How large are hysteresis effects? Estimates from a Keynesian growth model*

Steven Fazzari
Washington University in St. Louis
fazz@wustl.edu

Alejandro González
Washington University in St. Louis
alejandro.g@wustl.edu

October, 2022

Abstract: This paper estimates a demand-led model of macroeconomic growth and fluctuations in which the growth rate of the economy's supply side converges to the growth rate of demand. Convergence happens because labor supply and productivity growth respond to the degree of slack in the economy. Faster demand growth reduces slack and stimulates supply (and vice-versa). We estimate the model using simulated method of moments and find quantitatively important hysteresis effects: the elasticity of productivity growth and labor supply growth to the unemployment rate are 0.71 and 0.26, respectively. For an economy with labor market slack, these estimates imply an increase in the growth rate of demand of 1% would decrease long-run unemployment by 0.75%. Additionally, we show the model replicates several features of business cycles as well the response of the economy to autonomous demand shocks, providing further validation of this approach to understanding macroeconomic dynamics.

JEL Codes: *E32; E12; O41*
Keywords: *Hysteresis, Demand-Led Growth, Supermultiplier*

*We thank Fabio Ghironi, Brian Greany and Robert Blecker for suggestions of previous versions of this draft. We would also like to thank conference participants at the Keynesian Economics Working Group, the 2020 Annual Young Scholars Initiative Plenary, the 3rd Warsaw Money-Macro-Finance conference, the International Workshop on demand-led growth, the 33rd Annual EAEPE conference, and seminar participants at Washington University, St. Louis, the University of Washington, and Roma Tre university for their comments and suggestions. Remaining errors are our sole responsibility.

1 Introduction

What is the role of demand in driving macroeconomic activity? No one would deny demand is necessary to motivate production in a market economy; what is not sold will not be produced. However, mainstream macroeconomic models since (at least) Modigliani (1944) have relegated the role of demand to the short run, asserting that endogenous market-driven adjustments of wages and prices assure sufficient demand to purchase full-employment, potential output. These adjustments may not be instantaneous, and therefore autonomous fluctuations in aggregate demand may affect real variables for some time, but the long-run path of an economy is typically understood as independent of direct effects of spending choices, driven instead entirely by the supply side.

This perspective changed to some extent in the last two decades of the 20th century. In "new consensus" models, wise monetary policy replaced nominal adjustment as the instrument to restore aggregate demand to potential levels. This change elevated the importance of activist policy relative to "natural" market forces in generating convergence of output to a supply-driven path. But it still relegated to the role of autonomous demand dynamics to the short run, a period now defined by the speed of the transmission of monetary policy to spending decisions.

The events following the Great Recession have challenged this mainstream consensus, forcing macroeconomists to rethink the role of aggregate demand in shaping the medium- to long-run performance of an economy. In the US, output failed to regain its pre-recession trend at any point in the approximately 12 years from the beginning of the Great Recession to the COVID-19 pandemic lockdown. Some analysts, interpreting this disappointing outcome through the conventional new consensus lens, looked for some kind of a supply shock, independent of the demand collapse that caused the Great Recession (Fernald, 2015; Cette, Fernald, and Mojon, 2016). But another interpretation emerged raising the suspicion that weak demand can constrain the economy beyond the short run. Summers (2014), harking back to Alvin Hansen, labeled this phenomenon "secular stagnation."

In this paper, we present a simple framework to empirically study the role that autonomous aggregate demand — components of demand independent of economic conditions — plays in shaping medium- to long-term economic performance. Our model is based on the demand-led growth model from Fazzari, Ferri and Variato (2020, hereafter FFV). The demand side of this model follows the "supermultiplier" approach pioneered by Serrano (1995) and developed in a variety of recent Post-Keynesian contributions. [1] The main innovation in FFV is an explicit model of an endogenous supply side that can accommodate demand-led path. Supply accommodates demand growth (not just demand levels) through hysteresis effects. Stronger demand reduces slack in the economy, measured by the unemployment rate, and therefore induces faster growth of labor supply growth and labor productivity.

We estimate this Keynesian growth model through a minimum distance estimator and assess its empirical performance along three dimensions. First, we ask whether it can match the volatility, persistence and co-movements of seven macroeconomic variables at both business-cycle and lower frequencies. Second, we compare the estimated model impulse-response functions induced by a large demand shock to the empirical impulse responses obtained by Girardi, Paternesi and Stirati (2020) through local projections. Third, we test whether the key hysteresis effect present in the model, which acts through the adjustment of the steady-state unemployment rate to autonomous demand, is different from zero and quantitatively significant.

While there is an extensive literature on demand-led growth models - indeed, this is a pillar of Post-Keynesian Economics (Lavoie, 2014) - there is surprisingly little empirical work estimating a demand-led, dynamic macroeconomic model and comparing its predictions to the data. Recent studies have tested several implications of Post-Keynesian demand-led growth models (Girardi and Pariboni, 2020; Haluska, Braga and Summa, 2021; Girardi et. al, 2020). To the best of our knowledge, however, we are the first to use this kind of model to study quantitatively business-cycles and growth in a unified framework and to be explicit about the structural shocks driving the data.

Recent mainstream research analysing the persistent effects of aggregate demand slumps on economic performance usually requires a lower bound on the interest rate determined by monetary policy as the reason for demand effects beyond the short run (Benigno and Fornaro, 2018; Eggertsson, Mehrotra and Robbins, 2019; Anzoategui, Comin, Gertler and Martinez, 2019). We contend, however, that monetary policy may not be effective at restoring full employment and potential output even when short-term nominal interest rates are not pinned at zero.

Our main results are as follows: First, we find that estimating the model through simulated

[1] Early contributions include Allain (2015), Fazzari et al. (2013), Freitas and Serrano (2015), and Lavoie (2016). This literature has grown rapidly in recent years.

method of moments yields results that match the data reasonably well, especially regarding the co-movements and volatility of employment and labor productivity relative to output. In particular, the model outperforms real business cycle models augmented with search and matching frictions. It also yields correlations between employment, labor productivity, and unemployment that are significantly closer to the data than neoclassical models. Second, when simulating structural impulse-response functions for our model with the estimated parameters, we find our model reproduces, at least qualitatively, most of the responses to a plausibly exogenous autonomous demand shock over a decade-long horizon, as estimated by Girardi, et al. (2020) through local projections. Third, we find strong evidence for the hysteresis effects that the model identifies as necessary for the supply side to converge to the demand side. Quantitatively, we estimate that a one percentage point increase in the long-run annual growth rate of demand can be accommodated by the endogenous response of the supply side with a less than one percentage point drop in the steady-state unemployment rate. This result implies that it is empirically reasonable for persistent changes in aggregate demand dynamics to pull the supply side along with them, in both a positive and negative direction.

Related Literature. Our paper is connected to four strands of research. First, it is related to the empirical evaluation of the supermultiplier approach to growth and distribution. This literature has proceeded in three major ways; a first group of papers attempts to establish empirical candidates for autonomous demand, and seeks to show that autonomous demand Granger causes output, as implied by the model as well as testing that autonomous demand and output are co-integrated, which is also implied by the model. [2] A second group of papers seeks to show that the investment share responds strongly and positively to an increase in the growth rate of autonomous demand, another important implication of the supermultiplier model. [3] A third set of papers estimates the effects of autonomous demand shocks on output and other macroeconomic variables, either by the use of local projections (Girardi et. al, 2020) or structural vector autoregression with recursive restrictions (Goes and Deleidi, 2022). While this work has provided important support for the implications of the supermultiplier model, none of these papers attempt to estimate the underlying economic parameters; they all use reduced-forms to test implications of the model. Our paper contributes to this literature by providing explicit estimates of economically meaningful parameters which can be used to perform quantitative exercises and used to assess the stability conditions arising from

[2](See Girardi and Pariboni (2016); Perez-Montiel and Erbina (2020); Perez-Montiel and Manera(2022); and Perez-Montiel and Pariboni(2022)

[3]See Girardi and Pariboni (2020); Perez-Montiel and Erbina (2020); and Haluska, Braga and Summa (2021).

these models. In addition, we are the first to compute explicit impulse response functions in response to an autonomous demand shock, which can then be compared to those obtained from structural vector autoregressions or local projections.

There is a second strand of research that models demand-induced technical change in Post-Keynesian growth models. The mechanisms we present linking the supply side to the unemployment rate need not be explicitly confined to a supermultiplier model; they can be (and have) been incorporated in Neo-Kaleckian growth models, for example see Chapter 8 of Hein (2014) for a thorough discussion, and Tavani and Zamparelli (2018) for a broader discussion of induced technical change in Post-Keynesian growth models. The empirical counterpart to this literature usually focuses on estimates of Verdoorn's Law, which is the reduced-form correlation between output growth and productivity growth. We contribute to this literature by estimating Verdoorn's coefficient in a general-equilibrium framework along with all other model parameters, and we show how Verdoorn's coefficient can be derived as a reduced-form outcome from our underlying supply-side parameters. Our estimates imply a coefficient for Verdoorn's law that is slightly higher than what is found for OECD countries, but all of the Verdoorn effect comes from the impact of the unemployment rate on productivity growth. We cannot reject the hypothesis of a zero effect from the feedback of capital accumulation on productivity growth, which is consistent with the findings of Hein and Tarassow (2010).

The third strand of literature combines insights from endogenous growth theory with mainstream New Keynesian models to produce long-lived effects of demand shocks when the economy is at the zero lower bound.[4] In these models, a decline in demand when the economy is at the zero lower bound induces a large drop in profitability, which hampers innovation and R&D expenditures, reducing productivity growth and leaving long-lasting effects on output growth and permanent effects on the level of output. Our model produces these same long-lasting effects of demand shocks on productivity and output, but since interest rates and the price level are not incorporated in the model, it is silent on the issue of the zero lower bound.[5] We contribute to this literature by estimating our model on a sub-sample before the zero lower bound binds in the US, and we show that our hysteresis parameters remain economically and statistically significant in that period.

[4]See Benigno and Fornaro (2018); Eggertsson, Mehrotra and Robbins (2019); Anzoategui, and Comin, Gertler and Martinez (2019).

[5]Some Post-Keynesian research is sceptical about the ability of monetary policy to eliminate real effects of demand dynamics even if the interest rate is not constrained by a lower bound. See Fazzari (2020) for a summary of these arguments.

The fourth strand of literature makes use of identified vector autoregressions to asses whether transitory demand shocks have long-run effects on output. Furlanetto et. al (2021) and Maffei-Faccioli (2021) use sign restrictions on structural vector autoregressions to identify demand and supply shocks, and show demand shocks have permanent effects on the level of output. Anzoategui and Kim (2021) use a structural vector autoregression derived from a New Keynesian model to estimate the output gap, and show demand shocks have long-lasting but temporary effects on the output gap, which is a limited form of hysteresis. The framework that we present in this paper provides a simple way to account for the permanent effects of demand shocks on output growth and the supply-side of the economy, consistent with this SVAR evidence. It also shows qualitatively similar impulse-response functions to autonomous demand shocks can be generated from a simple Keynesian model.

Outline The rest of the paper is organised as follows: Section 2 presents a simpler version of FFV, which we use to develop intuition and characterise how the model responds to autonomous demand shocks. Section 3 presents the full model we use for estimation, which is a stochastic version of FFV. Section 4 presents the data used, our definition of autonomous demand, and the minimum distance estimator. Section 5 presents our main empirical results, including parameter estimates, a comparison between the moments generated by the model and those present in the data, and impulse-response functions. Section 5 concludes and outlines avenues for future research.

2 A simple demand-led growth model

This section presents a simple, stylized version of the model that delivers the key theoretical results: even in the long-run, macroeconomic performance depends on demand, demand shocks can have permanent effects on output and productivity, and the supply side accommodates the growth rate of demand. In addition, the simple version allows us to rigorously characterise the impulse-response functions to an autonomous demand shock.

In this simple model, there is no capital and labor supply is fixed and normalized to 1. In addition, while the simple model incorporates the dynamics of the supply side, we assume resources constraints are not binding, in the sense that the unemployment rate remains above a minimum level, which implies output at any point in time is determined by demand. The more general model in the next section relaxes these assumptions. However, the key results derived here continue to hold in a more general environment.

2.1 Steady-state results

We work in discrete time. Consumption depends on income, which equals output:

$$C_t = (1 - s)Y_t \tag{1}$$

Where s is the marginal propensity to save, C_t is consumption and Y_t is output. Demand also includes an *autonomous* component, F_t, independent from current and past output, growing at an exogenous rate, g^*. [6]Related research has considered many possible sources of such exogenous demand growth. Military spending is a natural candidate, since it can be argued that such expenditures are de-coupled from national economic performance (Allain, 2015). In small open economies, another natural candidate is exports, since at least part of the secular trend in exports is explained by the rate of growth of the rest of the world (Nah and Lavoie, 2017). Other authors propose that residential investment could be considered autonomous because it is largely determined by financial conditions, especially mortgate interest rates, rather than the current level of income (Fiebiger, 2018; Perez-Montiel and Pariboni, 2021). At this stage we do not take a stand on what constitutes autonomous demand; we discuss this issue further when we turn to estimating the model.

Autonomous demand has a constant exponential growth rate:

$$F_t = F_0 e^{g^* t} \tag{2}$$

Equilibrium in the goods market equates output (Y_t) to demand (Y_t^D):

$$Y_t = Y_t^D = C_t + F_t \tag{3}$$

which implies:

$$Y_t = \frac{1}{s} F_t. \tag{4}$$

Taking log differences and using equation 2:

$$g_Y = g^* \tag{5}$$

[6]Many macroeconomic models include government spending. The assumption that autonomous demand grows at an exogenous rate is less common. Supermultiplier models often make this assumption to generate analytical results. But when interpreting the relevance of the models for empirical analysis, authors discuss variations in autonomous demand growth (see Girardi et al. 2020 and FFV, section 6, for example). In the empirical analysis to follow we measure autonomous demand in the data and do not impose a constant growth rate.

where g_Y denotes the growth rate of output. This is simple a dynamic extension of the even more simple textbook Keynesian model. Equilibrium output at any point in time is autonomous demand times a multiplier that depends on the propensity to save. If the multiplier is constant, then the growth rate of demand-determined output equals the growth rate of autonomous demand.

What about the supply side of the economy? What guarantees that a balanced growth path exists such that supply accommodates the demand-led growth path? To answer these questions, begin by assuming that the ability of the economy to produce output, that is, supply-determined potential output, is a linear function of the state of technology (A_t) and labor (l_t).

$$Y_t^S = A_t l_t \tag{6}$$

The linearity assumption is made here for convenience; there is no loss of generality by considering a diminishing-returns technology. In addition, the size of the population is fixed to 1, and all workers are either employed or unemployed. Hence:

$$1 = u_t + l_t \tag{7}$$

The key assumption made in FFV, also present in Aghion and Howitt (1994) and Dutt (2006), is that productivity growth depends on the state of the economy, measured by the unemployment rate. Following Ester Boserup (1965), one could argue that necessity is the mother of invention, and labor-saving innovations depend on the tightness of the labor market and the effect of the labor market on wage costs. Alternatively, the degree of slack in the economy, partially determines cash flows, and if firms are financially constrained, cash flows impact investment in R&D (Brown, Fazzari and Petersen, 2009). These structural channels motivate a reduced-form link between unemployment and productivity growth:

$$g_A = \phi_0 - \phi_1 u_t. \tag{8}$$

With fixed labor supply, g_A determines the growth rate of Y_t^S.

For the simple model to deliver a balanced growth path, aggregate supply and demand must grow at the same rate. In steady state, demand growth is determined by the growth rate of autonomous spending. To equate supply growth to g^* use equation 8 to solve for the steady-state rate of unemployment:

$$u^* = \frac{\phi_0 - g^*}{\phi_1} \tag{9}$$

The unemployment rate along the balanced growth path depends on the growth rate of demand. Consider a positive shock to g^*. Faster growth of demand lowers the unemployment rate, which raises the growth rate of productivity, causing supply to accommodate the acceleration of demand. The lower bound on unemployment determines the maximum feasible growth rate: faster demand growth can increase supply growth only up to the point where there are no more hands or minds able to work. In this simple version of the model and with a zero lower bound on the unemployment rate, the maximum growth rate is ϕ_0, the intercept of the productivity growth equation. Equilibrium conditions do not determine a long-run unemployment rate independent of the demand side in this environment - supply constraints only impose a lower bound on the unemployment rate.

The diagram below in the (u, g) space illustrates the Keynesian steady-state in this model, and the effects of an increase in autonomous demand growth, when the lower bound on unemployment is not binding. The AD schedule depends on autonomous demand growth only. The AS schedule slopes downward because higher supply growth requires a lower unemployment rate to stimulate faster productivity growth. Acceleration in demand growth from g^* to g^{**} shifts the AD curve upward. Demand stimulus reduces unemployment and the stronger economy with less slack raises supply growth to match the faster rate of demand growth. To sustain faster supply growth in equilibrium, the unemployment rate must remain lower than its initial value, that is, faster demand growth permanently reduces the unemployment rate.

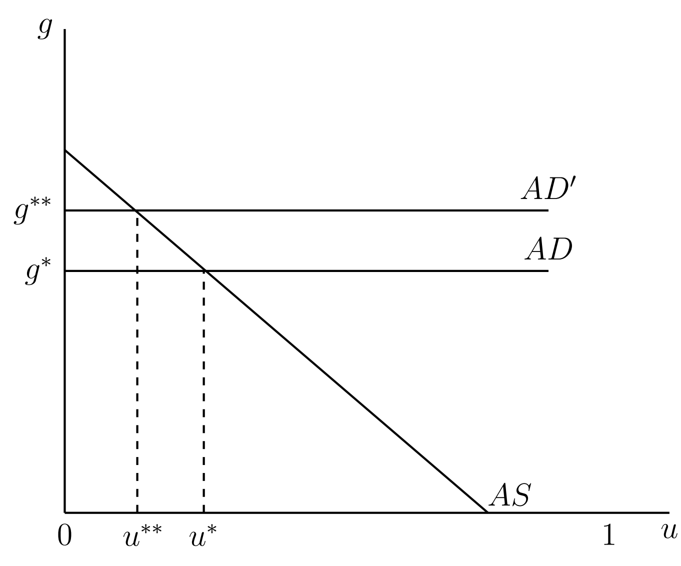

Figure 1: An increase in the rate of growth of demand

2.2 The effects of permanent demand shocks

This simple model can also be used to study macroeconomic fluctuations and persistent changes in the output and employment paths. To do so, we assume that log autonomous demand evolves according to a random walk:

$$\ln F_t = g^* + \ln F_{t-1} + \varepsilon_t \tag{10}$$

assuming the demand shock, ε_t is a white noise process. This specification implies that a unitary shock to log autonomous demand induces a permanent one percentage point increase in the level, but not the growth rate, of future autonomous demand. Using this specification, we show in Appendix 1 that the time-series processes for output, unemployment and growth rate of productivity are:

$$\ln Y_t = g^* + \ln Y_{t-1} + \varepsilon_t \tag{11}$$

$$u_t = \phi_0 - g^* + (1 - \phi_1)u_{t-1} - \varepsilon_t \tag{12}$$

$$g_{A,t} = \phi_1 g^* + (1 - \phi_1)g_{A,t-1} + \phi_1 \varepsilon_{t-1} \tag{13}$$

Therefore, log-output is a random walk with a drift, the unemployment rate is an AR(1) process, and log-productivity is an ARIMA(1,1,1) process. Given these simple univariate time-series process, it is straightforward to compute their impulse response functions after a temporary demand shocks. Let $IRF^h_{y,\varepsilon}$ be the impulse response function at horizon h of the endogenous variable y after being shocked, at time $t = 0$, with a unit shock on ε. Then, we have:

$$IRF^h_{\ln Y,\varepsilon} = 1 \quad \forall h > 0 \tag{14}$$

$$IRF^h_{u,\varepsilon} = (1 - \phi_1)^h \quad \forall h > 0 \tag{15}$$

$$IRF^h_{g_A,\varepsilon} = \phi_1(1 - \phi_1)^h \quad \forall h > 1 \tag{16}$$

Hence, a transitory demand shock has a permanent effect on the log-output, and a transitory effect on the level of the unemployment rate. It also has a transitory effect on the *growth* *rate* of labor productivity, but it has a permanent effect on it's level. These predictions are partially consistent with the empirical results obtained by Furlaneto et. al (2021), where, using a structural VAR with both sign and long-run restrictions, they find that a contractionary demand shock permanently reduces output and unemployment; however, they find that demand shocks have little effect on labor productivity.

3 A more complete model for estimation

The simple model transparently illustrates the theoretical intuition for how demand leads the economy and the supply side adjusts to the demand-led output path. We now extend the previous results to include capital, investment, endogenous labor supply, and expectations. Our purpose is to develop a model that is rich enough to be compared with time-series moments in macroeconomic data and can generate empirical parameter estimates to test the main implications of a demand-led growth model.

3.1 The demand side

The demand side is identical to FFV (2020) and will be summarized here. Agents take decisions based on their expectations of output growth; agents use adaptive expectations to forecast current output:

$$Eg_t = (1-\alpha)g_{t-1} + \alpha Eg_{t-1} \tag{17}$$

Agents update their forecasts with a weighted average of their previous growth rate expectation and the most recent observation of actual growth. This specification is equivalent to a forecast based on geometrically declining weights of past output growth. Although simple adaptive expectation models have been criticized, we interpret this specification as an approximation to a more general model of learning. [7].

The consumption equation now takes the form[8]:

[7] We draw support for this interpretation from Lawrence Summers who stated during a March 22, 2022 interview with Ezra Klein that "a consensus view would be that people learn from the past, because what else would you learn from? And so people form their expectations based on what they've observed recently and based on what they think is going to happen in the future" https://www.nytimes.com/2022/03/29/podcasts/transcript-ezra-klein-interviews-larry-summers.html

[8] Taxes are not modelled explicitly. A proportional income tax rate can be included in s without loss of generality.

$$C_t = (1 - s)(1 + Eg_t)Y_{t-1} \tag{18}$$

Business investment in period t (I_t) becomes productive capital in period $t+1$ (K_{t+1}) and targets an evolving capital-output ratio ($\hat{v}_t$) :

$$I_t = K_{t+1} - (1 - \delta_K)K_t = \hat{v}_t(1 + Eg_t)^2 Y_{t-1} - (1 - \delta_K)K_t \tag{19}$$

where δ_K is the geometric depreciation rate of the capital stock. The period-by-period target capital output ratio ($\hat{v}_t$) evolves according to a partial adjustment rule to eventually reach the long-run, technologically determined ratio (v^*)[9]:

$$\hat{v}_t = (1 - \lambda)v_{t-1} + \lambda v^*. \tag{20}$$

The parameter λ controls the speed of adjustment as in Freitas and Serrano (2017). Autonomous demand (F_t) follows a trend-stationary process in logs with growth rate g^*:

$$\ln F_{t+1} = g^* t + \rho \ln F_t + \varepsilon_t \tag{21}$$

where $\varepsilon_t \sim i.i.d.N(0, \sigma^2)$. These white-noise shocks ε_t, along with the propagation mechanisms inside the model, account for all the cyclical and low-frequency fluctuations generated by the model.

Aggregate demand is the sum of the demand components, and output equals demand (unless demand exceeds the supply-determined potential to produce):

$$Y_t = Y_t^D = C_t + I_t + F_t \tag{22}$$

Substituting specifications for the demand components into the current output equation and dividing by lagged output yields the law of motion for output growth:

$$1 + g_t = (1 - s)(1 + Eg_t) + \hat{v}_t(1 + Eg_t)^2 - v_t(1 - \delta)(1 + g_t) + f_t(1 + g_t) \tag{23}$$

where $g_t = (Y_t/(Y_{t-1}) - 1,$, $v_t = K_t/Y_t$, $f_t = F_t/Y_t$. If autonomous demand grows at a constant rate $g_t^* = g^*$, the model has a steady state given by:

$$g_t = Eg_t = g^* \tag{24}$$

[9]This specification may be interpreted as a reduced-form from a capital adjustment cost model.

$$v_t = \hat{v}_t = v^* \tag{25}$$

$$f_t = f^* = s - v^*(g^* + \delta) \tag{26}$$

Note that the steady-state share of autonomous demand in output is just the difference between the share of saving and the share of investment.

This steady state is unique; and, as shown in FFV(2020) the dynamics are stable for a wide range of empirically plausible parameter values.

3.2 The supply side

The economy's potential output now depends on capital as well as labor and the Leontief production function becomes

$$Y_t^S = \min\{A_t N_t, \frac{K_t}{\hat{v}_{t-1}}\}. \tag{27}$$

where N_t is labor supply. Since we now relax the assumption that labor supply is fixed, the resource constraint for labor is:

$$N_t = L_t + U_t \tag{28}$$

or, after some re-arranging, $L_t = (1 - u_t)N_t$. With an efficient choice of technique,that is, with no redundant labor, $Y_t = A_t L_t$. Capital may or may not be fully utilized.

Hysteresis effects connect supply growth to demand dynamics through two structural channels. First, labor force growth (g_t^{LS}) depends on the strength of the economy as measured by the unemployment rate:

$$g_t^{LS} = \theta_0 - \theta_1 u_{t-1} \tag{29}$$

Second, the growth rate of the labor productivity term (g_t^A) depends, again, on the current state of the economy through the unemployment rate as well as the rate at which the capital stock is renewed by new investment ($g_{t-1}^K + \delta$,the growth rate of the capital stock plus the depreciation rate):

$$g_t^A = \phi_0 - \phi_1 u_{t-1} + \phi_2(g_{t-1}^K + \delta) \tag{30}$$

Note that the productivity growth equation connects supply growth to both the *level* of the economy (represented by ϕ_1) and the *growth rate* of the economy (represented by ϕ_2). The growth-on-growth effect relates to the long-standing Kaldor-Verdoorn concept that productivity grows faster when the aggregate economy grows faster. However, the connection between the parameters in the productivity equation and empirical evidence for Kaldor-Verdoorn effects is somewhat nuanced, as we discuss when we present estimates of these parameters below. FFV (2020) presents more motivation and extensive references to justify these supply-side equations. González (2022) provides microfoundations for these productivity effects based on wage bargaining and learning-by-doing.

3.3 Steady-State Unemployment and Hysteresis

The key result from FFV (2020) is that supply growth will accommodate the demand growth path as long as the effect of unemployment on supply growth is positive ($\theta_1 + \phi_1 > 0$). Because demand growth is g^* in steady state, equations 29 and 30 can be used to solve for the steady-state unemployment rate that assures supply growth matches demand growth:

$$u^* = \frac{\theta_0 + \phi_0 - g^*(1 - \phi_2) + \phi_2\delta}{\theta_1 + \phi_1} \tag{31}$$

These results imply several interesting implications of the demand-led growth with accommodating supply approach. Clearly there is no "natural" rate of unemployment defined independently of the demand side of the economy. The engine of long-run demand growth (g^*) affects steady-state unemployment. As is the case in the simple model, the minimum unemployment rate limits how fast the economy can grow, but this rate is in no sense an equilibrium. Sluggish demand growth can trap the economy in a steady state with unemployment persistently higher than the level imposed by supply constraints.

The steady-state equation for the unemployment rate provides a framework to measure hysteresis effects. Suppose the growth rate of demand, g_z, increases. What will happen to the steady-state rate of unemployment? This is given by the following expression:

$$\frac{du}{dg_z} = -\frac{(1 - \phi_2)}{\theta_1 + \phi_1} \tag{32}$$

This derivative measures the response of the unemployment rate to the growth rate of demand, which is given by the inverse of the response of the supply-side parameters to demand. The intuition is simple: if there is a rise in autonomous demand growth, the unemployment rate falls and the supply-side of the economy reacts to this decrease in unemployment by accelerating productivity and labor supply growth. Hence, the supply side catches up to the demand-led growth path with faster steady-state supply growth induced by a lower unemployment rate, that is by an overall stronger economy.

Does this result mean that the economy can grow unboundedly faster by just increasing demand? Is the supply side, then, largely irrelevant? No. There is a lower bound on the unemployment rate; it cannot be negative. Suppose that the minimum unemployment rate is zero. Then, there is an upper bound on demand growth, $\bar{g}_z$ that the economy can accommodate, given by:

$$\bar{g}_z = \frac{\phi_0 + \theta_0 + \theta_2 \delta}{(1 - \phi_2)} \tag{33}$$

This equation shows that the exogenous growth rates of technical progress and labor supply determine the upper bound, or a *ceiling*, of the growth rate. If these parameters are too low, then autonomous demand growth will have little space to stimulate the economy. It is then natural to ask what is a reasonable empirical value for this ceiling, a question we tackle in our empirical exercise.

These results also relate to the interpretation of Verdoorn's law. Verdoorn found a statistical relationship between the growth rate of labor productivity and the growth rate of output, which was then made famous by Kaldor (1966). This reduced form, in our notation, is:

$$g^A = \alpha + \beta g^Y \tag{34}$$

where β is the Kaldor-Verdoorn coefficient. The β coefficient sometimes been interpreted as measuring the degree of demand-induced technical progress and the degree of increasing returns to scale present in the economy. Therefore, one could be tempted to use this statistical relationship as a building block of a macroeconomic model. However, as noted by (McCombie and Spreafico, 2015) and Basu and Budhiraja (2021), this is a statistical relationship, which must be derived from the underlying economic theory. To see the connection between the Kaldor-Verdorn coefficient and our model, write the steady-state version of our productivity growth equation, (27):

$$g^A = \phi_0 - \phi_1 u^* + \phi_2 (g^K + \delta) \tag{35}$$

Note that along the balanced growth path, it must be the case that $g^K = g^Y = g_z$, as we have shown above, and plugging in our equation for the steady-state unemployment rate, and some tedious but straightforward algebra:

$$g^A = \frac{\phi_0\theta_1 - \phi_1\theta_0 + \phi_2\delta(1 - \theta_1 - \phi_1)}{\theta_1 + \phi_1} + \frac{\phi_1 + \phi_2\theta_1}{\theta_1 + \phi_1} g_z \tag{36}$$

Therefore, along the balanced-growth path, the mapping between the Kaldor-Verdorn coefficient and our hysteresis parameters is:

$$\beta = \frac{\phi_1 + \phi_2\theta_1}{\theta_1 + \phi_1} \tag{37}$$

Given the parameter restrictions necessary for stability, that is $(\phi_1, \theta_1) \in (0, 1)^2$, the value of $\beta \in (0, \infty)$, β takes a value of 0 if $\phi_1 = \phi_2 = 0$ and $\theta_1 > 0$, that is, the growth rate of productivity is exogenous but there is some hysteresis in labor supply. On the other hand, if $\phi_2 = 0$ and $\theta_1 \to 0$, then $\beta \to \infty$. Hence, in contrast to Basu and Budhiraja (2021), our model economy implies that if the Kaldor-Verdoorn coefficient is strictly positive, then there is some influence of the demand side on the supply side of the economy. Note, however, that his coefficient bears no relation to the degree of static returns to scale: our Leontief production function assumes static constant returns. It is also clear that this statistical relationship is not informative about which parameters that govern the relationship between the supply side and the demand side are quantitatively important and makes it apparent that an estimation method that estimates simultaneously all the economically meaningful parameters that govern the supply-side hysteresis is needed, a task to which we now turn.

4 Estimating a Demand-Led Growth Model

The model presented in the previous section proposes simple relationships to describe the dynamics of consumption, investment, and autonomous demand, along with adjustment of the capital stock and expectations. Equations 29 and 30 introduce key hysteresis effects that create structural channels for demand dynamics to lead the supply side. In this section we develop a strategy to estimate the model parameters and quantitatively assess how well it can explain major features of US macroeconomic data.

4.1 Parameter Estimation

4.1.1 Data

We use data on the components of U.S. GDP for the period 1959:Q1 - 2019:Q4 obtained from the Federal Reserve Bank of St. Louis FRED data site. Nominal data are deflated by the same price index, the chain-weighted deflator for personal consumption expenditure. This index is the Federal Reserve's preferred measure of aggregate inflation. Using a common deflator is appropriate for studying demand-led dynamics because the objective is to capture expenditure flows. So-called "real" series for separate GDP components such as business fixed investment and personal consumption expenditure attempt to adjust for quality changes. For example, personal computer price indexes decline dramatically more than the actual purchase price of these machines reflecting the fact that they are faster and have more memory and storage than earlier products. These quality adjustments may reflect that new computers are somehow "better" than their earlier counterparts, but higher quality is not reflected, other things equal, in expenditure.

Our model identifies three components of demand that drive output: induced consumption, induced business investment, and autonomous demand. Business investment adds directly to productive capacity and therefore our business investment variable excludes residential construction and inventory changes. The definition of autonomous demand is central to the logic of the FFV model we study, but there are ambiguities in defining what is indeed "autonomous." We consider several measures to explore the robustness of results to this definition.

All definitions of autonomous demand include consumption and investment spending by federal, state, and local governments (Allain, 2015). In addition, all our autonomous demand definitions include government spending on social programs. In the U.S., this spending is almost entirely for health care (Medicare and Medicaid). In the national accounts, this spending is treated as a transfer to the private sector and then added, dollar-for-dollar, to personal consumption expenditure. But it is clear this spending is not induced by cash income flows to households. These spending categories are large injections of demand by government and are just as "autonomous" as other parts of government spending .

A second definition of autonomous demand adds Social Security, government transfer payments to retirees and the disabled. These payments are surely autonomous, in the sense they are steady, defined-benefit expenditures that do not depend on the current state of the economy. Ambiguity arises, however, because Social Security income is not necessarily spent and therefore does not add dollar-for-dollar to demand like government health care spending. However, Social Security is large and much of it goes to households of limited means. Most of this spending likely supports demand and that demand is largely autonomous.

The third definition of autonomous demand adds exports (Nah and Lavoie, 2017). For almost any country, exports are likely independent of domestic economic shocks. The U.S. is large enough, however, that its own exports could be affected by domestic shocks because those domestic shocks could spill over in non-trivial ways to its trading partners. Nonetheless, there is a good case that even US exports are largely autonomous, especially at lower data frequencies.

The fourth and final definition of autonomous demand adds residential construction. Some author's have argued that residential investment is driven fundamentally by credit flows and the accumulation of debt; hence, especially in recent U.S history, they have been largely detached from current output (Fiebiger, 2018; Pérez-Montiel and Pariboni, 2021). The same authors have also argued that, in the long-run, the growth rate of residential investment must be determined by the rate of population growth, which they presume as exogenous. As is well-known, however, both Malthusian and modern theories of fertility assume that population growth is dependent on the state of the economy (Barro and Becker, 1989).

Supply-side data come from the labor market. We use time series on the unemployment rate and employment, measured as total workers. Labor productivity is output divided by employment.

4.1.2 A Minimum Distance Estimator

The model described in section 3 poses numerous challenges for estimation. Because the dynamics are described by a system of non-linear stochastic difference equations, it is not possible to solve for a closed-form expression that links the endogenous variables to the lags of themselves and the stochastic term. Additionally, the model contains variables likely measured with error, including the capital stock, or which are unobserved, such as agents' expectations.

We address this issue by using a minimum distance estimator, in particular, the Simulated Method of Moments (SMM) estimator proposed originally by McFadden (1989) and Pakes and Pollard (1989) and extended to time-series models by Lee and Ingram (1991) and Duffie and Singleton (1993). To understand the SMM estimator consider, a fully-specified model with parameter vector Θ, of dimension q. The econometrician has a sample y_t of weakly stationary data, used to compute sample moments $T^{-1}\sum_t m(y_t)$, such as correlations between variables in the data. The economic model defined by Θ is used to simulate a data vector S and moments $m(y_s(\Theta))$, then a natural way to estimate the model is to minimize:

$$\min_{\theta} Q(\Theta) = M(\Theta)'WM(\Theta) \tag{38}$$

Where W is a positive-definite weighting matrix, and

$$M(\Theta) = T^{-1}\sum_{t=1}^{T} m(y_t) - S^{-1}\sum_{s=1}^{S} m(y_s(\Theta)) \tag{39}$$

That is, the vector of estimated parameters minimizes a measure of the distance between the moment generated by the model and those present in the data. As long as the moment conditions are correctly specified at the true parameter vector, and the vector of data y_t is weakly stationary, the estimator is consistent and asymptotically normal.

This estimator is powerful in our context: it allows consistent estimation when some variables are measured with error or unobserved, it does not require the economic model to have an explicit solution, only that it can be simulated, and it allows for model misspecification: we don't require all the implications of the model to be true, only those moment conditions which will be used for estimation. This last feature is an attractive feature of this approach, since we do not claim our model is sufficiently rich to be a full theory of growth and fluctuations; however, we do argue that it captures important transmission mechanisms of aggregate demand to the rest of the economy.

We estimate the full vector of parameters implied by the model, with the exception of g^* and δ). The growth rate of autonomous demand is set to equal the average quarterly growth rate of output over our sample, $g = 0.03/4$. In addition, we fix $\delta = 0.025$ at the quarterly level, as is conventional in much of the literature. The rest of the parameter vector $\Theta = (\sigma, \rho, s, \lambda, \alpha, v^*, \theta_0, \theta_1, \phi_0, \phi_1, \phi_2)$ is estimated by matching 23 moments, which contain information about business cycle fluctuations and long-run growth. To match business-cycle fluctuations, we use the cyclical correlation between output and consumption, investment, autonomous demand, unemployment, employment and labor productivity; the cyclical persistence of these variables, and the relative volatility between output and the remaining variables. To match the long-run behavior of the data, the moments include the average consumption share in output, which identifies s, the average investment share in output, which identifies v^* conditional on g and δ, the average unemployment rate, and the average growth rate of labor productivity.

Because SMM requires stationary y_t, we obtain business-cycle statistics by extracting frequencies between 6 and 32 quarters with the Baxter-King (1999) filter on both actual and simulated data. We do not apply this filter to the long-run moments, we simply compute sample averages of these ratios and growth rates.[10] Since the model has implications for the data at all frequencies, we also report estimates extracting frequencies at 32-80 quarters and between 6-80 quarters. These estimates correspond to what Comin and Gertler (2006) define as the 'medium-term' business cycles, and is a natural benchmark to understand how endogenous productivity and other dynamics behave in our environment over longer horizons.

Because our objective function is highly non-linear, we obtain point estimates numerically with the calibrated parameters from FFV (2020) as starting values. We obtain confidence intervals of our parameters and non-linear functions of our parameters by Monte Carlo simulations, simulating the path induced by a stochastic demand shock 1,000 times, and for each path, we obtain an estimator. The confidence intervals reported below correspond to the relevant percentiles of the empirical distribution function of these simulated estimates.

5 Empirical Results

5.1 Parameter Estimates and Second Moments

Table 1 presents our baseline estimation results, where we display the model estimated at business-cycle frequencies, the 'medium-run' frequencies, and both frequencies taken together. It is instructive to compare the results obtained here with those obtained in FFV through calibration and single-equation, reduced-form regressions. Starting with the demand side, the values for the savings rate and the capital-output ratio are remarkably close. Note that the FFV model was calibrated using annual data, while we have estimated the model using quarterly data; therefore, we multiply or divide the parameters by 4 where appropriate. The annualized value of the capital output ratio varies between 0.99 and 1.04 depending on

[10] Note that assuming these shares are stationary in the literature, or more stationary than variables in levels, is consistent with a wide class of growth models that imply a balanced growth path, including the one presented here. Stock and Watson (1999) present evidence to support this approach, although later studies have argued there are small trends in these series.

the frequency; while FFV calibrated a value of 1.2. The main differences are for the value of λ and α; our value for λ is not statistically different from 0 at the 5% significance level; however, the confidence intervals are wide. They include a value of $\lambda = 0.10$, which would imply 34% of adjustment between v_t and v^*. This implies that our estimates are not very informative when it comes to λ, a point to which we return below. The point estimate for the adaptive expectation parameter $\alpha = 0.94$, implies a yearly value[11] of 0.784, which implies expectations are less persistence than what was found in FFV, where $\alpha = 0.9$.

Table 1: Baseline Parameter Estimates

	6-32 Quarters		32-80 Quarters		6-80 Quarters	
Parameter	Coef.	C.I	Coef.	C.I	Coef.	C.I
σ	0.030	[0.003, 0.040]	0.040	[0.001, 0.040]	0.032	[0.001, 0.040]
ρ	0.989	[0.688, 0.990]	0.985	[0.081, 0.990]	0.990	[0.616, 0.989]
θ_0	0.008	[0.006,0.009]	0.008	[0.006,0.011]	0.010	[0.006,0.011]
θ_1	0.073	[0.053, 0.106]	0.087	[0.070, 0.140]	0.113	[0.076, 0.137]
ϕ_0	0.020	[0.012,0.024]	0.015	[0.006,0.020]	0.019	[0.014,0.024]
ϕ_1	0.261	[0.227, 0.324]	0.193	[0.162, 0.266]	0.258	[0.223, 0.328]
ϕ_2	0.0006	[0.000,0.197]	0.001	[0.000,0.247]	0.004	[0.000,0.104]
v^*	4.113	[3.508, 4.736]	3.982	[3.513, 4.723]	4.158	[3.525, 4.610]
s	0.441	[0.428, 0.500]	0.400	[0.391, 0.466]	0.419	[0.405, 0.475]
λ	0.003	[0.000, 0.132]	0.003	[0.000, 0.141]	0.008	[0.000, 0.131]
α	0.941	[0.925, 0.990]	0.950	[0.932, 0.996]	0.948	[0.926, 0.990]
J statistic	50.5		52.9		53.9	

There are also important differences in the supply side. The calibration in FFV along with the single-equation estimation methods suggested that hysteresis in labour supply and productivity where equal in importance. The SMM estimates of the full model yield a productivity effect ϕ_1 roughly three times greater than the labour supply effect θ_1. In this sense, hysteresis in labor productivity is quantitatively much more important than hysteresis in labor supply to explain the adjustment of the supply side of the economy to movements in autonomous demand. Another important difference is that the effect of the accumulation rate on productivity growth ϕ_2 is not statistically different from 0. As discussed before, this does not imply that the Kaldor-Verdorn coefficient in this economy is 0, but rather that the effect of the level of economic activity on growth is more important than the feedback of effect of growth on growth.

[11] To annualize these estimates, we use the formula $0.941^4 = 0.784$

It's interesting to note that to derive realistic fluctuations, the volatility of autonomous demand must be rather large. The process itself also needs to be very persistent to match the data, which means the propagation mechanisms inside the model are not very strong. All of our estimates are very similar across different frequencies; perhaps the most significant difference is that the hysteresis parameter on labor supply is 0.11 when using both business cycle and medium run frequencies, while it drops when only using a subset of frequencies. Due to the similarity of the estimates at different frequency ranges, in what follows we will use our estimates at business cycle frequencies to compare the moments predicted by the model to those found in the data, compute impulse-response functions, and compute some interesting non-linear functions of our parameter estimates.

To assess how well the model performs along various dimensions, tables 2 and 3 compare the business-cycle and growth statistics estimated from the model with those in the data.

Table 2: Business Cycle Statistics

	Correlations		Persistence		Volatility	
Variable	Data	Model	Data	Model	Data	Model
lnY	1.00	1.00	0.93	0.96	1.00	1.00
lnC	0.90	0.96	0.95	0.96	1.20	1.05
lnI	0.74	0.83	0.94	0.95	3.04	3.26
lnF	0.58	0.62	0.92	0.90	0.83	0.85
U	-0.83	-0.57	0.94	0.93	0.50	0.62
lnL	0.77	0.77	0.94	0.93	0.68	0.68
lnA	0.74	0.74	0.91	0.97	0.65	0.65

Our rather parsimonious model (just 11 parameters for 23 moments) matches all the business cycle statistics qualitatively, and it also matches them reasonably well quantitatively. The model generates the high persistence in all macroeconomic variables present in the data, it matches the large correlations between employment and output, and labor productivity and output, rather well. It also matches the fact that investment is highly volatile over the cycle, as well as matching closely the relative volatilities of employment and labor productivity relative to output.

The dimensions along which the model performs less well compared with the data are all related to the unemployment rate: the correlation between unemployment and output is 30%

below what we find in the data. Likewise, the volatility of unemployment estimated from model is somewhat higher than in the data. The volatility of consumption and the correlation between investment and output are also roughly 10% different from the data. The other differences between the model and the data are well below 10%. Somewhat paradoxically, although the moments concerning the unemployment rate match the data less well t, the moments concerning labor productivity and employment display an excellent match. These moments are arguably the most important in our study because they determine the critical transmission from demand shocks to the supply side.

Table 3: Great Ratios and Growth Rates

Variable	Data	Model
C/Y	49%	56%
I/Y	13%	14%
u	6.0%	6.0%
g_A	0.4%	0.4%

Table 3 displays our great ratios and growth rates of key variables used for estimation. The model does an excellent job of matching the average unemployment rate over the sample; the share of investment in output, and the average growth rate of productivity. The share of consumption in total output, however, is matched less successfully. Recall that we do not filter these moments, since we presume that these great ratios are stationary along the balanced growth path. The same holds for the growth rate of productivity and the unemployment share.

5.2 Identification

As is well know, non-linear models are often plagued by identification issues; in other words, there could be multiple linear combinations of our parameters that deliver the same match between the model and data. To address this issue, we check conditions for local identification. The Cragg-Donald (1993) statistic checks that the Hessian of our objective function is of full rank:

$$\text{rank}\left[E\left(\frac{\partial m(x_s(\theta))}{\partial\theta}\right)\right] = q \tag{40}$$

Hence, the Hessian is simply the Jacobian of our simulated moments with respect to the parameter estimates. Around our parameter estimates the matrix is indeed full rank, which demonstrates that our model is locally identified.

In addition to checking local identification, we can get a sense of how strong the identification is by plotting the percentage change of the objective function against the percentage change in parameter values relative to the point estimates. The graphs in figures 2 and 3 show substantial curvature around most parameter values; in particular, the hysteresis parameters on productivity and labor supply are sharply identified. Two parameters seem to be locally unidentified: the embedded technical progress coefficient (σ_2), and the adjustment speed of the capital stock (λ). This outcome is not surprising, given that we estimate these two parameters to be very close to zero, which means the point estimates are on the boundary of the parameter space. As we discussed earlier, our confidence intervals for λ were very wide and were compatible with the capital-output ratio being either a random wak or it having rather fast adjustment. It would thus seem that the model is misspecified along this dimension, and hence there is some need to go back to the drawing board. Also, the persistence of the demand shock parameter, though locally identified, has another local minimum between 0.93 and 0.96. Overall, identification issues that plague many large, dynamic, nonlinear stochastic models (see Canova and Sala, 2009) are not so pervasive in our application.

Figure 2: Identification: Autonomous Demand and Supply Side

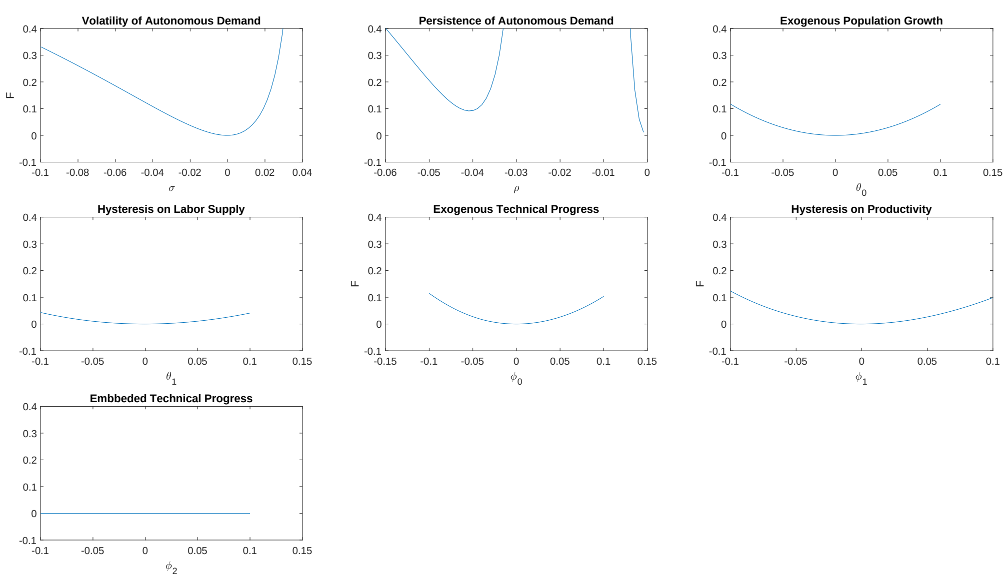

Figure 3: Identification: Demand Side

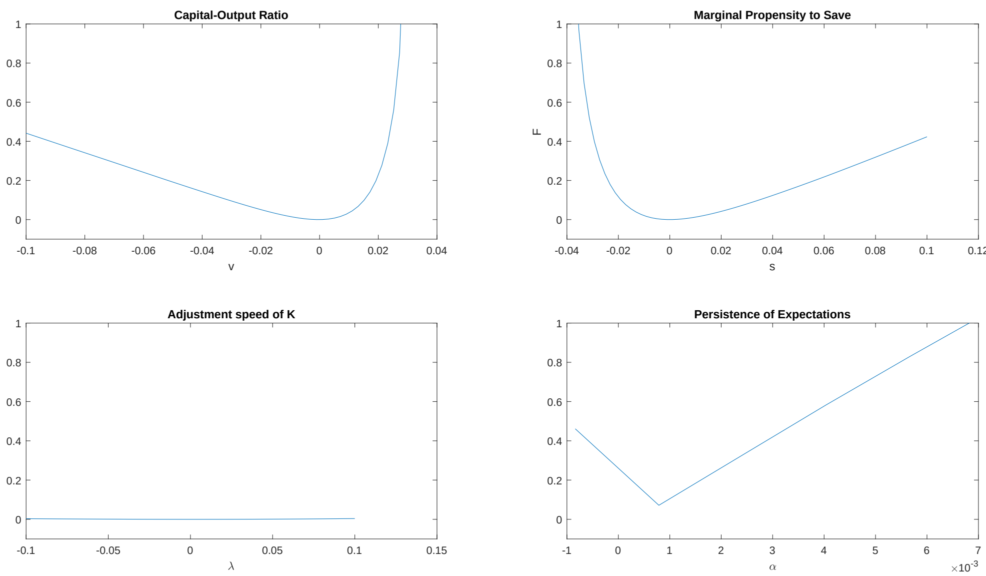

5.3 Hysteresis Effects and the Growth Ceiling

A key question in research on demand-led growth models is the extent to which demand can lead supply, that is, how strong are the hysteresis effects? In our context the answer to this question depends on in the long-run relationship between the growth rate of autonomous demand and the steady-state rate of unemployment. In particular, total differentiation of equation (15) gives:

$$\frac{du^*}{dg^*} = -\frac{(1-\phi_2)}{\theta_1 + \phi_1} \tag{41}$$

A plug-in estimator for this quantity can be obtained by replacing the parameters with our estimated values. Table 4 shows the results for this exercise for our baseline estimates, along with confidence intervals. The point for quarterly data is -3.05; an annualised estimate is obtained by dividing by 4, which gives -0.76, remarkably similar to the calibrated value of -0.7 obtained by FFV (2020). Furthermore, FFV reports a rather wide range of values for this effect (-0.3 to -1.7) by simulating across a plausible range of the parameter space. Our statistical estimate of the 95% confidence intervals reported in table 4 are much tighter. This result strongly supports the hypothesis of demand-led growth followed by accommodating supply. Parameter estimates that are too high would imply an excessive change in unemployment would be necessary for supply to accommodate a modest change in the autonomous demand growth rate. However, if the estimate is too low it would imply that supply is excessively sensitive to demand. With our estimates, if yearly autonomous demand growth were to increase by one percentage point, the unemployment rate would decline by 0.76 percentage points, which would bring down the sample average from 6% TO 5.2%. We believe this effect reflects a plausible quantitative magnitude. It implies that there may be substantial potential faster demand growth to allow the U.S economy to grow above what has been achieved historically.

Table 4: Estimates of Hysteresis Effect

	Coef.	95% C.I
6-32 Quarters	-0.75	[-0.83,-0.57]
32-80 Quarters	-0.89	[-0.99,-0.61]
6-80 Quarters	-0.67	[-0.91,-0.53]

As we discussed above, given our parameter estimates for the hysteresis effects, we can compute the Verdoorn coefficient implied by our model, β in the notation of equation (34) and (37). This yields:

$$\hat{\beta} = \frac{\hat{\phi}_1 + \hat{\phi}_2\hat{\theta}_1}{\hat{\theta}_1 + \hat{\phi}_1} = 0.786 \tag{42}$$

With confidence intervals in the range (0.72,0.83). This point estimate is slightly higher than what is commonly found in the literature that uses reduced-form methods with estimates ranging between 0.3 and 0.6 (Basu and Budhiraja, 2021). Importantly, we show the effect of capital accumulation on labor productivity (ϕ_2) - a key ingredient of the early formulations of Kaldor's technical progress function (Kaldor 1957, 1961) - has a point estimate near zero, with the upper end of the confidence interval still reasonably small near 0.2. This is the case even though there is substantial induced technical progress arising from the feedback from the unemployment rate to the growth rate of supply. Note that this outcome is also true at all frequencies - both across the 1.5 and 8 years that are usually thought as the typical duration of the business cycle, and in periods between 8 and 20 years, where one might expect to see some effect of capital accumulation on technical progress.

The fact that hysteresis effects are of a reasonable magnitude implies that the supply side adjusts in response to higher demand growth. But how much growth can the supply side actually accommodate? What growth rates are feasible in our economy? In our environment, the upper bound on growth, or the 'ceiling' growth rate is given when the economy reaches full employment. In particular, since we know that $u_t > 0$, we can use our steady-state condition to derive the following relationship:

$$\bar{g}_z = \frac{\phi_0 + \theta_0 + \theta_2 \delta}{1 - \phi_2} \tag{43}$$

Where $\bar{g}_z$ now denotes the ceiling, or the growth rate of output consistent with full employment. Naturally, it's rather extreme identify full employment with literally an unemployment rate of zero. Suppose instead, that there is a minimum rate of unemployment that the economy can reach, $\hat{u}$. In this case, the growth ceiling will be given by:

$$\bar{g}_z = \frac{\phi_0 + \theta_0 + \theta_2 \delta - \hat{u}(\phi_1 + \theta_1)}{1 - \phi_2} \tag{44}$$

Table 5 shows the point estimates of the growth ceiling along with confidence intervals for different values of the possible 'minimum' unemployment rates, again, we report annualised values. Our confidence intervals are surprisingly tight in each case.

Table 5: Estimates of the growth ceiling

$\hat{u}$	$\bar{g}_z$	95% C.I
0	0.112	[0.097,0.132]
0.02	0.085	[0.073,0.097]
0.04	0.059	[0.049,0.063]

To get a sense of how reasonable these estimates are, we plot in Figure 4 below the annualised quarterly growth rate of output, along with the confidence interval for the ceiling assuming the minimum unemployment rate is 2%. We choose this minimum unemployment as a compromise between what we saw as an exceedingly unrealistic lower bound (a 0% unemployment rate), and a minimum value that has been reached multiple times over the sample, and that implies a maximum growth rate that would have been surpassed many times during our sample (4%). A minimum unemployment rate of 2%, however, would imply that the U.S economy only reached the ceiling on four occasions during the whole sample: During 1959Q1-1959Q2, 1965Q4-1966Q2, 1973Q1 and then finally during 1983Q4-1984Q2.

Naturally, there a are a number of important caveats with these estimates. First, our model does not provide a theory of the minimum unemployment rate therefore there is substantial uncertainty over the actual growth ceiling. Second, the intercepts and hysteresis parameters governing productivity growth and labor supply growth could be time-varying. In particular, as figure 4 shows, peaks of the cyclical growth rate have declined since the middle 1980s and the maximum feasible rate may have also declined as well. Third, the hysteresis linear

Figure 4: The growth ceiling

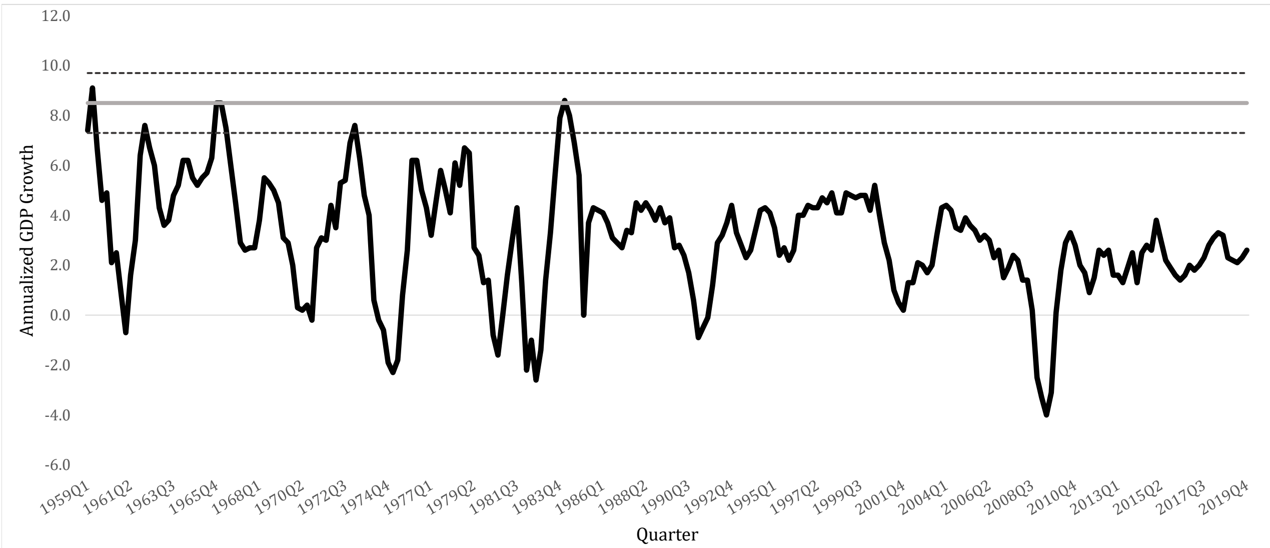

hysteresis relationships in our model could, in reality, be non-linear, particularly as the economy pushes up against a minimum unemployment constraint. Nevertheless, our estimates suggest that stronger demand growth - say, an increase of one percentage point - would not hit the growth ceiling, and hence could be accommodated with a lower rate of unemployment.

5.4 Untargeted Moments

We can compute a variety of other statistics of interest from our estimates. In particular, we can compute moments in the data that were not used in our parameter estimation - moments that were not targeted by our objective function - and ask how well the estimated model replicates them. In particular, we focus on two moments that have received widespread attention in the business cycle literature: the correlation between labor productivity and employment, and the correlation between unemployment and labor productivity. Matching the cyclical co-movement between employment and labor productivity has been perceived, in the words of Christiano and Eichenbaum (1992),*"a litmus test for aggregate economic models"*. The old vintage of Keynesian models proposed between 1930 and 1970, which invoked wage rigidities to explain the effects of aggregate demand shocks, predicted counter-cyclical real wages, while the data show real wage are either a-cyclical or weakly pro-cyclical. In contrast, standard RBC models usually imply that the real wage is strongly pro-cyclical, since technology shocks shift labor demand outward along an upward-sloping labor supply curve, increasing employment and wages. In any economic model where real wages are proportional to average labor productivity, as is the case with competitive neoclassical factor markets under a Cobb-Douglas production technology, these results carry over to the cyclical behavior of output and labor productivity. A large literature in the 1990s extended the standard RBC model to match the cyclical behavior of average labor productivity and real wages with moderate success; prominent examples that also summarize the research in this period are Hansen and Wright (1992) and Christiano and Eichenbaum (1992).

The correlation between unemployment and labor productivity has also received prominent attention in recent decades, starting with an observation by Shimer (2005) that search and matching labor markets in the tradition of Diamond-Mortensen-Pissarides usually produce a correlation between unemployment and labor productivity equal to -1, while this correlation is much weaker in the data. This observation has led to a rich literature extending or modifying the baseline model to account for this observation (Hall, 2005; Pissarides, 2009; Gertler and Trigari, 2009, among others). It thus seems natural to ask whether our model produces reasonable predictions for these correlations, and it hence describes the joint dynamics of productivity and unemployment better. Table 6 presents the results for both moments from our estimates and prominent estimates from real business cycle models.

Table 6: Untargeted Moments - Neoclassical vs Post Keynesian

Moment	Data	Neoclassical	Post-Keynesian
$\sigma_{A,L}$	0.14	0.57*	0.14
$\sigma_{A,u}$	-0.28	-0.95**	0.12

*=Christiano & Eichenbaum (1992); **=Shimer (2005)

The correlations produced by our model regarding employment and labor productivity are remarkably close to those implied by the data. Indeed, they are much closer than the correlation of one implied by a standard RBC model, but it also performs better than the various extensions proposed by Christiano and Eichenbaum (1992) and Hansen and Wright (1992).[12] These results carry over to some extent to the correlation between unemployment and labor productivity: a correlation of 0.12 is much close to -0.28 than -1. Nevertheless, there is still some room for improvement along this dimension.

[12]In those papers, the lower end of the correlations reported in tables 4 and 3, respectively for each paper, are of 0.57 and 0.49. Hence, our model outperforms these specifications as well.

The reason why our model produces an almost zero correlation between employment and labor productivity is simple. Suppose the economy is hit by a demand shock. Upon impact, production expands, which increases employment and hence decreases unemployment. This decrease in unemployment accelerates productivity growth (and it's level) in the next period: it's lagged unemployment that affects productivity growth, not contemporaneous unemployment. A natural interpretation for this is that newly employed workers must take some time to learn before they increase their productivity, as was suggested in our simple model.

5.5 The effects of autonomous demand shocks

A final exercise of interest is to compute the path of the endogenous variables in response of a demand shock; or in other words, to compute Impulse-Response Functions (IRFs) of the model at the estimated parameter values. The IRFs summarise the transitional dynamics of the model after being hit with an unexpected demand shock, and provide a way to think about the business-cycle co-movements after a demand shock. In addition, we can compare our IRFs to those estimated through reduced-form methods; in particular, to those obtained through local projections by Girardi, Paternesi and Stirati (2020). This provides another informal test of our estimated parameters.

To calculate impulse response functions, we proceed as follows. We initialise the model at the steady state, and compute a path of the vector of endogenous variables, $\mathbf{Y}$, in absence of shocks. We then initialise the model again at the steady state, and we shock the the first period, i.e, $\varepsilon_1^D = \delta$ and $\varepsilon_t^D = 0 \quad \forall t > 1$, and compute the path of the endogenous variables. We then subtract both vectors; $E[\mathbf{Y}_{t+h}|\varepsilon_1^D = \delta] - E[\mathbf{Y}_{t+h}|\varepsilon_1^D = 0] = IRF_h$. The size of the shock is δ; which we specify to be 5% of autonomous demand, consistent with the size of a large demand shock as defined by Girardi et. al (2020).

Figures 5 and 6 reproduce the response of the demand and supply side of the economy, respectively, to an autonomous demand shock. The figures also plot 68% and 95% confidence intervals. The peak effect of autonomous demand on output, consumption and investment is reached roughly at 8 quarters, after that, the effect monotonically declines and reaches 0 after 10 years. Given that all estimates are hump-shaped, this shows that autonomous demand shocks can generate the boom-and-bust patterns that are characteristic of business cycles. It also implies that all components of induced demand and output co-move strongly across horizons. Note that Girardi et. al find somewhat more persistent demand shocks data compared to our estimated model. Differences could be due to the fact that Girardi et al. compute average IRFss for a cross-section of countries.

Figure 5: IRF of Demand Side

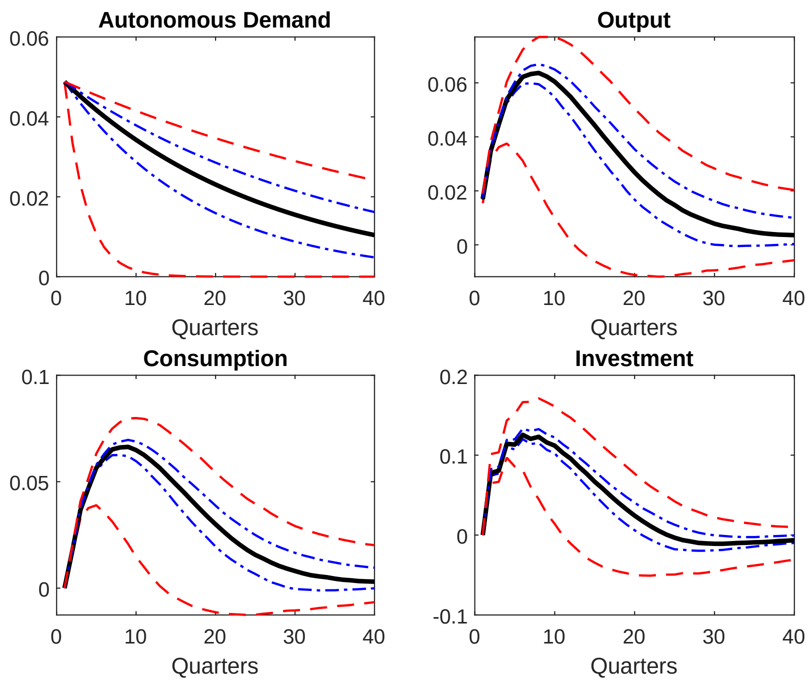

The response of the supply side also mimics the hump-shaped response that we find in the demand side. The peak response of unemployment and employment, however, come sooner, roughly 2 to 4 quarters after the demand shock hits. Note that the unemployment rate overshoots its steady-state value slightly after 10 quarters, and the economy experiences some slightly elevated unemployment before decreasing monotonically to the steady-state. The response of labor productivity is zero upon impact; however, as unemployment decreases sharply, demand-induced technical change starts to emerge slowly and the peak response occurs roughly at 10 quarters, which coincides with the peak response of the demand side. After that, it decreases monotonically to the steady-state. All in all, the supply side reacts in a way that also mimics business-cycle co-movements after a demand shock. It's interesting to note that our IRF for unemployment has the same qualitative shape than the one for the long-term unemployed estimated by Girardi et. al (2020). Interestingly, we find a stronger supply-side response, evaluated at the peak value of the IRF's, than what was found by Girardi et al. Our peak response is also quicker than what they found using international data.

Figure 6: IRF of Supply Side

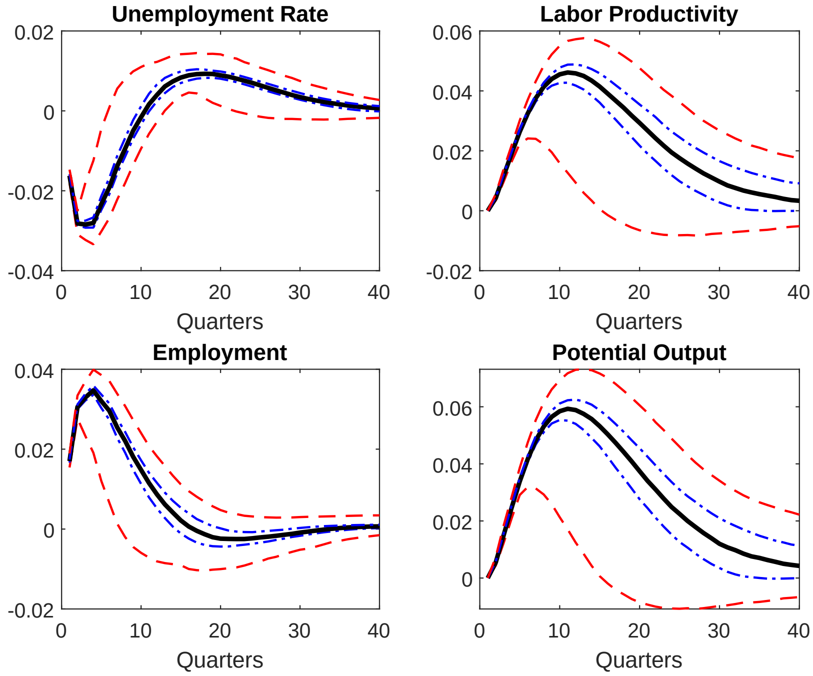

It should be clarified that the fact that the effects of the autonomous demand shock dies out does not mean that the supply side doesn't respond to the demand side in the long run. The IRFs answer the question of what happens if autonomous demand rises unexpectedly for a period and then return to it's steady-state value. Hence, it summarises the transitional dynamics of the economy in response to a demand shock. If the question is: what is the effect on the supply side of a permanent 1% increase in autonomous demand growth? Then this is answered by our hysteresis effect parameters, presented in table 4 above.

6 Conclusion

This paper presents a simple Post-Keynesian growth model where the growth rate of autonomous demand determines the growth rate of output and the steady-state level of unemployment, and it asks whether it can reproduce the most salient features of the U.S business cycles, as measured by the second moments of seven key macroeconomic variables, and the response of the economy to substantial autonomous demand shock. We find that the model performs at least as well as a real business cycle model with search and matching frictions along these dimensions, and hence, there is no immediate *empirical* reason to dismiss the model as a promising candidate to explain growth and business cycles in mature, capitalists economies. We hope this stimulates further refinement and extensions of this Keynesian growth framework.

7 Bibliography

Aghion, P., Howitt, P. (1994). Growth and unemployment. The Review of Economic Studies, 61(3), 477-494.

Allain, O. (2015). Tackling the instability of growth: a Kaleckian-Harrodian model with an autonomous expenditure component. Cambridge Journal of Economics, 39(5), 1351-1371.

Anzoategui, D., Comin, D., Gertler, M., Martinez, J. (2019). Endogenous technology adoption and RD as sources of business cycle persistence. American Economic Journal: Macroeconomics, 11(3), 67-110.

Anzoategui, D., & Kim, M. (2021). Re-estimating potential GDP: new evidence on output hysteresis. Mimeo.

Barro, R. J., Becker, G. S. (1989). Fertility choice in a model of economic growth. Econometrica: journal of the Econometric Society, 481-501.

Basu, D., & Budhiraja, M. (2021). What to make of the Kaldor-Verdoorn law?. Cambridge Journal of Economics, 45(6), 1243-1268.

Baxter, M., & King, R. G. (1999). Measuring business cycles: approximate band-pass filters for economic time series. Review of economics and statistics, 81(4), 575-593.

Benati, L., & Lubik, T. (2021). Searching for hysteresis.

Benigno, G., Fornaro, L. (2018). Stagnation traps. The Review of Economic Studies, 85(3), 1425-1470.

Brown, J. R., Fazzari, S. M., & Petersen, B. C. (2009). Financing innovation and growth: Cash flow, external equity, and the 1990s RD boom. The Journal of Finance, 64(1), 151-185.

Cette, G., Fernald, J., & Mojon, B. (2016). The pre-Great Recession slowdown in productivity. European Economic Review, 88, 3-20.

Comin, D., & Gertler, M. (2006). Medium-term business cycles. American Economic Review, 96(3), 523-551.

Christiano, L. J., & Eichenbaum, M. (1992). Current real-business-cycle theories and aggregate labor-market fluctuations. The American Economic Review, 430-450.

Cragg, J. G., & Donald, S. G. (1993). Testing identifiability and specification in instrumental variable models. Econometric Theory, 9(2), 222-240.

Duffie, D., Singleton, K. J. (1993). Simulated moments estimation of Markov models of asset prices. Econometrica, 61, 929-952.

Eggertsson, G. B., Mehrotra, N. R., Robbins, J. A. (2019). A model of secular stagnation: Theory and quantitative evaluation. American Economic Journal: Macroeconomics, 11(1), 1-48.

Boserup, Ester (1965). The conditions of agricultural growth: the economics of agrarian change under population pressure. London: Allen & Unwin. OCLC 231372

Fazzari, S. M. (2020). Was Keyneian economics ever dead? If so, has it been resurrected? Review of Keynesian Economics, 8(1), 46-60.
Fazzari, S. M., Ferri, P. E., Greenberg, E. G., Variato, A. M. (2013). Aggregate demand, instability, and growth. Review of Keynesian Economics, 1(1), 1-21.

Fazzari, S. M., Ferri, P., & Variato, A. M. (2020). Demand-led growth and accommodating supply. Cambridge Journal of Economics, 44(3), 583-605.

Fernald, J. G. (2015). Productivity and Potential Output before, during, and after the Great Recession. NBER macroeconomics annual, 29(1), 1-51.

Fiebiger, B. (2018). Semi-autonomous household expenditures as the causa causans of postwar US business cycles: the stability and instability of Luxemburg-type external markets. Cambridge Journal of Economics.

Freitas, F., & Serrano, F. (2015). Growth rate and level effects, the stability of the adjustment of capacity to demand and the Sraffian supermultiplier. Review of Political Economy, 27(3), 258-281.

Furlanetto, F., Lepetit, A., Robstad, Ø., Rubio Ramírez, J., & Ulvedal, P. (2021). Estimating hysteresis effects.

Gertler, M., & Trigari, A. (2009). Unemployment fluctuations with staggered Nash wage bargaining. Journal of political Economy, 117(1), 38-86.

Girardi, D., Paternesi Meloni, W., & Stirati, A. (2020). Reverse hysteresis? Persistent effects of autonomous demand expansions. Cambridge Journal of Economics, 44(4), 835-869.

Girardi, D., Pariboni, R. (2016). Long-run effective demand in the US economy: an empirical test of the sraffian supermultiplier model. Review of Political Economy, 28(4), 523-544.

Girardi, D., Pariboni, R. (2020). Autonomous demand and the investment share. Review of Keynesian Economics, 8(3), 428-453.

Góes, M. C. B., & Deleidi, M. (2022). Output determination and autonomous demand multipliers: An empirical investigation for the US economy. Economic Modelling, 116, 106004.

Haluska, G., Braga, J., Summa, R. (2021). Growth, investment share and the stability of the Sraffian Supermultiplier model in the US economy (1985–2017). Metroeconomica, 72(2), 345-364.

Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. American economic review, 95(1), 50-65.

Hansen, G. D., & Wright, R. (1992). The Labor Market in Real Business Cycle Theory. Federal Reserve Bank of Minneapolis. Quarterly Review-Federal Reserve Bank of Minneapolis, 16(2), 2.

Hein, E. (2014). Distribution and growth after Keynes: A Post-Keynesian guide. Edward Elgar Publishing.

Hein, E., Tarassow, A. (2010). Distribution, aggregate demand and productivity growth:

theory and empirical results for six OECD countries based on a post-Kaleckian model. Cambridge Journal of Economics, 34(4), 727-754.

Kaldor, N. (1957). A model of economic growth. The economic journal, 67(268), 591-624.

Kaldor, N. (1961). Capital accumulation and economic growth. In The theory of capital (pp. 177-222). Palgrave Macmillan, London.

Lee, B. S., & Ingram, B. F. (1991). Simulation estimation of time-series models. Journal of Econometrics, 47(2-3), 197-205.

Lavoie, M. (2014). Post-Keynesian economics: new foundations. Edward Elgar Publishing.

Lavoie, M. (2016). Convergence Towards the Normal Rate of Capacity Utilization in Neo-K aleckian Models: The Role of Non-Capacity Creating Autonomous Expenditures. Metroeconomica, 67(1), 172-201.

Maffei-Faccioli, N. (2021). Identifying the sources of the slowdown in growth: Demand vs. supply (No. 9/2021). Working Paper.

McCombie, J. S., & Spreafico, M. R. (2015). Kaldor's 'technical progress function'and Verdoorn's law revisited. Cambridge Journal of Economics, 40(4), 1117-1136.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. Econometrica: Journal of the Econometric Society, 995-1026.

Modigliani, F. (1944). Liquidity preference and the theory of interest and money. Econometrica, journal of the econometric society, 45-88.

Nah, W. J., & Lavoie, M. (2017). Long-run convergence in a neo-Kaleckian open-economy model with autonomous export growth. Journal of Post Keynesian Economics, 40(2), 223-238.

Pakes, A., Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. Econometrica: Journal of the Econometric Society, 1027-1057.

Pérez-Montiel, J. A., Erbina, C. M. (2020). Autonomous expenditures and induced investment: a panel test of the Sraffian supermultiplier model in European countries. Review of Keynesian Economics, 8(2), 220-239.

Pérez-Montiel, J. A., & Manera, C. (2022). Is autonomous demand really autonomous in the United States? An asymmetric frequency-domain Granger causality approach. Metroeconomica, 73(1), 78-92.

Pérez-Montiel, J. A., & Pariboni, R. (2022). Housing is NOT ONLY the business cycle: A Luxemburg-Kalecki external market empirical investigation for the United States. Review of Political Economy, 34(1), 1-22.

Pissarides, C. A. (2009). The unemployment volatility puzzle: Is wage stickiness the answer?. Econometrica, 77(5), 1339-1369.

Serrano, F. (1995). Long period effective demand and the Sraffian supermultiplier. Contributions to Political Economy, 14(1), 67-90.

Stock, J. H., & Watson, M. W. (1999). Business cycle fluctuations in US macroeconomic time series. Handbook of macroeconomics, 1, 3-64.

Summers, L. H. (2014). US economic prospects: Secular stagnation, hysteresis, and the zero lower bound. Business economics, 49(2), 65-73.

Tavani, D., Zamparelli, L. (2018). Endogenous technical change in alternative theories of growth and distribution. Analytical Political Economy, 139-174.

A The effects of permanent demand shocks

This section derives fully equations (x) through (xx) in section 2. We start from the assumption that autonomous demand is a random walk in logs:

$$(1-L)\ln F_t = g^* + \varepsilon_t \tag{45}$$

Where L is the lag operator, i.e, $Ly_t = y_{t-1}$. We first derive the time series process for output; we pre-multiply equation (z) by $(1-L)$ to get:

$$(1-L)\ln Y_t = (1-L)(-\ln s_t + \ln F_t) \tag{46}$$
$$= g^* + \varepsilon_t \tag{47}$$

This shows that output is a random walk with a drift, and it follows exactly the same process as autonomous demand. We next derive the stochastic process for unemployment. Start from the log-differenced production function:

$$(1-L)\ln Y_t = (1-L)\ln A_t - (1-L)u_t \tag{48}$$

Replace inside the equations for learning by doing, as well as the stochastic process for output to get:

$$g^* + \varepsilon_t = \phi_0 - \phi_1 u_{t-1} - u_t + u_{t-1} \tag{49}$$

Re arrange in terms of the unemployment rate to get:

$$u_t = \phi_0 - g^* + (1-\phi_1)u_{t-1} - \varepsilon_t \tag{50}$$

This shows that the unemployment rate is an AR(1) process; hence, permanent demand shocks have some persistent, but finite lived, effects on the unemployment rate. Finally, we solve out for by writing the process for the unemployment rate as:

$$[1-(1-\phi_1)L]u_t = \phi_0 - g^* - \varepsilon_t \tag{51}$$

Now take the learning by doing equation in growth from and pre-multiply it by $(1-L\phi_1)$ to get:

$$[1-(1-\phi_1)L]g_{A,t} = [1-(1-\phi_1)L](\phi_0 - \phi_1 u_{t-1}) \tag{52}$$

$$g_{A,t} = (1-(1-\phi_1))\phi_0 + (1-\phi_1)g_{A,t-1} - \phi_1(\phi_0 - g^* - \varepsilon_{t-1}) \tag{53}$$
$$= \phi_0\phi_1 + (1-\phi_1)g_{A,t-1} - \phi_0\phi_1 + \phi_1 g^* + \phi_1\varepsilon_{t-1} \tag{54}$$
$$= \phi_1 g^* + (1-\phi_1)g_{A,t-1} + \phi_1\varepsilon_{t-1} \tag{55}$$
$$\tag{56}$$

This shows that log-productivity is an ARIMA(1,1,1) process. To prove that the new shock $u_t = \phi_1 \varepsilon_t$ has no serial correlation, note that since it was assumed that ε_t is i.i.d, it follows that $E[\varepsilon_t \varepsilon_{t-j}] = 0 \quad \forall j = 1, ..., T$. Therefore, $E[u_t u_{t-j}] = E[\phi_1^2 \varepsilon_t \varepsilon_{t-j}] = 0$, which completes the proof.