Chapter 31

Keynesian perspectives on business cycles

Applying a Vector-Autoregression time series approach with two kinds of shocks interpreted as demand and supply shocks, respectively, Blanchard and Quah (1989) found on the basis of quarterly US data 1950-87 that demand shocks explain more that two thirds of the fluctuations in output and even more of the fluctuations in unemployment. Working with a somewhat larger system and quarterly US data for 1965-1986, Blanchard (1989) summarized the results this way:

(a) Demand shocks explain most of the short-run fluctuations in output.

(b) Positive demand shocks are associated with gradual increases in nominal prices and wages.

(c) Supply shocks dominate the medium and the long run, and positive supply shocks are associated with decreases in nominal prices and wages (relative to trend).

Demand shocks may arise from shifts in the state of confidence, shifts in exports, shifts in government spending, shifts in liquidity preference, a sudden tightening of credit, and similar. Points (a), (b), and (c) lead to a Keynesian interpretation of macroeconomic fluctuations. A prevalent interpretation of point (a) is that nominal and relative price rigidities are present. Then, point (b) supports the view that even though prices in the major sectors of the economy respond only sluggishly, they *do* respond to cost push from changes in the level of economic activity. Finally, point (c) says that durable influences on output come from supply factors, such as the labor force, capital, and technological change.

The Keynesian understanding of economic fluctuations is that they emanate from "large" specific events, often connected to the financial sector. Some of these events trigger *virtuous circles* in the economic system as a whole while others trigger *vicious circles*. In continuation of the emphasis on nominal price stickiness, a crucial element in this understanding is the *refutation of Say's law*. This is the "law" claiming that "supply creates its own demand", cf. Chapter 19. At the microeconomic level, refutation of this doctrine leads to replacement of the Walrasian budget constraint with an *effective budget constraint*, when trade occurs outside Walrasian equilibrium.

31.1 A minimalist Keynesian medium-run model in discrete time

Notation:

$$
\begin{aligned}
y &\equiv \ln Y, \\
m &\equiv \ln M, \\
p &\equiv \ln P, \\
\pi_t &\equiv p_t - p_{t-1}, \\
i_t &= \text{ policy rate}, \\
\mu &= \text{ autonomous demand (reflecting perhaps) the state of confidence}, \\
\omega(\mu) &= \text{ interest spread (interest differential)},
\end{aligned}
$$

and

x_t and z_t are exogenous stochastic variables.

Output market equilibrium in reduced form:

$$y_t = \alpha y^e_{t+1} - \beta(i_t + \omega(\mu) - \pi^e_{t+1}) + \mu + x_t, \quad \omega(\mu) \geq 0, \; \omega'(\mu) < 0, \; \alpha > 0, \; \beta > 0, \tag{IS}$$

Phillips curve with both forward- and backward-looking elements $+ z_t$, (Ph)

but not clear how it should precisely be specified (weight of forward- versus backward-looking elements, non-linearity, "natural rate" or "natural range"?).

Taylor rule (inflation targeting):

$$i_t = \max\left[0, \hat{\imath} + \alpha_1(y_t - y^*) + \alpha_2(\pi^e_{t+1} - \hat{\pi})\right], \quad \hat{\imath} > 0, \alpha_1 \geq 0, \alpha_2 > 1, \tag{Taylor rule}$$

and, finally, specification of

expectation formation, (exp)

but far from clear what gives a good approximation: rational, adaptive, extrapolative, or "natural"[1], or possibly a mix with shifting weights due to shifting circumstances. To think of a "mix" is appropriate if different categories of agents in the economy form their expectations in very different ways (heterogeneity).

A traditional approach to the Phillips curve There are alternative ways of modelling the details of a Phillips curve. Building on Blanchard and Katz (1999), we here give a broad picture, starting with a wage Phillips curve.

Macroeconometric evidence indicates, in particular for the US after the Second World War, a negative relationship between the rate of change of wages and the unemployment rate:

$$\begin{aligned} w_t - w_{t-1} &= a + (p_{t-1} - p_{t-2}) - bu_t + z_t \\ &= g + (p_{t-1} - p_{t-2}) - b(u_t - u^*) + z_t, \end{aligned} \tag{31.1}$$

where $w = \ln W$, $p = \ln P$, u is the unemployment rate, and a and b are positive constants, whereas z_t is an error term. This is a *wage Phillips curve.* One interpretation is this. As appears in the second line of (31.1), the parameter a can be split into a sum of two terms, g which indicates the long-run growth rate in labor productivity, Y/N, and a term bu^*, where $u^* \equiv (a-g)/b$ (to be interpreted below as the NAIRU rate of unemployment). A straightforward reading of the role of the (lagged) inflation term, $p_{t-1} - p_{t-2}$, in (31.1) is that it represents *expected inflation.* Let p_t^e denote the expected price level in period t as seen from the end of period $t-1$ and let π_t^e denote the expected inflation rate, i.e., $\pi_t^e \equiv (P_t^e - P_{t-1})/P_{t-1} \approx p_t^e - p_{t-1}$. Then, according to the hypothesis of *static expectations of the inflation rate* we have

$$p_t^e - p_{t-1} = p_{t-1} - p_{t-2}. \tag{31.2}$$

In fact, if inflation follows a random walk (which the data does not reject[2]), this hypothesis is consistent with rational expectations.

Substituting (31.2) into (31.1) and ordering gives the expected change of the real wage as a decreasing function of unemployment:

$$w_t - p_t^e - (w_{t-1} - p_{t-1}) = g - b(u_t - u^*) + z_t. \tag{31.3}$$

In this way the empirical Wage Phillips curve, (31.1), is seen as reflecting an *expected-real-wage Phillips curve.* If expectations are not systematically wrong and the trend rate of unemployment is close to u^*, this says that real wages tend

[1] On.the hypothesis of "natural expectations", see Fuster, Laibson, and Mendel (JEP, 2010).
[2] See Hendry (2008).

in the long run to grow at the same rate as labor productivity, Y/N. The data for the US roughly confirms this picture. Consequently, a first interpretation of u^* is that it is that rate of unemployment which is consistent with real wages tending to grow at the same rate as labor productivity.

Whatever the interpretation of (31.1), it can under a certain condition be transformed into a price Phillips curve. Suppose prices are formed by a more or less constant mark-up on marginal cost, $P_t = (1+\mu)W_t/A_t$, where A_t is labor productivity. Then roughly the price inflation rate equals the wage inflation rate minus the productivity growth rate,

$$p_t - p_{t-1} = w_t - w_{t-1} - g.$$

Substituting this into (31.1) gives a standard *backward-looking Phillips curve*

$$p_t - p_{t-1} = p_{t-1} - p_{t-2} - b(u_t - u^*) + z_t. \tag{31.4}$$

Thus, if $u_t < u^*$, inflation increases, and if $u_t > u^*$, inflation decreases. This corresponds to the interpretation of u^* as the NAIRU (for non-accelerating-inflation-rate of unemployment) in the sense of that rate of unemployment which is consistent with a constant inflation rate (other names for u^* are the "natural" or the "structural" rate of unemployment).

As discussed by Blanchard and Katz (1999), the wage Phillips curve (31.1) fits European data less well than US data. And at the theoretical level it is in fact not obvious why a Phillips curve should hold in the first place. According to the theories of the functioning of labor markets (efficiency wages, social norms, search theories, and bargaining) it is the *level* of the expected real wage, rather than the expected change in the real wage, that is negatively related to unemployment. Theory thus predicts a *wage curve:*

$$w_t - p_t^e = \beta v_t + (1-\beta)\alpha_t - bu_t + z_t, \tag{31.5}$$

where β is a constant $\in [0,1]$, v_t is the *reservation wage* (the minimum real wage at which the worker is willing to supply labor), and α_t a measure of labor productivity.

By reasonable hypotheses about how the reservation wage depends on the actual real wage (in the previous period) and on productivity, a *level* formulation as in (31.5) *may* be consistent with a *change* formulation as in (31.1). Blanchard and Katz (1999) find such consistency to be plausible for US labor markets, but not for the typical European labor market with more influential labor unions, more stringent hiring and firing regulations, and perhaps also a greater role of the underground economy. An interesting implication of this theory is that in Europe, the NAIRU should be sensitive to permanent shifts in factors such as the level of energy prices, payroll taxes, or real interest rates, whereas in the US it should not.[3]

31.2 Vicious and virtuous circles

As mentioned, a characteristic feature of the Keynesian approach to business cycle fluctuations is the emphasis on the sometimes *vicious*, sometimes *virtuous circles* that arise, due to production being in the short term demand-determined rather than supply-determined. A vicious circle may for example come about in the following way.

Suppose that during an economic boom a housing price bubble evolves. Sooner or later the bubble bursts, collaterals for bank loans loose value (the *balance sheet channel*), defaults occur, confidence is shaken, credit is *squeezed*, and further defaults occur.[4] The financial crisis spills over to the goods market in the form of an adverse demand disturbance leading to a contraction of production and employment. The fired workers with less income buy fewer consumption goods (in particular fewer durable consumption goods). The process tends to be self-reinforcing in that the fear of being fired increases *precautionary saving*.

Seeing their demand curves continue the inward movement, firms cut production further. The utilization rate of capital equipment falls and so does average and marginal q. The fall in consumption is thus not offset by firms' investment being stimulated, rather the opposite. Firms' access to credit is cut down further as the balance sheets deteriorate. An economic recession or depression may develop if not offset by contra-cyclical monetary and/or fiscal policy.

There are several self-reinforcement mechanisms that bring these "circles" forth, whether they are negative, as above, or positive. Below we list six examples of such mechanisms. We describe them in their negative mode, that is, when they lead to *vicious circles*. They could just as well, however, be described in their positive mode as when they lead to virtuous circles and thereby a boom.

1. *The spending multiplier* (Kahn 1931, Keynes 1936). Recall that a *multiplier* is the ratio of a change in an endogenous variable, here output or employment, to a change in an exogenous variable, for example an autonomous part of private investment or government spending. A decrease in an autonomous demand component leads to a decrease in production and income, which further reduces demand. The government spending multiplier is larger in a depression, especially in a liquidity trap because there will be no financial

[3] Different versions of the so-called *new-Keynesian Phillips curve* are presented in, e.g., Mankiw (2001) and Gali (2015), but not discussed here.

[4] Below we elaborate on the terms in italics.

crowding out and typically a lot of available labor. Households' and firms' *precautionary saving* (see Section 30.4) aggravates the downturn.[5]

2. *Destabilizing price flexibility* (Keynes, Mundell, Tobin). When some nominal price and wage rigidity is present, more flexibility may be destabilizing. Suppose there is an adverse shock to investor's and firms' general long-term confidence which leads to a downturn of investment, aggregate demand, production, and employment. Through the Phillips curve mechanism, inflation and expected inflation also go down. Will high price flexibility be a good or a bad thing? Under a "passive" monetary policy (the k percent rule), comparatively high price flexibility (though less than "full") may turn the incipient recession into a downward wage-price spiral rather than a transitory dip. This is because opposing effects on aggregate demand are in play, giving rise to a *centripetal force* and a *centrifugal force*. On the one hand, the fall in inflation increases real money supply and lowers the nominal rate of interest, thereby stimulating aggregate demand. And in an open economy net exports are stimulated. On the other hand, the fall in *expected* inflation raises the *real* rate of interest,

$$r = i + \omega - \pi^e,$$

for a given short-term nominal rate of interest i (the policy rate) and a given interest differential, $\omega \geq 0$, thereby reducing demand. Depending on the circumstances, this effect may be the strongest and lead to a self-sustaining economic contraction. In particular this may happen, when the nominal rate of interest is already low and therefore near its floor, the "slightly below zero" bound.[6]

3. *The balance sheet channel* (Kiyotaki and Moore, 1997, Bernanke et al., 1999, Eggertsson and Krugman, 2012). Suppose an adverse shock reduces the net worth of credit-constrained borrowers (entrepreneurs and households), whose assets serve as collateral for loans. This depresses aggregate demand

[5] Formally, a *multiplier* is the ratio of a change in an endogenous variable, here output or employment, to a change in an exogenous variable, for example autonomous government spending.

[6] Nominal interest rates cannot fall much below zero, since potential lenders would then prefer holding cash rather than assets paying a negative interest rate.

The scenario described may take the even more pregnant form of a *deflationary spiral* leading to ever-widening economic crisis. The Great Depression in the US in the 1930's is a conspicuous example and the problems in Japan since the early 1990s also have affinity with this. As witnessed by the repercussions of the global financial crisis 2007-09, also under a countracyclical monetary policy, like the Taylor rule, may the lower bound on the short-term nominal interest rate be reached and thereafter such a vicious spiral arise.

in two ways. Because of the reduced wealth and precautionary saving, consumption is decreased. In addition, if expected to persist, the reduced net worth is likely to lead to a credit contraction. In need of liquidity some agents are forced to sell illiquid assets at "fire sale" prices, thereby further reducing the net worth and credit worthiness of debtors. This means less borrowing, faster debt repayment, and thereby less capital investment and consumption.

4. *The bank lending channel* (Bernanke and Blinder, 1988, 1992). If an economic downturn is on the way, banks may perceive that the riskiness of loans has increased. A *credit squeeze* vis-a-vis other banks and the non-bank public may result whereby the spread between the short-term nominal interest rate on, say, government bonds, and the interest rate that the ultimate borrowers must pay is increased. This limits capital investment and spending on durable consumption goods, thus reinforcing the economic downturn.

5. *Coordination failures* and *multiple equilibria.* There are circumstances, e.g., "spillover complementarity", where more than one general equilibrium is possible. Universally held pessimistic expectations lead to prudent actions that sum to a low-level outcome, thus confirming the pessimistic expectations. But if agents held optimistic expectations, they would make confident upbeat decisions. Aggregate demand would boom, thus confirming the expectations that brought it about in the first place (see Heller 1986, Kiyotaki 1988, Xiao, 2004). As expressed by Wren-Lewis (2015):

   "The largest component of aggregate demand is consumption, and consumption depends on expected income, which can depend itself on actual output, and therefore on aggregate demand. The macroeconomy is therefore set up to allow self-fulfilling multiple equilibria".

6. *Hysteresis.* The described demand-side dynamics may interact with the supply side. This occurs when the initial creation of unemployment, through a de-qualification or discouragement effect on the unemployed or through insider-outsider wage-setting behavior, turns a spell of unemployment into long-term unemployment. Such a phenomenon, where unemployment in the longer run depends positively on unemployment in the short run, is called *unemployment hysteresis.*[7] This has implications for the trade-off between short-run benefits of a deficit-financed expansionary fiscal policy

[7] See Blanchard and Summers (1987), Blanchard (1990), and DeLong and Summers (2012). A corresponding *virtuous* hysteresis can arise through the qualification or learning-by-doing effect of being employed. More generally on hysteresis, see Fiorillo (1999).

in a liquidity trap and possible longer-run costs in the form of a higher government debt.

More generally we say that *medium-run hysteresis* is present if the current state of the economy affects the state in the medium run in “the same direction”. And *long-run hysteresis* is present if the current state of the economy affects the state in the long run in “the same direction”.

One factor contributing to the vicious circles under the headings 1 and 5 is the phenomenon of *precautionary saving* to which we now turn.

31.3 Precautionary saving

We say that *precautionary saving* is present if increased uncertainty, everything else equal, results in increased saving.

In the first years after the crash at the New York stock exchange in 1929 a sharp fall in private consumption and investment occurred. Many economists argue that this should be seen in the light of the fact that the consumption/saving decision is sensitive to increased uncertainty. Similarly, the international financial crisis, triggered by the subprime mortgage crisis in the US in 2007, created a massive worldwide economic recession 2008- (the “Great Recession”). In this downturn again precautionary saving is likely to have played an important role. If people feel more uncertain about what is going to happen, they tend to be more prudent and increase their saving in order to have a “buffer-stock”. But this may aggravate the negative spiral of falling aggregate demand and production.

To clarify the issue, we first consider a simple model of a household’s consumption/saving decision under uncertainty. Second, we discuss the possible macroeconomic implications and relate the discussion to the different business cycle “schools”.

31.3.1 Consumption/saving under alternative forms of uncertainty

Consider a given household facing uncertainty about future labor income and capital income. For simplicity, assume the household supplies one unit of labor inelastically each period. The household can never be sure whether it will be able to sell that amount of labor in the next period. As seen from period 0, the decision problem is:

$$\max E_0 U_0 = E_0[\sum_{t=0}^{T-1} u(c_t)(1+\rho)^{-t}] \quad \text{s.t.} \tag{31.6}$$

$$c_t \geq 0, \tag{31.7}$$

$$a_{t+1} = (1+r_t)a_t + w_t n_t - c_t, \qquad a_0 \text{ given}, \tag{31.8}$$

$$a_T \geq 0. \tag{31.9}$$

where $u' > 0$ and $u'' < 0$ (so there is risk aversion). The rate of time preference w.r.t. utility is $\rho > -1$ (usually $\rho > 0$ seems realistic, but here the sign of ρ is not important). We think of "period t" as the time interval $[t, t+1)$. Hence, the last period within the planning horizon T is period $T-1$. Real financial wealth is denoted a_t and w_t (> 0) is the real wage, whereas n_t is the exogenous amount of employment offered to the household by the labor market in period t, $0 \leq n_t \leq 1$.[8] The (net) real rate of return on financial wealth is called r_t (> -1). The symbol E_0 stands for the expectation operator, conditional on the information available in period 0. This information includes knowledge of all relevant variables up to and including period 0. There is uncertainty about future values of r, w, and n, but the household knows the stochastic processes that these variables follow.[9] The risk associated with the uncertainty is assumed to be not insurable.

There are two endogenous variables, the control variable, c_t, and the state variable, a_t. The constraint (31.7) defines the "control region", whereas (31.8) is the dynamic budget identity, and (31.9) is the solvency condition, given the finite planning horizon T. The decision as seen from period 0 is to choose a concrete *action* c_0 and a set of *contingency plans* $c(t, a_t)$ about what to do in the future periods. This decision is made so that expected discounted utility, $E_0 U_0$, is maximized. We call the function $c(t, a_t)$ a contingent plan because it tells what consumption will be in period t, *depending* on the realization of the as yet unknown variables up to period t, including the state variable a_t. To choose c_0 in a rational way, the household must take into account the whole future, including what the optimal conditional actions in the future will be.

Letting period t be an arbitrary period, i.e., $t \in \{0, 1, 2, ..., T-1\}$, we rewrite

[8] More generally, $w_t n_t$ could be replaced by y_t, interpreted as any kind of *exogenous* income, say an uncertain pension.

[9] Or at least the household has beliefs about these processes and calculates subjective conditional probability distributions on this basis.

U_0 in the following way

$$\begin{aligned} U_0 &= \sum_{s=0}^{t-1} u(c_s)(1+\rho)^{-s} + \sum_{s=t}^{T-1} u(c_s)(1+\rho)^{-s} \\ &= \sum_{s=0}^{t-1} u(c_s)(1+\rho)^{-s} + (1+\rho)^{-t} \sum_{s=t}^{T-1} u(c_s)(1+\rho)^{-(s-t)} \\ &\equiv \sum_{s=0}^{t-1} u(c_s)(1+\rho)^{-s} + (1+\rho)^{-t} U_t. \end{aligned}$$

When deciding the "action" c_0, the household knows that in every new period, it has to solve the remainder of the problem in a similar way, given the information revealed up to and including that period. As seen from period t, the objective function is

$$E_t U_t = u(c_t) + (1+\rho)^{-1} E_t[u(c_{t+1}) + u(c_{t+2})(1+\rho)^{-1} + ...] \tag{31.10}$$

To solve the problem as seen from period t we will use the substitution method. First, from (31.8) we have

$$\begin{aligned} c_t &= (1+r_t)a_t + w_t n_t - a_{t+1}, \qquad \text{and} \\ c_{t+1} &= (1+r_{t+1})a_{t+1} + w_{t+1} n_{t+1} - a_{t+2}. \end{aligned} \tag{31.11}$$

Substituting this into (31.10), the problem is reduced to an essentially unconstrained maximization problem, namely one of maximizing the function $E_t U_t$ w.r.t. $a_{t+1}, a_{t+2}, ..., a_T$ (thereby indirectly choosing $c_t, c_{t+1}, .., c_{T-1}$). Hence, we first take the partial derivative w.r.t. a_{t+1} in (31.10) and set it equal to 0:

$$\frac{\partial E_t U_t}{\partial a_{t+1}} = u'(c_t) \cdot (-1) + (1+\rho)^{-1} E_t[u'(c_{t+1})(1+r_{t+1})] = 0.$$

Reordering gives the stochastic Euler equation,

$$u'(c_t) = (1+\rho)^{-1} E_t[u'(c_{t+1})(1+r_{t+1})], \qquad t = 0, 1, 2, ..., T-2. \tag{31.12}$$

This first-order condition describes the trade-off between consumption in period t and period $t+1$, as seen from period t. The optimal plan must satisfy that the current utility loss by decreasing consumption by one unit is equal to the discounted expected utility gain next period by having $1 + r_{t+1}$ extra units available for consumption, namely the gross return on saving one more unit. Considering $\partial E_t U_t / \partial a_{t+i}$ for $i = 2, 3, ..., T-t-2$, we get similar first-order conditions, in expected value, for each i.

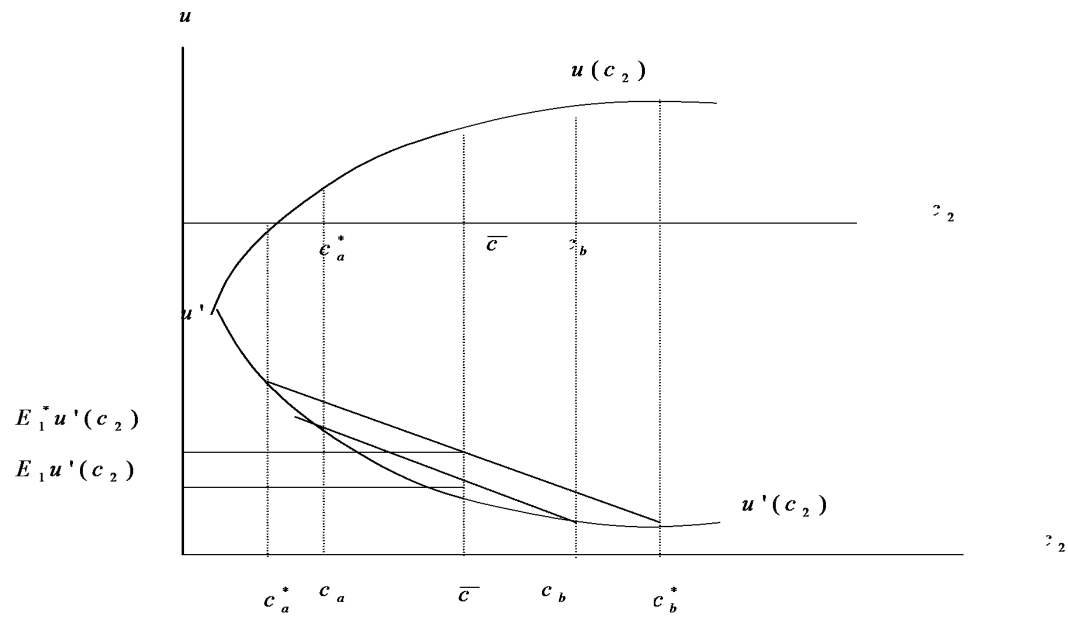

Figure 31.1: Graph of $u(c)$ (upper panel) and graph of $u'(c)$ (lower panel). The case $u'''(c) > 0$.

In the final period, given the solvency condition $a_T \geq 0$, the decision must be to choose $a_T = 0$ (the transversality condition). The alternative, $a_T > 0$, could always be improved upon by increasing c_{T-1} without violating the solvency condition. So, the optimal c_{T-1} satisfies

$$c_{T-1} = (1 + r_{T-1})a_{T-1} + w_{T-1}n_{T-1}. \tag{31.13}$$

First-order conditions only tell us about relative levels of consumption over time, however. The absolute level of consumption is determined by the condition that the current level of consumption, c_t, must be the highest possible consistent with: a) (31.12) for the given t; b) for t replaced by $t + i$, $i = 1, 2, ..., T - t - 2$, (31.12) in expected value as seen from period t, i.e., $E_t u'(c_{t+i}) = (1 + \rho)^{-1} E_t[u'(c_{t+i+1})(1+r_{t+i+1})]$; and c) (31.13) in expected value as seen from period t.

We will first consider the case where there is no uncertainty about the future real interest rates, only about future labor income.

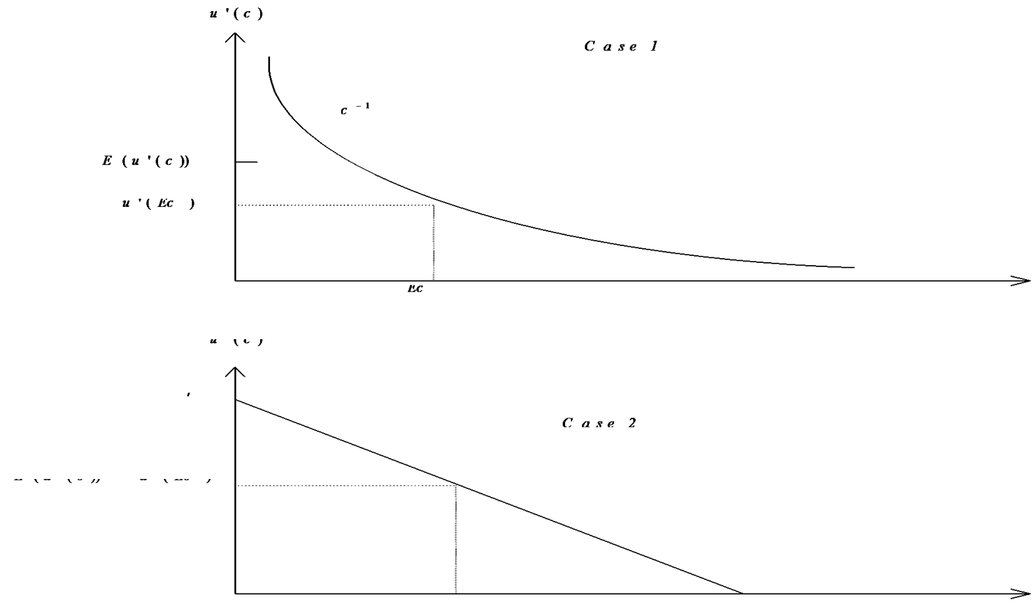

Figure 31.2: Graph of $u'(c)$ when $u(c) = \ln c$ (case 1) and when $u(c) = \eta c - \frac{1}{2}c^2$ (case 2).

Risk-free rate of return

Ruling out uncertainty about the future real interest rates, (31.12) reduces to

$$u'(c_t) = \frac{1 + r_{t+1}}{1 + \rho} E_t [u'(c_{t+1})], \qquad t = 0, 1, 2, ..., T - 2. \qquad (31.14)$$

It is natural to assume that higher wealth is associated with lower absolute risk aversion, $-u''/u'$. In that case, it can be shown that marginal utility u' is a strictly convex function of c, that is, $(u')'' > 0$.[10] But this implies that increased uncertainty in the form of a mean-preserving spread will lead to lower consumption "today" (more saving) than would otherwise be the case. This is what precautionary saving is about.

Fig. 30.1 gives an illustration. We can choose any utility function with $(u')'' > 0$. The often used logarithmic utility function is an example since $u(c) = \ln c$ gives $u'(c) = c^{-1}$, $u''(c) = -c^{-2}$ and $u'''(c) = 2c^{-3} > 0$. In the figure it is understood that $T = 3$ (so that the last period is period 2) and that we consider the decision problem as seen from period 1. There is uncertainty about labor income in period 2. It can be because the real wage is unknown or because employment is unknown

[10] See Section 31.4 below.

or both. Suppose, for simplicity, that there are only two possible outcomes for labor income y_t ($\equiv w_t n_t$), say y_a and y_b, each with probability $\frac{1}{2}$. That is, given a_2, there are, in view of (31.13), two possible outcomes for c_2:

$$c_2 = \begin{cases} c_a = (1+r_2)a_2 + y_a, & \text{with probability} = \frac{1}{2} \\ c_b = (1+r_2)a_2 + y_b & \text{with probability} = \frac{1}{2}. \end{cases} \tag{31.15}$$

Mean consumption will be $\bar{c} = (1+r_2)a_2 + \bar{y}$, where $\bar{y} = \frac{1}{2}(y_a + y_b)$.

Suppose c_1 has been chosen optimally. Then, with $t = 1$, (31.14) is satisfied, and a_2 is fixed, by (31.8) with $t = 1$. The lower panel of Fig. 30.1 shows graphically, how $E_1 u'(c_2)$ is determined, given this a_2.

Compare this outcome with a case of *higher uncertainty* in the form of a *mean-preserving spread.* By this is meant that the spread, $|y_b - y_a|$, is larger but the mean, $\bar{y}$, is unchanged. So, if a_2 remains unchanged, now the two possible outcomes for c_2 are c_a^* and c_b^*, while the average equals $\bar{c}$ as before. Fig. 31.1 illustrates. Owing to the strict convexity of marginal utility, the expected marginal utility of consumption is now greater than before, as indicated by $E_1 u'(c_2^*)$ in the figure. In order that (31.14) can still be satisfied, a lower value than before of c_1 must be chosen (since $u'' < 0$), hence, more saving occurs.

Yet, *this* lower value of c_1 is not the final outcome. Indeed, as soon as c_1 tends to be lowered, saving in period 1 tends to be raised. This means a higher a_2 so that the expected value of c_2 is in fact larger than $\bar{c}$ on the figure. This dampens, but does not eliminate, the effect of the mean-preserving spread on $E_1 u'(c_2)$. This expected value ends up somewhere between the old $E_1 u'(c_2)$ and $E_1 u'(c_2^*)$ in the figure. The conclusion is still that the new c_1 has to be lower than the original c_1 in order for the first-order condition (31.14) to be satisfied in the new situation.

If instead the increased uncertainty pertains to period 0, the effect is again to decrease current consumption to provide for a buffer.

What we see here is a manifestation of *precautionary saving:* increased saving as a result of increased uncertainty. In our example there is increased uncertainty about future labor income and as a result lower consumption "today". Consumption is postponed in order to have a buffer-stock. The intuition is that the household wants to be prepared for meeting bad luck, because it wants to avoid the risk of having to end up starving ("save for the rainy day").

Note that the mathematical basis for the phenomenon is the strict convexity of marginal utility, i.e., the assumption that $(u')'' > 0$. This implies $E(u'(c)) > u'(Ec)$ in view of Jensen's inequality. Case 1 in Fig. 30.2 shows the example $u(c) = \ln c$, i.e., $u'(c) = c^{-1}$.

If instead, $(u')'' = 0$, as with a quadratic utility function, then the graph for $u'(c_2)$ is a straight line (cf. case 2 in Fig. 30.2), and then precautionary saving can not occur. Indeed, a quadratic utility function can be written

$$u(c) = \begin{cases} \eta c - \frac{1}{2}c^2 & \text{if } 0 \leq c \leq \eta, \\ \frac{1}{2}\eta^2 & \text{if } c > \eta, \end{cases} \tag{31.16}$$

where $\eta > 0$. We have $u'(c) = \eta - c$ (a negatively sloped line), if $c < \eta$. At $c = \eta$, satiation occurs, and $u'(c) = 0$ for $c > \eta$. If in a given context we want the point of satiation to never be realized in practice, we may assume that η is "large".

The case of quadratic utility is an example of what is known as *certainty equivalence*. We say that certainty equivalence is present, if the decision under uncertainty follows the same rule as under certainty, only with actual values of the conditioning variables replaced by the expected values. Compare a situation where the relevant exogenous variables take on their expected values with probability one (certainty) with a situation where they do that with a probability *less* than one (uncertainty). If the decision is the same in the two situations, certainty equivalence is present. So, when there is certainty equivalence, the decision under uncertainty is independent of the degree of uncertainty, measured, say, by the variance of the relevant conditioning variable(s) for a fixed mean. Quadratic utility implies certainty equivalence. Yet, since (31.16) gives $u'' = -1 < 0$, a household with quadratic utility is risk averse. Hence, for precautionary saving to arise, more than risk aversion is needed.

What is needed for precautionary saving to occur is $u''' > 0$, i.e., "prudence". Just as the degree of (absolute) risk aversion is measured by $-u''/u'$ (i.e., the degree of concavity of the utility function), the degree of (absolute) prudence is measured by $-u'''/u''$ (i.e., the degree of convexity of marginal utility). The degree of risk aversion is important for the size of the *required compensation* for uncertainty, whereas the degree of prudence is important for how the household's *saving behavior* is affected by uncertainty.

Uncertain rate of return

We have just argued that strictly convex marginal utility is a necessary condition for precautionary saving. But, strictly speaking, it is not a sufficient condition. This is so because there may be uncertainty not only about future labor income, but also about the rate of return on saving.

Consider the case where, as seen from period t, r_{t+1} is unknown. Then the relevant first-order condition is (31.12), not (31.14). Now, at least at the theoretical level, the tendency for precautionary saving to arise may be dampened or even turned into its opposite by an offsetting factor. For simplicity, assume first that there is no uncertainty associated with future labor income so that the only uncertainty is about the rate of return, r_{t+1}. In this case it can be shown that there is positive precautionary saving if the *relative* risk aversion, $-cu''/u'$, is larger than 1 ("it is good to have a buffer in case of bad luck") and *negative* precautionary saving (a mean-preserving spread of the ex ante rate of return reduces saving) if the relative risk aversion is less than 1 ("get while the getting is good").

It is generally believed that the empirically relevant assumption from a macroeconomic point of view is that $-cu''/u' > 1$. Thus, increased uncertainty about the rate of return should lead to more saving. The resulting precautionary saving then *adds* to that arising from increased uncertainty about future labor income.

31.3.2 Precautionary saving in a macroeconomic perspective

Simple calculations as well as empirical investigations (for references, see Romer 2001, p. 357) indicate that precautionary saving is not only a theoretical possibility, but can be quantitatively important. A sudden increase in perceived uncertainty seems capable of creating a sizeable fall in consumption expenditure (in particular expenditure on durable consumption goods) and thereby in aggregate demand. According to a study by Christina Romer (1990), this played a major role for the economic downturn in the US after the crash at the stock market in 1929 (see also Blanchard, 2003, p. 471 ff.).

Note that the conception of precautionary saving as an important business cycle force does not fit equally well in all business cycle theories. In new-classical theories (since the 1980s, in practice the RBC theory) a lower propensity to consume is immediately and automatically compensated by higher investment demand and perhaps a larger labor supply and employment in the economy. According to the RBC model from the previous chapter, aggregate demand continues to be sufficient to absorb output at full capacity utilization. Higher uncertainty just leads to a change in the composition of demand, a manifestation of Say's law.

According to many empiricists, this story is contradicted by the data. Less consumption spending seems far form being automatically offset by higher investment spending. The Keynesian interpretation is that output is demand-determined in the short run. An adverse demand shock, triggered by a bursting housing price bubble, say, will, through precautionary saving, lead to a contraction of demand and therefore a downturn of production.

Also firms' behavior may in an economic crisis have aspects of precautionary financial saving. A deep crisis generates a lot of uncertainty: firms are unsure about what has happened and no one knows what actions to choose. The natural thing to do is to pause and wait until the situation becomes clearer. This entails a cutback in the plans for further purchase of investment goods. So on top of households' precautionary saving we have prudent investment behavior by the firms.

31.4 On the distinction between risk aversion and prudence*

We end this chapter with a more general account of basic concepts from the theory of decisions under uncertainty, including the concept of prudence. The aim is to clarify the distinction between the degree of risk aversion and the degree of prudence. We relate this distinction to commonly used utility functions such as the CARA, CRRA, and quadratic utility functions.

31.4.1 Risk aversion and risk premium

Let c be consumption and let E be the expectational operator. Consider a von Neumann-Morgenstern utility index $U = E[u(c)]$ where u is a twice continuously differentiable (sub-) utility function. Assume $u' > 0$. If $u'' < 0$, then the individual in question is said to be *risk averse*. Let $ARA(c)$ be the degree of *Absolute Risk Aversion* at consumption level c, i.e.,

$$ARA(c) \equiv -\frac{u''(c)}{u'(c)}.$$

For a risk-averse individual, this measure is a positive number.[11] As an example, suppose the utility function is

$$\text{CARA:} \qquad u(c) = -\alpha^{-1}e^{-\alpha c},$$

where α is a positive constant. For this function, $ARA(c) = \alpha > 0$, a constant (CARA stands for Constant Absolute Risk Aversion).

The economic significance of the ARA measure is that it is approximately proportional to the (required) *risk premium* (to be defined below). Let ℓ denote the "lottery" that the individual confronts, "lottery" in the sense of a random draw from the given probability distribution for c. For a risk-averse individual $u'' < 0$ (i.e., $u(c)$ is a strictly concave function) and therefore

$$E[u(c)] < u(Ec)$$

[11] The measure $ARA(c)$ is unaffected by an increasing linear transformation of u.

by Jensen's inequality. The *certainty equivalent* for the lottery ℓ is the number c^* satisfying

$$E[u(c)] = u(c^*). \tag{31.17}$$

In words, the certainty equivalent c^* is that *certain* consumption level which the individual is just willing to exchange for the lottery ℓ.

The *risk premium* for the lottery ℓ is defined as the number π satisfying

$$E(c) - \pi = c^*. \tag{31.18}$$

In words, the risk premium is the decrease in expected consumption that the individual is just willing to accept to get rid of the uncertainty and obtain a safe consumption level. Or, since $\pi = E(c) - c^*$, we may look at the matter from the opposite angle and define the risk premium as the increase in expected consumption that the individual requires to just accept an exchange of a safe consumption level c^* for the lottery ℓ.

Let $\bar{c} \equiv Ec$, i.e., $c = \bar{c} + \varepsilon$, where ε is white noise. Now, (31.17) and (31.18) imply

$$E[u(c)] = u[E(c) - \pi] = u(\bar{c} - \pi). \tag{31.19}$$

From this relation we can find an approximate value of π. As to the left-hand-side of (31.19), a *second*-order Taylor approximation of $u(c)$ gives

$$\begin{aligned} u(c) &\approx u(\bar{c}) + u'(\bar{c})\varepsilon + \frac{1}{2}u''(\bar{c})\varepsilon^2 \qquad \Rightarrow \\ E[u(c)] &\approx u(\bar{c}) + 0 + \frac{1}{2}u''(\bar{c})\sigma_\varepsilon^2, \end{aligned} \tag{31.20}$$

where $\sigma_\varepsilon^2 = Var(c) = Var(\varepsilon)$. As to the RHS of (31.19), a first-order Taylor approximation gives

$$u(\bar{c} - \pi) \approx u(\bar{c}) + u'(\bar{c})(-\pi) = u(\bar{c}) - \pi u'(\bar{c}).$$

Inserting this and (31.20) into (31.19) gives

$$\begin{aligned} & u(\bar{c}) + \frac{1}{2}u''(\bar{c})\sigma_\varepsilon^2 \approx u(\bar{c}) - \pi u'(\bar{c}) \ \Rightarrow \\ & \qquad \pi \approx -\frac{1}{2}\sigma_\varepsilon^2\frac{u''(\bar{c})}{u'(\bar{c})} = \frac{1}{2}\sigma_\varepsilon^2 ARA(\bar{c}) = \frac{1}{2}\sigma_\varepsilon^2 ARA(Ec). \end{aligned}$$

Hence, ARA evaluated at the consumption level $E(c)$ is approximately proportional to the risk premium.

It seems natural to suppose that as a individual becomes richer - higher $E(c)$ - she cares less and less about the risks she takes. This would say that $\pi'(E(c)) < 0$, i.e., π decreases - hence ARA decreases - as $E(c)$ increases. Therefore, the CARA utility function, defined above, does not seem very realistic. But CARA is just one member of a large family of convenient and more or less realistic utility functions that is called the HARA family.

The HARA family of utility functions is important for at least two reasons.[12] First, if labor income is "diversifiable" (so that the individual can sell shares against future labor income - which is not very realistic, it must be admitted), then it is possible to derive an explicit solution to standard optimum consumption and portfolio problems (as formulated in, e.g., Blanchard and Fischer, 1989, p. 280), if the utility function belongs to the HARA family. Second, the HARA family is the only class of concave utility functions which imply that the consumption function and the portfolio selection function become *linear* in financial wealth. The HARA family as a whole is described mathematically in Appendix B.

Here we shall just meet some prominent members of the family:

$$\text{Quadratic:} \qquad u(c) = \eta c - \frac{1}{2}c^2, \qquad 0 \leq c < \eta, \ \eta \text{ "large"}. \tag{31.21}$$

$$\text{CARA (or the exponential utility function):} \ \ u(c) = -\alpha^{-1}e^{-\alpha c}, \ \alpha > 0. \tag{31.22}$$

$$\text{CRRA with parameter } \theta > 0: \qquad u(c) = \begin{cases} \frac{c^{1-\theta}-1}{1-\theta}, & \text{if } \theta \neq 1, \\ \ln \ c & \text{if } \theta = 1, \end{cases} \tag{31.23}$$

where CRRA is an abbreviation for

31.4.2 The degree of prudence

The *degree of absolute prudence* is defined as the ratio

$$-u'''/u''.$$

As we saw, quadratic utility implies that marginal utility is linear in c (i.e., $u''' = 0$). Hence, in this case the degree of prudence is zero, and the phenomenon of precautionary saving does not arise. But still the quadratic function has $u'' = -1 < 0$, and therefore indicates risk aversion.

The CARA function features the desirable properties of risk aversion ($u'' < 0$) *and* prudence $-u'''/u'' > 0$. On the other hand, the CARA function implies that the required risk premium is constant (independent of wealth), which is probably not a realistic property. The CRRA function, however, has all three desirable properties (so there is no conflict between them).

[12]The name HARA stands for Hyperbolic Absolute Risk Aversion.

CLAIM 1. Given $u' > 0, u'' < 0$, and assuming u to be three times continuously differentiable, non-increasing ARA implies $u''' > 0$.

Proof. From the definition of ARA, we have

$$\begin{aligned} \frac{dARA}{dc} &= -\frac{u'u''' - (u'')^2}{(u')^2} = (\frac{u''}{u'})^2 - \frac{u'''}{u'} \qquad (31.24) \\ &= \frac{u''}{u'}\left(\frac{u''}{u'} - \frac{u'''}{u''}\right) = ARA(ARA + \frac{u'''}{u''}) \leq 0 \\ &\Rightarrow u''' > 0 \end{aligned}$$

since $ARA > 0$, and $u'' < 0$. □

We saw above that when the individual faces a larger future income risk, then, if $u''' > 0$, she has a tendency to consume less in the current period. In other words, precautionary saving tends to occur. The degree of absolute prudence, the ratio $(-u'''/u'')$, can be seen as a measure of the "degree of convexity" of marginal utility $u'(c)$.

The CRRA class of utility functions is characterized by the fact that the measure of *relative risk aversion*

$$RRA \equiv -c\frac{u''}{u'} \equiv c \cdot ARA = \theta$$

is constant (which explains the name CRRA for Constant Relative Risk Aversion). Obviously, this function has the property that ARA $(= \theta/c)$ is decreasing in c (as is desirable). Further, $u''' = (\theta + 1)\theta c^{-\theta-1} > 0$ (as expected from Claim 1).

The members of the CRRA class have the (sometimes inconvenient) property that when entering an additively time separable intertemporal utility index, the intertemporal elasticity of substitution becomes equal to $1/\theta$ and hence cannot vary without implying variation in the relative risk aversion measure RRA in the opposite way and in the same proportion. Unsatisfied with this property, ...TO BE CONTINUED

The HARA family is a much richer class, including the four standard cases shown in (31.21) - (31.23) above. By suitable adjustment of the parameters one can get a utility function with decreasing, increasing, or constant absolute or relative risk aversion. As an example, the *general* log utility function

$$u(c) = \ln(\eta + c) \qquad (31.25)$$

has decreasing, constant, or increasing RRA as η is negative, zero, or positive, respectively. Indeed, (31.25) has RRA $= c/(\eta + c)$. The case $\eta < 0$ may be interpreted in terms of a subsistence minimum, the subsistence minimum being $|\eta|$.[13]

The case $\eta > 0$ can be interpreted as that c refers to consumption of a luxury good.

31.5 Literature notes

(incomplete)

Paul Krugman's *The Return to Depression Economics* (Krugman 2000) reflects on the need for macroeconomic theory to include depression economics as one of its concerns.

The self-fulfilling prophesy investment theory by Kiyotaki (1988) and the inventory investment theory by Blinder ( ) are examples of business cycle theory emphasizing firms' investment.

Merton (1975).

31.6 Appendix

Jensen's inequality

Jensen's inequality is the proposition that when X is a stochastic variable, and the function f is *convex*, then

$$Ef(X) \geq f(EX)$$

with strict inequality, if f is *strictly* convex (unless X with probability 1 is equal to a constant). It follows that if f is *concave* (i.e., $-f$ is convex), then

$$Ef(X) \leq f(EX)$$

with strict inequality, if f is *strictly* concave (unless X with probability 1 is equal to a constant).

The HARA family of utility functions

Let $c \geq 0$ be consumption, and $u(c)$, $u' > 0$, $u'' < 0$, be a utility function entering a von Neumann-Morgenstern utility index. The measure of *absolute risk tolerance*, ART, is defined as the inverse of the measure of absolute risk aversion, ARA, that is

$$ART(c) \equiv \frac{1}{ARA(c)} \equiv -\frac{u'(c)}{u''(c)} > 0.$$

[13] We have $\lim_{c \to -\eta} \ln(\eta + c) = -\infty$.

A HARA utility function is defined as a utility function $u(c)$ with linear absolute risk tolerance, i.e., the requirement is that

$$ART(c) = \eta + \beta c, \tag{31.26}$$

where η and β are constant parameters[14]. Hence, we get the HARA family of utility functions by solving the second order differential equation

$$\frac{u''}{u'} = -\frac{1}{\eta + \beta c} \tag{31.27}$$

defined on the domain

$$\eta + \beta c > 0. \tag{31.28}$$

Depending on β, the solution is

$$u(c) = \begin{cases} \frac{(\eta+\beta c)^{1-1/\beta}}{\beta-1} + k, & \text{if } \beta \neq 0, \ \beta \neq 1 \\ \ln(\eta + c), & \text{if } \beta = 1, \\ -\eta e^{-c/\eta}, & \text{if } \beta = 0, \end{cases} \tag{31.29}$$

where k is an arbitrary constant (which can be chosen according to what is convenient).

(31.29) is the HARA family of utility functions. This family includes widely used functional forms as special cases: quadratic utility, the CRRA function, the log function, and the CARA function. Each of these, however, are often written in a slightly more convenient way. It is always allowed to add a constant to the function $u(c)$ and multiply by a positive constant (any increasing linear transformation of $u(c)$ will always represent the same von Neumann-Morgenstern preferences).

For example, when $\beta = -1$, $\eta > 0$, and $k = -\eta^2/2$, (31.29) gives

$$\text{the quadratic case:} \qquad u(c) = \eta c - \frac{1}{2}c^2, \qquad 0 \leq c < \eta.$$

When $\beta = 0$, hence $\eta > 0$ by (31.28), (31.29) gives

$$\text{the CARA or the exponential case:} \qquad u(c) = -\alpha^{-1}e^{-\alpha c}, \ \alpha \equiv 1/\eta > 0.$$

Letting $\theta \equiv 1/\beta$, where $\beta > 0$, $\beta \neq 1$, $\eta = 0$, and $k = -\beta^{-\theta}/(1-\theta)$, (31.29) gives (multiply through by β^{θ})

$$\text{the CRRA case:} \qquad u(c) = \frac{c^{1-\theta} - 1}{1-\theta}, \qquad \theta > 0, \ \theta \neq 1.$$

[14] The HARA definition can be generalized to include cases where η and β are functions of time.

When $\beta = 1$ and $\eta = 0$, (31.29) gives

$$\text{the (standard) logarithmic case:} \qquad u(c) = \ln\ c.$$

As seen by (31.26), the sign of β determines whether risk tolerance is increasing ($\beta > 0$), constant ($\beta = 0$), or decreasing ($\beta < 0$). Increasing risk tolerance - decreasing absolute risk aversion - is considered as the most realistic case. Hence, the CARA utility function (which has $\beta = 0$) should be interpreted as only a theoretical benchmark case which is sometimes mathematically convenient, but probably not realistic. The quadratic utility function is even less plausible (since it has β negative and, in contrast to the other standard functions, it has $u''' \equiv 0$).

Further unfinished notes: HARA $\Rightarrow$ Engel curves are linear $\Rightarrow$ Gorman's aggregation criteria are satisfied (see Basetto and Benhabib, RED 9, 211-23, 2006, and Pollak, Additive utility functions and linear Engel curves, RES, 38 (4), 401-14.

Key terms

mean-preserving spread
degree of risk aversion
risk premium
degree of prudence
precautionary saving
certainty equivalence
vicious circles
virtuous circles

31.7 Exercises