01	−0.01	0.01	−0.03
Panel C level shock						
AC(I/K)	0.89	0.79	0.70	0.61	0.54	0.47
AC(Δ(I/K))	−0.04	−0.04	−0.02	−0.07	−0.02	−0.02
AC(Δc)	0.09	0.18	0.12	0.14	0.12	0.09
AC(Δy)	−0.01	−0.02	−0.01	−0.05	−0.01	−0.02

3. The model

This section presents a dynamic general equilibrium model with a representative household and representative firm. The key departure from the existing investment literature (e.g., Caballero and Engel, 1999; Khan and Thomas, 2008; Bachmann et al., 2011, etc.) is that the underlying shocks that drive the business cycles are shocks to the growth rate of productivity, as opposed to the level of productivity as in standard macroeconomic models (Prescott, 1986). This implies permanent changes to the level of productivity as in the long run risk literature (Bansal et al., 2008; Kaltenbrunner and Lochstoer, 2010; Croce, 2012).

3.1. Households

There is a representative household with Epstein–Zin preferences given by

$$U_t = \{(1-\beta)[C_t^{\phi}(1-N_t)^{1-\phi}]^{1-1/\psi} + \beta E_t[U_{t+1}^{1-\gamma}]^{(1-1/\psi)/(1-\gamma)}\}^{1/(1-1/\psi)} \tag{1}$$

where ψ is the intertemporal elasticity of substitution, $\gamma*\phi$ is risk aversion, ϕ determines the relative importance of consumption and leisure, and β is the subjective discount factor. $C_t \geq 0$ denotes consumption, N_t denotes labor supply, and U_{t+1} denotes the continuation value of utility. In the model financial markets are complete, therefore the representative household receives labor income, chooses between consumption, leisure, and saving, and maximizes utility U_t. The IES determines preference for smooth consumption and leisure over time with low values of ψ implying stronger preferences for smoothness; as will be discussed below, it plays a crucial role in explaining the empirical facts.

Table 6
Correlations of output, investment, consumption, and hours. Panel A of this table reports correlations of output, investment, consumption, and hours. The first column is the data, the other columns are from the model for different intertemporal elasticities of substitution (ψ), which are listed in the second row. Panel B reports the lead and lag correlations between investment rate growth rate and log output growth rate, in particular, j indicates the correlation of Δy_t and $\Delta(\mathrm{I/K})_{t+j}$. The data comes from NIPA from 1958 to 2008. Hours come from BLS, the data starts in 1964. Bootstrapped standard errors for the data are in parentheses.

	Data	Growth shocks				Level shocks			
		0.10	0.25	0.50	1.50	0.1	0.25	0.50	1.50
Panel A									
YI	0.80	0.49	0.72	0.90	0.99	1.00	1.00	0.99	0.99
	(0.03)								
YC	0.89	0.80	0.84	0.91	0.96	0.99	0.98	0.96	0.73
	(0.01)								
YN	0.72	−0.06	0.84	0.89	0.97	0.99	0.99	0.99	0.98
	(0.04)								
IC	0.68	−0.13	0.24	0.65	0.91	0.98	0.96	0.93	0.61
	(0.04)								
IN	0.59	0.83	0.84	0.89	0.99	1.00	1.00	1.00	1.00
	(0.04)								
CN	0.66	−0.65	−0.33	0.23	0.85	0.98	0.95	0.91	0.55
	(0.04)								

Lag/Lead (j)	−3	−2	−1	0	1	2	3
Panel B							
Data	−0.10	0.06	0.14	0.47	0.33	0.21	0.15
	(0.08)	(0.10)	(0.07)	(0.07)	(0.06)	(0.06)	(0.07)
Growth rate shocks	0.01	0.07	0.30	0.94	0.36	0.03	−0.08
Level rate shocks	0.00	−0.01	0.00	0.98	−0.06	−0.07	−0.04

3.2. Firms

The representative firm has the standard Cobb–Douglas production function

$$Y_t = K_t^{\alpha}(X_t N_t)^{1-\alpha}, \quad 0<\alpha<1 \tag{2}$$

where Y_t is aggregate output, X_t is an exogenous, labor-enhancing technology level, N_t is hours worked supplied by the households, and K_t denotes the capital stock.

The representative firm's capital accumulation equation is given by

$$K_{t+1} = (1-\delta)K_t + I_t, \quad 0<\delta<1 \tag{3}$$

where δ is the capital depreciation rate and I_t is gross investment.

3.3. Technology

The key assumption that distinguishes this model from the existing literature is the formulation of productivity process. In particular, in the baseline model the process for TFP is given by

$$\log\left(\frac{X_{t+1}}{X_t}\right) = (1-\rho_G)*g + \rho_G*\log\left(\frac{X_t}{X_{t-1}}\right) + \sigma_G \epsilon_{t+1}^G \tag{4}$$

where g is the mean growth rate of the economy, ρ_G and σ_G are the persistence and conditional volatility of the growth rate, $\log(X_{t+1}/X_t)$. Note that ϵ_{t+1}^G is a transitory shock to the growth rate of productivity which permanently affects the level of TFP. Below, these shocks are referred to as growth rate shocks and denote them X_t^G. The baseline model with growth rate shocks will also be compared to a model where the level of productivity is trend stationary. In this case TFP follows:

$$\log(X_{t+1}) = g*(t+1) + \rho_L*\log(X_t) + \sigma_L \epsilon_{t+1}^L. \tag{5}$$

Note that in this case ϵ_{t+1}^L is a transitory shock to the level of productivity.[6] Below, these are referred to as level productivity shocks and denoted as X_t^L.

The resource constraint is standard given by

$$Y_t = C_t + I_t \tag{6}$$

[6] The estimate of the total factor productivity process includes both growth and level components together, i.e., $TFP_{t+1} = X_{t+1}^G X_{t+1}^L$ where X_{t+1}^G and X_{t+1}^L follow Eqs. (4) and (5). The persistence and conditional volatility of X_{t+1}^L, ρ_L and σ_L, are economically tiny. As such, the baseline model is solved only with growth rate shocks X_{t+1}^G and then compared to a model with level productivity X_{t+1}^L.

Here, and throughout the analysis we exploit the second welfare theorem and find the equilibrium allocations by solving relevant planning problem. The planning problem for this model is: maximize the household utility in Eq. (1) subject to Eqs. (2), (3), (4), and (6).

3.4. Discussion

At this stage it may be useful to compare this model with two closely related models: KT and BCE. The model written down so far differs from KT and BCE in three dimensions: (i) the model is a representative firm model and abstracts from micro frictions. This is because growth rate shocks are sufficient to give rise to nonlinearities while non-convex micro frictions are not necessary.[7] (ii) In this model, the growth rate of TFP is stationary making TFP itself non-stationary as in Bansal and Yaron (2004), whereas TFP is trend stationary in KT and BCE; (iii) preferences are recursive as in Epstein and Zin (1989), which allows for a separation between risk aversion and intertemporal elasticity of substitution of households, while both KT and BCE employ CRRA preference.

4. Main findings

This section analyzes how shocks to the growth rate of productivity affect aggregate investment and employment dynamics. It explains why the growth rate model is able to reproduce the features of aggregate investment described above, while the level model cannot. The aim is to understand why and how long-run productivity shocks affect aggregate investment dynamics and to evaluate how well a model calibrated to match the standard moments of macroeconomic variables is able to match these.

First, the calibrated benchmark model with growth rate technology shocks is presented. It is shown that the growth rate shocks lead to conditional heteroscedasticity and history dependence in the aggregate investment rate which are consistent with the data. Moreover, growth rate shocks generate positive higher order autocorrelations in aggregate investment growth and positive cross-correlations at various leads and lags between investment rate growth and output growth. The calibrated model with level technology shocks that fails to generate these features of investment is shown for comparison. Finally, it is shown that the model with growth shocks is able to match the positive co-movement between output, investment, employment, and consumption if the intertemporal elasticity of substitution is sufficiently high.

Although the focus is not on asset pricing, it is interesting to note that as in the data, the Sharpe Ratio in the growth model is high (approximately 0.38). This happens through the long run risk channel as in Bansal and Yaron (2004), Kaltenbrunner and Lochstoer (2010), and Croce (2012). On the other hand, the Sharpe Ratio in the level model is tiny.

4.1. Calibration

The model is solved numerically at a quarterly frequency by value function iteration. The solution method is discussed in the appendix.

Preferences. The intertemporal elasticity of substitution ψ is 1.5 following Bansal and Yaron (2004). The risk curvature parameter γ is 4, which, because of leisure in the utility function, implies that risk aversion is $\gamma\phi = 1.4$ (Swanson, 2012). This is within the range of reasonable values for the coefficient of relative risk aversion, as suggested by Mehra and Prescott (1985). ϕ determines the elasticity between consumption and leisure, it is set as 0.35 to deliver a reasonable volatility of hours. The subjective discount factor β is set at 0.995 to match the level of risk-free rate.

Technology. The production technology parameters are standard in the macroeconomic literature. Capital share α is 0.36, the quarterly depreciation rate δ is 0.025; and the quarterly log technology growth rate g is 0.5%; these values are identical to Jermann (1998).[8]

Productivity shocks. The persistence ρ_G and conditional volatility σ_G of the *growth rate* of TFP are set to 0.24 and 0.0093, respectively. To get these numbers, Eq. (4) is estimated by GMM using quarterly TFP data (not purified) provided by John Fernald.[9] Details of the GMM procedure are in the appendix. These numbers also lead to reasonable autocorrelation and volatility of output in the growth shock model. For the level shock model, autocorrelation ρ_L and conditional volatility σ_L of the *level* of TFP are set to 0.96 and 0.0066, which is estimated by regressing TFP on its lag. These allow the level shock model to roughly match the volatility and autocorrelation of output.

[7] In the previous draft, circulated as “Micro frictions, asset pricing, and aggregate implications”, the model included heterogenous firms facing non-convex adjustment costs. There too, growth rate shocks alone could explain investment nonlinearities and micro frictions were not necessary.

[8] There are no capital adjustment costs in the model because in the current calibration investment is already less volatile than in the data. However, if quadratic capital adjustment costs are included, the results on investment rate heteroscedasticity and history dependence are very similar to the current case. Furthermore, an earlier version of the paper contained a model with inelastic labor supply. That version did require capital adjustment costs to match investment volatility. The heteroscedasticity and history dependence in that version were also similar to the current version.

[9] If the purified series is used instead, TFP is estimated to have a lower volatility ($\sigma_G = 0.56$), but a higher persistence ($\rho_G = 0.61$). In the model, both heteroscedasticity and history dependence are stronger when volatility and persistence of TFP are higher. Therefore, the behavior of these features of investment rate is fairly similar in models solved with parameters estimated from the purified or the unpurified series.

4.2. Conventional RBC moments

Table 2 presents the standard business cycle moments for the data, the baseline model and the level shock model.[10] Overall, the baseline growth rate shock model matches these aggregate moments reasonably well. For example, the growth rate shock model does well with the volatilities and autocorrelations of output and the investment rate. The volatilities of consumption, investment and hours are somewhat low, yet still reasonably close to the data. A common problem in the standard RBC models is that the autocorrelation of consumption growth is too high, while its correlation with output growth is too low. The level shock model does not do well with the volatilities of consumption and the investment rate, both of which are lower than the data.

The following sections discuss the conditional heteroscedasticity and history dependence of the aggregate investment rate, the higher order autocorrelations of investment rate growth, the correlations at various leads and lags between the investment rate growth and output growth, and the correlations between employment, consumption, and investment.

As discussed earlier, KT argue that in general equilibrium micro-frictions do not affect aggregate quantities while BCE argue that modeling micro-frictions in a particular way can have an effect on aggregate quantities, in particular they help to match the heteroscedasticity of the aggregate investment rate. We do not take a stand on whether these frictions matter but confirm (consistent with BCE) that while heteroscedasticity does not exist in a standard model with level shocks (stationary TFP), it arises naturally in a growth rate shock model where TFP is consistent with long run risk. Moreover, history dependence, the higher order autocorrelations and the cross-correlations are all consistent with a long run risk model but not with a level shock model.

4.3. Heteroscedasticity

Section 2.1 followed Bachmann et al. (2011) to compute the heteroscedasticity range for investment rate in the data and showed that innovations in investment rate are more volatile following investment rate being high. This exact procedure is repeated on simulated data from the model, these results are in Table 3. For the growth rate shocks model, the heteroscedasticity range is 0.21 and 0.16 for the two respective specifications compared to 0.16 and 0.22 in the data. However, a model identical to the baseline model but with stationary level shocks predicts that conditional volatility of the investment rate is near zero.[11] Moreover, additional measures of heteroscedasticity, such as $\log(\sigma_{90}/\sigma_{10})$ and $\log(\sigma_{Max}/\sigma_{Min})$, are also close to the data in the growth rate shock model. Overall, the growth rate shock model quantitatively captures the conditional heteroscedasticity in the aggregate investment rate but the model with level shocks does not.

The heteroscedasticity ranges for the output-to-capital and consumption-to-capital ratios have also been computed, these are in Panels B and C of Table 3. Unlike the investment rate, these are insignificant and the point estimates are smaller in magnitude. Although the point estimates in the baseline model have opposite sign compared to the data for these two quantities, they are also smaller in magnitude than for the investment rate, and are insignificant. In the model with level shocks these quantities are also insignificant.

Below it is argued that a preference for smoothing consumption over time, combined with permanent (or highly persistent) productivity shocks can lead to heteroscedasticity of the investment rate. For intuition, consider the problem with inelastic labor ($\phi=1$) and TFP being $X^G X^L$ where X^G and X^L are defined by Eqs. (4) and (5) respectively. Define $A_{t+1}=X^G_{t+1}/X^G_t$ to be the growth rate of the non-stationary component of TFP.

4.3.1. Analytic results

First, consider a special case with inelastic labor $\phi=1$, logarithmic preferences ($\theta=1/\psi=1$), and full depreciation ($\delta=1$).

Proposition 1. *If $\phi=1$, $\theta=1/\psi=1$, and $\delta=1$ then an analytic solution exists. The investment rate and its conditional volatility are given by*

$$\frac{I_{t+1}}{K_{t+1}}=\frac{X^L_{t+1}}{X^L_t}A_{t+1}\left(\frac{I_t}{K_t}\right)^{\alpha}$$
$$\sigma_t\left[\frac{I_{t+1}}{K_{t+1}}\right]=\sigma[e^{\sigma_L\epsilon^L_{t+1}}A_{t+1}]e^{g}(X^L_t)^{\rho_L-1}\left(\frac{I_t}{K_t}\right)^{\alpha} \quad (7)$$

Proof. See Appendix.

Thus, in this special case, high investment rate (low capital) amplifies shocks to the growth rate of productivity. In the case of growth shocks, even if the growth rate of technology (A_{t+1}) is i.i.d., high investment rate will forecast high volatility.

[10] All standard errors are bootstrapped. This includes HP-filtered quantities. Note that due to high persistence of HP-filtered quantities, standard errors may be underestimated, however these quantities are not the main focus of the paper.

[11] In addition to the heteroscedasticity range being small, the volatility of I/K is actually somewhat higher when past I/K is low in the level model. This is indicated by the sign of the coefficient η in Table 3. This is the reverse of the data and the growth model.

However, this will not necessarily be the case with level shocks because I_t/K_t is positively correlated with X_t^L, and therefore for $\rho^L<1$, the two will offset. That is, because the conditional expectation of X_{t+1}^L/X_t^L is negatively correlated with I_t/K_t, heteroscedasticity appears only if shocks are very persistent or permanent. In other words, with transitory shocks, a high investment rate signals both a high level of target capital and a low productivity growth rate.

Now, consider cases with more general utility and depreciation. First, consider what happens when there is no preference for smoothing ($\psi=\infty$) and all shocks are i.i.d.

Proposition 2. *If $\psi=\infty$, X_t^L is i.i.d., and A_t is i.i.d. then capital choice is independent of past capital stock, investment rate is $I_t/K_t=-(1-\delta)+A_{t+1}^{1/(1-\alpha)}$ and cannot exhibit heteroscedasticity.*

Proof. See Appendix.

This result is intuitive, if households do not care about smoothing consumption then their capital choice exactly follows the trend and does not depend on past capital.

4.3.2. Intuition

Even if there is no exogenously built in heteroscedasticity through the shocks (or through frictions as in Bachmann et al., 2011), household preferences for smoothing consumption can lead to heteroscedasticity. If households wish to smooth consumption over time, then capital will not immediately move to its long run level so that optimal choice of capital K_{t+1} will positively depend on K_t (when K_t is low relative to trend, investing all the way to trend would result in very low consumption today and very high next period). For simplicity consider a world with only non-stationary i.i.d. shocks. Suppose that next period's capital choice is a linear combination of the optimal choice of capital[12] $K_t^*=\overline{k}(X_t^G)^{1/(1-\alpha)}$ and of today's capital K_t: $K_{t+1}=\chi^*K_t^*+\chi K_t$. Although this is not the optimal solution to our problem, but rather a simple heuristic example, it is a first order approximation to the true solution. In this case it is easy to show[13] that the volatility of the investment rate will positively depend on the level of capital despite no heteroscedasticity in the shocks

$$\frac{I_{t+1}}{K_{t+1}}=\frac{A_{t+1}^{1/(1-\alpha)}}{1+\frac{\chi}{\chi^*}\frac{k_t}{\overline{k}}}-(1-\delta-\chi)$$
$$\sigma_t\left[\frac{I_{t+1}}{K_{t+1}}\right]=\frac{\sigma[A_{t+1}^{1/(1-\alpha)}]}{1+\frac{\chi}{\chi^*}\frac{k_t}{\overline{k}}} \tag{9}$$

where $k_t=K_t/(X_t^G)^{1/(1-\alpha)}$ is detrended capital. When $\chi=0$, just as in the previous paragraph, there is no preference for smoothing and the investment rate inherits the properties of $A_{t+1}^{1/(1-\alpha)}$, which does not depend on past variables. When $\chi>0$, low capital (high investment rate) at t amplifies $t+1$ permanent shocks.[14]

While level shocks could, in principle also lead to heteroscedasticity in the investment rate, the same logic can be applied to see why under most parameterizations level shocks will not lead to heteroscedasticity. Consider a world with level shocks only and suppose that $\log(X^L)$ is AR(1) with persistence ρ^L so that $X_{t+1}^L/X_t^L=(X_t^L)^{\rho^L-1}e^{\sigma^L\epsilon_{t+1}^L}$. Despite level shocks being temporary, if they are persistent they will also lead to an increase in capital so that the optimal choice of capital will still be (approximately) a linear combination of current capital and productivity: $K_{t+1}=\chi^*\overline{k}X_t^L+\chi K_t$. Just as before, the investment rate can be expressed as a function of the shock and the state variables

$$\frac{I_{t+1}}{K_{t+1}}=\frac{X_{t+1}^L/X_t^L}{1+\frac{\chi}{\chi^*}\frac{k_t}{\overline{k}}}-(1-\delta-\chi)=e^{\sigma^L\epsilon_{t+1}^L}\frac{(X_t^L)^{\rho^L-1}}{1+\frac{\chi}{\chi^*}\frac{k_t}{\overline{k}}}-(1-\delta-\chi)$$
$$\sigma_t\left[\frac{I_{t+1}}{K_{t+1}}\right]=\sigma[e^{\sigma^L\epsilon^L}]\frac{(X_t^L)^{\rho^L-1}}{1+\frac{\chi}{\chi^*}\frac{k_t}{\overline{k}}} \tag{10}$$

The intuition here is similar to the logarithmic utility case above. Holding X_t^L constant, low k_t (high investment rate) still amplifies the shock. However, unconditionally, k_t is positively correlated with X_t^L, therefore, since $\rho^L-1<0$, high X_t^L dampens the shock. The net effect depends on the persistence X_t^L, and certainly, as $\rho^L\to1$ this model becomes a permanent shock

[12] Here $\overline{k}$ is the average level of detrended capital and determines the optimal capital to productivity ratio. Since there is a balanced growth path, $\overline{k}$ exists.

[13] The derivation is:

$$\frac{I_{t+1}}{K_{t+1}}=\frac{K_{t+2}-(1-\delta)K_{t+1}}{K_{t+1}}=\chi^*\frac{\overline{k}(X_{t+1}^G)^{1/(1-\alpha)}}{K_{t+1}}-(1-\delta-\chi)$$
$$=\chi^*\frac{\overline{k}(X_{t+1}^G)^{1/(1-\alpha)}}{\chi^*\overline{k}(X_t^G)^{1/(1-\alpha)}+\chi K_t}-(1-\delta-\chi)=\frac{A_{t+1}^{1/(1-\alpha)}}{1+\frac{\chi}{\chi^*}\frac{k_t}{\overline{k}}}-(1-\delta-\chi) \tag{8}$$

[14] Note that the above logic suggests that while investment rate is heteroscedastic, its logarithm is not. Indeed, in both the model and the data, the heteroscedasticity range for $\log(I/K)$ is close to zero and insignificant. For example, if $\log(\sigma_{95}/\sigma_5)$ is recomputed for the logarithm of investment rate, instead of for investment rate, then the heteroscedasticity range would be −0.10 for the data and −0.07 for the model (specification one), or −0.04 for the data and −0.08 for the model (specification two).

Table 7
Heteroscedasticity mechanism. For this table, the model (without labor) is solved for several utility specifications and productivity shock specifications. The productivity shock X_t always takes the form $\log(X_{t+1}) = \rho^L \log(X_t) + \sigma^L \epsilon_{t+1}$. For each ρ^L, σ^L is selected so as to keep $\log(X_{t+1}) - \log(X_t)$ constant. This implies that as ρ^L rises, this process approaches a random walk with infinite unconditional volatility. The utility specifications are (1) no preference for smoothing and risk neutrality ($\psi = \infty$, $\theta = 0$); (2) no preference for smoothing and risk aversion ($\psi = \infty$, $\theta = 4$); (3), (4), (5) preference for smoothing and risk neutrality ($\psi < \infty$, $\theta = 0$); (6) preference for smoothing and risk aversion ($\psi = 0.5$, $\theta = 4$); and (7) logarithmic utility with full depreciation. Note that (1) and (7) are solved analytically, the others are solved numerically.

ρ^L	σ^L $\times 100$	$\psi = \infty$ $\theta = 0$ 1	$\psi = \infty$ $\theta = 4$ 2	$\psi = 2.0$ $\theta = 0$ 3	$\psi = 1.5$ $\theta = 0$ 4	$\psi = 0.5$ $\theta = 0$ 5	$\psi = 0.5$ $\theta = 4$ 6	log util $\delta = 1$ 7
0	0.48	0	0.00	−0.02	−0.02	−0.02	−0.01	−0.01
0.500	0.58	0	0.05	−0.01	−0.01	−0.03	0.03	−0.01
0.750	0.63	0	0.01	0.02	0.00	0.07	0.06	0.00
0.960	0.66	0	0.02	0.03	0.05	0.09	0.12	0.01
0.990	0.67	0	−0.02	0.07	0.14	0.19	0.18	0.01
0.999	0.67	0	0.00	0.18	0.28	0.29	0.30	0.02
1.000	0.67	0	0.00	0.19	0.29	0.30	0.31	0.02

model where the investment rate is heteroscedastic. However, in the numerical exercise, for estimated values of ρ^L, the investment rate looks very close to homoscedastic.

4.3.3. Numerical results

While the above heuristic approach is useful for intuition, in general optimal capital choice will not exactly be equal to a linear combination with constant weights. For additional intuition, Table 7 presents numerical results highlighting the smoothing mechanism. The model above is solved for stationary TFP of differing persistence. In particular, let $\log(X_{t+1}^L) = \rho^L \log(X_t^L) + \sigma^L \epsilon_{t+1}$. For each ρ^L, σ^L is chosen such that $\log(X_{t+1}) - \log(X_t) = 0.0067$ (this is equal to the estimated volatility of the process). Note that as ρ^L gets larger, the unconditional volatility of X^L grows; as $\rho^L \to \infty$, this process approaches a non-stationary process. In the first (no smoothing, risk neutrality) and last (log utility and full depreciation) columns are the two cases with analytic solutions. As discussed above, when there is no preference for smoothing, the heteroscedasticity range is zero, while with log utility it grows with persistence. The case with no preference for smoothing but positive risk aversion looks very similar to the case with no preference for smoothing and risk neutrality. Cases with risk neutrality and progressively stronger preference for smoothing (lower ψ) are in columns three through five. In all of these cases the heteroscedasticity range grows with persistence, and this effect is stronger when the smoothing preference is stronger, that is when the IES is smaller.

Note that it is the preference for smoothing and not risk aversion that matters for heteroscedasticity. Epstein and Zin (1989) preferences allow for a separation of risk aversion and the IES. The sixth column presents a case with the same smoothing preference ($\psi = 0.5$) as the fifth column, but with higher risk aversion. The heteroscedasticity range is similar in both cases.

4.4. History dependence

Section 2.2 showed that longer expansions are associated with larger rises in the investment rate while longer recessions are associated with larger falls. This was summarized by regressing the rise in investment rate on the length of the associated recession or expansion with the slope being positive for expansions (b^E) and negative for recessions (b^R). These results for the model and data are in Table 4.

In the baseline model with growth rate shocks, as in the data, the investment rate rises through an expansion, and falls through a recession: $b^E = 0.10$, and $b^R = -0.19$ in the model compared to 0.10 and −0.27 in the data. Moreover, both b^E and b^R are statistically significant. On the other hand, a model identical to the baseline model, but with stationary level shocks has a much weaker pattern: the investment rate does not rise or fall much as the expansion or recession get longer: $b^E = 0.03$ and $b^R = -0.11$, which are both about 1/3 of their data counterparts. The level shock model fails to quantitatively capture the history dependence in the aggregate investment rate. The results also look similar when recessions are defined as two consecutive quarters of negative growth.

History dependence is also computed for the output-to-capital and consumption-to-capital ratios, these too, are in Table 4. As in the data, in the growth rate model, these ratios follow a similar pattern as investment-to-capital: larger rises for longer recessions and larger falls for longer recessions. However, the level shock model is not able to match these patterns.[15]

[15] The magnitude of coefficients in the data is much bigger than in the model. This happens because as in Bachmann et al. (2011), capital in the data is constructed from gross private nonresidential investment rather than total investment resulting in average C/K and Y/K being bigger in the data than in the model. If instead, capital from the fixed asset tables is used, then average C/K and Y/K are similar to the data, as are the magnitudes of the coefficients.

4.4.1. Intuition

As with heteroscedasticity, history dependence is naturally implied by growth rate shocks. Fig. 3 plots impulse responses of the investment rate to positive shocks in the random variable governing productivity for a short and a long expansion. In particular, let $i \in (L, G)$ indicate the model with level shocks or growth shocks. The model is simulated for 500 periods setting $\epsilon_t^i = 0, t = -499, \ldots 0$ in either Eq. (5) (level shocks) or (4) (growth shocks). For a short expansion, $\epsilon_t^i = \sigma^i$, $t = 1$ and $\epsilon_t^i = 0, t > 1$. For a long expansion, $\epsilon_t^i = \sigma^i, t = 1, \ldots 5$ and $\epsilon_t^i = 0, t > 5$. Note that additional positive shocks extend the length of a recession, however because of the autoregressive nature of productivity (level model) and productivity growth (growth model), their marginal effect on the level or growth rate of productivity decreases as the expansion gets longer.[16]

In a model with shocks to the level of TFP (bottom panel in Fig. 3), following one positive shock, the investment rate rises upon impact and then immediately falls in the second quarter. For a longer expansion of five consecutive positive shocks, the investment rate rises for three consecutive quarters (less than the length of an expansion) then starts to fall in the fourth quarter. This is because each level of productivity is associated with an optimal level of capital. A temporary one quarter rise in productivity causes the firm to increase its target capital level and increase investment immediately; but when productivity level mean-reverts towards the long-run average in the second quarter, the investment rate falls immediately. For a longer expansion, despite additional positive shocks, the autoregressive forces begin to dominate even before the end of the expansion; this results in the investment rate falling even before the expansion ends (note that although falling, the investment rate remains above average).

The impulse response of investment in the growth rate shock model is quite different (top panel). Following one positive shock to the growth rate of productivity, the investment rate rises upon impact and then continues to rise for an additional quarter despite no positive shock in the second quarter. For a longer expansion of five consecutive positive shocks, the investment rate increases consecutively through quarter six despite positive shocks stopping in quarter five. Compare this to the level shock model, where the investment rate began to fall after quarter three despite five consecutive positive shocks. The reason for this is straight forward: growth rate shocks are permanent shocks to the level of productivity. A single quarter positive shock increases optimal capital dramatically and causes the firm to invest up to two quarters to reach the optimal capital level. If the growth rate of productivity is high for a longer number of periods, the level of productivity, and therefore the optimal level of capital continues to grow. Therefore an unexpected lengthening of an expansion leads the firm to ramp up investment rather than slow it down.

Thus, in the baseline model, shocks to the growth rate of TFP imply that aggregate investment is history dependent while shocks to the level of TFP leads to a much weaker effect, as is shown in Table 4.

4.5. Higher order autocorrelations and cross-correlations

In Section 2.3 and Table 5 it was shown that investment rate, investment rate growth, output growth, and consumption growth are all fairly persistent in the data. Like in the data, both models have investment rates that are highly persistent. However the model with level shocks behaves very differently from the data once growth rates are considered. Investment rate growth and output growth are each negatively autocorrelated at all lags while consumption growth has small positive autocorrelation. On the contrary, the growth rate shock model looks very much like the data. Investment rate growth is persistent with declining autocorrelations at higher orders; output and consumption growth have a similar positive and declining pattern with consumption growth being somewhat more persistent at higher orders.

Panel B of Table 6 reports the correlations at three leads and lags between investment rate growth and output growth. As in the data, in the growth shocks model investment rate growth is positively correlated with output at various leads and lags. However, the level shock model generates *zero* correlation between investment rate growth and output growth at all lags and even predicts a negative correlation at different leads, exactly opposite to the data.

The differences in the persistence of investment rate growth, consumption growth and output growth between the growth rate shock model and the level shock model are directly implied by the difference between the two productivity shocks. Growth rate productivity X_{t+1}^G is a first-order difference stationary process while level shock productivity X_{t+1}^L is a trend stationary process, i.e., growth rate shocks (ϵ_{t+1}^G) make the growth rate of TFP $\log(X_{t+1}^G/X_t^G)$ persistent. Level shocks (ϵ_{t+1}^L) actually result in a negative autocorrelation in the growth rate; this can be seen by taking first differences of Eq. (5)

$$\log X_{t+1}^L - \log X_t^L = g + (\rho_L - 1) * \log X_t^L + \sigma_L \epsilon_{t+1}^L, \tag{11}$$

so that if X_t^L is high due to past positive growth rates, future growth rate is expected to be low. Numerical simulation also shows that this process leads to negative correlation in the growth rate.

Despite households' attempts to smooth consumption over time, the growth rates of macro aggregates are closely linked to the growth rates of the underlying shocks. Thus, the growth rates of investment, consumption, and output are all positively autocorrelated in the growth model whereas the growth rates of investment and output are weakly negatively autocorrelated in the level model (consumption growth rate has weak positive autocorrelation due to smoothing).

[16] Let x be either the level (level model) or the growth rate (growth model) of TFP; x is the relevant state variable in each model. For one positive shock, x is given by the following: $x_0 = 0$, $x_1 = \sigma^i$, $x_2 = \rho_i \sigma^i$, $x_3 = \rho_i^2 \sigma^i$, etc. For five consecutive shocks, x is given by the following: $x_0 = 0, x_1 = \sigma^i$, $x_2 = \sigma^i + \rho_i \sigma^i$, $x_3 = \sigma^i + \rho_i \sigma^i + \rho_i^2 \sigma^i$, $x_4 = \sigma^i + \rho_i \sigma^i + \rho_i^2 \sigma^i + \rho_i^3 \sigma^i$, $x_5 = \sigma^i + \rho_i \sigma^i + \rho_i^2 \sigma^i + \rho_i^3 \sigma^i + \rho_i^4 \sigma^i$.

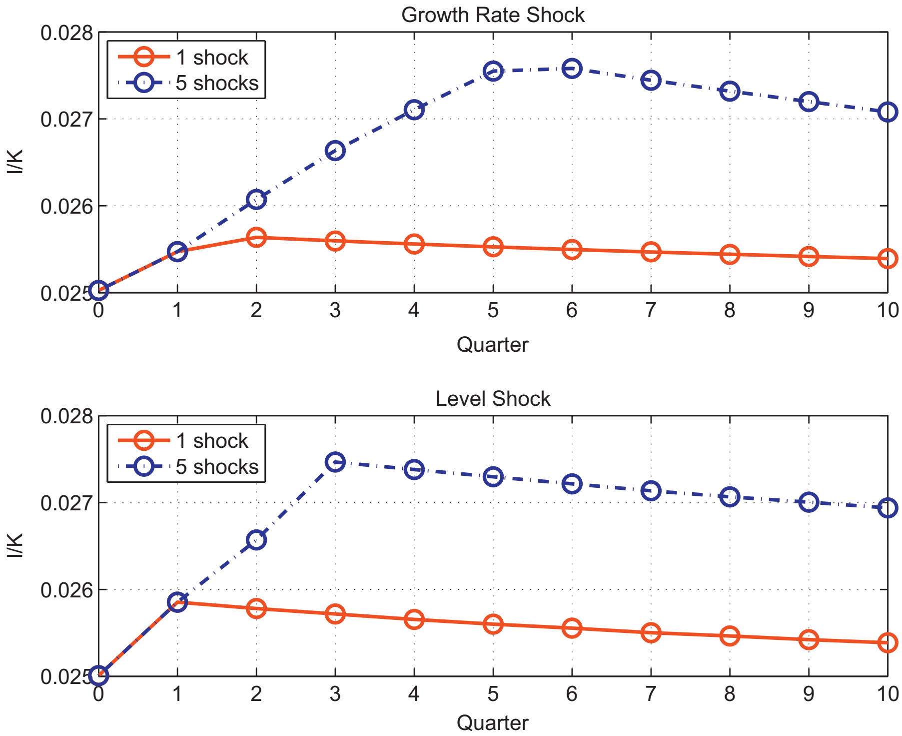

Fig. 3. Impulse responses of investment rate. This figure plots the impulse responses of investment rate (I/K) to 1 positive TFP shock (solid line) or 5 consecutive positive TFP shocks (dashed line). The top panel presents results from the growth shock model, where shocks are to the growth rate of TFP according to Eq. (4). The top panel presents results from the level shock model, where shocks are to the level rate of TFP according to Eq. (5).

Similarly, in the growth rate shock model, because output and investment cointegrate with TFP (which is non-stationary), the growth rates of these two series are positively correlated at various leads and lags. However, in the level shock model, output growth and investment rate growth do not co-move at any leads because there is no long term non-stationary trend other than deterministic growth.

4.6. The IES, hours, and consumption

This paper primarily focuses on the behavior of investment in the presence of alternative types of risk, however it is also interesting to explore the behavior of other quantities, such as labor hours and consumption. Up to now we have argued for the importance of growth (permanent) shocks to match various investment behavior. It was also shown that to match heteroscedasticity in investment rate, the intertemporal elasticity of substitution should not be too high, that is, there should be some preference for smoothing consumption over time. In this section it is shown that to match the behavior of employment in a model with growth shocks the IES should not be too low, that is the preference for smoothing should not be too strong. Together, the two imply rough upper and lower bounds for the IES.[17]

Panel A of Table 6 reports the correlations of output, consumption, investment, and hours for the data, compared to the model, for different levels of intertemporal elasticity of substitution. In the data all four quantities exhibit fairly high correlations. Both models are able to replicate this, however the growth shock model requires that the IES not be too low. Note that a low IES implies a strong preference for smoothing over time. When output is expected to be high in the future, low IES agents wish to raise consumption and leisure today; this is the wealth effect. Conversely, for high IES agents the preference for smoothing is tempered by the possibility of efficiently investing today to consume more in the future; this is the substitution effect.

After a permanent positive shock to TFP the long-run expected level of capital rises and today's capital is low relative to the long run. High IES households work more hours, to take advantage of higher productivity. Output rises immediately since both hours and productivity rose, allowing households to consume more. Over time, capital rises which further increasing output and allows consumption to continue rising. Eventually, as capital and output reach their long-run level, hours (which are stationary) fall back towards the long run mean while consumption stays permanently raised. Thus,

[17] This statement is, of course, conditional on the limitations of the model. In particular, with additional parameters, it is possible to generalize both the trade-off between consumption and leisure in the utility function, and the trade-off between capital and labor in the production function. The extra parameters may allow the model to match investment and employment behavior without restrictions on the IES. This question is left for future research.

following a positive productivity shock in a high IES world, output, investment, consumption, and hours all rise together. If shocks are also persistent ($\rho^G > 0$), then TFP is expected to further grow in the future implying that today's TFP is low relative to the long-run; this can reinforce this effect.

Now consider a permanent positive shock to TFP if the IES is low. In expectation of high future output consumption rises; hours do not rise as much as in the high IES case and may even fall if the IES is low enough as the preference for smoothing (wealth effect) dominates. Since TFP rises immediately, total output and investment may still rise despite the fall in hours, but investment rises much slower than in the high IES case. With IES=0.1, the growth shock model produces negative correlations between hours and output, hours and consumption, and investment and consumption (second column of Panel A in Table 6). Even with IES=0.25 (which corresponds to CRRA utility since $\theta = 1/\psi$) the correlation between consumption and investment is −0.33. Although the exact lower bound for the IES would depend on a particular parametrization, these results suggest that the IES cannot be too low if growth shocks are the primary drivers of business cycles.

Although we believe that growth shocks are the empirically relevant case to consider, note that with level shocks higher IES (as opposed to lower IES in the growth shocks model) leads to lower correlations. This is because a positive shock does not affect the long run values of TFP or capital. Immediately following a positive shock, capital is still equal to the long-run value while TFP is now above the long-run value causing consumption to rise. If the shock is persistent, over the next several periods households will invest more to take advantage of higher productivity causing capital to rise. However now that capital and output are above their long-run means, households will work fewer hours. The higher the IES, the more sensitive hours are to capital; this leads to lower correlations.

4.6.1. Relationship to news shocks

It is interesting to contrast these findings to the literature on news shocks where the positive co-movement between output, investment, hours, and consumption observed in the data is difficult to replicate in a model (Beaudry and Portier, 2007; Jaimovich and Rebelo, 2007; Lorenzoni, 2011). News shocks are shocks that affect expectations of future productivity without affecting today's productivity. Since after a positive news shock, current capital and productivity are unaffected, it is impossible for both consumption and investment to increase unless hours increase too. However, if consumption increases due to the wealth effect, the same wealth effect increases demand for leisure so hours should fall. Thus, in standard models with a low IES, the wealth effect dominates and though consumption rises after a positive shock, hours and investment fall. If, on the other hand, the IES is high, the substitution effect dominates which leads to hours and investment rising but consumption falling.

The co-movement problem does not arise for two reasons. First, when the IES is high enough, the substitution effect dominates and it is optimal in the model to raise investment and output in response to a positive shock. Second, the growth shocks are not pure news shocks. The process for X^G implies that after a positive shock there is (i) an instantaneous rise in productivity, and (ii) a rise in expectations about future productivity (if $\rho^G > 0$); a news shock is (ii) only. For this reason this economy is able to produce more output immediately after a positive shock even if there were no rise in employment allowing for a rise in consumption.

5. Conclusion

Shocks to the growth rate of TFP, which make TFP non-stationary, are the type of shocks needed to improve a model's asset pricing performance through the long run risk channel (Bansal and Yaron, 2004. We explore the implications of such shocks for aggregate investment behavior by solving a general equilibrium production economy. In addition to the well known asset pricing implications of such shocks, these shocks also improve the model's ability to explain the behavior of aggregate investment. In particular, these shocks can generate the nonlinearities in aggregate investment, including conditional heteroscedasticity and history-dependence, emphasized in the literature (e.g., Caballero and Engel, 1999; Bachmann et al., 2011). Previous literature, e.g., Bachmann et al. (2011) has been able to match heteroscedasticity only by using non-convex frictions. The model provides an alternative explanation for nonlinearities in aggregate investment. In addition, it is shown that to match the joint cyclical behavior of output, employment, consumption, and investment in a model with shocks to the growth rate of TFP, the IES cannot be too low.

Acknowledgments

We would like to thank Sudipto Bhattacharya, Frederico Belo, Bob Goldstein, Vito Gala (EFA discussant), Francois Gourio, Ravi Jagannathan, Lars Lochsoter (UBC Winter Finance Conference discussant), Stavros Panageas, Dimitris Papanikolaou (NBER discussant), Monika Piazzesi, Christopher Polk, Stijn Van Nieuwerburgh, Dimitri Vayanos, Neng Wang, Amir Yaron, and Lu Zhang for helpful comments. We thank the seminar participants at the Higher School of Economics, London Business School, London School of Economics, Ohio State University, University of Minnesota, University of Nottingham, EFA 2011 meetings, NBER 2010 Asset Pricing Summer Institute, SED 2011 meetings, and UBC 2011 Winter Finance Conference. This work supersedes our working paper previously circulated as “Micro Frictions, Asset Pricing, and Aggregate Implications”. All remaining errors are our own.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at http://dx.doi.org.10.1016/j.jmoneco.2013.05.002.

References

Alvarez, F., Jermann, U., 2005. Using asset prices to measure the persistence of the marginal utility of wealth. Econometrica 73 (6), 1977–2016.
Attanasio, O., Vissing-Jorgensen, A., 2003. Stock-market participation, intertemporal substitution, and risk-aversion. American Economic Review 93 (2), 383–391.
Attanasio, O., Weber, G., 1989. Intertemporal substitution, risk aversion and the Euler equation for consumption. Economic Journal 99, 59–73.
Bachmann, R., Caballero, R., Engel, E., 2011. Lumpy investment in dynamic general equilibrium. Working Paper.
Bansal, R., Khatchatrian, V., Yaron, A., 2005. Interpretable asset markets. European Economic Review 49 (3), 531–560.
Bansal, R., Kiku, D., Yaron, A., 2008. Risks for the long run: estimation and inference. Working Paper.
Bansal, R., Shaliastovich, I. A long-run risks explanation of predictability puzzles in bond and currency markets. Review of Financial Studies, forthcoming.
Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. Journal of Finance 59, 1481–1509.
Beaudry, P., Portier, F., 2007. When can changes in expectations cause business cycle fluctuations in neo-classical settings? Journal of Economic Theory 135 (1), 458–477.
Bry, G., Boschan, C., 1971. Cyclical Analysis of Time Series: Procedures and Computer Programs. NBER.
Caballero, R., Engel, E., 1999. Explaining investment dynamics in us manufacturing: a generalized (s,s) approach. Econometrica 67 (4), 741–782.
Caballero, R., Engel, E., Haltiwanger, J., 1995. Plant-level adjustment and aggregate investment dynamics. Brookings Papers on Economic Activity 2, 1–54.
Campbell, J.Y., 1999. Asset prices, consumption and the business cycle. Handbook of Macroeconomics.
Campbell, J.Y., Mankiw, G., 1989. Consumption, income, and interest rates: reinterpreting the time series evidence. NBER Macroeconomic Annual, pp. 185–216.
Cooper, R., Haltiwanger, J., Power, L., 1999. Machine replacement and the business cycle: lumps and bumps. American Economic Review 89 (4), 921–946.
Croce, M., 2012. Long-run productivity risk: a new hope for production-based asset pricing? Working Paper.
Epstein, L., Zin, S., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. Econometrica 57, 937–969.
Hall, R., 1988. Intertemporal substitution in consumption. Journal of Political Economy 96 (2), 339–357.
Hansen, L., Singleton, K., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica 50 (5), 1269–1286.
Jaimovich, N., Rebelo, N., 2007. Can news about the future drive business cycles. American Economic Review 99 (4), 1097–1118.
Jermann, U.J., 1998. Asset pricing in production economies. Journal of Monetary Economics 41, 257–275.
Kaltenbrunner, G., Lochstoer, L., 2010. Long-run risk through consumption smoothing. Review of Financial Studies 23, 3141–3189.
Khan, A., Thomas, J., 2008. Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. Econometrica 76 (2), 395–436.
Lorenzoni, G., 2011. News and aggregate demand shocks. Annual Review of Economics 3, 537–557.
Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15 (2), 145–161.
Prescott, E.C., 1986. Theory ahead of business cycle measurement. Federal Reserve Bank of Minneapolis Staff Report (102).
Swanson, E., 2012. Risk aversion and the labor margin in dynamic equilibrium models. American Economic Review 102 (4), 1663–1691.
Thomas, J., 2002. Is lumpy investment relevant for the business cycle? Journal of Political Economy 110 (3), 508–534.