"Cost of Business Cycles"
from Models of Business Cycles
by Robert Lucas

III

If we are to move beyond the general formalism I have been discussing to models that account for specific features of business cycles, it will be necessary to be more specific about the agents in the system, about the technology at their disposal, and about the way they interact. In this section, I will consider consumer preferences only, saying nothing about the other, much more difficult, aspects of the problem. It is remarkable how much one can say about the importance of macroeconomic questions on the basis of preferences alone.

Any economic model is going to have at its center a collection of hypothetical consumers whose decisions, together with the technology and market structure, determine the operating characteristics of the system and whose welfare is the explicit subject of normative analysis. A typical household will consume a collection c_t of goods at date t, possibly contingent on probabilistically-determined events between dates zero and t, and will evaluate an entire sequence, or process, $\{c_t\}$ of

consumption according to a utility function such as (1), specialized here to:

$$E\left\{\sum_{t=0}^{\infty}\beta^t U(c_t)\right\} \tag{3}$$

If we are to think about economic policy starting from this viewpoint, what we mean, in the first place, is that we want the ability to be able to determine how different policies will induce different consumption sequences $\{c_{it}\}$ for each agent i in this economy. In the second place, we mean to evaluate policies normatively according to their effects on agents' welfare as measured by (3).

We can move to some sharper, quantitative conclusions for macroeconomic problems if we abstract from issues involving the *mix* of consumption goods at date t and limit discussion to policies affecting consumption of goods-in-general. At the simplest level, let us identify c_t with real consumption at date t, and specialize preferences to:

$$E\left\{\sum_{t=0}^{\infty}\beta^t \frac{1}{1-\sigma}(c_t^{1-\sigma}-1)\right\} \tag{4}$$

where $\beta \in (0, 1)$ is a constant discount factor and $\sigma > 0$ is the constant coefficient of relative risk-aversion. This two-parameter preference family will not be adequate for every problem I will want to address, but it will serve to get the discussion going in a concrete way.

Before situating this consumer (4) in a model of the sort outlined in Section II, let us examine his attitudes toward some purely hypothetical consumption streams by simply *asking* him about them. Since I am particularly interested in his attitudes toward growth and fluctuations, it will be useful to work with a class of consumption streams with 'trend' and 'cycle' components, such as:

$$c_t = (1+\lambda)(1+\mu)^t e^{-\frac{1}{2}\sigma_z^2} z_t, \quad t = 0, 1, \ldots, \qquad (5)$$

where $\{z_t\}$ is a stationary stochastic process with a stationary distribution given by:

$$\ln(z_t) \sim N(O, \sigma_z^2).$$

Then $E(e^{-\frac{1}{2}\sigma_z^2} z_t) = 1$, so that mean consumption under these assumptions is $(1+\lambda)(1+\mu)^t$.[1] Setting

1 The assumption that mean consumption follows a deterministic trend, and hence is perfectly predictable, is not innocuous in this context, and it would be desirable to work through the calculations below under alternative assumptions. Charles R. Nelson and Charles I. Plosser, 'Trends and random walks in macroeconomic time series, some evidence and implications', *Journal of Monetary Economics* **10** (1981), pp. 139–62, have recently argued that most year-to-year variability in US real GNP can be attributed to a random walk, or stochastic trend, component. It is likely that similar methods applied to consumption would lead to a similar conclusion. I agree with John H. Cochrane, 'How big is the random walk in GNP?' (University of Chicago Working Paper, 1986) that Nelson and Plosser's methods have considerably overstated the importance of the random walk component, but even so it seems clear that something inter-

the parameter λ is just a matter of units: I will use it later on to provide 'compensation' for variations in the parameters μ and σ_z^2. For the USA the annual growth rate in total consumption is about 3 per cent, so $\mu = 0.03$ can serve as a benchmark value. For the post-war period in the USA the standard deviation of the log of consumption about trend is about 0.013 (so that average consumption that is as much as 2.5 per cent above or below trend is a 'two-sigma' event), so that $\sigma_z^2 = (0.013)^2$ is a benchmark value.[2] Thus we may take (5) with $(\lambda, \mu, \sigma_z^2) = (0, 0.03, (0.013)^2)$ as a rough description of the consumption behavior the average American family is used to, and examine its attitudes toward changes.

Given any choice of $(\lambda, \mu, \sigma_z^2)$ we could simply calculate the value of (4) under the consumption

mediate to Nelson and Plosser's model and (5) would provide a better description of consumption behavior than (5) does.

2 This figure is from Finn E. Kydland and Edward C. Prescott, 'Time to build and aggregate fluctuations', *Econometrica* **50** (1982), pp. 1345–70, table IV, p. 1365. Since 0.013 is a quarterly figure, it overstates the standard deviation of annual consumption, but not by much, since consumption is highly serially correlated. This number and others used in this section are intended to give a rough idea of the relative importance of certain issues. The reader will agree, I think, upon reaching the end of the section that its conclusions do not hinge on delicate questions of measurement.

behavior (5) and call the indirect utility function so defined $U(\lambda, \mu, \sigma_z^2)$. But we will obtain a measure that is easier to think about if we use *compensating* variations in λ to evaluate various μ and σ_z^2 changes. To evaluate changes in the growth rate μ, for example, let us define $f(\mu, \mu_0)$ by:

$$U(f(\mu, \mu_0), \mu, \sigma_z^2) = U(0, \mu_0, \sigma_z^2), \qquad (6)$$

so that $f(\mu, \mu_0)$ is the percentage change in consumption, uniform across all dates and values of the shocks, required to leave the consumer indifferent between the growth rates μ and μ_0. (In general, σ_z^2 would appear as an argument of f, too, but under these 'constant relative risk-aversion' preferences it drops out.) A direct calculation gives:

$$f(\mu, \mu_0) = \left(\frac{1+\mu_0}{1+\mu}\right)^{\beta/(1-\beta)} - 1.$$

Here is a table of this function f, which I will call simply the *cost* of reduced growth, for $\beta = 0.95$ and a base growth rate of $\mu_0 = 0.03$.

At the parameters used in table 1, then, consumers would require a 20 per cent across-the-board consumption increase to accept voluntarily a reduction in the consumption growth rate from 0.03 to 0.02, and would surrender 42 per cent across the board to obtain an increase in the growth rate from 0.03 to 0.06.

I hasten to add that I have said nothing about the *feasibility* of these growth rate changes; I am

TABLE 1 Cost of reducing growth from $\mu_0 = 0.03$ when $\beta = 0.95$

μ	$f(\mu, \mu_0)$
0.01	0.45
0.02	0.20
0.03	0.00
0.04	−0.17
0.05	−0.31
0.06	−0.42

simply running experiments on this fictional consumer. Indeed, under a standard, neoclassical technology, policies that affect growth usually do so only over a transient period only, not permanently as in table 1. But the range of growth rates in table 1 is not large relative to what we observe across countries, and the welfare consequences of 'small' changes are *enormous*, relative to anything we will see in what follows. I shall return to this later on.

The costs of economic instability can be measured in a way that is identical conceptually to this way of measuring the costs of reduced growth. To this end, define the function $g(\sigma_z^2)$ by

$$U(g(\sigma_z^2), \mu, \sigma_z^2) = U(0, \mu, 0). \qquad (7)$$

That is, $g(\sigma_z^2)$ is the percentage increase in consumption, uniform across all dates and values of the shocks, required to leave the consumer indifferent between consumption instability of σ_z^2 and a *perfectly* smooth consumption path. I will call $g(\sigma_z^2)$ the cost of consumption instability.

By direct calculation, and using the approximation $\ln(1+\lambda)=\lambda$ (which is entirely safe in this context), g is given by:

$$g(\sigma_z^2) \cong \tfrac{1}{2}\sigma \cdot \sigma_z^2. \tag{8}$$

Table 2 shows the function g, for various σ and σ_z^2 values. The coefficient σ of risk-aversion can be estimated from a variety of different samples: estimates vary widely. A value of unity means logarithmic preferences; people appear to be more risk-averse than this. No available estimates are as large as 20, but some do exceed 10.

The value 0.013 is the standard deviation of the log of US real quarterly consumption, expressed as a deviation from fitted trend, over the period following the Second World War. Eliminating

TABLE 2 Cost of consumption instability: $g(\sigma_z^2)$

σ	σ_z		
	0.013	*0.039*	*0.120*
1	0.00008	0.00072	0.0065
5	0.00042	0.0038	0.034
10	0.00084	0.0076	0.068
20	0.0017	0.015	0.136

aggregate consumption variability of this magnitude entirely, would, from table 2, be the equivalent in utility terms of an increase in *average* consumption of something less than one tenth of a percentage point. (Total US consumption in 1983 was $2 trillion, so one-tenth of 1 per cent is $2 billion, which sounds like a sizeable free lunch. But there were 234 million people to feed, so lunch will have to run about $8.50 per person.) I want to propose taking these numbers seriously as giving the order-of-magnitude of the potential marginal social product of additional advances in business cycle theory – or more accurately, as a loose upper bound, since there is no reason to think that eliminating *all* consumption variability is either a feasible or a desirable objective of policy. But I imagine that even one-tenth of a percentage point will seem to many to be an extremely low estimate of the costs of economic instability – at least, it did to me – so it will be useful to digress to discuss some aspects of this estimate.[3]

3 The estimates in table 2 appear somewhat less surprisingly low when compared to estimates of the welfare gains from other (also purely hypothetical) policy changes. For example, Arnold C. Harberger, 'Monopoly and resource allocation', *American Economic Review* **44** (1954), pp. 77–87, found one-tenth of 1 per cent of income to be an upper bound on the welfare gain from the (costless) elimination of *all* product market monopoly in the US economy. Harberger, too, was led to 'confess that I was amazed at this result' (p. 86), but it has not been revised upward by any

The last two columns of table 2 set up what seem to me the two most important qualifications or elaborations of this cost estimate. First, in the period prior to the Second World War, and extending as far back in time as we have usable data, the standard deviation (logarithmic deviations from trend) of consumption was about three times its post-war level.[4] Since this number is squared in the

quantitative research I have seen in the 30 years since his study was published. Perhaps we should be open to the possibility that the intrinsic importance of substantive economic questions is not accurately reflected by the number of journal pages devoted to them.

It is not quite accurate to identify instability in goods consumption with economic instability in general, since consumption of leisure also fluctuates. But since hours worked and goods consumption are positively correlated cyclically, I would guess that taking leisure fluctuations into account more carefully would *reduce* the estimates in the text still further.

4 In a recent paper, Christina Romer, 'Spurious volatility in historical unemployment data', *Journal of Political Economy* **94** (1986, 1–37), has argued that pre-First (not second) World War variability in unemployment rates was, correctly measured, no larger than post-Second World War variability. It appears from related work of hers that the amplitudes of other series (probably including real consumption) were also badly overstated in pre-First World War data. The statement in the text may then rest much more heavily on the experience of the 1930s than I would previously have thought. It would be hard to overstate the importance of the questions on which these findings bear, but I have not attempted to incorporate them into my illustrative calculations.

formula (8), the implied cost estimates are multiplied by nine, becoming something like one-half of 1 per cent of total consumption. As deadweight losses go, *this* is a large number. Second, fluctuations in total consumption do not affect all households equiproportionally, so that variability of total consumption does not capture anything like all of the consumption risk faced by the typical household. Perhaps correcting for this effect would lead to another tripling of the relevant standard deviation, and hence another multiplication by nine in the cost estimates. This is the last column of table 2.

But as a measure of the possible gains from improvements in aggregative policy, this last column is way too high. In so far as the absence of income-risk pooling reflects 'imperfections' in capital markets, and I think it does, the cost of *individual* income variability measures the potential or actual gain from social insurance, not from stabilization policy. Aggregate income variability is but one source of individual income risk, and reduction of aggregate variability – which is all that stabilization policies can accomplish – cannot be expected to eliminate more than a small part of the uninsurable risk borne at the individual level. I will return to the issue of social insurance later on, in section V, when we have invested in a framework more suitable for posing questions about individual earnings risk and ways of dealing with it.

An economic system is a collection of people and serious evaluation of economic policy involves tracing the consequences of policies back to the welfare of the individuals they affect. Without saying much more about the nature or workings of the economy than this, we can get a good if rough idea of the potential benefit of policies that alter individual consumption streams in various ways. I have run through two exercises to assess the potential welfare gains of policies that affect the growth of consumption and policies that affect the variability of consumption about its trend, not by describing policies that would have these effects, but simply by imagining that these effects somehow come about. It is worth re-emphasizing that these calculations rest on assumptions about preferences *only*, and not about any particular mechanism – equilibrium or disequilibrium – assumed to generate business cycles.

I find the exercise instructive, for it indicates that economic instability at the level we have experienced since the Second World War is a minor problem, even relative to historically experienced inflation[5] and certainly relative to the costs

5 In Robert E. Lucas, Jr, 'Discussion of: Stanley Fischer, "Towards an understanding of the costs of inflation: II"', *Carnegie–Rochester Conference Series on Public Policy* **15** (1981), pp. 43–50, I estimated the annual social cost of a sustained 10 per cent inflation in the USA to be 0.5 per cent of national income. The estimate follows the method used in Martin J. Bailey, 'The welfare cost of inflationary finance', *Journal of Political Economy* **64** (1956), pp. 93–110, and uses the interest elasticity of money demand estimates in

of modestly reduced rates of economic growth. This is not to say that economic fluctuations are a trivial problem, for fluctuations at the pre-Second World War level, especially combined as they were with an absence of adequate programs for social insurance, were associated with large costs in welfare. But it suggests that the main social gains from a deeper understanding of business cycles, whatever form this deeper understanding may take, will be in helping us to see how to avoid large mistakes with policies that have minimally inefficient side-effects, not in devising ever more subtle policies to remove the residual amount of business-cycle risk.

Phillip Cagan, 'The monetary dynamics of hyperinflation', in Milton Friedman (ed.) *Studies in the Quantity Theory of Money* (University of Chicago Press, Chicago, 1956).