Data consistent modelling of medium-frequency cycles and their origins.

Tom Holden[1], School of Economics, University of Surrey

Abstract: *This paper presents four stylized facts on medium-frequency cycles, then builds and estimates a model capable of replicating both these facts and standard business-cycle ones. We show that GDP returns to trend at long lags, that aggregate mark-ups always lead output, and are only counter-cyclical at low frequencies, and that medium-frequency cycles are larger in countries with longer patent protection. Since traditional dynamic endogenous growth models generate large trend-breaks following business-cycle shocks, our model is based on that of Holden (2013a). After estimation, a financial-type shock to the stock of ideas emerges as the key driver of the medium-frequency cycle.*

Keywords: *medium frequency cycles, patent protection, mark-ups, unit roots*

JEL Classification: *E32, E37, L16, O31, O33, O34*

[1] Contact address: School of Economics, FBEL, University of Surrey, Guildford, GU2 7XH, England
Contact email: thomas.holden@gmail.comWeb-site: http://www.tholden.org/
The author would particularly like to thank his supervisor, Simon Wren-Lewis. Further thanks are due to Guido Ascari, Eric Bartelsman, Paul Beaudry, Pedro Bento, George Bitsakakis, Beatriz de Blas, Cristiano Cantore, Diego Comin, Ana-Maria Dumitru, Martin Ellison, Fabio Ghironi, Bart Hobijn, Tim Kehoe, Tobias Kitlinski, Paul Levine, Omar Licandro, Ioana Moldovan, Peter Neary, Joaquim Oliveira Martins, Christopher Nekarda, Walter Park, Pietro Peretto, James Reade, Torsten Schmidt, Diego Restuccia, Nicole Tabasso, David Vines, Klaus Wälde and the participants in numerous seminars at which versions of this paper were presented. Financial support provided to the author by the ESRC is also greatly appreciated. The author gratefully acknowledges the use of the Oxford e-Research Centre's Windows Compute Cluster and the University of Surrey FBEL cluster for the numerical computations performed in this paper, and the use of data from the Nuffield College Data Library.

In this paper, we take steps towards building a fully data consistent model of both business cycles and cycles at medium frequencies. We begin by presenting four new stylised facts on medium frequency business cycles, contributing to the empirical literature began by Comin and Gertler (2006). All point to the importance of the mechanisms that we present in our theoretical companion paper (Holden 2013a). We then embed the model of that paper into a medium-scale framework, producing a fully featured model we may estimate. Doing so, reveals that a single financial-type shock explains much of the variance of all key variables, thanks to the strength of the transmission mechanism embedded in the original theoretical model.

In section 1, firstly, we present new evidence that real GDP per capita returns to trend at long lags. This low frequency return to trend is large enough that we may come close to rejecting the null of a unit root. Since traditional models of endogenous growth and business cycles generate large trend breaks following standard business cycle shocks, this provides strong evidence in favour of our model (Holden 2013a) in which firm dynamics absorb such shocks giving an endogenously near constant frontier growth rate.

Secondly, we show that Nekarda and Ramey's (2010) evidence on the cross-correlation of mark-ups (as measured by inverse labour shares) and GDP continues to hold at medium frequencies, with mark-ups always leading output, and only becoming counter-cyclical if very low frequency components are included. This supports the model of our companion paper (Holden 2013a) which generates exactly this cross-correlation, in contrast to the existing literature (e.g. Comin and Gertler 2006) that generates counter-cyclical mark-ups even at high to medium frequencies.

Finally, we show that in countries with longer patent protection, medium frequency cycles account for a larger fraction of output's variance. This points to a key role for patent protection in generating medium frequency cycles, in line with our theoretical model (Holden 2013a).

To demonstrate the power of our theoretical model in capturing these facts, in section 2 we produce a medium-scale version of it, augmented with additional shocks, physical and research capital stocks, habits, assorted adjustment costs, nominal wage rigidity and augmented Taylor rule monetary policy, which we then estimate in section 3. Our estimation methodology is a development of that of Canova (2009), and avoids contorting the (inevitably misspecified) model to fit the data. By using maximum a posteori estimation with flat priors on almost all parameters, we come close to the agnosticism of classical maximum likelihood methods, while still recovering reasonable parameter estimates. The lack of prior on shock variances also frees the estimation to switch off most of the shocks in the model, resulting in a model in which just a few shocks are able to account for much of the data, at both business cycle and medium frequencies.

1. Empirics

1.1. The near trend stationarity of output

We begin by presenting evidence that real GDP per capita returns to trend at long lags. Since statistical tests on regressions with large numbers of lags tend to suffer from a lack of power, we have to find a sparsely parameterised way of capturing this long-run behaviour. It seems implausible that a high-frequency spike in GDP should lead to another spike in GDP many periods later. Instead, if GDP responds at all to its own past fluctuations at long lags, it will only respond to the low frequency (i.e. smoothed) fluctuations. We would like to smooth the data then at a range of frequencies, and regress output on the lags of these smoothed series. It will also help the interpretability of results if each lag of the data affects at most one of these smoothed series, which suggests taking moving averages. We choose then to regress log US quarterly GDP per-capita on a linear trend, the first lag of its one period moving average (i.e. its first lag), the second lag of its two period moving average, the fourth lag of its four period moving average, and so on up to the 32^{nd} lag of its 32 period moving average.

I.e. we run the regression:

$$\begin{aligned} y_t = \mu + \delta t + \phi_1 y_{t-1} + \phi_2 \frac{1}{2}(y_{t-2} + y_{t-3}) \\ + \phi_3 \frac{1}{4}(y_{t-4} + y_{t-5} + y_{t-6} + y_{t-7}) + \cdots \\ + \phi_6 \frac{1}{32}(y_{t-32} + \cdots + y_{t-63}) + \varepsilon_t. \end{aligned} \tag{1.1}$$

The full results of this regression are given in Table 1. The key facts to note here though are that ϕ_2, ϕ_3, ..., ϕ_6 are all negative, and that ϕ_6 is comfortably significant at 5%, suggesting that GDP is indeed returning towards trend at long lags. ϕ_6 corresponds to a period of eight to sixteen years, which includes the principal band of medium-frequency cycles, as is shown in Figure 2.

Variable	Coefficient	Std. Error	t-value	t-prob.	Part R^2
μ	-1.20281	0.3603	-3.34	**0.0010**	0.0574
δ	0.000572088	0.0001751	3.27	**0.0013**	0.0551
ϕ_1	1.21142	0.06323	19.2	**0.0000**	0.6673
ϕ_2	-0.251229	0.08649	-2.90	**0.0041**	0.0441
ϕ_3	-0.0272064	0.05389	-0.505	*0.6143*	0.0014
ϕ_4	-0.00266296	0.03332	-0.0799	*0.9364*	0.0000
ϕ_5	-0.0139299	0.02365	-0.589	*0.5566*	0.0019
ϕ_6	-0.0531785	0.02489	-2.14	**0.0339**	0.0243

Table 1: Results of the regression (1.1).
Run on log US quarterly real GDP (from NIPA) over X12 seasonally adjusted civilian non-institutional population (CNP16OV from FRED). 1948:1-2011:2.

We would like to know whether the magnitude of ϕ_6 is sufficient to pull GDP completely back to trend, or equivalently, whether log-GDP has a unit root. We can test for this if we transform (1.1) into Augmented Dickey-Fuller (ADF) form (Said and Dickey 1984), giving:

$$\begin{aligned} \Delta y_t = \mu + \delta t + \left[\sum_{i=1}^{6} \phi_i - 1\right] y_{t-1} - \phi_2 \frac{1}{2}(2\Delta y_{t-1} + \Delta y_{t-2}) - \cdots \\ - \phi_6 \frac{1}{32}(\dots). \end{aligned} \tag{1.2}$$

Since this is an equivalent model, no parameter estimates or standard errors change. However, we can now use the t-value on the y_{t-1} coefficient (-3.36) to perform an ADF test. Our Monte-Carlo experiments[2] indicate that there is only

[2] With 2^{20} replications, where in each case the regression (1.2) was run on the second half of a sample from a unit variance random walk, started at zero and twice the length of our data sample. This is broadly the

an 11.1% chance we would observe a result as extreme as this if the true data generating process were a random walk.[3] We do not wish to claim because of this that GDP is unambiguously trend-stationary. However, it does suggest that the size of the unit root in US GDP is (at most) very small, reinforcing the findings of Cochrane (1988).

1.2. Mark-ups

Nekarda and Ramey (2010) found that mark-ups were pro-cyclical both when the data was filtered with a standard ($\lambda = 1600$) HP-filter, and when it was filtered by taking first differences. However, Comin and Gertler (2006) report that mark-ups are counter-cyclical when the data is filtered via a band pass filter that keeps cycles of periods from one to fifty years.[4] Given that Comin and Gertler find that the medium-frequency variance of output is concentrated on cycles taking around ten years, the natural question is whether the counter-cyclicality of mark-ups they observe is a consequence of behaviour around these frequencies, or whether it is driven by counter-cyclicality at lower frequencies. Nekarda and Ramey (2010) also found that at business cycle frequencies, mark-ups were strongly correlated with future output, and negatively correlated with past output. Again, we would like to know if this still holds at plausible medium frequencies. The plot in Figure 1 below answers both of these questions.

Each vertical slice of this plot shows the cross-correlation[5] of quarterly log output and log mark-ups[6] when both are filtered by a high pass filter[7] with a cut-off given by the x-axis's value. (Shaded areas indicate positive correlations, with

methodology used by Cheung and Lai (1995) in their study of the finite sample properties of the ADF test with varying lag-order.

[3] Standard asymptotic critical values suggest a p-value close to 5%, but given the large number of lags and fairly small sample, it is unsurprising these are inaccurate.

[4] Using annual data, they also find that mark-ups are counter-cyclical at business cycle frequencies, though less so than at medium ones; however, their measure of the mark-up relies on many more questionable assumptions about utility and production functions than the Nekarda and Ramey one does. Additionally, Nekarda and Ramey find that the use of annual data always biases observed correlations towards counter-cyclicality.

[5] Fractional lags are evaluated via linear interpolation.

[6] Mark-ups are measured by the inverse labour share (following Nekarda and Ramey (2010)). Data is from NIPA, 1947:Q1-2011Q2.

[7] Implemented by setting the lower cut-off of a Christiano and Fitzgerald (2003) band-pass filter to two quarters.

the darker area being significantly different from zero at 5%. The cross-hatched area is negative but insignificantly different from zero at 5%.) We see immediately that Nekarda and Ramey's finding that mark-ups are positively correlated with future output and negatively correlated with past output holds particularly strongly at medium frequencies.[8]

Additionally, tracing along the lead=0 line we see that mark-ups are pro-cyclical when the data is filtered by a high-pass filter with a cut-off less than 16.5 years, suggesting that the Comin and Gertler's medium-frequency counter-cyclicality result was indeed driven by behaviour below the main frequencies of medium-frequency cycles. Indeed, from the spectral decomposition[9] of output growth shown in Figure 2, we see that mark-ups are significantly pro-cyclical when filtered at any frequency corresponding to a peak in the spectral decomposition, including the medium-frequency peak at twelve years. This establishes that the relevant medium-frequency cycles feature pro-cyclical movements in mark-ups.

[8] Given these two apparent directions of causation, it is tempting to try to explain away one as merely a product of the cyclical behaviour of the variables concerned. (Recessions inevitably follow expansions, which inevitably follow recessions, etc.) Along these lines, Tutino and Cheremukhin (2012) argue that only the negative association between output and the leads of mark-ups is indicative of a causal connection, backing this up with evidence of Granger (1969) causality from output to mark-ups, but not in the reverse direction. Unfortunately, given the relatively large amount of (near-) idiosyncratic noise in mark-ups, Granger causality tests are unlikely to be reliable. Indeed, we show in the online appendix (Holden 2013b sec. 9) that the presence of idiosyncratic noise in mark-ups can lead Granger causality to go in the opposite direction to the true flow of information in the model, particularly when the two variables have such tight cyclic associations. In our model, there is certainly the possibility of bi-directional information flow, due to both the cross-industry and within-industry composition effects, although the former will tend to dominate. Given the aforementioned problems with Granger causality tests, reduced-form empirical work is little use in deciding between our story and that of Tutino and Cheremukhin (2012).

[9] Constructed using an entirely parameter free method. We first filter the data with a Christiano and Fitzgerald (2003) band-pass filter with a lower cut-off of two quarters and a higher cut-off equal to the data length, in order to remove the influence of structural change and ensure stationarity. We then use the Hurvich (1985) cross-validation procedure to choose the bandwidth for the spectral-decomposition of the data, with his Stuetzle-derived estimator of the mean integrated squared error, the standard Blackman-Tukey lag-weights estimate, and the Quadratic Spectral Kernel recommended by Andrews (1991) amongst others.

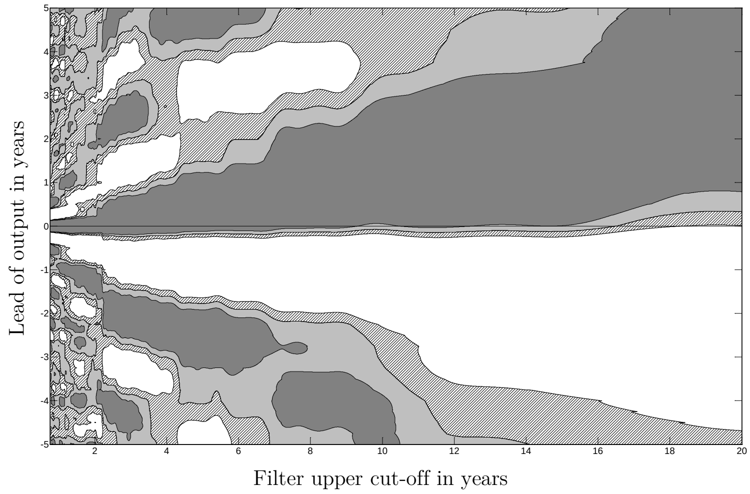

Figure 1: The cross correlation of US output and mark-ups, as a function of filter cut-off. (Dark grey is a significantly positive correlation (at 5%), light grey is a positive but insignificant one, cross-hatched is a negative but insignificant one and white is a significantly negative one.)

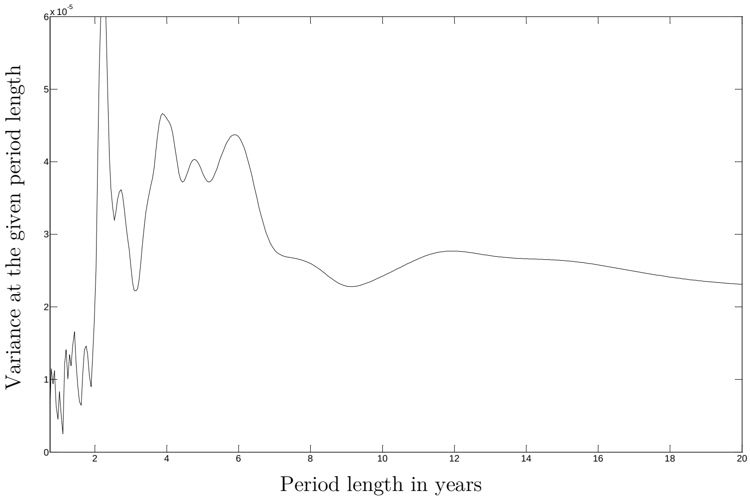

Figure 2: The spectral decomposition of US output growth.

1.3. GDP variance

The model of our companion paper (Holden 2013a) predicts that the length of patent-protection should be positively correlated with the observed size of medium-frequency cycles, at least for durations of patent-protection around those we observe in reality. In Table 2, we exploit cross-country variation in effective patent duration to demonstrate the presence of this correlation in the data, even when we control for GDP, legal origins and various measures of political stability and risk.[10]

[10] Full details of the data are given in footnotes to the table.

Variable	**Spec. 1**	**Spec. 2**	**Spec. 3**	**Spec. 4**	**Spec. 5**
Constant	**-2.09811 (0.0300)**	**-2.14048 (0.0206)**	**-1.91285 (0.0180)**	**-2.70784 (0.0000)**	**-2.18372 (0.0009)**
English legal origin[11]	*-0.0506172 (0.8567)*				-0.448554 (0.0810)
French legal origin[11]	*-0.0557074 (0.8394)*				*-0.350747 (0.1653)*
German legal origin[11]	*-0.151587 (0.6364)*				*-0.325196 (0.3154)*
Log GDP per effective adult[12]	*0.0715242 (0.3620)*	*0.0707845 (0.3501)*			
GDP per effective adult growth[12]	*7.39306 (0.1647)*	*7.24517 (0.1606)*			
Socioeconomic Conditions (ICRG)[13]	**-0.224159 (0.0078)**	**-0.229358 (0.0044)**	**-0.170029 (0.0107)**		
Law and order (ICRG)[13]	-0.154013 (0.0856)	-0.150749 (0.0818)	-0.148729 (0.0856)		
Logit overall political risk (ICRG)[13,14]	**0.806772 (0.0013)**	**0.811630 (0.0006)**	**0.823980 (0.0003)**		
Index of patent duration, 1960[15]	**0.357215 (0.0336)**	**0.363052 (0.0242)**	**0.384211 (0.0131)**	**0.395486 (0.0044)**	**0.396382 (0.0060)**
Index of patent duration, 2005[15]	**1.79391 (0.0223)**	**1.79854 (0.0197)**	**1.88715 (0.0140)**	**1.66419 (0.0053)**	**1.50279 (0.0133)**
Observations	100	100	100	111	111
Specification test p-values[16]	*0.50, 0.31, 0.58*	*0.51, 0.20, 0.63*	*0.58,* 0.08, *0.74*	*0.31,* 0.06, **0.05**	*0.32, 0.12,* 0.06

Table 2: The impact of patent duration on the strength of medium frequency cycles. Coefficients from assorted regression specifications. (P-values in brackets.) In all cases, the dependent variable is a logit transform of the proportion of GDP per effective adult growth variance that is at frequencies with periods greater than eight years[17].

[11] All countries which neither have English, French or German legal origins have Scandinavian legal origin in our sample. Data is from La Porta, Lopez-de-Silanes and Shleifer (2008).

[12] The intercept and the slope from running a regression of log GDP per effective adult on time. Data from the Penn World Tables (Heston, Summers, and Aten 2011), samples identical to those used to construct the dependent variable.

[13] International Country Risk Guide, The PRS Group. Data provided by the Nuffield College Data Library. Variables are means of annual data from 1986-2007 (the largest span available for all countries in the sample).

[14] This is the sum of the two components mentioned above, along with measures of government stability, the investment profile, internal/external conflict, corruption, the military/religion in Politics, ethnic tensions, democratic accountability and bureaucracy quality. The logit transform was taken after the mean. We ran regressions including all components separately and our results were almost identical (p-values on patent duration of 0.0192 and 0.0172 respectively), but to save space here we focus on the components found to be most relevant.

[15] Data kindly provided by Walter Park, updated from Ginarte and Park (1997).

[16] Respectively, a normality test (Doornik and Hansen 2008), the White heteroskedasticity test (White 1980) and the reset test with squares and cubes (Ramsey 1969).

[17] Data is from the Penn World Tables (Heston, Summers, and Aten 2011) and spans 1950-2009, though many countries have shorter samples. The shortest sample (of growth rates) is 23 years. We ran regressions including the sample length as a regressor, but it consistently came out insignificant. Medium frequency variance shares are constructed from spectral decompositions, following Levy and Dezhbakhsh (2003), where the spectral decomposition is performed using the parameter free method outlined in footnote 9, with the initial filter set to accept period lengths between 2 and 59 years (the length of the largest samples).

Patent duration in both 1960 and 2005 has a significantly positive effect (at 5%) on the strength of medium frequency cycles in all our five specifications, and only in in the specification with no controls is there marginal evidence of misspecification (at 5%). Concerns about endogeneity mean some restraint must be exerted in interpreting these results, but they are nonetheless suggestive of a role for patent protection in the mechanism generating medium frequency cycles in the data.

2. The model

The model we present here is a close descendent of that of our companion paper (Holden 2013a). The broad structure of both models is as follows. There is a continuum of industries, each producing a different product. New products are the result of a costly invention process, with a free entry condition determining the quantity of invention, and the rewards from it stemming from the (stochastic-length) patents awarded to inventors. Inventors lack the necessary human capital to produce their product at scale themselves, so they instead offer licenses to a finite number of manufacturing firms to do this, with a free entry condition determining how many firms produce in each industry. Manufacturing firms benefit from free process technology transfer within an industry, but must spend resources in order to catch-up to the highest technology across industries ("appropriation"). They spend further resources on research, which may increase their productivity beyond that of the frontier. Without loss of generality, firms exist for two periods, with entry, appropriation and research in the first, and production in the second. Asymptotically, the cost of appropriation is dominated by that of research, leading to instantaneous process technology transfer across patent-protected industries. However, asymptotically non-patent-protected industries contain infinitely many firms, and do not perform any research or appropriation.

In order to build a model that we may seriously compare to the data, we extend the model of our companion paper to incorporate habits, imperfect competition in labour markets, a variety of shocks, including stochastic movements in key parameters, and both physical and research capital. Research capital will be used in both research, appropriation and invention, and may be thought of as capturing (variously) education, creativity, ideas, knowledge and advanced physical capital. Additionally, we include intermediate goods as a factor of production, where intermediate goods are produced one-for-one from final goods. This may be necessary in order to reconcile the low mark-ups found in micro-evidence with the higher mark-ups implied by aggregate evidence.

We also allow for sticky nominal wages in line with the micro-evidence of Barattieri, Basu, and Gottschalk (2010), and to enable us to make preliminary remarks about the possible medium-term impact of monetary policy. In all of the impulse responses presented below though, we will show the model's performance both with and without this feature. We do not include sticky prices for several reasons. Firstly, it is hard to reconcile the highly sophisticated behaviour of firms in our model with the naïve behaviour of firms in the Calvo (1983) model. Secondly, introducing sticky prices would make solving for firm behaviour very complicated, unless the sticky prices were only introduced to a separate retail sector, further increasing the size of our model. Finally, as is well known, introducing sticky prices results in counter-cyclical mark-ups, contrary to the evidence of Nekarda and Ramey (2010). The observed frequency of price adjustment can perhaps be reconciled with pro-cyclical mark-ups using a consumer search model as in Head et al. (2011). We do not pursue this avenue here.

We now give the detailed structure of the model.

2.1. Households and investment good producers

There is a unit mass of households, each of which contains N_t members in period t. Household $h \in [0,1]$ maximises:

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s N_{t+s} \Theta_{t+s} \left[\log \tilde{C}_{t+s}(h) - \frac{\Phi_{t+s}}{1+\nu} \tilde{L}_t^{\mathrm{S}}(h)^{1+\nu} \right],$$

where β is the discount rate, ν is the inverse of the Frisch elasticity of labour supply to wages, Θ_t is a demand shock, Φ_t is a labour supply shock, $\tilde{C}_t(h) :=$

$\frac{C_t(h)}{N_t} - \hbar \frac{(1-\hbar^{\mathrm{INT}})C_{t-1} + \hbar^{\mathrm{INT}} C_{t-1}(h)}{N_{t-1}}$ is habit adjusted consumption per head,[18] and where $\tilde{L}_t^{\mathrm{S}}(h) := \frac{L_{t+s}^{\mathrm{S}}(h)}{N_{t+s}} - \hbar^{\mathrm{LS}} \frac{L_{t+s-1}^{\mathrm{S}}}{N_{t+s-1}}$ is habit adjusted labour supply per head.[19] Each household supplies a different type of labour $L_t^{\mathrm{S}}(h)$ and potentially receives a different real wage, $W_t(h)$. They face the budget constraint: $C_t + I_t^{\mathrm{KP}} + I_t^{\mathrm{KR}} + B_t = L_t^{\mathrm{S}}(h)W_t(h) + R_t^{\mathrm{KP}} u_t^{\mathrm{P}} K_{t-1}^{\mathrm{P}} + R_t^{\mathrm{KR}} u_t^{\mathrm{R}} K_{t-1}^{\mathrm{R}} + B_{t-1}R_{t-1} + \Pi_t$, where B_t is the aggregate number of (zero net supply) bonds bought by households in period t, R_{t-1} is the period t sale price of a (unit cost) bond bought in period $t-1$, Π_t is the households' period t dividend income, I_t^{KP} and I_t^{KR} is investment in physical and research capital, respectively,[20] and $u_t^{\mathrm{P}} K_{t-1}^{\mathrm{P}}$ and $u_t^{\mathrm{R}} K_{t-1}^{\mathrm{R}}$ are the quantities of these stocks that households make available to firms, with u_t^{P} and u_t^{R} their chosen utilisation rates and K_{t-1}^{P} and K_{t-1}^{R} the level of the capital stocks at the end of period $t-1$. The utilisation of research capital decision may be thought of as capturing the incentives to bunch the implementation of ideas, as stressed by Francois and Lloyd-Ellis (2008; 2009).

We model sticky nominal wages in the standard Calvo (1983) fashion, following Erceg, Henderson and Levin (2000). Each household is able to set its wage optimally with probability $1 - \boldsymbol{v}$. We assume that those households that cannot adjust their wage optimally will fully index their wage to its steady-state growth rate.

Following Schmitt-Grohé and Uribe (2011), investment goods of type $\mathrm{V} \in \{\mathrm{P,R}\}$ are produced from consumption goods using the technology $I_t^{\mathrm{KV}*} = A_t^{*\xi_{\mathrm{KV}}} E_t^{\mathrm{KV}} I_t^{\mathrm{KV}}$ where I_t^{KV} is investment in units of consumption goods and $A_t^{*\xi_{\mathrm{KV}}} E_t^{\mathrm{KV}}$ captures investment specific technological change, as a short-cut alternative to modelling separate endogenous growth processes in a multi-sector model. As in Schmitt-Grohé and Uribe (2011), the productivity of the frontier industry (A_t^*, the underlying trend in productivity, A_t) enters into this expression in order to capture the cointegration between the relative price of investment

---

[18] With $\hbar \in [0,1)$ controlling the strength of consumption habits and $\hbar^{\mathrm{INT}} \in [0,1]$ controlling whether consumption habits are internal or external.

[19] With $\hbar^{\mathrm{LS}}$ determining the strength of these external labour habits.

[20] We assume a complete set of nominal state contingent securities, meaning B_t, C_t, Π_t, I_t^{KP} and I_t^{KR} will not differ across households.

and productivity that is observed in the data. It may be justified as reflecting improvements in installation technologies, or improvements to the allocation of new capital across firms, both of which come as a side-effect of the increase in general knowledge following an increase in A_t^*. Explicitly modelling a role for human capital in physical capital production would generate very similar results, while adding unnecessary complications.

Both capital stocks evolve according to:

$$K_t^{\mathrm{V}} = \left(1 - \delta_t^{\mathrm{V}}(u_t^{\mathrm{V}})\right) K_{t-1}^{\mathrm{V}} + \Gamma_t I_t^{\mathrm{KV}*} \left[1 - Q^{\mathrm{KV}} \left(\frac{I_t^{\mathrm{KV}*}}{I_{t-1}^{\mathrm{KV}*}}\right)\right]$$

for V ∈ {P,R}, where $\delta_t^{\mathrm{V}}(\cdot)$ for V ∈ {P,R} are increasing functions capturing the effect of utilisation on depreciation, locally convex at the steady-state, $Q^{\mathrm{KV}}(\cdot)$ for V ∈ {P,R} are convex functions capturing adjustment costs to the rate of investment (following Christiano, Eichenbaum, and Evans (2005)), which attain their minimum value of zero at the steady state rate of growth of investment, and where Γ_t is a shock to the marginal efficiency of investment, which, following Justiniano, Primiceri, and Tambalotti (2011), we will identify with a decreasing function of Moody's BAA-AAA bond spreads.[21] $\delta_t^{\mathrm{V}}(\cdot)$ has a time subscript since we allow for a shock to depreciation to capture some of the volatility in depreciation shares that we observe in the data.[22] There is a single shock across both capital types, which we call $\tilde{\delta}_t$, and it is constrained to weakly increase both the levels and the first derivatives of $\delta_t^{\mathrm{P}}(\cdot)$ and $\delta_t^{\mathrm{R}}(\cdot)$.[23] Depreciation shocks have been shown to be important by Dueker, Fischer, and Dittmar (2007), Liu, Waggoner, and Zha (2011) and Furlanetto and Seneca (2011) amongst others, and will turn out to be important here too. As these authors note, they may be

[21] Justiniano, Primiceri, and Tambalotti (2011) used the high yield to AAA spread. We choose the BAA-AAA one due to increased data availability. The difference between Γ_t and E_t^{KV} is that only the latter will appear in the measured relative price of investment, and only the former is common to both processes.

[22] Our measure of depreciation is the consumption of fixed capital from NIPA. If anything, this will underestimate the true variance of depreciation, since the NIPA measure omits variation in depreciation rates within individual product categories. We thank Martin Seneca for this observation.

[23] We additionally constrain the response of $\delta_t^{\mathrm{V}}(\cdot)$ to the shock such that in its linearised version, with utilisation at its steady-state level, both $\delta_t^{\mathrm{V}}(\cdot)$ and $\delta_t^{\mathrm{V}\prime}(\cdot)$ are positive with at least 95% probability. This is true automatically in the source non-linear specification in which $\delta_t^{\mathrm{V}}(\cdot)$ and $\delta_t^{\mathrm{V}\prime}(\cdot)$ are log-linear in $\tilde{\delta}_t$ when utilisation is at its steady-state, but in preliminary estimates the linearised $\delta_t^{\mathrm{P}\prime}(\cdot)$ turned negative a high proportion of the time, in the absence of this additional constraint.

interpreted as proxying for a combination of product specific capital, heterogeneity in capital quality across products, and changes in consumer preferences across these products. With this interpretation allowing depreciation shocks to affect the first derivative of $\delta_t^{\mathrm{V}}(\cdot)$ as well as its level is natural, since low quality capital will both break faster on average, and be more sensitive to heavy usage. This will also aid us in matching the negative correlation between depreciation and utilisation that is observed in the data.

2.2. Aggregators

The consumption good is produced by a perfectly competitive industry from the aggregated output $Y_t(i)$ of each industry $i \in [0, I_{t-1}]$, using the following Dixit-Stiglitz-Ethier (Dixit and Stiglitz 1977; Ethier 1982) style technology:

$$Y_t = I_{t-1}^{-\lambda} \left[\int_0^{I_{t-1}} Y_t(i)^{\frac{1}{1+\lambda}} \, di \right]^{1+\lambda}$$

where $\frac{1+\lambda}{\lambda}$ is the elasticity of substitution between goods and where the exponent on the measure of industries (I_{t-1})[24] has been chosen to remove any preference for variety in consumption. We normalize the price of the aggregate good to 1.

Similarly, each industry aggregate good $Y_t(i)$ is produced by a perfectly competitive industry from the intermediate goods $Y_t(i,j)$ for $j \in \{1, \dots, J_{t-1}(i)\}$,[25] using the technology:

$$Y_t(i) = J_{t-1}(i)^{-\eta\lambda} \left[\sum_{j=1}^{J_{t-1}(i)} Y_t(i,j)^{\frac{1}{1+\eta\lambda}} \right]^{1+\eta\lambda}$$

where $\eta \in (0,1)$ controls the degree of differentiation between firms, relative to that between industries.

Aggregate labour services to firms are provided by a competitive industry of labour packers using the technology:

$$L_t^{\mathrm{T}} = A_t^{*\xi_{\mathrm{L}}} E_t^{\mathrm{L}} \left[\int_0^1 L_t^{\mathrm{S}}(h)^{\frac{1}{1+\lambda_{\mathrm{L}}}} \, dh \right]^{1+\lambda_{\mathrm{L}}},$$

where E_t^{L} is an exogenous stationary labour productivity shock. (In the absence of research and development, this E_t^{L} shock would act exactly like a classical

[24] The $t-1$ subscript here reflects the fact that industries are invented one period before their product is available to consumers.

[25] Again, the $t-1$ subscript reflects the fact that firms enter one period before production.

TFP shock.) The productivity of the frontier (A_t^*) enters our expression for labour services in order to capture the improvements in labour productivity that arise from the higher knowledge levels after an increase in frontier productivity. Again, explicitly modelling human capital evolution would add little to our model's performance. However, following Jaimovich and Rebelo (2008) we do include labour adjustment costs. In particular, we assume that in sector V $\in$ {P,R} there is a perfectly competitive industry that transforms aggregate labour services into sector specific labour services using the technology $L_t^{\mathrm{EV}} = L_t^{\mathrm{TV}} \left[1 - Q^{\mathrm{LV}} \left(\frac{L_t^{\mathrm{TV}}}{L_{t-1}^{\mathrm{TV}}}\right)\right]$, where $Q^{\mathrm{LV}}(\cdot)$ is a monotone increasing function that is zero at the steady state rate of growth of L_t^{TV}. The aggregate labour market clearing condition is then $L_t^{\mathrm{T}} = L_t^{\mathrm{TP}} + L_t^{\mathrm{TR}}$. In the absence of labour adjustment costs, there is a risk that the capital share of R&D would be biased upwards since there are adjustment costs to capital. Labour adjustment costs also help generate plausible business cycles in response to news about future productivity (Jaimovich and Rebelo 2008), which may be important here due to the endogenous movements in future productivity that our model generates.

2.3. Intermediate firms

Firm j in industry i has access to the Cobb-Douglas production technology $Y_t(i,j) = A_t(i,j) X_t^{\mathrm{P}}(i,j)^{\iota_{\mathrm{P}}} [K_t^{\mathrm{P}}(i,j)^{\alpha_{\mathrm{P}}} L_t^{\mathrm{P}}(i,j)^{1-\alpha_{\mathrm{P}}}]^{1-\iota_{\mathrm{P}}}$ where $X_t^{\mathrm{P}}(i,j)$ is their level of intermediate good input,[26] $L_t^{\mathrm{P}}(i,j)$ is their input of production specific labour service, and $K_t^{\mathrm{P}}(i,j)$ is the quantity of capital they hire from households, at a cost of R_t^{KP} per unit. We use a Hicks-neutral specification here since it minimises the changes necessary to the model of Holden (2013a).[27] Pricing decisions and cross industry aggregation are the same as in that paper.

The presence of intermediates in production will amplify shocks in our economy, as it implies that an increase in the proportion of industries that are patent-protected means intermediate inputs are cheaper for non-protected industries, increasing their output too. However, the two spill-overs from frontier productivity growth previously mentioned (that to the technology for producing

[26] Recall that intermediate goods are identical to the final aggregate consumption good.
[27] In particular, profits take the same form, and so research incentives are identical.

investment goods, and the technology for aggregating labour) will help dampen our model's then overly powerful amplification mechanism, since they permit lower levels of firm-level productivity growth, while still matching aggregate growth rates.[28]

In the period prior to production, firms pay a fixed entry cost, and possibly costs for license rents, appropriation and research. Asymptotically, the fixed entry cost is irrelevant, and appropriation is effectively free for firms in patent-protected industries, while being too expensive for all firms in non-patent-protected industries. Given it is the asymptotic behaviour of the model we simulate, we refer the reader to our companion paper (Holden 2013a) for the full non-asymptotic specification of these two costs.[29]

Also asymptotically, thanks to the firm free-entry condition, bargaining between inventors and firms will lead patent holders to set license rents to $\frac{1-\wp}{\wp}$ times the expenditure of each firm in their industry on research, where $\wp \in (0,1)$ is the bargaining power of the firm, in the sense of the generalized Nash bargaining solution. Again, we refer the reader to our companion paper (Holden 2013a) and the online appendix (Holden 2013b sec. 3) for the full details. This just leaves research expenditures to specify.

If firm j in patent protected industry i employs in research, in period t, $X_t^{\mathrm{R}}(i,j)$ units of intermediate goods, $L_t^{\mathrm{R}}(i,j)$ units of research specific labour

[28] These spillovers mean that the steady-state growth rate of real output per capita is given by $\frac{g_A}{(1-\alpha_{\mathrm{P}})(1-\iota_{\mathrm{P}})} + \left(\xi_{\mathrm{L}} + \frac{\alpha_{\mathrm{P}}}{1-\alpha_{\mathrm{P}}}\xi_{\mathrm{KP}}\right) g_{A^*}$.

[29] In this richer set-up, we may ensure the asymptotic structure of the model is the same under identical conditions by assuming that the input to fixed costs and appropriation is produced using the same production function as that to research and invention. Away from this special case the lower bound on ζ would be non-zero, and possibly negative.

services, and $K_t^{\mathrm{R}}(i,j)$ units of hired research capital, then its productivity level in period $t+1$ will be given by:[30]

$$A_{t+1}(i,j) = A_t^*\big(1 + \gamma Z_{t+1} A_t^{*-\zeta} X_t^{\mathrm{R}}(i,j)^{\iota_R}[K_t^{\mathrm{R}}(i,j)^{\alpha_R} L_t^R(i,j)^{1-\alpha_R}]^{1-\iota_R}\big)^{\frac{1}{\gamma}},$$

where ζ controls the extent to which research is getting harder over time, Z_{t+1} is a shock representing the luck component of research (common[31] across firms and industries), and $\gamma > 0$ controls the "parallelizability" of research. In equilibrium, all firms in patent-protected industries will perform the same amount of research, meaning $A_{t+1}^* = A_{t+1}(i,j)$ for any firm j in a patent-protected industry i. Firms choose research levels to maximize expected discounted profits, leading to a solution for research levels of much the same form as that in our companion paper (Holden 2013a).[32]

2.4. Inventors

Each new industry is controlled by an inventor who owns the patent rights to the product the industry produces. Until the inventor's product goes on sale, the patent holder can successfully protect their revenue stream through contractual arrangements, such as non-disclosure agreements. This means that even in the absence of patent-protection a patent holder will receive one period of license rents. The inventor of a new product has a probability of $1-q$ of being granted a patent to enable them to extract rents for a second period. After this, if they have a patent at t, then they face a constant probability of $1-q$ of having a patent at $t+1$.

Thanks to instantaneous catch-up amongst firms in patent-protected industries,[33] we may assume that new products begin life with a production process at the frontier. However, since the frontier is growing over time, this suggests invention too may be becoming harder over time. We assume that inventing a new product requires $\mathcal{L}_t^{\mathrm{I}} A_t^{*\zeta}$ units of "invention output", where $\mathcal{L}_t^{\mathrm{I}}$ is

[30] Again this expression is strictly only valid asymptotically, since it assumes instantaneous cross-industry catch-up.

[31] The common shock assumption is justified in Holden (2013a).

[32] The online appendix (Holden 2013b sec. 5) contains the new solution.

[33] See Holden (2013a) for the full conditions for this.

a shock determining the difficulty of invention, and where ζ captures the fact that invention is getting harder over time.[34] For simplicity, we assume invention output is produced using the same form of Cobb-Douglas production function as research (i.e. $X_t^{\mathrm{R}}(\cdot)^{\iota_R}[\mathrm{K}_{\mathrm{t}}^{\mathrm{R}}(\cdot)^{\alpha_R} L_t^{\mathrm{R}}(\cdot)^{1-\alpha_R}]^{1-\iota_R}$).

Inventors will enter while the expected discounted profits from doing so are larger than the costs of inventing a new product, and it will always be the case that $I_t \geq I_{t-1}$. This means there may possibly be some asymmetry in the model's response to shocks, though in practice this does not appear empirically important.

2.5. Closing the model

To close the model, we specify a $\log \mathrm{AR}(1)$ form for all the model's shocks, with the exception of Z_t, the true technology shock which is uncorrelated across time as in Holden (2013a).[35] The data will be allowed to choose which, if any, of these shocks might be important drivers of business cycles, at high, or medium frequencies.

We also specify an augmented Taylor rule form for the setting of nominal interest rates. We allow the central bank to respond to all prices in the economy (i.e. the price of consumption, production investment, research investment and labour), four proxies for the real interest rate (the return on production and research investment, the demand shock and the depreciation shock), as well as both output's deviation from trend and its growth rate.[36] In the absence of endogenous productivity, the optimal policy would fully stabilise nominal wages,

---

[34] The fact this exponent is equal to the equivalent exponent on research difficulty is an assumption made for convenience, it is not necessary for the model to be well behaved. See Holden (2013a).

[35] Z_t is assumed i.i.d. log-normal. The shocks driving the various AR(1) processes are assumed i.i.d. normal.

[36] In particular, $\frac{R_t^{\mathrm{NOM}}}{R^{\mathrm{NOM}}} = \left(\frac{R_{t-1}^{\mathrm{NOM}}}{R^{\mathrm{NOM}}}\right)^{\rho_{R^{\mathrm{NOM}}}} \left[\left(\frac{G_{P,t}}{G_{P,t}^*}\right)^{\mathcal{M}_{\mathrm{P}}} \left(\frac{E_{t-1}^{\mathrm{KP}} G_{A^*}^{\xi_{\mathrm{KP}}}}{E_t^{\mathrm{KP}} G_{A^*,t}^{\xi_{\mathrm{KP}}}}\right)^{\mathcal{M}_{\mathrm{PKP}}} \left(\frac{E_{t-1}^{\mathrm{KR}} G_{A^*}^{\xi_{\mathrm{KR}}}}{E_t^{\mathrm{KR}} G_{A^*,t}^{\xi_{\mathrm{KR}}}}\right)^{\mathcal{M}_{\mathrm{PKR}}} \left(\frac{G_{W,t}}{G_W}\right)^{\mathcal{M}_{\mathrm{W}}}\right]^{1-\rho_{R^{\mathrm{NOM}}}} \cdot \left[\left(\frac{R_t^{\mathrm{KP}}}{A_t^{*\,\xi_{\mathrm{KP}}}}\right)^{\mathcal{M}_{\mathrm{RKP}}} \left(\frac{R_t^{\mathrm{KR}}}{A_t^{*\,\xi_{\mathrm{KR}}}}\right)^{\mathcal{M}_{\mathrm{RKR}}} \Theta_t^{\mathcal{M}_\Theta} \tilde{\delta}_t^{-\mathcal{M}_{\tilde{\delta}}}\right]^{1-\rho_{R^{\mathrm{NOM}}}} \left[\left(\frac{Y_t}{N_t A_t^{a} A_t^{*e}}\right)^{\mathcal{M}_{\mathrm{Y}}} \left(\frac{G_{Y,t}/G_{N,t}}{G_Y/G_N}\right)^{\mathcal{M}_{\mathrm{G}}}\right]^{1-\rho_{R^{\mathrm{NOM}}}} \cdot \exp \epsilon_{R^{\mathrm{NOM}},t}$, where R_t^{NOM} is the gross nominal interest rate, $G_{P,t}$ is the (gross) growth rate of the nominal price of the consumption good, $G_{P,t}^*$ is the stochastic target for this growth rate, $\frac{E_{t-1}^{\mathrm{KV}}}{E_t^{\mathrm{KV}} G_{A^*,t}^{\xi_{\mathrm{KV}}}}$ is the growth rate of the *real* price of investment goods of type $\mathrm{V} \in \{\mathrm{P},\mathrm{R}\}$, $G_{W,t}$ is the growth rate of the real wage, Y_t/N_t is log real GDP, $G_{Y,t}/G_{N,t}$ is the real per capita GDP growth rate and $R_t^{\mathrm{SHOCK}} := \exp \sigma_{R^{\mathrm{NOM}}} \epsilon_{R^{\mathrm{NOM}},t}$ is a monetary policy shock. Variables without time subscripts are steady-state values, and the constants a and e are defined in the online appendix (Holden 2013b sec. 5).

completely removing the Calvo distortion, thus it is important to allow wages to enter the Taylor rule.[37] It turns out however that the only significant terms in the estimated Taylor rule are the lag, the price response, and the response to the depreciation shock and the rental rate of production capital (which are tightly correlated with the Wicksellian real interest rate (Woodford 2001)), so the estimated rule takes a more standard form.

The model's full set of de-trended equations is given in the online appendix (Holden 2013b sec. 5).

3. Empirical tests

3.1. Data and estimation

The model is estimated on logs of quarterly U.S. series for nominal output growth,[38] consumption price inflation, investment price inflation, population growth, labour supply per capita, the R&D share, the consumption share, the labour share, the depreciation share, nominal interest rates, capacity utilisation and the BAA-AAA spread. The longest samples are from 1947Q1 to 2011Q2, though some series are shorter. (Our estimation method can cope with an uneven sample.) Most series comes from NIPA or the FRB. Full details of the sources and construction methods of the data are given in the online appendix (Holden 2013b sec. 6), and the full data set is available from the author on request.

In order to remove any structural change, we filter the data before estimation, with a high-pass filter that allows frequencies with periods below the sample length (258 quarters). We adjust the level of the filtered data so that the mean of the filtered series matches that of the original data. (Broadly) following Canova (2009) we also include IID, AR(1) and repeated-root AR(2)[39] "measurement

---

[37] There is no guarantee though that this prescription carries over into our model with endogenous productivity. We intend to investigate optimal policy in this model in future work.

[38] We use nominal output as there should be less measurement error in the nominal series than in the real series.

[39] Our justification for going up to a repeated-root AR(2) process is that as the auto-regressive parameter of such a process tends to one, the process becomes an I(2) trend, which is exactly the type of trend removed by the widely used HP-filter (Hodrick and Prescott 1997). In order to avoid implicitly removing an I(3) trend from the series in differences (nominal output growth, consumption price inflation, investment price inflation and population growth) we suppose that the measurement error enters the observation equations for these series with the over-differenced moving average form $\mathrm{me}_t - \mathrm{me}_{t-1}$.

error" shocks in each observation equation, to prevent our model from being contorted to fit the data. (Canova (2009) advocates the inclusion of IID, I(1) and I(2) shocks.)

In standard DSGE models, there are usually enough degrees of freedom that almost any set of first moments may be matched without impacting the model's ability to match second moments. The presence of endogenous growth in our model, though, means this is no longer true for us. In our model, almost all first moments are tightly coupled both to each other (e.g. the labour-share, mark-ups and growth) and to the model's dynamics. This raises the possibility that our model's inevitable misspecification may mean it is impossible for our model to match simultaneously all first moments without grossly compromising its dynamics. The Canova (2009) approach is to discard all information about first moments, and to assume the "measurement error" has a unit root, but this necessitates the use of strong priors, something that is infeasible here since the dimensionality of our model rules out MCMC based estimation. Additionally, allowing unit roots in measurement error would prevent us using the variance share of measurement error as a measure of the quality of our model. Instead, we allow for a mean term in the measurement error to prevent misspecification of the kind described from severely biasing other parameters. However, to ensure the means of the data series remain informative, we follow Lee et al. (2010) and Candès, Wakin, and Boyd (2008) in imposing a sparsity inducing "adaptive lasso" (generalized t) prior on these mean measurement error terms.[40]

Since we want our model to rely on its internal persistence mechanism, rather than the persistence of shocks, and since we want all shocks to be stationary, we impose a prior on all the "ρ" parameters of our model (these include the persistence of shocks, the persistence of AR(1) and repeated-root AR(2)

[40] In the notation of Lee et al. (2010), in this prior we set a_j to the length of the data to the power of $^1\!/_3$ (to ensure the method possesses the oracle property), and b_j is chosen so that the expected absolute measurement error mean term is 1%. To reduce the dimensionality of the state space, we force these measurement error mean terms to the level at which the model's steady state for observable variables exactly matches their mean in the data.

measurement errors, and the persistence of monetary policy). We use a logit-normal distribution that is scaled to $[-1,1]$ then truncated to $[0,1]$.[41] We set the mean of the underlying normal distribution to 0 and its variance to 2, which are the unique values which result in a density which has zero first, second and third derivatives at the origin, ensuring small to medium values of ρ are not biased.

We fix the discount factor (β) at 0.99 following standard practice for quarterly models. We also bound the inverse-Frisch elasticity (ν) to be above 0.25, which is a lower bound on standard macro calibrations as reported by Peterman (2011). All the other parameters of our model are given flat priors. We then estimate by the "maximum a posteriori" method (which is very close to maximum-likelihood since the majority of parameters have flat priors), subject to[42]:

- all variables being stationary,
- a unique (determinate) solution existing for both the simple model and this extended one, (with an identical number of firms per industry in both, and with all parameters identical except possibly $\mathcal{L}^{\mathrm{I}}$),
- all parameters being in the region in which the model is well behaved asymptotically,[43]
- the steady-state value of the average mark-up (μ_t) equalling 0.056 (to 3 decimal places), in line with the micro-evidence of Boulhol (2007),[44]
- patent protected industries being 17% (to 0 decimal places) more productive than non-protected industries in steady-state, in line with the micro-evidence of Balasubramanian and Sivadasan (2011),[45]

[41] I.e. if Z is normally distributed, $\frac{1-\exp(-Z)}{1+\exp(-Z)}\Big| Z > \frac{1}{2}$ has our distribution.

[42] We also constrained the share of medium frequency variance (as measured by applying a perfect filter to the spectral density generated by the transition matrices, with accepted band between 8 and 60 years) to decrease when the mean length of patent protection is reduced by one quarter, in line with the evidence of section 1.3, but this constraint did not bind at the optimum.

[43] $\mathcal{p}\lambda\mathbb{E}\gamma_t < 1$, $\lambda\mathbb{E}\eta_t\gamma_t \geq 1$, $\mathbb{E}\boldsymbol{d}_t > \mathcal{p}\mathbb{E}\gamma_t\mu_t^{\mathrm{P}}$, $\mathbb{E}g_{I,t} > 0$ and $\mathbb{E}J_t^P > 1$.

[44] This is implemented by adding the steady-state mark-up as an additional observation variable to the model, with an NIID(0,0.0005) shock (added both to the data and to the model, with known standard deviation).

[45] Similarly, this is implemented by adding the steady-state value of $\log \widehat{A}_t^{\mathrm{N}}$ as an additional observation variable, with an $\mathrm{NIID}\big(0, 1/2\,(\log(1/1.165) - \log(1/1.175))\big)$ shock (as before).

- the correlation of log mark-ups (as measured by the inverse labour share) and log output, being positive when the data is filtered by a cut-off of one, five or eleven years and *negative* when the data is filtered by a filter with a cut-off of twenty years, in line with the evidence of Figure 1.[46]

By disciplining mark-ups and relative-productivity from micro-evidence, we hope to go some way to answering the concerns about the introduction of free-parameters raised by Chari, Kehoe and McGrattan (2009). For technical reasons, we ignore the positivity constraint on the growth rate of the stock of products during estimation.

The maximisation is carried out using the CMA-ES algorithm (Hansen et al. 2009), which is known to have good global search performance, particularly when run with large populations, as we do. However, although the dimensionality of our model is much smaller than that of a VAR(1) run on the same series, we still cannot absolutely guarantee that a global maximum has been found. This is a standard problem in estimating large models.

3.2. Estimation results

The full list of estimated parameters is given in the online appendix (Holden 2013b sec. 7). We briefly discuss a few key parameters here however. In the below, approximate posterior standard errors are given in brackets. (These are generated from the optimisation algorithm, which gives the inverse hessian of a robust quadratic approximation to the upper envelope of the maximand. Our Monte Carlo experiments indicate that the resulting standard errors are moderately biased upwards, meaning that parameters may be estimated more precisely than they appear to be.[47])

[46] More specifically, we begin by generating 2^{10} simulated runs from the model, each the same length as the data, using the same random seed for each set of runs, for the sake of variance reduction. We then take the correlation of the given variables at each filter cut-off, for each of the runs. We require that the proportion of the runs for which these are of the correct sign is both greater than one-half and significantly different from one-half at 5%. (We use a two-sided test in order to preserve comparability with Figure 1.)

[47] Our estimate of the Hessian of the maximand may be affected by the inclusion of exact bound constraints, since these will tend to reduce the variance of parameters that lead the bound constraint to be violated. However, our procedure estimates the scale of the hessian separately, so still on average over all parameters we expect posterior standard errors to be upward biased.

$\mathcal{p}$ was estimated at 0.0427 (0.00021), implying that manufacturing firms have very little bargaining power in dealing with patent holders. The large bargaining power of patent holders suggests that they may be bargaining simultaneously with all firms keen to licence their product, rather than bargaining with each independently as in our model. In future work we intend to study the strategic interactions in this simultaneous bargaining and entry process more rigorously.

$\mathcal{q}$ was estimated at 0.0374 (0.00030), which implies that only 4.9% of patents last twenty years. This is consistent with some patented products not being commercialised until long after their patent was granted, and others having their patent challenged in court prior to their expiry. It is also consistent with a broader interpretation of "patent protection" within the model, since some inventors are able to exclude entry to their industry for a while, even in the absence of patent protection, via obfuscation or contractual arrangements.

Our estimates imply that γ, which controls the "parallelizability" of research is 18.6 (0.054). In line with the evidence of Siliverstovs and Kancs (2012), this results in an R&D elasticity of productivity that is increasing and concave in research, providing some validation for our chosen functional form. At the steady-state, it implies an R&D elasticity of output of around 0.015. This is not significantly different to estimates using firm level data from e.g. Hall and Mairesse (1995), Bartelsman et al. (1996), Los and Verspagen (2000), Griffith, Harrison, and van Reenen (2006), or Doraszelski and Jaumandreu (2008), though there are certainly higher estimates in the literature as well.[48]

The inverse Frisch elasticity of labour supply was driven to its lower bound of $\nu = 0.25$ by the estimation procedure.[49] While older studies suggested that such highly elastic labour supplies were difficult to reconcile with the micro-data, recent studies (e.g. Peterman (2011) and Keane & Rogerson (2012)) have concluded that highly elastic labour supplies are consistent with the micro evidence when that data includes a broad range of individuals, and is interpreted in light of e.g. human capital accumulation. Our model also includes labour

[48] See the comprehensive survey article by Hall, Mairesse, and Mohnen (2010).

[49] When this bound was not imposed, the estimated value was below 0.01.

adjustment costs, which make aggregate labour supply appear less elastic. Consequently, a standard RBC calibration of the Frisch elasticity based on simulated data from our model would produce a much lower Frisch elasticity than 4. In light of this, we do not consider our estimated elasticity to be implausible. Nonetheless, in future work we intend to investigate the performance of our model when it is augmented by employment search and participation decisions.

α_{P} was estimated to be 0.201 (0.00040), much lower than the traditional value for the capital share of around 0.3. In line with this low value, the consumption share generated by our model was about 10.9% higher than the true value, and the labour share was around 34.5% higher. The treatment here of net exports as investment may be one factor that is biasing down the capital share, due to the US's persistent trade deficit. Another explanation is the existence of some missing heterogeneity across sectors in the real world, with the sectors that are driving growth (e.g. services) tending to be more labour intensive. There is further evidence of missing sectoral heterogeneity in the estimated intermediate goods share in production of 0.0534 (0.0026), (standard estimates are around 0.4), however, this is most likely just a function of the absence of a retail sector in our model. Allowing for the possibility that consumption of intermediate goods in R&D is measured as investment, rather than intermediate consumption, would also help fix these shares as it would decrease the numerators and increase the denominators ($\iota_{\mathrm{R}} = 0.178$ (0.0032)).

However, the low value for the capital share of output is at least partially balanced by a very high estimated value for the capital share of R&D ($\alpha_{\mathrm{R}} = 0.996$ (7.4×10^{-6})). Further insight into the nature of this research-capital comes from the very high adjustment costs to increasing the growth rate of its stock ($Q^{\mathrm{R}\prime\prime}(G_{I^{\mathrm{KR}*}}) = 62.6$ (4.0), in comparison, $Q^{\mathrm{P}\prime\prime}(G_{I^{\mathrm{KP}*}}) = 0.00533$ (0.0012)). These values suggest our interpretation of research-capital as being an external "idea-stock" may be correct. Additional evidence for this comes from the fact that depreciation shocks knock large amounts off the level of the research capital stock (ideas we thought were good turned out to be not so great), whereas they only affect the sensitivity of production-capital depreciation to utilisation (machines we thought to be reliable turned out to be quite sensitive).

In estimating our model, we allowed the data to specify whether investment in R&D capital was measured in the standard national accounts, or whether it was only measured in the R&D satellite account data, since it was not obvious a priori that those producing the accounts can distinguish investment to help future R&D from investment to help future production. Our estimates suggests that 49.4% (1.3%) of all R&D investment is actually captured by the standard national accounts, with the rest measured in the satellite accounts. This level of mismeasurement seems plausible given the difficulties in ascertaining for what a piece of physical capital will be used.

The frictions in our model take plausible values, with households able to update their wage optimally in 17.4% (0.42%) of quarters, which is not statistically different (at 5%) from the probability of a wage change for hourly workers found in micro data by Barattieri et al. (2010) (18%). Recall, too, that when households in our model cannot optimally update their wage, they instead index to steady-state inflation, so the welfare costs of this friction are likely to be small. As observed previously, there is virtually no adjustment cost on production capital, however we find a substantial adjustment cost to production labour ($Q^{\mathrm{LP}\prime}(G_{L^{\mathrm{TP}}}) = 0.0875$ (0.0047)). As shown by Jaimovich and Rebelo (2009), this enables the model to produce co-movement in response to news about future productivity, which is provided in our model by almost any standard shock, thanks to the endogenous growth mechanism. Consumption habits are estimated as being predominately external ($\hbar^{\mathrm{INT}} = 0.0151$ (0.0032)), and much less strong than in many DSGE models ($\hbar = 0.253$ (0.0041)). Estimated habits in labour are negligible. This lesser role for habits of both kinds stems from the much stronger persistence mechanism in our model.

We now turn to the estimated sources of growth. Core (Hicks-neutral) frontier productivity is estimated to grow at 1.11% per year, which is further scaled up by the roles of intermediates and capital, along with the various spillovers, to arrive at an aggregate real growth rate (in units of the consumption good) of

1.57% per year, only slightly lower than that found in the data (1.76% per year[50]).[51] The importance of spillovers for growth has been stressed extensively in the empirical literature before (Griliches 1998; Eaton and Kortum 1999; Forni and Paba 2002; Klenow and Rodríguez-Clare 2005).

Finally, on the sources of cycles, we find that all variables are primarily driven by the depreciation shock, with lesser contributions from the labour supply shock and the population shock. The monetary policy shock plays an even smaller role (contributing to less than 1% of each variable's non measurement error variance), and all other shocks make a negligible contribution. (The full variance decomposition is given in the online appendix (Holden 2013b sec. 7).) Of note is the fact that all shocks have a persistence parameter of less than 0.9, suggesting that the model is able to generate the observed persistence in macroeconomic time series on its own.

The depreciation shock is estimated as having two distinct effects here. Firstly, it increases the sensitivity of the production-capital depreciation rate to increased utilisation. Since the derivative of the depreciation rate with respect to utilisation enters directly into the investment and utilisation equations, even under a first order approximation this can have a large effect on investment and utilisation, by increasing the costs of using capital. Secondly, it increases the depreciation rate of the stock of research-capital, independent of utilisation. The natural interpretation for the shock then is as a proxy for the financial wedge. Indeed, the correlation between the estimated series for $\tilde{\delta}_t$ and the BAA-AAA spread is 0.296 (with a p-value of less than 0.00001), confirming this interpretation.

In a time of great uncertainty, or low asset values, such as the aftermath of the recent crisis, if capital is "put to work" there is a risk it will disappear completely. This is in the spirit of the Kiyotaki-Moore model (Kiyotaki and Moore 1997), and captures the first of these two effects. (For an example that

[50] The low figure comes from deflating by the consumption price, rather than by a consumption-investment price aggregate.

[51] It is likely that there is some downwards bias in real GDP growth rate estimates, due to the difficulty of valuing new products (Broda and Weinstein 2010), so in future work we intend to examine the robustness of our results to correcting for this in the data, at least approximately.

makes clear the effect is on the sensitivity of the depreciation rate to utilisation, consider the incentives of a mortgage-holder in negative-equity to maintain their house.) That financial shocks should result in an increase in the depreciation rate of the stock of ideas is equally clear. In the absence of sufficiently valuable collateral, inventors may be unable to finance the commercialisation of their invention, and by the time asset values recover, it may no longer be "timely" enough to warrant that expense. Obviously, this calls for the inclusion of structurally modelled financial frictions in our model. We intend to pursue this avenue in future work.

3.3. Model evaluation

We use the estimated amount of measurement error to quantify the model's performance. Aside from the two series previously discussed (the labour and capital shares), all of our series had mean levels of measurement error below 0.05%, implying the model is well able to capture the rest of the data's first moments. This leaves the data's second moments to discuss. Since our model is designed to explain cycles at business and medium frequencies, but is unlikely to be able to match either very high frequency noise, or low-frequency structural change, we report measurement error variance in a range of frequency bands. (These are produced by applying perfect filters to the measurement error and observation variable series.) The results of this may be seen in Table 3 below.

Data series	**High frequency 0-1 years**	**Business cycles 1-8 years**	**Medium frequency 8-50 years**	**Low frequency >50 years**
Nominal output growth	2.2%	9.8%	44.1%	1.3%
Consumption price inflation	89.0%	94.0%	66.1%	2.4%
Investment price inflation	97.6%	99.0%	93.2%	17.7%
Population growth	6.6%	37.2%	89.9%	80.1%
Labour supply per capita	44.6%	24.7%	48.8%	82.6%
R&D share	0.0%	0.0%	0.0%	0.0%
Consumption share	67.3%	22.3%	16.7%	35.2%
Labour share	100.0%	100.0%	99.6%	99.1%
Depreciation share	5.5%	37.4%	83.1%	89.9%
Nominal interest rates	86.6%	89.2%	54.3%	15.0%
Capacity utilisation	47.2%	87.8%	89.1%	87.7%
BAA-AAA Spread	100.0%	100.0%	100.0%	100.0%

Table 3: Proportion of variance attributed to measurement error in the unconstrained model.

Significantly, our model explains much of the variance in nominal GDP, labour supply, and the R&D and consumption shares, suggesting it is capturing well the linkages between research and the business cycle. Indeed, from summing the percentages our model explains (i.e. 100% minus the measurement error share), we see that the model is fully explaining the equivalent of 5.0 variables at business cycle frequencies and 4.2 variables at medium frequencies. Given there are only four shocks given any weight by the estimation procedure (with one of those given a tiny weight), the model is fully explaining more variables than there are driving shocks. Note too that the interpretation of these percentages is somewhat different to the percentages of explained variance given in traditional business cycle analysis. Whereas for us, explaining a high percentage of the variance means that the model's response is preferred by the data to the general measurement error process (i.e. it is a claim about the full covariance structure of the model), the claim in the business cycle literature is really only about the variance of each variable, and covariances across variables or time need not be plausible.

Nonetheless, the model's poor performance along other axes deserves comment. Its difficulties matching inflation rates and nominal interest rates at business cycle frequencies most likely reflect the absence of short run price-rigidity in our model. The model also does spectacularly poorly in matching the variance of the labour share. However, we will see below that the labour share our model generates has a similar correlation structure with GDP across frequencies as we observe in the data. This suggests that the pro-cyclical movements in mark-ups generated by our model are too small relative to those in the data, which is not too surprising given that at the estimated parameters, there are 6.47 firms even in patent protected industries, meaning even these industries will have quite low mark-ups. Now, certainly our model can generate larger swings in mark-ups over the cycle with alternative parameterizations, but these parameterizations will imply even larger movements in productivity. One way of dampening down these excessively large movements in productivity would be to consider the non-asymptotic version of our model in which it takes several periods for new firms to catch-up to the frontier. Producing a non-asymptotic version of the model that may be feasibly simulated is left for future work.

As an additional test of the model, we re-estimated the model under the constraint that $q = 0$.[52] Doing this reduced the log posterior density by 14.14[53] which with flat priors would mean we could reject the null of the validity of the $q = 0$ constraint at even 0.01% significance. Now, with $q = 0$, patent protection is indefinite, so there cannot be any of the movement in the share of patent protected industries that was previously seen to drive our model's behaviour, and so the model collapses to a medium scale variant of the Jaimovich (2007) model. Hence, our ability to reject the null of $q = 0$ provides strong evidence of the macroeconomic importance of our key mechanism.

We can further statistically test our model by looking for evidence of misspecification. Under the null hypothesis of no misspecification, the estimated shock residuals should be NIID(0,1). In the online appendix (Holden 2013b sec. 7), we report the p-values of LM tests for the presence of auto-correlation in these residuals. We are unable to reject the null of no auto-correlation (at 1%) for six shocks, including the depreciation shock, the population shock and the monetary policy shock. Given these last three shocks together explain more than 50% of the non-measurement-error variance in ten out of the twelve variables (including output and prices), and given that the estimated shocks from DSGE models tend to be highly auto-correlated, this is a further strong vindication of our model.

A final natural test of the model is its ability to replicate the empirical results of section 1.

By varying q and calculating the medium-frequency variance share using the spectral density implied by the transition matrices, we can verify that the model does indeed predict that increasing the duration of patent-protection increases the share of variance attributable to medium-frequency cycles, in line with the evidence of section 1.3. As expected, increasing q from its estimated value (i.e.

---

[52] And without any constraint on the effect of increasing q on the share of medium-frequency variance.
[53] The log posterior density decreased from 13462.01 to 13447.86.

shortening patent protection) results in a smooth decline in the medium-frequency variance share.[54] A plot of this may be found in the online appendix (Holden 2013b sec. 8). The mechanism here is that with longer patent-protection (i.e. a smaller value of q), following a boom in invention the share of patent-protected industries will be above its steady-state level for longer, implying that productivity too will be above trend for longer.

Additionally, output per capita is near trend stationary in our model, just as we found in the data in section 1.1. By construction, there is only one potential source of non-stationarity in output per capita: the non-stationarity of A_t^*. However, the standard deviation of g_{A^*} is only 0.00186%, meaning that A_t^* is very close to being deterministic. Thus in the long run in our model, log-output will always return to its original linear trend. The low variance of g_{A^*} comes from the fact that fluctuations in the number of industries and the number of firms absorb almost all demand variations in the long and short runs, meaning each individual firm faces roughly constant incentives to perform research. Despite this long-run return to trend however, our model still generates sizeable medium-frequency cycles, as may be seen in the impulse responses shown in the next section.

[54] Although our estimation constrained the model to have an initial decrease in medium-frequency variance share, it may be seen for the graph that this constraint does not bind, since its left hand axis corresponds to the estimated value of q.

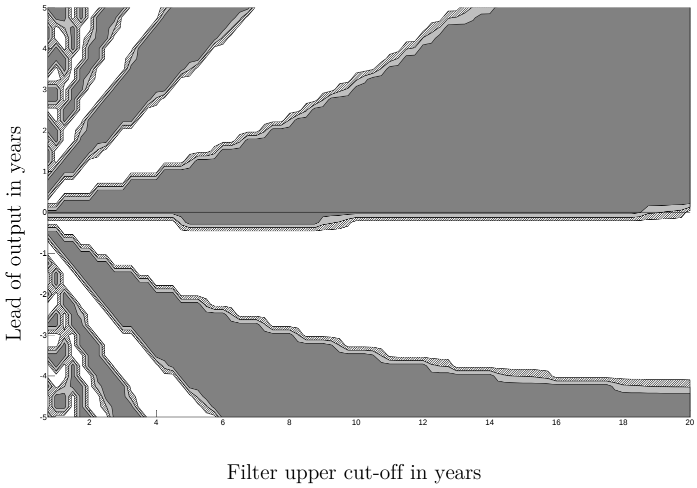

Figure 3: The cross correlation of model output and mark-ups, as a function of filter cut-off.
(Dark gray is a significantly positive correlation (at 5%), light grey is a positive but insignificant one, cross-hatched is a negative but insignificant one and white is a significantly negative one.)

Finally, although our estimation procedure guarantees that mark-ups (inverse labour-shares) are pro-cyclical when the model's output is filtered with a cut-off of one, five or eleven years and counter-cyclical when the output is filtered by a filter with a cut-off of twenty years, the estimation procedure does not impose anything about the cross-correlation of output and mark-ups at lags or leads. In Figure 3, we replicate Figure 1 on simulated data from the estimated model. Immediately, we see that only the bound at twenty years is actually binding, meaning our model is not being contorted in order to produce pro-cyclicality at high to medium frequencies. Indeed, the similarity between the figures is remarkable. Just as in reality, the model predicts that mark-ups are pro-cyclical for small lags or leads, unless the data is filtered with a very low frequency lower cut-off. Again, as in reality, the model predicts that mark-ups are positively correlated with leads of output, and negatively correlated with its lags.

This pro-cyclicality is not driven by sticky wages. Indeed, with fully flexible wages we get pro-cyclicality whatever our filter cut-off. Instead, the pro-cyclicality is driven by the fact that increases in the proportion of industries producing patent protected products both increase mark-ups and productivity. This also explains why mark-ups should lead output; the increase in mark-ups is rapid, however due to the assorted real rigidities in our model, the increase in output will only occur gradually.

3.4. Impulse responses

In Figure 4, we present the impulse responses to the four key driving shocks. Each graph is given in terms of per cent deviations from the value the variable would have taken had the shock never arrived, and the horizontal axis shows time in years, though this is a quarterly model. For no shocks was there an asymmetric positive and negative response, so the lower bound on invention is irrelevant. Each shock is in a different column, and the key response variables are in rows. Solid lines show the response with the estimated degree of wage stickiness, dashed lines show responses under flexible wages.

To show the magnitude of the effects of these shocks on productivity, we include the implied Solow residual[55] in the third row. Our chief driving shock, that to depreciation, has both a direct effect on the Solow residual through reduced utilisation, and an indirect one through the consequent reduction in invention and transfer away from new, highly productive industries, both of which operate in the same direction initially. However, the indirect effect far outlasts the direct one, with aggregate productivity still negative nearly twenty years after the original shock. It then slightly overshoots due to our model's real rigidities, producing a medium frequency cycle in productivity.

[55] The Solow residual is given by $\frac{Y_t}{K_{t-1}^{\alpha_P} L_t^{1-\alpha_P}} = \frac{\hat{Y}_t A_t^{\frac{1}{1-\iota_P}} A_t^{*(1-\alpha_P)\xi_L}}{\widehat{K}_{t-1}^{\alpha_P} \widehat{L}_t^{1-\alpha_P}}$ in the notation of the online appendix (Holden 2013b sec. 5).

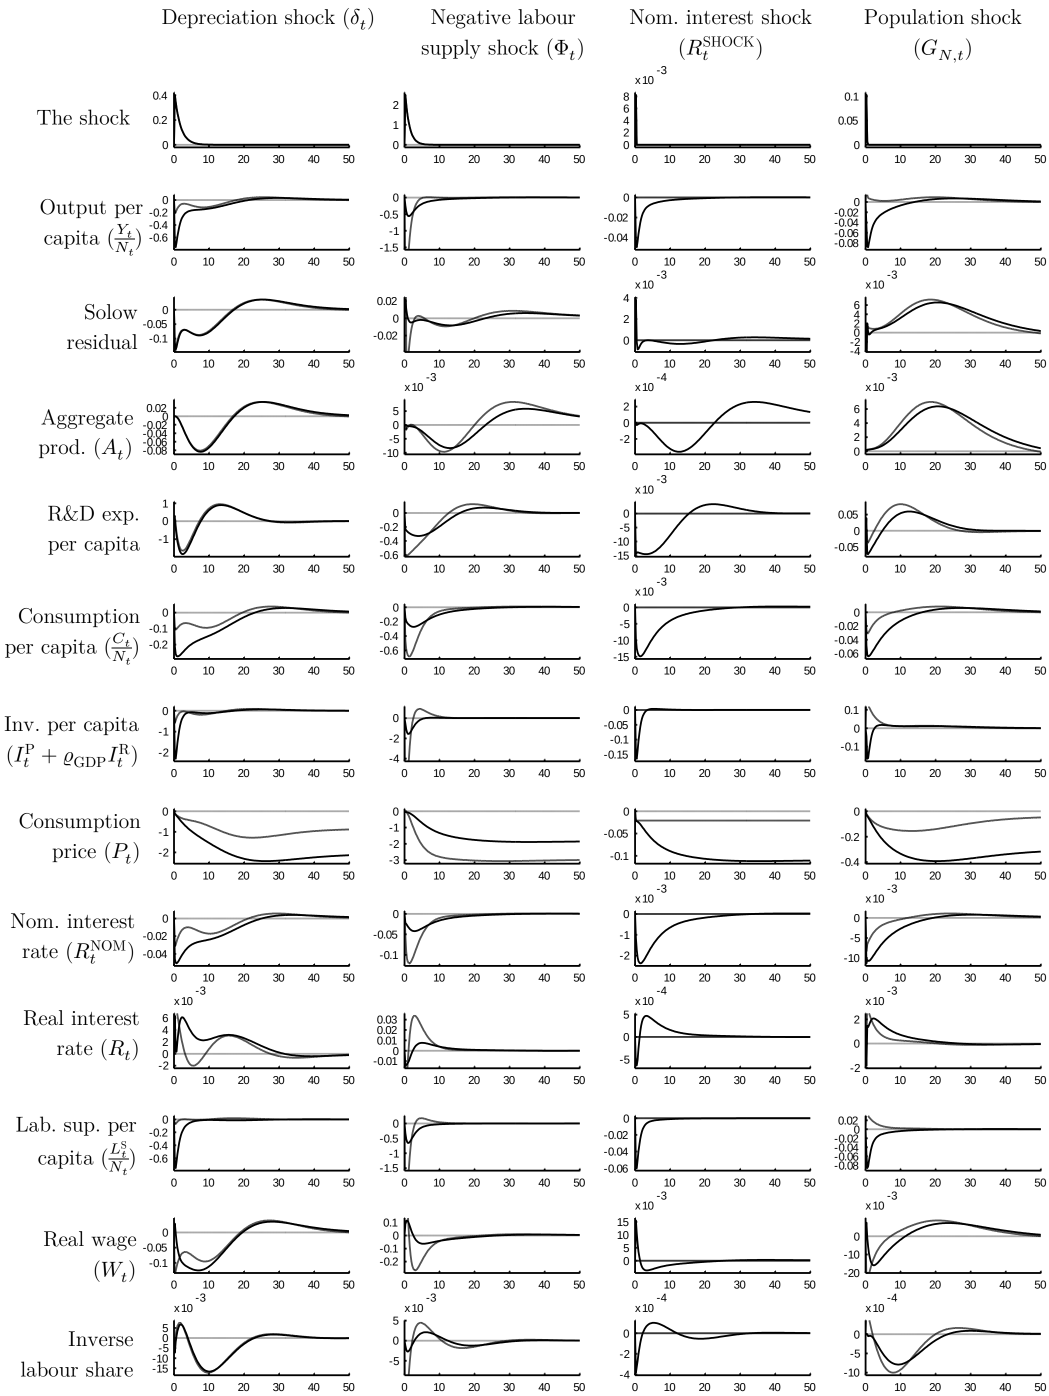

Figure 4: Impulse responses from the core model.

(Vertical axes are in percent, horizontal axes are in years. Solid lines are with nominal wage rigidity, dashed lines are with flexible wages.)

In fact, thanks to the model's endogenous growth component, the Solow residual moves following each of the four shocks, so in a sense all shocks are TFP shocks. Most interesting of these is our monetary policy shock, as a large medium term impact of monetary policy on productivity would substantially alter prescriptions for optimal monetary policy. However, at the estimated parameters the movement in productivity following a monetary policy shock is miniscule, so (perhaps unsurprisingly) the medium term impacts of monetary policy on productivity are not something that policy makers need to factor in to their decisions.

4. Conclusion

We have presented four new stylized facts on the nature of medium frequency cycles. All of them point to the deficiencies of the prior dynamic endogenous growth literature, and reinforce the need for a model like the one presented in this paper and its companion (Holden 2013a). We went on to show that our full model is capable of accurately matching both these stylized facts and others, providing a statistically significant improvement in model fit.

We showed that all shocks lead to changes in the rate of product invention that have significant consequences for aggregate productivity and mark-ups at medium-frequency, due to fluctuations in the proportion of industries that are producing patent-protected products. Our model's propagation mechanisms thus lend persistence to all shocks, not just shocks to the invention or research process.

The fact we are able to combine a plausible growth model with a business cycle model also enables us to get much tighter estimates of the strength of externalities (for example) than is possible from traditional growth models, since these parameters have an impact on the dynamics as well as on the long run growth rate. This will enable the testing of hypotheses about the mechanics of endogenous growth that were previously near impossible to test.

5. References

Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59 (3) (May 1): 817–858.

Balasubramanian, Natarajan, and Jagadeesh Sivadasan. 2011. "What Happens When Firms Patent? New Evidence from U.S. Economic Census Data." *Review of Economics and Statistics* 93 (1) (July 7): 126–146.

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk. 2010. *Some Evidence on the Importance of Sticky Wages*. NBER Working Papers. National Bureau of Economic Research, Inc.

Bartelsman, Eric Jan, George van Leeuwen, Henry Nieuwenhuijsen, and Kees Zeelenberg. 1996. "R&D and Productivity Growth: Evidence from Firm-level Data for the Netherlands." *Netherlands Official Statistics* 11 (3): 52–69.

Boulhol, Hervé. 2007. "The Convergence of Price–cost Margins." *Open Economies Review* 19 (2) (May): 221–240.

Broda, Christian, and David E Weinstein. 2010. "Product Creation and Destruction: Evidence and Price Implications." *American Economic Review* 100 (3) (June): 691–723.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–398.

Candès, Emmanuel, Michael Wakin, and Stephen Boyd. 2008. "Enhancing Sparsity by Reweighted ℓ_1 Minimization." *Journal of Fourier Analysis and Applications* 14 (5) (December 1): 877–905.

Canova, Fabio. 2009. "Bridging Cyclical DSGE Models and the Raw Data."

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2009. "New Keynesian Models: Not Yet Useful for Policy Analysis." *American Economic Journal: Macroeconomics* 1 (1): 242–66.

Cheung, Yin-Wong, and Kon S. Lai. 1995. "Lag Order and Critical Values of the Augmented Dickey-Fuller Test." *Journal of Business & Economic Statistics* 13 (3) (July 1): 277–280.

Christiano L., and Fitzgerald T. 2003. “The Band Pass Filter.” *International Economic Review* 44 (May): 435–465.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy* 113 (1): 1–45.

Cochrane, John H. 1988. “How Big Is the Random Walk in GNP?” *Journal of Political Economy* 96 (5) (October 1): 893–920.

Comin, Diego, and Mark Gertler. 2006. “Medium-Term Business Cycles.” *American Economic Review* 96 (3) (June): 523–551.

Dixit, Avinash K., and Joseph E. Stiglitz. 1977. “Monopolistic Competition and Optimum Product Diversity.” *The American Economic Review* 67 (3) (June): 297–308.

Doornik, Jurgen A, and Henrik Hansen. 2008. “An Omnibus Test for Univariate and Multivariate Normality*.” *Oxford Bulletin of Economics and Statistics* 70 (December 1): 927–939.

Doraszelski, Ulrich, and Jordi Jaumandreu. 2008. *R&D and Productivity: Estimating Production Functions When Productivity Is Endogenous*. CEPR Discussion Papers. C.E.P.R. Discussion Papers.

Dueker, Michael, Andreas Fischer, and Robert Dittmar. 2007. “Stochastic Capital Depreciation and the Co-movement of Hours and Productivity.” *The B.E. Journal of Macroeconomics* 0 (3): 6.

Eaton, Jonathan, and Samuel Kortum. 1999. “International Technology Diffusion: Theory and Measurement.” *International Economic Review* 40 (3) (August 1): 537–570.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. 2000. “Optimal Monetary Policy with Staggered Wage and Price Contracts.” *Journal of Monetary Economics* 46 (2) (October): 281–313.

Ethier, Wilfred J. 1982. “National and International Returns to Scale in the Modern Theory of International Trade.” *The American Economic Review* 72 (3) (June): 389–405.

Forni, Mario, and Sergio Paba. 2002. “Spillovers and the Growth of Local Industries.” *The Journal of Industrial Economics* 50 (2) (June 1): 151–171.

Francois, Patrick, and Huw Lloyd-Ellis. 2008. “Implementation Cycles, Investment, and Growth.” *International Economic Review* 49 (3) (August 1): 901–942.

———. 2009. “Schumpeterian Cycles with Pro-cyclical R&D.” *Review of Economic Dynamics* 12 (4) (October): 567–591.

Furlanetto, Francesco, and Martin Seneca. 2011. *New Perspectives on Depreciation Shocks as a Source of Business Cycle Fluctuations*. Working Paper. Norges Bank.

Ginarte, Juan C., and Walter G. Park. 1997. “Determinants of Patent Rights: A Cross-national Study.” *Research Policy* 26 (3) (October): 283–301.

Granger, C. W. J. 1969. “Investigating Causal Relations by Econometric Models and Cross-spectral Methods.” *Econometrica* 37 (3) (August): 424.

Griffith, Rachel, Rupert Harrison, and John van Reenen. 2006. “How Special Is the Special Relationship? Using the Impact of U.S. R&D Spillovers on U.K. Firms as a Test of Technology Sourcing.” *The American Economic Review* 96 (5): pp. 1859–1875.

Griliches, Zvi. 1998. *The Search for R&D Spillovers*. NBER Chapters. National Bureau of Economic Research, Inc.

Hall, Bronwyn H., and Jacques Mairesse. 1995. “Exploring the Relationship Between R&D and Productivity in French Manufacturing Firms.” *Journal of Econometrics* 65 (1) (January): 263–293.

Hall, Bronwyn H., Jacques Mairesse, and Pierre Mohnen. 2010. *Measuring the Returns to R&D*. UNU-MERIT Working Paper Series. United Nations University, Maastricht Economic and social Research and training centre on Innovation and Technology.

Hansen, N., A.S.P. Niederberger, L. Guzzella, and P. Koumoutsakos. 2009. “A Method for Handling Uncertainty in Evolutionary Optimization With an Application to Feedback Control of Combustion.” *Evolutionary Computation, IEEE Transactions On* 13 (1): 180–197.

Head, Allen, Lucy Qian Liu, Guido Menzio, and Randall Wright. 2011. *Sticky Prices: A New Monetarist Approach*. NBER Working Paper. National Bureau of Economic Research, Inc.

Heston, Alan, Robert Summers, and Bettina Aten. 2011. "Penn World Table Version 7.0". Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

Hodrick, Robert J., and Edward C. Prescott. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking* 29 (1) (February): 1–16.

Holden, Tom. 2013a. "Reconciling Near Trend-stationary Growth with Medium-frequency Cycles."

———. 2013b. "Online Appendices to 'Reconciling Near Trend-stationary Growth with Medium-frequency Cycles' and 'Data Consistent Modelling of Medium-frequency Cycles and Their Origins'."

Hurvich, Clifford M. 1985. "Data-Driven Choice of a Spectrum Estimate: Extending the Applicability of Cross-Validation Methods." *Journal of the American Statistical Association* 80 (392) (December 1): 933–940.

Jaimovich, Nir. 2007. "Firm Dynamics and Markup Variations: Implications for Sunspot Equilibria and Endogenous Economic Fluctuations." *Journal of Economic Theory* 137 (1) (November): 300–325.

Jaimovich, Nir, and Sergio Rebelo. 2008. "News and Business Cycles in Open Economies." *Journal of Money, Credit and Banking* 40 (8): 1699–1711.

———. 2009. "Can News About the Future Drive the Business Cycle?" *American Economic Review* 99 (4) (September): 1097–1118.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2011. "Investment Shocks and the Relative Price of Investment." *Review of Economic Dynamics* 14 (1) (January): 102–121.

Keane, Michael, and Richard Rogerson. 2012. "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom." *Journal of Economic Literature* 50 (2) (June): 464–476.

Kiyotaki, Nobuhiro, and John Moore. 1997. “Credit Cycles.” *Journal of Political Economy* 105 (2) (April 1): 211–248.

Klenow, Peter J., and Andrés Rodríguez-Clare. 2005. “Chapter 11 Externalities and Growth.” In *Handbook of Economic Growth*, Volume 1, Part A:817–861. Elsevier.

Lee, Anthony, Francois Caron, Arnaud Doucet, and Chris Holmes. 2010. “A Hierarchical Bayesian Framework for Constructing Sparsity-inducing Priors.” *arXiv:1009.1914* (September 9).

Levy, Daniel, and Hashem Dezhbakhsh. 2003. “International Evidence on Output Fluctuation and Shock Persistence.” *Journal of Monetary Economics* 50 (7) (October): 1499–1530.

Liu, Zheng, Daniel F Waggoner, and Tao Zha. 2011. “Sources of Macroeconomic Fluctuations: A Regime-switching DSGE Approach.” *Quantitative Economics* 2 (2) (July 1): 251–301.

Los, Bart, and Bart Verspagen. 2000. “R&D Spillovers and Productivity: Evidence from U.S. Manufacturing Microdata.” *Empirical Economics* 25 (1): 127–148.

Nekarda, Christopher, and Valerie Ramey. 2010. “The Cyclical Behavior of the Price-Cost Markup.”

Peterman, W. B. 2011. “Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity.”

La Porta, Rafael, Florencio Lopez-de-Silanes, and Andrei Shleifer. 2008. “The Economic Consequences of Legal Origins.” *Journal of Economic Literature* 46 (2). Journal of Economic Literature: 285–332.

Ramsey, J. B. 1969. “Tests for Specification Errors in Classical Linear Least-Squares Regression Analysis.” *Journal of the Royal Statistical Society. Series B (Methodological)* 31 (2) (January 1): 350–371.

Said, Said E., and David A. Dickey. 1984. “Testing for Unit Roots in Autoregressive-Moving Average Models of Unknown Order.” *Biometrika* 71 (3) (December 1): 599–607.

Schmitt-Grohé, Stephanie, and Martín Uribe. 2011. "Business Cycles with a Common Trend in Neutral and Investment-specific Productivity." *Review of Economic Dynamics* 14 (1) (January): 122–135.

Siliverstovs, Boriss, and D'Artis Kancs. 2012. *R&D and Non-Linear Productivity Growth of Heterogeneous Firms*. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Tutino, Antonella, and Anton Cheremukhin. 2012. *Asymmetric Firm Dynamics Under Rational Inattention*. 2012 Meeting Paper. Society for Economic Dynamics.

White, Halbert. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica* 48 (4) (May 1): 817–838.

Woodford, Michael. 2001. "The Taylor Rule and Optimal Monetary Policy." *The American Economic Review* 91 (2) (May 1): 232–237.