GRESHAM COLLEGE

01 OCTOBER 2015

THE COSTS OF BUSINESS CYCLES

PROFESSOR JAGJIT CHADHA

"When the next recession arrives you will find written on its bottom Made in Washington, just as was the case with the last one. This is not because the Fed is a sadist or an ignoramus; nor is it because the bedrooms in the White House are occupied by politicians eager to lose elections. Rather it is like the fact that hospitals are where most people die: If the central bank and fiscal authorities did not step on the brakes of an over exuberant economy now, they might well have to overdo that later. There is never a guarantee that intelligent and feasible policies can be discovered which will lead to perpetual soft landings at high employment and a steady price level....So let me make clear that, like the below-median poor, economic instability we have always with us."

Paul Samuelson (1998)

Introduction[1]

Macroeconomists can be cut up in two ways. Though some argue they should be cut into many more pieces. But the two cuts I have in mind are those interested in the short run and the long-run. The long run is concerned with economic growth, secular trends in income, wealth distribution and demographics. The short run is the stuff of unemployment, inflation, boom, bust and concentrated attention on interest rates: what we might think of as temporary fluctuations about the trend. These are often called business cycles and have been defined as recurrent, but not identical, persistent movements in output and other aggregate time series around an uncertain trend.

The study of business cycles has a long and honourable history in economics. Indeed prior to the advent of our science, the Bible seems clear with Joseph telling the Pharaoh to expect seven years of plenty and seven lean years. Well, since the collapse of Lehman Brothers in 2008, we might be entitled to feel that we have had those lean years. The serious study of business cycles can perhaps be traced to the work by the NBER by Burns and Mitchell on the US economy but also by a famous volume on the UK economy, also produced by US economists, Gayer, Rostow and Schwartz (GRS) called the Growth and Fluctuation in the British Economy, 1790-1850.

The GRS period can be identified closely with the classic period of industrial revolution, which propelled Britain to become the first industrial nation and produce sustained economic growth in a manner we have come to expect. But even in this period we can observe sharp temporary fluctuations in economic production. So the question might be just how costly were these and subsequent fluctuations? Our immediate answer shaped by the collective memory of unemployment in the 1930s, long term unemployment in the 1970s and 1980s, from the house price crash of 1990-92 and the pain of the global financial crisis is that the costs must be high. It would thus logically follow that we would be prepared to pay a high premium, in terms of consumption forgone today, in order to avoid such costs. In this lecture we shall try to understand how we might want to think about the

---

[1] This version of the lecture should be treated as a very preliminary and incomplete draft on which comments are very welcome.

costs of the business cycle, in terms of the variance in average consumption per head (or household) and why when we do the calculate the costs of business cycle fluctuations they do not look terribly high at all: typically a small fraction of 1% of average consumption per household.

Naturally, such a radically counter-intuitive result can be challenged and has been in many quarters. Rather than concentrating on the rebuttal in this lecture we shall try carefully and slowly to understand the calculation from first principles. There are a number of legs to the argument. First, that the business cycle is whatever the source essentially a temporary shock to income. Households will therefore try to maintain their level of consumption by cashing in on risk sharing schemes that they have with other members of their households, other households, firms, financial institutions, the government or monetary policy makers. To that extent that such hedges provide partial risk sharing and prevent any large changes in consumption from these temporary shocks, households will not be willing to pay any further premium to insure themselves from consumption risk. If, on the other hand, there are large welfare costs from business cycle fluctuations then there be a strong case for both stabilisation policy and for reforms to fill in missing markets.

In Section two, I will try to define the business cycle. In Section three, I will try to think about risk as the number of bits of consumption we are willing to give up, to eliminate uncertainty in consumption. In Section four, we shall try to measure and estimate this risk. In Section five, we shall argue against this calculation. Finally I conclude.

The Business Cycle

The key point about a business cycle is that it represents a temporary deviation from the economy's normal trend. We should not think of these deviations as uniformly periodic fluctuations as they are likely to result from a variety of shocks interacting with the economy and will depend on people's understand of these shocks and the anticipated economic policy response. In the early to mid-20th century business cycles became a study in their own right and economists such as Schumpeter, although clear in their acceptance of the phenomenon, also accepted that 'business cycles are not uniform in frequency, amplitude, or duration'. So if we identify a business cycle, we do not know for how long it will last but to qualify as such it cannot last forever, even though we might think that it will!

Figure two shows a stylised depiction of a business cycle with a peak and a trough representing turning points and with the bits between the turning points corresponding to expansions (trough to peak) and contractions (peak to trough), where the former has positive growth and the latter negative growth. A long expansion might lead to a boom with growth persistently higher than trend whereas a stuttering expansion with growth below trend might be termed a growth recession. As explained in last year's lecture the former boom may be responsible for the latter recession. We cannot know the trend of the economy with any degree of certainty in real time, although our estimates of the distant past may improve: ONS willing. The long run of UK data prior (see Chadha *et al.* 2002) to this current cycle, from 1871-1997 suggested that the business cycle could be expected to last some 20 quarters from peak to peak with a standard deviation of some nine quarters. As is usual in economics, the cycle of 1992 to now confounded that prediction rather well.

If I were to discuss the causes of business cycles, this lecture would quickly go to a quite different topic. Business cycles may have many, many primitive causes: from over investment, to financial crises, to imports of gold, too much or too little money, misperceptions and fiscal policy that are either too loose or too tight or too plain vanilla changes in total factor productivity. We could even try to uncover the primitive causes of business cycle fluctuations in the long run or across countries but to misquote Tolstoy: 'All happy expansions are alike; each unhappy recession is unhappy in its own way'.

Ultimately though, if we are to treat business cycles as temporary events that do not change the present value of economic wealth or permanent income, they may be less important than we might first have thought. The wind may blow and the rain may fall but if your house remains standing and you remain dry then what does it matter? Well it matters to the extent that the weather means you have to fix your house and so give up a trip to the

cinema or a pint or buying clothes for your children. If we treat all shocks that drive the business cycle as income or endowment then households have the unenviable task of trying to split those shocks into those that reflect changes in wealth and so might affect consumption and those that do not and so should be ignored from the perspective of the optimal path of the consumption. So let us assume that households do try to split all the temporary movements economy into two parts: those that reflect shocks to output and the ongoing responses of the economy to those shocks and do not have any substantial impact on permanent income those that do impact on wealth, permanent income, and so effect the pattern of consumption. Once that arduous task has been undertaken what next?

Understanding Risk

Macroeconomists think that a given representative household makes plans that maximise consumption over its lifetime or if, linked generationally, over infinite lives subject to a long run budget constraint in which all income from labour, land and capital is spent over that lifetime with a wish to the minimise the period to period variance of that consumption. The income in any period can be used to smooth consumption by borrowing when it is temporarily below the expected long run, or permanent, level of income and saving when it is above that level. Even if we do not know household by household the exact nature of the temporary versus permanent income, we might argue that over the whole population any one household's mistakes are ironed out by others, so on average this kind of process might get at aggregate fluctuations in the right manner.

Actually there may also be formal methods of risk sharing. Consider two households and suppose that income falls from the sky to these two households with a negative covariation. So that when the first household gets manna, the second gets none and vice versa. These two households could separately consume or starve. On average they will eat half of the manna but in reality every other time period will be quite difficult. So why not agree to share risk? Write a contract so that the household with the manna in any period gives half to the one without. This way both households consume half the manna every period and even though both households still consume on average half the manna, now they actually both consume half every period. Risk is shared across the two households in every period, or intratemporally.

Now consider that there are time-based fluctuations in the quantity of manna. That in some periods the households gets two, in some one and some none. On average over the long run, they get one. But if they consume up to the endowment, then the path of consumption although at an average of one, will display some variance and will be less preferable to one of period-to-period stability. Ideally, the household would want to store the excess manna so that it would be used in periods when the manna count is low. The representative household would like to smooth intertemporally. By sharing risk with other households in every period and then again sharing risk with the future itself, the household can in principle receive sufficient manna in every period to have a stable consumption path.

To the extent that intratemporal contracts may not be complete and inter-temporal risk sharing may be limited by the perishability of manna we may need a government policy. Indeed economic policy might also be thought of as a way of risk sharing. Increasing indebtedness from fiscal policy may allow impoverished households today to borrow from the taxes of future generations. Lower interest rates may allow the same set of impoverished households to borrow from their own future richer-selves. In this context, even movements in the exchange rate can be thought of as risk sharing because when Sterling depreciates in a recession, we invite foreigners to buy our goods and so provide some payoffs to beleaguered UK households.

If we then boil down all households to the average one, we can then see that they will be willing to trade away some consumption today for consumption variance. Along some indifference curve, which maps equivalent levels of utility (or happiness) we can think in terms of a mean-variance trade-off. That given an increase in the variance of consumption, our representative household will require a higher level of consumption in order to feel equally well off. If the household does not care about the variance of consumption and only cares about the mean, the indifference curves will be flat. But the steeper they are, the more compensation in terms of consumption is required for every unit increase in risk, as measured by the variance of consumption. Indeed we

can go the other way and can see how much consumption this representative household is willing to give up for a more stable life. This consumption offering is the heart of the idea of a risk premium and the costs of a business cycle.

Measuring Risk

We now need to understand how to quantify risk. I am not talking about uncertainty, which is some notion of helplessness in the face of all the problems that we might have in life. But, simply, how we can arrive at a notion of the cost of the willingness to avoid well specified variation in consumption payoffs? Our first problem is to draw a mapping from any given level of consumption into a measure of its value to a household. Let us call this a utility function and one is drawn in Figure three. We can see that any level of consumption translates into a unique level of welfare. Rather than a line, I have drawn this as an increasing function at a decreasing rate. This means that we get higher utility from a higher level of consumption but at a rate that falls for every extra unit of consumption.

Now consider the appearance of £100 of consumption units with certainty. Perhaps the payoff from a government bond. We can map the level of utility and if we are sure about the pay-off we do not need to pay a risk premium: if we are going to get £100 no matter what, there is no risk to quantify and to price. But now consider the possibility that we will get £50 half the time and £150 half the time: the former is a bad state of nature and the latter a good state of nature. The expected payoff is unchanged but now I have introduced some risk because in any one period I cannot be sure what I will receive. Let us assume that we can cannot write some contract with an agent who gets £150 when we get £50 and gets £50 when we get £150, which would provide complete risk sharing. The average utility from this kind of pay-off can be found by adding the utility from each state of nature and dividing by two.

We can see immediately that because the utility function is neither linear nor convex: it is concave; the average utility of any two points is always below the utility level derived from the mean level of consumption. Of course, this result derives from our assumptions about the mapping between consumption and utility but it does tell us that in there somewhere is the statement that the household prefers a certain payoff to an uncertain payoff with the same expected return. Would you prefer £100 in your pocket or a coin toss for £50 or £150? The more concave the function the more you would prefer the certain payoff because the more risk averse you are. But there is another aspect, the higher the variance of the payoff, for example, if the amount varied between £0 and £200, even though you have the same attitude to risk, the more you would be prepared to pay to avoid the risk. Let me illustrate with Figure 4.

The average utility gained from the two uncertain outcomes is the same as the household would receive from a lower level of consumption with certainty. The difference between the certain level of consumption and expected level of consumption under uncertainty is the amount the household would be prepared to pay to receive consumption with certainty rather than under uncertainty: it is the risk premium. Indeed we can also see that if we increase the variance of risk, the household will bid up its risk premium accordingly. We are now in a position to understand how to calculate the costs of business cycles: it is simply the amount of consumption the representative household will give up in order to eliminate any variation in consumption per household.

The calculation has the following parts. What is the steady-state or certainty-equivalent amount of household consumption? Let's say that is something like $30,000 in constant 2005 prices.[2] This calculation is rather much of guess, as we do not really know the correct long run exchange rate and we are using final consumption expenditure by household. But this is the kind of number we get in the US or the UK. What we now need is a measure of the standard deviation in consumption, which is how much do the consumption draws change over time according to what state of nature we receive. Again the US or UK data over the long or short run is clear the standard deviation is small at around 2-3% per year. Much smaller than the thought experiment we ran a bit

---

[2] http://data.worldbank.org/indicator/NE.CON.PRVT.PC.KD

earlier. If we take some guess as the curvature of the utility function, we then get a number of around £50 per year: a fraction of the insurance for a car.

The Meaning of the Median

Figure 5 below show the long time series in consumption in the US from 1930 to present compared to output and we can note that consumption is around half as variable as output with a standard deviation of 2.5% compared to 5%. The average household thus is able to filter away around 50% of business cycle noise. The Figure 6 shows the distribution of changes in consumption growth for the UK since the 1990s to present. The mean growth rate in this period has been 0.6% per year and the standard deviation of that growth has been just over 2% per year. Neither of this facts change the basic calculation in the previous section. Yet it is not entirely correct.

This calculation treats the average or median household as the unit of account. For reasons that I have already touched upon, I have not really bothered here to examine the impact of business cycle fluctuations on different income cohorts. This representative agent or household might mislead us. It might be that certain sections of the income cohort, probably lower, suffer considerably more variation in consumption as a result of business cycle shocks, as they may be more likely to lose their jobs and/or hit borrowing constraints with little or no form of personal household safety net. On the other hand, higher income groups may find they have access to many kinds of reinforcing instruments to absorb temporary income shocks such as high net worth, collaterisable wealth and high levels of human capital, which allows consumption to remain stable in the face of business cycle risks. Researchers such as Krusell and Smith (1999) have examined the relaxation of the representative agent assumption and examine households in which idiosyncratic and aggregate risk can only be smoothed with household wealth. Even though they are able to match the wealth distribution, and do find considerable differences across households, only for a tiny fraction of households are costs even as high as 2% of consumption.

Another possibility is that a considerable amount of the risk has already been reduced for all households from the operation of stabilisation policies. Even if business cycle shocks are large, their impact is mitigated by both actual and expected policies. Fiscal and monetary policies can offset both the intra-household and intertemporal impact of shocks and if designed correctly may reduce our measure of the risk that households face. We can though examine some data on consumption prior to the construction of the modern set of stabilisation policies and even then the variance in consumption does not seem significantly larger.

There is a further problem. We cannot treat the mean consumption growth rate and the dispersion around it over time as our measure of the true dispersion in consumption. Because when we sum up every household in the population there are two separate parts of the variance calculation. We need to understand the average variance for every household as well as the simple dispersion in the mean household's consumption over time. The true variance is the sum of both average household variance and the variance in the average. To that extent that the average variance in each household matters, this calculation is an underestimate.

Finally, in everything we have looked at today, the cycle has been treated as a separate entity from the trend. But what if the cycle affects the trend? Or the variance itself impacts on the mean. If that were the case, the costs of the business cycle are once again leveraged up. This effect could matter for the representative consumer so that the business cycle has permanent effects. But, perhaps more importantly, might matter for individual households (Atkeson and Phelan, 1994) within the distribution. So that even if there is no effect on the median household in the long run one any one household, if it gets a bad draw may end up facing permanent costs. Some estimates have placed costs for the very poorest households at up to 8% of their consumption.

Concluding Remarks

William Beveridge's 1944 call for 'Full Employment' demanded that the State take responsibility for shoring up aggregate demand deficiencies.[3] The underlying analysis was that business cycle fluctuations were not only costly but also that they would be reduced by appropriate stabilisation policies. Given the continuing incidence of words like recession, boom and bust in our common language, it seems odd that the latter day economists consider that the rational cost of business cycle fluctuations is low. It seems that the median household will only be prepared to give up a very small fraction of their annual consumption to eliminate their risk. One possibility is that they have already given up much of their income in taxes and stabilisation polices are already stabilising much of the fluctuations. Have we not paid the government to insure us already? This is one possible explanation that deserves more investigation.

There are two more. First, it could be that we should not be analysing the median household. It is quite possible that intra-household risk may not be fully shared and there may be quite different levels of risk to be absorbed by different income and wealth cohorts. So that even if the willingness of the median household to pay for insuring business cycle risk might be low, the risk faced by certain households may be significant and the risk premium not affordable. Secondly, and this kind of idea has increased in popularity in the past couple of decades, households may not plan according to strict rationality. I may buy a lottery ticket even if the expected return is negative and I may buy insurance even if the expected payoff is less than the premium. We may well not think of ourselves as average but exceptional. We may be more likely to believe that good things will happen and that bad things will happen to us more than they actually will.

References:

[1] Atkeson, A. and C. Phelan, (1994). Reconsidering the Costs of Business Cycles with Incomplete Markets, in *NBER Macroeconomics Annual*, S. Fischer and J. Rotemberg (eds) Cambridge, MA: MIT Press.

[2] Barlevy, G., (2005). The Costs of Business Cycles and the Benefits of Stabilisation, *Economic Perspectives*, Federal Reserve Bank of Chicago.

[3] Beveridge, W., (1944). *Full Employment in a Free Society*. London: Allen and Unwin.

[4] Burns, A. F. and W. C. Mitchell, (1946). *Measuring Business Cycles*, NBER Studies in Business Cycles No. 2.

[5] Chadha, J. S. and C. Nolan, (2002). A Long View of the UK Business Cycle, *National Institute Economic Review*, 182, pp.72-89.

[6] Gayer, A. D., W. W. Rostow and A. J. Schwartz, (1953). *The Growth and Fluctuation of the British Economy*, 1790-1850. Oxford: Clarendon Press.

[7] Imrohoroglu, A., (2008). Welfare Costs of Business Cycles. in *New Palgrave Dictionary of Economics*, S. N. Durlauf and L. E. Blume (eds) Basingstoke, Hampshire New York: Palgrave Macmillan.

[8] Krussell, P. and A. Smith, (1999). Revisiting the Welfare Effects of Eliminating Business Cycles, Review of Economic Dynamics, 2(1). pp. 245-272.

[9] Lucas, R. E., (1985). *Models of Business Cycles*. Yrjö Jahnsson Lectures, Oxford: Blackwell.

---

[3] We discussed this analysis in the lecture given in November 2014: The Great Depression and its Legacy at Gresham College.

[10] Samuleson, P., (1998). Opening Address, *Beyond Shocks: What Causes Business Cycles?* Federal Reserve Bank of Boston Conference Series 42, J. C. Fuhrer and S. Schuh (eds).

[11] Schumpeter, J. A., (1939). *Business Cycles*. New York: McGraw-Hill Book Company.

Gresham College
Barnard's Inn Hall
Holborn
London
EC1N 2HH
www.gresham.ac.uk