Wedges and Business Cycle Accounting

Kevin Kotzé

Table of contents

Introduction

- Most modern dynamic macroeconomic models include different frictions, adjustment costs, and shocks
- These represent deviations from the basic RBC framework
- But how do we know which of these are important?
- Chari, Kehoe, & McGrattan (2007) develop a methodology that they call "business cycle accounting" to provide insight into this question

Introduction

- They start with a basic one sector real business cycle model
- Then introduce four exogenous stochastic variables - called wedges
- These wedges are introduced as reduced form accounting devices
- Hence the wedges could emerge because of exogenous shocks, or because of some friction or adjustment cost

Introduction

- Through the lens of the basic RBC model they empirically measure the size of the four different wedges
- They then use this exercise to talk about:
  - which exogenous shocks are the most promising candidates for drivers of the business cycle
  - which kind of frictions and adjustment costs researchers ought to build into their models

Introduction

- The four wedges on which CKM focus are called the efficiency wedge, the labour wedge, the investment wedge, and the government consumption wedge
- The efficiency wedge is similar to a productivity shock - it simply measures TFP
- The government consumption wedge is the residual output component not explained by consumption or investment
- Therefore it is equivalent to government spending in a simple model
- It could also include net exports in an open economy model
- The labour and investment wedges are similar to distortionary tax rates on labour income and investment

Introduction

- In what follows we look at the basic features of the model and look to measure the wedges in the data
- These take the form of residuals from the first order conditions from the model
- Then we feed the observed wedges back into the prototypical model and do a variance decomposition of output and other variables
- To conclude we discuss some of the potential uses (and abuses) of this methodology

Business cycle accounting model

Household

- Consider what appears to be a standard RBC model in the household problem
- For simplicity, assume that there are no bonds available to the household
- This implies that the only way transfer resources intertemporally is through capital accumulation

Business cycle accounting model

Household

- The household faces the budget constraint:

$$C_t + (1+\tau_t^I)I_t \leq (1-\tau_t^N)w_t N_t + R_t K_t + \Pi_t - T_t$$

- Where τ_t^I is like a tax on investment
- Alters the relative price between consumption and investment
- Similarly τ_t^N is a tax on labour income
- We refer to $1-\tau_t^N$ as the labour wedge
- While $1+\tau_t^I$ is the investment wedge

Business cycle accounting model

Household

- The household takes Π_t as given and T_t is a lump sum tax
- Capital accumulates according to:

$$K_{t+1} = I_t + (1-\delta)K_t$$

- In this case, preferences are standard:

$$u(C_t, N_t) = \log C_t - \theta \frac{N_t^{1+\gamma}}{1+\gamma}$$

- The household discounts future utility flows by $0 < \beta < 1$

Business cycle accounting model

Household

- A Lagrangian for the household may take the form

$$
\begin{aligned}
\mathcal{L} = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \Bigg\{ & \left(\log C_{t+j} - \theta \frac{N_{t+j}^{1+\gamma}}{1+\gamma} \right) \\
\ldots + \lambda_{t+j} & \left[\left(1 - \tau_{t+j}^{N}\right) w_{t+j} N_{t+j} + R_{t+j} K_{t+j} \right. \\
& \ldots + \Pi_{t+j} - T_{t+j} - C_{t+j} - (1 + \tau_{t+j}^{I}) I_{t+j} \big] \\
& \ldots + \mu_{t+j} \left[I_{t+j} + (1-\delta) K_{t+j} - K_{t+j+1} \right] \Bigg\}
\end{aligned}
$$

Business cycle accounting model

Household

- The first order conditions are:

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial C_t} &= \lambda_t = \frac{1}{C_t} = 0 \\
\frac{\partial \mathcal{L}}{\partial N_t} &= \theta N_t^{\gamma} = \lambda_t (1 - \tau_t^N) w_t = 0 \\
\frac{\partial \mathcal{L}}{\partial I_t} &= (1 + \tau_t^I)\lambda_t = \mu_t = 0 \\
\frac{\partial \mathcal{L}}{\partial K_{t+1}} &= \mu_t = \beta \mathbb{E}_t[\lambda_{t+1} R_{t+1} + \mu_{t+1}(1 - \delta)] = 0
\end{aligned}
$$

Business cycle accounting model

Household

- Re-arranging so as to eliminate multipliers yields:

$$\theta N_t^{\gamma} = \frac{1}{C_t}(1 - \tau_t^N) w_t$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} \left[R_{t+1} + (1 + \tau_{t+1}^I)(1 - \delta) \right]$$

Business cycle accounting model

Firm

- The firm problem is standard
- Using a Cobb-Douglass production function for the output constraint

$$Y_t = A_t K_t^{\alpha} N_t^{1-\alpha}$$

- The optimal amounts of capital and labour are:

$$\begin{aligned} w_t &= (1-\alpha) A_t K_t^{\alpha} N_t^{-\alpha} \\ R_t &= \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \end{aligned}$$

- CKM refer to A_t as the efficiency wedge, which is just the standard productivity variable (or Solow residual)

Business cycle accounting model

Government

- Assume that the government chooses spending, G_t and the lump sum tax is chosen to balance the government's budget
- Therefore

$$T_t = G_t - \tau_t^I I_t - \tau_t^N w_t N_t$$

- In effect this would imply that the distortionary tax revenue is remitted back to the household

Business cycle accounting model

- Then the aggregate resource constraint is:

$$Y_t = C_t + I_t + G_t$$

- CKM refer to G_t as the government consumption wedge

Business cycle accounting model

Equilibrium

- Not counting the four exogenous wedges, the equilibrium of the model economy is then summarize by:

$$
\begin{aligned}
\theta N_t^{\gamma} &= \frac{1}{C_t}(1-\tau_t^N)w_t \\
(1+\tau_t^I)\frac{1}{C_t} &= \beta \mathbb{E}_t \frac{1}{C_{t+1}}\left[R_{t+1}+(1+\tau_{t+1}^I)(1-\delta)\right] \\
Y_t &= A_t K_t^{\alpha} N_t^{1-\alpha} \\
w_t &= (1-\alpha)A_t K_t^{\alpha} N_t^{-\alpha} \\
R_t &= \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \\
Y_t &= C_t + I_t + G_t \\
K_{t+1} &= I_t + (1-\delta)K_t
\end{aligned}
$$

Log-Linearisation

- Let's log-linearise the model about some non-stochastic steady state
- For the purposes of linearisation define the labour wedge as

$$\psi_t^N = (1 - \tau_t^N)$$

- And similarly the investment wedge is defined as

$$\psi_t^I = (1 + \tau_t^I)$$

Log-Linearisation

- The linearised labour supply curve is:

$$\gamma \tilde{N}_t = -\tilde{C}_t + \tilde{\psi}_t^N + \tilde{w}_t$$

- The linearised labour demand condition can be written:

$$\tilde{w}_t = \tilde{Y}_t - \tilde{N}_t$$

Log-Linearisation

- We get this by noting that the marginal product of labour can be written as proportional to the average product of labour
- Combining these provides:

$$\gamma \tilde{N}_t = -\tilde{C}_t + \tilde{\psi}_t^N + \tilde{Y}_t - \tilde{N}_t$$

- Or, simplifying further:

$$\tilde{\psi}_t^N = (1+\gamma)\tilde{N}_t + \tilde{C}_t - \tilde{Y}_t$$

Log-Linearisation

- Hence, given data on consumption, GDP, and labour input, as well as a value of γ, we can measure $\tilde{\psi}_t^N$ as the residual
- It measures the extent to which the deterministic labour FOC does not hold in the data
- To derive the linearised production function is straightforward and can be re-arranged to yield:

$$\tilde{A}_t = \tilde{Y}_t - \alpha \tilde{K} - (1 - \alpha) \tilde{N}_t$$

- In other words, the efficiency wedge can be measured as TFP

Log-Linearisation

- The linearised resource constraint can be written:

$$\tilde{G}_t = \frac{Y}{G}\tilde{Y}_t - \frac{C}{G}\tilde{C}_t - \frac{I}{G}\tilde{I}_t$$

- Hence, the government consumption wedge can be measured as the residual of output that is not accounted for by consumption and investment

Log-Linearisation

- Now let's linearise the Euler equation for capital
- This one is a bit more involved so we can start by taking logs

$$\log \psi_t^I - \log C_t = \log \beta - \log C_{t+1} + \log R_{t+1} + \psi_{t+1}^I (1-\delta)$$

- Totally differentiate:

$$\tilde{\psi}_t - \tilde{C}_t = -\tilde{C}_{t+1} + \frac{\beta}{\psi^*}[dR_{t+1} + (1-\delta) d\tilde{\psi}_{t+1}]$$

- This follows because $R + (1-\delta)\psi = \frac{\psi^*}{\beta}$

Log-Linearisation

- Writing this in tilde notation:

$$\tilde{\psi}_t^I - \tilde{C}_t = -\tilde{C}_{t+1} + \frac{\beta R^*}{\psi^*} \tilde{R}_{t+1} + (1-\delta)\beta \tilde{\psi}_{t+1}^I$$

- Note that we can write:

$$R^* = \psi \left(\frac{1}{\beta} - (1-\delta) \right)$$

Log-Linearisation

- This means we can write:

$$\tilde{\psi}_t^I - \tilde{C}_t = -\tilde{C}_{t+1} + (1 - \beta(1-\delta))\tilde{R}_{t+1} + (1-\delta)\beta\tilde{\psi}_{t+1}^I$$

- Now, lets re-arrange this slightly to isolate the $\tilde{\psi}_t$ terms on one side
- This yields:

$$\beta(1-\delta)\tilde{\psi}_t^I - \tilde{\psi}_t^I = \mathbb{E}_t\tilde{C}_{t+1} - \tilde{C}_t - (1 - \beta(1-\delta))\mathbb{E}_t\tilde{R}_{t+1}$$

Log-Linearisation

- Now, lets assume that $\tilde{\psi}^I_{t+1}$ follows an AR(1) process
- Therefore $\mathbb{E}_t \tilde{\psi}^I_{t+1} = \rho_I \tilde{\psi}^I_t$ in a deterministic setting
- And the left hand side becomes:

$$(\beta(1-\delta)\rho_I - 1)\tilde{\psi}^I_t = \mathbb{E}_t \tilde{C}_{t+1} - \tilde{C}_t - (1-\beta(1-\delta))\mathbb{E}_t \tilde{R}_{t+1}$$

Log-Linearisation

- Now, where does this get us?
- From the firm's FOC, we can get $\mathbb{E}_t \tilde{R}_{t+1}$ as:

$$\mathbb{E}_t \tilde{R}_{t+1} = \tilde{Y}_{t+1} - \tilde{K}_{t+1}$$

- Define the transformed investment wedge as $\tilde{\psi}^I_{2,t} = (\beta(1-\delta)\rho_I - 1)\tilde{\psi}^I_t$ then we can write:

$$(\beta(1-\delta)\rho_I - 1)\tilde{\psi}^I_{2,t} = \mathbb{E}_t \tilde{C}_{t+1} - \tilde{C}_t - (1 - \beta(1-\delta))\left(\mathbb{E}_t \tilde{Y}_{t+1} - \tilde{K}_{t+1}\right)$$

- So we measure a (scaled) investment wedge as the residual from expected consumption growth and the expected marginal product of capital
- This allows for us to get the investment wedge

Log-Linearisation

- In other words, if we have data for output, capital, consumption, investment, and labour input, and are willing to take a stand on a few parameter values
- We can back out the wedges essentially as residuals from the relevant first order conditions

Measuring the Wedges in the Data

- In the example I've use data on output, the capital stock, consumption, investment, and labour hours
- Consumption is the real personal consumption expenditure series (which is not ideal as it includes durable goods purchases, which are usually termed investment)
- Real private fixed investment is the investment series
- After taking the natural logarithm of each series, we remove a linear time trend
- This would ensure that we are consistent with the idea of deviations from the steady state in the equations above
- The sample period is 1947q1 - 2016q4

Measuring the Wedges in the Data

- We then measure the wedges using the log-linearised equations and the data
- We assume that $\beta = 0.99$, $\delta = 0.02$, $\alpha = 1/3$, and $\gamma = 1$
- One can then investigate the plots of the time series of the four wedges
- Shaded gray regions are recessions as defined by the NBER
- One thing that sticks out is that both the efficiency and labour wedges are quite persistent and very procyclical
- They tend to be low in periods identified by the NBER as recessions

Measuring the Wedges in the Data

- The investment wedge seems to have quite a bit of high frequency volatility (but the scale is small) and not very persistent
- If anything seems to be countercyclical - i.e. high during the Great Recession
- The government consumption wedge moves around quite a bit, with a very big spike in the early 1950s, which is associated with the Korean War and the start of the Cold War
- It's not obviously pro or countercyclical from eyeballing the series

Measuring the Wedges in the Data

Time Series of the Wedges

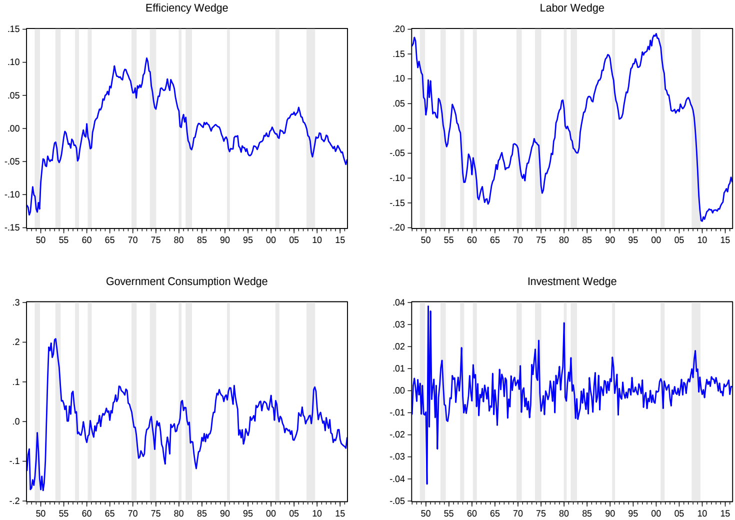

Measuring the Wedges in the Data

- The summary statistics for the wedges are provided below
- Note that the labour wedge is the most volatile and the investment wedge is least volatile
- All are positively correlated with output
- Note these are correlations are based on linearly detrended data (not HP filtered data) so they may not line up with what you might see elsewhere

Measuring the Wedges in the Data

Moments of the Wedges

Wedge	Standard Deviation	Correlation w/ Output
Efficiency	0.0462	0.610
Labour	0.0992	0.5178
Government	0.0645	0.2398
Investment	0.008	0.1087

Measuring the Wedges in the Data

- That the efficiency wedge is procyclical should not be surprising
- The efficiency wedge is just measured TFP
- What is interesting is that the labour wedge is both
  - quite volatile and very procyclical
- Recall that the labour wedge above is defined as $\psi_t^N = 1 - \tau_t^N$
- In other words, the labour wedge is low when output is low
- This means that it is as though there is a tax rate of labour income that is high (i.e. during the Great Recession)
- Hence, the procyclical labour wedge implies that the labour market seems to be countercyclically distorted
- Alternatively we could say it is as if there is a tax on labour income that is high in recessions and low in expansions

Measuring the Wedges in the Data

- Simply looking at unconditional moments of the wedges doesn't really tell us much about how important they are or are not for understanding fluctuations
- To do that, we need to take the model and feed these wedges into it and look at how much of the volatility of output (and other variables) can be explained by these wedges
- Hence, we need to parameterize stochastic processes for the wedges that could then be used for a variance decomposition

Measuring the Wedges in the Data

- In a fairly simplistic manner we could consider a fairly simple exercise
- Would involve simply estimating AR(1) processes on each wedge

$$
\begin{aligned}
\log \tilde{A}_t &= \rho_A \log \tilde{A}_{t-1} + s_A \varepsilon_{A,t} \\
\tilde{G}_t &= \rho_G \tilde{G}_{t-1} + s_G \varepsilon_{G,t} \\
\tilde{\psi}_t^N &= \rho_N \tilde{\psi}_{t-1}^N + s_N \varepsilon_{N,t} \\
\tilde{\psi}_t^I &= \rho_N \tilde{\psi}_{t-1}^I + s_I \varepsilon_{I,t}
\end{aligned}
$$

- Thereafter we report the AR coefficient and the standard error of the innovation

Measuring the Wedges in the Data

AR(1) Coefficients from Wedge Regressions

Wedge	AR Parameter	SD of Innovation
Efficiency	0.9739	0.0083
Labour	0.9854	0.0151
Government	0.9193	0.0245
Investment	0.064	0.0079

Measuring the Wedges in the Data

- There are a couple of things to note about these regressions
- First, there should be no constant (while I included a constant) because the wedges are based on linearly detrended data, and hence are mean zero
- Second, the wedges are correlated with one another in the data
- When expressing them as independent AR(1) processes we stipulate that these wedges are uncorrelated processes
- CKM actually write down a VAR process for the wedges
- This allows for the innovations to the wedges to be correlated as well as for lagged values of one wedge to impact other wedges

Measuring the Wedges in the Data

- This example is relatively straightforward, but it does not affect the main message of the exercise to follow
- consistent with a visual inspection of the series, the efficiency, labour, and government wedges are quite persistent, while the investment wedge is not
- What we can do next is to solve the model using these specified AR(1) processes for the wedges
- In so doing, I assume that the efficiency, labour, and investment wedges are all in fact mean zero

Measuring the Wedges in the Data

- We also need to specify a non-stochastic steady state value for G
- Hence we assume that government spending, which may include output not accounted for by consumption and investment
- This includes both government spending and net exports and is about 20% of steady state output

Measuring the Wedges in the Data

- The table shows the unconditional variance decomposition of output, hours, investment, and consumption to the four wedge shocks in the model
- The results are pretty stark and are consistent with what CKM claim
- In particular, the efficiency and labour wedges explain virtually all of the variance of these variables

Measuring the Wedges in the Data

- Government spending and investment wedges are essentially completely irrelevant
- Loosely speaking, the efficiency wedge accounts for about 60% of the variance of output
- While the labour wedge accounts for the other 40%

Measuring the Wedges in the Data

AR(1) Coefficients from Wedge Regressions

	Percent of Unconditional Variance Due to			
Variable	Efficiency	Labour	Government	Investment
Output	61.7	37.8	0.2	0.3
Labour hours	3.5	92.9	1.9	1.7
Investment	62.8	26.8	3.2	7.3
Consumption	56.5	40.7	1.8	1

Interpreting the Wedges

- The efficiency wedge can be interpreted as exogenous stochastic shock to productivity
- The government consumption wedge can be interpreted as exogenous stochastic shock to government spending
- The labour and investment wedges do not have a clear interpretation as exogenous shocks
- These could reflect exogenous shocks, or they could represent some mis-specification of the model along some dimension

Interpreting the Wedges

- What we have laid out above is a simple accounting exercise
- Hence the title of the CKM paper
- Through the lens of a RBC model, we examine where the model needs wedges to explain the data
- Going from these wedges to structural economic shocks, with the potential exception of the efficiency and government consumption wedges, involves a bit of liberal interpretation

Interpreting the Wedges

- One key message that comes out of the CKM paper is that we need to come up with a way to account for the labour wedge
- The static first order condition for labour supply fails very badly in the data
- This suggests that there is some important mis-specification or some important missing shock
- Since the publication of their paper there has been a substantial amount of research aimed at explaining the labour wedge

Interpreting the Wedges

- There are several potential alternative interpretations of what gives rise to the labour wedge
- In the model that has been presented one possibility is a time-varying tax on labour income
- Statutory tax rates do not vary that much in the data to take this potential explanation seriously
- Though one could argue that effective tax burdens vary in a countercyclical way over the business cycle

Interpreting the Wedges

- An alternative explanation of the labour wedge is simply a shock to the disutility from labour:

$$\upsilon_t \theta N_t^{\gamma} = \frac{1}{C_t} w_t$$

- Here, after log-linearisation υ_t would be similar to a time-varying tax on labour income

Interpreting the Wedges

- Since the model-implied labour wedge moves around so much in the data, several papers have sought to formally estimate the stochastic properties of different structural shocks
- Very often they find that labour supply shocks, like υ_t as modeled above, are very important drivers of the business cycle
- Smets & Wouters (2007, AER) estimate a structural New Keynesian DSGE model and argue that a "wage markup shock" is one of the key drivers of the business cycle
- To a first order approximation, the wage markup shock looks identical to a labour supply preference shock

Interpreting the Wedges

- In a separate paper, Chari, Kehoe, & McGrattan (2009, AEJ Macro) make this point, and hence argue that the Smets-Wouters wage markup shock cannot be considered structural
- There are very different policy implications if something is causing the labour market to be more distorted in recessions (the wage markup interpretation) than if people simply dislike working more in recessions (the labour supply interpretation)
- Time-varying monopoly power in product markets could also help account for the labour wedge
- Either because of explicit time-variation in monopoly power or because of price stickiness resulting in countercyclical price markups

Conclusion

- It is well-accepted that the static first order condition for labour supply from the base- line model is not consistent with the data
- In that sense, CKM's paper makes an important point that macroeconomists need to better understand the labour wedge and why the first order condition for labour supply fails relative to the data
- Hall (1997, Journal of Labour Economics) and Gali, Gertler, & Lopez-Salid (2007, ReStat) previously made a very similar argument

Conclusion

- The more provocative claim in CKM's paper is that the lack of importance of the investment wedge means that research focusing on financial shocks and frictions is not likely to be a fruitful avenue for future research
- Many feel that this claim is too strong, and it is not too difficult to write down a model with a type of financial constraint that manifests directly as a labour wedge
- For example, one could also envision a model in which a financial shock leads to a poor allocation of factors of production across firms in a way that would manifest as an efficiency wedge