The Welfare Cost of Business Cycles

Jennifer La'O*

1 The Welfare Cost of Business Cycles

We now use our model to compute the welfare cost of business cycles as in Lucas (1987).

First: what do we mean by the cost of business cycles? Although we use utils as a measure of happiness, this measure only has relative meaning–it has no meaning in terms of levels in and of itself. Lucas thereby proposed to measure the cost of business cycles in terms of a proportional upward shift in the consumption process that would be required to make the representative consumer indifferent between its random consumption allocation and a nonrandom consumption allocation with its current mean.

That is, suppose the representative agent has the following consumption path:

$$c_t = \bar{c}\exp(gt)\exp(\varepsilon_t)$$

or equivalently,

$$\log c_t = \log \bar{c} + gt + \varepsilon_t$$

where $g > 0$ is the trend and ε_t is the business-cycle variation in consumption that follows some stochastic process.

Let us assume a hypothetical no-cycle path $\hat{c}_t$ given by

$$\hat{c}_t = \bar{c}\exp(gt)$$

Then let us define the welfare cost of business cycles as the fraction λ that satisfies

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u((1+\lambda)c_t) = \sum_{t=0}^{\infty} \beta^t u(\hat{c}_t) \tag{1}$$

Thus λ represents the upward shift in consumption required to make the representative consumer indifferent between its consumption path and a hypothetical nonrandom consumption path.

To compute this, let us make some simplifying assumptions. First let us assume the standard

*Department of Economics, Columbia University, jenlao@columbia.edu.

CRRA preferences

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}.$$

Next for simplicity let's assume i.i.d. consumption shocks

$$\varepsilon_t \sim \mathcal{N}\left(\mu, \sigma^2\right)$$

with μ normalized so that

$$\mathbb{E}\exp\varepsilon_t = \exp\left\{\mu + \frac{1}{2}\sigma^2\right\} = 1$$

Therefore we set $\mu = -\frac{1}{2}\sigma^2$.

The right hand side of equation (1) is thus given by

$$\begin{aligned}
\sum_{t=0}^{\infty}\beta^t u(\hat{c}_t) &= \sum_{t=0}^{\infty}\beta^t \frac{(\bar{c}\exp(gt))^{1-\gamma}}{1-\gamma} \\
&= \frac{\bar{c}^{1-\gamma}}{1-\gamma}\sum_{t=0}^{\infty}\beta^t\left(\exp(gt)\right)^{1-\gamma} \\
&= \frac{1}{1-\beta\exp(g\left(1-\gamma\right))}\frac{(\bar{c})^{1-\gamma}}{1-\gamma}
\end{aligned}$$

The left hand side of equation (1) is given by

$$\begin{aligned}
\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t u((1+\lambda)c_t) &= \mathbb{E}_0\sum_{t=0}^{\infty}\beta^t \frac{((1+\lambda)\bar{c}\exp(gt)\exp(\varepsilon_t))^{1-\gamma}}{1-\gamma} \\
&= \frac{\bar{c}^{1-\gamma}}{1-\gamma}(1+\lambda)^{1-\gamma}\sum_{t=0}^{\infty}\beta^t\left(\exp(gt)\right)^{1-\gamma}\mathbb{E}_0\left(\exp(\varepsilon_t)\right)^{1-\gamma}
\end{aligned}$$

Note that

$$\begin{aligned}
\mathbb{E}_0\left(\exp(\varepsilon_t)\right)^{1-\gamma} &= \exp\left\{(1-\gamma)\,\mu + \frac{1}{2}\left(1-\gamma\right)^2\sigma^2\right\} \\
&= \exp\left\{-\left(1-\gamma\right)\tfrac{1}{2}\sigma^2 + \frac{1}{2}\left(1-\gamma\right)^2\sigma^2\right\} \\
&= \exp\left\{-\gamma\left(1-\gamma\right)\tfrac{1}{2}\sigma^2\right\}
\end{aligned}$$

Therefore

$$\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t u((1+\lambda)c_t) = \frac{\bar{c}^{1-\gamma}}{1-\gamma}(1+\lambda)^{1-\gamma}\exp\left\{-\gamma\left(1-\gamma\right)\tfrac{1}{2}\sigma^2\right\}\sum_{t=0}^{\infty}\beta^t\left(\exp(gt)\right)^{1-\gamma}$$

Substituting these back into our indifference condition (1) gives us

$$\frac{\bar{c}^{1-\gamma}}{1-\gamma}(1+\lambda)^{1-\gamma}\exp\left\{-\gamma\left(1-\gamma\right)\tfrac{1}{2}\sigma^2\right\}\sum_t\beta^t\left(\exp(gt)\right)^{1-\gamma}=\frac{\bar{c}^{1-\gamma}}{1-\gamma}\sum_t\beta^t\left(\exp(gt)\right)^{1-\gamma}$$

Cancelling terms gives us,

$$(1+\lambda)^{1-\gamma}\exp\left\{-\gamma\left(1-\gamma\right)\tfrac{1}{2}\sigma^2\right\}=1.$$

Solving this expression for λ gives us

$$\begin{aligned}(1+\lambda)^{1-\gamma} &= \exp\left\{\gamma\left(1-\gamma\right)\tfrac{1}{2}\sigma^2\right\}\\ 1+\lambda &= \exp\left\{\gamma\tfrac{1}{2}\sigma^2\right\}\end{aligned}$$

Therefore, with some simple algebra we get that

$$\lambda\approx\frac{1}{2}\gamma\sigma^2$$

Let's consider this cost in terms of U.S. data. We have that the standard deviation of consumption growth lies somewhere between 1%-2%

$$\sigma\in(.01,.02)$$

Let's just use the higher number of $\sigma=.02$. Then

$$\lambda\approx\frac{1}{2}\gamma\left(.02\right)^2=\frac{1}{2}\gamma\left(.0004\right)$$

Finally, consider a reasonable range of γ the CRRA parameter.

- Suppose we take $\gamma=1$ (log preferences). Then

  $$\lambda\approx\frac{1}{2}\left(.02\right)^2=1/2\left(.0004\right)=.0002$$

  Thus the implied welfare cost is given by

  $$\lambda\approx0.02\%$$

- Suppose instead we take $\gamma=10$ (huge risk aversion). Then the implied welfare cost is given by

  $$\lambda\approx0.2\%$$

Conclusion. Lucas (1987)'s simple back-of-the-envelope calculation tells us that even with very high risk aversion, the welfare cost of business cycles seems to be miniscule!

Do we believe that the cost of business cycles is so small? Depends on who you ask. Many people would say that like the equity premium puzzle, Lucas's welfare costs calculation simply tells us that there is something deeply wrong with our model. Some caveats/extensions:

- even higher risk aversion? still seems implausible
- persistent shocks to consumption
- incomplete markets, distributional effects of the business cycle

2 Using Asset Prices to infer the Cost of Business Cycles

Lucas (1987) is a seminal paper which informs us that the cost of business cycles through the lens of our model is small. However, in order to calculate this cost, one had to make certain assumptions on preferences (above we assumed CRRA) and on the consumption process. Several studies have proposed estimates of this cost of business cycles under alternative assumptions on preferences and the consumption processes. Primarily as a function of the specification and parameterization of preferences, these estimates can vary widely across studies.

Following this work, Alvarez and Jermann (2004) instead try to infer information about the cost of business cycles from financial asset prices. The advantage of this approach is that they may obtain a measure of the welfare cost of business cycles without fully specifying consumer preferences nor the stochastic process for consumption; see also Chapter 8 of Ljungqvist and Sargent (2004).

Consider a one-period consumption strip as a claim to the random payoff c_t sold at date $t-1$. The price in terms of $t-1$ consumption of this one-period consumption strip is

$$p_{t-1} = \mathbb{E}_{t-1}\left[m_t c_t\right]$$

where m_t is the one-period stochastic discount factor.

We thus have that

$$p_{t-1} = \mathbb{E}_{t-1} m_t \mathbb{E}_{t-1} c_t + Cov_{t-1}\left(c_t, m_t\right). \tag{2}$$

where recall that $Cov_{t-1}\left(c_t, m_t\right)$ is the risk correction of the asset. Given the CCAPM, we have that the stochastic discount factor $m_t = \beta u'\left(c_t\right)/u'\left(c_{t-1}\right)$ which implies that

$$Cov_{t-1}\left(c_t, m_t\right) < 0$$

as long as utility is concave.

Now consider the price of another asset: a one period risk-free claim to expected consumption $\hat{c}_t \equiv \mathbb{E}_{t-1} c_t$. The price of this security is simply given by

$$\hat{p}_{t-1} = \mathbb{E}_{t-1}\left[m_t \hat{c}_t\right] = \mathbb{E}_{t-1} m_t \mathbb{E}_{t-1} c_t$$

Thus the negative covariance in (2) is a discount on the risky claim relative to the risk-free claim on a payout with the same mean.

Let us define the multiplicative risk premium on the consumption strip as

$$1+\mu_{t-1} \equiv \frac{\hat{p}_{t-1}}{p_{t-1}} = \frac{\mathbb{E}_{t-1}m_t\mathbb{E}_{t-1}c_t}{\mathbb{E}_{t-1}\left[m_t c_t\right]} > 1 \tag{3}$$

The idea is to then coax attitudes about the cost of business cycles from asset prices.

Alvarez and Jermann (2004) again define the *total* cost of the business cycle in terms of a stochastic process of adjustments to consumption λ_{t-1} that satisfies

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u((1+\lambda_{t-1})\, c_t) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(\mathbb{E}_{t-1}c_t)$$

The idea is to compensate the consumer for the one-period-ahead risk in consumption that he faces.

The time t component of the marginal cost of business cycles is defined as follows. Let $\alpha \in [0,1]$ be a parameter that indexes consumption processes given by

$$\alpha\mathbb{E}_{t-1}c_t + (1-\alpha)\, c_t$$

and let us define $\lambda_{t-1}(\alpha)$ implicitly by

$$\mathbb{E}_{t-1}u\left((1+\lambda_{t-1}(\alpha))\, c_t\right) = \mathbb{E}_{t-1}u\left(\alpha\mathbb{E}_{t-1}c_t + (1-\alpha)\, c_t\right) \tag{4}$$

Note that at $\alpha = 0$

$$\mathbb{E}_{t-1}u\left((1+\lambda_{t-1}(0))\, c_t\right) = \mathbb{E}_{t-1}u\left(c_t\right).$$

Therefore $\lambda_{t-1}(0) = 0$.

Let's differentiate equation (4) with respect to α to get

$$\mathbb{E}_{t-1}\left[u'\left((1+\lambda_{t-1}(\alpha))\, c_t\right)\lambda'_{t-1}(\alpha)\, c_t\right] = \mathbb{E}_{t-1}\left[u'\left(\alpha\mathbb{E}_{t-1}c_t + (1-\alpha)\, c_t\right)\left(\mathbb{E}_{t-1}c_t - c_t\right)\right]$$

Evaluating this at $\alpha = 0$ gives us

$$\lambda'_{t-1}(0)\,\mathbb{E}_{t-1}\left[u'(c_t)\, c_t\right] = \mathbb{E}_{t-1}\left[u'(c_t)\left(E_{t-1}c_t - c_t\right)\right]$$

Thus

$$\lambda'_{t-1}(0) = \frac{\mathbb{E}_{t-1}\left[u'(c_t)\left(\mathbb{E}_{t-1}c_t - c_t\right)\right]}{\mathbb{E}_{t-1}\left[u'(c_t)\, c_t\right]} \tag{5}$$

Alvarez and Jermann (2004) define this as the marginal cost of consumption fluctuations. Lucas' cost of business cycles measures the welfare gain from removing all the business cycle risk; it can be thought of as a total cost. Instead the marginal cost measures the welfare benefits from reduced consumption fluctuations at the margin.

This definition has two useful features: (i) Because it is a marginal cost, we can use asset prices to estimate the cost of business cycles for a representative agent, and (ii) Given that most economic policies would not be intended to eliminate business cycle fluctuations entirely, knowing the potential benefits at the margin may be useful in and of itself.

Multiplying both the numerator and denominator of (5) by $\beta/u'(c_{t-1})$ gives us

$$\lambda'_{t-1}(0) = \frac{\mathbb{E}_{t-1}\left[\frac{\beta u'(c_t)}{u'(c_{t-1})}\left(\mathbb{E}_{t-1}c_t - c_t\right)\right]}{\mathbb{E}_{t-1}\left[\frac{\beta u'(c_t)}{u'(c_{t-1})}c_t\right]}$$

therefore

$$\begin{aligned}\lambda'_{t-1}(0) &= \frac{\mathbb{E}_{t-1}\left[m_t\left(\mathbb{E}_{t-1}c_t - c_t\right)\right]}{\mathbb{E}_{t-1}\left[m_t c_t\right]} \\ &= \frac{\mathbb{E}_{t-1}\left[m_t \mathbb{E}_{t-1}c_t - m_t c_t\right]}{\mathbb{E}_{t-1}\left[m_t c_t\right]} \\ &= \frac{\mathbb{E}_{t-1}(m_t)\,\mathbb{E}_{t-1}(c_t)}{\mathbb{E}_{t-1}\left[m_t c_t\right]} - 1\end{aligned}$$

Therefore

$$1 + \lambda'_{t-1}(0) = \frac{\mathbb{E}_{t-1}(m_t)\,\mathbb{E}_{t-1}(c_t)}{\mathbb{E}_{t-1}\left[m_t c_t\right]} \tag{6}$$

Comparing equation (6) to (3) we infer that

$$1 + \lambda'_{t-1}(0) = 1 + \mu_{t-1} = \frac{\hat{p}_{t-1}}{p_{t-1}}$$

Therefore the marginal cost of business cycles at time t is equal to the multiplicative risk premium on the one-period consumption strip. Thus, these authors show that the marginal cost of business cycles can be coaxed from asset market data! Pretty clever.

Again, using asset market data for short-term (business cycle) they find this marginal cost of risk to be as low as 0.3%. See Alvarez and Jermann (2004) for details.

References

Alvarez, Fernando and Urban Jermann, "Using Asset Prices to Measure the Cost of Business Cycles," *Journal of Political Economy*, 2004, *112* (6), 1223–1256.

Ljungqvist, Lars and Thomas J. Sargent, *Recursive macroeconomic theory*, Cambridge, Mass: MIT Press, 2004.

Lucas, Robert E., *Models of Business Cycles*, New York: Basil Blackwell, 1987.