

Three essays in dynamic macroeconomics

*A thesis submitted for the degree of
Doctor of Philosophy, at the University of Oxford,
Trinity Term 2012.*

Thomas Holden
Balliol College
University of Oxford

*Viva passed without corrections,
on the 22nd of May, 2013.*

With thanks to my supervisor, Simon Wren-Lewis,
and my examiners, Campbell Leith and
Rick van der Ploeg.

Three essays in dynamic macroeconomics

A thesis submitted for the degree of Doctor of Philosophy, at the University of Oxford, Trinity Term 2012.

Tom Holden¹, Balliol College, University of Oxford

Abstract: *This thesis presents three papers within the field of dynamic macroeconomics.*

The first paper, entitled “Medium-frequency cycles and the remarkable near trend-stationarity of output”, presents a dynamic stochastic general equilibrium model with endogenous growth, capable of reconciling the observed large medium-frequency fluctuations in output, with its long run (near) trend-stationarity. This requires a model in which standard business cycle shocks lead to highly persistent movements around trend, without significantly altering the trend itself. The robustness of the trend also requires that scale effects are eliminated both in the long and short runs. In an estimated version of the model, a financial-type shock to the stock of ideas emerges as the key driver of the medium frequency cycle.

The second paper, entitled “Learning from learners”, is an intervention into two long running debates: the first, on whether learnability may be used to rule out explosive paths for inflation in New Keynesian models, and the second, into whether Taylor rule parameters may be identified from observing the data. We find that in an economy populated with traditional macroeconomic learners, Taylor rule parameters can always be identified by sophisticated econometric techniques. Furthermore, when all agents in the economy use such sophisticated techniques, stationary sunspot solutions are readily learnable, and there is no guarantee of convergence to a stationary solution even in the “determinate” case. This implies that learnability cannot be used for equilibrium selection.

Finally, in the third paper, “Efficient simulation of DSGE models with inequality constraints” (joint with Michael Paetz), we present a new algorithm for the simulation of models subject to inequality constraints, such as the zero lower bound on nominal interest rates. Our algorithm is shown to deliver higher accuracy than all other non-global algorithms, and leading speed. We go on to provide a number of applications of our algorithm.

Effective word count: *Approx. 65,000*

Actual word count: *Approx. 45,000*

Supervisor: *Simon Wren-Lewis*

¹ Contact address: School of Economics, FBEL, University of Surrey, Guildford, GU2 7XH, England
Contact email: thomas.holden@gmail.com Web-site: <http://www.tholden.org/>

The author would particularly like to thank his supervisor, Simon Wren-Lewis. Financial support provided to the author by the ESRC is also greatly appreciated. Further acknowledgements are included in the title page to each chapter.

Table of Contents

Foreword	Page 3
References	Page 5
Chapter 1:	Page 7
Medium-frequency cycles and the remarkable near trend-stationarity of output	
1. Introduction	Page 8
2. Empirics	Page 13
3. The model	Page 19
4. Extended model and empirical tests	Page 42
5. Conclusion	Page 65
6. References	Page 66
7. Appendices	Page 71
Chapter 2:	Page 93
Learning from learners	
1. Introduction	Page 94
2. FREE solutions	Page 96
3. Learning (and identifying) from unsophisticated learners	Page 101
4. Learning from sophisticated learners	Page 119
5. Conclusion	Page 130
6. References	Page 130
7. Appendices	Page 132
Chapter 3:	Page 147
Efficient simulation of DSGE models with inequality constraints	
1. Introduction	Page 148
2. The numerical method	Page 149
3. Accuracy	Page 160
4. Sample applications	Page 163
5. Conclusion	Page 170
6. References	Page 170
7. Appendices	Page 172

Foreword

In the years since the seminal paper of Kydland and Prescott (1982), the field of dynamic macroeconomics has exploded. Estimated dynamic stochastic general equilibrium (DSGE) models are now used in all of the world's major central banks, with models often derived from the medium-scale New Keynesian DSGE model of Smets and Wouters (2003). In this thesis, I present three contributions to this literature.

The first paper, "Medium-frequency cycles and the remarkable near trend-stationarity of output" arose out of a concern with the source of fluctuations in productivity in DSGE models, with early versions of the paper attempting to endogenise aggregate productivity shocks from the firm-level upwards (Holden 2010a). However, it soon became clear that firm-level shocks could not possibly explain the large medium frequency fluctuations in output that we observe in the data. Consequently, in the included version of the paper, the movement in productivity stems rather from changes in the proportion of industries that are highly productive, which, in my paper, are those producing patent protected products.

In a boom, inventing a new product is more profitable, leading to a higher proportion of relatively new industries, and so to a higher proportion producing patent protected products. The rents firms must pay to patent holders in our model lead entry into these industries to be more expensive, resulting in higher mark-ups, and consequently greater incentives for improvements in production processes. Thus industries producing patent protected products are endogenously more productive, so booms will lead to increases in aggregate productivity, further amplifying the expansion.

I present a wide range of empirical evidence to support this story, and conclude with a demonstration of the performance of an estimated version of the model.

The second paper included here, "Learning from learners", is a distant ancestor of earlier, ultimately unsuccessful, work on fully rational macroeconomic learning (Holden 2008). Here

I step back slightly from the ambitious goal of full rationality, to produce a learning algorithm that is merely close enough to rationality that the users of the algorithm would not detect the misspecification in their econometric method. With traditional macroeconomic learning, the misspecification is readily detectable, meaning that competent econometricians would switch to a more sophisticated method. I argue that it is this misspecification that leads the traditional macroeconomic learning literature to find that some equilibria are not learnable, and I show that with a more sophisticated method, in fact almost all equilibria are learnable. As a result we agree with Cochrane (2011) that there is no justification for ruling out solutions to New Keynesian models in which inflation explodes.

I also show that if everyone is using the traditional macroeconomic learning method for some reason, then the expectational errors they will make enable the identification of Taylor rule parameters, contrary to the claims of Cochrane (2009).

The last paper I present here is a development of an idea for solving bounded DSGE models, first described in an appendix to an early version of the “Medium frequency cycles” paper (Holden 2010b). In that model, there is a lower bound on the growth of the stock of products, since products cannot be un-invented, which necessitated the creation of an algorithm for solving such models. Together with my co-author, Michael Paetz, the algorithm has been developed into the paper “Efficient simulation of DSGE models with inequality constraints” presented here.

We show that our algorithm beats all others in terms of speed, and only slow, poorly scaling, global methods beat it in terms of accuracy. We go on to apply it to a range of dynamic macroeconomic models, including to the zero lower bound on nominal interest rates in the Smets and Wouters (2003) model.

I hope all of these papers will have a lasting impact on the field of dynamic macroeconomics.

References

- Cochrane, John H. 2009. 'Can Learnability Save new-Keynesian Models?' *Journal of Monetary Economics* 56 (8) (November): 1109–1113.
doi:10.1016/j.jmoneco.2009.10.012.
- . 2011. 'Determinacy and Identification with Taylor Rules'. *Journal of Political Economy* 119 (3): pp. 565–615.
- Holden, Tom. 2008. *Rational Macroeconomic Learning in Linear Expectational Models*. MPRA Paper. University Library of Munich, Germany.
<http://ideas.repec.org/p/pra/mprapa/10872.html>.
- . 2010a. 'Business Cycles and Business Stealing: A Model of Incremental Endogenous Growth'.
- . 2010b. *Products, Patents and Productivity Persistence: A DSGE Model of Endogenous Growth*. Economics Series Working Paper. University of Oxford, Department of Economics. <http://ideas.repec.org/p/oxf/wpaper/512.html>.
- Kydland, Finn E, and Edward C Prescott. 1982. 'Time to Build and Aggregate Fluctuations'. *Econometrica* 50 (6) (November): 1345–70.
- Smets, F., and R. Wouters. 2003. 'An Estimated Dynamic Stochastic General Equilibrium Model for the Euro Area'. *Journal of the European Economic Association* 1 (5).

Chapter 1: Medium-frequency cycles and the remarkable near trend-stationarity of output

Tom Holden¹, Balliol College, University of Oxford

Abstract: *This paper builds a dynamic stochastic general equilibrium (DSGE) model of endogenous growth that generates large medium-frequency cycles while robustly matching the near trend-stationary path of observed output. This requires a model in which standard business cycle shocks lead to highly persistent movements around trend, without significantly altering the trend itself. The robustness of the trend also requires that we eliminate the scale effects and knife edge assumptions that plague most growth models. In our model, when products go out of patent protection, the rush of entry into their production destroys incentives for process improvements. Consequently, old production processes are enshrined in industries producing non-protected products, and shocks that affect invention rates change the proportion of industries with advanced technologies. In an estimated version of our model, 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, scale effects*

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.com Web-site: <http://www.tholden.org/>

The author would particularly like to thank his supervisor, Simon Wren-Lewis. Further thanks are due to Eric Bartelsman, Paul Beaudry, Pedro Bento, George Bitsakakis, Beatriz de Blas, Ana-Maria Dumitru, Martin Ellison, Fabio Ghironi, Tim Kehoe, Paul Levine, Omar Licandro, Ioana Moldovan, Peter Neary, Joaquim Oliveira Martins, Christopher Nekarda, Walter Park, Pietro Peretto, James Reade, Diego Restuccia, Nicole Tabasso, David Vines 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.

1. Introduction

Viewed from a distance, a log-plot of the last one hundred years of US GDP looks very near linear. However, closer inspection reveals large medium frequency fluctuations around this linear trend. Generating this combination of remarkably near trend-stationary long run growth, and large cycles around the trend, is a challenge for traditional models of endogenous growth. The near linear trend requires scale effects to be removed not just in the long run, but in the shorter run as well. Models that remove these scale effects via knife-edge assumptions will usually fail this test, as temporary business cycle shocks will knock the model away from perfectly removing the scale effect, leading to a permanent break in the trend of the GDP. Equally, models that remove scale effects via new product creation will tend to produce such trend breaks in GDP if the stock of new products can only respond slowly following a shock. On the other hand, if the stock of products can adjust instantly following a shock, then, (in standard models) there would be no movement in productivity at all, let alone the large, persistent medium frequency cycles that Comin and Gertler (2006) document in the data, and that may be seen in our Figure 1 below. In this paper, we present a mechanism capable of reconciling this apparently contradictory low and medium frequency behaviour of output, while also matching the cyclicity of mark-ups: the key determinant of research and invention decisions.

Our story is as follows. The returns to inventing a new product are higher in a boom due to the higher demand. As a result, during periods of expansion, the rate of creation of new products increases, in line with the evidence of Broda and Weinstein (2010). Due to a first mover advantage, patent protection, or reverse-engineering difficulties, the inventors of these new products will be able to extract rents from them, increasing the costs manufacturing firms face if they wish to produce the new product. These higher costs lead to lower competition in new industries, increasing mark-ups and thus increasing firms' incentives to perform the R&D necessary to catch-up with and surpass the frontier, for

Medium-frequency cycles and the remarkable near trend-stationarity of output.

basically Schumpeterian reasons. Consequently, the higher proportion of industries that are relatively new in a boom will lead to higher aggregate productivity, lower dispersion of both productivity levels and growth rates, as well as higher mark-ups. Since the length of time for which inventors can extract rents will be determined by the effective duration of patent-protection, this effect will naturally work at medium frequencies. However, since we allow both for the creation of new industries (producing new products) and for a varying number of firms within each industry, even in the short-run the demand faced by any given firm will be roughly constant, meaning that our model will not produce large deviations from linear growth.




Figure 1: The results of modelling quarterly log real US GDP per capita as a sum of a random walk, an AR(2) process and an idiosyncratic shock. The solid line in the second graph is a crude representation of the medium-frequency cycle.

Evidence for the pro-cyclicality of TFP has been presented by [Bils \(1998\)](#) and [Campbell \(1998\)](#) amongst others, with [Comin and Gertler \(2006\)](#) showing that the evidence is particularly clear at medium-frequencies. The counter-cyclicality of productivity dispersion has been shown by [Kehrig \(2011\)](#), with evidence on the counter-cyclicality of the dispersion of productivity growth rates provided by e.g. [Eisfeldt and Rampini \(2006\)](#) and [Bachmann and](#)

Bayer (2009). Evidence for the pro-cyclicality of aggregate mark-ups has been presented by Boulhol (2007) and Nekarda and Ramey (2010). Nekarda and Ramey also show that mark-ups lead output at business-cycle frequencies, we will present further evidence in section 2 below that this relationship continues to hold at medium-frequencies. Boulhol (2007) also shows that although aggregate mark-ups are pro-cyclical, the mark-ups in any particular industry tend to be counter-cyclical. This apparent contradiction will be readily explained in our model since the increase in competition in any particular industry will lead to a decline in mark-ups in that industry (much as in the models of Bilbiie, Ghironi, and Melitz (2012) and Jaimovich (2007)), despite the fact that aggregate mark-ups have increased due to the greater proportion of industries with relatively high mark-ups. Formal evidence on the small size of the unit root in output (i.e. its near trend stationarity) was presented by Cochrane (1988), and we will present further evidence in the next section that GDP returns to trend at long lags.

Direct evidence for the importance of our mechanism comes from a number of sources. Balasubramanian and Sivadasan (2011) find that firms holding patents have 17% higher TFP levels on average, and additionally find that firms that go from not holding a patent to holding one experience a 7.4% increase in a fixed effects measure of productivity, suggesting that industries producing patent-protected products are indeed significantly more productive. Serrano (2007) finds that although aggregate patenting is only weakly correlated with aggregate TFP, a measure of the number of patents whose ownership is transferred is strongly related to productivity. He argues that there is a great deal of noise in measures of total patent activity, since so many patents are never seriously commercialised. Patent transfers are usually observed though when their purchaser intends to begin exactly such a commercialisation. Thus, patent transfers provide a proxy for the commencement of production of new patented-products, one that is found to be highly pro-cyclical. Finally, we will present new evidence that longer patent protection significantly increases the share of GDP variance attributable to cycles of medium frequency.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

Previous papers have introduced endogenous productivity improvement into business cycle models (e.g. Comin and Gertler (2006), Comin (2009), Comin, Gertler, and Santacreu (2009), Phillips and Wrase (2006), Nuño (2008; 2009; 2011)), or looked at cycles in growth models (e.g. Bental and Peled (1996), Matsuyama (1999), Wälde (2005), Francois and Lloyd-Ellis (2008; 2009), Comin and Mulani (2009)). However, all of these papers have problems with scale effects, either in the long-run, or in the short-run, and thus all of them would predict counter-factually large unit roots in output in the presence of standard DSGE shocks. Furthermore, it is not obvious how these scale effects could be removed without destroying the papers' mechanisms for generating aggregate TFP movements. For example, the papers of Wälde (2005) and Phillips and Wrase (2006) rely on there being a small finite number of sectors. Removing the scale effect would mean allowing this number to grow over time with population, meaning the variance of productivity would rapidly go to zero. Indeed, this happens endogenously in the model of Horii (2011). Many models of endogenous mark-up determination (e.g. Bilbiie, Ghironi, and Melitz (2012) or Jaimovich (2007)) have a similar problem, with the presence of a small finite number of industries being crucial for explaining the observed variance of mark-ups. Indeed, Bilbiie, Ghironi, and Melitz (2011) write that "reconciling an endogenous time-varying markup with stylized growth facts (that imply constant markups and profit shares in the long run) is a challenge to growth theory". By disentangling the margins of firm entry and product creation, we will be able to answer this challenge.

The paper of most relevance to our work is Comin and Gertler (2006), as they made the important contribution of bringing the significance of medium-frequency cycles to the attention of the profession. Additionally, their theoretical model, like ours, stresses the effects of mark-up variations on productivity growth. Unfortunately, however, it counter-factually predicts that increases in mark-ups lead to falls in output, contrary to the empirical evidence of Nekarda and Ramey (2010). Furthermore, its only major sources of productivity persistence are the persistence of the driving mark-up shock, and the counter-factual trend

break in productivity following such a shock. We conclude then, that the literature still lacks a model of productivity capable of explaining both its short run and its long run behaviour.

In section 3, we present a model capable of doing this. In order to remove both the long run and the short run scale effect, as discussed above it will feature a varying number of industries, each of which will contain a varying number of firms. We do not wish to make any exogenous assumptions on the differences between industries producing patented products versus those producing unpatented ones, so in order to match the medium-frequency behaviour of productivity and mark-ups it is important that our model allow endogenous variation in these quantities across industries. Were we to assume free transfer of technologies across industries there would be too little difference in productivity between patent-protected and un-patent-protected industries, and hence we would not be able to generate medium-frequency cycles. Equally, were we to assume technology transfer across industries was impossible then it would be legitimate to inquire whether the difference between these industry types was implausibly large, as perhaps firms in non-protected industries would find it optimal to perform technology transfer even if they did not find it optimal to perform any research. Consequently, in modelling the endogenous productivity in each industry we will allow firms both to perform research, and to perform a costly process of catch-up to the frontier we shall term appropriation. To make clear the strength of the amplification and persistence mechanism presented here, we initially omit capital from the model, and we focus on the impulse responses to non-persistent shocks when we discuss our model's qualitative behaviour in section 3.5. Finally, in section 4, we add a few standard additional features to the model (habits, capital with adjustment costs, variable capacity utilisation, sticky wages, Taylor rule monetary policy) and we show that this model matches the data well at low, medium and high frequencies, with financial-type shocks to the stock of ideas playing the key role in driving medium-frequency fluctuations.

2. Empirics

2.1. The near trend stationarity of output

We begin by presenting evidence that GDP 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 32nd lag of its 32 period moving average. I.e. we run the regression:

$$\begin{aligned}
 y_t = & \mu + \delta t + \phi_1 y_{t-1} + \phi_2 \frac{1}{2} (y_{t-2} + y_{t-3}) \\
 & + \phi_3 \frac{1}{4} (y_{t-4} + y_{t-5} + y_{t-6} + y_{t-7}) + \dots \quad (2.1) \\
 & + \phi_6 \frac{1}{32} (y_{t-32} + \dots + y_{t-63}) + \varepsilon_t.
 \end{aligned}$$

The full results of this regression are given in Table 1. The key facts to note here though are that $\phi_2, \phi_3, \dots, \phi_6$ 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 3.

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 (2.1). 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 (2.1) into Augmented Dickey-Fuller (ADF) form (Said and Dickey 1984), giving:

$$\Delta y_t = \mu + \delta t + \left[\sum_{i=1}^6 \phi_i - 1 \right] y_{t-1} - \phi_2 \frac{1}{2} (2\Delta y_{t-1} + \Delta y_{t-2}) - \dots - \phi_6 \frac{1}{32} (\dots). \quad (2.2)$$

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 experiments² indicate that there is only 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).

² With 2²⁰ replications, where in each case the regression (2.2) 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 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.

2.2. 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-cyclicity of mark-ups they observe is a consequence of behaviour around these frequencies, or whether it is driven by counter-cyclicity 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 2 below answers both of these questions.

Each vertical slice of this plot shows the cross-correlation⁵ of quarterly log output and log mark-ups⁶ when both are filtered by a high pass filter⁷ with a cut-off given by the x-axis's value. (Shaded areas indicate positive correlations, with 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. Additionally, tracing along the lead=0 line we

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-cyclicity.

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.

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.

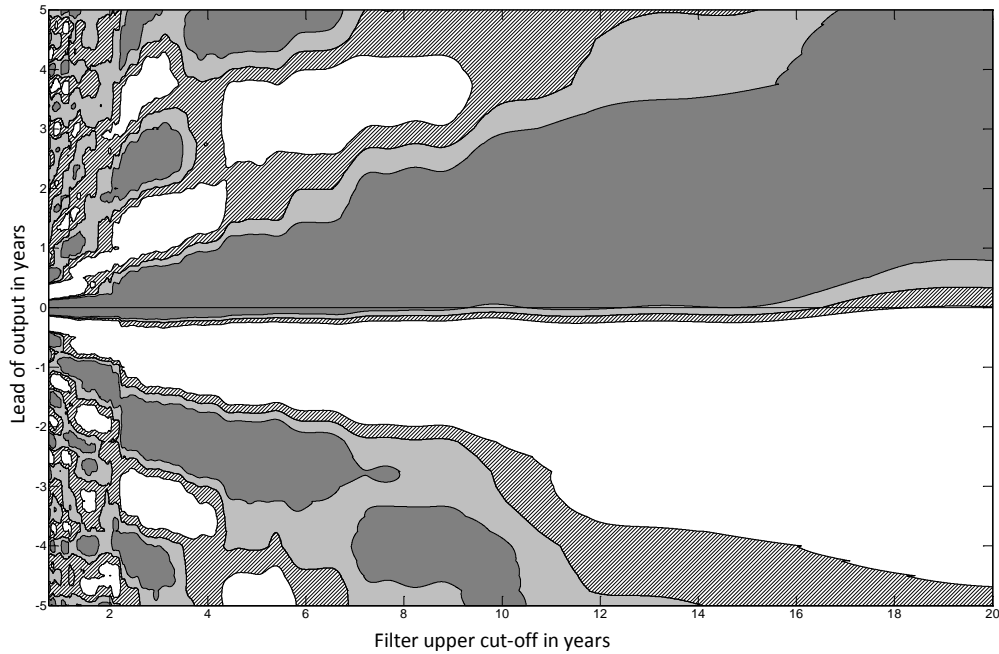


Figure 2: The cross correlation of US output and mark-ups, as a function of filter cut-off. (Dark grey is a significantly positive correlation (at 5%), light grey is a positive but insignificant one, cross-hatched is a negative but insignificant one and white is a significantly negative one.)




Figure 3: The spectral decomposition of US output growth.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

Indeed, from the spectral decomposition⁸ of output growth shown in Figure 3, 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.

2.3. GDP variance

Our model 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. (Full details of the data are given in footnotes to the table.) Patent duration in both 1960 and 2005 has a significantly positive effect (at 5%) on the strength of medium frequency cycles⁹ 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.

⁸ 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.

⁹ 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 8, with the initial filter set to accept period lengths between 2 and 59 years (the length of the largest samples).

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 ¹⁰	-0.0506172 (0.8567)				-0.448554 (0.0810)
French legal origin ¹⁰	-0.0557074 (0.8394)				-0.350747 (0.1653)
German legal origin ¹⁰	-0.151587 (0.6364)				-0.325196 (0.3154)
Log GDP per effective adult ¹¹	0.0715242 (0.3620)	0.0707845 (0.3501)			
GDP per effective adult growth ¹¹	7.39306 (0.1647)	7.24517 (0.1606)			
Socioeconomic Conditions (ICRG) ¹²	-0.224159 (0.0078)	-0.229358 (0.0044)	-0.170029 (0.0107)		
Law and order (ICRG) ¹²	-0.154013 (0.0856)	-0.150749 (0.0818)	-0.148729 (0.0856)		
Logit overall political risk (ICRG) ^{12,13}	0.806772 (0.0013)	0.811630 (0.0006)	0.823980 (0.0003)		
Index of patent duration, 1960 ¹⁴	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 ¹⁴	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-values¹⁵	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.⁹

¹⁰ 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).

¹¹ 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.

¹² 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).

¹³ 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.

¹⁴ Data kindly provided by Walter Park, updated from Ginarte and Park (1997).

¹⁵ Respectively, a normality test (Doornik and Henrik Hansen 2008), the White heteroskedasticity test (White 1980) and the reset test with squares and cubes (Ramsey 1969).

3. The model

Our base model is a standard quarterly real business cycle (RBC) model without capital, augmented by the addition of models of endogenous competition, research, appropriation and invention. The lack of capital means the underlying RBC model has no endogenous propagation mechanism, making clearer the contribution of our additions.

Our model has a continuum of narrow industries, each of which contains finitely many firms producing a unique product. The measure of industries is increased by the invention of new products, which start their life patent-protected. However, we assume that product inventors lack the necessary human capital to produce their product at scale themselves, and so they must licence out their patent to manufacturing firms. The duration of patent-protection is given by a geometric distribution, in line with Serrano's (2010) evidence on the large proportion of patents that are allowed to expire early, perhaps because they are challenged in court or perhaps because another new product is a close substitute. An earlier working-paper version of this model (Holden 2011) considered the fixed duration case, which is somewhat less tractable. Allowing for a distribution of protection lengths also allows us to give a broader interpretation to protection within our model. Even in the absence of patent protection, the combination of contractual agreements such as NDAs, and difficulties in reverse engineering, is likely to enable the inventor of a new product to extract rents for a period.

Our model of endogenous competition within each industry is derived from Jaimovich (2007). We chose the Jaimovich model as it is a small departure from the standard Dixit-Stiglitz (1977) set-up, and leads to some particularly neat expressions. Similar results could be attained with Cournot competition, or the Translog form advocated by Bilbiie, Ghironi, and Melitz (2012). One important departure from the Jaimovich model is that in our model entry decisions take place one period in advance. This is natural as we wish to model research as taking place after entry but before production.

Productivity within a firm is increased by performing research or appropriation. We regard process research as incremental, with regular small changes rather than the unpredictable jumps found in Schumpeterian models (Aghion and Howitt 1992; Wälde 2005; Phillips and Wrase 2006).

Throughout, we assume that only products are patentable,¹⁶ and so by exerting effort firms are able to “appropriate” process innovations from other industries to aid in the production of their own product. This appropriation is costly since technologies for producing other products will not be directly applicable to producing a firm’s own product. We assume that technology transfer within an industry is costless however, due to intra-industry labour flows and the fact that all firms in an industry are producing the same product. This is important for preserving the tractability of the model, as it means that without loss of generality we may think of all firms as just existing for two periods, in the first of which they enter and perform research, and in the second of which they produce.

The broad timing of our model is as follows. At the beginning of period t invention takes place, creating new industries. All holders of current patents (including these new inventors) then decide what level of licence fee to charge. Then, based on these licence fees and the level of overhead costs, firms choose whether to enter each industry. Next, firms perform appropriation, raising their next-period productivity towards that of the frontier, then research, further improving their productivity next period. In period $t + 1$, they then produce using their newly improved production process. Meanwhile, a new batch of firms will be starting this cycle again.

¹⁶ This is at least broadly in line with the law in most developed countries: ideas that are not embedded in a product (in which category we include machines) generally have at most limited patentability. In the U.S., the most recent Supreme court decision found that the following was “a useful and important clue” to the patentability of processes (*Bilski v. Kappos*, 561 U.S. ____ (2010)): “a method claim is surely patentable subject matter if (1) it is tied to a particular machine or apparatus, or (2) it transforms a particular article into a different state or thing” (*In re Bilski*, 545 F.3d 943, 88 U.S.P.Q.2d 1385 (Fed. Cir. 2008)). This “machine or transformation” test was widely believed at the time to have ended the patentability of business processes (The Associated Press 2008), and this position was only slightly softened by *Bilski v. Kappos*.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

We now give the detailed structure of the model.

3.1. Households

There is a unit mass of households, each of which contains N_t members in period t . The representative household maximises:

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s N_{t+s} \Theta_{t+s} \left[\log \frac{C_{t+s}}{N_{t+s}} - \frac{\Phi_{t+s}}{1+\nu} \left(\frac{L_{t+s}^S}{N_{t+s}} \right)^{1+\nu} \right]$$

where C_t is aggregate period t consumption, L_t^S is aggregate period t labour supply, Θ_t is a demand shock, Φ_t is a labour supply shock, β is the discount rate and ν is the inverse of the Frisch elasticity of labour supply to wages, subject to the aggregate budget constraint that $C_t + B_t = L_t^S W_t + B_{t-1} R_{t-1} + \Pi_t$, where B_t is the aggregate number of (zero net supply) bonds bought by households in period t , W_t is the period t wage, R_{t-1} is the period t sale price of a (unit cost) bond bought in period $t-1$, and Π_t is the households' period t dividend income. In the following, where we refer to preference shocks we mean either a shock to Θ_t or a shock to Φ_t . However, both of these shocks may be interpreted as proxying for real changes in the economy that are independent of preferences. For example, Θ_t will capture changes in government consumption demand coming from wars, and Φ_t will pick up changes in marginal tax rates and in the degree of imperfect competition in labour markets.

Let $\beta \Xi_{t+1}$ be the households' period t stochastic discount factor, then the households' first order conditions imply:

$$\Xi_t = \frac{\Theta_t N_t C_{t-1}}{\Theta_{t-1} N_{t-1} C_t}, \quad \Phi_t L_t^{S\nu} = N_t^{1+\nu} \frac{W_t}{C_t}, \quad \beta R_t \mathbb{E}_t[\Xi_{t+1}] = 1.$$

3.2. 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 $(I_{t-1})^{17}$ has been chosen to remove the preference for variety in consumption.¹⁸

Normalising the price of the aggregate consumption good to 1, and writing $P_t(i)$ for the price of the aggregate good from industry i in period t , we have that:

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

Similarly, each industry aggregate good $Y_t(i)$ is produced by a perfectly competitive industry from the intermediate goods $Y_t(i, j)$ for $j \in \{1, \dots, J_{t-1}(i)\}$,¹⁹ 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.

¹⁷ The $t - 1$ subscript here reflects the fact that industries are invented one period before their product is available to consumers.

¹⁸ Incorporating a preference for variety would not change the long-run stability of our model.

¹⁹ Again, the $t - 1$ subscript reflects the fact that firms enter one period before production.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

This means that if $P_t(i, j)$ is the price of intermediate good j in industry i :

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

3.3. Intermediate firms

3.3.1. Pricing

Firm j in industry i has access to the linear production technology $Y_t(i, j) = A_t(i, j)L_t^P(i, j)$ for production in period t . As in Jaimovich (2007), strategic profit maximisation then implies that in a symmetric equilibrium $P_t(i) = P_t(i, j) = (1 + \mu_{t-1}(i)) \frac{W_t}{A_t(i, j)} = (1 + \mu_{t-1}(i)) \frac{W_t}{A_t(i)}$, where $\mu_t(i) := \lambda \frac{\eta J_t(i)}{J_t(i) - (1-\eta)} \in (\eta\lambda, \lambda]$ is the industry i mark-up in period $t + 1$ and $A_t(i) = A_t(i, j)$ is the productivity shared by all firms in industry i in symmetric equilibrium.

From aggregating across industries we have that $W_t = \frac{A_t}{1 + \mu_{t-1}}$ where:

$$\frac{1}{1 + \mu_t} = \left[\frac{1}{I_t} \int_0^{I_t} \left[\frac{1}{1 + \mu_t(i)} \right]^{\frac{1}{\lambda}} di \right]^\lambda$$

determines the aggregate mark-up μ_{t-1} and where:

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

is a measure of the aggregate productivity level.²⁰

²⁰ Due to the non-linear aggregation, it will not generically be the case that aggregate output is aggregate labour input times A_t . However, the aggregation chosen here is the unique one under which aggregate mark-ups are known one period in advance, as industry mark-ups are.

3.3.2. Sunk costs: rents, appropriation and research

Following Jaimovich (2007), we assume that the number of firms in an industry is pinned down by the zero profit condition that equates pre-production costs to production period revenues. Firms borrow in order to cover these upfront costs, which come from four sources.

Firstly, firms must pay a fixed operating cost L^F that covers things such as bureaucracy, human resources, facility maintenance, training, advertising, shop set-up and capital installation/creation. Asymptotically, the level of fixed costs will not matter, but including it here will help in our explanation of the importance of patent protection for long run growth.

Secondly, if the product produced by industry i is currently patent-protected, then firms must pay a rent of $\mathcal{R}_t(i)$ units of the consumption good to the patent-holder for the right to produce in their industry. Since all other sunk costs are paid to labour, for convenience we define $L_t^{\mathcal{R}}(i) := \frac{\mathcal{R}_t(i)}{w_t}$, i.e. the labour amount equivalent in cost to the rent.

Thirdly, firms will expand labour effort on appropriating the previous process innovations of the leading industry. We define the level of the leading technology within industry i by $A_t^*(i) := \max_{j \in \{1, \dots, J_{t-1}(i)\}} A_t(i, j)$ and the level of the best technology anywhere by $A_t^* := \sup_{i \in [0, I_{t-1}]} A_t^*(i)$. Due to free in-industry transfer, even without exerting any appropriation

effort, firms in industry i may start their research from $A_t^*(i)$ in period t . By employing appropriation workers, a firm may raise this level towards A_t^* .

We write $A_t^{**}(i, j)$ for the base from which firm $j \in \{1, \dots, J_t(i)\}$ will start research in period t , and we assume that if firm j employs $L_t^A(i, j)$ units of appropriation labour in period t then:

$$A_t^{**}(i, j) = \left[A_t^*(i)^\tau + (A_t^*{}^\tau - A_t^*(i)^\tau) \frac{A_t^*(i)^{-\zeta^A} Y L_t^A(i, j)}{1 + A_t^*(i)^{-\zeta^A} Y L_t^A(i, j)} \right]^{\frac{1}{\tau}}, \quad (3.1)$$

Medium-frequency cycles and the remarkable near trend-stationarity of output.

where Υ is the productivity of appropriation labour, $\zeta^A > 0$ controls the extent to which appropriation is getting harder over time (due, for example, to the increased complexity of later technologies) and where $\tau > 0$ controls whether the catch-up amount is a proportion of the technology difference in levels ($\tau = 1$), log-levels ($\tau = 0$) or anything in between or beyond. This specification captures the key idea that the further a firm is behind the frontier, the more productive will be appropriation. Allowing for appropriation (and research, and invention) to get harder over time is both realistic, and essential for the tractability of our model, since it will lead our model to have a finite dimensional state vector asymptotically, despite all the heterogeneity across industries.

Fourthly and finally, firms will employ labour in research. If firm $j \in \{1, \dots, J_t(i)\}$ employs $L_t^R(i, j)$ units of research labour in period t , its productivity level in period $t + 1$ will be given by:

$$A_{t+1}(i, j) = A_t^{**}(i, j) \left(1 + \gamma Z_{t+1}(i, j) A_t^{**}(i, j)^{-\zeta^R} \Psi L_t^R(i, j) \right)^{\frac{1}{\gamma}},$$

where Ψ is the productivity of research labour, $\zeta^R > \zeta^A$ controls the extent to which research is getting harder over time, $Z_{t+1}(i, j) > 0$ is a shock representing the luck component of research, and $\gamma > 0$ controls the “parallelizability” of research.²¹ If $\gamma = 1$, research may be perfectly parallelized, so arbitrarily large quantities may be performed within a given period without loss of productivity, but if γ is large, then the productivity of research declines sharply as the firm attempts to pack more into one period. The restriction that $\zeta^R > \zeta^A$ means that the difficulty of research is increasing over time faster than the difficulty of appropriation. This is made because research is very much specific to the industry in which it is being conducted, whereas appropriation is a similar task across all

²¹ Peretto (1999) also looks at research that drives incremental improvements in productivity, and chooses a similar specification. The particular one used here is inspired by Groth, Koch, and Steger (2009).

industries attempting to appropriate the same technology, and hence is more likely to have been standardised, or to benefit from other positive spillovers.

In the following, we will assume that $Z_t(i, j) := Z_t$ so that all firms in all industries receive the same “idea” shock. We make this assumption chiefly for simplicity, but it may be justified by appeal to common inputs to private research, such as university research output or the availability of new tools, or by appeal to in-period labour market movements carrying ideas with them. We will see in the following that allowing for industry-specific shocks has minimal impact on our results, providing there are at least correlations across industries (plausible if they are producing similar products). For concreteness, we assume that $Z_t := \exp(\sigma_Z \epsilon_{Z,t})$, where $\sigma_Z > 0$ and $\epsilon_{Z,t} \sim \text{NIID}(0,1)$.

3.3.3. Research and appropriation effort decisions

Firms are owned by households and so they choose research and appropriation to maximize:

$$\beta \mathbb{E}_t \left[\mathbb{E}_{t+1} \left(P_t(i, j) - \frac{W_{t+1}}{A_{t+1}(i, j)} \right) Y_t(i, j) \right] - [L_t^R(i, j) + L_t^A(i, j) + L_t^R(i) + L^F] W_t$$

It may be shown that, for firms in frontier industries (those for which $A_t^*(i) = A_t^*$), if an equilibrium exists, then it is unique and symmetric within an industry; but we cannot rule out the possibility of asymmetric equilibria more generally.²² However, since the

²² The equilibrium concept we use is that of pure-strategy subgame-perfect *local* Nash equilibria (SPLNE) (i.e. only profitable local deviations are ruled out). We have no reason to believe the equilibrium we find is not in fact a subgame-perfect Nash equilibria (SPNE). Indeed, if there is a pure-strategy symmetric SPNE then it will be identical to the unique pure-strategy symmetric SPLNE that we find. Furthermore, our numerical investigations suggest that at least in steady-state, at our calibrated parameters, the equilibrium we describe is indeed an SPNE. (Code available on request.) However, due to the analytic intractability of the second stage pricing game when productivities are asymmetric, we cannot guarantee that it remains an equilibrium away from the steady-state, or for other possible calibrations. However, SPLNE’s are independently plausible since they only require firms to know the demand curve they face in the local vicinity of an equilibrium, which reduces the riskiness of the experimentation they must perform to find this demand curve (Bonanno 1988). It is arguable that the coordination required to sustain asymmetric equilibria and the computational demands of mixed strategy equilibria render either of these less plausible than our SPLNE.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

coordination requirements of asymmetric equilibria render them somewhat implausible, we restrict ourselves to the unique equilibrium in which all firms within an industry choose the same levels of research and appropriation. Let us then define effective research performed by firms in industry i by $\mathcal{L}_t^R(i) := A_t^{**}(i)^{-\zeta^R} \Psi L_t^R(i, j)$ (valid for any $j \in \{1, \dots, J_{t-1}(i)\}$) and effective appropriation performed by firms in that industry by $\mathcal{L}_t^A(i) := A_t^*(i)^{-\zeta^A} \Upsilon L_t^A(i, j)$ (again, valid for any $j \in \{1, \dots, J_{t-1}(i)\}$).

Providing $\frac{1}{\mu_t(i)} < \min\{\gamma, \tau\}$, $\gamma > \zeta^R$ and $\lambda < 1$ (for the second order conditions²³ and for uniqueness), combining the first order and free entry conditions then gives us that, in the limit as $\sigma_Z \rightarrow 0$:²⁴

$$\mathcal{L}_t^R(i) = \max \left\{ 0, \frac{d_t(i) A_t^{**}(i)^{-\zeta^R} \Psi (L_t^A(i, j) + L_t^R(i) + L_t^F) - \mu_t(i)}{\gamma \mu_t(i) - d_t(i)} \right\} \quad (3.2)$$

and:
$$\mathcal{L}_t^A(i) = \max \left\{ 0, \mathfrak{f}_t(i) + \sqrt{\max\{0, \mathfrak{f}_t(i)^2 + \mathfrak{g}_t(i)\}} \right\}, \quad (3.3)$$

where $d_t(i) \in (0, 1)$ ²⁵ is small when firm behaviour is highly distorted by firms' incentives to deviate from choosing the same price as the other firms in their industry, off the equilibrium path (so $d_t(i) \rightarrow 1$ as $J_t(i) \rightarrow \infty$), and $\mathfrak{f}_t(i)$ and $\mathfrak{g}_t(i)$ ²⁶ are increasing in an industry's distance from the frontier, as the further behind a firm is, the greater are the returns to appropriation.

²³ The second order condition for research may be derived most readily by noting that when $d_t(i) \rightarrow 1$, (i.e. $J_t(i) \rightarrow \infty$) the first order condition for research is identical to the one that would have been derived had there been a continuum of firms in each industry with exogenous elasticity of substitution $\frac{1+\mu_t(i)}{\mu_t(i)}$. That it holds more generally follows by continuity. Since $A_t^{**}(i, j)$ is bounded above, no matter how much appropriation is performed the highest solution of the appropriation first order condition must be at least a local maximum.

²⁴ The first order and zero profit conditions are reported in an appendix, section 7.1, where we also derive these solutions. We do not assume $\sigma_Z = 0$ when simulating, but it leads here to expressions that are easier to interpret.

²⁵ Defined in the appendix, section 7.1.

²⁶ $\mathfrak{f}_t(i) := \frac{1}{2} \left[1 + \frac{d_t(i)}{\tau \mu_t(i)} \frac{1+(\gamma-\zeta^R)\mathcal{L}_t^R(i)}{1+\gamma\mathcal{L}_t^R(i)} \right] \left[1 - \left(\frac{A_t^*(i)}{A_t^*} \right)^\tau \right] - 1$, $\mathfrak{g}_t(i) := \frac{d_t(i)}{\tau \mu_t(i)} \frac{1+(\gamma-\zeta^R)\mathcal{L}_t^R(i)}{1+\gamma\mathcal{L}_t^R(i)} A_t^*(i)^{-\zeta^A} \Upsilon [L_t^R(i) + L_t^R(i) + L_t^F] \left[1 - \left(\frac{A_t^*(i)}{A_t^*} \right)^\tau \right] - \left(\frac{A_t^*(i)}{A_t^*} \right)^\tau$.

Equations (3.2) and (3.3) mean that research and appropriation levels are increasing in the other sunk costs a firm must pay prior to production, but decreasing in mark-ups. They also mean that the strategic distortions caused by there being a small number of firms within an industry tend to reduce research and appropriation levels. Other sunk costs matter for research levels because when other sunk costs are high, entry into the industry is lower, meaning that each firm receives a greater slice of production-period profits, and so has correspondingly amplified research incentives.

Why mark-up increases decrease research incentives is clearest when those mark-up increases are driven by exogenous decreases in the elasticity of substitution. When products are close substitutes, then by performing research (and cutting its price) a firm may significantly expand its market-share, something that will not happen when the firm's good is a poor substitute for its rivals. When $d_t(i) \approx 1$ (i.e. there are a lot of firms in the industry) firms act as if they faced an exogenous elasticity of substitution $\frac{1+\mu_t(i)}{\mu_t(i)}$, and so when mark-ups are high they will want to perform little research. When $d_t(i)$ is small (i.e. there are only a few firms) then firms' behaviour is distorted by strategic considerations. Each firm realises that if they perform extra research today then their competitors will accept lower mark-ups the next period. This reduces the extent to which research allows market-share expansion, depressing research incentives.

Perhaps counter-intuitively, the minimum value of $d_t(i)$ occurs when there is a strictly positive number of firms in the industry. It is certainly true that if there is a single firm in an industry, then, as you would expect, very little research will be performed (because the firm's only incentive to cut prices comes from competition from other industries, competition which is very weak, since those industries are producing poor substitutes to its own good). However, this drop in research incentives is working entirely through the mark-up channel, and so in fact we also have that $d_t(i) \rightarrow 1$ as $J_t(i) \rightarrow 1$. One intuition for this is that there can be no strategic behaviour when there is only a single firm.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

The key thing to note about (3.2) and (3.3) is that research and appropriation are independent of the level of demand, except insofar as demand affects mark-ups and the level of the strategic distortion. This is because when demand is high there is greater entry, so each firm still faces roughly the same demand. This is essential for removing the short-run scale effect.

In industries that are no longer patent-protected, rents will be zero (i.e. $L_t^R(i) \equiv 0$). Since research is getting harder at a faster rate than appropriation ($\zeta^R > \zeta^A$), at least asymptotically, no research will be performed in these industries. This is because $A_t^{**}(i)^{-\zeta^R} \Psi[L_t^A(i) + L^F] - \mu_t(i)$ is asymptotically negative since $\mu_t(i) \in (\eta\lambda, \lambda]$. For growth to continue forever in the absence of patent protection, we would require that the overhead cost (L^F) was growing over time at exactly the right rate to offset the increasing difficulty of research. This does not seem particularly plausible. However, it will turn out that optimal patent rents grow at exactly this rate, so with patent protection we will be able to sustain long run growth even when overhead costs are asymptotically dominated by the costs of research. In the presence of sufficiently-severe financial frictions of the “pledgibility constraint” form (Hart and Moore 1994), it may be shown that long run growth is sustainable even without patent protection. We leave the details of this for future work.

Appropriation is performed in an industry if and only if $g_t(i) > 0$, which, for a non-patent protected industry no longer performing research, is true if and only if:

$$\frac{A_t^*(i)}{A_t^*} < \left(\frac{A_t^*(i)^{-\zeta^A} \Upsilon L^F}{A_t^*(i)^{-\zeta^A} \Upsilon L^F + \tau \frac{\mu_t(i)}{d_t(i)}} \right)^{\frac{1}{\tau}}.$$

The left hand side of this equation is the relative productivity of the industry compared to the frontier. The right hand side of this equation will be shrinking over time at roughly $\frac{\zeta^A}{\tau}$ times the growth rate of the frontier, meaning the no-appropriation cut-off point is also declining over time. Indeed, we show in an appendix, section 7.2, that asymptotically the

relative productivity of non-protected firms shrinks at $\frac{\zeta^A}{\tau} \left[1 + \frac{\zeta^A}{\tau}\right]^{-1}$ times the growth rate of the frontier. This is plausible since productivity differences across industries have been steadily increasing over time,²⁷ and is important for the tractability of our model since it enables us to focus on the asymptotic case in which non-protected firms never perform appropriation. It is also in line with the long delays in the diffusion of technology found by Mansfield (1993) amongst others.

3.4. 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 revenues. In this period, and each subsequent one for which they have a patent, the inventor optimally chooses the rent $\mathcal{R}_t(i)$ (or equivalently $L_t^{\mathcal{R}}(i)$) to charge all the firms that wish to produce their product. We are supposing inventors lack the necessary human capital to produce their product at scale themselves.

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 $t + 1$.

The reader should have a firm such as Apple in mind when thinking about these inventors. Apple has no manufacturing plants and instead maintains its profits by product innovation and tough bargaining with suppliers.

²⁷ Some indirect evidence for this is provided by the increase in wage inequality, documented in e.g. Autor, Katz, and Kearney (2008). Further evidence is provided by the much higher productivity growth rates experienced in manufacturing, compared to those in services (mostly unpatented and unpatentable), documented in e.g. Duarte and Restuccia (2009).

3.4.1. Optimal rent decisions

Inventor's businesses are also owned by households; hence, an inventor's problem is to choose $L_{t+s}^{\mathcal{R}}(i)$ for $s \in \mathbb{N}$ to maximise their expected profits, which are given by:

$$\pi_t := \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s (1 - q)^s \left[\prod_{k=1}^s \mathbb{E}_{t+k} \right] L_{t+s}^{\mathcal{R}}(i) W_{t+s} J_{t+s}(i),$$

subject to an enforceability constraint on rents. If the rents charged by a patent-holder go too high, a firm is likely to ignore them completely in the hope that either they will be lucky, and escape having their profits confiscated from them by the courts (since proving patent infringement is often difficult), or that the courts will award damages less than the licence fee. This is plausible since the relevant U.S. statute states that *"upon finding for the claimant the court shall award the claimant damages adequate to compensate for the infringement but in no event less than a reasonable royalty for the use made of the invention by the infringer, together with interest and costs as fixed by the court"*.^{28,29} The established legal definition of a "reasonable royalty" is set at the outcome of a hypothetical bargaining process that took place immediately before production,³⁰ so patent-holders may just as well undertake precisely this bargaining process before production begins.³¹

²⁸ 35 U.S.C. § 284 Damages.

²⁹ The reasonable royalty condition is indeed the relevant one for us since our assumption that the patent-holder lacks the necessary human capital to produce at scale themselves means it would be legally debatable if they had truly "lost profits" following an infringement (Pincus 1991).

³⁰ *Georgia-Pacific*, 318 F. Supp. at 1120 (S.D.N.Y. 1970), *modified on other grounds*, 446 F.2d 295 (2d Cir.), *cert. denied*, 404 U.S. 870 (1971), cited in Pincus (1991), defines a reasonable royalty as *"the amount that a licensor (such as the patentee) and a licensee (such as the infringer) would have agreed upon (at the time the infringement began) if both had been reasonably and voluntarily trying to reach an agreement; that is, the amount which a prudent licensee—who desired, as a business proposition, to obtain the licence to manufacture and sell a particular article embodying the patented invention—would have been willing to pay as a royalty and yet be able to make a reasonable profit and which amount would have been acceptable by a prudent patentee who was willing to grant a licence."*

³¹ In any case, if we allow for idiosyncratic "idea shocks" firms will wish to delay bargaining until this point anyway, since with a bad shock they will be less inclined to accept high rents. Patent-holders also wish to delay till this point because the more sunk costs the firms have already expended before bargaining begins, the greater the size of the "pie" they are bargaining over.

This leads patent-holders to set:

$$L_t^R(i) = \frac{1-p}{p} [L_t^R(i) + L_t^A(i) + L^F], \quad (3.4)$$

at least for sufficiently large t , where $p \in (0,1)$ is the bargaining power of the firm, in the sense of the generalized Nash bargaining solution. The simple form of this expression comes from the fact that a firm's production period profits (which is what is being bargained over) are precisely equal to the costs they face prior to production, thanks to the free entry condition. A full description of the legally motivated bargaining process is contained in an appendix, section 7.3, along with a discussion of some technical complications pertaining to off equilibrium play.

From combining (3.2) and (3.4) then, at least for sufficiently large t , in the limit as $\sigma_Z \rightarrow 0$, we have that:

$$\mathcal{L}_t^R(i) = \frac{p\mu_t(i) - d_t(i)A_t^{**}(i)^{-\zeta^R}\Psi(L_t^A(i) + L^F)}{d_t(i) - \gamma p\mu_t(i)}.$$

For there to be growth in the long run then, we require $d_t(i) > \gamma p\mu_t(i)$, which together with the second order and appropriation uniqueness conditions means that it must at least be true that $p\gamma < \frac{1}{\mu_t(i)} < \min\{\gamma, \tau\}$.³² We see that, once optimal rents are allowed for, research is no longer decreasing in mark-ups within an industry, at least for firms at the frontier. Mathematically, this is because the patent-holder sets rents as such a steeply increasing function of research levels. More intuitively, you may think of the patent-holder as effectively controlling how much research is performed by firms in their industry, and as taking most of the rewards from this research. It is then unsurprising that we reach these Schumpeterian conclusions.

³² If the number of firms in protected industries is growing over time then $d_t(i) \rightarrow 1$, so asymptotically these conditions are equivalent.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

The empirical evidence (Scott 1984; Richard C. Levin, Cohen, and Mowery 1985; Aghion et al. 2005; Tingvall and Poldahl 2006) suggests that the cross-industry relationship between competition and research takes the form of an inverted-U. Based on the fact that strategic distortions are maximised (i.e. $d_t(i)$ is minimised) when there is a small finite number of firms, one might perhaps hope that this holds in our model too. Unfortunately, the maximum of $\frac{\mu_t(i)}{d_t(i)}$ (and hence of research) as a function of $J_t(i)$ may be shown to always occur at some $J_t(i) < 1$. While fractional entry may be a legitimate way of modelling niche products that are never fully commercialised, we prefer to explain the inverted-U in the data with reference to the cross-sectional distribution of industries. New industries will start with a production process behind that of the frontier, and thus firms in them will wish to perform large amounts of appropriation and relatively small amounts of research, since appropriation is a cheaper means of increasing productivity for a firm behind the frontier. In the presence of a luck component to appropriation (not included above, for simplicity) this leads new industries to have the highest degree of productivity dispersion, as older industries remain close to the frontier. As a result of this high productivity dispersion, there will be firms in new industries setting both very high, and very low mark-ups, which, combined with the fact they are performing less research than more mature patent-protected industries, would generate an inverted-U.

3.4.2. Invention and long-run stability

We consider invention as a costly process undertaken by inventors until the expected profits from inventing a new product fall to zero. New products appear at the end of the product spectrum. Additionally, once a product has been invented, it cannot be “un-invented”. Therefore, the product index i always refers to the same product, once it has been invented.

There is, however, no reason to think that newly invented products will start off with a competitive production process. A newly invented product may be thought of as akin to a prototype: yes, identical prototypes could be produced by the same method, but doing this

is highly unlikely to be commercially viable. Instead, there will be rapid investment in improving the product's production process until it may be produced as efficiently as its rivals can be. In our model, this investment in the production process is performed not by the inventor but by the manufacturers. Prototyping technology has certainly improved over time;³³ in light of this, we assume that a new product i is invented with a production process of level $A_t^*(i) = E_t A_t^*$, where $E_t \in (0,1)$ controls initial relative productivity.

Just as we expect process research to be getting harder over time, as all the obvious process innovations have already been discovered, so too we may expect product invention to be getting harder over time, as all the obvious products have already been invented. In addition, the necessity of actually finding a way to produce a prototype will result in the cost of product invention also being increasing in $A_t^*(i)$, the initial productivity level of the process for producing the new product. As a result of these considerations, we assume that the labour cost is given by $\mathcal{L}_t^1 I_{t-1}^\chi A_t^*(i)^{\zeta^1}$, where $\mathcal{L}^1 > 0$ determines the difficulty of invention and where $\chi \in \mathbb{R}$ and $\zeta^1 > 0$ control the rate at which inventing a new product gets more difficult because of, respectively, an increased number of existing products or an increased level of productivity.

We are assuming there is free entry of new inventions, so the marginal entrant must not make a positive profit from entering. That is, $I_t \geq I_{t-1}$ must be as small as possible such that:

$$\mathcal{L}_t^1 I_{t-1}^\chi A_t^*(i)^{\zeta^1} W_t \geq \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s (1 - q)^s \left[\prod_{k=1}^s \Xi_{t+k} \right] L_{t+s}^{\mathcal{R}}(I_t) W_{t+s} J_{t+s}(I_t).$$

If, after a shock, invention can satisfy this equation with equality without the growth rate of the stock of products turning negative, then the number of firms per industry will not have

³³ Examples of recent technologies that have raised the efficiency of prototype production include 3D printing and computer scripting languages such as Python.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

to adjust significantly. However, if the $I_t \geq I_{t-1}$ constraint binds, then the number of firms per industry will have to adjust instead, meaning there may be an asymmetry in the response of mark-ups to certain shocks.

It may be shown that, in the long run, $g_I = \frac{1}{1+\chi}(g_N - \zeta^I g_{A^*})$ (where g_V is the asymptotic growth rate of the variable V_t). Therefore, if $\chi = \zeta^I = 0$ the stock of products will grow at exactly the same rate as population, and away from this special case it will be growing more slowly. If invention were to stop asymptotically, eventually there would be no protected industries, and hence no productivity growth. Therefore, for long-run growth, we either require that $g_N \geq \zeta^I g_{A^*}$ (which will hold providing research is getting more difficult sufficiently slowly, as long as population growth continues), or that there is sufficiently fast depreciation of the stock of products.³⁴ Even without product depreciation, productivity growth may be sustained indefinitely in the presence of a declining population if the government offers infinitely renewable patent-protection.

The existence of a solution for our model, at all time periods, requires the number of firms in a protected industry to be bounded below asymptotically. The previous result on the growth rate of the stock of products implies it is sufficient that $(\zeta^R - \frac{\zeta^I}{1+\chi})g_{A^*} \leq \frac{\chi}{1+\chi}g_N$ for this to hold. This inequality is guaranteed to be satisfied providing $\zeta^R - \frac{\zeta^I}{1+\chi}$ is sufficiently small. To do this while also ensuring that $g_I > 0$ requires that $\max\{\zeta^I, \zeta^R + \frac{1}{\chi}(\zeta^R - \zeta^I)\} < \frac{g_N}{g_{A^*}}$, which will hold for a positive measure of parameter values providing population growth is strictly positive.³⁵

³⁴ Bilbiie, Ghironi, and Melitz (2012) include such product depreciation in their model. We have chosen not to model it here.

³⁵ More generally, when population is stable, providing there is sufficiently fast (proportional) depreciation of the stock of products, we just require that $\zeta^R < \frac{\zeta^I}{1+\chi}$.

Assuming this condition holds, we may show³⁶ that providing the growth rate of the productivity of newly invented products is sufficiently close to the frontier growth rate (i.e. E_t does not decline too quickly³⁷), asymptotically catch-up to the frontier is instantaneous in protected industries, and the frontier growth rate is stationary. This instantaneous catch-up to the frontier means that, had we allowed for industry-specific shocks, all other protected industries would “inherit” the best industry shock, the period after it arrived. This justifies our focus on aggregate “idea” shocks. Additionally, instantaneous catch-up to the frontier means that providing there is population growth or product depreciation, asymptotically, long-run growth may be sustained even in the absence of patent-protection (i.e. when $q_t = 0$), as the one period in which the inventor has a first mover advantage is sufficient for their industry to surpass the existing frontier.

If the number of firms in protected industries were asymptotically infinite, then our simulations would tell us nothing about the consequences of the variations in this number that we might see non-asymptotically. Therefore, it will be helpful if it is additionally the case that this number is asymptotically finite. To guarantee this will, unfortunately, require a knife-edge assumption, namely that $\left(\zeta^R - \frac{\zeta^I}{1+\chi}\right) g_{A^*} = \frac{\chi}{1+\chi} g_N$. To satisfy this without restricting population growth rates means $\chi = 0$ (so invention is not made more difficult by the number of existing products) and $\zeta^R = \zeta^I$ (so prototype production is increasing in difficulty at the same rate as research). The former assumption may be justified by noting that many situations in which invention is apparently getting harder over time because of

³⁶ Suppose $(i_t)_{t=0}^{\infty}$ is a sequence of industries, all protected at t , whose productivity grows at rate $\tilde{g} \leq g_{A^*}$ asymptotically. We conjecture that $\lim_{t \rightarrow \infty} A_t^{**} (i_t)^{-\zeta^R} \Psi L_t^A (i_t) = 0$ and verify. This assumption implies that effective research is asymptotically bounded, since mark-ups are. Hence from (3.3), since $\zeta^R > \zeta^A$, effective appropriation is growing at a rate in the interval $\left(\frac{\zeta^R \tilde{g} - \zeta^A \tilde{g}}{2}, \frac{\zeta^R g_{A^*} - \zeta^A \tilde{g}}{2}\right) \subseteq (0, \infty)$. Therefore $A_t^{**} (i_t)^{-\zeta^R} \Psi L_t^A (i_t)$ is growing at a rate in the interval $\left(-\zeta^R g_{A^*} + \zeta^A \tilde{g} + \frac{\zeta^R \tilde{g} - \zeta^A \tilde{g}}{2}, -\zeta^R \tilde{g} + \zeta^A \tilde{g} + \frac{\zeta^R g_{A^*} - \zeta^A \tilde{g}}{2}\right)$. For our claim to be verified we then just need that $\frac{\zeta^R}{2\zeta^R - \zeta^A} g_{A^*} < \tilde{g}$, which certainly holds when $\tilde{g} = g_{A^*}$ as $\zeta^R > \zeta^A$.

³⁷ As $\zeta^A \rightarrow 0$ it is sufficient that E_t is declining at less than half the rate that A_t^* is growing.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

congestion effects may equally well be explained by production-process-difficulty effects. The latter assumption is immediately plausible, since both parameters are measuring the complexity of working with a given production process. However, unlike with knife-edge growth models whereby relatively slight departures from the stable parameter values results in growth that could not possibly explain our observed stable exponential growth, here, away from the knife-edge case we will have slowly decreasing mark-ups, consistent with Ellis's (2006) evidence of a persistent decline in UK whole economy mark-ups over the last thirty years and Kim's (2010) evidence of non-stationarity in mark-ups.

We assume then that $0 = \chi < \zeta^A < \zeta^R = \zeta^I$. Since asymptotically non-protected industries perform no research or appropriation under these assumptions, their entry cost to post-entry industry profits ratio is tending to zero, meaning their number of firms will tend to infinity as $t \rightarrow \infty$. This is in line with our motivating intuition that excess entry in non-protected industries kills research and appropriation incentives.

3.5. Simulations

With $0 = \chi < \zeta^A < \zeta^R = \zeta^I$, as $t \rightarrow \infty$ the behaviour of our model tends towards stationarity in the key variables. It is this asymptotically stationary model that we simulate. For convenience we define $\zeta := \zeta^R = \zeta^I$. The full set of equations of the de-trended model are given in an appendix, section 7.4. The definition of equilibrium here is entirely standard.

When $\lambda = \nu = \gamma = 1$, it may be shown analytically that the equations determining the model's steady-state have at most two solutions with more than one firm in each industry. However, only one of these two solutions exists for large values of \mathcal{L}^I , i.e. when invention is costly. Since we think that in reality invention is getting harder over time due to congestion effects (i.e. $\chi > 0$), any solution that only exists for small values of \mathcal{L}^I is non-feasible. Our numerical investigations suggest that the model always has at most these two equilibria,

and that always at most one of them exists for large values of \mathcal{L}^1 .³⁸ However, since the existence of multiple-equilibria is indicative that linear approximations may be inaccurate in that region, rather than just picking the solution that exists for arbitrarily large \mathcal{L}^1 , we instead restrict the parameter space to regions in which there is a unique solution. This ensures that the value of \mathcal{L}^1 we use is indeed large, in this sense.

Since $\Psi E^{\zeta} \mathcal{L}_t^1$ always occurs as a group, without loss of generality we may make the normalization $\Psi := E := 1$. We fix all of the model's other parameters, except \mathcal{L}^1 , to the values estimated for our extended model in section 4. \mathcal{L}^1 is set such that the number of firms in patent-protected industries in this model is equal to that of the estimated extended model. The full parameterisation is reported in an appendix (section 7.7). At the chosen parameters, the model has a unique solution, which will exist for arbitrarily high values of \mathcal{L}^1 .

3.5.1. Simulation method

We take a first-order perturbation approximation around the non-stochastic steady state, perturbing in the variance of shocks, and solve for the rational expectations solution of the linearized model.³⁹ As we have previously mentioned, the zero lower bound on net product creation (i.e. on $g_{I,t}$) means there may be an asymmetric response to sufficiently large shocks, but in fact we do not find that the bound is hit with shocks of the magnitudes we consider.

3.5.2. Impulse responses

In Figure 4 we present the impulse responses that result from 0.1% IID (hence non-persistent) shocks to “ideas” (Z_t), labour supply (Φ_t), demand (Θ_t) and population growth

³⁸ It may be shown analytically that the complete model may always be solved by solving a single nonlinear equation, which was always concave for all the parameters we examined.

³⁹ This was performed using Dynare (Adjemian et al. 2011).

Medium-frequency cycles and the remarkable near trend-stationarity of output.

$(G_{N,t})$. 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 remains a quarterly model. Each shock is in a different column, and the key response variables are in rows.

The principle mechanism of our paper is shown most clearly by the population growth rate shock, shown in the final column. (We do not wish to advance population shocks as a key driver of business cycles though, since real rigidities will significantly reduce their impact.) Following a permanent increase in population, demand is permanently higher, so, in the long run, the number of industries must grow to balance this out. Given sufficiently inelastic labour supply, this long-run increase in the measure of industries requires a short-run substitution of labour from production to invention, pushing down consumption and pushing up wages, and so moderating the rate at which invention will grow. Consequently, in the short run some of the additional demand is absorbed by fluctuations in the number of firms in each industry. Without this additional margin of adjustment, this shock would have led to a large increase in average firm sizes, with a consequent increase in the frontier growth rate and counter-factually large unit root in output.

Despite the tiny movement in frontier productivity (less than 0.01%), there is still however a substantial movement in aggregate productivity in the medium-term. Following the shock, more new products are being invented each period, meaning that a greater proportion of industries are relatively new, and so a greater proportion are patent-protected. But because patent-protected industries have such strong incentives to catch-up to the frontier, patent-protected industries are more productive than non-protected ones, so an increase in the proportion of industries that are patent protected means an increase in aggregate productivity. Patent-protected industries also have higher mark-ups due to the cost of

paying licence fees, and so we also see a rise in mark-ups over the medium-term. It is this mechanism that generates medium-frequency pro-cyclical mark-ups in our model.⁴⁰

This mechanism also underlies our model's response to the other shocks we consider. Following a negative labour supply shock, invention is temporarily more expensive, meaning fewer new industries and consequently lower productivity and mark-ups.⁴¹ Following a demand shock labour is transferred away from research and invention towards production, in order to satisfy the temporary higher demand. This drop in invention on impact means that demand shocks actually reduce output in subsequent periods. This is no longer the case when the shock has some persistence, or when there are real rigidities.

An idea shock permanently increases the productivity of patent-protected industries. Over time, these industries fall out of patent-protection, carrying their higher productivity with them, and thus increasing the average productivity of non-protected firms too. Consequently, aggregate productivity slowly rises towards its permanently higher long run level. However, since the magnitude of the original shock was very small, this will not result in a large unit root in output. Following the shock, patent-protected industries are relatively more productive than normal, and so they are also relatively more profitable. This means patent holders can extract higher rents, and so we see an increase in invention with a corresponding increase in mark-ups over the medium-term.

⁴⁰ Pavlov and Weder (2012) also develop a business cycle model capable of generating pro-cyclical mark-ups, via the changing importance of different types of *buyers* over the business cycle. The properties of these buyers are exogenous in their model however, whereas the properties of the different types of *sellers* that drive our results are endogenous.

⁴¹ Were the number of firms per protected industry to absorb the cost-cut instead, then next-period mark-ups would rise and so future wages would fall. However, an expected fall in wages will increase invention today, since inventor returns are increasing in the expected future wage. Hence, invention must fall in the period of the shock.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

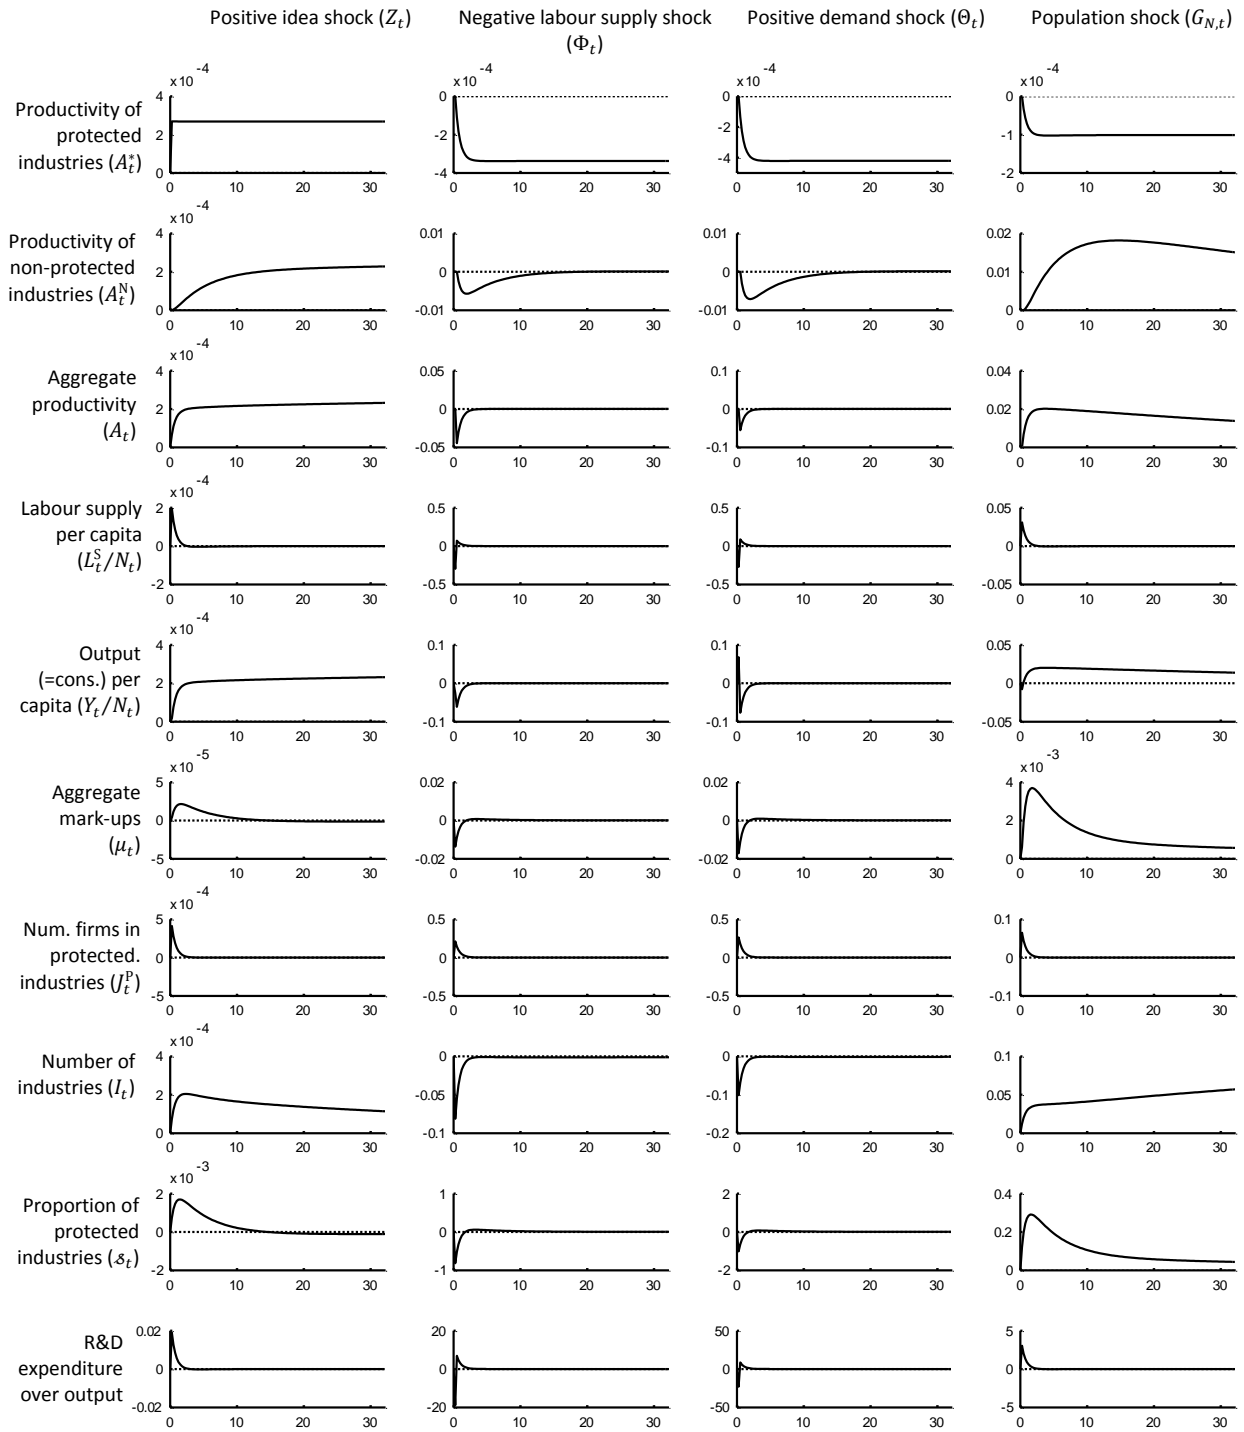


Figure 4: Impulse responses from the core model.
(Vertical axes are in percent, horizontal axes are in years.)

4. Extended model and empirical tests

In order to compare our model to the data seriously, we incorporate habits, capital, and imperfect competition in labour markets. We allow for the possibility of stochastic movements in the key parameters \mathcal{L}^1 , γ and η ,⁴² (though it turns out that the data favours constant values for these parameters), and we specify an AR(1) form for these and all other shocks, with the exception of Z_t , the true technology shock which remains uncorrelated across time. 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.

Additionally, we include intermediate goods as a factor of production, which may be necessary in order to reconcile the low mark-ups found in micro-evidence with the higher mark-ups implied by the inverse labour share. 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. To potentially dampen our model's overly powerful amplification mechanism, we include some spill-overs from frontier productivity growth; these mean that the variance of TFP may be less than that of A_t .

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,

⁴² These parameters are assumed to be known before the entry decision at t , for production in period $t + 1$.

Medium-frequency cycles and the remarkable near trend-stationarity of output. 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.

4.1. Model changes

We assume that firm i in industry j has access to the production technology:

$$Y_t(i, j) = A_t(i, j)X_t^P(i, j)^{\iota_P} [K_t^P(i, j)^{\alpha_P} L_t^P(i, j)^{1-\alpha_P}]^{1-\iota_P}$$

where $X_t^P(i, j)$ is their level of intermediate good input and $K_t^P(i, j)$ 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 without capital. (In particular, profits take the same form, and so research incentives are identical.)

Research, appropriation and invention will also use capital, but we assume that the capital they use is from a different stock. This research/invention capital may be thought of as capturing (variously) education, creativity, ideas, knowledge and advanced physical capital. Rather than the input to the research function for firm j in industry i being $L_t^R(i, j)$, as it was originally, it is now:

$$X_t^R(i, j)^{\iota_R} [K_t^R(i, j)^{\alpha_R} L_t^R(i, j)^{1-\alpha_R}]^{1-\iota_R},$$

where $X_t^R(i, j)$ are the intermediates they use in research, and $K_t^R(i, j)$ is the research/invention capital they use. This will not significantly change research incentives as we can decompose the research problem into a research level one and a cost minimisation one. Additionally, rather than invention requiring a stochastic amount of invention labour, it now requires a stochastic amount of invention output, which is produced using the same production function as research (chiefly for simplicity).

Households' preferences are now given by:

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

for each household $h \in [0,1]$, where

$$\tilde{C}_t(h) := \frac{C_t(h)}{N_t} - \hbar \frac{(1 - \hbar^{\text{INT}})C_{t-1} + \hbar^{\text{INT}}C_{t-1}(h)}{N_{t-1}}$$

is habit adjusted consumption per head, (with $\hbar \in [0,1]$ controlling the strength of consumption habits and $\hbar^{\text{INT}} \in [0,1]$ controlling whether consumption habits are internal or external), and where:

$$\tilde{L}_t^S(h) := \frac{L_{t+s}^S(h)}{N_{t+s}} - \hbar^{\text{LS}} \frac{L_{t+s-1}^S}{N_{t+s-1}}$$

is habit adjusted labour supply per head (with \hbar^{LS} determining the strength of these external labour habits). Each household now supplies a different type of labour $L_t^S(h)$ and potentially receives a different real wage, $W_t(h)$. They face the budget constraint: $C_t + I_t^{\text{KP}} + I_t^{\text{KR}} + B_t = L_t^S(h)W_t(h) + R_t^{\text{KP}}u_t^{\text{P}}K_{t-1}^{\text{P}} + R_t^{\text{KR}}u_t^{\text{R}}K_{t-1}^{\text{R}} + B_{t-1}R_{t-1} + \Pi_t$, where I_t^{KP} and I_t^{KR} is investment in the two capital stocks,⁴³ and $u_t^{\text{P}}K_{t-1}^{\text{P}}$ and $u_t^{\text{R}}K_{t-1}^{\text{R}}$ 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/invention decision may be thought of as capturing the incentives to bunch the implementation of ideas, as stressed by Francois and Lloyd-Ellis (2008; 2009).

⁴³ We assume a complete set of nominal state contingent securities, meaning C_t , I_t^{KP} and I_t^{KR} will not differ across households.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

Following Schmitt-Grohé and Uribe (2011), investment goods of type $V \in \{P,R\}$ are produced from consumption goods using the technology:

$$I_t^{KV*} = A_t^{*\xi_{KV}} E_t^{KV} I_t^{KV}$$

where I_t^{KV} is investment in units of consumption goods and $A_t^{*\xi_{KV}} E_t^{KV}$ 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 (i.e. the underlying trend in A_t) enters into this expression in order to capture the cointegration between the relative price of investment 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:

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

for $V \in \{P,R\}$, where $\delta_t^V(\cdot)$ for $V \in \{P,R\}$ are increasing functions capturing the effect of utilisation on depreciation, locally convex at the steady-state, $Q^{KV}(\cdot)$ for $V \in \{P,R\}$ 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.⁴⁴ (The difference between Γ_t

⁴⁴ Justiniano, Primiceri, and Tambalotti (2011) used the high yield to AAA spread. We choose the BAA-AAA one due to increased data availability.

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.) $\delta_t^V(\cdot)$ has a time subscript since we allow for a shock to depreciation to capture some of the volatility in depreciation shares we observe in the data.⁴⁵ 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 $\delta_t^P(\cdot)$ and $\delta_t^R(\cdot)$.⁴⁶ 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 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 $\delta_t^V(\cdot)$ 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.

Aggregate labour services to firms are now provided by a competitive industry of labour packers using the technology $L_t^T = A_t^{\xi_L} E_t^L \left[\int_0^1 L_t^S(h)^{\frac{1}{1+\lambda_L}} dh \right]^{1+\lambda_L}$, 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 TFP shock.) The productivity of the frontier 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,

⁴⁵ 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.

⁴⁶ We additionally constrain the response of $\delta_t^V(\cdot)$ to the shock such that in its linearised version, with utilisation at its steady-state level, both $\delta_t^V(\cdot)$ and $\delta_t^{V'}(\cdot)$ are positive with at least 95% probability. This is true automatically in the source non-linear specification in which $\delta_t^V(\cdot)$ and $\delta_t^{V'}(\cdot)$ are log-linear in $\tilde{\delta}_t$ when utilisation is at its steady-state, but in preliminary estimates the linearised $\delta_t^{P'}(\cdot)$ turned negative a high proportion of the time, in the absence of this additional constraint.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

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 $L_t^{EV} = L_t^{TV} \left[1 - Q^{LV} \left(\frac{L_t^{TV}}{L_{t-1}^{TV}} \right) \right]$, where $Q^{LV}(\cdot)$ 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 $L_t^T = L_t^{TP} + L_t^{TR}$. 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.

The two positive spillovers from frontier productivity growth mean that the steady-state growth rate of real output per capita is given by $\frac{g_A}{(1-\alpha_p)(1-\iota_p)} + \left(\xi_L + \frac{\alpha_p}{1-\alpha_p} \xi_{KP} \right) g_{A^*}$. If ξ_L and ξ_{KP} are positive then g_{A^*} will not need to be as high, meaning the variance of g_A (and hence that of output) will be lower. Providing the technology for producing overheads takes the same form as that for producing the input to research and invention, neither these spillovers nor the presence of capital and intermediate goods in the production function will change the criterion for no appropriation to be performed asymptotically in non-protected industries. (Away from this special case the lower bound on ζ would be non-zero, and possibly negative.)

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 $1 - \nu$. We assume that those households that cannot adjust their wage optimally will fully index their wage to its steady-state growth rate.

Monetary policy takes an augmented Taylor rule form. 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. In particular:

$$\begin{aligned} \frac{R_t^{\text{NOM}}}{R^{\text{NOM}}} &= \left(\frac{R_{t-1}^{\text{NOM}}}{R^{\text{NOM}}} \right)^{\rho_{R^{\text{NOM}}}} \left[\left(\frac{G_{P,t}}{G_{P,t}^*} \right)^{\mathcal{M}_P} \left(\frac{E_{t-1}^{\text{KP}} G_{A^*,t}^{\xi_{\text{KP}}}}{E_t^{\text{KP}} G_{A^*,t}^{\xi_{\text{KP}}}} \right)^{\mathcal{M}_{\text{PKP}}} \left(\frac{E_{t-1}^{\text{KR}} G_{A^*,t}^{\xi_{\text{KR}}}}{E_t^{\text{KR}} G_{A^*,t}^{\xi_{\text{KR}}}} \right)^{\mathcal{M}_{\text{PKR}}} \left(\frac{G_{W,t}}{G_W} \right)^{\mathcal{M}_W} \right]^{1-\rho_{R^{\text{NOM}}}} \\ &\cdot \left[\left(\frac{R_t^{\text{KP}}}{A_t^{\xi_{\text{KP}}}} \right)^{\mathcal{M}_{\text{RKP}}} \left(\frac{R_t^{\text{KR}}}{A_t^{\xi_{\text{KR}}}} \right)^{\mathcal{M}_{\text{RKR}}} \Theta_t^{\mathcal{M}_\Theta} \tilde{\delta}_t^{-\mathcal{M}_\delta} \right]^{1-\rho_{R^{\text{NOM}}}} \\ &\cdot \left[\left(\frac{Y_t}{N_t A_t^a A_t^{*e}} \right)^{\mathcal{M}_Y} \left(\frac{G_{Y,t}/G_{N,t}}{G_Y/G_N} \right)^{\mathcal{M}_G} \right]^{1-\rho_{R^{\text{NOM}}}} \exp \epsilon_{R^{\text{NOM}},t}, \end{aligned}$$

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, $\frac{E_{t-1}^{\text{KV}}}{E_t^{\text{KV}} G_{A^*,t}^{\xi_{\text{KV}}}}$ is the growth rate of the *real* price of investment goods of type $V \in \{P, R\}$, $G_{W,t}$ 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 $R_t^{\text{SHOCK}} := \exp \sigma_{R^{\text{NOM}}} \epsilon_{R^{\text{NOM}},t}$ is a monetary policy shock. Variables without time subscripts are steady-state values, and the constants a and e are defined in the appendix, section 7.5. In the absence of endogenous productivity, the optimal policy would fully stabilise nominal wages, completely removing the Calvo distortion, thus it is important to allow wages to enter the Taylor rule. 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.) It turns out however that the only significant terms in the 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.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

The model's full set of de-trended equations is given in an appendix (section 7.5).

4.2. Data and estimation

The model is estimated on logs of quarterly U.S. series for nominal output growth,⁴⁷ 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 an appendix, section 7.6, 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)⁴⁸ "measurement 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

⁴⁷ We use nominal output as there should be less measurement error in the nominal series than in the real series.

⁴⁸ 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 $me_t - me_{t-1}$.

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.⁴⁹

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) measurement errors, and the persistence of monetary policy). We use a logit-normal distribution that is scaled to $[-1,1]$ then truncated to $[0,1]$ (i.e. if Z is normally distributed, $\frac{1-\exp(-Z)}{1+\exp(-Z)} \Big| Z > \frac{1}{2}$ has our distribution). We set the mean of the underlying normal distribution to 0 and its variance to 2, which are the unique

⁴⁹ 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.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

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. 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:

- 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}^1),
- all parameters being in the region in which the model is well behaved asymptotically,⁵⁰
- 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),⁵¹
- 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),⁵²
- 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,⁵³

⁵⁰ $p\lambda E\gamma_t < 1$, $\lambda E\eta_t\gamma_t \geq 1$, $E d_t > p E\gamma_t \mu_t^P$, $E g_{I,t} > 0$ and $E J_t^P > 1$.

⁵¹ 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).

⁵² Similarly, this is implemented by adding the steady-state value of $\log \hat{A}_t^N$ as an additional observation variable, with an NIID(0, 1/2 (log(1/1.165) – log(1/1.175))) shock (as before).

- the share of medium frequency variance⁵⁴ decreasing when the mean length of patent protection is reduced by one quarter.

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 $g_{i,t}$ during estimation.

The maximisation is carried out using the CMA-ES algorithm (N. 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.

4.3. Estimation results

The full list of estimated parameters is given in an appendix, section 7.7. 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

⁵³ 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 2.)

⁵⁴ As measured by applying a perfect filter to the spectral density generated by the transition matrices, with accepted band between 8 and 60 years.

Medium-frequency cycles and the remarkable near trend-stationarity of output. moderately biased upwards, meaning that parameters may be estimated more precisely than they appear to be.⁵⁵⁾

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.

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.

The inverse Frisch elasticity of labour supply was driven to its lower bound of $\nu = 0.25$ by the estimation procedure.⁵⁶ 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 adjustment costs, which make aggregate labour supply appear less elastic. Consequently, a standard RBC calibration of the Frisch elasticity based on simulated data

⁵⁵ 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.

⁵⁶ When this bound was not imposed, the estimated value was below 0.01.

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 ($\iota_R = 0.178$ (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 ($\alpha_R = 0.996$ (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 ($Q^{R''}(G_{IKR^*}) = 62.6$ (4.0), in comparison, $Q^{P''}(G_{IKP^*}) = 0.00533$ (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

Medium-frequency cycles and the remarkable near trend-stationarity of output 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 mis-measurement 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 ($Q^{LP'}(G_{LTP}) = 0.0875 (0.0047)$). 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 ($h^{INT} = 0.0151 (0.0032)$), and much less strong than in many DSGE models ($h = 0.253 (0.0041)$). 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 year⁵⁷). 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). 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.

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 Table 6 in the appendix.) 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,

⁵⁷ The low figure comes from deflating by the consumption price, rather than by a consumption-investment price aggregate.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

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 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.

4.4. Model evaluation

As previously mentioned, 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%
Cons. 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

Medium-frequency cycles and the remarkable near trend-stationarity of output. 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 $q_j = 0$, (and without the constraint on the effect of increasing q_j on the share of medium-frequency variance). Doing this reduced the log posterior density by 14.14⁵⁸ which with flat priors would mean we could reject the null of the validity of the $q_j = 0$ constraint at even 0.01% significance. Now, with $q_j = 0$, 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 the model collapses to a medium scale variant of the Jaimovich (2007) model. Hence, our ability to reject the null of $q_j = 0$ 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 Table 5 in the appendix, we report the p-values of LM tests for the presence of

⁵⁸ The log posterior density decreased from 13462.01 to 13447.86.

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 results of section 2.




Figure 5: The effect of patent duration on the importance of medium-frequency cycles.

In Figure 5 we verify that the model does indeed predict that increasing the duration of patent-protection increases the share of variance attributable to medium-frequency cycles. Each dot represents an estimated variance share using the spectral density implied by the transition matrices. 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. Consequently, we see in Figure 3 that increasing patent duration (reducing q) does indeed increase the share of medium-frequency variance. The left hand axis of this graph corresponds to the estimated value of q , so of course at that point it was constrained to have negative slope, but its continual decrease across the range was not a product of a constraint imposed in estimation.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

Additionally, output per capita is near trend stationary in our model, just as in the data. 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.




Figure 6: 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 6 we replicate Figure 2 with 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. 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 instant, however due to the assorted real rigidities in our model, the increase in output will only occur gradually.

4.5. Impulse responses

In Figure 7, we present the impulse responses to the four key driving shocks. As in the previous section, 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.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

To show the magnitude of the effects of these shocks on productivity, we include the implied Solow residual⁵⁹ 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.

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.

⁵⁹ The Solow residual is given by $\frac{Y_t}{K_{t-1}^{\alpha_p} L_t^{1-\alpha_p}} = \frac{\hat{Y}_t A_t^{\frac{1}{1-\alpha_p}} A_t^{*(1-\alpha_p)} \xi_L}{\hat{K}_{t-1}^{\alpha_p} \hat{L}_t^{1-\alpha_p}}$ in the notation of the appendix, section 7.5.

Chapter 1

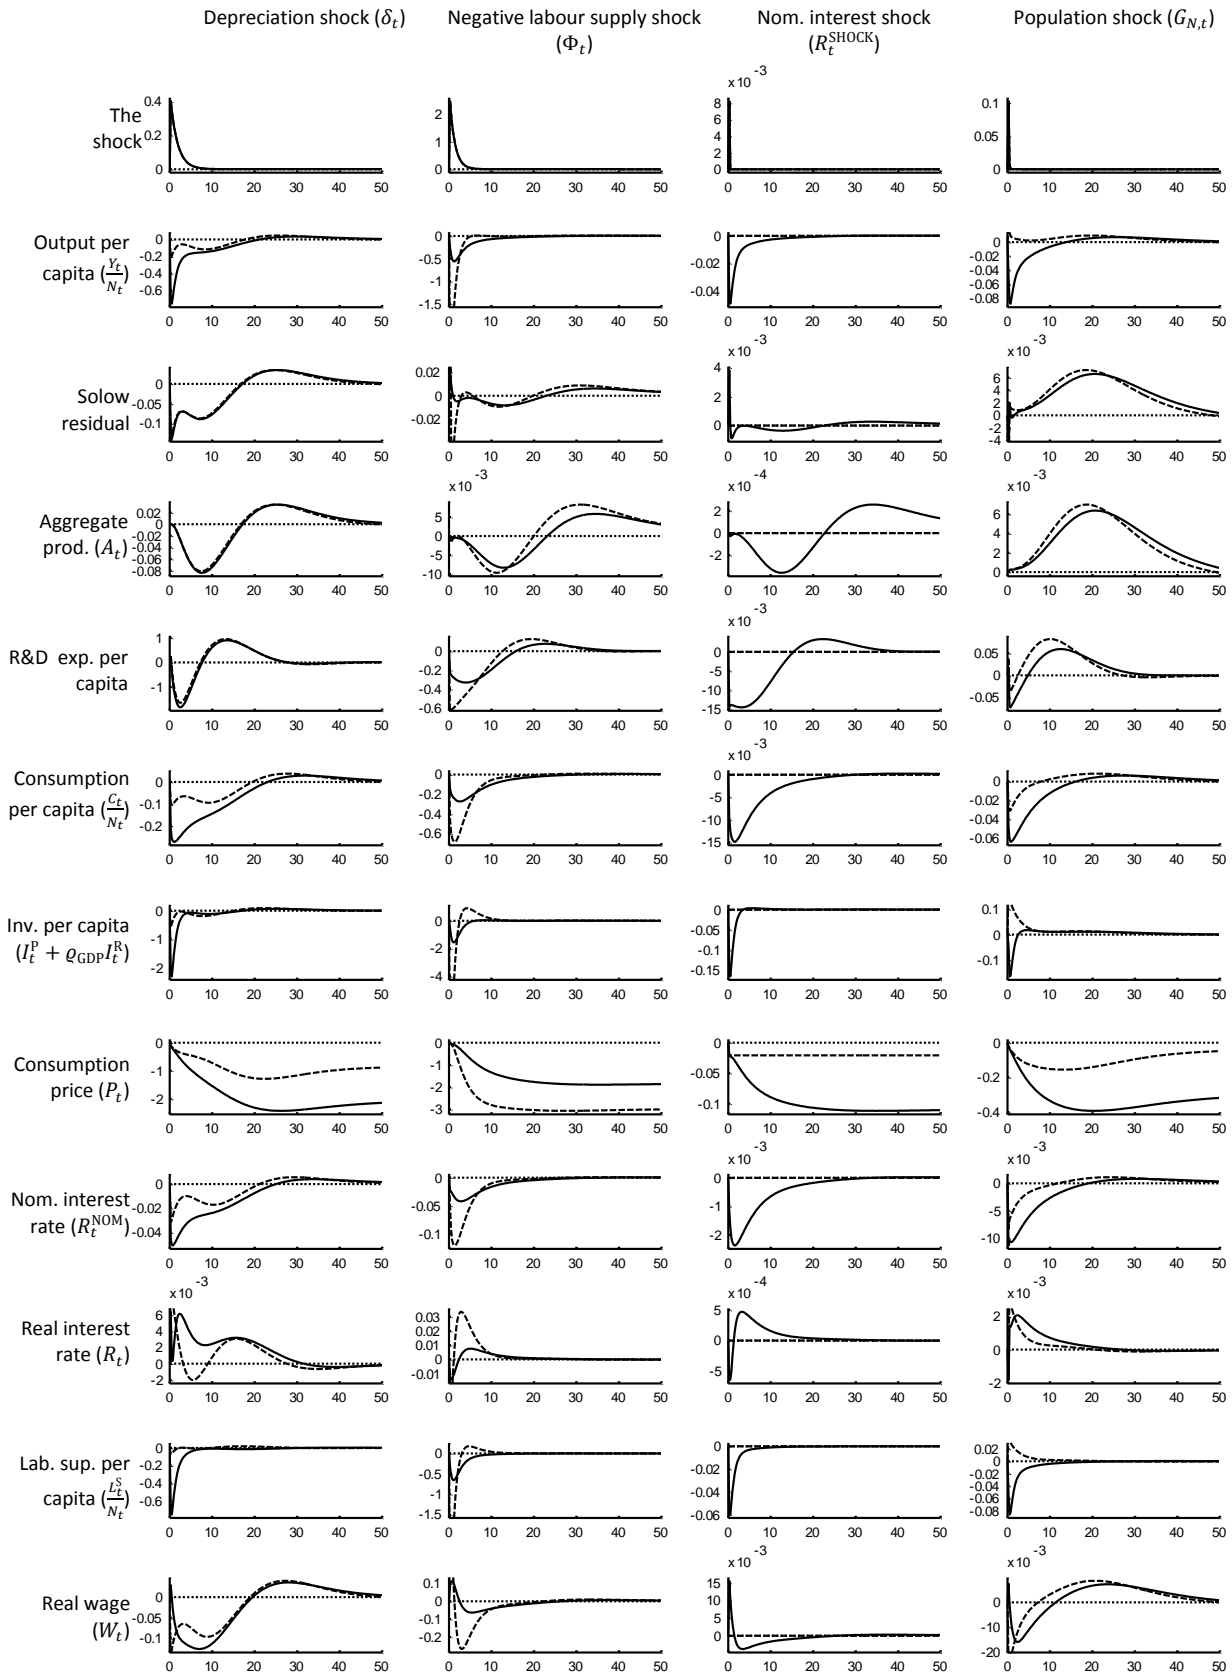


Figure 7: Impulse responses from the core model.

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

5. Conclusion

Many have expressed the worry that “the apparent fit of the DSGE model [has] more to do with the inclusion of suitable exogenous driving processes than with the realism of the model structure itself”⁶⁰. In this paper, we have demonstrated that if productivity is endogenized through research, appropriation and invention then even a frictionless RBC model is capable of generating rich persistent dynamics from uncorrelated shocks, and a full medium-scale model is capable of accurately matching key moments, 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. Furthermore, this improvement in the model’s propagation mechanism does not come at the expense of counter-factual movements in mark-ups, implausibly large unit roots in output, or the use of a growth model that we can reject thanks to the absence of strong scale effects in the data. In all of these respects, then, our model presents a substantial advance on the prior literature.

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.

⁶⁰ Del Negro et al. (2007) paraphrasing Kilian (2007).

Our model suggests that a switch to indefinite patent protection would result in significant welfare improvements. Such a switch would both permanently increase the level of aggregate productivity, and substantially lessen its variance and persistence, while only slightly increasing mark-ups and efficiency losses due to research duplication. Indeed, it may be shown that in our model increasing patent protection even slightly increases growth rates, as industry profits are decreasing in aggregate productivity, and so with indefinite patent protection each (protected) industry has fewer firms meaning higher mark-ups and higher research. However, it is clear that the structure of our model has “stacked-the-deck” in favour of finding a beneficial role for patent protection. Patents in our model are less broad than in the real world, and they do not hinder future research or invention. One minimal conclusion we can draw on patent protection is that product patents should at least be long enough that by the end of patent protection, production process have reached frontier productivity. In our model, this time goes to zero asymptotically. A less radical policy change might be to grant temporary extensions to patents that would otherwise expire during a recession. We intend to explore the full policy implications of this model in future work.

6. References

- Adjemian, Stéphane, Houtan Bastani, Michel Juillard, Ferhat Mihoubi, George Perendia, Marco Ratto, and Sébastien Villemot. 2011. *Dynare: Reference Manual, Version 4*. Dynare Working Papers. CEPREMAP, April.
<http://ideas.repec.org/p/cpm/dynare/001.html>.
- Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. Competition and Innovation: An Inverted-U Relationship. *Quarterly Journal of Economics* 120, no. 2 (May): 701–728.
- Aghion, Philippe, and Peter Howitt. 1992. A Model of Growth Through Creative Destruction. *Econometrica* 60, no. 2 (March): 323–351.
- Andrews, Donald W. K. 1991. Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica* 59, no. 3 (May 1): 817–858.
- Autor, David H, Lawrence F Katz, and Melissa S Kearney. 2008. Trends in U.S. Wage Inequality: Revising the Revisionists. *Review of Economics and Statistics* 90, no. 2: 300–323.

- Bachmann, Ruediger, and Christian Bayer. 2009. *The Cross-section of Firms over the Business Cycle: New Facts and a DSGE Exploration*. CESifo Group Munich.
http://ideas.repec.org/p/ces/ceswps/_2810.html.
- 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, no. 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. <http://ideas.repec.org/p/nbr/nberwo/16130.html>.
- Bental, Benjamin, and Dan Peled. 1996. The Accumulation of Wealth and the Cyclical Generation of New Technologies: A Search Theoretic Approach. *International Economic Review* 37, no. 3 (August): 687–718.
- Bilbiie, Florin O., Fabio Ghironi, and Marc J. Melitz. 2011. *Endogenous Entry, Product Variety, and Business Cycles*. CEPR Discussion Papers. C.E.P.R. Discussion Papers, September.
<http://ideas.repec.org/p/cpr/ceprdp/8564.html>.
- . 2012. Endogenous Entry, Product Variety, and Business Cycles. *Journal of Political Economy* 120, no. 2 (April 1): 304–345.
- Bils, Mark. 1998. Discussion of Basu (1998). *Conference Series ; [Proceedings]*, no. Jun.
http://www.bos.frb.org/economic/conf/conf42/con42_14.pdf.
- Bonanno, Giacomo. 1988. Oligopoly Equilibria When Firms Have Local Knowledge of Demand. *International Economic Review* 29, no. 1 (February 1): 45–55.
- Boulhol, Hervé. 2007. The Convergence of Price–cost Margins. *Open Economies Review* 19, no. 2 (May): 221–240.
- Broda, Christian, and David E Weinstein. 2010. Product Creation and Destruction: Evidence and Price Implications. *American Economic Review* 100, no. 3 (June): 691–723.
- Calvo, Guillermo A. 1983. Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics* 12, no. 3: 383–398.
- Campbell, Jeffrey R. 1998. Entry, Exit, Embodied Technology, and Business Cycles. *Review of Economic Dynamics* 1, no. 2 (April): 371–408.
- Candès, Emmanuel, Michael Wakin, and Stephen Boyd. 2008. Enhancing Sparsity by Reweighted ℓ_1 Minimization. *Journal of Fourier Analysis and Applications* 14, no. 5 (December 1): 877–905.
- Canova, Fabio. 2009. Bridging cyclical DSGE models and the raw data. August 28.
[http://www.crei.cat/people/canova/dsge_trend\(1\).pdf](http://www.crei.cat/people/canova/dsge_trend(1).pdf).
- 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, no. 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, no. 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, no. 1: 1–45.
- Cochrane, John H. 1988. How Big Is the Random Walk in GNP? *Journal of Political Economy* 96, no. 5 (October 1): 893–920.
- Comin, Diego. 2009. On the integration of growth and business cycles. *Empirica* 36, no. 2 (May 1): 165–176.

- Comin, Diego, and Mark Gertler. 2006. Medium-Term Business Cycles. *American Economic Review* 96, no. 3 (June): 523–551.
- Comin, Diego, Mark Gertler, and Ana Maria Santacreu. 2009. *Technology Innovation and Diffusion as Sources of Output and Asset Price Fluctuations*. National Bureau of Economic Research, Inc, June. RePEc. <http://ideas.repec.org/p/nbr/nberwo/15029.html>.
- Comin, Diego, and Sunil Mulani. 2009. A theory of growth and volatility at the aggregate and firm level. *Journal of Monetary Economics* 56, no. 8 (November): 1023–1042.
- Diewert, W. E. 1978. Superlative Index Numbers and Consistency in Aggregation. *Econometrica* 46, no. 4 (July 1): 883–900.
- Dixit, Avinash K., and Joseph E. Stiglitz. 1977. Monopolistic Competition and Optimum Product Diversity. *The American Economic Review* 67, no. 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.
- Duarte, Margarida, and Diego Restuccia. 2009. The Role of the Structural Transformation in Aggregate Productivity. Working Papers tecipa-348. University of Toronto, Department of Economics. <http://ideas.repec.org/p/tor/tecipa/tecipa-348.html>.
- 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, no. 3: 6.
- Eaton, Jonathan, and Samuel Kortum. 1999. International Technology Diffusion: Theory and Measurement. *International Economic Review* 40, no. 3 (August 1): 537–570.
- Eisfeldt, Andrea L., and Adriano A. Rampini. 2006. Capital reallocation and liquidity. *Journal of Monetary Economics* 53, no. 3 (April): 369–399.
- Ellis, Colin. 2006. *Elasticities, Markups and Technical Progress: Evidence from a State-Space Approach*. SSRN. <http://ssrn.com/paper=932580>.
- 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, no. 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, no. 3 (June): 389–405.
- Forni, Mario, and Sergio Paba. 2002. Spillovers and the growth of local industries. *The Journal of Industrial Economics* 50, no. 2 (June 1): 151–171.
- Francois, Patrick, and Huw Lloyd-Ellis. 2009. Schumpeterian cycles with pro-cyclical R&D. *Review of Economic Dynamics* 12, no. 4 (October): 567–591.
- Francois, Patrick, and Huw Lloyd-Ellis. 2008. Implementation cycles, investment, and growth. *International Economic Review* 49, no. 3 (August 1): 901–942.
- Furlanetto, Francesco, and Martin Seneca. 2011. *New perspectives on depreciation shocks as a source of business cycle fluctuations*. Working Paper. Norges Bank, March. http://ideas.repec.org/p/bno/worpaper/2011_02.html.
- Ginarte, Juan C., and Walter G. Park. 1997. Determinants of patent rights: A cross-national study. *Research Policy* 26, no. 3 (October): 283–301.
- Griliches, Zvi. 1998. *The Search for R&D Spillovers*. NBER Chapters. National Bureau of Economic Research, Inc. <http://www.nber.org/chapters/c8349.pdf>.
- Groth, Christian, Karl-Josef Koch, and Thomas Steger. 2009. When economic growth is less than exponential. *Economic Theory* (July 4). <http://dx.doi.org/10.1007/s00199-009-0480-y>.
- 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, no. 1: 180–197.
- Hart, Oliver, and John Moore. 1994. A Theory of Debt Based on the Inalienability of Human Capital. *The Quarterly Journal of Economics* 109, no. 4 (November 1): 841–879.
- 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. <http://ideas.repec.org/p/nbr/nberwo/17520.html>.
- 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, May. http://pwt.econ.upenn.edu/php_site/pwt_index.php.
- Hodrick, Robert J., and Edward C. Prescott. 1997. Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking* 29, no. 1 (February): 1–16.
- Holden, Tom. 2011. *Products, patents and productivity persistence: A DSGE model of endogenous growth*. Dynare Working Papers. CEPREMAP, May. <http://ideas.repec.org/p/cpm/dynare/004.html>.
- Horii, Ryo. 2011. Wants and past knowledge: Growth cycles with emerging industries. *Journal of Economic Dynamics and Control* In Press, Accepted Manuscript. <http://www.sciencedirect.com/science/article/pii/S0165188911001436>.
- 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, no. 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, no. 1 (November): 300–325.
- Jaimovich, Nir, and Sergio Rebelo. 2008. News and Business Cycles in Open Economies. *Journal of Money, Credit and Banking* 40, no. 8: 1699–1711.
- . 2009. Can News about the Future Drive the Business Cycle? *American Economic Review* 99, no. 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, no. 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, no. 2 (June): 464–476.
- Kehrig, Matthias. 2011. *The Cyclicalities of Productivity Dispersion*. Center for Economic Studies, U.S. Census Bureau, May. <http://ideas.repec.org/p/cen/wpaper/11-15.html>.
- Kilian, Lutz. 2007. Comment. *Journal of Business & Economic Statistics* 25. *Journal of Business & Economic Statistics*: 156–159.
- Kim, Bae-Geun. 2010. Identifying a permanent markup shock and its implications for macroeconomic dynamics. *Journal of Economic Dynamics and Control* (April 14). <http://dx.doi.org/10.1016/j.jedc.2010.04.003>.
- Kiyotaki, Nobuhiro, and John Moore. 1997. Credit Cycles. *Journal of Political Economy* 105, no. 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. Vol. Volume 1, Part A. Elsevier. <http://www.sciencedirect.com/science/article/pii/S1574068405010117>.
- Lee, Anthony, Francois Caron, Arnaud Doucet, and Chris Holmes. 2010. A Hierarchical Bayesian Framework for Constructing Sparsity-inducing Priors. *arXiv:1009.1914* (September 9). <http://arxiv.org/abs/1009.1914>.

- Levin, Richard C., Wesley M. Cohen, and David C. Mowery. 1985. R & D Appropriability, Opportunity, and Market Structure: New Evidence on Some Schumpeterian Hypotheses. *The American Economic Review* 75, no. 2 (May): 20–24.
- Levy, Daniel, and Hashem Dezhbakhsh. 2003. International evidence on output fluctuation and shock persistence. *Journal of Monetary Economics* 50, no. 7 (October): 1499–1530.
- Litterman, Robert B. 1983. A Random Walk, Markov Model for the Distribution of Time Series. *Journal of Business & Economic Statistics* 1, no. 2. *Journal of Business & Economic Statistics*: 169–73.
- Liu, Zheng, Daniel F Waggoner, and Tao Zha. 2011. Sources of macroeconomic fluctuations: A regime-switching DSGE approach. *Quantitative Economics* 2, no. 2 (July 1): 251–301.
- Mansfield, Edwin. 1993. The diffusion of industrial robots in Japan and the United States. *Research Policy* 22, no. 2 (April): 105.
- Matsuyama, Kiminori. 1999. Growing Through Cycles. *Econometrica* 67, no. 2 (March): 335–347.
- Del Negro, Marco, Frank Schorfheide, Frank Smets, and Rafael Wouters. 2007. Rejoinder. *Journal of Business and Economic Statistics* 25, no. 2: 159.
- Nekarda, Christopher, and Valerie Ramey. 2010. The Cyclical Behavior of the Price-Cost Markup. <http://econ.ucsd.edu/~vramey/research/markupcyc.pdf>.
- Nuño, Galo. 2008. *Schumpeterian Foundations of Real Business Cycles*. Working Papers. International Economics Institute, University of Valencia, December. <http://ideas.repec.org/p/iei/wpaper/0805.html>.
- . 2009. *Technology, convergence and business cycles*. Banco de España, October. <http://ideas.repec.org/p/bde/wpaper/0922.html>.
- . 2011. Optimal research and development and the cost of business cycles. *Journal of Economic Growth* 16, no. 3 (April 26): 257–283.
- Pavlov, Oscar, and Mark Weder. 2012. Variety matters. *Journal of Economic Dynamics and Control* 36, no. 4 (April): 629–641.
- Peretto, Pietro F. 1999. Cost reduction, entry, and the interdependence of market structure and economic growth. *Journal of Monetary Economics* 43, no. 1 (February 19): 173–195.
- Peterman, W. B. 2011. Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity.
- Phillips, Kerk L., and Jeff Wrase. 2006. Is Schumpeterian ‘creative destruction’ a plausible source of endogenous real business cycle shocks? *Journal of Economic Dynamics and Control* 30, no. 11 (November): 1885–1913.
- Pincus, Laura B. 1991. Computation of Damages in Patent Infringement Actions. *Harvard Journal of Law & Technology* 5: 95.
- La Porta, Rafael, Florencio Lopez-de-Silanes, and Andrei Shleifer. 2008. The Economic Consequences of Legal Origins. *Journal of Economic Literature* 46, no. 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, no. 2 (January 1): 350–371.
- Rubinstein, Ariel. 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica* 50, no. 1 (January): 97–109.
- Said, Said E., and David A. Dickey. 1984. Testing for Unit Roots in Autoregressive-Moving Average Models of Unknown Order. *Biometrika* 71, no. 3 (December 1): 599–607.

- Schmitt-Grohé, Stephanie, and Martín Uribe. 2006. *Comparing Two Variants of Calvo-Type Wage Stickiness*. NBER Working Paper. National Bureau of Economic Research, Inc. <http://ideas.repec.org/p/nbr/nberwo/12740.html>.
- — —. 2011. Business cycles with a common trend in neutral and investment-specific productivity. *Review of Economic Dynamics* 14, no. 1 (January): 122–135.
- Scott, John. 1984. Firm versus Industry Variability in R&D Intensity. In *R & D, Patents, and Productivity*. NBER, January 1. <http://www.nber.org/chapters/c10051>.
- Serrano, C. J. 2007. Technology Transfer and the Business Cycle: Evidence from Patent Reassignments. <http://homes.chass.utoronto.ca/~serrano/papers/RBCpatent.pdf>.
- Serrano, Carlos J. 2010. The dynamics of the transfer and renewal of patents. *The RAND Journal of Economics* 41, no. 4 (December 1): 686–708.
- The Associated Press. 2008. Court Rules Business Concept Cannot Be Patented. *The New York Times*, October 31, sec. Business. http://www.nytimes.com/2008/10/31/business/31patent.html?_r=1.
- Tingvall, Patrik Gustavsson, and Andreas Poldahl. 2006. Is there really an inverted U-shaped relation between competition and R&D? *Economics of Innovation and New Technology* 15, no. 3: 101.
- Wälde, Klaus. 2005. Endogenous Growth Cycles. *International Economic Review* 46, no. 3: 867–894.
- White, Halbert. 1980. A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica* 48, no. 4 (May 1): 817–838.
- Woodford, Michael. 2001. The Taylor Rule and Optimal Monetary Policy. *The American Economic Review* 91, no. 2 (May 1): 232–237.

7. Appendices

7.1. The free-entry and first order conditions

When deciding how much research and appropriation to perform, firms realise that if they perform a non-equilibrium amount then in the next period they will have an incentive to set a different mark-up to the other firms in their industry. (The clearest example of this is when we have perfect competition, in which case the most productive firm would want to price just below the second most productive firms' marginal cost.) It may be seen that in non-symmetric equilibrium the optimal price satisfies:

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

Since we are looking for a symmetric equilibrium, it is sufficient to approximate this locally around $P_t(i) = P_t(i, j)$ in order to calculate firms' research and appropriation incentives.

Taking a log-linear approximation of $\log P_t(i, j)$ in $\frac{P_t(i, j)}{P_t(i)}$ gives us that:

$$P_t(i, j) \approx \frac{W_t}{A_t(i, j)} (1 + \mu_{t-1}(i)) \left(\frac{P_t(i, j)}{P_t(i)} \right)^{-\omega_{t-1}(i)}$$

where $\omega_t(i) := \frac{J_t(i)(1-\eta)}{(J_t(i)-(1-\eta))^2(1+\mu_t(i))}$ captures the strength of these incentives to deviate

from setting the same mark-up as all other firms in their industry. Therefore $P_t(i) \approx$

$\frac{W_t}{A_t(i)} (1 + \mu_{t-1}(i))$ and $P_t(i, j) \approx \frac{W_t}{A_t(i, j)} (1 + \mu_{t-1}(i)) \left(\frac{A_t(i, j)}{A_t(i)} \right)^{\frac{\omega_{t-1}(i)}{1+\omega_{t-1}(i)}}$ where:

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

Therefore, up to a first order approximation around the symmetric solution, profits are given by:

$$\begin{aligned} & \beta \frac{1}{I_t J_t(i)} \left(\frac{1 + \mu_t}{1 + \mu_t(i)} \right)^{\frac{1}{\lambda}} \mathbb{E}_t \mathbb{E}_{t+1} Y_{t+1} \left[\left(\frac{A_{t+1}(i, j)}{A_{t+1}(i)} \right)^{\frac{\omega_t(i)}{1+\omega_t(i)}} \right. \\ & \quad \left. - \frac{1}{1 + \mu_t(i)} \right] \left(\frac{A_{t+1}(i, j)}{A_{t+1}(i)} \right)^{\frac{1-\eta\lambda\omega_t(i)}{\eta\lambda(1+\omega_t(i))}} \left(\frac{A_{t+1}(i)}{A_{t+1}} \right)^{\frac{1}{\lambda}} \\ & \quad - [L_t^R(i, j) + L_t^A(i, j) + L_t^R(i) + L^F] W_t. \end{aligned}$$

Note that if $J_t(i) > \frac{2\sqrt{2}(3-\sqrt{2})}{1+2\sqrt{2}} \approx 1.17$, then $1 - \eta\lambda\omega_t(i) > 0$ (by tedious algebra), so providing there are at least two firms in the industry, this expression is guaranteed to be increasing and concave in $A_{t+1}(i, j)$.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

Let $m_t^R(i, j)W_t$ be the Lagrange multiplier on research's positivity constraint and $m_t^A(i, j)W_t$ be the Lagrange multiplier on appropriation's positivity constraint. Then in a symmetric equilibrium the two first order conditions and the free entry condition (respectively) mean:

$$\begin{aligned} & \beta \frac{1}{I_t J_t(i)} \frac{\mu_t(i)}{1 + \mu_t(i)} \left(\frac{1 + \mu_t}{1 + \mu_t(i)} \right)^{\frac{1}{\lambda}} \mathbb{E}_t \Xi_{t+1} Y_{t+1} \left(\frac{A_{t+1}(i)}{A_{t+1}} \right)^{\frac{1}{\lambda}} \frac{d_t(i)}{\mu_t(i)} \frac{Z_{t+1} A_t^{**}(i)^{-\zeta^R} \Psi}{1 + \gamma Z_{t+1} A_t^{**}(i)^{-\zeta^R} \Psi L_t^R(i)} \\ & = W_t (1 - m_t^R(i)) \\ & \beta \frac{1}{I_t J_t(i)} \frac{\mu_t(i)}{1 + \mu_t(i)} \left(\frac{1 + \mu_t}{1 + \mu_t(i)} \right)^{\frac{1}{\lambda}} \mathbb{E}_t \Xi_{t+1} Y_{t+1} \left(\frac{A_{t+1}(i)}{A_{t+1}} \right)^{\frac{1}{\lambda}} \frac{d_t(i)}{\mu_t(i)} \frac{1 + (\gamma - \zeta^R) Z_{t+1} A_t^{**}(i)^{-\zeta^R} \Psi L_t^R(i)}{1 + \gamma Z_{t+1} A_t^{**}(i)^{-\zeta^R} \Psi L_t^R(i)} \\ & \cdot \frac{1}{\tau} \frac{A_t(i)^{-\zeta^A} \Upsilon (A_t^{*\tau} - A_t(i)^\tau)}{A_t^{**}(i)^\tau (1 + A_t(i)^{-\zeta^A} \Upsilon L_t^A(i))^2} = W_t (1 - m_t^A(i)) \\ & \beta \frac{1}{I_t J_t(i)} \frac{\mu_t(i)}{1 + \mu_t(i)} \left(\frac{1 + \mu_t}{1 + \mu_t(i)} \right)^{\frac{1}{\lambda}} \mathbb{E}_t \Xi_{t+1} Y_{t+1} \left(\frac{A_{t+1}(i)}{A_{t+1}} \right)^{\frac{1}{\lambda}} \\ & = [L_t^R(i, j) + L_t^A(i, j) + L_t^R(i) + L_t^F] W_t \end{aligned}$$

where:

$$d_t(i) := 1 - \frac{\omega_t(i)}{1 + \omega_t(i)} \frac{(\lambda - \mu_t(i))(\mu_t(i) - \eta\lambda)}{\lambda(1 - \eta)\mu_t(i)} < 1$$

and where we have dropped j indices on variables which are the same across the industry.

We also have that:

$$\frac{(\lambda - \mu_t(i))(\mu_t(i) - \eta\lambda)}{\lambda(1 - \eta)\mu_t(i)} \leq \frac{\lambda(1 - \sqrt{\eta})(\sqrt{\eta} - \eta)}{\sqrt{\eta}} < \lambda$$

so providing $\lambda < 1$, $d_t(i) > 0$.

That the solution for research when $Z_{t+1} \equiv 1$ is given by equation (3.2) is a trivial consequence of the complementary slackness condition and the facts that $\frac{1}{\mu_t(i)} < \gamma$ and $d_t(i) < 1$. Deriving (3.3) is less trivial though.

Begin by defining $\kappa_t(i) := \frac{1+(\gamma-\zeta^R)\mathcal{L}_t^R(i)}{1+\gamma\mathcal{L}_t^R(i)}$, and note that since we are assuming $\gamma > \zeta^R \geq 0$, we have that $0 < \kappa_t(i) \leq 1$.

Also define:

$$n_t(i) := \frac{d_t(i)\kappa_t(i)}{\tau\mu_t(i)} A_t^*(i)^{-\zeta^A} \Upsilon \left[\left(\frac{A_t^*}{A_t^*(i)} \right)^\tau - 1 \right] [L_t^R(i) + L_t^R(i) + L_t^F] \geq 0,$$

which is not a function of $L_t^A(i)$, given $L_t^R(i)$.

We can then combine the appropriation first order condition with the free entry condition to obtain:

$$\frac{1}{(1 + \mathcal{L}_t^A(i))^2} \left(\frac{A_t^*(i)}{A_t^{**}(i)} \right)^\tau \left[\frac{d_t(i)\kappa_t(i)}{\tau\mu_t(i)} \left[\left(\frac{A_t^*}{A_t^*(i)} \right)^\tau - 1 \right] \mathcal{L}_t^A(i) + n_t(i) \right] = 1 - m_t^A(i).$$

Since the left hand side is weakly positive, from the dual feasibility condition we know $m_t^A(i) \in [0,1]$. Now when $L_t^A(i) = 0$, this becomes:

$$n_t(i) = 1 - m_t^A(i),$$

since in this case $A_t^*(i) = A_t^{**}(i)$. Therefore when $L_t^A(i) = 0$, $n_t(i) \leq 1$.

We now prove the converse. Suppose then for a contradiction that $L_t^A(i) > 0$, but $n_t(i) \leq$

1. By complementary slackness, we must have $m_t^A(i) = 0$, hence:

$$\begin{aligned} 1 \geq n_t(i) &= (1 + \mathcal{L}_t^A(i))^2 \left(\frac{A_t^{**}(i)}{A_t^*(i)} \right)^\tau - \frac{d_t(i)\kappa_t(i)}{\tau\mu_t(i)} \left[\left(\frac{A_t^*}{A_t^*(i)} \right)^\tau - 1 \right] \mathcal{L}_t^A(i) \\ &\geq (1 + \mathcal{L}_t^A(i))^2 \left(\frac{A_t^{**}(i)}{A_t^*(i)} \right)^\tau - \left[\left(\frac{A_t^*}{A_t^*(i)} \right)^\tau - 1 \right] \mathcal{L}_t^A(i) \\ &= (1 + \mathcal{L}_t^A(i)) \left[(1 + \mathcal{L}_t^A(i)) + \mathcal{L}_t^A(i) \left[\left(\frac{A_t^*}{A_t^*(i)} \right)^\tau - 1 \right] \right] - \left[\left(\frac{A_t^*}{A_t^*(i)} \right)^\tau - 1 \right] \mathcal{L}_t^A(i), \end{aligned}$$

where we have used the facts that $d_t(i)\kappa_t(i) \leq 1$ and $\frac{1}{\mu_t(i)} < \tau$ to derive the second inequality.

Expanding the brackets then gives that:

$$1 \geq 1 + 2\mathcal{L}_t^A(i) + \left(\frac{A_t^*}{A_t^*(i)}\right)^\tau \mathcal{L}_t^A(i)^2,$$

i.e. that $0 \geq 2 + \left(\frac{A_t^*}{A_t^*(i)}\right)^\tau \mathcal{L}_t^A(i)$ which is a contradiction as $\left(\frac{A_t^*}{A_t^*(i)}\right)^\tau \mathcal{L}_t^A(i) \geq 0$.

We have proven then that providing $\frac{1}{\mu_t(i)} < \tau$, $L_t^A(i) = 0$ if and only if $n_t(i) \leq 1$. It just remains for us to solve for $L_t^A(i)$ when it is strictly positive. From the above, we have that, in this case:

$$\left(\frac{A_t^*(i)}{A_t^*}\right)^\tau [n_t(i) - 1] = 2 \left[1 - \frac{1}{2} \left[1 + \frac{d_t(i)\kappa_t(i)}{\tau\mu_t(i)} \right] \left[1 - \left(\frac{A_t^*(i)}{A_t^*}\right)^\tau \right] \right] \mathcal{L}_t^A(i) + \mathcal{L}_t^A(i)^2.$$

Hence:

$$\begin{aligned} \mathcal{L}_t^A(i) = & - \left[1 - \frac{1}{2} \left[1 + \frac{d_t(i)\kappa_t(i)}{\tau\mu_t(i)} \right] \left[1 - \left(\frac{A_t^*(i)}{A_t^*}\right)^\tau \right] \right] \\ & + \sqrt{\left[1 - \frac{1}{2} \left[1 + \frac{d_t(i)\kappa_t(i)}{\tau\mu_t(i)} \right] \left[1 - \left(\frac{A_t^*(i)}{A_t^*}\right)^\tau \right] \right]^2 + \left(\frac{A_t^*(i)}{A_t^*}\right)^\tau [n_t(i) - 1]}, \end{aligned}$$

since the lower solution is guaranteed to be negative as $n_t(i) > 1$ when $L_t^A(i) > 0$.

7.2. The steady state for non-patent-protected industries

In an industry i which is not patent-protected and in which appropriation, but no research, is performed, from (3.1) and (3.3):

$$\mathcal{f}_t(i) + \sqrt{\mathcal{f}_t(i)^2 + \mathcal{g}_t(i)} = \mathcal{L}_t^A(i) = \left[1 - \frac{\left(\frac{A_{t+1}^*(i)}{A_t^*(i)}\right)^\tau - 1}{1 - \left(\frac{A_t^*(i)}{A_t^*}\right)^\tau} \left(\frac{A_t^*(i)}{A_t^*}\right)^\tau \right]^{-1} - 1.$$

If we treat $p_1 := \tau \frac{\mu_t(i)}{d_t(i)} - 1 \approx 0$, $p_2 := A_t^*(i)^{-\zeta^A} \Upsilon_t L_t^F \approx 0$ and $p_3 := \left(\frac{A_{t+1}^*(i)}{A_t^*(i)}\right)^\tau - 1 \approx 0$ as fixed, this leaves us with a cubic in $\left(\frac{A_t^*(i)}{A_t^*}\right)^\tau$, for which only one solution will be feasible (i.e.

strictly less than 1). Taking a second order Taylor approximation of this solution in p_1 , p_2 and p_3 , reveals (after some messy computation), that:

$$\left(\frac{A_t^*(i)}{A_t^*}\right)^\tau \approx p_2(1 - (p_1 + p_2)) = A_t^*(i)^{-\zeta^A} Y_t L_t^F \left(2 - \tau \frac{\mu_t(i)}{d_t(i)} - A_t^*(i)^{-\zeta^A} Y_t L_t^F\right)$$

(The effect of p_3 on $\left(\frac{A_t^*(i)}{A_t^*}\right)^\tau$ is third order and hence it does not appear in this expression.)

From this approximate solution for $\left(\frac{A_t^*(i)}{A_t^*}\right)^\tau$ then, we have that the relative productivity of a non-protected industry is decreasing in its mark-up. Furthermore, from dropping to a first order approximation, we have that $A_t^*(i)^{1+\frac{\zeta^A}{\tau}} \approx A_t^*(Y_t L_t^F)^{\frac{1}{\tau}}$, so asymptotically non-protected industries are growing at $\left[1 + \frac{\zeta^A}{\tau}\right]^{-1}$ times the growth rate of the frontier.

7.3. The inventor-firm bargaining process

We model the entire process of setting and paying rents as follows:

- 1) Firms enter, paying the fixed cost.
- 2) Firms who have entered conduct appropriation, then research.
- 3) The “idea shock” for next period’s production, Z_{t+1} , is realised and firms and patent holders learn its level.
- 4) Finally, firms arrive at the patent-holder to conduct bargaining, with these arrivals taking place sequentially but in a random order. (For example, all firms phone the patent-holder sometime in the week before production is to begin.) In this bargaining we suppose that the patent-holder has greater bargaining power, since they have a longer

Medium-frequency cycles and the remarkable near trend-stationarity of output.

outlook⁶¹ and since they lose nothing if bargaining collapses⁶². We also suppose that neither patent-holders nor firms are able to observe or verify either how many (other) firms paid the fixed cost, or what research and appropriation levels they chose. This is plausible because until production begins it is relatively easy to keep such things hidden (for example, by purchasing the licence under a spin-off company), and because it is hard to ascertain ahead of production exactly what product a firm will be producing. We assume bargaining takes an alternating offer form, (Rubinstein 1982) but that it happens arbitrarily quickly (i.e. in the no discounting limit).

- 5) Firms pay the agreed rents if bargaining was successful. Since this cost is expended before production, we continue to suppose firms have to borrow in the period before production in order to cover it. Firms will treat it as a fixed cost, sunk upon entry, since our unobservability assumptions mean bargaining's outcome will not be a function of research and appropriation levels.
- 6) The next period starts, other aggregate shocks are realised and production takes place.
- 7) The patent-holder brings court cases against any firms who produced but decided not to pay the rent. For simplicity, we assume the court always orders the violating firm to pay damages to the patent-holder, which are given as follows:
 - a) When the courts believe rents were not reasonable (i.e. $L_t^{\mathcal{R}}(i) > L_t^{\mathcal{R}^*}(i)$, where $L_t^{\mathcal{R}^*}(i)W_t$ is the level courts determine to be "reasonable royalties"), they set damages greater than $L_t^{\mathcal{R}^*}(i)W_t$, as *"the infringer would have nothing to lose, and everything to gain if he could count on paying only the normal, routine royalty non-*

⁶¹ Consider what happens as the time gap between offers increases. When this gap is large enough only one offer would be made per-period, meaning the patent-holder would make a take-it-or-leave-it offer giving (almost) nothing to the firm, which the firm would then accept.

⁶² The firm owner may, for example, face restrictions from starting businesses in future if as a result of the bargaining collapse they are unable to repay their creditors.

*infringers might have paid*⁶³. We assume excess damages over $L_t^{\mathcal{R}^*}(i)W_t$ are less than the patent-holder's legal costs however.

- b) When the courts consider the charged rent to have been reasonable (i.e. $L_t^{\mathcal{R}}(i) \leq L_t^{\mathcal{R}^*}(i)$) the courts award punitive damages of more than $\max\left\{L_t^{\mathcal{R}^*}(i)W_t, \left(\frac{1}{1-\varphi}\right)L_t^{\mathcal{R}}(i)W_t\right\}$, where φ is the bargaining power of the firm, in the sense of the generalized Nash bargaining solution.⁶⁴

Under this specification:

$$L_t^{\mathcal{R}}(i) = \min\{L_t^{\mathcal{R}^*}(i), (1 - \varphi)[L_t^{\mathcal{R}}(i) + L_t^{\mathcal{A}}(i) + L_t^{\mathcal{R}}(i) + L^{\mathcal{F}}]\}$$

since entry is fixed when bargaining takes place, since patent-holders know that bargaining to a rent level any higher than $L_t^{\mathcal{R}^*}(i)W_t$ will just result in them having to pay legal costs,⁶⁵ and since $[L_t^{\mathcal{R}}(i) + L_t^{\mathcal{A}}(i) + L_t^{\mathcal{R}}(i) + L^{\mathcal{F}}]W_t$ is equal to the production period profits of each firm in industry i , by the free entry condition.⁶⁶ Therefore, in equilibrium:

$$L_t^{\mathcal{R}}(i) = \min\{L_t^{\mathcal{R}^*}(i), L_t^{\mathcal{R}^\dagger}(i)\}, \quad (7.1)$$

where $L_t^{\mathcal{R}^\dagger}(i)$ is a solution to equations (3.2), (3.3) along with equation (3.4), (i.e. $L_t^{\mathcal{R}}(i) = \frac{1-\varphi}{\varphi} [L_t^{\mathcal{R}}(i) + L_t^{\mathcal{A}}(i) + L^{\mathcal{F}}]$) if one exists, or $+\infty$ otherwise. Because damages are always greater than $L_t^{\mathcal{R}^*}(i)W_t$, these rents will be sufficiently low to ensure firms are always prepared to licence the patent at the bargained price in equilibrium.

⁶³ *Panduit Corp. v. Stahl Brothers Fibre Works, Inc.*, 575 F.2d 1152, 1158 (6th Circuit 1978), cited in Pincus (1991).

⁶⁴ The level $\left(\frac{1}{1-\varphi}\right)L_t^{\mathcal{R}}(i)W_t$ is chosen to ensure that, with equilibrium rents, firms prefer not to produce at all rather than to produce without paying rents.

⁶⁵ The disagreement point is zero since it is guaranteed that $L_t^{\mathcal{R}}(i) \leq L_t^{\mathcal{R}^*}(i)$ and so punitive damages would be awarded were the firm to produce without paying rents, which, by construction, leaves them worse off than not producing.

⁶⁶ A similar expression can also be derived if we assume instead that courts guarantee infringers a fraction φ of production profits, or if we assume courts always award punitive damages but firms are able to hide a fraction φ of their production profits.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

Now suppose we are out of equilibrium and fewer firms than expected have entered. Since neither the patent-holder nor firms can observe how many firms have entered, and since firms arrive at the patent-holder sequentially, both sides will continue to believe that the equilibrium number of firms has entered and so rents will not adjust. On the other hand, suppose that (out of equilibrium) too many firms enter. When the first unexpected firm arrives at the patent-holder to negotiate, the patent-holder will indeed realise that too many firms have entered. However, since the firm they are bargaining with has no way of knowing this,⁶⁷ the patent-holder can bargain for the same rents as in equilibrium. Therefore, even out of equilibrium:

$$L_t^{\mathcal{R}}(i) = \min\{L_t^{\mathcal{R}*}(i), L_t^{\mathcal{R}\dagger}(i)\}$$

where we stress $L_t^{\mathcal{R}\dagger}(i)$ is not a function of the decisions any firm happened to take. This ensures that any solution of equations (3.2), (3.3) and (7.1) for research, appropriation and rents will also be an equilibrium, even allowing for the additional condition that the derivative of firm profits with respect to the number of firms must be negative at an optimum.

We now just have to pin down “reasonable royalties”, $L_t^{\mathcal{R}*}(i)W_t$. Certainly it must be the case that $L_t^{\mathcal{R}*}(i) \leq L_t^{\bar{\mathcal{R}}}(i)$, where $L_t^{\bar{\mathcal{R}}}(i)$ is the level of rents at which $J_t(i) = 1$, since rents so high that no one is prepared to pay them must fall foul of the courts’ desire to ensure licensees can make a profit.⁶⁸ However, since when $J_t(i) = 1$ the sole entering firm (almost) may as well be the patent-holder themselves, where possible the courts will set $L_t^{\mathcal{R}*}(i)$ sufficiently low to ensure that $J_t(i) > 1$ in equilibrium, again following the idea that

⁶⁷ Either they are a firm that thinks the equilibrium number of firms has entered, or they are a firm that thinks more than the equilibrium number of firms has entered, but that does not know whether the patent-holder has yet realised this.

⁶⁸ “...the very definition of a reasonable royalty assumes that, after payment, the infringer will be left with a profit.” *Georgia-Pacific Corp. v. U.S. Plywood-Champion Papers Corp.*, 446 F.2d 295, 299 & n.1 (2d Cir.), *cert. denied*, 404 U.S. 870 (1971), cited in Pincus (1991).

licensees ought to be able to make a profit. When there is a $J_t(i) > 1$ solution to equations (3.2), (3.3) and (3.4) already (i.e. $L_t^{\mathcal{R}\dagger}(i) < \infty$), the courts will just set $L_t^{\mathcal{R}*}(i)$ at the rent level that would obtain in that solution, thus preventing the possibility of $J_t(i) = 1$ being an equilibrium. It may be shown that for sufficiently large t such a solution is guaranteed to exist, so in this case $L_t^{\mathcal{R}*}(i) = L_t^{\mathcal{R}\dagger}(i) = L_t^{\mathcal{R}}(i)$.⁶⁹

7.4. The de-trended model

Below we give the equations of the stationary model to which the model described in section 3 converges as $t \rightarrow \infty$.

7.4.1. Households

- **Stochastic discount factor:** $\Xi_t = \frac{\Theta_t \hat{C}_{t-1}}{\Theta_{t-1} \hat{C}_t G_{A,t}}$, where $\hat{C}_t := \frac{C_t}{N_t A_t}$ is consumption per person in labour supply units and $G_{V,t}$ is the exponent of the growth rate of the variable V_t at t .
- **Labour supply:** $\Phi_t \hat{L}_t^{S^V} = \frac{\hat{W}_t}{\hat{C}_t}$, where $\hat{L}_t^S := \frac{L_t^S}{N_t}$ is labour supply per person and $\hat{W}_t := \frac{W_t}{A_t}$ is the wage per effective unit of labour supply.
- **Euler equation:** $\beta R_t \mathbb{E}_t[\Xi_{t+1}] = 1$, where R_t is the real interest rate.

7.4.2. Aggregate relationships

- **Aggregate mark-up pricing:** $\hat{W}_t = \frac{1}{1+\mu_{t-1}}$ where μ_{t-1} is the aggregate mark-up in period t .

⁶⁹ There may still be multiple solutions for rents (as (3.2), (3.3) and (3.4) might have multiple solutions), but of these only the one with minimal entry is really plausible, since this is both weakly Pareto dominant (firms always make zero profits and it may be shown that the patent-holder prefers minimal entry) and less risky for entering firms (if entering firms are unsure if the patent-holder will play the high rent or the low rent equilibrium, they are always better off assuming the high rent one since if that assumption is wrong they make strict profits, whereas had they assumed low rents but rents were in fact high they would make a strict loss).

Medium-frequency cycles and the remarkable near trend-stationarity of output.

- **Mark-up aggregation:** $\left(\frac{1}{1+\mu_t}\right)^{\frac{1}{\lambda}} = \left(\frac{1}{1+\mu_t^p}\right)^{\frac{1}{\lambda}} s_t + \left(\frac{1}{1+\eta\lambda}\right)^{\frac{1}{\lambda}} (1 - s_t)$, where $\mu_t^p = \mu_t(I_t)$ is the mark-up in any protected industry at $t + 1$, and $s_t := (1 - q)\frac{s_{t-1}}{G_{I,t}} + 1 - \frac{1}{G_{I,t}}$ is the proportion of industries that will produce a patent protected product in period $t + 1$.

- **Productivity aggregation:** $\left(\frac{\hat{A}_t}{1+\mu_{t-1}}\right)^{\frac{1}{\lambda}} = \left(\frac{1}{1+\mu_{t-1}^p}\right)^{\frac{1}{\lambda}} s_{t-1} + \left(\frac{\hat{A}_t^N}{1+\eta\lambda}\right)^{\frac{1}{\lambda}} (1 - s_{t-1})$, where $\hat{A}_t := \frac{A_t}{A_t^*}$ is aggregate productivity relative to the frontier⁷⁰ and

$$\hat{A}_t^N := \left[\left(\frac{1}{G_{A^*,t}}\right)^{\frac{1}{\lambda}} \left(\frac{q}{1/s_{t-2} - (1-q)}\right) + \left(\frac{\hat{A}_{t-1}^N}{G_{A^*,t}}\right)^{\frac{1}{\lambda}} \left(1 - \frac{q}{1/s_{t-2} - (1-q)}\right) \right]^{\lambda}$$

is the aggregate relative productivity of non-protected industries.

7.4.3. Firm decisions

- **Strategic in-industry pricing:** $\mu_t^p = \lambda \frac{\eta j_t^p}{j_t^p - (1-\eta)}$, where $j_t^p := J_t(I_t)$ is the number of firms in a protected industry performing research at t .

- **Firm research decisions:** $\frac{d_t}{p\mu_t^p} \mathbb{E}_t \Xi_{t+1} G_{Y,t+1} \hat{A}_{t+1}^{-\frac{1}{\lambda}} \frac{Z_{t+1} \hat{\mathcal{L}}_t^R}{1+\gamma Z_{t+1} \hat{\mathcal{L}}_t^R} = (1 - m_t^R) \mathbb{E}_t \Xi_{t+1} G_{Y,t+1} \hat{A}_{t+1}^{-\frac{1}{\lambda}}$,

where $\hat{\mathcal{L}}_t^R := A_t^{*\zeta} \Psi L_t^R$ is the amount of effective research conducted by firms in protected industries d_t is the value of $d_t(i)$ in protected industries and Z_t is the aggregate research-return shock. (This equation means that $\hat{\mathcal{L}}_t^R \approx \frac{p\mu_t^p}{d_t - p\gamma\mu_t^p}$.)

- **Research and appropriation payoff:** $G_{A^*,t} = (1 + \gamma Z_t \hat{\mathcal{L}}_{t-1}^R)^{\frac{1}{\gamma}}$.

⁷⁰ As a consequence, we have that $G_{A,t} = \frac{\hat{A}_t}{\hat{A}_{t-1}} G_{A^*,t}$.

- **Free entry of firms:** $\beta \frac{1}{\hat{l}_t \hat{j}_t^p} \frac{\mu_t^p}{1+\mu_t^p} \left(\frac{1+\mu_t}{1+\mu_t^p} \right)^{\frac{1}{\lambda}} \mathbb{E}_t \Xi_{t+1} G_{Y,t+1} \hat{A}_{t+1}^{-\frac{1}{\lambda}} = \frac{1}{p} \hat{\mathcal{L}}_t^R \frac{\hat{W}_t}{\hat{Y}_t}$, where $\hat{l}_t := \frac{l_t}{N_t A_t^{*-\zeta} \Psi}$

is the measure of products relative to its trend,⁷¹ and $\hat{Y}_t := \frac{Y_t}{N_t A_t}$ is output per person in labour supply units.

7.4.4. Inventor decisions

- **Inventor profits:** are given recursively by:

$$\hat{\pi}_t = \frac{1-p}{p} \hat{\mathcal{L}}_t^R \hat{W}_t \hat{j}_t^p + \beta(1-q) \mathbb{E}_t \Xi_{t+1} G_{A,t+1} G_{A^*,t+1}^{\zeta} \hat{\pi}_{t+1}, \text{ where } \hat{\pi}_t := \frac{\pi_t}{A_t A_t^{*\zeta}}.$$

- **Free entry of inventors:** Either $G_{I,t} \geq 1$ binds or $\Psi E^{\zeta} \mathcal{L}^1 \hat{W}_t \geq \hat{\pi}_t$ does.

7.4.5. Market clearing

- **Labour market clearing:**

$$\begin{aligned} \hat{\mathcal{L}}_t^S = & \Psi E^{\zeta} \mathcal{L}_t^1 \hat{l}_t \left(1 - \frac{1}{G_{I,t}} \right) + \hat{l}_t \mathcal{S}_t \hat{j}_t^p \hat{\mathcal{L}}_t^R \\ & + \hat{Y}_t \left[\left(\frac{1}{\hat{A}_t} \right)^{\frac{1}{\lambda}} \left(\frac{1+\mu_{t-1}}{1+\mu_{t-1}^p} \right)^{\frac{1+\lambda}{\lambda}} \mathcal{S}_{t-1} + \left(\frac{\hat{A}_t^N}{\hat{A}_t} \right)^{\frac{1}{\lambda}} \left(\frac{1+\mu_{t-1}}{1+\eta\lambda} \right)^{\frac{1+\lambda}{\lambda}} (1 - \mathcal{S}_{t-1}) \right]. \end{aligned}$$

- **Goods market clearing:** $\hat{Y}_t = \hat{C}_t$.

7.5. The extended de-trended model

Define $a := \frac{1}{(1-\alpha_p)(1-\iota_p)}$, $b := (1-\alpha_R)(1-\iota_R)$, $c := \left(\frac{1-\alpha_R}{1-\alpha_p} \alpha_p \xi_{KP} - \alpha_R \xi_{KR} \right) (1-\iota_R)$, $e :=$

$\xi_L + \frac{\alpha_p}{1-\alpha_p} \xi_{KP}$ and make the normalisation $\Psi = E = 1$.

⁷¹ This means $G_{I,t} = G_{N,t} G_{A^*,t}^{-\zeta} \frac{l_t}{\hat{l}_{t-1}}$.

7.5.1. Households

- **Budget constraint Lagrange multiplier:** $\frac{1}{\hat{C}_t} = \hat{m}_t^C + \beta \mathcal{h} \mathcal{h}^{\text{INT}} \mathbb{E}_t \frac{N_{t+1} \Theta_{t+1}}{N_t \Theta_t} \frac{1}{G_{A,t+1}^a \hat{C}_{t+1}}$, where

$\frac{\hat{m}_t^C}{A_t^a A_t^{*e} N_t}$ is the Lagrange multiplier on the budget constraint and $\hat{C}_t := \frac{\tilde{C}_t}{A_t^a A_t^{*e}} = \hat{C}_t -$

$$\mathcal{h} \frac{\hat{C}_{t-1}}{G_{A,t}^a G_{A^*,t}^e}.$$

- **Stochastic discount factor:** $\Xi_t = \frac{\Theta_t \hat{m}_t^C}{\Theta_{t-1} \hat{m}_{t-1}^C G_{A,t}^a G_{A^*,t}^e}$.

- **Labour supply:** $(1 + \lambda_L) \hat{w}_{1,t} = \hat{W}_t^{1+\nu} \frac{1+\lambda_L}{\lambda_L} \hat{w}_{2,t}$, where:

$$\hat{W}_t := \left[\frac{\hat{W}_t^{-\frac{1}{\lambda_L} - \nu} \left(\frac{G_P}{G_{P,t} G_{A,t}^a G_{A^*,t}^e} \hat{W}_{t-1} \right)^{-\frac{1}{\lambda_L}}}{1 - \nu} \right]^{-\lambda_L} \quad (\hat{W}_t A_t^a A_t^{*e} \text{ is the real wage set by a household}$$

that updates its wage at t), $\hat{W}_t := \frac{W_t}{A_t^a A_t^{*e}}$, and where $\hat{w}_{1,t}$ and $\hat{w}_{2,t}$ are the sums of costs

and benefits respectively from the wage setting first order conditions.⁷²

- **Euler equation:** $\beta R_t \mathbb{E}_t [\Xi_{t+1}] = 1$.

- **Investment decisions:** for $V \in \{P, R\}$:

$$\frac{1}{E_t^{\text{KV}}} = \Gamma_t \frac{\hat{R}_t^{\text{KV}}}{\delta^{\nu'} (u_t^{\nu})} \left[1 - Q^{\text{KV}}(G_{I^{\text{KV}^*},t}) - G_{I^{\text{KV},t}} Q^{\text{KV}'}(G_{I^{\text{KV}^*},t}) \right] +$$

$$\beta \mathbb{E}_t \Xi_{t+1} \Gamma_{t+1} \frac{\hat{R}_{t+1}^{\text{KV}}}{G_{A^*,t}^{\xi_{\text{KV}}} \delta^{\nu'} (u_{t+1}^{\nu})} G_{I^{\text{KV}^*},t+1}^2 Q^{\text{KV}'}(G_{I^{\text{KV}^*},t+1}), \text{ where } \hat{R}_t^{\text{KV}} := R_t^{\text{KV}} A_t^{\xi_{\text{KV}}} \text{ and } G_{I^{\text{KV}^*},t} =$$

$$G_{A^*,t}^{\xi_{\text{KV}}} \frac{E_t^{\text{KV}}}{E_{t-1}^{\text{KV}}} G_{I^{\text{KV},t}}$$

⁷² $\hat{w}_{1,t} = \Phi_t \left(\hat{W}_t^{\frac{1+\lambda_L}{\lambda_L}} \bar{L}_t^S \right)^\nu + \beta \nu \mathbb{E}_t \frac{\Theta_{t+1} N_{t+1}}{\Theta_t N_t} \left(\frac{G_P}{G_{P,t+1}} \frac{G_W}{G_{W,t+1}} \right)^{-\frac{1+\lambda_L}{\lambda_L}} \frac{\bar{L}_{t+1}^S}{\bar{L}_t^S} \left(\frac{G_P}{G_{P,t+1}} \frac{G_W}{G_{A,t+1}^a G_{A^*,t+1}^e} \right)^{-\nu \frac{1+\lambda_L}{\lambda_L}} \hat{w}_{1,t+1}$, $\hat{w}_{2,t} =$

$\hat{m}_t^C + \beta \nu \mathbb{E}_t \frac{\Theta_{t+1} N_{t+1}}{\Theta_t N_t} \left(\frac{G_P}{G_{P,t+1}} \frac{G_W}{G_{W,t+1}} \right)^{-\frac{1+\lambda_L}{\lambda_L}} \frac{\bar{L}_{t+1}^S}{\bar{L}_t^S} \frac{G_P}{G_{P,t+1}} \frac{G_W}{G_{A,t+1}^a G_{A^*,t+1}^e} \hat{w}_{2,t+1}$. This formulation avoids any explicit log-linearization and allows us to compute arbitrarily high order approximations to the model, for robustness checks. A similar formulation is used in Schmitt-Grohé and Uribe (2006).

- **Utilisation decisions:** for $V \in \{P, R\}$: $\frac{\hat{R}_t^{KV}}{\delta^{Vl}(u_t^V)} = \beta \mathbb{E}_t \Xi_{t+1} \frac{\hat{R}_{t+1}^{KV}}{G_{A^*,t}^{\xi_{KV}}} \left[u_{t+1}^V + \frac{1-\delta^V(u_{t+1}^V)}{\delta^{Vl}(u_{t+1}^V)} \right]$.
- **Capital accumulation:** for $V \in \{P, R\}$: $\hat{K}_t^V = \left(1 - \delta^V(u_t^V) \right) \frac{\hat{K}_{t-1}^V}{G_{N,t} G_{A,t}^a G_{A^*,t}^{e+\xi_{KV}}} + \Gamma_t E_t^{KV} \hat{I}_t^{KV} \left[1 - Q^{KV}(G_{I^{KV},t}) \right]$, where $\hat{K}_t^V := \frac{K_t^V}{N_t A_t^a A_t^{*e+\xi_{KV}}}$ and $\hat{I}_t^{KV} = \frac{I_t^{KV}}{N_t A_t^a A_t^{*e}}$ (hence $G_{I^{KV},t} = G_{N,t} G_{A,t}^a G_{A^*,t}^e \frac{\hat{I}_t^{KV}}{\hat{I}_{t-1}^{KV}}$).

7.5.2. Aggregate relationships

- **Aggregate mark-up pricing:** $\frac{[R_t^{KP \alpha_P} \hat{W}_t^{EP 1-\alpha_P}]^{1-l_P}}{l_P^{l_P} (1-l_P)^{1-l_P} [\alpha_P^{\alpha_P} (1-\alpha_P)^{1-\alpha_P}]^{1-l_P}} = \frac{1}{1+\mu_{t-1}}$ where $\hat{W}_t^{EP} := \frac{\hat{W}_t}{E_t^L \left[1 - Q^{LP} \left(\frac{\hat{L}_t^{TP}}{\hat{L}_{t-1}^{TP}} G_{N,t} G_{A^*,t}^{\xi_L} \right) \right]}$ and $\hat{L}_t^{TP} = \frac{L_t^{TP}}{N_t A_t^{*\xi_L}}$, where $L_t^T := A_t^{*\xi_L} E_t^L L_t^S$.
- **Mark-up aggregation:** $\left(\frac{1}{1+\mu_t} \right)^{\frac{1}{\lambda}} = \left(\frac{1}{1+\mu_t^P} \right)^{\frac{1}{\lambda}} s_t + \left(\frac{1}{1+\eta\lambda} \right)^{\frac{1}{\lambda}} (1-s_t)$, where $\mu_t^P = \mu_t(I_t)$ and $s_t := (1-q) \frac{s_{t-1}}{G_{I,t}} + 1 - \frac{1}{G_{I,t}}$.
- **Productivity aggregation:** $\left(\frac{\hat{A}_t}{1+\mu_{t-1}} \right)^{\frac{1}{\lambda}} = \left(\frac{1}{1+\mu_{t-1}^P} \right)^{\frac{1}{\lambda}} s_{t-1} + \left(\frac{\hat{A}_{t-1}^N}{1+\eta\lambda} \right)^{\frac{1}{\lambda}} (1-s_{t-1})$, where $\hat{A}_t := \frac{A_t}{A_t^*}$ and $\hat{A}_t^N := \left[\left(\frac{1}{G_{A^*,t}} \right)^{\frac{1}{\lambda}} \left(\frac{q}{1/s_{t-2} - (1-q)} \right) + \left(\frac{\hat{A}_{t-1}^N}{G_{A^*,t}} \right)^{\frac{1}{\lambda}} \left(1 - \frac{q}{1/s_{t-2} - (1-q)} \right) \right]^{\lambda}$.

7.5.3. Firm decisions

- **Strategic in-industry pricing:** $\mu_t^P = \lambda \frac{\eta \hat{f}_t^P}{\hat{f}_t^P - (1-\eta)}$, where $\hat{f}_t^P = J_t(I_t)$.
- **Firm research decisions:** $\frac{d_t}{\rho \mu_t^P} \mathbb{E}_t \Xi_{t+1} G_{Y,t+1} \hat{A}_{t+1}^{-\frac{1}{\lambda}} \frac{Z_{t+1} \hat{L}_t^R}{1+\gamma Z_{t+1} \hat{L}_t^R} = (1-m_t^R) \mathbb{E}_t \Xi_{t+1} G_{Y,t+1} \hat{A}_{t+1}^{-\frac{1}{\lambda}}$, where $\hat{L}_t^R := A_t^{*\zeta} X_t^{Rl_R} \left[K_t^{R\alpha_R} L_t^{R1-\alpha_R} \right]^{1-l_R}$ is the amount of effective research conducted by firms in protected industries.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

- **Research and appropriation payoff:** $G_{A^*,t} = (1 + \gamma Z_t \hat{L}_{t-1}^R)^{\frac{1}{\gamma}}$.
- **Free entry of firms:** $\beta \frac{1}{\hat{I}_t \hat{J}_t^P} \frac{\mu_t^P}{1 + \mu_t^P} \left(\frac{1 + \mu_t}{1 + \mu_t^P} \right)^{\frac{1}{\lambda}} \mathbb{E}_t \Xi_{t+1} G_{Y,t+1} \hat{A}_{t+1}^{-\frac{1}{\lambda}} = \frac{1}{p} \hat{L}_t^R \hat{C}_t$, where $\hat{I}_t := \frac{I_t}{N_t A_t^{a(1-\theta)} A_t^{*e-(c+\zeta)}}$ is the measure of products relative to its trend,⁷³ $\hat{Y}_t^{\text{GROSS}} := \frac{Y_t^{\text{GROSS}}}{N_t A_t^a A_t^{*e}}$ is gross output relative to trend and $\hat{C}_t := \frac{[R_t^{\text{KR}\alpha_R} \hat{W}_t^{\text{ER}^{1-\alpha_R}}]^{1-\iota_R}}{\iota_R^{\iota_R} (1-\iota_R)^{1-\iota_R} [\alpha_R^{\alpha_R} (1-\alpha_R)^{1-\alpha_R}]^{1-\iota_R}}$ is the marginal cost of a unit of research or invention, divided by $A_t^{a\theta} A_t^{*c}$ (where $\hat{W}_t^{\text{ER}} := \frac{\hat{W}_t}{E_t^L \left[1 - Q^{\text{LR}} \left(\frac{\hat{L}_t^{\text{TR}}}{\hat{L}_{t-1}^{\text{TR}}} G_{N,t} G_{A^*,t}^{\xi_L} \right) \right]}$ and $\hat{L}_t^{\text{TR}} = \frac{L_t^{\text{TR}}}{N_t A_t^{*\xi_L}}$).

7.5.4. Inventor decisions

- **Inventor profits:** are given recursively by:

$$\hat{\pi}_t = \frac{1-p}{p} \hat{L}_t^R \hat{C}_t \hat{J}_t^P + \beta(1-q) \mathbb{E}_t \Xi_{t+1} G_{A,t+1}^{a\theta} G_{A^*,t+1}^{c+\zeta} \hat{\pi}_{t+1}, \text{ where } \hat{\pi}_t := \frac{\pi_t}{A_t^{a\theta} A_t^{*c+\zeta}}.$$

- **Free entry of inventors:** Either $G_{I,t} \geq 1$ binds or $\mathcal{L}_t^I \hat{C}_t \geq \hat{\pi}_t$ does.

7.5.5. Market clearing

- **R&D expenditure:** $\text{RND}_t := \hat{C}_t \hat{I}_t \left[\mathcal{L}_t^I \left(1 - \frac{1}{G_{I,t}} \right) + \hat{L}_t^R \mathcal{S}_t \hat{J}_t^P \right]$.
- **Labour market clearing:** $E_t^L \hat{L}_t^S = \hat{L}_t^{\text{TY}} + \hat{L}_t^{\text{TR}}$, where $\hat{L}_t^S := \frac{L_t^T}{A_t^{*\xi_L} N_t E_t^L}$.
- **Production labour market clearing:** $\hat{W}_t \hat{L}_t^{\text{TY}} = E_t^L (1 - \alpha_P) (1 - \iota_P) \mathcal{J}_t \hat{Y}_t^{\text{GROSS}}$ where $\mathcal{J}_t := \frac{s_{t-1}}{1 + \mu_{t-1}^P} \left(\frac{1 + \mu_{t-1}}{\hat{A}_t} \frac{1 + \mu_{t-1}^P}{1 + \mu_{t-1}^P} \right)^{\frac{1}{\lambda}} + \frac{1 - s_{t-1}}{1 + \eta \lambda} \left(\frac{\hat{A}_t^N}{\hat{A}_t} \frac{1 + \mu_{t-1}}{1 + \eta \lambda} \right)^{\frac{1}{\lambda}}$ is a weighted measure of average inverse gross mark-ups.

⁷³ This means $G_{I,t} = G_{N,t} G_{A,t}^{a(1-\theta)} G_{A^*,t}^{e-(c+\zeta)} \frac{I_t}{\hat{I}_{t-1}}$.

- **R&D labour market clearing:** $\widehat{W}_t \widehat{L}_t^{\text{TR}} = E_t^L (1 - \alpha_R)(1 - \iota_R) \text{RND}_t$.
- **Capital markets clearing:** $u_t^P \widehat{K}_{t-1}^P \widehat{R}_t^{\text{KP}} = \alpha_P (1 - \iota_P) J_t \widehat{Y}_t^{\text{GROSS}}$, $u_t^R \widehat{K}_{t-1}^R \widehat{R}_t^{\text{KR}} = \alpha_R (1 - \iota_R) \text{RND}_t$
- **Goods market clearing:** $\widehat{Y}_t = \widehat{Y}_t^{\text{GROSS}} (1 - \iota_P J_t) - \iota_R \text{RND}_t - (1 - \varrho_{\text{GDP}}) \widehat{I}_t^{\text{R}} = \widehat{C}_t + \widehat{I}_t^{\text{P}} + \varrho_{\text{GDP}} \widehat{I}_t^{\text{R}}$, where \widehat{Y}_t is GDP over $N_t A_t^a A_t^{*e}$ and ϱ_{GDP} specifies the proportion of R&D capital investment that is measured in GDP. (Given R&D itself is not measured in GDP it is not obvious that this equals 1.)
- **Monetary rule:**

$$\frac{R_t^{\text{NOM}}}{R^{\text{NOM}}} =$$

$$\left(\frac{R_t^{\text{NOM}}}{R^{\text{NOM}}} \right)^{\rho_{R^{\text{NOM}}}} \left[\left(\frac{G_{P,t}}{G_{P,t}^*} \right)^{\mathcal{M}_P} \left(\frac{E_{t-1}^{\text{KP}} G_{A^*,t}^{\xi_{\text{KP}}}}{E_t^{\text{KP}} G_{A^*,t}^{\xi_{\text{KP}}}} \right)^{\mathcal{M}_{\text{PKP}}} \left(\frac{E_{t-1}^{\text{KR}} G_{A^*,t}^{\xi_{\text{KR}}}}{E_t^{\text{KR}} G_{A^*,t}^{\xi_{\text{KR}}}} \right)^{\mathcal{M}_{\text{PKR}}} \left(\frac{G_{W,t}}{G_W} \right)^{\mathcal{M}_W} \left(\frac{\widehat{R}_t^{\text{KP}}}{\widehat{R}^{\text{KP}}} \right)^{\mathcal{M}_{\text{RKP}}} \left(\frac{\widehat{R}_t^{\text{KR}}}{\widehat{R}^{\text{KR}}} \right)^{\mathcal{M}_{\text{RKR}}} \Theta_t^{\mathcal{M}_\Theta} \delta_t^{-\mathcal{M}_\delta} \right]^{1 - \rho_{R^{\text{NOM}}}}$$

$$\left[\left(\frac{\widehat{Y}_t}{\widehat{Y}} \right)^{\mathcal{M}_Y} \left(\frac{G_{Y,t}/G_{N,t}}{G_Y/G_N} \right)^{\mathcal{M}_G} \right]^{1 - \rho_{R^{\text{NOM}}}} \exp \epsilon_{R^{\text{NOM}},t}$$

7.5.6. Observation equations

- **Nominal output growth:** $g_{Y,t} + g_{P,t} + \text{me}_{Y,t} - \text{me}_{Y,t-1}$, where $g_{Y,t} = \log \left(\frac{\widehat{Y}_t}{\widehat{Y}_{t-1}} G_{N,t} G_{A,t}^a G_{A^*,t}^e \right)$.
- **Consumption price inflation:** $g_{P,t} + \text{me}_{\text{PC},t} - \text{me}_{\text{PC},t-1}$.
- **Investment price inflation:** $g_{P,t} + g_{P^I,t} + \text{me}_{\text{PI},t} - \text{me}_{\text{PI},t-1}$, where:

$$G_{P^I,t} = \sqrt{\frac{\left(\frac{E_{t-1}^{\text{KP}} \widehat{I}_{t-1}^{\text{KP}}}{E_t^{\text{KP}} G_{A^*,t}^{\xi_{\text{KP}}}} + \varrho_{\text{GDP}} \frac{E_{t-1}^{\text{KR}} \widehat{I}_{t-1}^{\text{KR}}}{E_t^{\text{KR}} G_{A^*,t}^{\xi_{\text{KR}}}} \right) \left(\widehat{I}_t^{\text{KP}} + \varrho_{\text{GDP}} \widehat{I}_t^{\text{KR}} \right)}{\left(\widehat{I}_{t-1}^{\text{KP}} + \varrho_{\text{GDP}} \widehat{I}_{t-1}^{\text{KR}} \right) \left(\frac{E_t^{\text{KP}} \widehat{I}_t^{\text{KP}} G_{A^*,t}^{\xi_{\text{KP}}}}{E_{t-1}^{\text{KP}}} + \varrho_{\text{GDP}} \frac{E_t^{\text{KR}} \widehat{I}_t^{\text{KR}} G_{A^*,t}^{\xi_{\text{KR}}}}{E_{t-1}^{\text{KR}}} \right)}}$$

- **Population growth:** $g_{N,t} + \text{me}_{N,t} - \text{me}_{N,t-1}$.
- **Demeaned labour supply:** $l_t^{\text{S}} + \text{me}_{\text{LS},t}$.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

- **R&D share:** $\log\left(\frac{\text{RND}_t + \varrho_{\text{RND}} \hat{I}_t^R}{\hat{Y}_t}\right) + \text{me}_{\text{RND},t}$, where ϱ_{RND} is the proportion of R&D capital investment that is measured in the NIPA R&D measure. ($\varrho_{\text{GDP}} + \varrho_{\text{RND}} \leq 1$).
- **Consumption share:** $\log\left(\frac{\hat{C}_t}{\hat{Y}_t}\right) + \text{me}_{\text{C},t}$.
- **Labour share:** $\log\left(\frac{\hat{W}_t \hat{L}_t^S}{\hat{Y}_t}\right) + \text{me}_{\text{L},t}$.
- **Depreciation share:** $\log\left(\frac{\delta^Y(u_t^Y) \hat{R}_{t-1}^Y}{\hat{Y}_t (G_{N,t} G_{A,t}^a G_{A^*,t}^{e+\xi_{\text{KY}} E_t^{\text{KY}}})} + \varrho_{\text{GDP}} \frac{\delta^R(u_t^R) \hat{R}_{t-1}^R}{\hat{Y}_t (G_{N,t} G_{A,t}^a G_{A^*,t}^{e+\xi_{\text{KR}} E_t^{\text{KR}}})}\right) + \text{me}_{\text{D},t}$.
- **Demeaned nominal interest rates:** $\log\left(\frac{R_t^{\text{NOM}}}{R^{\text{NOM}}}\right) + \text{me}_{\text{R},t}$.
- **Capacity utilisation:** $\frac{u_t^Y \frac{\hat{R}_{t-1}^Y}{G_{A^*,t}^{\xi_{\text{KY}} E_t^{\text{KY}}} + u_t^R \varrho_{\text{GDP}} \frac{\hat{R}_{t-1}^R}{G_{A^*,t}^{\xi_{\text{KR}} E_t^{\text{KR}}}}}{\frac{\hat{R}_{t-1}^Y}{G_{A^*,t}^{\xi_{\text{KY}} E_t^{\text{KY}}} + \varrho_{\text{GDP}} \frac{\hat{R}_{t-1}^R}{G_{A^*,t}^{\xi_{\text{KR}} E_t^{\text{KR}}}}} + \text{me}_{\text{U},t}$. (The capital stocks enter here in order to correctly weight to produce the average utilisation.)
- **BAA-AAA Spread:** $\zeta_0 - \zeta_1 \log \Gamma_t + \text{me}_{\text{S},t}$.

7.6. Data details

- **Nominal output growth** (1947Q2 – 2011Q2), from NIPA table 1.1.5.
- **Consumption price inflation** (1947Q2 – 2011Q2), including non-durables and durables (from NIPA table 1.1.4) and government consumption⁷⁴ (from NIPA table 3.9.4) and excluding education⁷⁵ (from NIPA tables 2.4.4⁷⁶ and 3.15.4⁷⁷).

⁷⁴ We are implicitly making the optimistic assumption that government consumption is a perfect substitute for private consumption. This is a simplifying shortcut to save us modelling government consumption.

⁷⁵ Removing education from the consumption share brings it substantially closer to stationarity, so it is important to do the same for the price level too. The price disaggregation necessary to remove education was performed by inverting the Fisher formula, which, due to its approximate aggregation property (Diewert 1978) is sufficiently accurate.

- **Investment price inflation** (1947Q2 – 2011Q2), including education (data sources as for consumption price inflation).
- **Population growth** (1948Q2 – 2011Q2), X-12 seasonally adjusted, from the BLS’s Civilian Non-institutional Population Over 16 series.
- **Labour supply per capita** (1948Q1 – 2011Q2), from NIPA table 6.9, interpolated to quarterly using the Litterman (1983) method, with “Business Sector: Hours of All Persons” from the BEA as a high frequency indicator.
- **R&D share** (1959Q1 – 2007Q4), given by R&D expenditure from NIPA R&D Satellite Account (1959-2007) table 2.1, over GDP from NIPA table 1.1.5, interpolated to quarterly using the Litterman (1983) method with GDP as the high frequency indicator.
- **Consumption share** (1947Q1 – 2011Q2), given by consumption of durables and non-durables (from NIPA table 1.1.5) plus government consumption (from NIPA table 3.9.5) minus education expenditure (from NIPA table 2.4.5⁷⁸ and NIPA table 3.15.5⁷⁹) all over GDP (from NIPA table 1.1.5).⁸⁰
- **Labour share** (1947Q1 – 2011Q2), given by compensation of employees paid from NIPA table 1.10, over GDP (from NIPA table 1.1.5).

⁷⁶ Interpolated to quarterly using the Litterman (1983) method, with consumption and investment prices as indicators (from NIPA table 1.1.4).

⁷⁷ Extrapolated back to 1947 using the Litterman (1983) method, with government consumption and investment prices (from NIPA table 3.9.4) and private education prices (from NIPA table 2.4.4) as indicators, then interpolated to quarterly using the same method with government consumption and investment prices (from NIPA table 3.9.4) as high frequency indicators.

⁷⁸ Interpolated to quarterly using the Litterman (1983) method, with consumption and investment as indicators (from NIPA table 1.1.5).

⁷⁹ Extrapolated back to 1947 using the Litterman (1983) method with log-linearly interpolated data from the National Centre for Education Statistics, Digest of Education Statistics 2010, table 29 as an indicator, along with government consumption and investment (from NIPA table 3.9.5) and private education expenditure (from NIPA table 2.4.5). Then interpolated using the same method with government consumption and investment (from NIPA table 3.9.5) as high frequency indicators.

⁸⁰ In fitting this to the model, we are implicitly treating net exports as investment.

Medium-frequency cycles and the remarkable near trend-stationarity of output.

- **Depreciation share** (1947Q1 – 2011Q2), given by consumption of fixed capital from NIPA table 1.10, over GDP (from NIPA table 1.1.5).
- **Nominal interest rates** (1947Q1 – 2011Q2), in particular, the 3-month Treasury bill secondary market rate, from the FRB, release H.15.
- **Capacity utilisation** (1967Q1 – 2011Q2), (total industry) from the FRB, release G.17, table 7.
- **BAA-AAA Spread** (1947Q1 – 2011Q2), from the FRB, release H.15.

7.7. Estimated parameters

Any parameters in bold are fixed rather than estimated. All values are reported to three significant figures, except those below 10^{-4} which are rounded down to zero, those which are of the form $1 + x$, with $|x| < 0.1$ in which case we give x to three significant figures, percentages, which are given to one decimal place, and approximate standard errors (in brackets) which are given to two significant figures.

Variable	Value	Variable	Value
ν	0.250 (0.0056)	β	0.99
h	0.253 (0.0041)	h^{INT}	0.0151 (0.0032)
h^{LS}	0 (0)	ν	0.826 (0.0042)
λ	0.320 (0.00054)	λ_L	0.170 (0.0041)
ρ	0.0427 (0.00021)	q	0.0374 (0.00030)
ρ_R^{NOM}	0.615 (0.013)	\mathcal{M}_P	1.0275 (0.0059)
\mathcal{M}_{PKP}	0 (0)	\mathcal{M}_{PKR}	0 (0)
\mathcal{M}_{RKP}	0.0509 (0.0016)	\mathcal{M}_{RKR}	0 (0)
\mathcal{M}_Θ	0 (0)	$\mathcal{M}_{\bar{\delta}}$	0.0108 (0.0074)
\mathcal{M}_Y	0 (0)	\mathcal{M}_G	0 (0)
\mathcal{M}_W	0 (0)		
$\exp \zeta_0$	2.57 (2.9×10^{-5})	ζ_1	872 (880)
ϱ_{GDP}	0.494 (0.013)	ϱ_{RND}	0.506 (0.013)
ζ	0 (0)	ξ_L	0.0859 (0.0012)
ξ_{KP}	0.0828 (0.00053)	ξ_{KR}	2.73 (0.0094)
α_P	0.201 (0.00040)	α_R	0.996 (7.4×10^{-6})
ι_P	0.0427 (0.0011)	ι_R	0.178 (0.0032)
$\delta^P(u^P)$	0.0189 (7.5×10^{-5})	$\delta^R(u^R)$	0.0284 (0.00062)
$\delta^{P'}(u^P)$	0.0413 (0.00011)	$\delta^{R'}(u^R)$	0.0501 (0.00063)
$\delta^{P''}(u^P)$	1.64 (0.035)	$\delta^{R''}(u^R)$	133 (9.4)
$\frac{d}{d\bar{\delta}} \log \delta^P(u^P)$	1	$\frac{d}{d\bar{\delta}} \log \delta^R(u^R)$	64.2 (1.5)
$\frac{d}{d\bar{\delta}} \log \delta^{P'}(u^P)$	64.2 (1.5)	$\frac{d}{d\bar{\delta}} \log \delta^{R'}(u^R)$	0 (0)
$Q^{P''}(G_{I_{\text{KP}}^*})$	0.00533 (0.0012)	$Q^{R''}(G_{I_{\text{KR}}^*})$	62.6 (4.0)
$Q^{\text{LP}'}(G_{L_{\text{TP}}})$	0.0875 (0.0047)	$Q^{\text{LR}'}(G_{L_{\text{TR}}})$	0 (0)

Table 4: Estimated parameters, excluding shocks.

Variable	V (i.e. steady-state)	ρ_V	$100\sigma_V$	p-value on 1 lag LM-test ⁸¹
Φ	1.0349 (0.0047)	0.815 (0.010)	2.46 (0.16)	0
Θ	1	0.443 (0.0056)	0.0231 (0.0114)	0.0318
G_N	1.00372 (1.4×10^{-5})	0.0675 (0.019)	0.103 (0.0021)	0.146
\mathcal{L}^I	7.26 (0.034)	0 (0)	0 (0)	0
Z	1	0	0 (0)	0
Γ	1	0 (0)	0 (0)	0.000926
E^L	1	0.614 (0.0056)	0 (0)	0.725
E^{KP}	1	0 (0)	0 (0)	0
E^{KR}	1	0.664 (0.0071)	0.000360 (0.00012)	0.000148
$G_{P,t}^*$	1.00851 (6.1×10^{-6})	0.887 (0.00027)	0 (0)	0.161
η	0.169 (0.00024)	0.0605 (0.2)	0.0147 (0.012)	0
γ	18.6 (0.054)	0 (0)	0 (0)	0
$\exp \tilde{\delta}$	1	0.862 (0.0027)	0.403 (0.011)	0.958
R_t^{SHOCK}	1	0	0.00824 (0.00075)	0.464

Table 5: Estimated parameters from non-measurement error shocks, tests of misspecification of their residuals.

Each shock takes the form $\log V_t = (1 - \rho_V) \log V + \rho_V \log V_{t-1} + \sigma_V \epsilon_{v,t}$, where $\epsilon_{v,t} \sim \text{NIID}(0,1)$.

⁸¹ Bold values indicate the cases in which we cannot reject the null hypothesis of no auto-correlation at 1%. The test uses heteroskedasticity robust standard errors. The lag length of 1 was preferred by the AIC, AICc and BIC criteria for all variables.

Variable	Φ	Θ	G_N	\mathcal{L}^I	Z	Γ	E^L	E^{KP}	E^{KR}	$G_{P,t}^*$	η	γ	R_t^{SHOCK}	$\exp \tilde{\delta}$
Nom. output growth	17.9	0.0	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.5	81.1
Con. price inflation	37.5	0.0	2.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	59.6
Inv. price inflation	37.1	0.0	2.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4	60.1
Population growth	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lab. supply per capita	60.4	0.0	0.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	38.8
R&D share	2.1	0.0	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	97.6
Consumption share	45.0	0.0	0.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	54.3
Labour share	1.8	0.0	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	97.9
Depreciation share	45.3	0.0	1.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	52.7
Nom. interest rates	41.5	0.0	2.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	56.4
Capacity utilisation	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	99.9
BAA-AAA Spread	0.0	0.0	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 6: Percentage non-measurement error variance decomposition of the observation variables.⁸²

⁸² Bold values are larger than 1%.

Chapter 2: Learning from learners

Tom Holden¹, Balliol College, University of Oxford

Abstract: *Traditional macroeconomic learning algorithms are misspecified when all agents are learning simultaneously. In this paper, we produce a number of learning algorithms that do not share this failing, and show that this enables them to learn almost any solution, for any parameters, implying learning cannot be used for equilibrium selection. As a by-product, we are able to show that when all agents are learning by traditional methods, all deep structural parameters of standard new-Keynesian models are identified, overturning a key result of Cochrane (2009; 2011). This holds irrespective of whether the central bank is following the Taylor principle, irrespective of whether the implied path is or is not explosive, and irrespective of whether agents' beliefs converge. If shocks are observed then this result is trivial, so following Cochrane (2009) our analysis is carried out in the more plausible case in which agents do not observe shocks.*

Keywords: *Identification, Learnability, Limited Information, Indeterminacy, Taylor Rules*

JEL Classification: *E3, E52, D83*

¹ Contact address: School of Economics, FBEL, University of Surrey, Guildford, GU2 7XH, England

Contact email: thomas.holden@gmail.com Web-site: <http://www.tholden.org/>

The author would like to thank Charles Brendon, Ana-Maria Dumitru, Martin Ellison, Giovanni Forchini, Bruce Hansen, Paul Levine, Sophocles Mavroeidis, Bent Nielsen, Joseph Pearlman, David Vines and Simon Wren-Lewis, for assistance and helpful comments on versions of this paper and its antecedents. Financial support provided to the author by the ESRC is also greatly appreciated. The author gratefully acknowledges the use of the University of Surrey FBEL cluster for the numerical computations performed in this paper.

1. Introduction

There is a contradiction at the heart of the traditional approach to macroeconomic learning (Marcet and Sargent (1989), Evans and Honkapohja (2001)). In this literature, each of the agents in an economy is supposed to run a regression that is correctly specified when all the other agents know the true law of motion. Were it indeed the case that only one agent in the economy had partial information about the economy's law of motion, then this agent's regression would always converge to the true law of motion, meaning that "learnability" in this weak sense is of no use for equilibrium selection. The literature supposes instead that all agents are learning at the same time, yet they continue to run a regression that is only correctly specified when everyone else has full information. As a result, these agents would be readily able to detect the misspecification in their regression, through evidence of serially correlated errors, or parameter non-constancy. This misspecification is most clear precisely when learning fails, meaning a finding of non-learnability via the traditional method only implies that agents would switch from that traditional method to a more sophisticated one. In this paper, we demonstrate the existence of a family of learning mechanisms that remain correctly specified when all agents are learning simultaneously.

Along the way, we will answer three challenges raised by Cochrane (2009) (directly or otherwise). Firstly, we will show that the non-observability of shocks does not pose any fundamental challenges either to learning, or to the formation of rational expectations, and we give general conditions under which a rational expectations equilibrium is precisely implementable without observing shocks.² Secondly, we show that serially correlated monetary policy shocks do not prevent Taylor-rule parameter identification, at least when everyone is learning at the same time, whether or not the central bank is following active policy. Finally, we demonstrate a learning mechanism capable of learning stationary minimal

² In general a Kalman filter must be used as in Pearlman, Currie, and Levine (1986) or Ellison and Pearlman (2011), and impulse responses will differ.

state variable (McCallum 1983) solutions whenever they exist, and another that may converge towards any sunspot solution, including explosive ones,³ though a simple extension of our mechanism will rule out the latter when (and only when) they are prohibited by transversality or non-explosiveness constraints. Since, new-Keynesian models generally have no such constraints ruling out explosive paths for inflation (Cochrane 2011), in such models there is no guarantee that the stationary minimal state variable solution will be learnt, meaning that Cochrane (2009) was correct to conclude that learnability could not “save” the standard logic of new-Keynesian models.

The structure of our paper is as follows. In section 2 and the first appendix (7.1), we derive the general solution of a rational expectations model, under determinacy and indeterminacy, when shocks are unobserved. The resulting reduced form solution will be the basis of all of the learning mechanisms considered. The presence of sunspot shocks in the general solution will be key to our proof of structural parameter identification when agents are learning. In section 3, we show that an awareness that everyone else is learning is sufficient to achieve identification even when other agents are learning using a traditional method. Then in section 4, we introduce our family of sophisticated learning algorithms under which everyone in the economy realises everyone else is learning at the same time.

³ We cannot guarantee asymptotic convergence to explosive solutions, nonetheless beliefs will at least initially approach these solutions, and they will certainly diverge from beliefs under the stationary minimal state variable solution.

2. FREE solutions

2.1. Motivating example

Suppose, following Cochrane (2009), that the central bank follows the Taylor rule:

$$i_t = r + \frac{1}{\beta}(x_t - \gamma - \sigma s_t), \quad (2.1)$$

where x_t is the inflation rate, r is the constant real interest rate, $\frac{\gamma}{1-\beta}$ is the inflation target

and s_t is the monetary policy shock which is given by:

$$s_t = \rho s_{t-1} + \varepsilon_t,$$

with $\varepsilon_t \sim \text{NIID}(0,1)$. From the Fisher equation, we also have that:

$$i_t = r + \mathbb{E}_t x_{t+1}.^4 \quad (2.2)$$

Hence, from combining (2.1) and (2.2):

$$x_t = \beta \mathbb{E}_t x_{t+1} + \gamma + \sigma s_t.$$

More generally, there might also be a lag term in the model. Here, this would emerge if the central bank used the rule:

$$i_t = r + \frac{1}{\beta} \left((1-\alpha)x_t + \alpha \Delta x_t - \gamma - \sigma s_t \right)$$

which punishes accelerating inflation, and leads to the general univariate model:

$$x_t = \alpha x_{t-1} + \beta \mathbb{E}_t x_{t+1} + \gamma + \sigma s_t. \quad (2.3)$$

We work with this general model not because we believe central banks respond to inflation acceleration, but because in its multivariate version this structure encompasses all linear macroeconomic models, and we wish to make clear nothing we say is specific to the $\alpha = 0$ case.

⁴ Throughout this document, variables with t subscripts are in the information set under which \mathbb{E}_t is taken.

The crucial thing to note about (2.3) is that since the transversality conditions of the consumer's optimisation problem do not restrict inflation, when solving this model there is no justification for restricting ourselves to stationary solutions.⁵

2.2. Solution

For the time being, we suppose that all the agents in the economy have full knowledge of the values of α , β , γ , ρ and σ , and may observe x_t (and its lags), and $\mathbb{E}_{t-1}x_t$ (and its lags), at t . In our motivating example, the observability of expectations just requires nominal interest rates to be observable, thanks to the constant real interest rate, and the Fisher equation, (2.2). In reality, expectations may still be observed thanks to the survey of professional forecasters (or, more plausibly, media reports based on economic pundit's expectations). Expectations are also effectively observable if agents have access to prices from futures markets, or if they know that all other agents are forming expectations via the same mechanism. The traditional learning literature usually assumes homogeneous beliefs across agents, and we will continue to do so here, so in the models we work with, even in the absence of observable nominal and real interest rates, or observable futures contracts, aggregate expectations will always be observable.

We do not assume however that agents may observe s_t or ε_t . As pointed out by Cochrane (2009), that most shocks in DSGE models should be observable is rather implausible, thus ruling out rational expectations equilibria (REE) which require the observability of shocks seems like a minimal sensible restriction. We call the set of resulting equilibria the feasible rational expectations equilibria (FREE) of the original model. The key trick that enables agents to form expectations without seeing shocks is the fact that current news about past

⁵ If the Taylor rule is the result of optimal policy on behalf of the central bank, then there will in general be a transversality constraint coming from the central bank's optimisation problem that restricts inflation. But since it is consumer inflation expectations that determine the solution picked, the central bank's transversality constraint does not rule out explosive solutions, conditional on them using a Taylor rule.

expectational errors is informative about the current shock. Thus, in general, agents will form expectations as a linear function of their lagged expectations.

To see this, let us begin by defining the expectational error by $\eta_t := x_t - \mathbb{E}_{t-1}x_t$. Now, normally when solving rational expectations models we choose η_t to rule out explosive solutions, but here this is not justified, due to the lack of a consumer transversality condition on inflation. Thus there is a REE to the model for any η_t satisfying $\mathbb{E}_{t-1}\eta_t = 0$. Without loss of generality then, we may assume (following Lubik and Schorfheide (2003)) that $\eta_t = m_{\varepsilon,t-1}\varepsilon_t + m'_{\zeta,t-1}\zeta_t$, for some sunspot shock ζ_t (possibly a vector) satisfying $\mathbb{E}_{t-1}\zeta_t = 0$, $\mathbb{E}_{t-1}\varepsilon_t\zeta_t = 0$ and $\mathbb{E}_{t-1}\zeta_t\zeta_t' = I$, and some possibly time-varying belief parameters $m_{\varepsilon,t-1}$ and $m_{\zeta,t-1}$, known at $t - 1$. (There is no reason why agents should always believe in the same set of sunspot shocks.)

Under the assumption then that $m_{\varepsilon,t-1} \neq 0$ for all t , subtracting ρ times the first lag of (2.3) from (2.3), gives:

$$\begin{aligned} x_t &= (\alpha + \rho)x_{t-1} - \alpha\rho x_{t-2} + \beta\mathbb{E}_t x_{t+1} - \beta\rho\mathbb{E}_{t-1}x_t + (1 - \rho)\gamma + \sigma\varepsilon_t \quad (2.4) \\ &= (\alpha + \rho)x_{t-1} - \alpha\rho x_{t-2} + \beta\mathbb{E}_t x_{t+1} - \beta\rho\mathbb{E}_{t-1}x_t + (1 - \rho)\gamma \\ &\quad + \sigma \frac{x_t - \mathbb{E}_{t-1}x_t - m'_{\zeta,t-1}\zeta_t}{m_{\varepsilon,t-1}}. \end{aligned}$$

Hence providing $\beta \neq 0$ ⁶:

$$\begin{aligned} \mathbb{E}_t x_{t+1} &= \frac{1}{\beta} \left(1 - \frac{\sigma}{m_{\varepsilon,t-1}} \right) x_t - \frac{1}{\beta} (\alpha + \rho)x_{t-1} + \frac{1}{\beta} \alpha\rho x_{t-2} \\ &\quad + \left(\rho + \frac{1}{\beta} \frac{\sigma}{m_{\varepsilon,t-1}} \right) \mathbb{E}_{t-1}x_t - \frac{1}{\beta} (1 - \rho)\gamma + \frac{1}{\beta} \frac{\sigma}{m_{\varepsilon,t-1}} m'_{\zeta,t-1}\zeta_t, \end{aligned} \quad (2.5)$$

which enables agents to form rational expectations without observing the value of shocks (i.e. s_t or ε_t). Thus providing $\beta \neq 0$, almost all of the model's REE are FREE.

⁶ Automatic in the particular case under consideration, but in other models there may be particular parameters for which expectations cease to matter, and in the multivariate case, β may not be invertible.

When $|\alpha + \beta| < 1$, the unique stationary minimal state variable (MSV) solution corresponds to setting $m_{\varepsilon,t} \equiv m_{\varepsilon}^{\text{MSV}} := \sigma \left[\frac{1}{2} - \beta\rho + \frac{1}{2}\sqrt{1 - 4\alpha\beta} \right]^{-1}$ and $m_{\zeta,t} \equiv m_{\zeta}^{\text{MSV}} := 0$. To see this, let us first define:

$$v_t := \mathbb{E}_t x_{t+1} - a_1^{\text{MSV}} x_t - a_2^{\text{MSV}} x_{t-1} - c^{\text{MSV}}$$

where $a_1^{\text{MSV}} := \rho + \frac{1 - \sqrt{1 - 4\alpha\beta}}{2\beta}$, $a_2^{\text{MSV}} := -\rho \frac{1 - \sqrt{1 - 4\alpha\beta}}{2\beta}$ and $c^{\text{MSV}} := \frac{2(1-\rho)\gamma}{1 - 2\beta + \sqrt{1 - 4\alpha\beta}}$, so $\mathbb{E}_t x_{t+1} = a_1^{\text{MSV}} x_t + a_2^{\text{MSV}} x_{t-1} + c^{\text{MSV}} + v_t$ for all t .

Then, when $m_{\varepsilon,t} \equiv m_{\varepsilon}^{\text{MSV}}$ and $m_{\zeta,t} \equiv m_{\zeta}^{\text{MSV}}$, from (2.5):

$$\mathbb{E}_t x_{t+1} = a_1^{\text{MSV}} x_t + a_2^{\text{MSV}} x_{t-1} + c^{\text{MSV}} + \frac{1 + \sqrt{1 - 4\alpha\beta}}{2\beta} v_{t-1},$$

i.e. $v_t = \frac{1 + \sqrt{1 - 4\alpha\beta}}{2\beta} v_{t-1}$. Now when $|\alpha + \beta| < 1$ and $\alpha\beta < 1/4$ (so x_t is real), $\frac{1 + \sqrt{1 - 4\alpha\beta}}{2\beta} > 1$,

therefore x_t is stationary if and only if $v_t = 0$ for all t , i.e. if and only if expectations always take this minimum state variable form. However, since current expectations are not constrained to render past expectations rational, if agents find themselves off the $v_t = 0$ path, it is still rational for them to jump back onto it, at least if x_t is constrained to be stationary.

Linear models such as this have two MSV solutions, however only one of them will be stationary under determinacy. In the below we refer to the MSV solution that is stationary under determinacy as the SMSV solution.

2.3. Generalization

All our analysis in the body of this paper will be confined to the univariate case; however, the tricks used above to express expectations as a function of observables carry over to the multivariate case, and the case in which at least some combinations of variables are constrained by transversality. This is discussed in the first appendix, section 7.1, where we

provide a range of necessary and/or sufficient conditions for the existence of FREE solutions in multivariate models. Particularly intuitive results include the facts that:

- if the model is completely indeterminate (perhaps because of a lack of transversality conditions), so there are as many degrees of freedom in expectations as there are variables, and there are at most as many shocks as variables, then almost all REE are FREE;
- there is always a REE with the form $\mathbb{E}_t x_{t+1} = T_{-1,21}x_{t-1} + T_{-1,22}\mathbb{E}_{t-1}x_t + T_{\mu,2} + T_{s,2}s_t$, which is always a FREE when $\dim s_t = 1$, and is a FREE more generally providing:
 - $T_{s,2}$ has linearly independent columns,
 - the number of explosive (or transversality violating) roots is greater or equal to $\dim s_t$,
 - a further technical condition is satisfied;
- if the unobserved shocks are not serially correlated, and if for any linear combination of shocks which does not appear in the transversality-violating block, that same linear combination does not appear anywhere in the model (i.e. agents can back out the value of relevant shocks from observing jump variables), then the model has at least one FREE, and a continuum under indeterminacy.

In all cases, the FREE solution to the model takes the form:

$$\mathbb{E}_t x_{t+1} = \mathcal{A}_1 x_t + \mathcal{A}_2 x_{t-1} + \mathcal{A}_3 x_{t-2} + \mathcal{B}_1 \mathbb{E}_{t-1} x_t + \mathcal{B}_2 \mathbb{E}_{t-2} x_{t-1} + c + d'_{1,t-1} \zeta_t,$$

which is identical to the univariate case, except for the extra lag on expectations.

These results hopefully go some way to reassuring the reader that although from here on in we will be focussing on the univariate case, the non-observability of shocks does not cause any additional problems when we generalise to the multivariate case.⁷

3. Learning (and identifying) from unsophisticated learners

We now turn to the formation of expectations when the values of $\alpha, \beta, \gamma, \rho, \sigma, m_{\varepsilon,t}$ and $m_{\zeta,t}$ are not common knowledge. Before introducing our misspecification free learning methods in section 4, we address the issue of parameter identification when the agents in an economy are using a traditional learning method. For the duration of this section, we also assume it is common knowledge that $m_{\varepsilon,t}$ and $m_{\zeta,t}$ are constant across time, since the traditional models of macroeconomic learning cannot deal with actual laws of motion (ALMs) with time varying parameters.

3.1. Set-up

Under the saddle-path learning method of Ellison and Pearlman (2011), agents learn using the same rule they use to form expectations. Under the FREE solution to (2.3), given in equation (2.5), this suggests that agents should learn by estimating the regression model:

$$x_{t+1} = a_1 x_t + a_2 x_{t-1} + a_3 x_{t-2} + b \mathbb{E}_{t-1}^* x_t + c + d'_1 \zeta_t + \eta_{t+1}, \quad (3.1)$$

$$\eta_{t+1} \sim \text{NIID}(0, \sigma_\eta^2),$$

where $\mathbb{E}_{t-1}^* x_t$ is lagged aggregate (not-necessarily rational) expectations, which are observable for the reasons given previously.

⁷ These results are closely related to the conditions derived by Levine et al. (2012) for solutions under imperfect information to be identical to solutions under perfect information. The results of Levine et al. (2012) are at once more general than our results (as they allow for arbitrary informational assumptions, rather than assuming that only shocks are unobserved) and less general (as they are restricted to the solutions of determinate models, and depend on assorted strong invertability assumptions).

If agents observed shocks, then by replacing η_{t+1} with $m_\varepsilon \varepsilon_{t+1} + m'_\zeta \zeta_{t+1}$, this would become an exact line fitting exercise, rather than a regression problem: after a finite number of periods agents would know the value of all parameters, thanks to the observability of $\mathbb{E}_{t-1}^* x_t$. (We also need that there is at least some variation in $\mathbb{E}_{t-1}^* x_t$ that is independent of the other terms, this will be true providing initial beliefs about a_3 and/or d_1 are non-zero.) Thus when shocks are observed, learning is trivial. This further justifies our focus on the non-observable shock case in this paper.

3.2. (Non-)Identification via OLS

Given that it is common knowledge that $m_{\varepsilon,t}$ and $m_{\zeta,t}$ are constant, the “true” model has $6 + \dim \zeta_t$ free parameters $(\alpha, \beta, \gamma, \rho, \sigma, m_\varepsilon, m_\zeta)$, and by running the regression (3.1) agents will also learn $6 + \dim \zeta_t$ parameters $(a_1, a_2, a_3, b, c, \sigma_\eta^2, d_1)$, which is a necessary condition for the identification of all of the model’s parameters. This also means that if any variables are omitted from this regression (as they are in the traditional regressions used in the literature) then agents will have no information about at least one of the model’s parameters.

Providing $\rho \neq 1$ and $\sigma > 0$, equating terms reveals that all the model’s parameters are uniquely identified if any only if either $\alpha = \rho = 0$, or the following equation for β has a unique solution:⁸

$$\beta^3 a_3 = (-\beta^2 a_2 - (\beta b - 1 + \beta a_1))(\beta b - 1 + \beta a_1).$$

Tedious algebra reveals that this in turn holds if any only if $\alpha \neq 0$, $\rho \neq 0$ and $\alpha\beta > \frac{1}{4}$, which implies there is no non-explosive, real, minimal state variable solution for x_t . This confirms Cochrane’s (2009) result that Taylor rule parameters are not identified under determinacy

⁸ The equations also have a unique solution when either $\alpha = 0$ and $\rho = \frac{1}{\beta}$, or when $\rho = 0$. However, these two cases are observationally equivalent.

via this simple form of OLS learning. Away from this case, there will either be two or three discrete solutions for the model's parameters.

However, we previously argued that sunspots were observable to agents. Hence, agents using the perceived law of motion (PLM) (3.1) are not using all available information. If they instead run the regression:

$$x_{t+1} = a_1x_t + a_2x_{t-1} + a_3x_{t-2} + b\mathbb{E}_{t-1}^*x_t + c + d_1'\zeta_t + d_0'\zeta_{t+1} + m_\varepsilon\varepsilon_{t+1}, \quad (3.2)$$

$$\varepsilon_{t+1} \sim \text{NIID}(0,1),$$

then all parameters will apparently be identified, providing $d_0 \neq 0$. For example, in the case where $\dim \zeta_t = 1$ we have: $\frac{1}{\beta} = a_1 + \frac{d_1}{d_0}$ and $\rho = b - \frac{d_1}{d_0}$. We also have the over-identifying restriction $a_3 + \left(a_1 + \frac{d_1}{d_0}\right)\left(b - \frac{d_1}{d_0}\right)^2 = -a_2\left(b - \frac{d_1}{d_0}\right)$. When $\dim \zeta_t > 1$, these equalities must hold for each non-zero component of d_0 and the corresponding component of d_1 , giving further over-identifying restrictions. Unfortunately, since the estimated value of d_0 will be non-zero with probability one (even under a MSV solution with $m_\zeta = 0$), under (3.2) although it may seem like we have identified a non-MSV solution, we must continue to place positive probability on being in a MSV solution, so the identification here is illusory. Furthermore, agents generally have no grounds for believing that $m_{\varepsilon,t}$ and $m_{\zeta,t}$ are indeed constant. This means that the standard errors on their parameter estimates should be bounded away from zero even asymptotically, further dashing any hope of identification.

3.3. Identification by learning from learners

Although agents cannot identify structural parameters via running either of the regressions given in the last section, if one sophisticated agent realises that everyone else is running these regressions in order to form expectations then that sophisticated agent will be able to identify parameters.

Since we did not use the rationality of expectations in deriving equation (2.4), it must always be the case that:

$$x_t = (\alpha + \rho)x_{t-1} - \alpha\rho x_{t-2} + \beta\mathbb{E}_t^*x_{t+1} - \beta\rho\mathbb{E}_{t-1}^*x_t + (1 - \rho)\gamma + \sigma\varepsilon_t. \quad (3.3)$$

The only thing stopping us from running a regression of this form in order to identify β is the endogeneity of $\mathbb{E}_t^*x_{t+1}$. But if agents are forming expectations using (3.1) or (3.2) then we know that $d'_{1,t-1}\zeta_t$ is a valid instrument for $\mathbb{E}_t^*x_{t+1}$ (where $d_{1,t-1}$ is the estimated values of d_1 using information up to period $t - 1$ at the latest)⁹, since ζ_t is uncorrelated with ε_t by assumption. Hence, one potential way of achieving identification would be to run a standard IV-regression. However, this is unlikely to be very efficient as it discards a lot of information.

We can do considerably better here by considering the structure of the implied actual law of motion (ALM). Note that if everyone is forming expectations by running the regression (3.1) or (3.2), then:

$$x_t = (1 - \beta a_{1,t-1})^{-1} [(\alpha + \rho + \beta a_{2,t-1})x_{t-1} + (\beta a_{3,t-1} - \alpha\rho)x_{t-2} + \beta(b_{t-1} - \rho)\mathbb{E}_{t-1}^*x_t + [(1 - \rho)\gamma + \beta c_{t-1}] + \beta d'_{1,t-1}\zeta_t + \sigma\varepsilon_t],$$

where time subscripts on the regression coefficients again refer to agents' estimates using information up to period $t - 1$ at the latest. We do not specify at this point if these estimates are the result of recursive least squares (RLS—equivalent to OLS), constant gain least squares (CGLS), or some other estimation method. In the appendix, section 7.2 we analyse e-stability, which will determine convergence of the naïve agents' beliefs under RLS; but this will not be important for the analysis of the convergence of the beliefs of our one sophisticated agent.

⁹ We are assuming that the OLS agents adopt the standard convention of forming expectations using parameter estimates from previous periods' observations. When they are allowed to use current observations then we can proxy the estimates with current observations by the estimates with lagged ones to avoid further endogeneity issues.

Using the ALM above, we can estimate the model's structural parameters by conditional maximum likelihood (ML). The conditional log-likelihood is given by:

$$\begin{aligned} & \log f(x_1, \dots, x_T | x_0, x_{-1}, \mathbb{E}_0^* x_1, \zeta_1, \dots, \zeta_T, h_0, \theta) \\ &= \sum_{t=1}^T \log f(x_t | x_{t-1}, x_{t-2}, \mathbb{E}_{t-1}^* x_t, \zeta_t, h_0, \dots, h_{t-1}, \theta) \\ &= -\frac{T}{2} \log 2\pi + \sum_{t=1}^T \left[\log |1 - \beta a_{1,t-1}| - \log \sigma - \frac{1}{2\sigma^2} (x_t - \mu_t)^2 \right] \end{aligned}$$

where $h_t = [a_{1,t} \ a_{2,t} \ a_{3,t} \ b_t \ c_t \ d'_{1,t}]'$, $\theta = [\alpha \ \beta \ \gamma \ \rho \ \sigma]'$,

$$\mu_t := (\alpha + \rho)x_{t-1} - \alpha\rho x_{t-2} + \beta\mathbb{E}_t^* x_{t+1} - \beta\rho\mathbb{E}_{t-1}^* x_t + (1 - \rho)\gamma,$$

and:

$$\mathbb{E}_t^* x_{t+1} = a_{1,t-1}x_t + a_{2,t-1}x_{t-1} + a_{3,t-1}x_{t-2} + b_{t-1}\mathbb{E}_{t-1}^* x_t + c_{t-1} + d'_{1,t-1}\zeta_t. \quad (3.4)$$

Note that in introducing the conditioning on h_0, \dots, h_{t-1} in the first equality we have used the fact that h_0, \dots, h_{t-1} are deterministic functions of x_{-1}, \dots, x_{t-1} .

The first order conditions then imply that¹⁰:

$$\begin{aligned} 0 &= \sum_{t=1}^T (x_{t-1} - \hat{\rho}x_{t-2})(x_t - \hat{\mu}_t) \\ 0 &= \sum_{t=1}^T \left[\mathbb{E}_t^* x_{t+1} (x_t - \hat{\mu}_t) - \frac{a_{1,t-1}\hat{\sigma}^2}{1 - \hat{\beta}a_{1,t-1}} \right] \\ 0 &= \sum_{t=1}^T (x_{t-1} - \hat{\alpha}x_{t-2} - \hat{\beta}\mathbb{E}_{t-1}^* x_t - \gamma)(x_t - \hat{\mu}_t) \\ 0 &= \sum_{t=1}^T (x_t - \hat{\mu}_t), \quad \hat{\sigma}^2 = \frac{1}{T} \sum_{t=1}^T (x_t - \hat{\mu}_t)^2 \end{aligned} \quad (3.5)$$

Since the second equation is a polynomial of at least order T in β , in general these equations will have to be solved numerically. However, providing parameters are indeed

¹⁰ As usual, hats denote estimates.

identified, the resulting estimates will have all the usual desirable properties of ML estimates (consistency, efficiency, asymptotic normality).

To show that the ML estimator does indeed identify parameters, we give an alternative estimator that we are able to prove to be consistent. Since the existence of a consistent estimator implies identification (Gabrielsen 1978), this is sufficient for the consistency and asymptotic normality of the ML estimator. This alternative estimator will also have a recursive form, making it convenient for the case in which everyone realises everyone else is learning.

Let $\theta := [\theta_1 \ \theta_2 \ \theta_3 \ \theta_4 \ \theta_5]' = [(1 - \rho)\gamma \ \alpha + \rho \ -\alpha\rho \ \beta \ -\beta\rho]'$ be a vector of parameters to be estimated, and let:

$$z_t := \begin{bmatrix} 1 & x_{t-1} & x_{t-2} & \begin{pmatrix} a_{2,t-1}x_{t-1} + a_{3,t-1}x_{t-2} + \\ b_{t-1}\mathbb{E}_{t-1}^*x_t + c_{t-1} + d'_{1,t-1}\zeta_t \end{pmatrix} & \mathbb{E}_{t-1}^*x_t \end{bmatrix}'. \quad (3.6)$$

Suppose for the moment that an oracle told us the value of β . Then by running the regression:

$$(1 - \beta a_{1,t-1})x_t = z_t'\theta + \sigma\varepsilon_t, \quad \varepsilon_t \sim \text{NIID}(0,1), \quad (3.7)$$

we could identify all parameters, even if we forgot what the oracle had told us as soon as the regression had been run. In particular $\hat{\sigma}$ is the standard deviation of the shock, $\hat{\beta} = \hat{\theta}_4$,

$$\hat{\rho} = -\frac{\hat{\theta}_5}{\hat{\beta}} = -\frac{\hat{\theta}_5}{\hat{\theta}_4}, \hat{\gamma} = \frac{\hat{\theta}_1}{1-\hat{\rho}} = \frac{\hat{\theta}_1\hat{\theta}_4}{\hat{\theta}_4+\hat{\theta}_5}, \text{ and } \hat{\alpha} \text{ is given by either } \hat{\theta}_2 - \hat{\rho} = \frac{\hat{\theta}_2\hat{\theta}_4+\hat{\theta}_5}{\hat{\theta}_4} \text{ or } -\frac{\hat{\theta}_3}{\hat{\rho}} = \frac{\hat{\theta}_3\hat{\theta}_4}{\hat{\theta}_5}.$$

(The two estimates of α may be near-optimally combined to give $\hat{\alpha} =$

$$\frac{\hat{\theta}_3\hat{\theta}_4\hat{\theta}_5 s_{\theta,22} + \hat{\theta}_4(\hat{\theta}_2\hat{\theta}_4 + \hat{\theta}_5) s_{\theta,33} - (\hat{\theta}_2\hat{\theta}_4\hat{\theta}_5 + \hat{\theta}_3\hat{\theta}_4^2 + \hat{\theta}_5^2) s_{\theta,23}}{\hat{\theta}_5^2 s_{\theta,22} + \hat{\theta}_4^2 s_{\theta,33} - 2\hat{\theta}_4\hat{\theta}_5 s_{\theta,23}}, \text{ where } \begin{bmatrix} s_{\theta,22} & s_{\theta,23} \\ s_{\theta,32} & s_{\theta,33} \end{bmatrix} \text{ is the estimated}$$

covariance matrix of $\begin{bmatrix} \hat{\theta}_2 \\ \hat{\theta}_3 \end{bmatrix}$.)

Now let $Z_T := \begin{bmatrix} z_1' \\ \vdots \\ z_T' \end{bmatrix}$, $x := \begin{bmatrix} x_1 \\ \vdots \\ x_T \end{bmatrix}$ and $y := \begin{bmatrix} a_0x_1 \\ \vdots \\ a_{T-1}x_T \end{bmatrix}$. Then the (OLS) estimated value of θ is

given by:

$$\hat{\theta} = (Z_T'Z_T)^{-1}Z_T'(x - y\beta).$$

To show consistency of this estimator, let us begin by defining a vector of “pseudo-instruments” (variables that we would like to use in place of z_t , were they observable):

$$d_t := \left[1 \quad \frac{\sigma \varepsilon_{t-1}}{1 - \beta a_{1,t-2}} \quad \frac{\sigma \varepsilon_{t-2}}{1 - \beta a_{1,t-3}} \quad d'_{1,t-1} \zeta_t \quad \frac{d'_{1,t-2} \zeta_{t-1}}{1 - \beta a_{1,t-2}} \right]'$$

Denote by $\mathbb{E}^+ V$ the unconditional expectation of V that would have obtained were $a_{1,t}$, $a_{2,t}$, $a_{3,t}$, b_t , c_t and $d_{1,t}$ non-stochastic for all t . Then if $J_t := \mathbb{E}^+ d_t d'_t$,

$$J_t = \text{diag} \left[1 \quad \frac{\sigma^2}{(1 - \beta a_{1,t-2})^2} \quad \frac{\sigma^2}{(1 - \beta a_{1,t-3})^2} \quad d'_{1,t-1} d_{1,t-1} \quad \frac{d'_{1,t-2} d_{1,t-2}}{(1 - \beta a_{1,t-2})^2} \right],^{11}$$

and if $K_t := (\mathbb{E}^+ d_t d'_t)^{-1} \mathbb{E}^+ d_t z'_t$,

$$K_t = \begin{bmatrix} 1 & ? & ? & ? & ? \\ 0 & 1 & 0 & a_{2,t-1} + b_{t-1} a_{1,t-2} & a_{1,t-2} \\ 0 & q_{t-2} & 1 & (a_{2,t-1} + b_{t-1} a_{1,t-2}) q_{t-2} + a_{3,t-1} + b_{t-1} a_{2,t-2} & a_{2,t-2} + a_{1,t-2} q_{t-2} \\ 0 & 0 & 0 & 1 & 0 \\ 0 & \beta & 0 & \beta a_{2,t-1} + b_{t-1} & 1 \end{bmatrix},$$

where $q_{t-2} = \frac{\alpha + \rho + \beta a_{2,t-2} + \beta (b_{t-2} - \rho) a_{1,t-3}}{1 - \beta a_{1,t-2}}$, and ? denotes a term omitted for the sake of

space. We also define $\tilde{J}_T := \sum_{t=1}^T J_t$, and $\tilde{K}_T := \tilde{J}_T^{-1} \sum_{t=1}^T J_t K_t$, so if $D := \begin{bmatrix} d'_1 \\ \vdots \\ d'_T \end{bmatrix}$, $\tilde{J}_T = \mathbb{E}^+ D' D$

and $\tilde{K}_T = (\mathbb{E}^+ D' D)^{-1} \mathbb{E}^+ D' Z_T$. These definitions are valid as \tilde{J}_T is diagonal, with a strictly positive diagonal, for all t . (Though the elements of the diagonal may tend to 0 asymptotically.) A sufficient condition for the invertability of both K_T and \tilde{K}_T , for all T , is that $\beta \neq 1$, in which case the eigenvalues of K_t and \tilde{K}_T must be bounded away from 0 asymptotically.

¹¹ The diag operator maps vectors to diagonal matrices with a diagonal with the same elements as the vector, and maps matrices to a vector with the same elements as their diagonal.

If we go on to define:

$$U_T := Z_T - D(D'D)^{-1}D'Z_T,$$

then $D'U = 0$ and:

$$Z_T'Z_T = Z_T'D(D'D)^{-1}D'D(D'D)^{-1}D'Z_T + U_T'U_T.$$

If it were valid to drop the \mathbb{E}^+ operators from our expressions for \check{J}_T and \check{K}_T , asymptotically, then we would have:

$$\Pr\left(\lim_{T \rightarrow \infty} (\check{K}_T' \check{J}_T \check{K}_T + U_T'U_T - Z_T'Z_T) = 0\right) = 1. \quad (3.8)$$

Dropping the \mathbb{E}^+ operators in this way might be valid, for example, if agents were learning a sunspot solution via RLS, and eventually the dependence between their estimates was sufficiently weak that $a_{1,t}$, $a_{2,t}$, etc. were “near exogenous”, in some loose sense. However, rather than making such specific assumptions, we will instead just assume the validity of (3.8), since (3.8) encompasses many other cases, including ones in which $\text{plim}_{T \rightarrow \infty} \check{K}_T$ does not even exist, as it will not under constant gain learning.

Given (3.8), by applying Theorem 1 of Lai and Wei (1982) to the regression (3.7), providing:

- 1) there exists $\delta > 0$ such that $\limsup_{t \rightarrow \infty} \frac{\max\{1, a_{1,t}^2\}}{t^{1-\delta} \min\{1, d_{1,t}'d_{1,t}\}} < \infty$,¹² and
- 2) there exists $\delta' \geq 0$ such that $\limsup_{t \rightarrow \infty} \frac{z_t'z_t}{t^{\delta'}} < \infty$,¹³

then $\hat{\theta} \xrightarrow{a.s.} \theta$. Note that 2) already covers all sub-exponential explosion in $z_t'z_t$. We do not as yet have a proof of consistency for the case with an exponential (or super-exponential) explosion, but our simulation results below certainly suggest that β can still be consistently estimated in this case (though obviously γ cannot be).

¹² Sufficient as $\sum_{t=1}^{\infty} t^{-(1-\delta)} = \infty$ for all $\delta \geq 0$.

¹³ Sufficient as $\lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T t^{\delta'}}{T^{1+\delta'}} < \infty$, $\lim_{T \rightarrow \infty} \frac{\log T}{\sum_{t=1}^T t^{-(1-\delta)}} = 0$ for all $\delta > 0$, and since $\text{tr } z_t z_t'$ is guaranteed to be between the largest eigenvalue of $z_t z_t'$ and 5 times this quantity.

Furthermore, under slightly stronger assumptions $\tilde{J}_T^{-1/2} \tilde{K}_T \hat{\theta}$ will be asymptotically normally distributed, implying that we have $\frac{1}{\sqrt{\log T}}$ convergence in the worst case.

It is easy to see that these sufficient conditions will hold under any non-exponentially-explosive learning algorithm, with slower than $\frac{1}{\sqrt{T}}$ convergence, such as constant gain least squares, or stochastic gradient learning. Under recursive least squares, there exists $\delta \geq 0$ such that $t^{\frac{1-\delta}{2}} d_{1,t}$ converges in distribution to a normal, (Marcet and Sargent 1992), with $\delta = 0$ only if the real parts of the eigenvalues of the “ T ” matrix are all less than $1/2$.¹⁴ When $\delta > 0$ here, our sufficient conditions will be satisfied, but in the other case, Theorem 1 of Lai and Wei (1982) no longer applies. From their reasoning, we do however have that $\limsup_{t \rightarrow \infty} (\hat{\theta} - \theta)' (\hat{\theta} - \theta) < \infty$, even here, so at worst, beyond a certain point in time standard errors on θ would cease improving. Additionally, we note that a sufficient condition for consistency in this case is that:

$$\limsup_{T \rightarrow \infty} \left\| \tilde{J}_T^{-1/2} \tilde{K}_T^{-1} U_T' U_T \tilde{K}_T^{-1} \tilde{J}_T^{-1/2} \right\| < \infty, \quad (3.9)$$

by Theorem 3 of Lai and Wei (1982). This will hold, for example, if $\alpha = \rho = 0$, so it may be thought of as an additional weak-dependency condition.

We have demonstrated then a range of conditions under which $\hat{\theta}$ is a consistent estimator of θ , in our oracle-aided regression, equation (3.7). Now suppose there is no oracle, but we have received infinitely many periods of data. If we guessed a value for β , we could repeat the “oracle” exercise with the guessed value and we would end up with an alternative estimate for β (namely $\hat{\theta}_3$). We can thus think of this as a fixed-point problem. In general our guess of β and the estimated value will not coincide, but we know that they must coincide at least once, namely when our guess is the true value. Thus if the (infinite-data)

¹⁴ These eigenvalues are given in the appendix, 7.2.

fixed-point problem has a unique solution for β , then we know that value must be the true value. Hence, if in finite samples this fixed-point problem also has a unique solution, that solution must be a consistent estimator of β , at least when the conditions discussed above hold.

We proceed to establish the uniqueness of the solution to the fixed-point problem, by establishing a closed form solution. Let $e_4 := [0 \ 0 \ 0 \ 1 \ 0]'$. Then the fixed-point problem may be expressed as finding the value of $\hat{\beta}$ for which:

$$\hat{\beta} = e_4'(Z'Z)^{-1}Z'(x - y\hat{\beta}).$$

Consequently:

$$\hat{\beta} = \frac{e_4'(Z'Z)^{-1}Z'x}{1 + e_4'(Z'Z)^{-1}Z'y}.$$

Armed with a consistent estimator of $\hat{\beta}$, all other parameters may be estimated consistently by following our oracle procedure. In particular, the consistent estimator of θ is:

$$\begin{aligned} \hat{\theta}^{2SLS} &= (Z'Z)^{-1}Z' \left[x - y \frac{e_4'(Z'Z)^{-1}Z'x}{1 + e_4'(Z'Z)^{-1}Z'y} \right] \\ &= (I + (Z'Z)^{-1}Z'ye_4')^{-1}(Z'Z)^{-1}Z'x \\ &= (Z'Z + Z'ye_4')^{-1}Z'x, \end{aligned} \tag{3.10}$$

which turns out to be equal to the 2SLS-IV estimator when $(a_{2,t-1}x_{t-1} + a_{3,t-1}x_{t-2} + b_{t-1}\mathbb{E}_{t-1}^*x_t + c_{t-1} + d'_{1,t-1}\zeta_t)$ is used as an instrument for $\mathbb{E}_t^*x_{t+1}$.

This gives us the following proposition:

Proposition 1: *Suppose the economy is made up of agents that are all forming expectations through running regressions of the form of (3.1) or (3.2), with $\dim \zeta_t > 0$. Let $\hat{\theta}^{2SLS}$ be the estimator defined by equation (3.10), and suppose that:*

- 1) *the weak-dependence condition (3.8) holds,*

2) there exists $\delta > 0$ such that $\limsup_{t \rightarrow \infty} \frac{\max\{1, a_{1,t}^2\}}{t^{1-\delta} \min\{1, d'_{1,t} d_{1,t}\}} < \infty$, and

3) there exists $\delta' \geq 0$ such that $\limsup_{t \rightarrow \infty} \frac{z'_t z_t}{t^{\delta'}} < \infty$,

Then if one of the following conditions holds:

a) the agents learn by any algorithm with slower than $1/\sqrt{t}$ convergence, such as constant gain least squares, stochastic gradient learning, or recursive least squares in the case in which the eigenvalues of the “T” matrix (defined in appendix 7.2) are greater than $1/2$,

b) the agents learn a sunspot solution,

c) the agents learn by recursive least squares, or another algorithm under which $\sqrt{t}d'_{1,t}$ converges in distribution, and the second weak-dependence condition (3.9) holds,

then the 2SLS-like estimator $\hat{\theta}^{2SLS}$ is consistent.

Since the existence of a consistent estimator implies parameter identification under maximum likelihood, we have the following immediate corollary:

Corollary 1.1: Under the conditions of **Proposition 1**, the maximum likelihood estimator given by the solution to the FOCs, (3.5) is consistent.

Note that the consistency of these estimators is in spite of the convergence of $a_{1,t}$, $a_{2,t}$, etc. rather than because of this convergence. Indeed, the worse the learning process that is determining $a_{1,t}$, $a_{2,t}$, etc., the faster this more sophisticated agent will learn the structural parameters of the model. So for example, if almost all agents are using stochastic gradient learning or constant gain least squares, then learning structural parameters is likely to be particularly easy. Likewise if $a_{1,t}$, $a_{2,t}$, etc. never converge then learning the structural parameters is again likely to be fast. This result is related to Cochrane’s (2009) claim that

with unsophisticated learning it is only in the explosive case that structural parameters may be identified, but here we have identification quite generally.

3.4. Learning from MSV learners

It is natural to wonder the extent to which our results are driven by the fact that the agents in the economy are learning and forming expectations using equation (3.1) or (3.2), rather than the more traditional MSV form:

$$x_{t+1} = a_1 x_t + a_2 x_{t-1} + c + m_\varepsilon \varepsilon_{t+1}, \quad \varepsilon_{t+1} \sim \text{NIID}(0,1). \quad (3.11)$$

Since many REE do not have a representation in this form, by estimating (3.11) the agents in the economy are already putting a prior probability of zero on any non-fundamental solution, which is certainly not justified in the absence of transversality constraints limiting x_t to asymptotic stationarity. Nonetheless, even given these priors, when agents observe a stationary realisation of x_t they will still not be able to work out the value of β , as there are observationally equivalent MSV solutions. So, it remains an interesting question whether or not β can be identified from examining these learners.

The argument of the previous section would suggest using $a_{2,t-1}x_{t-1} + c_{t-1}$ as an instrument for $\mathbb{E}_t^* x_{t+1}$. Proving the general validity of this instrument in the MSV set-up is tricky, however. This is clearest when $\alpha = \rho = 0$, in which case, asymptotically $x_{t+1} = m_\varepsilon \varepsilon_{t+1}$, if parameters converge. With no serial correlation in x_t , finding “pseudo-instruments” (i.e. potential elements of d_t) that are correlated with $\mathbb{E}_t^* x_{t+1}$ and $\mathbb{E}_{t-1}^* x_t$, but not with ε_{t-1} or ε_{t-2} is non-trivial.

Suppose that $\frac{1}{\kappa_a(t)} \begin{bmatrix} a_{1,t} - a_{1,\infty} \\ a_{2,t} - a_{2,\infty} \end{bmatrix}$ tends in distribution to some non-degenerate distribution, as $t \rightarrow \infty$, for some function $\kappa_a(t)$, and some constants $a_{1,\infty}$ and $a_{2,\infty}$. Then under any “reasonable” estimator (including the RLS, CGLS etc. estimators):

$$\begin{aligned} \liminf_{t \rightarrow \infty} \kappa_a(t)^2 \text{cov}(a_{1,t}, \varepsilon_t \varepsilon_{t-1}) &> 0, \\ \liminf_{t \rightarrow \infty} \kappa_a(t)^2 \text{cov}(a_{2,t}, \varepsilon_t \varepsilon_{t-2}) &> 0, \text{ \&} \\ \limsup_{t \rightarrow \infty} \kappa_a(t)^2 \text{cov}(a_{2,t}, \varepsilon_t \varepsilon_{t-1}) &= 0. \end{aligned} \tag{3.12}$$

Thus if we define:

$$d_t := \left[1 \quad \frac{\sigma \varepsilon_{t-1}}{1 - \beta a_{1,t-2}} \quad \frac{\sigma \varepsilon_{t-2}}{1 - \beta a_{1,t-3}} \quad \frac{\sigma \kappa_a(t)^2 \varepsilon_{t-1}^2 \varepsilon_{t-3}}{1 - \beta a_{1,t-2}} \quad \frac{\sigma \kappa_a(t)^2 \varepsilon_{t-1} \varepsilon_{t-2} \varepsilon_{t-4}}{1 - \beta a_{1,t-2}} \right]', \tag{3.13}$$

then providing $\liminf_{t \rightarrow \infty} t^{1-\delta} \kappa_a(t)^2 > 0$ for some $\delta > 0$, the previous proof goes through.¹⁵

Of course, under recursive least squares learning $\kappa_a(t) = \frac{1}{\sqrt{t}}$ when the eigenvalues of the “ T ” matrix are less than $1/2$, so this sufficient condition does not hold.

While the second weak-dependence condition (3.9) could be generalised to this case, it seems highly implausible that it would hold here, due to the convoluted nature of our “pseudo-instruments”.¹⁶ The convoluted nature of these pseudo-instruments also suggests that our actual-instrument vector, z_t may be a rather poor instrument. One other possibility that could be used as an additional instrument is $a_{1,t-1}$, since it is correlated with the first term of $\mathbb{E}_t^* x_{t+1}$. Indeed, it is easy to see that whether agents are learning from (3.11), or one of our more general laws, (3.1) or (3.2), the asymptotically optimal choice of instruments is:

$$z_t^* := \begin{bmatrix} z_t \\ a_{1,t-1} z_t \end{bmatrix}$$

¹⁵ We also need to adjust the definition of \mathbb{E}_t^+ so that only the $a_{1,t}$ in the denominator of the ALM of x_t is treated as non-stochastic.

¹⁶ Since completing this paper, we discovered the results of Christopheit and Massmann (2010) who were able to prove consistency in an RLS learning of the MSV solution context, for a simple model, using a more direct technique. In future work we intend to investigate whether their proof techniques may be generalised to cover regressions such as these.

since $\mathbb{E}_t x_{t+1} = Fz_t^* + a_{1,t-1}\varepsilon_t$ for some non-stochastic, constant matrix F , and this is not true for any proper subset of these instruments. We then have the following generalisation of **Proposition 1** and **Corollary 1.1** for this choice of instruments:

Proposition 2: *Suppose the economy is made up of agents that are all forming expectations through running regressions of the form of (3.1), (3.2) or (3.11). Let $z_t^* = [z_t' \quad a_{1,t-1}z_t']'$, where z_t is defined by equation (3.6), and let $Y := Z + ye_4'$, $Z^* := [z_1^* \quad \dots \quad z_T^*]'$, and:*

$$\hat{\theta}_T^{\text{AEIV}} := \left(Y'Z^*(Z^{*\prime}Z^*)^{-1}Z^{*\prime}Y \right)^{-1} Z^*(Z^{*\prime}Z^*)^{-1}Z^{*\prime}x.$$

Then if either:

- i) (3.1) or (3.2) is being used, and conditions 1), 2) and 3) of **Proposition 1** hold, or:
- ii) (3.11) is being used and:

1) the weak-dependence condition (3.8) holds (with d_t defined by (3.13)), and,

2) there exists $\delta' \geq 0$ such that $\limsup_{t \rightarrow \infty} \frac{a_{1,t}^2}{t^{\delta'}} < \infty$ and $\limsup_{t \rightarrow \infty} \frac{z_t'z_t}{t^{\delta'}} < \infty$,

and one of the following further conditions holds also:

- a) the agents learn by any reasonable¹⁷ algorithm which converges in distribution, but slower than $1/\sqrt{t}$, such as stochastic gradient learning, or recursive least squares in the case in which the eigenvalues of the “T” matrix (defined in appendix 7.2) are greater than $1/2$,
- b) the agents learn a sunspot solution,

¹⁷ Where a reasonable algorithm is defined as one for which (3.12) is satisfied.

c) *the agents learn by recursive least squares on regression (3.1) or (3.2), or another algorithm under which $\sqrt{t}d'_{1,t}$ converges in distribution, $\dim \zeta_t > 0$ and the second weak-dependence condition (3.9) holds,*

then the estimator $\hat{\theta}_T^{\text{AEIV}}$ is consistent and asymptotically efficient.

Corollary 2.1: *Under the conditions of **Proposition 2**, the maximum likelihood estimator given by the solution to the FOCs, (3.5) is consistent.*

3.5. Simulation evidence

In light of the slightly obscure nature of some of our theoretical conditions, particularly in the recursive least squares (RLS) case, we now present some simulation evidence of the estimator's success in identifying the key β parameter. Figure 1 gives results for economies populated with RLS learners estimating equation (3.2), and Figure 2 gives results for economies populated with RLS learners estimating the MSV form, equation (3.11).

In order to show the estimator's performance, for each parameterisation (different rows of the two figures) we generate 2^{14} simulation paths (each of length 2^8), and then apply each estimator considered to each of the resulting paths.

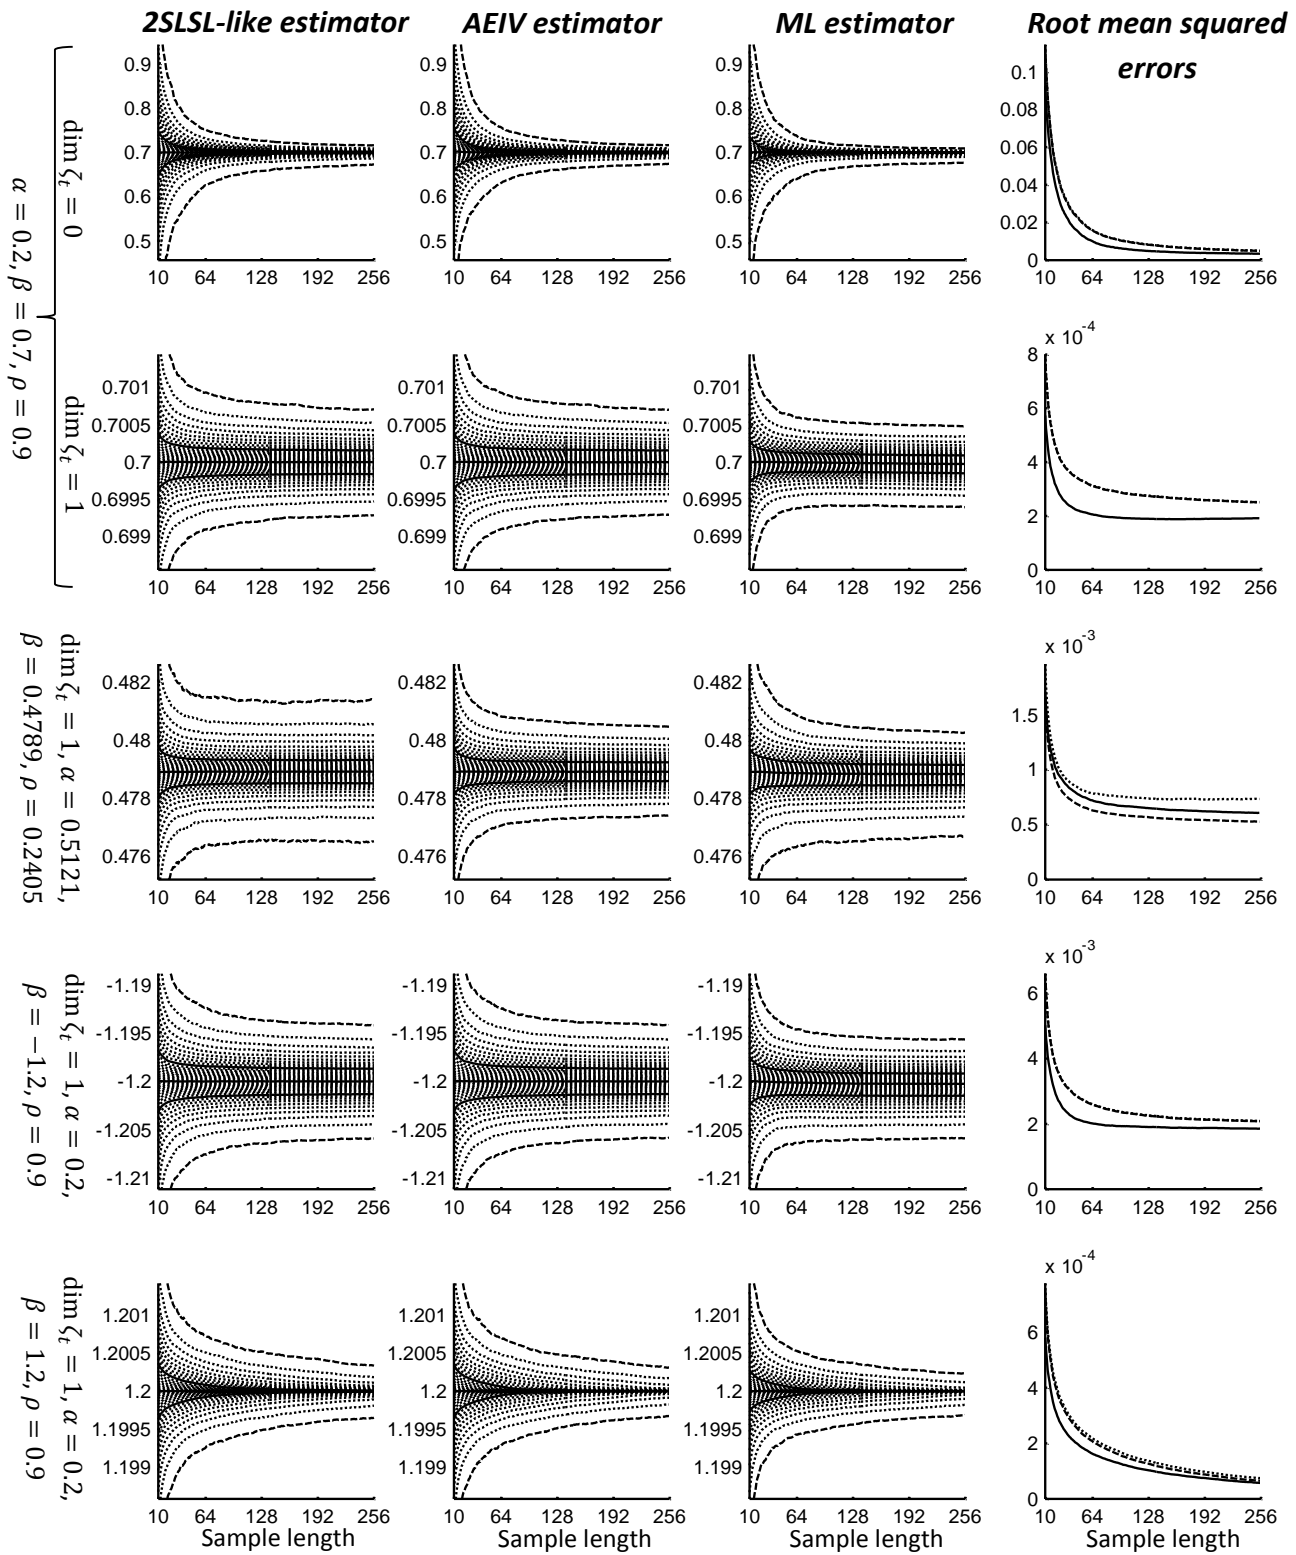


Figure 1: Distribution properties of the estimates of β , from 2^{14} runs, when agents estimate equation (3.2) using OLS. See text (section 3.5) for full details.

In both figures, each of the first three columns corresponds to a different estimator. For both figures, column 1 is our original 2SLS estimator, column 2 is the asymptotically efficient

IV one (henceforth, AEIV) and column 3 is the ML estimator¹⁸. In each graph of the first three columns, we plot the 2.5%, 5.0%, 7.5%, ..., 97.5% percentiles of the estimator's distribution. For convenience, the quartiles are given in solid rather than dotted lines. The final column of both figures gives the 95% trimmed root mean squared error (RMSE) of the estimators.¹⁹ In this column, the dotted line corresponds to the 2SLS estimator, the dashed to the AEIV one, and the solid to the ML one.

In each simulation run, there was a "burn-in" time of 32 periods during which time expectations were set to their value under the SMSV solution (defined in section 2.2), plus $\sum_{i=1}^{\dim \zeta_t} \zeta_{t,i} + \zeta_t^B$, where ζ_t^B is an additional, unobservable, NIID(0,1) shock. This was done purely in order to help the OLS learners converge, and our estimators were only run on simulated data from the end of the burn-in period. Additionally, the OLS learners' estimates were constrained to have each parameter in $[-1000,1000]$, to prevent numerically unstable hyper-explosions with super-exponential growth. This is in the spirit of the "projection facility" invoked by Marcet and Sargent (1989).

The first two rows of graphs in Figure 1, and the first row in Figure 2, are all generated with $\alpha = 0.2$, $\beta = 0.7$, $\rho = 0.9$, $\sigma = 0.001$ and $\mathbb{E}x_t = 0.005$. These parameters mean there is a unique stationary MSV solution, which is also the only e-stable MSV solution. The graphs in the first row of Figure 1 are with $\dim \zeta_t = 0$, while those in the second have $\dim \zeta_t = 1$. Obviously, in Figure 2 we always set $\dim \zeta_t = 0$. As was expected, the ML estimator dominates the other two, which are practically indistinguishable here. The initial rate of convergence is very quick for all three estimators, but beyond a certain point, convergence

¹⁸ Obtaining a global solution to the numerical maximum likelihood was too slow to permit us to perform as many replications as necessary. Instead then, we start the local maximisation algorithm at the AEIV solution, denoted $\hat{\beta}_t^{\text{AEIV}}$, and constrain the ML estimate of β to be between $\max[\{0\} \cup \{1/a_{1,t-1} \mid \hat{\beta}_t^{\text{AEIV}} > 1/a_{1,t-1}, 1 \leq t \leq T\}]$ and $\min[\{0\} \cup \{1/a_{1,t-1} \mid \hat{\beta}_t^{\text{AEIV}} < 1/a_{1,t-1}, 1 \leq t \leq T\}]$.

¹⁹ i.e. the RMSE after first discarding any observations below the 2.5% percentile or above the 97.5% percentile. These outliers are trimmed to limit the damage caused by the numerical errors that are introduced by the occasional explosive, or near-explosive, path.

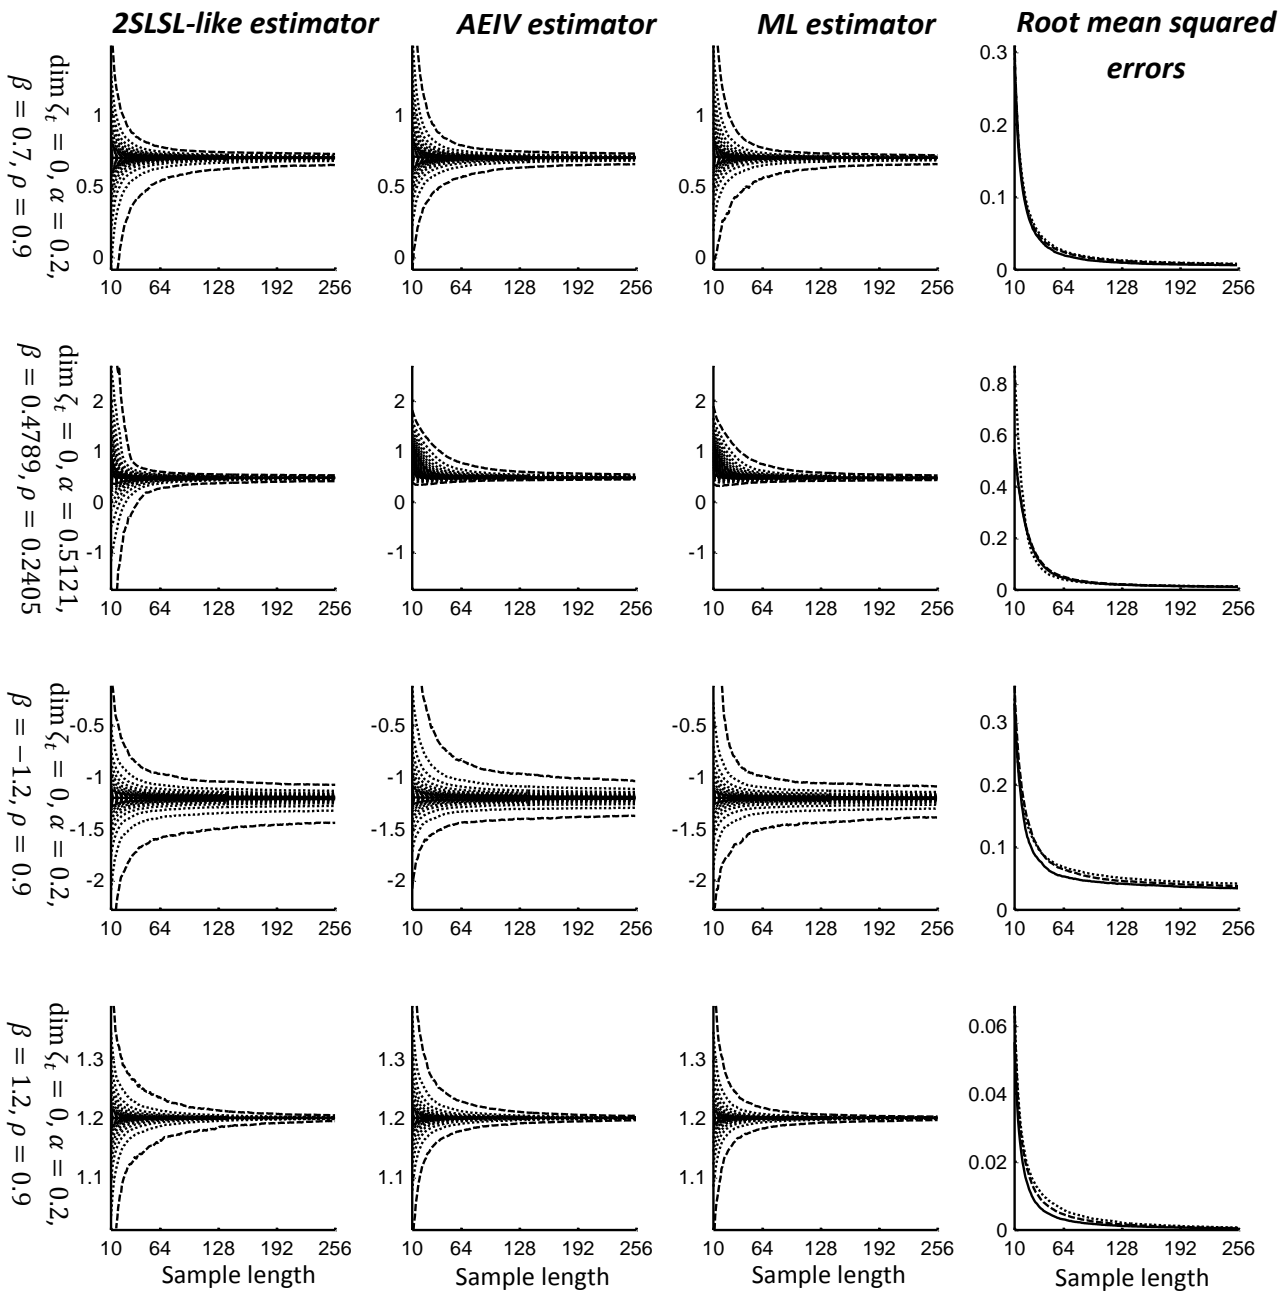


Figure 2: Distribution properties of the estimates of β , from 2^{14} runs, when agents estimate equation (3.11) using OLS. See text (section 3.5) for full details.

certainly seems to slow, in line with our $\frac{1}{\sqrt{\log T}}$ convergence finding. However, although the rate of improvement is slow, the level of the RMSE is low enough that this is unlikely to be a problem in practice.

In the next row of both figures, we repeat the exercise with $\alpha = 0.5121, \beta = 0.4789$ and $\rho = 0.2405$. These values were selected as they result in dynamics under full-information that are observationally equivalent to our original ones. Convergence here is slower since

two of the eigenvalues of the “ T ” map are now greater than $1/2$. There is also clearly large upwards bias in finite samples when agents are estimating (3.11). Surprisingly, it appears the AEIV estimator dominates the ML one in this case, whichever equation is being estimated. Nonetheless, asymptotically our estimators appear to have very similar properties.

In the penultimate row of the figures we show the results when $\alpha = 0.2$, $\beta = -1.2$ and $\rho = 0.9$. This is in the indeterminate region of the parameter space, but still in a region in which the MSV solution is e-stable. Performance appears similar to performance in the $\beta = 0.7$ case.

Finally, in the last row of both figures we show the behaviour of our estimators in an indeterminate region of the parameter space in which the SMSV is not e-stable. (In particular we set $\alpha = 0.2$, $\beta = 1.2$ and $\rho = 0.9$.) The underlying instability of the system makes identification easier for our sophisticated agent, giving us better performance than in any other case, whichever equation is being estimated.

The graphs make clear that even in small samples, when agents are estimating (3.2) all three estimators are approximately unbiased, whatever the true parameters, and whatever the value of $\dim \zeta_t$. Moreover, the estimators are highly peaked around the true value, meaning that the RMSE significantly overstates the median absolute error. Hence, people using these estimators can expect their estimated values to be closer to the truth than is suggested by the standard errors.

4. Learning from sophisticated learners

Having established that our ML and 2SLS-like estimators can successfully identify the structural parameters of the model, we now use these techniques to describe our family of misspecification free learning algorithms. Under these algorithms, each agent in the

economy will realise that everyone else is learning at the same time as them, and indeed, they will take advantage of this fact to identify the model's structural parameters. By learning these structural parameters, rather than a reduced form equation, agents will be able to disentangle learning which particular solution to the model is being used from the time variation in reduced form parameters caused by simultaneous learning.

4.1. General results

Suppose for the moment that $m_{\varepsilon,t}$ and $m_{\zeta,t}$ are public knowledge and hence do not have to be estimated, even when no one knows any of the other structural parameters.

Suppose further that everyone is learning using the ML or 2SLS-like estimator from section 3.3. Providing agents continue to use an expression of the form of (3.4) to form expectations, where now $a_{1,t}$ etc. will be functions of estimated structural parameters, this will be valid. In particular, we might suppose that agents treat their estimate of structural parameters as the true values and set:

$$\begin{aligned} a_{1,t} &= \frac{1}{\hat{\beta}_t} \left(1 - \frac{\hat{\sigma}_t}{m_{\varepsilon,t}} \right), & a_{2,t} &= -\frac{1}{\hat{\beta}_t} (\hat{\alpha}_t + \hat{\rho}_t), & a_{3,t} &= \frac{1}{\hat{\beta}_t} \hat{\alpha}_t \hat{\rho}_t, \\ b_t &= \hat{\rho}_t + \frac{1}{\hat{\beta}_t} \frac{\hat{\sigma}_t}{m_{\varepsilon,t}}, & c_t &= -\frac{1}{\hat{\beta}_t} (1 - \hat{\rho}_t) \hat{\gamma}_t, & d_{1,t} &= \frac{1}{\hat{\beta}_t} \frac{\hat{\sigma}_t}{m_{\varepsilon,t}} m_{\zeta,t}. \end{aligned} \quad (4.1)$$

(4.1) is reasonable since the actual law of motion implied by equations (2.4) and (3.4) is:

$$\begin{aligned} x_{t+1} &= (1 - \beta a_{1,t})^{-1} [(\alpha + \rho + \beta a_{2,t} + \beta (b_t - \rho) a_{1,t-1}) x_t \\ &\quad + (\beta a_{3,t} - \alpha \rho + \beta (b_t - \rho) a_{2,t-1}) x_{t-1} + \beta (b_t - \rho) a_{3,t-1} x_{t-2} \\ &\quad + \beta (b_t - \rho) b_{t-1} \mathbb{E}_{t-1}^* x_t + [(1 - \rho) \gamma + \beta c_t + \beta (b_t - \rho) c_{t-1}] + \beta d'_{1,t} \zeta_{t+1} \\ &\quad + \beta (b_t - \rho) d'_{1,t-1} \zeta_t + \sigma \varepsilon_{t+1}], \end{aligned}$$

and so when agents use (4.1), if the agents estimates of structural parameters converge in probability to their true values, then $\mathbb{E}_t x_{t+1} - \mathbb{E}_t^* x_{t+1}$ converges in probability to zero.

If agents believe in the SMSV for some reason, then we might suppose they set:

$$\begin{aligned} \hat{f}_t &= \sqrt{\max\{0, 1 - 4\hat{\alpha}_t\hat{\beta}_t\}}, & a_{1,t} &= \hat{\rho}_t + \frac{1 - \hat{f}_t}{2\hat{\beta}_t}, & a_{2,t} &= -\hat{\rho}_t \frac{1 - \hat{f}_t}{2\hat{\beta}_t}, \\ a_{3,t} &= 0, & b_t &= 0, & c_t &= \frac{2(1 - \hat{\rho}_t)\hat{\gamma}_t}{1 - 2\hat{\beta}_t + \hat{f}_t}, & d_{1,t} &= 0. \end{aligned} \quad (4.2)$$

If they do this, again as estimates of structural parameters converge in probability to their true values, $\mathbb{E}_t x_{t+1} - \mathbb{E}_t^* x_{t+1}$ will converge in probability to zero.

Furthermore, from **Proposition 2** we immediately have the following two corollaries:

Corollary 2.2: *Suppose that $m_{\varepsilon,t}$ and $m_{\zeta,t}$ are in all agent's period t information set, and $m_{\varepsilon,t} \neq 0$ for all t . Then if:*

- 1) *all agents form expectations using (3.4) and (4.1),*
- 2) *conditions 1), 2) and 3) of **Proposition 1** hold,*
- 3) *there exists $\delta > 0$ such that $\liminf_{t \rightarrow \infty} t^{1-\delta} m_{\zeta,t} > 0$, and,*
- 4) *agents estimate structural parameters using either the AEIV estimator defined in **Proposition 2**, or the ML estimator given by the solution to the FOCs, (3.5),*

*then all estimates of **structural** parameters will converge in probability to the true values, and agents' expectations will converge in probability to their values under the full information, rational expectations solution.*

Corollary 2.3: *If:*

- 1) *all agents form expectations using (3.4) and (4.2),*
- 2) *conditions 1) and 2) of **Proposition 2** hold, and,*
- 3) *agents estimate structural parameters using either the AEIV estimator defined in **Proposition 2**, or the ML estimator given by the solution to the FOCs, (3.5),*

then all estimates of **reduced form** parameters will converge in probability to the true values, and agents' expectations will converge in probability to their values under the full information, rational expectations, SMSV solution.

Note that **Corollary 2.3** only guarantees convergence of reduced form parameters, not structural ones. This is because if reduced form parameters converge too quickly, **Proposition 2** does not apply. Since there are more structural parameters than reduced form ones in the MSV case, it is quite possible for the reduced form parameters to converge without the structural ones converging. Guaranteeing convergence of reduced form parameters is sufficient for expectations to converge to the SMSV solution, however.

To guarantee the existence of a learning algorithm that will learn an arbitrary solution, we need the following supplemental corollary of **Corollary 2.2**:

Corollary 2.4: Suppose that agents do not know $m_{\varepsilon,t}$ and $m_{\zeta,t}$, and each agent i forms the estimate $\hat{m}_{\varepsilon,t}(i)$ and $\hat{m}_{\zeta,t}(i)$ (respectively) of these parameters at t . Suppose further that the mechanism they use for learning these parameters means that either:

1) there exists some $T \in \mathbb{Z}$ such that for all $t \geq T$, and all agents i and j , $\hat{m}_{\varepsilon,t}(i) = \hat{m}_{\varepsilon,t}(j)$ and $\hat{m}_{\zeta,t}(i) = \hat{m}_{\zeta,t}(j)$, or,

2) for all agents i and j $\text{plim}_{t \rightarrow \infty} \frac{\hat{m}_{\varepsilon,t}(i)}{\hat{m}_{\varepsilon,t}(j)} = 1$ and

$$\text{plim}_{t \rightarrow \infty} \frac{(\hat{m}_{\zeta,t}(i) - \hat{m}_{\zeta,t}(j))' (\hat{m}_{\zeta,t}(i) - \hat{m}_{\zeta,t}(j))}{\hat{m}_{\zeta,t}(i)' \hat{m}_{\zeta,t}(i)} = 0,^{20}$$

then if $m_{\varepsilon,t}(i) \neq 0$ for all t and i , and conditions 1), 2) and 4) of **Corollary 2.2** are satisfied, then all estimates of **reduced form** parameters will converge in probability to the true values, and agents' expectations will converge in probability to their values under the full

²⁰ Condition 1) is strictly encompassed by condition 2), but the former will be more useful in practice.

information, rational expectations solution. If in addition condition 3) of **Corollary 2.2** is satisfied, then all estimates of **structural** parameters will also converge.

The proof of the result under condition 1) of this proposition follows from **Proposition 2**. Under condition 2) the result follows from the fact that condition 2) implies that asymptotically the measurement error induced by treating an idiosyncratic estimate as an aggregate one is dominated by the signal, so the estimates will remain consistent, at least when x_t is non-explosive.

The set of learning mechanisms covered by **Corollary 2.3** and **Corollary 2.4** includes a very large number of plausible learning mechanisms. In the below, we mention three of particular interest.

4.2. Guaranteed learning of SMSV solutions

Corollary 2.3 guarantees convergence to any SMSV solution, given minimal conditions. Again, since these technical conditions are a little opaque, in Figure 3 we present simulation evidence demonstrating the broad convergence of our algorithm. The rows of Figure 3 correspond to the same rows of Figure 2 (identical parameters were used).

As in section 3.5, we make 2^{14} simulation runs, each of length 2^8 . For the sake of numerical stability, we again use a projection facility, with all reduced form and structural parameters constrained to lie in the interval $[-1000,1000]$. We also have an eight period burn-in, during which expectations are given by their SMSV solution, plus ζ_t (always a scalar). For all simulations, we use the ML algorithm for parameter estimation, due to its greater efficiency.²¹

²¹ Again, we only search for a local maximum, using the constraints as set up in footnote 18. To further increase the chance of finding a global maximum however, each period we try starting the optimisation routine at two different points: last period's estimate, and the AEIV solution.

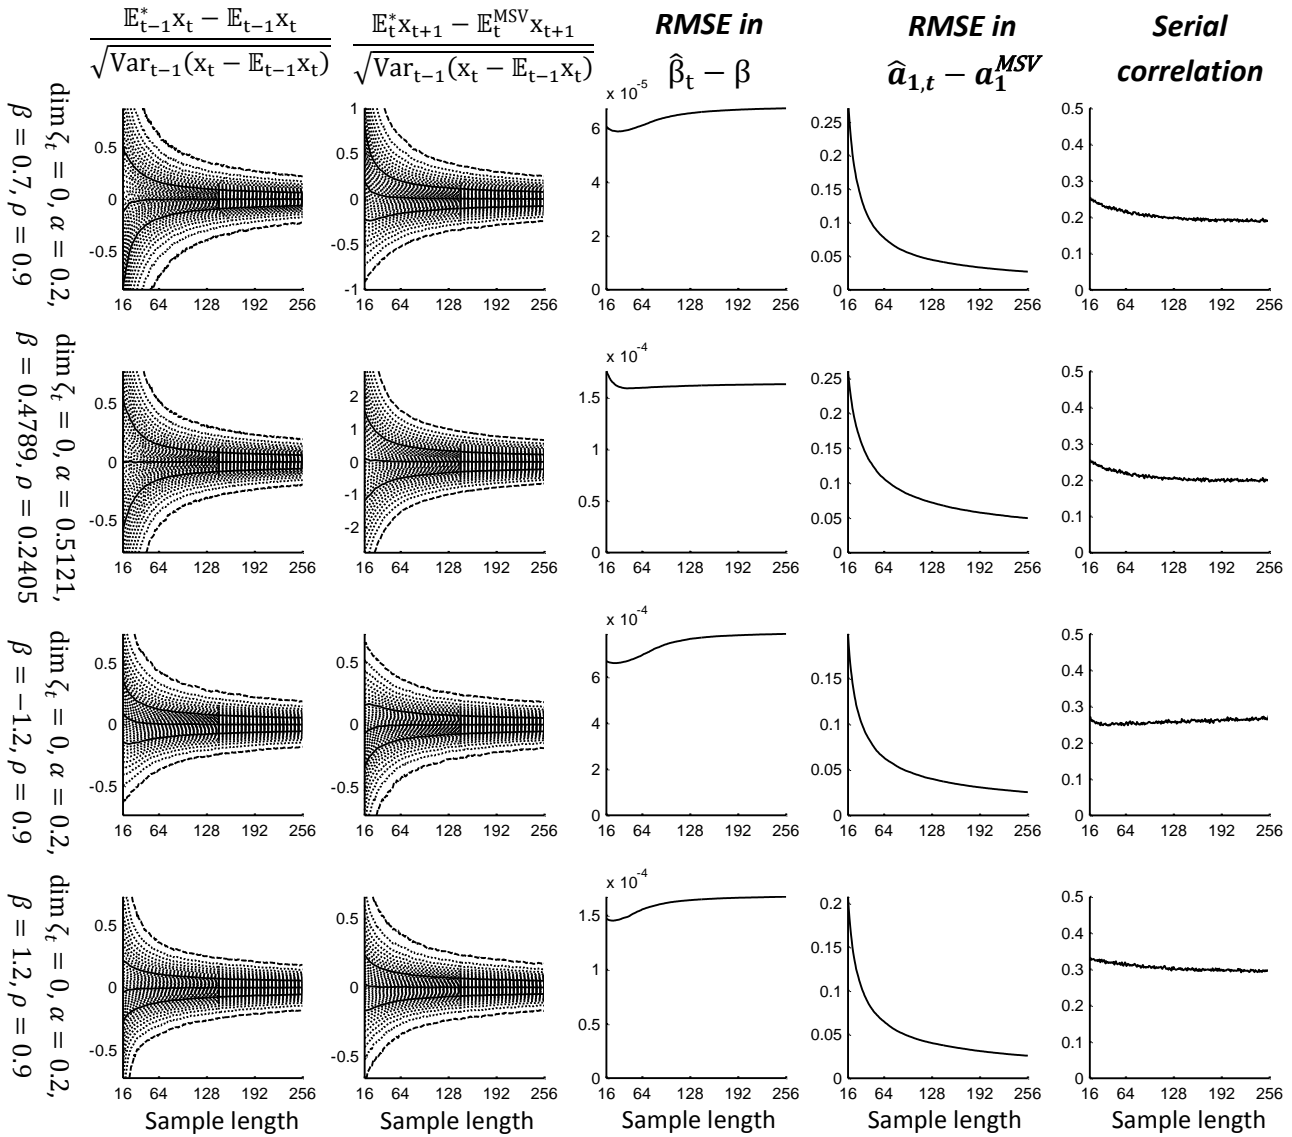


Figure 3: Results from simulations of sophisticated SMSV learners, from 2^{14} runs.

See text (section 4.2) for full details.

The first column of Figure 3 presents the distribution of the difference between the expectations formed by our sophisticated agents, and the expectations that would be formed by fully informed, fully rational agents in the same economy, normalised by the full information one-step ahead standard deviation. The second column presents the distribution of the difference between our agents' expectations and the SMSV solution,²² with the same normalisation. In all cases, it is clear that we have rapid convergence to the SMSV solution, and even faster convergence to rationality.

²² Given by $E_t^{MSV} x_{t+1} = a_1^{MSV} x_t - a_2^{MSV} x_{t-1} - c^{MSV}$

The third column presents the 95% trimmed RMSE in agents' estimates of β , and the fourth column does the same for a_1^{MSV} . In line with our theoretical results, while β does not appear to converge, agents' estimates of a_1^{MSV} converge to the truth in all cases. (The RMSE in β is nonetheless very small.) Finally, the fifth column presents the mean p-value from a (one-sample) bootstrapped LM test of serial correlation in expectational errors, at one lag. If information is being used fully efficiently, there should be no serial correlation, and these mean p-values should be equal to 0.5. While our found p-values are not quite so high, in all cases they are comfortably above 0.2 at all lags, so an econometrician would not reject the null of no serial correlation, at any standard significance level. Thus although this sophisticated learning algorithm is still not quite fully rational, it is close enough to rationality that users of it could not detect their own deviations from rationality.

Under standard OLS learning, there are non-learnable stationary MSV solutions such as the one in the final row of Figure 3, so by this measure the present learning algorithm is an improvement. However, it is in no sense an answer to Cochrane's (2009) challenge for learnability to "save new-Keynesian models". This learning algorithm is only reasonable if agents already believe that the solution is of the SMSV form, an assumption that is not justified by anything in the model. That dramatically different results may obtain with different learning mechanisms is made clear by the next one presented.

4.3. Learning any sunspot solution (with positive density)

Suppose, that agent i believes that as well as having access to all the same information as them, everyone else in the economy also had access to the additional information that $m_{\varepsilon,t} \equiv m_{\varepsilon,0}$ and $m_{\zeta,t} \equiv m_{\zeta,0}$, where $m_{\varepsilon,0}$ and $m_{\zeta,0}$ are constants, unknown to agent i .

Let us define:

$$e_t := \frac{\hat{\beta}_{t-1}}{\hat{\sigma}_{t-1}} \left[\frac{1}{\hat{\beta}_{t-1}} x_t + \hat{a}_{2,t-1} x_{t-1} + \hat{a}_{3,t-1} x_{t-2} + \hat{\rho}_{t-1} \mathbb{E}_{t-1}^* x_t + \hat{c}_{t-1} - \mathbb{E}_t^* x_{t+1} \right], \quad (4.3)$$

then:

$$\begin{bmatrix} e_t & \zeta_t' \end{bmatrix} \begin{bmatrix} m_{\varepsilon,t-1} \\ m_{\zeta,t-1} \end{bmatrix} \approx x_t - \mathbb{E}_{t-1}^* x_t =: \eta_t^*,$$

where the approximation is exact when $m_{\varepsilon,t-1} = m_{\varepsilon,0}$. (Away from this point, agent i 's estimate of $a_{1,t}$ will differ from the true value, introducing error into their estimates of α_t , etc..) The natural estimate of $m_{\varepsilon,t}$ and $m_{\zeta,t}$ is then:

$$\begin{bmatrix} \hat{m}_{\varepsilon,t} \\ \hat{m}_{\zeta,t} \end{bmatrix} = \begin{bmatrix} e_t & \zeta_t' \\ e_t & \zeta_t' \end{bmatrix}^+ \begin{bmatrix} \eta_t^* \\ \eta_t^* \end{bmatrix},$$

where superscript $+$ denotes the Moore-Penrose pseudo-inverse.²³ By the standard properties of least squares estimates, this will converge on the truth, and indeed despite the presence of the approximation in the previous equation this will happen exactly in finite time, providing estimates of other parameters are updated recursively.²⁴

In the case we are chiefly concerned with, everyone is learning simultaneously, so by the properties of the Moore-Penrose pseudo-inverse, we will have $\hat{m}_{\varepsilon,t} \equiv \hat{m}_{\varepsilon,1} = \frac{e_1}{e_1^2 + \zeta_1' \zeta_1} \eta_1^*$ and $\hat{m}_{\zeta,t} \equiv \hat{m}_{\zeta,1} = \frac{\zeta_1}{e_1^2 + \zeta_1' \zeta_1} \eta_1^*$, ex-post justifying the constancy assumption that motivated the learning method. By varying initial beliefs we may attain any value for η_1^* , and hence any value for $\hat{m}_{\varepsilon,1}$ and $\hat{m}_{\zeta,1}$. So with stochastic initial beliefs (a public signal perhaps), any

²³ This is of course the standard linear regression formula when $t \geq \dim \zeta_t$.

²⁴ In this situation, agent i should update their estimates of $a_{1,1}$ in all periods t with $t \geq 1$. I.e. in period t , they should estimate $a_{1,1}$ as $\frac{1}{\hat{\beta}_1} \left(1 - \frac{\hat{\sigma}_1}{\hat{m}_{\varepsilon,t}} \right)$. Based on this revised estimate of $a_{1,1}$, they can then re-estimate α_2 , etc., and then $a_{1,2}$, etc., and so on. Armed with this set of new estimates, they can then re-estimate $m_{\varepsilon,t}$ and $m_{\zeta,t}$, repeating the entire procedure until they converge on a fixed point. After $1 + \dim \zeta_t$ periods have elapsed, there may possible be multiple such fixed points, however, the next period, with probability 1 only one will remain.

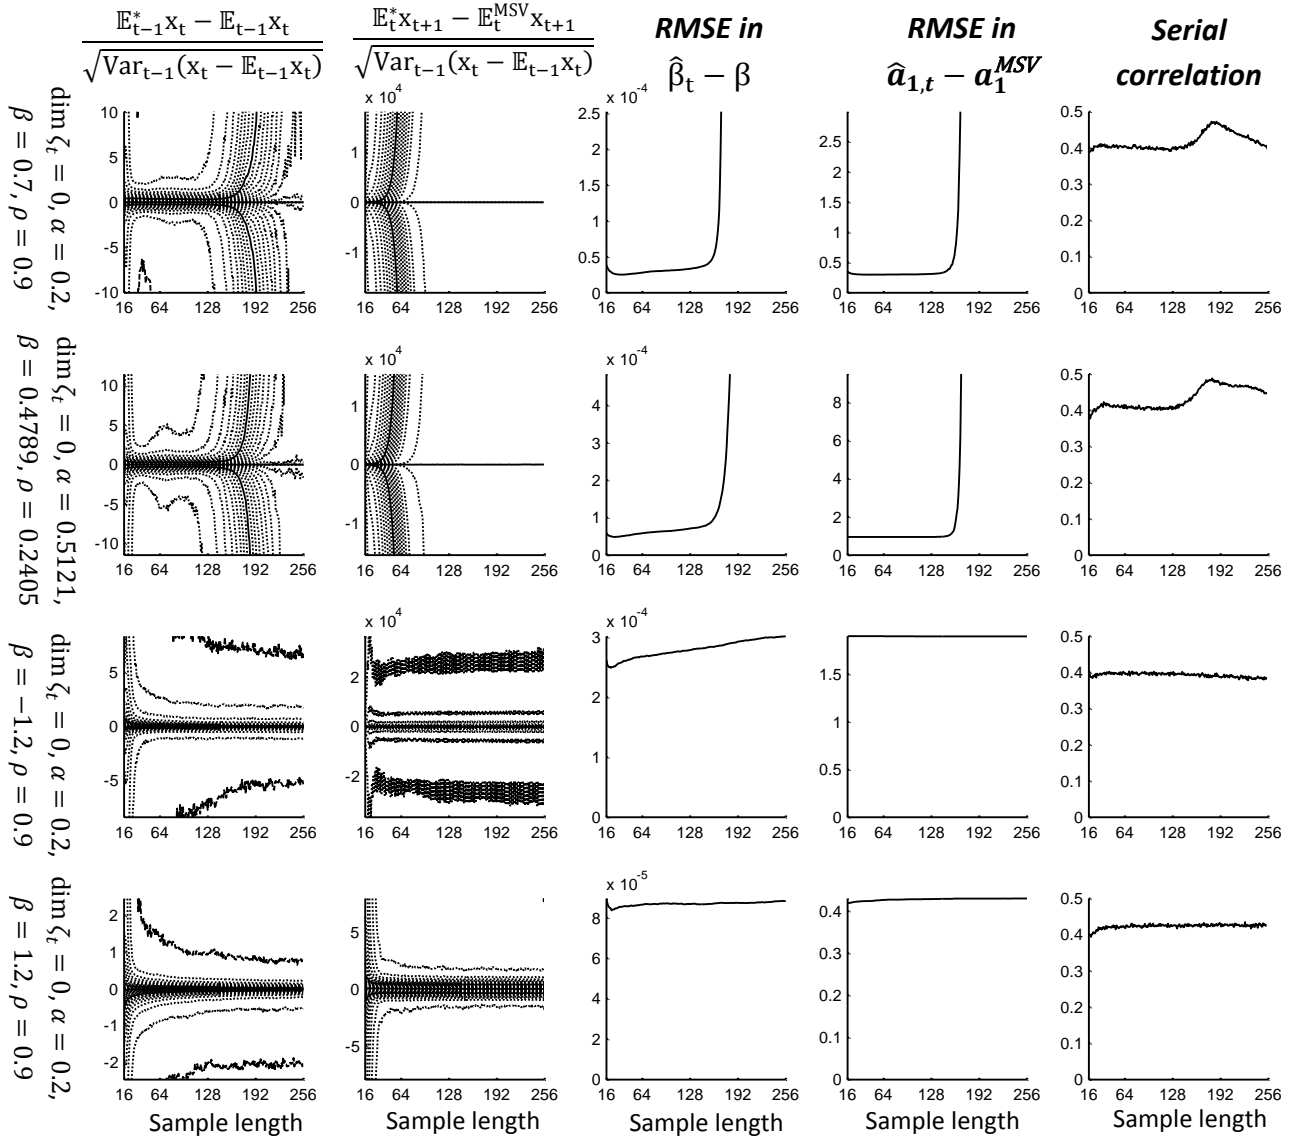


Figure 4: Results from simulations of sophisticated sunspot learners, from 2^{14} runs.
See text (section 4.3) for full details.

solution is attainable with positive density, and expectations will converge to rationality with probability one (at least given the relevant technical conditions), by **Corollary 2.4**.²⁵

This learning method is readily extended to the case in which agents believe that $m_{\varepsilon,t}$ and $m_{\zeta,t}$ are constant until a certain event occurs. Possible candidates for these events include changes of central bank governors, changes of governments, financial crashes and natural disasters. In this case, each time the event occurs, a new draw for $\hat{m}_{\varepsilon,t}$ and $\hat{m}_{\zeta,t}$ will be

²⁵ The solutions with $\hat{m}_{\zeta,1} = 0$ are not guaranteed to converge, but the set of such solutions is of measure zero in the whole space.

taken, and they will remain fixed at those values until the event occurs again. In the extreme case in which the event occurs every period, we have that $\hat{m}_{\varepsilon,t} = \frac{e_t}{e_t^2 + \zeta_t' \zeta_t} \eta_t^*$ and $\hat{m}_{\zeta,t} = \frac{\zeta_t}{e_t^2 + \zeta_t' \zeta_t} \eta_t^*$. Since $\mathbb{E}_{t-1} \eta_t^{*2} - (\hat{m}_{\varepsilon,t-1}^2 + \hat{m}_{\zeta,t-1}' \hat{m}_{\zeta,t-1})$ tends to 0 as $t \rightarrow \infty$, this means $\mathbb{E}_{t-1}(\eta_{t+1}^{*2}) - \mathbb{E}_{t-1}(\eta_t^{*2})$ tends to 0 too, so the variance of expectational errors follows a random walk asymptotically, providing endogenous stochastic volatility.

In Figure 4 we show simulations of this learning method, with the exact same set-up as in section 4.2. (We do not bound $\hat{m}_{\varepsilon,t}$ or $\hat{m}_{\zeta,t}$ however.) Since initial estimates of e_t are highly inaccurate, we assume all agents update their estimates of $\hat{m}_{\varepsilon,t}$ and $\hat{m}_{\zeta,t}$ in each of the first 8 periods after the end of the burn-in (i.e. periods 9 to 16), but not in any future period.

In the two cases in which only the SMSV solution is stationary, expectations asymptotically diverge from rationality. However, there is an initial period of rapid convergence, so it is hard to know if this divergence is merely driven by the numerical errors stemming from the explosive behaviour of x_t . (Either hypothesis would be consistent with our theoretical results, as these do not cover cases in which x_t grows exponentially or faster.) In the two “indeterminate” cases, expectations rapidly converge to rationality, though not to the MSV solution, implying a sunspot solution has been learnt. While structural parameter estimates are very close to the truth in all cases, they do not appear to be converging. This again is consistent with our theoretical results if reduced form parameters have converged too quickly. Finally, note that there is even less evidence of serial correlation in this sunspot case, so again the agents in the model would not be able to detect their own departure from rationality.

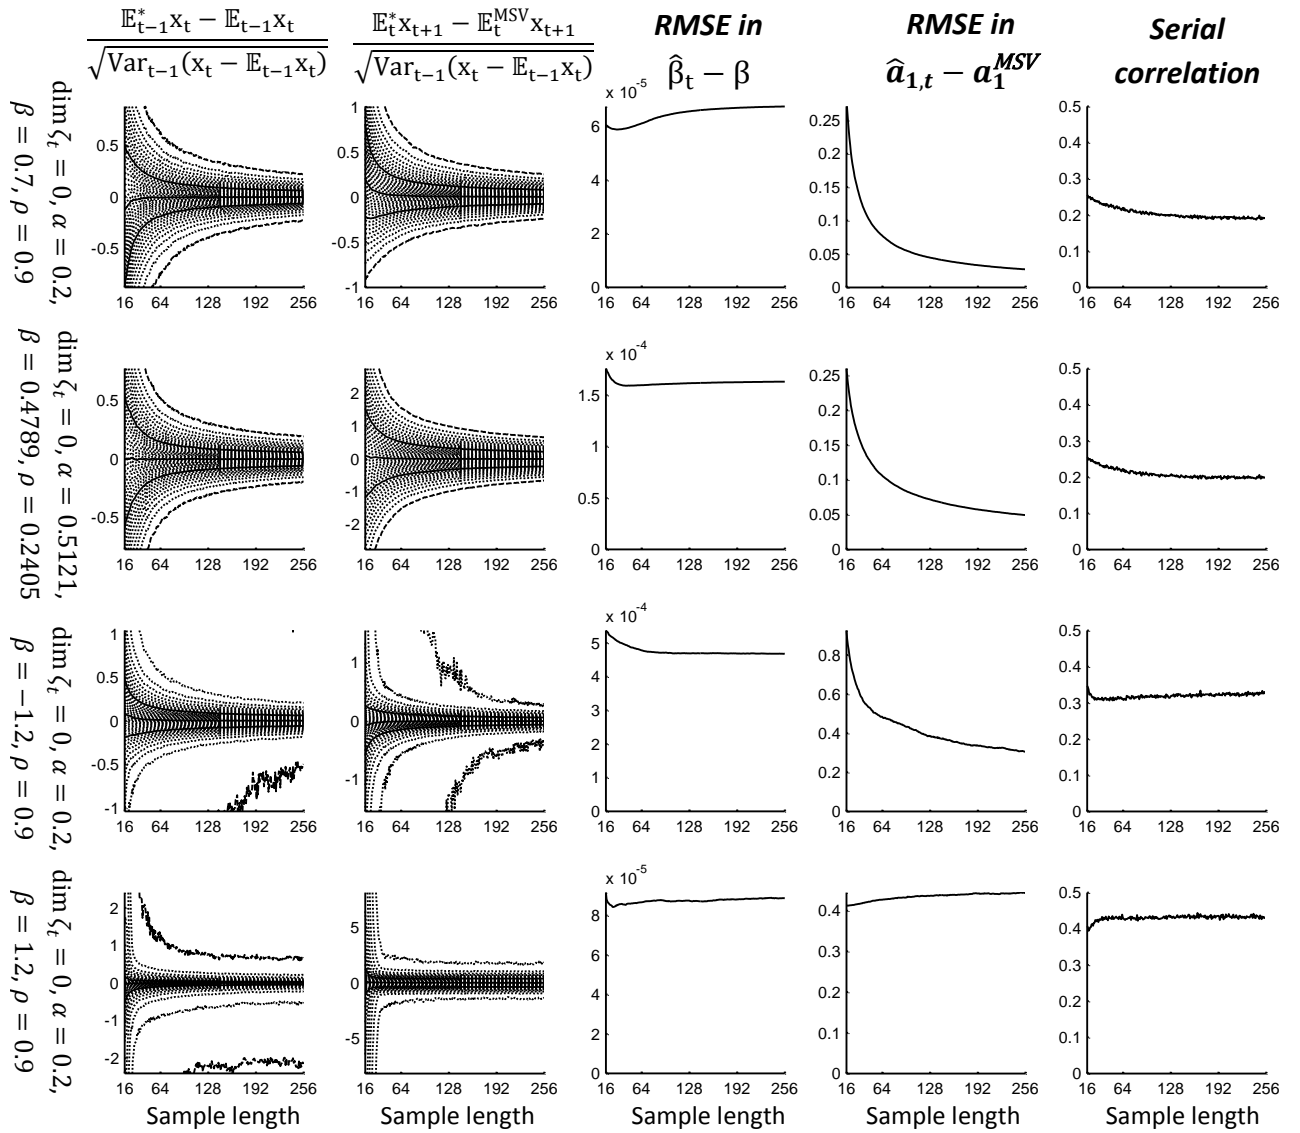


Figure 5: Results from simulations of sophisticated transversality learners, from 2^{14} runs. See text (section 4.4) for full details.

4.4. Learning in the presence of transversality constraints

Finally, suppose that in the model under consideration, x_t is restricted by a transversality constraint. (To recap, this is not the case for inflation.) Then if agents are ever confident they are in an indeterminate region of the parameter space, they should switch to the SMSV solution. This suggests that agents should begin using the sunspot learning method from the previous section. If however their estimates ever imply that $|\hat{a}_t + \hat{\beta}_t| < 1$, then they should switch to forming MSV expectations. If at a later date they again come to believe that

$|\hat{\alpha}_t + \hat{\beta}_t| > 1$, they should switch back to the general sunspot solution, with updated values for $\hat{m}_{\varepsilon,t}$ and $\hat{m}_{\zeta,t}$.

Figure 5 presents simulations of this learning method. Performance is an amalgam of the previous two cases, with convergence to the SMSV solution under determinacy, and convergence to a sunspot solution otherwise.

5. Conclusion

This paper has set forward a family of macroeconomic learning algorithms that are correctly specified, even along the transition path. Our simulations and theoretical results imply that vastly more equilibria are learnable via these algorithms than via traditional learning methods, implying that learnability cannot be used for equilibrium selection. We have also demonstrated that from observing traditional macroeconomic learners we may identify all a model's structural parameters, providing those traditional learners are running a regression that encompasses the general solution to the model.

The new estimators produced in this paper have many practical applications. In future empirical work we hope to use them to assess whether the Federal Reserve has ever pursued a policy satisfying the Taylor principle, something that was not possible until now due to the non-identification of the key parameter given unobserved, auto-correlated monetary policy shocks. We also hope to look for empirical evidence on whether real world macroeconomic learning is best described by the traditional algorithm or one of our new, misspecification-free methods.

6. References

- Blanchard, Olivier Jean, and Charles M. Kahn. 1980. The Solution of Linear Difference Models under Rational Expectations. *Econometrica* 48, no. 5: 1305–1311.
- Christopeit, Norbert, and Michael Massmann. 2010. Consistent Estimation of Structural Parameters in Regression Models with Adaptive Learning. *SSRN eLibrary* (August 23). http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1653325.
- Cochrane, John H. 2009. Can learnability save new-Keynesian models? *Journal of Monetary Economics* 56, no. 8 (November): 1109–1113.
- . 2011. Determinacy and Identification with Taylor Rules. *Journal of Political Economy* 119, no. 3: pp. 565–615.
- Ellison, Martin, and Joseph Pearlman. 2011. Saddlepath learning. *Journal of Economic Theory* 146, no. 4 (July): 1500–1519.
- Evans, George, and Seppo Honkapohja. 2001. *Learning and expectations in macroeconomics*. Frontiers of Economic Research. Princeton and Oxford: Princeton University Press.
- Gabrielsen, Arne. 1978. Consistency and identifiability. *Journal of Econometrics* 8, no. 2 (October): 261–263.
- Holden, Tom. 2008. *Rational macroeconomic learning in linear expectational models*. MPRA Paper. University Library of Munich, Germany, May. <http://ideas.repec.org/p/pra/mprapa/10872.html>.
- Horn, Roger A., and Charles R. Johnson. 1985. *Matrix analysis*. Cambridge [Cambridgeshire]; New York: Cambridge University Press.
- Lai, Tze Leung, and Ching Zong Wei. 1982. Least Squares Estimates in Stochastic Regression Models with Applications to Identification and Control of Dynamic Systems. *The Annals of Statistics* 10, no. 1 (March): 154–166.
- Levine, Paul, Joseph Pearlman, George Perendia, and Bo Yang. 2012. Endogenous Persistence in an estimated DSGE Model Under Imperfect Information*. *The Economic Journal*: no–no.
- Lubik, Thomas A., and Frank Schorfheide. 2003. Computing sunspot equilibria in linear rational expectations models. *Journal of Economic Dynamics and Control* 28, no. 2 (November): 273–285.
- Marcet, Albert, and Thomas J. Sargent. 1989. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, no. 2: 337–368.
- . 1992. *Speed of convergence of recursive least squares learning with ARMA perceptions*. Economics Working Paper. Department of Economics and Business, Universitat Pompeu Fabra. <http://ideas.repec.org/p/upf/upfgen/15.html>.
- Mavroeidis, Sophocles, and Yori Zwols. 2007. LiRE - An Ox package for solving Linear Rational Expectations Models - Manual Version 3.0. December 7. http://www.econ.brown.edu/fac/Sophocles_Mavroeidis/code/lire3.pdf.
- McCallum, Bennett T. 1983. On non-uniqueness in rational expectations models: An attempt at perspective. *Journal of Monetary Economics* 11, no. 2: 139–168.
- Pearlman, Joseph, David Currie, and Paul Levine. 1986. Rational expectations models with partial information. *Economic Modelling* 3, no. 2: 90–105.
- Quarteroni, Alfio, Riccardo Sacco, and Fausto Saleri. 2000. *Numerical mathematics*. Texts in applied mathematics, 37. New York: Springer.
- Sims, Christopher A. 2002. Solving linear rational expectations models. *Computational Economics* 20, no. 1: 1–20.

7. Appendices

7.1. FREE solutions for arbitrary linear models

We now extend the structure of (2.3) to the general multivariate case:

$$Kx_t = Ax_{t-1} + B\mathbb{E}_t x_{t+1} + \gamma + \Sigma_s s_t$$

where:

$$s_t = Ps_{t-1} + \Sigma_\varepsilon \varepsilon_t$$

for the arbitrary matrices K, A, B, P, Σ_s and Σ_ε , the vector γ and the shock $\varepsilon_t \sim \text{NIID}(0, I)$. Initially, we suppose that there are no transversality conditions restricting any of the components of x_t .

Again defining the expectational error by $\eta_t := x_t - \mathbb{E}_{t-1} x_t$, when B and Σ_s have linearly independent columns, from the properties of the Moore-Penrose pseudoinverse (denoted by $+$), we have that:

$$\begin{aligned} \mathbb{E}_t x_{t+1} = & B^+(K + \Sigma_s P \Sigma_s^+ B)x_t - B^+(A + \Sigma_s P \Sigma_s^+ K)x_{t-1} + B^+ \Sigma_s P \Sigma_s^+ A x_{t-2} \\ & - B^+ \Sigma_s (I - P) \Sigma_s^+ \gamma - B^+ \Sigma_s \Sigma_\varepsilon \varepsilon_t - B^+ \Sigma_s P \Sigma_s^{-1} B \eta_t. \end{aligned}$$

As before, without loss of generality we may assume that $\eta_t = M_{\varepsilon, t-1} \varepsilon_t + M_{\zeta, t-1} \zeta_t$, for some sunspot shock ζ_t uncorrelated with ε_t (and satisfying $\mathbb{E}_{t-1} \zeta_t = 0$, $\mathbb{E}_{t-1} \zeta_t \zeta_t' = I$).

Then, if $M_{\varepsilon, t-1}$ has linearly independent columns:

$$\begin{aligned} \mathbb{E}_t x_{t+1} = & B^+(K - \Sigma_s \Sigma_\varepsilon M_{\varepsilon, t-1}^+) x_t - B^+(A + \Sigma_s P \Sigma_s^+ K)x_{t-1} + B^+ \Sigma_s P \Sigma_s^+ A x_{t-2} \\ & + B^+(\Sigma_s P \Sigma_s^+ B + \Sigma_s \Sigma_\varepsilon M_{\varepsilon, t-1}^+) \mathbb{E}_{t-1} x_t - B^+ \Sigma_s (I - P) \Sigma_s^+ \gamma \\ & + B^+ \Sigma_s \Sigma_\varepsilon M_{\varepsilon, t-1}^+ M_{\zeta, t-1} \zeta_t. \end{aligned}$$

This expression no longer contains either ε_t or s_t . Thus, when B and Σ_s have linearly independent columns, almost all rational expectations solutions to the original model are FREE, i.e. they are implementable by agents who cannot observe the model's fundamental shocks.

More generally, there will be transversality conditions restricting some variables, and B and Σ_s will not have linearly independent columns. To solve this case, we closely follow Mavroeidis and Zwols's (2007) presentation of Lubik and Schorfheide's (2003) extension to the irregular case of Sims's (2002) method for solving rational expectations models, which is itself more general than that of Blanchard and Kahn (1980). The majority of the results here that are not due to Mavroeidis, Zwols, Lubik, Schorfheide or Sims were first shown in an earlier working paper by this author (Holden 2008).

With the model set-up as before, let us define $v_t := \begin{bmatrix} x_t \\ \mathbb{E}_t x_{t+1} \end{bmatrix}$, $\Gamma_0 := \begin{bmatrix} K & -B \\ I & 0 \end{bmatrix}$, $\Gamma_1 := \begin{bmatrix} A & 0 \\ 0 & I \end{bmatrix}$, $\mu := \begin{bmatrix} \gamma \\ 0 \end{bmatrix}$, $\Psi := \begin{bmatrix} \Sigma_s \\ 0 \end{bmatrix}$ and $\Pi := \begin{bmatrix} 0 \end{bmatrix}$. We then have the general canonical form we will solve here:

$$\Gamma_0 v_t = \Gamma_1 v_{t-1} + \mu + \Psi s_t + \Pi \eta_t.$$

In deriving the conditions for the existence of a rational expectations equilibria (REE) below, we will not assume anything about the structure of v_t , η_t , Γ_0 , Γ_1 , μ , Ψ , Π , P , Σ_s or Σ_ε (beyond the fact that η_t must be chosen subject to $\mathbb{E}_{t-1} \eta_t = 0$). We will also be able to derive sufficient conditions for the existence of a FREE in this fully general case. However, in deriving necessary conditions we will assume that $v_t = \begin{bmatrix} x_t \\ \mathbb{E}_t x_{t+1} \end{bmatrix}$ and $\eta_t = x_t - \mathbb{E}_{t-1} x_t$, as in the above.

By the generalized complex Schur decomposition (also known as the QZ decomposition) (Quarteroni, Sacco, and Saleri 2000) of the matrices Γ_0 and Γ_1 , there always exist possibly complex matrices Q , Z , $\Lambda = (\lambda_{i,j})_{i,j}$ and $\Omega = (\omega_{i,j})_{i,j}$ such that $Q^H \Lambda Z^H = \Gamma_0$, $Q^H \Omega Z^H = \Gamma_1$, Q and Z are unitary, Λ and Ω are upper triangular and a superscript H denotes conjugate transpose.

Now let $w_t = Z^H v_t$ for all $t \in \mathbb{Z}$, then if we pre-multiply the canonical form by Q we have:

$$\Lambda w_t = \Omega w_{t-1} + Q(\mu + \Psi s_t + \Pi \eta_t).$$

Providing Γ_0 and Γ_1 do not have zero eigenvalues corresponding to the same eigenvector²⁶

the QZ decomposition always exists and the set $\left\{ \left| \frac{\omega_{ii}}{\lambda_{ii}} \right| \mid i \in \{1, \dots, \dim v_t\} \right\} \subseteq \mathbb{R} \cup \{\infty\}$ is

unique even though the decomposition itself is not (Sims 2002). Thus, without loss of

generality we may assume that for $i < j$, $\left| \frac{\omega_{ii}}{\lambda_{ii}} \right| < \left| \frac{\omega_{jj}}{\lambda_{jj}} \right|$. Let \bar{u} be the number of i for which

$\left| \frac{\omega_{ii}}{\lambda_{ii}} \right| \leq 1$ and consider a partition of the matrices under consideration in which in each case

the top left block is of dimension $\bar{u} \times \bar{u}$ ²⁷.

We may then write:

$$\begin{bmatrix} \Lambda_{11} & \Lambda_{12} \\ 0 & \Lambda_{22} \end{bmatrix} \begin{bmatrix} w_{1,t} \\ w_{2,t} \end{bmatrix} = \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ 0 & \Omega_{22} \end{bmatrix} \begin{bmatrix} w_{1,t-1} \\ w_{2,t-1} \end{bmatrix} + \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} (\mu + \Psi s_t + \Pi \eta_t). \quad (7.1)$$

The second block of this equation is purely explosive by construction. More generally, we

may follow Sims (2002) and allow explosive combinations of variables that do not violate

transversality to enter into the upper block. In New-Keynesian models, inflation rates will

generally be such a variable.

²⁶ This means that there is one or more equation that places no restrictions on either v_t or v_{t-1} . This will create an additional source of indeterminacy in v_t and may also imply that one or more components of ε_t and η_t are linear combinations of the others. We, like both Sims and Lubik & Schorfheide, will not further investigate this avenue.

²⁷ This means that we are not treating unit roots as explosive. Doing this avoids some minor technical complications.

If agents expect a non-transversality violating path for v_t , from solving forward, following Sims (2002) and Mavroeidis and Zwols (2007), we must have:

$$\begin{aligned} w_{2,t} = \mathbb{E}_t w_{2,t} &= -\mathbb{E}_t \sum_{k=1}^{\infty} (\Omega_{22}^{-1} \Lambda_{22})^{k-1} \Omega_{22}^{-1} Q_2 \cdot (\mu + \Psi s_{t+k} + \Pi \eta_{t+k}) \\ &= -\sum_{k=0}^{\infty} (\Omega_{22}^{-1} \Lambda_{22})^k \Omega_{22}^{-1} Q_2 \cdot \Psi P^{1+k} s_t - \left[\sum_{k=0}^{\infty} (\Omega_{22}^{-1} \Lambda_{22})^k \right] \Omega_{22}^{-1} Q_2 \cdot \mu \\ &= SP s_t + (\Lambda_{22} - \Omega_{22})^{-1} Q_2 \cdot \mu, \end{aligned}$$

where S is the solution to the Stein equation²⁸:

$$\Omega_{22}^{-1} \Lambda_{22} SP - S = \Omega_{22}^{-1} Q_2 \cdot \Psi$$

and where the sums are well defined since the eigenvalues of $\Omega_{22}^{-1} \Lambda_{22}$ are strictly in the unit circle by construction (and Ω_{22} is invertible by construction). Note that for S to have linearly independent columns, it is necessary that $\dim w_{2,t} \geq \dim s_t$.

Consequently (following Mavroeidis and Zwols (2007)), $\mathbb{E}_{t+1} w_{2,t} = \mathbb{E}_t w_{2,t}$, and so:

$$\begin{aligned} -\mathbb{E}_{t+1} \sum_{k=1}^{\infty} (\Omega_{22}^{-1} \Lambda_{22})^{k-1} \Omega_{22}^{-1} Q_2 \cdot (\mu + \Psi s_{t+k} + \Pi \eta_{t+k}) \\ = -\mathbb{E}_t \sum_{k=1}^{\infty} (\Omega_{22}^{-1} \Lambda_{22})^{k-1} \Omega_{22}^{-1} Q_2 \cdot (\mu + \Psi s_{t+k} + \Pi \eta_{t+k}) \end{aligned}$$

i.e.
$$\Omega_{22} S \Sigma_{\varepsilon} \varepsilon_{t+1} = Q_2 \cdot \Pi \eta_{t+1} \tag{7.2}$$

(using the fact that Ω_{22} is of full rank and the definition of S). This is the key constraint limiting expectations. If $P = 0$, then $S = -\Omega_{22}^{-1} Q_2 \cdot \Psi$ so under the normalisation $\Sigma_{\varepsilon} = I$, it collapses to the expression given in Lubik and Schorfheide (2003).

²⁸ This equation has a unique solution providing none of the eigenvalues of P are in the set $\left\{ \left| \frac{\omega_{ii}}{\lambda_{ii}} \right| \mid i \in \{\bar{u} + 1, \dots, \dim v_t\} \right\}$, which holds automatically providing the autoregressive process for ε_t is non-explosive. The (non-numerically robust) solution is given by: $\text{vec } S = (P' \otimes \Omega_{22}^{-1} \Lambda_{22} - I)^{-1} \text{vec } \Omega_{22}^{-1} Q_2 \cdot \Psi$.

By the singular value decomposition (SVD) (Horn and Johnson 1985) of $Q_2 \cdot \Pi$ and $\Omega_{22} S \Sigma_\varepsilon$ we

can write $Q_2 \cdot \Pi = UDV^H = [U_{\cdot 1} \ U_{\cdot 2}] \begin{bmatrix} D_{11} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_{\cdot 1}^H \\ V_{\cdot 2}^H \end{bmatrix} = U_{\cdot 1} D_{11} V_{\cdot 1}^H$ and $\Omega_{22} S \Sigma_\varepsilon = \hat{U} \hat{D} \hat{V}^H =$

$[\hat{U}_{\cdot 1} \ \hat{U}_{\cdot 2}] \begin{bmatrix} \hat{D}_{11} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{V}_{\cdot 1}^H \\ \hat{V}_{\cdot 2}^H \end{bmatrix} = \hat{U}_{\cdot 1} \hat{D}_{11} \hat{V}_{\cdot 1}^H$ where U, V, \hat{U} and \hat{V} are unitary and D_{11} and \hat{D}_{11}

have strictly positive diagonals and zeroes elsewhere, and where H denotes the Hermitian transpose. Pre-multiplying the constraint (7.2) by $U_{\cdot 1} U_{\cdot 1}^H$ then gives that:

$$\begin{aligned} U_{\cdot 1} U_{\cdot 1}^H \Omega_{22} S \Sigma_\varepsilon \varepsilon_{t+1} &= U_{\cdot 1} U_{\cdot 1}^H Q_2 \cdot \Pi \eta_{t+1} = U_{\cdot 1} U_{\cdot 1}^H U_{\cdot 1} D_{11} V_{\cdot 1}^H \eta_{t+1} = U_{\cdot 1} D_{11} V_{\cdot 1}^H \eta_{t+1} = Q_2 \cdot \Pi \eta_{t+1} \\ &= \Omega_{22} S \Sigma_\varepsilon \varepsilon_{t+1} \end{aligned}$$

(by the constraint and the unitarity of U). Thus since ε_{t+1} may take the value $\hat{V}_{\cdot 1} \hat{D}_{11}^{-1} v$ for any v , by the unitarity of \hat{V} , we must have:

$$U_{\cdot 1} U_{\cdot 1}^H \hat{U}_{\cdot 1} = \hat{U}_{\cdot 1}. \quad (7.3)$$

This condition is also sufficient for the existence of a solution, which we now demonstrate by exhibiting an explicit solution.

Let $q := \text{rank } Q_2 \cdot \Pi$, so that D_{11} is of dimension $q \times q$. Then following Lubik and Schorfheide (2003), we posit the following set of solutions for the forecast errors η_t :

$$\eta_t = [V_{\cdot 1} \ V_{\cdot 2}] \begin{bmatrix} D_{11}^{-1} U_{\cdot 1}^H \Omega_{22} S \Sigma_\varepsilon \\ M_{\varepsilon, t-1} \end{bmatrix} \varepsilon_t + [V_{\cdot 1} \ V_{\cdot 2}] \begin{bmatrix} 0 \\ M_{\zeta, t-1} \end{bmatrix} \zeta_t, \quad (7.4)$$

where ζ_t is an arbitrary vector of sunspot shocks, uncorrelated with ε_t , and $M_{\varepsilon, t-1}$ and $M_{\zeta, t-1}$ are arbitrary matrices of size $(\dim \eta_t - q) \times \dim \varepsilon_t$ and $(\dim \eta_t - q) \times \dim \zeta_t$ respectively, known at $t - 1$. (The possibility of time variation in $M_{\varepsilon, t-1}$ and $M_{\zeta, t-1}$ was not noticed by Lubik and Schorfheide (2003).)

When the condition (7.3) holds, by the unitarity of V we have that:

$$\begin{aligned} Q_2 \cdot \Pi \eta_t &= U_{\cdot 1} D_{11} V_{\cdot 1}^H \eta_t \\ &= (U_{\cdot 1} D_{11} V_{\cdot 1}^H V_{\cdot 1} D_{11}^{-1} U_{\cdot 1}^H \Omega_{22} S \Sigma_\varepsilon + U_{\cdot 1} D_{11} V_{\cdot 1}^H V_{\cdot 2} M_{\varepsilon, t-1}) \varepsilon_t + U_{\cdot 1} D_{11} V_{\cdot 1}^H V_{\cdot 2} M_{\zeta, t-1} \zeta_t \\ &= U_{\cdot 1} U_{\cdot 1}^H \Omega_{22} S \Sigma_\varepsilon \varepsilon_t = U_{\cdot 1} U_{\cdot 1}^H \hat{U}_{\cdot 1} \hat{D}_{11} \hat{V}_{\cdot 1}^H \varepsilon_t = \hat{U}_{\cdot 1} \hat{D}_{11} \hat{V}_{\cdot 1}^H \varepsilon_t = \Omega_{22} S \Sigma_\varepsilon \varepsilon_t \end{aligned}$$

and so the constraint (7.2) does indeed hold. It is immediate from this solution for the forecast errors that the model has a unique solution if and only if $q = \dim \eta_t$.

In order for there to be a FREE solution, we must be able to express ε_t as a function of η_t

and ζ_t . If we pre-multiply the above solution for η_t by $\begin{bmatrix} \Omega_{22}^{-1} U_{\cdot 1} D_{11} & 0 \\ 0 & I \end{bmatrix} V^H$, using condition

(7.3) and the unitarity of V we have that:

$$\begin{bmatrix} S \Sigma_\varepsilon \\ M_{\varepsilon, t-1} \end{bmatrix} \varepsilon_t = \begin{bmatrix} \Omega_{22}^{-1} U_{\cdot 1} D_{11} & 0 \\ 0 & I \end{bmatrix} V^H \eta_t - \begin{bmatrix} 0 \\ M_{\zeta, t-1} \end{bmatrix} \zeta_t.$$

Therefore, a FREE solution will certainly exist if $\begin{bmatrix} S \Sigma_\varepsilon \\ M_{\varepsilon, t-1} \end{bmatrix}$ has linearly independent columns

for all t , since when this holds, from standard results on the Moore-Penrose pseudo-inverse we have that:

$$\varepsilon_t = \begin{bmatrix} \Sigma_\varepsilon^H S^H S \Sigma_\varepsilon \\ M_{\varepsilon, t-1}^H M_{\varepsilon, t-1} \end{bmatrix}^{-1} \left[\begin{bmatrix} \Sigma_\varepsilon^H S^H \Omega_{22}^{-1} U_{\cdot 1} D_{11} & M_{\varepsilon, t-1}^H \end{bmatrix} V^H \eta_t - M_{\varepsilon, t-1}^H M_{\zeta, t-1} \zeta_t \right]$$

and so it is as if ε_t is in even the limited information set.

When $\dim \eta_t - q \geq \dim \varepsilon_t$, $\begin{bmatrix} S \Sigma_\varepsilon \\ M_{\varepsilon, t-1} \end{bmatrix}$ will have linearly independent columns for almost all

$M_{\varepsilon, t-1}$.²⁹ More generally, we require that $\text{rank } S \Sigma_\varepsilon + \dim \eta_t - q \geq \dim \varepsilon_t$.

²⁹ With $q = 0$, this gives a generalisation of our initial result to the case in which B and Σ_s do not have full rank.

Now by (7.3), $Q_2 \cdot \Pi V_1 D_{11}^{-1} U_1^H \Omega_{22} S \Sigma_\varepsilon = \Omega_{22} S \Sigma_\varepsilon$, thus $\text{span } S \Sigma_\varepsilon = \text{span } \Omega_{22} S \Sigma_\varepsilon \subseteq \text{span } Q_2 \cdot \Pi$

and so $\text{rank } S \Sigma_\varepsilon \leq \text{rank } Q_2 \cdot \Pi = q$. Thus, if it is to be the case that $\begin{bmatrix} S \Sigma_\varepsilon \\ M_{\varepsilon, t-1} \end{bmatrix}$ has linearly independent columns, we must have that:

$$\dim \varepsilon_t - (\dim \eta_t - q) \leq \text{rank } S \Sigma_\varepsilon \leq \text{rank } Q_2 \cdot \Pi = q,$$

which implies $\dim \varepsilon_t \leq \dim \eta_t$. In the special case in which $\dim \varepsilon_t = \dim \eta_t$, these inequalities become equalities, meaning that we must have $\text{span } \Omega_{22} S \Sigma_\varepsilon = \text{span } Q_2 \cdot \Pi$, and hence $\widehat{U}_1 \widehat{U}_1^H U_1 = U_1$, by (7.3).

The fact that $\begin{bmatrix} S \Sigma_\varepsilon \\ M_{\varepsilon, t-1} \end{bmatrix}$ having linearly independent columns implies $\dim \varepsilon_t \leq \dim \eta_t$ makes clear that this condition is not necessary for the existence of a FREE. For example, suppose $\Sigma_\varepsilon = 0$, then a FREE must exist independently of the dimension of $\dim \varepsilon_t$ when $M_{\varepsilon, t-1} \equiv 0$.

In order to derive necessary conditions (and tighter sufficient ones) we must first solve for v_t . We begin by pre-multiplying (7.1) by $[I \quad -Q_1 \cdot \Pi V_1 D_{11}^{-1} U_1^H]$, which gives:

$$\begin{aligned} & [\Lambda_{11} \quad \Lambda_{12} - Q_1 \cdot \Pi V_1 D_{11}^{-1} U_1^H \Lambda_{22}] \begin{bmatrix} w_{1,t} \\ w_{2,t} \end{bmatrix} \\ &= [\Omega_{11} \quad \Omega_{12} - Q_1 \cdot \Pi V_1 D_{11}^{-1} U_1^H \Omega_{22}] \begin{bmatrix} w_{1,t-1} \\ w_{2,t-1} \end{bmatrix} + (Q_1 \cdot - Q_1 \cdot \Pi V_1 D_{11}^{-1} U_1^H Q_2 \cdot) (\mu + \Psi_{S_t} + \Pi \eta_t) \\ &= [\Omega_{11} \quad \Omega_{12} - Q_1 \cdot \Pi V_1 D_{11}^{-1} U_1^H \Omega_{22}] \begin{bmatrix} w_{1,t-1} \\ w_{2,t-1} \end{bmatrix} + (Q_1 \cdot - Q_1 \cdot \Pi V_1 D_{11}^{-1} U_1^H Q_2 \cdot) (\mu + \Psi_{S_t}) \\ &\quad + Q_1 \cdot \Pi V_2 (M_{\varepsilon, t-1} \varepsilon_t + M_{\zeta, t-1} \zeta_t) \end{aligned}$$

(using the unitary of U and V , and equation (7.4)).

Hence, if we stack the equation above with the solution for the transversality-violating terms, and pre-multiply by:

$$[Z_{\cdot 1}\Lambda_{11}^{-1} \quad Z_{\cdot 2} - Z_{\cdot 1}\Lambda_{11}^{-1}(\Lambda_{12} - Q_{1\cdot}\Pi V_{\cdot 1}D_{11}^{-1}U_{\cdot 1}^H\Lambda_{22})],$$

(valid as Λ_{11} is invertible by construction) we have:

$$\begin{aligned} v_t &= Z_{\cdot 1}\Lambda_{11}^{-1}[\Omega_{11}Z_{\cdot 1}^H + (\Omega_{12} - J\Omega_{22})Z_{\cdot 2}^H]v_{t-1} \\ &\quad + [Z_{\cdot 1}\Lambda_{11}^{-1}(Q_{1\cdot} - JQ_{2\cdot}) + [Z_{\cdot 2} - Z_{\cdot 1}\Lambda_{11}^{-1}(\Lambda_{12} - J\Lambda_{22})](\Lambda_{22} - \Omega_{22})^{-1}Q_{2\cdot}]\mu \\ &\quad + [Z_{\cdot 1}\Lambda_{11}^{-1}(Q_{1\cdot} - JQ_{2\cdot})\Psi + [Z_{\cdot 2} - Z_{\cdot 1}\Lambda_{11}^{-1}(\Lambda_{12} - J\Lambda_{22})]SP]s_t \\ &\quad + Z_{\cdot 1}\Lambda_{11}^{-1}Q_{1\cdot}\Pi V_{\cdot 2}V_{\cdot 2}^H\eta_t, \end{aligned}$$

where Z has been partitioned conformably with w_t and where $J := Q_{1\cdot}\Pi V_{\cdot 1}D_{11}^{-1}U_{\cdot 1}^H$.

For brevity, we rewrite this solution for v_t as:

$$v_t = T_{-1}v_{t-1} + T_{\mu} + T_s s_t + T_{\eta}\eta_t, \quad (7.5)$$

where T_{-1} , T_{μ} , T_s and T_{η} are defined by matching terms.

Let us assume then that $v_t = \begin{bmatrix} x_t \\ \mathbb{E}_t x_{t+1} \end{bmatrix}$ and $\eta_t = x_t - \mathbb{E}_{t-1}x_t$, as in the general linear expectational model we presented at the start of this appendix. Then if we define $T_{\varepsilon,t-1} := T_{\eta}V_{\cdot 2}M_{\varepsilon,t-1}$ and $T_{\zeta,t-1} := T_{\eta}V_{\cdot 2}M_{\zeta,t-1}$ and partition all the T . matrices conformably with v_t , we have:

$$\begin{aligned} \mathbb{E}_t x_{t+1} &= T_{-1,21}x_{t-1} + T_{-1,22}\mathbb{E}_{t-1}x_t + T_{\mu,2} + T_{s,2}s_t + T_{\eta,2}\eta_t \\ &= T_{-1,21}x_{t-1} + T_{-1,22}\mathbb{E}_{t-1}x_t + T_{\mu,2} + T_{s,2}s_t + T_{\varepsilon,t-1,2}\varepsilon_t + T_{\zeta,t-1,2}\zeta_t. \end{aligned} \quad (7.6)$$

When either $P = 0$, or when s_{t-1} is observed, the feasibility of this solution requires that agents can work out $(T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2})\varepsilon_t$, given knowledge of x_t, η_t and ζ_t . By taking the SVD of $(T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2})$ and $\begin{bmatrix} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix}$ it is straightforward to show that a sufficient condition for feasibility is that:

$$\ker S\Sigma_\varepsilon \cap \ker M_{\varepsilon,t-1} \subseteq \ker [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}], \quad (7.7)$$

in which case:

$$[T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}]\varepsilon_t = [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \begin{bmatrix} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix}^+ \left[\begin{bmatrix} \Omega_{22}^{-1}U_{\cdot 1}D_{11} & 0 \\ 0 & I \end{bmatrix} V^H \eta_t - \begin{bmatrix} 0 \\ M_{\zeta,t-1} \end{bmatrix} \zeta_t \right],$$

and:

$$\begin{aligned} \mathbb{E}_t x_{t+1} &= [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \begin{bmatrix} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix}^+ \begin{bmatrix} \Omega_{22}^{-1}U_{\cdot 1}D_{11} & 0 \\ 0 & I \end{bmatrix} V^H x_t + T_{-1,21}x_{t-1} + T_{s,2}Ps_{t-1} \\ &+ \left[T_{-1,22} - [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \begin{bmatrix} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix}^+ \begin{bmatrix} \Omega_{22}^{-1}U_{\cdot 1}D_{11} & 0 \\ 0 & I \end{bmatrix} V^H \right] \mathbb{E}_{t-1}x_t + T_{\mu,2} \\ &+ \left[T_{\zeta,t-1,2} - [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \begin{bmatrix} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix}^+ \begin{bmatrix} 0 \\ M_{\zeta,t-1} \end{bmatrix} \right] \zeta_t, \end{aligned}$$

which is in a ‘‘semi’’-FREE form.

In fact, when $P = 0$, we can provide a more intuitive sufficient condition, under the normalisation that $\Sigma_\varepsilon = I$. In this case, $\ker S\Sigma_\varepsilon = \ker S = \ker Q_2 \cdot \Psi$ and so for $v \in \ker S\Sigma_\varepsilon \cap \ker M_{\varepsilon,t-1}$, $\Psi v = Q_1^H Q_1 \cdot \Psi v$ and hence:

$$Q_1^H \Lambda_{11} Z_{\cdot 1}^H (T_s \Sigma_\varepsilon + T_{\varepsilon,t-1}) v = Q_1^H \Lambda_{11} Z_{\cdot 1}^H Z_{\cdot 1} \Lambda_{11}^{-1} Q_1 \cdot \left[\Psi + \Pi V \begin{bmatrix} D_{11}^{-1} U_{\cdot 1}^H \Omega_{22} S \Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix} \right] v = \Psi v$$

Hence if $v \in \ker S\Sigma_\varepsilon \cap \ker M_{\varepsilon,t-1} \cap \ker \Psi$, $Q_1^H \Lambda_{11} Z_{\cdot 1}^H (T_s \Sigma_\varepsilon + T_{\varepsilon,t-1}) v = 0$ which (from pre-multiplying by $[0 \ I] Z_{\cdot 1} \Lambda_{11}^{-1} Q_1 \cdot$) implies $(T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2})v = 0$. Thus, a sufficient condition for feasibility is that:

$$\ker Q_2 \cdot \Psi \cap \ker M_{\varepsilon,t-1} \subseteq \ker Q_2 \cdot \Psi \cap \ker M_{\varepsilon,t-1} \cap \ker \Psi = \ker \Psi \cap \ker M_{\varepsilon,t-1}.$$

Consequently, a sufficient condition for feasibility for any $M_{\varepsilon,t-1}$ is that:

$$\ker Q_2 \Psi = \ker \Psi.$$

This states that if there is some linear combination of shocks which does not appear in the transversality-violating block, then that same linear combination does not appear anywhere in the model. This reveals that it is deviations from the saddle path that enable agents to back out the values of shocks.

We now turn to the general case in which we do not assume that $P = 0$ or that s_t is observed even with a lag. Our first claim is that (7.7) is a necessary condition for the existence of a FREE. Suppose for a contradiction that (7.7) does not hold, but that:

$$\mathbb{E}_t x_{t+1} = \mathcal{R}_{t-1} x_t + \mathcal{S}_{t-1} \zeta_t + \text{other terms known at } t - 1,$$

so the expectation can be formed without knowing the value of ε_t . Since $\ker S\Sigma_\varepsilon \cap \ker M_{\varepsilon,t-1} \not\subseteq \ker (T_{s,2} + T_{\varepsilon,t-1,2})$, there must exist some $\nu \neq 0$ such that $S\Sigma_\varepsilon \nu = M_{\varepsilon,t-1} \nu = 0$, but $(T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2})\nu \neq 0$.

Then from (7.6) and the fact that ζ_t is uncorrelated with ε_t , $\text{Cov}_{t-1}(\mathcal{R}_{t-1} x_t, \nu \nu^H \varepsilon_t | s_{t-1}) = \text{Cov}_{t-1}(\mathbb{E}_t x_{t+1}, \nu \nu^H \varepsilon_t | s_{t-1}) = (T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2})\mathbb{E}_{t-1} \varepsilon_t \varepsilon_t^H \nu \nu^H = (T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,2})\nu \nu^H \neq 0$. Hence, by our assumption:

$$\begin{aligned} 0 \neq \text{Cov}_{t-1}(\mathcal{R}_{t-1} x_t, \nu \nu^H \varepsilon_t | s_{t-1}) &= \text{Cov}_{t-1}(\mathcal{R}_{t-1}(\eta_t + \mathbb{E}_{t-1} x_t), \nu \nu^H \varepsilon_t | s_{t-1}) \\ &= \mathbb{E}_{t-1} \mathcal{R}_{t-1} \eta_t \varepsilon_t^H \nu \nu^H = \mathcal{R}_{t-1} \mathbb{E}_{t-1} \left[V \begin{bmatrix} D_{11}^{-1} U_1^H \Omega_{22} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{bmatrix} \varepsilon_t + V \begin{bmatrix} 0 \\ M_{\zeta,t-1} \end{bmatrix} \zeta_t \right] \varepsilon_t^H \nu \nu^H = 0 \end{aligned}$$

(using equation (7.4)), as $S\Sigma_\varepsilon \nu = M_{\varepsilon,t-1} \nu = 0$ and ζ_t is uncorrelated with ε_t . This gives the required contradiction.

Finally, we show that (7.7) and $\ker T_{s,2} = \{0\}$ are jointly sufficient. First note that if $\ker T_{s,2} = \{0\}$, then $T_{s,2}^+ T_{s,2} = I$. Then, from substituting $\mathbb{E}_{t-1} x_t$ out of the top line of (7.6), using the definition of η_t , subtracting $T_{s,2} P T_{s,2}^+$ times the equation's lag, then using again the definition of η_t :

$$\begin{aligned} \mathbb{E}_t x_{t+1} &= [T_{-1,22} + T_{s,2} P T_{s,2}^+] x_t + [T_{-1,21} - T_{s,2} P T_{s,2}^+ T_{-1,22}] x_{t-1} - T_{s,2} P T_{s,2}^+ T_{-1,21} x_{t-2} \\ &\quad + [I - T_{s,2} P T_{s,2}^+] T_{\mu,2} + T_{s,2} \Sigma_\varepsilon \varepsilon_t + [T_{\eta,2} - T_{-1,22} - T_{s,2} P T_{s,2}^+] \eta_t \\ &\quad - T_{s,2} P T_{s,2}^+ [T_{\eta,2} - T_{-1,22}] \eta_{t-1}, \end{aligned}$$

or equivalently (again by the definition of η_t):

$$\begin{aligned} \mathbb{E}_t x_{t+1} &= T_{\eta,2} x_t + [T_{-1,21} - T_{s,2} P T_{s,2}^+ T_{\eta,2}] x_{t-1} - T_{s,2} P T_{s,2}^+ T_{-1,21} x_{t-2} + [I - T_{s,2} P T_{s,2}^+] T_{\mu,2} \\ &\quad + [T_{-1,22} + T_{s,2} P T_{s,2}^+ - T_{\eta,2}] \mathbb{E}_{t-1} x_t + T_{s,2} P T_{s,2}^+ [T_{\eta,2} - T_{-1,22}] \mathbb{E}_{t-2} x_{t-1} \\ &\quad + T_{s,2} \Sigma_\varepsilon \varepsilon_t. \end{aligned}$$

Hence, since $V_{s,2}^H \eta_t = M_{\varepsilon,t-1} \varepsilon_t + M_{\zeta,t-1} \zeta_t$:

$$\begin{aligned} \mathbb{E}_t x_{t+1} &= [T_{-1,21} - T_{s,2} P T_{s,2}^+ T_{\eta,2}] x_{t-1} - T_{s,2} P T_{s,2}^+ T_{-1,21} x_{t-2} + [I - T_{s,2} P T_{s,2}^+] T_{\mu,2} \\ &\quad + [T_{-1,22} + T_{s,2} P T_{s,2}^+ - T_{\eta,2}] \mathbb{E}_{t-1} x_t + T_{s,2} P T_{s,2}^+ [T_{\eta,2} - T_{-1,22}] \mathbb{E}_{t-2} x_{t-1} \\ &\quad + [T_{s,2} \Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \varepsilon_t + T_{\zeta,t-1,2} \zeta_t. \end{aligned}$$

By (7.7) then we have the FREE solution:

$$\begin{aligned}
 \mathbb{E}_t x_{t+1} &= [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \left[\begin{array}{c} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{array} \right]^+ \left[\begin{array}{cc} \Omega_{22}^{-1}U_{.1}D_{11} & 0 \\ 0 & I \end{array} \right] V^H x_t + [T_{-1,21} - T_{s,2}PT_{s,2}^+T_{\eta,2}]x_{t-1} \\
 &\quad - T_{s,2}PT_{s,2}^+T_{-1,21}x_{t-2} + [I - T_{s,2}PT_{s,2}^+]T_{\mu,2} \\
 &\quad + \left[\begin{array}{c} T_{-1,22} + T_{s,2}PT_{s,2}^+ - T_{\eta,2} \\ \\ \\ \end{array} \right. \\
 &\quad \quad \left. - [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \left[\begin{array}{c} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{array} \right]^+ \left[\begin{array}{cc} \Omega_{22}^{-1}U_{.1}D_{11} & 0 \\ 0 & I \end{array} \right] V^H \right] \mathbb{E}_{t-1} x_t \\
 &\quad + T_{s,2}PT_{s,2}^+[T_{\eta,2} - T_{-1,22}] \mathbb{E}_{t-2} x_{t-1} \\
 &\quad + \left[\begin{array}{c} T_{\zeta,t-1,2} - [T_{s,2}\Sigma_\varepsilon + T_{\varepsilon,t-1,2}] \left[\begin{array}{c} S\Sigma_\varepsilon \\ M_{\varepsilon,t-1} \end{array} \right]^+ \left[\begin{array}{c} 0 \\ M_{\zeta,t-1} \end{array} \right] \end{array} \right] \zeta_t,
 \end{aligned}$$

which establishes the result.

A final remark is that the condition (7.7) holds if and only if:

$$\ker S\Sigma_\varepsilon \cap \ker M_{\varepsilon,t-1} \subseteq \ker T_{s,2}\Sigma_\varepsilon$$

by the definition of $T_{\varepsilon,t-1,2}$. Under determinacy, this in turn holds if and only if $\ker S \subseteq \ker T_{s,2}$.

7.2. E-stability analysis

Following Marcet and Sargent (1989) and Evans and Honkapohja (2001), we calculate the eigenvalues of the Jacobian of the mapping from the PLM (3.1) to the actual law of motion (ALM) (2.3). This mapping takes the form:

$$T \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ b \\ c \\ d'_1 \end{bmatrix} = \frac{1}{1 - \beta a_1} \begin{bmatrix} \alpha + \rho + \beta(a_2 + (b - \rho)a_1) \\ \beta(a_3 + (b - \rho)a_2) - \alpha\rho \\ \beta(b - \rho)a_3 \\ \beta(b - \rho)b \\ (1 - \rho)\gamma + \beta c(1 + b - \rho) \\ \beta(b - \rho)d'_1 \end{bmatrix},$$

since:

$$\begin{aligned}
 (1 - \beta a_1)x_{t+1} &= (\alpha + \rho + \beta(a_2 + (b - \rho)a_1))x_t + (\beta(a_3 + (b - \rho)a_2) - \alpha\rho)x_{t-1} \\
 &+ \beta(b - \rho)a_3x_{t-2} + \beta(b - \rho)b\mathbb{E}_{t-1}x_t + ((1 - \rho)\gamma + \beta c(1 + b - \rho)) \\
 &+ \beta(b - \rho)d_1\zeta_t + \sigma\varepsilon_{t+1} + \beta d_1\zeta_{t+1}.
 \end{aligned}$$

The set of fixed points of T comprises three discrete islands, two of which are single points with $a_3 = b = d_1 = 0$ (i.e. the MSV solutions). These only exist when $\alpha\beta \leq \frac{1}{4}$. The third island is of dimension $1 + \dim \zeta_t$, capturing the degrees of freedom under indeterminacy.

If we define $\mathcal{f} := \sqrt{\max\{0, 1 - 4\alpha\beta\}}$, then the real-parts of the eigenvalues in the three cases are:

- For the two MSV solutions, indexed by $\lambda \in \left\{\frac{1 \pm \mathcal{f}}{2\beta}\right\}$ (and assuming $\alpha\beta \leq \frac{1}{4}$):

$$0, \quad \frac{\beta(1 - \rho)}{1 - \beta(\rho + \lambda)}, \quad \frac{\beta(\alpha - \rho(1 - \beta\rho))}{(1 - \beta(\rho + \lambda))^2}, \quad -\frac{\beta\rho}{1 - \beta(\rho + \lambda)}$$

- For the sunspot solution (where b is a free parameter):

$$1, \quad -\frac{b}{\rho - b}, \quad 1 - \frac{1}{\rho - b}, \quad 1 - \frac{|\rho - b| \pm (\rho - b)\mathcal{f}}{2\beta|\rho - b|(\rho - b)}.$$

By the results of Evans and Honkapohja (2001), least squares learning will not converge if any of the eigenvalues' real parts are greater than one. These are similar to, but not identical to, the conditions Evans and Honkapohja (2001) derive for the MSV PLM in their proposition 8.3, under the assumption that the shock is observable.

For convergence in the sunspot case, we at last need the following conditions to hold: $b \leq \rho$, $0 \leq \rho$, $0 \leq \alpha$, $0 < \beta$. Providing these conditions hold, the T map will not have any eigenvalues with real parts greater than one, and those eigenvalues for which the real part equals one will have zero complex parts (a further necessary condition for convergence, without this there may be stable cycles under learning). Note that these parameter

restrictions include the most economically relevant case from our motivating example of the Taylor rule, where we would expect $0 \leq \rho < 1$, $\alpha = 0$ and $\beta > 0$. However, they also includes many explosive regions (when α is large), and regions exhibiting stable cycles in which ρ is fully identified (i.e. $\alpha\beta > \frac{1}{4}$, which requires large β).

Define $\phi := [a_1 \ a_2 \ a_3 \ b \ c \ d_1]'$. The system is weakly e-stable at the solution $[\frac{1}{\beta} + (\rho - \tilde{b}) \ -\frac{\alpha+\rho}{\beta} \ \frac{\alpha\rho}{\beta} \ \tilde{b} \ -\frac{\gamma(1-\rho)}{\beta} \ \tilde{d}]'$ for fixed \tilde{b} and \tilde{d} if and only if the differential equation $\dot{\phi} = T\phi - \phi$ is locally stable at this solution, where the dot denotes a derivative with respect to “virtual-time” τ .

Defining:

$$\psi := \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \phi - \begin{bmatrix} \frac{1}{\beta} + (\rho - \tilde{b}) \\ -\frac{\alpha + \rho}{\beta} \\ \frac{\alpha\rho}{\beta} \\ \tilde{b} \\ -\frac{\gamma(1-\rho)}{\beta} \\ \tilde{d}' \end{bmatrix} = \begin{bmatrix} a_1 - \frac{1}{\beta} - (\rho - b) \\ a_2 + \frac{\alpha + \rho}{\beta} \\ a_3 - \frac{\alpha\rho}{\beta} \\ b - \tilde{b} \\ c + \frac{\gamma(1-\rho)}{\beta} \\ d_1' - \tilde{d}' \end{bmatrix},$$

we then have that:

$$\dot{\psi} = -\frac{1}{\psi_1 - \psi_4 + (\rho - \tilde{b})} \begin{bmatrix} \psi_2 + \psi_1 \left(\psi_1 + \frac{1}{\beta} + \rho \right) \\ \psi_3 + \psi_1 \left(\psi_2 - \frac{\alpha + \rho}{\beta} \right) \\ \psi_1 \left(\psi_3 + \frac{\alpha\rho}{\beta} \right) \\ \psi_1 (\psi_4 + \tilde{b}) \\ \psi_5 + \psi_1 \left(\psi_5 - \frac{\gamma(1-\rho)}{\beta} \right) \\ \psi_1 (\psi_6 + \tilde{d}') \end{bmatrix}.$$

Combining the third and fourth equations then gives that:

$$\frac{\psi_4(\tau) + \tilde{b}}{\psi_3(\tau) + \frac{\alpha\rho}{\beta}} = \frac{\psi_4(0) + \tilde{b}}{\psi_3(0) + \frac{\alpha\rho}{\beta}} = \frac{b(0)}{a_3(0)}.$$

Using this equation, we can substitute ψ_4 out of the above differential equation. We can also ignore the final equation since it is the only one containing ψ_6 , meaning that if the other components converge to something, so will ψ_6 . The resulting four-equation system has real eigenvalues components:

$$\frac{\beta}{C\alpha - \beta}, \quad \frac{\beta}{\rho(C\alpha - \beta)}, \quad \frac{1 \pm \sqrt{\max\{0, 1 - 4\alpha\beta\}}}{2\rho(C\alpha - \beta)}$$

when evaluated at the (zero) steady-state, where $C := \frac{b(0)}{a_3(0)}$. Given the necessary conditions already derived ($\tilde{b} \leq \rho$, $0 \leq \rho$, $0 \leq \alpha$ and $0 < \beta$), for these real eigenvalues components to be strictly negative, we require that $\alpha C - \beta \leq 0$. However, since we only require local convergence, we may assume that b and a_3 begin close enough to their steady state for us to have $C = \frac{\beta\tilde{b}}{\alpha\rho} + \beta\epsilon$ for some ϵ , small in magnitude. Then $\alpha C - \beta \leq 0$ if and only if $\tilde{b} \leq \rho(1 - \alpha\epsilon)$. We can always find an ϵ for which this holds (i.e. start sufficiently close to the solution) providing $\tilde{b} < \rho$ or $\alpha = 0$ and $\tilde{b} \leq \rho$.

We now turn to the second PLM, (3.2). Since the two PLMs only differ in a term that is unknown at t , period t expectations of x_{t+1} are identical under both PLMs, meaning that the T -map is just as before, but with one extra component, taking d'_0 to $\frac{\beta}{1-\beta a_1} d'_1$. Consequently, a solution is weakly (strongly) e-stable under the PLM (3.2) if and only if it is weakly (strongly) e-stable under the PLM (3.1).³⁰

³⁰ This follows from integrating the corresponding differential equation, to give $d_0(\tau) = e^{-\tau} \int_0^\tau \frac{\beta}{1-\beta a_1(t)} d_1(t) e^t dt + d_0(0) e^{-\tau}$. Hence as $\tau \rightarrow \infty$, $d_0(\tau) \rightarrow \lim_{t \rightarrow \infty} \frac{\beta}{1-\beta a_1(t)} d_1(t)$.

Chapter 3: Efficient simulation of DSGE models with inequality constraints

Tom Holden¹, Balliol College, University of Oxford

Michael Paetz², Univeristy of Hamburg, Department of Economics

Abstract: *This paper presents a fast, simple and intuitive algorithm for simulation of dynamic stochastic general equilibrium models with inequality constraints. The algorithm handles both the computation of impulse responses, and stochastic simulation, and can deal with arbitrarily many bounded variables. Furthermore, the algorithm is able to capture the precautionary motive associated with the risk of hitting such a bound. To illustrate the usefulness and efficiency of this algorithm we provide a variety of applications including to models incorporating a zero lower bound (ZLB) on nominal interest rates. Our procedure is much faster than comparable methods and can readily handle large models. We therefore expect this algorithm to be useful in a wide variety of applications.*

Keywords: *Inequality constraints, Zero lower bound, Precautionary motives, Two-country models*

JEL Classification: *C63, E32, E43, E52*

¹ Contact address: School of Economics, FBEL, University of Surrey, Guildford, GU2 7XH, England
Contact email: thomas.holden@gmail.com Web-site: <http://www.tholden.org/>

² Contact address: Fakultät für Wirtschafts- und Sozialwissenschaften, Universität Hamburg, Von-Melle-Park 9, 20146 Hamburg, Deutschland

Contact email: michael.paetz@wiso.uni-hamburg.de

The authors would like to acknowledge helpful comments on earlier drafts of the paper from Michael Funke, William T. Gavin, Matteo Iacoviello, Michel Juillard, Paul Levine, Jesper Lindé, Antonio Mele, Søren Ravn and Simon Wren-Lewis, along with those from other participants in the 8th Dynare Conference.

1. Introduction

New Keynesian (NK) Dynamic Stochastic General Equilibrium (DSGE) models are today's standard framework for analysing central bank policies.³ The nominal and real rigidities in these models mean central banks may improve welfare through monetary policy. Traditionally, DSGE models are log-linearised, which results in both computational and analytic tractability. However, the simplicity of this approach also neglects important non-linearities, not-least inequality constraints, of which the zero lower bound (ZLB) on interest rates is the most prominent example. In this paper, we present an efficient algorithm for simulating DSGE models subject to arbitrarily many inequality constraints, at arbitrary accuracy, at least away from the bound.

Macroeconomic analysis had ignored the zero lower bound almost completely before the experience of Japan in the 1990s, since the bulk of macroeconomists believed the constraint would bind only for a short time span (if at all). Under this presumption, the effects would be negligible, and so ignoring the bound seemed to be a reasonable simplification. However, the interest rates in Japan during the 1990s, as well as those in the US over the last few years, disabused researchers and policymakers of this popular fallacy.

In the aftermath of the crisis, the transmission of monetary policy under the ZLB became an important matter for central banks and academia. To cope with the bound, researchers either solve such non-linear models using global approximation methods (which come at a dramatic increase in computational costs, and scale exceptionally poorly) or use deterministic setups. This paper provides a fast, simple and intuitive algorithm to deal with inequality constraints in perturbation approximations to DSGE models, which correctly captures the precautionary motives associated with such bounds. The code is designed to work with Dynare (Adjemian et al. 2011), so incorporating it into existing models is trivial.

³ See Clarida et al. (1999) for an early literature review on NK models.

The method is not only useful for deriving impulse response functions (IRFs), but can also be used for stochastic simulations, opening up the possibility of particle filter based estimation of models with inequality constraints. The method endogenously determines when the constraint will bind, and can handle constraints that may bind in multiple disjoint runs, or that may not begin to bind until long after the initial impulse. The general idea is to introduce “shadow price shocks”, which hit the bounded variables every time the constraint is violated, and “push” these variables back to zero. To ensure the solution is consistent with rational expectations, these shocks are expected by agents in advance, so they may be thought of as a kind of endogenous news shock.

This algorithm is not solely useful for modelling a ZLB on interest rates, but can be used for any model including constrained variables. Holden (2010), for example, uses it to constrain invention rates to be positive in a model of endogenous growth. In addition, Funke and Paetz (2012) use this technique to evaluate threshold loan-to-value policies in Hong Kong, where policymakers decrease the loan-to-value ratio, when property price inflation exceeds a certain value. In Chen et al. (2012), the same method is used to model the People’s Bank of China’s interest rate corridor on retail lending and deposit rates.

The rest of the paper is organized as follows. In section 2, the algorithm is described and related to the existing literature. Section 3 assesses the accuracy of the procedure in a variety of models for which high accuracy solutions are available, and section 4 goes on to provide some sample applications to larger models. The final section concludes.

2. The numerical method

2.1. The existing literature

Due to the recent experience of the US and Europe, the literature on (stochastic) simulation of models with a ZLB has grown rapidly in the past few years. The most important contributions include Eggertsson and Woodford (2003), Erceg and Lindé (2010), Braun and

Körber (2011), Christiano et al. (2011) and Fernández-Villaverde et al. (2012). In what follows, we highlight the similarities and differences between the approaches employed in these papers and the method presented in this paper.⁴

The first generation of papers used variants of the method proposed by Eggertsson and Woodford (2003), in an appendix. This relied on a piecewise linear approximation to the model, with the model being driven by a two-state Markov chain with an absorbing state. Once the ZLB is hit, there is a positive probability in each period that the discount factor jumps to its long run value, at which point the ZLB will never be hit again. Obviously, this is a highly restrictive assumption. A version of this algorithm without the restriction has been proposed by Jung et al. (2005), and implemented in full generality in Dynare by Guerrieri and Iacoviello (2012). Nonetheless, the algorithm still relies on a linear approximation, which the results of Braun et al. (2012) suggest may lead to unreliable conclusions in the presence of the ZLB.

The next generation of papers used nonlinear perfect foresight solvers. These include Coenen et al. (2004), Braun and Körber (2011) and the “extended path” method of Adjemian and Juillard (2011). These solve the model’s nonlinear equations, under the assumption that eventually (e.g. after 100 periods) the model is guaranteed to have returned to steady state. Such methods fully capture the nonlinearities of the model, but because they solve under perfect foresight, they omit any “precautionary motives” including those that arise from the risk of hitting the ZLB. Furthermore, since the model has to have returned to steady state up to machine precision by the final period considered, they require a very large number of nonlinear equations to be solved. This means they tend to both be prohibitively slow, and unstable, with the algorithm frequently failing to find a solution to the equations.

⁴ See the introduction in Braun et al. (2012) for a survey on models including a zero lower bound. An early analysis of the zero lower bound in a deterministic model can also be found in Fuhrer and Madigan (1997).

A third strand of the literature considers global approximations to models containing inequality constraints, with Fernández-Villaverde et al. (2012) doing this for a small scale NK model, using the Smolyak collocation method of Krueger et al. (2011). Global methods successfully capture both the model's nonlinearities and precautionary motives; however, they are subject to a curse of dimensionality that renders them infeasible in the medium scale NK models we usually consider. While global methods that avoid the curse of dimensionality have been developed by Maliar et al. (2011), these rely on an endogenous grid constructed from the model's ergodic set, which is likely to lead to low accuracy at the ZLB if this bound is only hit occasionally.

Our method represents a compromise between the accuracy of global methods, and the speed and scalability of linear ones, much like standard high order perturbation approximations. The paper that is probably most closely related to our work is Erceg and Lindé (2010), which we were not aware of until after the completion of the first version of our algorithm in Holden (2010). The authors rely on techniques, explained in an unpublished mimeo of James Hebden, Jesper Lindé and Lars Svensson, which, like our algorithm, are based on the idea of adding shocks to the bounded variable. Since we have not seen this mimeo, we are unable to relate our work to this algorithm, but we are confident that our method is novel in several respects. Firstly, it is designed to take advantage of existing algorithms both for simulating DSGE models (e.g. those of Dynare), and algorithms for quadratic programming, leading to its high speed. Secondly, it is generalized to permit any number of constrained variables. Thirdly, it is extended for use in stochastic simulations, permitting us to derive average IRFs, and opening up the possibility of estimating bounded models. Finally, it is generalized to perturbation approximations of arbitrary order, which is what enables the algorithm to capture precautionary incentives, thus improving on the accuracy of nonlinear perfect foresight algorithms.

2.2. Our basic IRF algorithm with a single bound

Suppose we have a rational expectations model in the variables $x_{1,t}, \dots, x_{n,t}$, and we are interested in the response to the shock, ϵ_t . Initially, we will suppose further that all of the model's equations are linear, except one that takes the form:

$$x_{1,t} = \max\{0, \mu_1 + \phi'_{-1}x_{t-1} + \phi'_0x_t + \phi'_1\mathbb{E}_t x_{t+1} - (\phi'_{-1} + \phi'_0 + \phi'_1)\mu\}, \quad (2.1)$$

where x_t is the vector $[x_{1,t}, x_{2,t}, \dots, x_{n,t}]'$, $\mu = [\mu_1, \mu_2, \dots, \mu_n]'$ is a vector stacking the variables' steady state values and $\mu_1 > 0$. We can transform any linear model with a bound into this form through the addition of appropriate auxiliary variables.⁵

Now, a shock that drives $x_{1,t}$ to 0 for some number of periods is like a combination of the original shock, and a news shock stating $x_{1,t}$ will be higher than expected for some duration. We call these news shocks "shadow price shocks" as in a model with bounded assets, they represent the Lagrange multiplier on the constraint. The key to the simplicity of our algorithm is the fact that in linear models, the IRF to a linear combination of shocks is equal to the same linear combination of each shock's IRF.⁶

Let us start by defining v_j to be the column vector containing the relative impulse response of variable x_j to the shock, ignoring the ZLB. Let T^* be the number of periods after which the constraint is no longer expected to bind. Note that this will in general be much smaller than the time it takes to return to steady state. We assume that the IRF vectors are of length T , where $T \geq T^*$.

⁵ So, if the bounded equation stated that $x_t = \max\{\tilde{x}_t, y_t\}$ (where $\tilde{\cdot}$ denotes unconstrained variables), with $x_t > y_t$ in steady state, we would add an auxiliary variable defined as $\tilde{x}_t - y_t$, noting that $x_t = y_t + \max\{0, \tilde{x}_t - y_t\}$. The models of Funke and Paetz (2012) and Chen et al. (2012) include variables that are bounded at a positive value, for example. When $x_t = y_t$ in steady state, we instead add the variable $y_t - \tilde{x}_t$, and note that $x_t = \tilde{x}_t + \max\{0, y_t - \tilde{x}_t\}$. If the model is in levels, rather than in logs, it may be preferable to define auxiliary variables as ratios rather than differences. In this case, rather than adding shadow shocks, we must multiply by their exponentials.

⁶ The algorithm presented here was first explained in the appendix of Holden (2010). The various extensions we describe were not covered in that paper, however.

The first step of our algorithm is to replace equation (2.1) of the model with:

$$x_{1,t} = \mu_1 + \phi'_{-1}x_{t-1} + \phi'_0x_t + \phi'_1\mathbb{E}_t x_{t+1} - (\phi'_{-1} + \phi'_0 + \phi'_1)\mu + \sum_{s=0}^{T^*-1} \epsilon_{s,t-s}^{\text{SP}} \quad (2.2)$$

where $\epsilon_{s,t}^{\text{SP}}$ is a newly introduced shadow price shock at horizon s , for $s = 0, \dots, T^* - 1$. Since $\epsilon_{s,t}^{\text{SP}}$ is known at t , but does not “hit” until $t + s$, it will function as a news shock, as required. In an impulse response exercise, all news arrives in period 0, so we want to find values for these shadow shocks such that $\epsilon_{s,t}^{\text{SP}} = 0$ except perhaps when $t = 0$.

In order to exploit the linearity of IRFs in linear models, we also store the impulse responses to these new shadow price shocks. In particular, we let $m_{j,s}$ be a column vector containing the relative impulse response of variable $x_{j,t}$ to the shock $\epsilon_{s,t}^{\text{SP}}$. We then horizontally stack these vectors into the matrix $M_j = [m_{j,0}, m_{j,1}, \dots, m_{j,T^*-1}]$. With this matrix, we can now calculate the impulse responses to a simultaneous shock to ϵ_t of magnitude 1 and to $\epsilon_{s,t}^{\text{SP}}$ of magnitude α_s (for each $s = 0, \dots, T^* - 1$), for an arbitrary $\alpha = [\alpha_0, \alpha_1, \dots, \alpha_{T^*-1}]'$, without any further simulation of the model. In particular, the IRF to this combination of shocks for variable x_j will be equal to $\mu_j + v_j + M_j\alpha$.

Our task then is just to find a value for α that is consistent with the model and with rational expectations. For the constraint to be satisfied, it certainly has to be the case that $\mu_1 + v_1 + M_1\alpha \geq 0_T$ (where 0_T is a length T vector of 0s). Additionally, we require that $\alpha \geq 0_{T^*}$, as shadow shocks must increase the (lower) bounded variable.⁷ Finally, a complementary slackness type condition must hold: the shadow shock at horizon s can only be non-zero if the bound binds at that horizon.

⁷ We assume throughout that the diagonal of the M_1 matrix is strictly positive, so a shadow shock at horizon s increases the bounded variable at horizon s . When this does not hold, (as it may not, for example, at long horizons in sufficiently rich NK models), the sign with which the shadow shock enter equation (2.2) must be flipped.

We can express this as:

$$\alpha'[\mu^* + v^* + M^* \alpha] = 0_{T^*}, \quad (2.3)$$

where $\mu^* := \mu_1 1_{T^*}$, v^* is the first T^* elements of v_1 and M^* is the upper $T^* \times T^*$ sub-matrix of M_1 . To solve for an α that satisfies these constraints, we run the following quadratic programming problem:

$$\alpha^* := \arg \min_{\substack{\alpha \geq 0_{T^*} \\ \mu^* + v^* + M^* \alpha \geq 0_{T^*}}} \left[\alpha'(\mu^* + v^*) + \frac{1}{2} \alpha'(M^* + M^{*'})\alpha \right], \quad (2.4)$$

where the solution is considered admissible if the minimand is 0 at the optimum (i.e. equation (2.3) is satisfied).⁸ Since there are well-established, fast, robust algorithms for quadratic programming, this is then a straightforward problem.⁹

The standard properties of quadratic programming problems imply that a sufficient condition for the existence of a unique solution to (2.4) is that $M^* + M^{*'}$ is positive definite. In our experience, this is satisfied in only the simplest models. When $M^* + M^{*'}$ is not positive definite, we cannot rule out the existence of multiple solutions. In these cases, which solution is returned will depend on the precise properties of the quadratic programming algorithm used. However, for most models the construction of our problem will lead the algorithm to select the solution in which the components of α are as small as possible, which will also tend to minimise the amount of time the constraint binds. If desired, explicit guarantees on the solution selected may be enforced via homotopy methods,¹⁰ though this will increase the time cost of our algorithm.

⁸ The constraint $\mu^* + v^* + M^* \alpha \geq 0_{T^*}$ may also be replaced with $\mu_1 + v_1 + M_1 \alpha \geq 0_T$ to check there are no bound violations after T^* .

⁹ In MATLAB, these are provided by the “quadprog” command.

¹⁰ To guarantee selecting the equilibrium in which $\|\alpha\|_2$ is minimal, we replace $M^* + M^{*'}$ by $M^* + M^{*' + \lambda I}$, where $\lambda \rightarrow 0$ is the homotopy parameter. To guarantee selecting the equilibrium in which $\|\alpha\|_1$ is minimal, we replace $\mu^* + v^*$ by $\mu^* + v^* + \lambda 1_{T^*}$, where $\lambda \rightarrow 0$ is the homotopy parameter.

It is also possible that there will be no admissible solution to (2.4), at least for sufficiently large shocks. This will happen if there are “complementarities” between constraints: e.g. hitting one constraint increases the chance of hitting another constraint *and vice versa*. Obviously, this is much more likely when there are multiple bounds, rather than merely multiple horizons. A necessary condition for the existence of an admissible solution to (2.4) for arbitrarily large shocks is that there exists some $\alpha \geq 0_{T^*}$ such that $M^* \alpha \geq 1_{T^*}$. In simple models the anticipation effects of hitting the bound in future are weak, so this condition will be satisfied, but in medium scale models, we will generally not be able to provide such a guarantee.

2.3. Dealing with multiple bounds

The algorithm previously described may be readily generalised to cases with multiple bounds. Suppose that in the set-up above, rather than just $x_{1,t}$ being bounded, each of the variables $x_{1,t}, x_{2,t}, \dots, x_{n^*,t}$ is bounded, with a corresponding equation taking the form of (2.1). Much as before, we add shadow shocks to each of these equations (giving a total of n^*T^* extra shocks), and we horizontally concatenate the impulse responses of variable $x_{j,t}$ to each of the shadow shocks in the equation for $x_{l,t}$ into the matrix $M_{j,l}$. We then define $M_{j,l}^*$ to be the upper $T^* \times T^*$ sub-matrix of $M_{j,l}$, and M^* to be the $n^*T^* \times n^*T^*$ block-matrix with $(j,l)^{\text{th}}$ block $M_{j,l}^*$ for $j, l \in \{1, \dots, n^*\}$. Likewise, we define $\mu_l^* := \mu_l 1_{T^*}$, v_l^* to be the first T^* elements of v_l , and μ^* and v^* to be the length n^*T^* block vectors with l^{th} block μ_l^* and v_l^* , respectively. With these (re-)definitions, an admissible solution to (2.4) again gives us the required combination of shadow shocks to enforce all bounds, and the uniqueness and existence conditions are identical as well.

2.4. Stochastic simulation

The algorithm of the previous sections may be readily extended to the stochastic simulation of bounded, linear models.¹¹ As before we begin by adding shadow price shocks to the equations defining bounded variables. Now suppose we have simulated up to period $t - 1$. In a linear model, agents' expectations at t of the state of the economy at $t + 1, t + 2, \dots$ are the same as they would be were the variance of all shocks equal to 0 from $t + 1$ onwards. This will not be exactly true in a model with bounds, since the bounds will tend to increase the means of lower bounded variables. However, since (log-)linearisation has already removed any effects of uncertainty on the mean, if we solve assuming there are no shocks after period t , our approximation error is likely to be of the same order as that of a linearised model without constraints,¹² at least providing the precautionary incentives stemming from the risk of hitting the bound are fairly weak.

Thus, much as in the IRF case, we first simulate the model to find the path by which it would return to the steady-state, in the absence of bounds, and with no shocks arriving after period t . If the constraints are not violated along this path, then our simulated value for period t is fine, and we may move on to period $t + 1$. Otherwise, shadow shocks must be added. The algorithm for doing this is identical to that described above for IRFs, except that the simulated return paths of the economy's variables take the role of $\mu^* + v^*$ above. (We again use the $M_{j,l}^*$ matrices formed from the impulse responses to shadow shocks.) The found solution to the quadratic programming problem gives a valid, new value for variables at t , enabling us to go on to the next period. Note that it is now no longer the case that the "news" contained in shadow shocks is guaranteed to come true, since other shocks may arrive in the meantime pushing us away from the bounds. Consequently, α_0 no longer represents the found value of today's shadow price shock. Rather, it is equal to the

¹¹ The algorithm described here was first publicly described by Tom Holden at <http://bit.ly/10nAHf>.

¹² An identical approximation is made in the non-linear case by Adjemian and Juillard (2011).

cumulated history of shadow price shocks that hit in period 0 (i.e. $\epsilon_{0,0}^{\text{SP}} + \epsilon_{1,-1}^{\text{SP}} + \dots + \epsilon_{T^*-1,-T^*+1}^{\text{SP}}$).

2.5. Average IRFs

Using our simulation algorithm we can go on to produce average IRFs, i.e. IRFs that give a measure of the average response of the model to a one standard deviation shock, rather than a measure of the response in steady-state. (The two measures agree in the absence of bounds.) To do this, we run a stochastic simulation of the model, and then rerun the same simulation with the same shocks in all periods except one, in which we add 1 to the shock of interest. The difference between these two simulations gives one sample IRF, and the average of many such sample IRFs gives us our average IRF.¹³

2.6. Estimating Bounded Models

Our method for simulating models incorporating a zero lower bound naturally leads to an algorithm for particle filter based estimation of them. Indeed, since the observation equations are still linear in the state, and the transition equations are near linear, this is likely to be far more efficient than particle filter estimation of second order approximations to standard DSGE models (as in Fernández-Villaverde and Rubio-Ramírez (2010)). Indeed, since our solution method readily delivers last period's expectation of today's shadow price shock, we can write down a close approximation to the model for which the transition equations are linear in today's shock. For this approximated model, the optimal particle filter "proposal distribution" may be derived analytically, giving us a near optimal proposal distribution for the actual model, and enabling us to get high accuracy out of a small number of particles. We intend to assess the practical performance of this method in future work.

¹³ This is the algorithm used by Dynare for constructing IRFs in non-linear models. See <http://www.dynare.org/DynareWiki/IrFs>.

2.7. Generalisation to approximations of arbitrary order

The recent work of Braun et al. (2012) brought to light some serious problems with log-linearised solutions to models with a ZLB. They illustrate that the log-linearised equilibrium conditions can be misleading with respect to the existence and uniqueness of equilibrium, and may lead to “wrong” dynamics under the ZLB. For example, they show that in a simple NK model with Rotemberg (1996) quadratic price adjustment cost, the “paradox of toil”¹⁴ disappears in the fully nonlinear model, at least when solved under perfect foresight. The authors pinpoint the resource cost of price adjustment, which is zero in a linearised model, as being the key to this discrepancy. The paper shows that these costs work as automatic stabilizers that reduce the variation in marginal costs and inflation, and decrease the government spending multiplier at the ZLB.

In fact, even fully nonlinear perfect foresight solutions may be misleading at the ZLB. For example, Adjemian and Juillard (2011) evaluate the accuracy of the (fully non-linear perfect foresight) extended path approach in a small-scale DSGE model and show that the accuracy drops significantly when the ZLB is hit regularly. Furthermore, using global methods, Fernández-Villaverde et al. (2012) find that when the interest rate stays at the ZLB for a prolonged time period, the government spending multiplier does indeed become large, contrary to the claims of Braun et al. (2012), again suggesting that the assumption of perfect foresight may itself be a source of substantial inaccuracy in the presence of a ZLB. Our algorithm provides an answer to these worries, as it may be readily generalised to “pruned” perturbation approximations (Kim et al. 2008) of arbitrary order, and these approximations may be so constructed as to capture the precautionary motive arising from the risk of hitting the ZLB in future.

¹⁴ A fall in employment after a labour tax cut at the ZLB. See Eggertsson and Woodford (2003).

The key to ensuring our algorithm works at higher orders, is that in a d^{th} order pruned perturbation approximation, shocks of the form ϵ_t^d only enter linearly,¹⁵ hence, if we use shocks of the form $(\epsilon_{s,t}^{\text{SP}})^d$ as shadow shocks, then our algorithm will work much as before, with the expected path implied by the pruned perturbation approximation taking the place of the expected path under perfect foresight. Using shocks of the highest available order as shadow shocks is also consistent with perturbation approximation theory, since with Gaussian shocks the probability of hitting the ZLB is $o(\sigma^d)$ for any $d \in \mathbb{N}$, where σ is the perturbation parameter controlling the variance of shocks.

To capture the precautionary motive arising from the risk of hitting the ZLB in future we require d to be even, since in that case $(\epsilon_{s,t}^{\text{SP}})^d$ is positive in expectation.¹⁶ This enables us to capture the effects of the ZLB on each series' mean, removing the deficiency in our stochastic simulation algorithm previously noted. To do this requires us to first solve a fixed-point problem to ensure that the variance of $\epsilon_{s,t}^{\text{SP}}$ used in constructing the perturbation approximation actually agrees with the variance that is implied by the simulated mean values of α . In practice, the fixed-point problem is solved by standard numerical optimisation algorithms, at low accuracies since each residual evaluation requires the computation of simulated moments. We also found it helpful to treat the non-stochastic steady-state inflation target as an additional parameter to be optimised, to hold the mean level of inflation constant.

It is also possible to approximate around the risky steady state, following Juillard (2011), within our algorithm. This enables the model's responses to reflect better the differences in all equations' slope close to the ZLB.

¹⁵ See equation (36) of (Kim et al. 2008).

¹⁶ Our algorithm will not always generate positive values for $(\epsilon_{s,t}^{\text{SP}})^d$ as news may not be realised. However, since $\epsilon_{s,t}^{\text{SP}}$ only enters into the transition equations when raised to the power of d , this will not result in complex numbers entering into the simulated paths.

At sufficiently high (even) orders of perturbation approximation, our algorithm (using either the non-stochastic steady state or the risky steady state) must beat any perfect foresight method in terms of accuracy, since it captures the effects of uncertainty on the variables' means. Indeed, we show in the next section that even at order 2 our algorithm beats the extended path in standard models. Our algorithm is likely to be particularly useful for the analysis of "paradox of toil"-type effects, since it enables the analysis of the effects of the ZLB even in second order approximations to large models.

3. Accuracy

3.1. A borrowing constraints model

We begin with a very simple model taken from Guerrieri and Iacoviello (2012).

An individual's income follows the process $\log Y_t = \rho \log Y_{t-1} + \sigma \sqrt{1 - \rho^2} \varepsilon_t$, where $\varepsilon_t \sim \text{NIID}(0,1)$, $\sigma = 0.03$ and $\rho = 0.9$. This income may be used for consumption C_t or saving, and it also acts as collateral. There is a collateral constraint on total borrowing, B_t , that states that $B_t \leq MY_t$, with $M = 2$. Consumers maximise the utility function:

$$U = \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \log C_{t+s},$$

with $\beta = 0.94$, subject to the collateral constraint and subject to the budget constraint $C_t = Y_t + B_t - RB_{t-1}$, with $R = 1.05$. In normal times, the collateral constraint binds, so in place of the standard Euler equation we use the equations:

$$A_t = \max \left\{ 0, \frac{1}{(1+M)Y_t - RB_{t-1}} - \mathbb{E}_t \frac{\beta R}{C_{t+1}} \right\}$$

$$\frac{1}{C_t} = \mathbb{E}_t \frac{\beta R}{C_{t+1}} + A_t$$

where A_t is an auxiliary variable.

	$\Pr\left(\frac{B_t}{Y_t} > 1.98\right)$	$\mathbb{E} \log C_t$	$\sqrt{\text{Var} \log C_t}$	$\text{Corr}(\log C_t, \log Y_t)$
Linear ¹⁷	100%	-0.1054	0.0480	0.75
Log-linear ¹⁸	100%	-0.1043	0.0475	0.77
Piecewise linear ¹⁷	88%	-0.1051	0.0426	0.82
Extended path	89%	-0.1045	0.0431	0.84
Our algorithm order 1 ¹⁸	88%	-0.1040	0.0427	0.85
Our algorithm order 2 ¹⁸	85%	-0.1046	0.0420	0.86
Value function ¹⁷	79%	-0.1045	0.0403	0.86

Table 1: An accuracy comparison in the simple borrowing constraints model of Guerrieri and Iacoviello (2012).

In Table 1 above we present a comparison of the results of our algorithms with the piecewise linear algorithm of Guerrieri and Iacoviello (2012) and the extended path method of Adjemian and Juillard (2011). All values except those taken from Guerrieri and Iacoviello (2012) were generated from a run of 10000 periods. The shaded cells with bold text show the results coming closest to the value function iteration results, and the shaded cells without bold text show the next closest values. Our second order algorithm comes closest to the value function iteration results along three out of the four metrics considered, and is a runner up in that fourth case. Indeed, for this model, even our first order algorithm beats the extended path algorithm in three cases out of four. However, it is only really the second order version that can come close to matching the percentage of time in which the constraint binds, by virtue of taking into account the incentive to save to avoid it.

3.2. An NK model

We now turn to the NK model of Fernández-Villaverde et al. (2012). This is a standard nonlinear NK model with sticky prices, labour choice, flexible wages, Taylor rule monetary policy and a stochastic share of government spending in output. The equilibrium conditions for this model are given in the first appendix, section 7.1, and we calibrate to the same standard values as Fernández-Villaverde et al. (2012).

¹⁷ Taken from Guerrieri and Iacoviello (2012).

¹⁸ With Y_t and C_t simulated in logs.

To assess our accuracy, we treat the solution of Fernández-Villaverde et al. (2012) as the “truth”, and measure the deviations between their simulated paths¹⁹ and the paths generated by our algorithms (with all variables in logs) and the extended path method of Adjemian and Juillard (2011). All paths were length 30000 periods. We report the norms of these errors for consumption, nominal interest rates and inflation in Tables 2, 3 and 4 respectively. (Cells are formatted as before.)

	1 norm	2 norm	∞ norm
Log-linear	364.0	1.011	0.03816
Extended path	311.0	0.7575	0.01792
Our algorithm order 1	336.7	0.8120	0.03593
Our algorithm order 2	276.1	0.6924	0.01780

Table 2: Norms of the approximation errors in log consumption in the NK model of Fernández-Villaverde et al. (2012).

	1 norm	2 norm	∞ norm
Capped log-linear ²⁰	314.4	0.8344	0.01096
Extended path	311.6	0.8240	0.01064
Our algorithm order 1	314.4	0.8344	0.01096
Our algorithm order 2	313.8	0.8144	0.01050

Table 3: Norms of the approximation errors in log gross nominal interest rates in the NK model of Fernández-Villaverde et al. (2012).

	1 norm	2 norm	∞ norm
Log-linear	217.7	0.6021	0.01576
Extended path	206.5	0.5555	0.01387
Our algorithm order 1	212.8	0.5695	0.01381
Our algorithm order 2	208.2	0.5621	0.009203

Table 4: Norms of the approximation errors in log inflation in the NK model of Fernández-Villaverde et al. (2012).

At second order, our algorithm beats the extended path algorithm for consumption whichever norm is used. This is unsurprising since consumption is sensitive to risk. For the nominal interest rate, it beats the extended path method with respect to the 2 norm or the ∞ norm, but not with respect to the 1 norm. This implies the extended path algorithm is

¹⁹ We would like to thank Grey Gordon for providing us with this data.

²⁰ We replace the interest rate generated by the log-linear simulation with the maximum of 0 and the generated value.

capable of closely following interest rates a lot of the time, but occasionally it is a long way off, perhaps because of its difficulties in tracking consumption. Finally, for the inflation rate, the extended path method is more accurate with respect to all norms except the ∞ norm. It is also only with respect to the ∞ norm of the inflation error that our first order algorithm is capable of beating the extended path one.

The evidence then is a little mixed here, with six instances in which our second order algorithm wins, and three instances in which the extended path one does. However, whereas the extended path algorithm took 7 hours and 25 minutes, ours completed in only 2 hours and 34 minutes on an identical system, with the vast majority of that time in the “fixed cost” stage of finding the correct variance for the shadow shocks. In summary, then, we find that our second order algorithm is both marginally more accurate, and significantly faster, and so it seems right to conclude that our second order algorithm is the natural choice for NK models.

4. Sample applications

In order to illustrate the usefulness of the algorithm provided in section 2, we apply it to two popular linear DSGE models. The first is the stylised two-country setting of Clarida et al. (2002), which we choose in order to show that our method can easily handle multiple constraints. The second is the estimated Smets and Wouters (2003) model of the Euro area, which acts as an illustration of the speed of our algorithm, even for large models. In future versions of this paper we will also analyse second order approximations to these models.

4.1. The ZLB in the two-country model of Clarida et al. (2002)

The framework used in the following is a completely symmetric version of the Clarida et al. (2002) model with sticky prices in both countries, and perfect risk-sharing. We add domestic and foreign preference shocks as exogenous drivers, and introduce a fraction of backward-looking price-setters as in Galí and Gertler (1999). The model is described in full in the

appendix, section 7.2, along with its calibration. This calibration is standard with the exception of the elasticity of intertemporal substitution which takes a high value (i.e. a low value for risk aversion) in order to generate strong co-movement across countries. As a result, we do not claim that the simulations provide a realistic story of the transmission of shocks during the recent (nor any other) financial crises; we provide these simulations purely for illustrative reasons.

In Figure 1, we show the IRFs of the benchmark model, ignoring the ZLB (solid line), and the model imposing the ZLB via our algorithm (dashed line), to a negative domestic preference shock of magnitude 0.65. We choose such a strong shock to ensure that both countries' ZLBs are hit.




Figure 1: Perfect foresight IRFs to a domestic demand shock in the two-country model of Clarida et al. (2002). (Dashed line imposes the ZLB, solid line does not.)²¹

The fall in domestic demand induces output in both countries to fall. This leads to falling inflation in both countries and hence the central banks decrease the interest rates to boost demand. However, in the presence of the ZLB, the central banks are unable to decrease the interest rate strongly enough to generate a negative real interest rate, and consequently the recession is even more severe. This is further amplified by even larger downturns of both inflation rates. Since the foreign nominal interest rate hits the ZLB, a strong foreign recession is generated, despite the model's otherwise feeble transmission mechanism.

Producing Figure 1 took roughly a half second on a standard desktop PC, illustrating just how low is the computation cost of imposing the ZLB.

In Figure 2, we go on to simulate this model (with the standard deviation of both shocks set equal to 0.2), for 220 quarters, dropping the first 100, a process that took less than two seconds. The shaded areas of the figure show periods when both bounds are hit simultaneously, illustrating that recessions become substantially more severe in these situations.

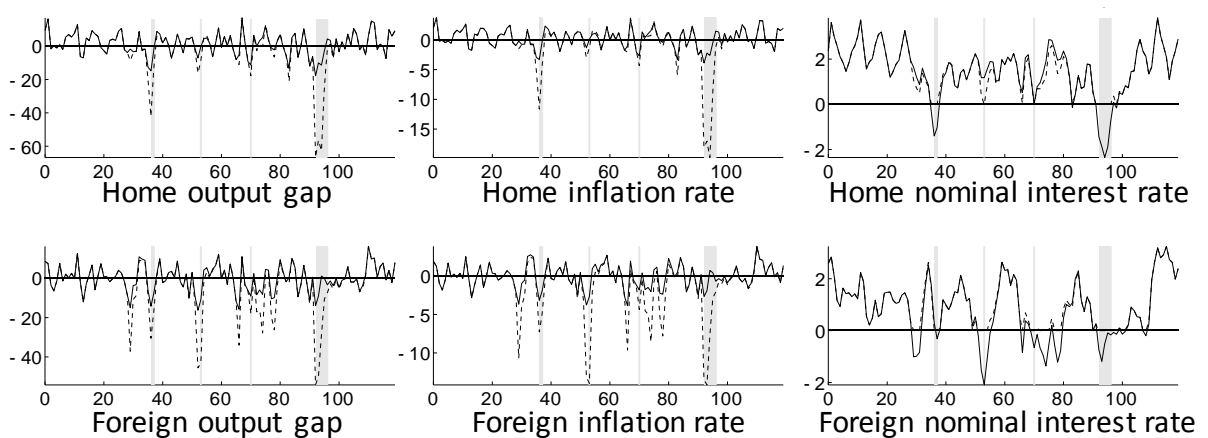


Figure 2: Simulations from the two-country model of Clarida et al. (2002).
(Dashed line imposes the ZLB, solid line does not.)²¹

Finally, in Table 5 we show the simulated moments of both the benchmark model (ignoring the ZLBs) and the constrained one, evaluated from a sample of 50,000 periods. Using our algorithm this took only 7 minutes and 24 seconds, whereas the extended path algorithm took 37 minutes and 19 seconds to produce identical results, running on the same machine.²²

²¹ Output gaps and inflation are measured in percentage deviations from equilibrium, and interest rates are measured in percent.

²² The relative speed of our algorithm becomes even more apparent if we decrease the standard deviations of both shocks to 0.1, as then our algorithm needs only 4m36s, while the extended path method still takes 31m52s.

	Home country			Foreign country		
	Output	Inflation	Nom. int.	Output	Inflation	Nom. int.
Benchmark mean	-0.07	-0.01	0.97	0.01	0.00	1.01
Constrained mean	-3.75	-1.17	1.12	-3.55	-1.10	1.14
Benchmark std. dev.	6.90	1.58	1.33	6.88	1.57	1.31
Constrained std. dev.	14.80	4.13	1.06	14.54	4.03	1.06

Table 5: Moments from the two-country model of Clarida et al. (2002).

Table 5 makes clear the magnitude of the increase in the volatility of both output gaps and inflation when we impose the ZLB on interest rates, as well as the large recessionary impact.²³ We stress again that we do not claim these values to be realistic. Rather, we presented the previous graphs and Table 5 to make clear the importance of the ZLB, and underline the tiny cost of imposing it using the method proposed in this paper.

Having shown the workings of the algorithm in a very stylized example, we now turn to a more realistic application, using the seminal Smets and Wouters (2003) model of the Euro Area.

4.2. The ZLB in the Smets and Wouters (2003) model

Modern macroeconomic models are getting progressively larger, so it is important that our algorithm can handle big models, of which the Smets and Wouters (2003) model is the canonical example. It also enables us to get a more realistic impression of the importance of the ZLB for the analysis of DSGE models.

We briefly recap the features of the model, but refer the reader to the paper for further details. It features sticky prices and wages, capital adjustment costs, variable capital utilisation and habit formation in consumption. The stochastic side of the model consists of ten exogenous shocks: six persistent shocks (technology, investment, preferences, labor supply, government spending and an inflation objective shock), all modelled as standard

²³ Obviously, the assumption of completely symmetric countries implies that the moments of domestic and foreign variables should converge to equal values as we increase the number of periods.

AR(1) -processes, and four short-run i.i.d. shocks (wage mark-ups, price mark-ups, Tobin's Q, and a monetary policy shock).

Since the original paper uses detrended data for estimation, the inflation target is set to zero. As this implies a very low steady state interest rate, we would find ourselves at the ZLB implausibly often. Hence, we add a positive annual inflation target of 1.8% to the Taylor rule, implying a quarterly steady state interest rate of 1.451%. Since the official target of the ECB is below, but near 2%, we believe this to be an adequate representation of Europe's monetary policy. All other model parameters are calibrated according to the reported posterior mode in Smets and Wouters (2003).²⁴

Smets and Wouters (2003) find that preference and productivity shocks are the chief drivers of fluctuations in nominal interest rates. It makes sense then to look at productivity shocks to illustrate how things change when the ZLB is imposed. In Figure 3, we plot the IRF to a large positive productivity shock (4.9 times the shock's estimated standard deviation). Since natural output increases faster than actual output, this produces a large negative output gap, and the ZLB is hit. Hitting the ZLB dampens the response of investment, and so rather than output returning to trend within four years (as it would in the absence of the ZLB), instead it takes around eight years for the output gap to close.

²⁴ The model code is taken from the extensive model database described in Cwik et al. (2012).




Figure 3: Impulse response to a large positive productivity shock in the Smets and Wouters (2003) model. (Dashed line imposes the ZLB, solid line does not.)²⁵

Obviously a realisation of a shock that is 4.9 times the shock's standard deviation is not particularly plausible, so one may wonder about the importance of the ZLB within the Smets and Wouters (2003) model. To assess this we ran a simulation of length 250,000 periods and noted when the ZLB was hit. With our preferred specification for the inflation target (1.8%), the implied probability of hitting the ZLB was only 0.01%. With an inflation target of 1.5% this increased to 0.02%, and with a target of 1.2% it hit 0.04%. This most likely reflects the absence of a financial accelerator mechanism from the Smets and Wouters (2003) model, and the fact that the data on which they estimated excluded the recent financial crisis.

However, although the probability of hitting the ZLB is very low in the Smets and Wouters (2003) model, the implications when it is hit can be severe. This is illustrated in Figure 4 which presents simulated paths for a period of 15 years during which the ZLB was hit. This figure also illustrates the increasing effect of the ZLB as the inflation target is decreased.

²⁵ The output gap, employment, investment and consumption are measured in percentage deviations from equilibrium, the interest rate and inflation are measured in percent; inflation and the interest rate are annualised.

Efficient simulation of DSGE models with inequality constraints




Figure 4: Simulations from the Smets and Wouters (2003) model. (Dark line does not impose the ZLB, others correspond to imposing the ZLB with inflation targets of 1.8%, 1.5% and 1.2%)²⁴

The solid black line represents the benchmark scenario, ignoring the bound; the grey lines represent scenarios with a ZLB for different inflation targets (the lighter the line, the lower the inflation target). For a 1.8% percent inflation target, the ZLB is hit in the 13th quarter, and holds for only two periods. In this case, the simulated path is nearly identical to the benchmark one ignoring the ZLB. However, the lower the inflation target, the more severe is the fall in prices and aggregate activity. For an inflation target of 1.2%, the ZLB binds for one year. During this period, monetary policy is unable to prevent a drop in investment, resulting in investment staying below equilibrium for eleven and a half years. Moreover, employment stays below its long-term steady state for about nine years. The conclusion of this exercise is that with a high enough inflation target, hitting the ZLB is incredibly unlikely, and so the expected welfare loss from the bound is minimal. However, once the bound is hit, the impact can be very strong and highly persistent, as we are presently seeing in reality.

Finally, the Smets and Wouters (2003) model gives us another opportunity to present our algorithm's speed advantage. With the first two inflation targets, our algorithm only took 10 minutes and 33 seconds, for 250,000 periods, rising to 22 minutes and 17 seconds with an inflation target of 1.2%. By contrast, the extended path algorithm of Adjemian and Juillard

(2011) took 3 hours and 45 minutes just to simulate 50000 periods, after which it crashed. Clearly, only our algorithm is viable on models of the scale of Smets and Wouters (2003).

5. Conclusion

This paper provides a fast, simple and intuitive method for the simulations of DSGE models with inequality constraints, based on the introduction of “shadow price shocks” which hit the bounded variables whenever the constraints are violated. We showed that at second order, the algorithm is more accurate than all methods except fully global ones, and we showed that the second order algorithm also leads in terms of speed, at least when compared to other algorithms of similar accuracy. At first order we illustrated that the speed was even greater, enabling it to be reliably used on the largest models around today.

We believe our algorithm will prove useful to a very wide variety of models including bounded variables, and we hope to investigate its application to the estimation of such models in future work.

6. References

- Adjemian, Stéphane, Houtan Bastani, Michel Juillard, Ferhat Mihoubi, George Perendia, Marco Ratto, and Sébastien Villemot. 2011. *Dynare: Reference Manual, Version 4*. Dynare Working Papers. CEPREMAP, April. <http://ideas.repec.org/p/cpm/dynare/001.html>.
- Adjemian, Stéphane, and Michel Juillard. 2011. *Accuracy of the Extended Path Simulation Method in a New Keynesian Model with a Zero Lower Bound on the Nominal Interest Rate*. Mimeo.
- Braun, R. Anton, and Lena Mareen Körber. 2011. New Keynesian Dynamics in a Low Interest Rate Environment. *Journal of Economic Dynamics and Control* 35, no. 1: 2213–2227.
- Braun, R. Anton, Lena Mareen Körber, and Yuichiro Waki. 2012. *Some unpleasant properties of log-linearized solutions when the nominal rate is zero*. Working Paper. Federal Reserve Bank of Atlanta. <http://ideas.repec.org/p/fip/fedawp/2012-05.html>.
- Chen, Qianying, Michael Funke, and Michael Paetz. 2012. *Market and Non-Market Monetary Policy Tools in a Calibrated DSGE Model for Mainland China*. BOFIT Discussion Papers. Bank of Finland, Institute for Economies in Transition, July. http://ideas.repec.org/p/hhs/bofitp/2009_014.html.
- Christiano, L. J., M. Eichenbaum, and S. Rebelo. 2011. When is the Government Multiplier Large. *Journal of Political Economy* 119: 78–121.

- Clarida, R., J. Galí, and M. Gertler. 1999. The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature* 37: 1661–1701.
- . 2002. A Simple Framework for International Monetary Policy Analysis. *Journal of Monetary Economics* 49, no. 5: 879–904.
- Coenen, G., A. Orphanides, and V. Wieland. 2004. Price Stability and Monetary Policy Effectiveness when Nominal Interest Rates are Bounded at Zero. *The B.E. Journal of Macroeconomics* 4, no. 1: -.
- Cwik, Tobias, Gernot Mueller, Sebastian Schmidt, Volker Wieland, and Maik H. Wolters. 2012. A New Comparative Approach to Macroeconomic Modeling and Policy Analysis. *forthcoming in Journal of Economic Behavior and Organization*.
- Eggertsson, Gauti B., and Michael Woodford. 2003. The Zero Bound on Interest Rates and Optimal Monetary Policy. *Brookings Papers on Economic Activity* 34, no. 1: 139–235.
- Erceg, Christopher J., and Jesper Lindé. 2010. *Is there a fiscal free lunch in a liquidity trap?* International Finance Discussion Papers. Board of Governors of the Federal Reserve System (U.S.). <http://ideas.repec.org/p/fip/fedgif/1003.html>.
- Fernández-Villaverde, Jesús, Grey Gordon, Pablo A. Guerrón-Quintana, and Juan Rubio-Ramírez. 2012. *Nonlinear Adventures at the Zero Lower Bound*. NBER Working Papers. National Bureau of Economic Research, Inc, May. <http://ideas.repec.org/p/nbr/nberwo/18058.html>.
- Fernández-Villaverde, Jesús, and Juan Rubio-Ramírez. 2010. *Macroeconomics and Volatility: Data, Models, and Estimation*. NBER Working Papers. National Bureau of Economic Research, Inc, December. <http://ideas.repec.org/p/nbr/nberwo/16618.html>.
- Fuhrer, Jeffrey C., and Brian F. Madigan. 1997. Monetary Policy When Interest Rates Are Bounded At Zero. *The Review of Economics and Statistics* 79, no. 4 (November): 573–585.
- Funke, Michael, and Michael Paetz. 2012. *A DSGE-Based Assessment of Nonlinear Loan-to-Value Policies: Evidence from Hong Kong*. BOFIT Discussion Papers. Bank of Finland, Institute for Economies in Transition, May. http://ideas.repec.org/p/hhs/bofitp/2012_011.html.
- Galí, J., and M. Gertler. 1999. Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics* 44: 195–222.
- Guerrieri, Luca, and Matteo Iacoviello. 2012. A Toolkit to Solve Models with Occasionally Binding Constraints Easily. September 21.
- Holden, Tom. 2010. *Products, patents and productivity persistence: A DSGE model of endogenous growth*. Economics Series Working Paper. University of Oxford, Department of Economics. <http://ideas.repec.org/p/oxf/wpaper/512.html>.
- Juillard, Michel. 2011. *Local approximation of DSGE models around the risky steady state*. Wp.comunite. Department of Communication, University of Teramo. <http://ideas.repec.org/p/ter/wpaper/0087.html>.
- Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe. 2005. Optimal Monetary Policy at the Zero-Interest-Rate Bound. *Journal of Money, Credit and Banking* 37, no. 5 (October 1): 813–835.
- Kim, Sunghyun Henry, Jinill Kim, Ernst Schaumburg, and Christopher A. Sims. 2008. Calculating and Using Second Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models. *Journal of Economic Dynamics and Control* 32, no. 11: 3397–3414.
- Krueger, Dirk, Felix Kubler, and Benjamin A. Malin. 2011. Solving the multi-country real business cycle model using a Smolyak-collocation method. *Journal of Economic Dynamics and Control* 35, no. 2 (February): 229–239.

- Maliar, Serguei, Lilia Maliar, and Kenneth Judd. 2011. Solving the multi-country real business cycle model using ergodic set methods. *Journal of Economic Dynamics and Control* 35, no. 2 (February): 207–228.
- Rotemberg, Julio J. 1996. Prices, output, and hours: An empirical analysis based on a sticky price model. *Journal of Monetary Economics* 37, no. 3 (June): 505–533.
- Smets, F., and R. Wouters. 2003. An Estimated Dynamic Stochastic General Equilibrium Model for the Euro Area. *Journal of the European Economic Association* 1, no. 5.

7. Appendices

7.1. The NK model of Fernández-Villaverde et al. (2012)

Households choose consumption C_t and labour supply L_t to maximise their utility, given the wage W_t , the rate of inflation Π_t and the nominal interest rate R_t . This leads to the FOCs:

$$\frac{1}{C_t} = \mathbb{E}_t \frac{\beta_{t+1} R_t}{C_{t+1} \Pi_{t+1}}, \quad \psi L_t^\vartheta C_t = W_t.$$

Firms choose prices to maximise profits. This leads them to set a relative price Π_t^* satisfying:

$$\begin{aligned} \text{MC}_t &= \frac{W_t}{A_t}, \quad \varepsilon x_{1,t} = (\varepsilon - 1)x_{2,t}, \\ x_{1,t} &= \frac{\text{MC}_t Y_t}{C_t} + \theta \mathbb{E}_t \beta_{t+1} \Pi_{t+1}^\varepsilon x_{1,t+1}, \quad x_{2,t} = \Pi_t^* \left(\frac{Y_t}{C_t} + \theta \mathbb{E}_t \beta_{t+1} \frac{\Pi_{t+1}^{\varepsilon-1}}{\Pi_{t+1}^*} x_{2,t+1} \right), \end{aligned}$$

where A_t is productivity, MC_t is marginal costs and Y_t is output.

The central bank follows a standard Taylor rule, with monetary policy shock M_t , and government spending is a stochastic fraction $S_{g,t}$ of total output:

$$R_t = \max \left\{ 1, R^{1-\rho_r} R_{t-1}^{\rho_r} \left[\left(\frac{\Pi_t}{\Pi} \right)^{\phi_\pi} \left(\frac{Y_t}{Y} \right)^{\phi_y} \right]^{1-\rho_r} M_t \right\}, \quad G_t = S_{g,t} Y_t.$$

Inflation and price dispersion evolve according to:

$$1 = \theta \Pi_t^{\varepsilon-1} + (1 - \theta) (\Pi_t^*)^{1-\varepsilon}, \quad v_t = \theta \Pi_t^\varepsilon + (1 - \theta) (\Pi_t^*)^{-\varepsilon}.$$

Finally, market clearing in goods and labour markets imply:

$$Y_t = C_t + G_t, \quad Y_t = \frac{A_t L_t}{v_t}.$$

The stochastic processes β_t , A_t , M_t and $S_{g,t}$ are all log AR(1).

7.2. Our modified version of the two-country model of Clarida et al. (2002)

Let \tilde{y}_t , π_t and i_t represent the domestic output gap, inflation rate and interest rate respectively, and let \tilde{y}_t^* , π_t^* and i_t^* denote their foreign equivalents. Let v_t and v_t^* be home and foreign demand shocks. When both countries are symmetric, the linearised model is described by the following four equations, and another four in which the roles of home and foreign country are swapped:

$$\begin{aligned} \tilde{y}_t &= \mathbb{E}_t \tilde{y}_{t+1} - \sigma_0^{-1} (i_t - \mathbb{E}_t \pi_{t+1} - \kappa_0 \mathbb{E}_t \Delta \tilde{y}_{t+1}^* + \kappa_0 (1 - \rho_v) v_t - \bar{i}), \\ \pi_t &= \Phi (\theta \beta \mathbb{E}_t \pi_{t+1} + \tau \pi_{t-1}) + \lambda \tilde{y}_t, \\ i_t &= \max\{0, \bar{i} + (1 - \phi_i) (\phi_\pi \pi_t + \phi_y \tilde{y}_t) + \phi_i (i_{t-1} - \bar{i})\}, \\ v_t &= \rho_v v_{t-1} + \varepsilon_t^v. \end{aligned}$$

These represent the Euler equation, the NK Phillips curve, the Taylor rule and the evolution of the exogenous shock, respectively.

The parameters used in the equations above are the following functions of the structural parameters:

$$\kappa_0 = \frac{\sigma - 1}{2}, \quad \sigma_0 = \sigma - \kappa_0, \quad \kappa = \sigma - \kappa_0 + \varphi, \quad \lambda = \frac{(1 - \theta)(1 - \beta\theta)}{\theta} \kappa, \quad \bar{i} = \frac{1}{\beta}$$

$$\Phi = \frac{1}{\theta + \tau(1 - \theta(1 - \beta))}$$

where $\sigma := 1/3$ is the inverse elasticity of intertemporal substitution in consumption,²⁶ $\varphi := 1$ is the inverse Frisch elasticity of labour supply, $\beta := 0.99$ is the discount factor, $\theta := 3/4$ is the probability a firm's price must remain fixed and $\tau := 1/10$ is the fraction of backwards looking price setters. We also set $\phi_\pi := 1.5$, $\phi_y := 0.125$ and $\phi_i := 0.8$ in the central bank's reaction function and the shocks' persistence as $\rho_v := 0.7$.

²⁶ The model implicitly assumes a unitary substitution elasticity between domestic and foreign goods. Hence, for values of σ smaller than one, any increase in foreign production is accompanied by an increase in domestic production, since the negative effect due to the implied real appreciation is dominated by the fall in the real interest rate stemming from an increase in domestic consumer price inflation. With $\sigma = 1$ (i.e. log utility) both effects balance each other out, since in that case $\kappa_0 = 0$, meaning there is no transmission at all. We could derive similar results for lower substitution elasticities and higher values of σ , but we maintain this parameterisation to keep the model close to that of Clarida et al. (2002).