Long-term policy-making

Ivar Ekeland
Canada Research Chair in Mathematical Economics, UBC

July 11, 2008

Abstract

These are lecture notes for the first four lessons I gave at the PIMS Summer School on "Perceiving, measuring and managing risk" held at UBC, June 30 - July 11,m 2008. The fifth lecture is available separately. A complete version of these lectures, including a detailed bibliograpy and proofs, will be available by July 31, 2008.

Acknowledgement 1 *The support of PIMS and the discussions with the Summer School participants is gratefully acknowledged.*

1 The general purpose

1.1 What is the long term ?

These lectures deal with the economic aspects of long-term policy-making. As the historical notes will show, the basic problem has been around for many years, in fact since the beginnings of economic theory, and there is a vast literature on this subject. To state it as simply as possible, a decision-maker, either an individual or a collective entity (the government), is to make a decision today (time $t = 0$), the consequences of which will kick in only at a (much) later time $t = T$, and he/she has to weigh the immediate benefit of that decision against the future costs. Alternatively, the cost occurs today and the benefits at time T, and the question is then how much cost the decision-maker is willing to bear today in order to reap the benefits at time T.

What has changed, however, since these early days, is the time horizon. Up to very recently, what the economist, the engineer or the politician would consider long-term would be in the range 10 to 30 years (note for instance that the US Treasury does not issue bonds with a longer maturity than 30 years). Anything beyond that was considered beyond the horizon - just like accelerating galaxies slip beyond the boundary of the observable universe. This has changed in recent years, where the consequences of our actions beyond that horizon have become part of the agenda. Here are two examples:

- the lifetime of a nuclear plant is 40 to 60 years, after which it will have to be decommissioned and the site reclaimed, at a considerable cost, which has to be factored in the investment decision
- the Stern Review on Climate Change states that the course for the next 50 years is set: the inertia of the physical and biological system governing the Earth climate is such that the consequences of any policy we enact today will not be felt before 50 years have elapsed. The question is what happens after that, and the Stern Review depicts alternate scenarios spanning the 50 to 200 years period.

In these lecture notes, we will define the long term as the 50 to 200 years range.

1.2 What are the difficulties of long-term policy-making ?

There are two main features which set long-term decisions apart from short and middle-term ones. The first one is *high uncertainty*. This comes in two different guises:

- the predictable outcomes, that is those for which probabilities can be set, have a very high dispersion. For instance, the Stern Review states that, with a probability of 95%, under the business as usual scenario, the loss of GDP to the world econony 200 years from now will be in the range 2% to 35%
- but there also non-predictable events, which Stern calls "bifurcations", such as the cessation or the thermo-haline convector which runs the Gulf Stream. or the melting of the West Antarctic ice sheet. We are in no position to assign probabilities to such event but we know (a) that they can occur,and (b) that their consequences would be catastrophic

Let me mention "en passant" that the fact of global warming as such no longer is part of the uncertainty: it is now certain that it is occuring. At the time of this writing, there is a 50% chance that, for the first time in recorded history, the North Pole will be on open water.

The second difficulty is *non-commitment*. Whatever policy we enact today will presumably have to be adhered to until the desired consequences are achieved, 50 to 200 years from now, when we (or whoever has decided on these policies) no longer is there to carry them out. This means that we have to rely on future generations (and future governments) to carry them out when we are gone. There is no way we can commit unborn generations and whoever rules the planet one hundred from now to anything: they will do as they please. Whatever policy we design now for the long term has to answer the question: one hundred years from now, when the powers that be are supposed to implement these policies, is there a reasonable chance that they will do it ?

1.3 Is economic theory relevant ?

There are, of course, the economic approach to these problems is not the only possible or relevant one. Clearly, there are ethical considerations. As Keynes famously said, in the long term we are all dead. Do we care what happens after that ? Some people don't: this is the "après moi le déluge" philosophy, which has quite a number of proponents in academic and government circles. Most people do, most famously Adam Smith. because of ethical considerations towards future generations and the planet itself. In addition to ethical considerations, there are political agendas, with immediate gains or losses for decision-makers which clearly preempt any long-term concern.

Even in the presence of ethical or political considerations, traditional cost-benefit analysis has shown itself to be a useful tool, if only to clarify the issues. To attach figures to policies does make a lot of difference. Part of the impact of the Stern Review is due to the fact that it came up with the conclusion that, under a policy of business as usual, climate change would cost 10% of GDP per year, while the cost of prevention stood at less than 1%. This is the kind of argument to which politicians and business leaders pay attention, and the scientific community should make every effort to speak to them in their own language.

Even so, applying cost-benefit analysis to long-term policies is by no means straightforward. In practice, this means discounting future benefits at a certain rate, say r. A natural choice for r would be the market rate of interest, especially for the longest maturity, which today stands at 4.6% (rate of 30-years US Treasury bond). In the following table, we give the present value of $1,000,000$ \$ at 50, 100, and 200 years, for interest rates of 10%, 4.6% and 1.4%, which is the value that the Stern Review took:

	50 ys	100 ys	200 ys
10%	$8,519$	73	< 0.00
4.6%	$105,540$	$11,140$	124
$1,4\%$	$499,000$	$249,000$	$62,000$

Clearly, an interest rate of 10% just wipes out the long term. The current market rate of 4.6% does somewhat better, but the rate of $1,4\%$ really make future events loom large. It is evident that the conclusions of the Stern report heavily depend on this choice of the interest rate, and that they would have been entirely different if, for instance, Stern had chose to discount at market rates. So the question now is: what justification, if any, is there for choosing such a low rate ? This is the question which we will address now.

We will proceed by first giving an exposition of the standard model of economic growth, and then modifying it to incorporate the specific constraints of long-term policy (high uncertainty, non-commitment), and the special concerns associated with environmental issues and intergenerational equity.

2 The standard model of economic growth

2.1 Firms and consumers

This model is described in the opening chapters of most graduate textbooks in macroeconomics. There is *a single good* in the economy, which can be either consumed (in which case it is denoted by C) or used to produce more of the same good (in which case it is called capital and denoted by K) . This good is produced by a large (fixed) number of identical firms in perfect competition, so that if functions as a single firm which is a price-taker. The global production function is:

$$Y = F(K, AL)$$

where K is the total capital invested in the economy, L is the labour force, and A is the productivity of labour (which will eventually depend on time, $A(t)$, to reflect technological progress). It will be assumed that there are *constant returns to scale*, so that setting $y = Y/AL$ (production per unit of effective labour) and $k = K/AL$ (capital per unit of productive labour), we have:

$$y = f(k)$$

where $f(k) := F(k, 1)$ is the reduced production function. It will be assumed to be concave, increasing, with:

$$\begin{aligned} k &\geq 0, \quad f(k) > 0, \quad f(0) = 0 \\ f'(0) &= +\infty, \quad f'(k) \longrightarrow 0 \text{ when } k \longrightarrow \infty \end{aligned}$$

The population consists of identical individuals. Total consumption is C; up to a constant factor, it is also the consumption of each individual. Using the same scaling for production and for consumption, we find that the relevant variable is $c = C/AL$, the consumption per unit of effective labour. The consumption per individual is $C/L = Ac$, and this is the variable which enters the individual's utility function u. In the sequel, we will take the following specification:

$$\begin{aligned} u(x) &= \frac{x^{1-\theta}}{1-\theta} \text{ for } \theta > 0, \ \theta \neq 1 \\ u(x) &= \ln x \text{ for } \theta = 1 \end{aligned}$$

so that the utility of each individual alive at time t is:

$$u(A(t)\, c(t)) = \frac{A(t)^{1-\theta}}{1-\theta} c(t)^{1-\theta}$$

2.2 The growth of the economy

We will now consider (and compare) different consumption scenarios $C(t)$. To do that, we assume that the population consists of N *identical dynasties*, each of which will be treated as a single, infinite-lived individual, who consumes $A(t)\,c(t) = C(t)\,/L(t)$ at time t. Note that the dynasty considers, not how much it consumes at time t, which would be $C(t)\,/N$, but how much its average member consumes at time t, which is indeed $C(t)\,/L(t)$; the difference could be quite significant, since the dynasties grow at the same rate as the population.

Each dynasty, or rather this single, infinite-lived individual, has a *pure rate of time preference* ρ. This means that, given the choice between consuming $c(0)$ at time $t = 0$ and $c(t)$ at time $t > 0$, he/she will be indifferent if and only if:

$$u(c(0)) = e^{-\rho t} u(c(t))$$

Let us note right now (and we will expand on this later) that ρ is NOT an interest rate.

The economy is driven exogeneously by technological progress and population growth, which happen at the constant rates g and n. We set:

$$\begin{aligned} A(t) &= A(0)\,e^{gt} \\ L(t) &= L(0)\,e^{nt} \end{aligned}$$

so that the utility of the representative consumer at time t is:

$$\frac{A(0)^{1-\theta}}{1-\theta} c^{1-\theta} e^{g(1-\theta)t}$$

We now imagine a benevolent and omniscient planner, who wants to maximize the intertemporal welfare of the representative citizen. He/she will consider the following problem:

$$\begin{aligned} & \max \frac{1}{1-\theta} \int_0^{\infty} c^{1-\theta} e^{g(1-\theta)t} e^{-\rho t} dt \qquad \text{(Ramsey)} \\ \frac{dk}{dt} &= f(k) - c - (n+g)\,k, \quad k(0) = k_0 \end{aligned}$$

The second equation represent the balance equation between savings and consumption. It states that at every moment t, (scaled) production $f(k)$ is fully allocated between immediate (scaled) consumption c and (scaled) capital investment dk/dt, the correction term $(n+g)\,k$ being there to take into account growth of population and technological progress. Of course k_0 is the initial capital.

Solving (**??**) leads to the following result:

(a) There is a single k_∞, called the equilibrium value of capital, which solves the equation:

$$f'(k_\infty) = \rho - \theta g - n$$

(b) Problem (**??**) has a single solution $k(t)$, which has the property:

$$k(t) \longrightarrow k_\infty \text{ when } t \longrightarrow \infty$$

(c) The corresponding consumption $c(t)$ along the optimal path also converges:

$$c(t) \longrightarrow c_\infty = f(k_\infty) - (n+g)\, k_\infty$$

In equilibrium, when $k(0) = k_\infty$ and $c(0) = c_\infty$, we have $C(t) = C(0)\, e^{(n+g)t}$, $K(t) = K(0)\, e^{n+g}$ and $C(t)/L(t) = A(0)\, e^{gt}$, so that total consumption and total production are growing at the rate $n+g$, while consumption per head is growing at the rate g.

2.3 Some history

From the beginning of economics as a separate science, it was apparent that individual choices between present and future rewards were driven by some kind of time preference: individuals prefer to enjoy goods sooner rather than later. That theme was developed by John Rae (1834), Boehm-Bawerk (I889), Irving Fisher (1930), as a psychological trait of human nature. In 1960, Tjalling Koopmans showed that impatience can be derived from benign assumptions on preferences. In other words, preferences in economic theory are usually ascribed to immediate consumption c, resulting in utility functions $u(c)$. If one now considers consumption schedules, $c(t)$ for $t \geq 0$ (possibly discrete), and tries to write down a reasonable set of axioms that preferences should satisfy, one is inevitably led to time preference as a logical consequence.

Since time preference is well established, the next question is how to translate it into a mathematical model. This is done in the standard model by introducing the parameter $\rho > 0$: the larger ρ, the more impatient the consumer. If $\rho = 0$, the consumer is indifferent between immediate and deferred consumption; in other words, he exhibits no impatience at all.

The idea of setting up the question of economic growth as an optimisation problem is due to Frank Ramsey (1928). Interestingly, he chose $\rho = 0$ as his preferred option: "we do not discount later enjoyments in comparison with earlier ones, a practice that is ethically indefensible and arises merely from the weakness of the imagination". He did, however, treat the case $\varrho > 0$ as well, and that became the standard of the industry, following Samuelson (1937)

References

[1] Aghion, Philippe and Howitt, Peter (1998) "Endogeneous growth theory", MIT Press

[2] Barro, Robert and Sala-i-Martin, Xavier (1995) "Economic growth", McGraw and Hill

[3] Koopmans, Tjalling (1960), "Stationary utility and impatience", Econometrica 28 (2), p. 287-290

[4] Ramsey, Frank (1928) "A mathematical theory of saving", Economic Journal 38, p. 543-559

[5] Romer, David (1996), "Advanced macroeconomics", McGraw and Hill

[6] Samuelson, Paul (1937) ""A note on measurement of utility", Review of Economic Studies 4, p. 155-61

A Determinants of the interest rate.

We shall use the standard model as a benchmark, and introduce successive modifications.

A.1 The classical theory

Consider an infinite-lived individual, with utility function u, pure rate of time preference $\rho > 0$, and who is facing a schedule of consumption $c(t)$, leading to an overall utility of:

$$\int_0^\infty e^{-\rho t} u(c(t))\, dt$$

Let us ask ourselves how much immediate consumption $\Delta c(0)$ he/she would be willing to forgo in order to increase its consumption by $\Delta c(t)$ at some later time $t > 0$. Assuming these are small quantities, we can work on the margin, and we get the relation:

$$\begin{aligned} e^{-\rho t} u'(c(t))\, \Delta c(t) - u'(c(0))\, \Delta c(0) &= 0 \\ \frac{\Delta c(t)}{\Delta c(0)} &= \frac{u'(c(0))}{u'(c(t))} e^{\rho t} \end{aligned}$$

We define the *interest rate on consumption*, $r(t)$, by the following relation:

$$\frac{\Delta c(t)}{\Delta c(0)} = \exp \int_0^t r(s)\, ds$$

which leads to the following:

$$r(t) = \rho - \frac{d}{dt} \ln u'(c(t)) = \rho - \frac{u''(c(t))}{u'(c(t))} \frac{dc}{dt}(t)$$

This is most conveniently rewritten as follows:

$$\begin{aligned} r(t) &= \rho + \left(-c(t) \frac{u''(c(t))}{u'(c(t))} \right) \left(\frac{1}{c(t)} \frac{dc}{dt}(t) \right) \\ &= \rho + \eta(c(t))\, G(t) \end{aligned}$$

where:

- $\eta(c) := -cu''(c)/u'(c)$ is a positive parameter (because u is concave), usually called the relative risk aversion; in this context, it would be more relevant to call it the relative satiation. It usually depends on the level of consumption c. In the special case of power utilities, $u(c) = c^{1-\theta}/(1-\theta)$, it is constant and equal to θ

- $G(t) := (dc/dt)/c(t)$ is the rate of growth of the economy

Putting all this together, in the framework of the standard model, where the rate of growth of average consumption is constant and equal to g, we find that the interest rate on consumption is constant and equal to:

$$r = \rho + \theta g \qquad \text{(Interest Rate)}$$

This important formula is the benchmark for determining the interest rate on consumption, and is generally accepted in the economic literature. For instance, the Stern report takes $\rho = 0.1\%$, $\theta = 1$ and $g = 1.3\%$, yielding $r = 1.4\%$. Most of its critics claim that it is too low, and take $\rho = 2\%$, $\theta = 2$ and $g = 2\%$ as more reasonable numbers, yielding $r = 6\%$. We will discuss these claims, and bring more economic arguments to bear, in the sequel. Meanwhile, let us make some observations:

- as soon as there is growth in the economy ($g > 0$), we have $r > \rho$. For instance, as Ramsey found out, we can have positive interest rate $r > 0$ even if the pure rate of time preference is zero, $\rho = 0$.

- the interest rate *rises* with the growth rate g. For instance, setting $\rho = 2\%$ and $\theta = 2$, we get $r = 6\%$ if $g = 2\%$ and $r = 10\%$ if $g = 4\%$. Why is that so ? Well, ask yourself the following question. Historically, growth has been around 2% for the past two hundred years. Now, imagine how your own ancestors were living 200 years ago - probably in conditions which you would consider of extreme need and poverty. Would you want such miserable people to have set something aside for you ? Probably not - quite the opposite, if you were able to do something for them, you would do it. Well, it growth continues at the same rate, this is the way that our descendants will look upon us; they will be richer than we can imagine. So why should we make sacrifices for such people ? Hence the high interest rate that we are in fact charging them.

- on the other tack, the interest rate *falls* with the growth rate. For instance, setting $\rho = 2\%$ and $\theta = 2$ again, we get $r = -2\%$ if $g = -2\%$, that is, if the economy contracts at the rate of 2% a year. So *negative interest rates* are not unthinkable - they might actually be needed in periods of negative growth. Think for instance of an economy where the only good is the environment, which cannot be produced, and actually has to decrease as the population growth - in such an economy, the interest rate would have to be negative. This leads us to the idea that one would actually have to use different rates for environmental goods and for consumption (manufactured) goods. The proper setting for exploring this idea is a two-goods model, and this is what we will be doing next.

A.2 Modifications 1- The ecological discount rate

We complement the standard model by adding a environment good E, which is not produced, and available in a fixed quantity $\bar{E}$. The utility function of the representative consumer is:

$$u(C,E) = \frac{1}{1-\theta} v(C,E)^{1-\theta}$$

with

$$v(C,E) = \left(\frac{1}{C^{\alpha}} + \frac{1}{E^{\alpha}}\right)^{-(1-\theta)/\alpha}$$

The parameter $\alpha > -1$ denotes the extent to which the environment good E and the consumption good C are substitutes. If a simultaneous and marginal changes $C \longrightarrow C - \Delta C$ and $E \longrightarrow E + \Delta E$ is to leave the total utility invariant, then we must have:

$$-\frac{\partial v}{\partial C}\Delta C + \frac{\partial v}{\partial E}\Delta E = 0$$

so that:

$$\frac{\Delta E}{\Delta C} = \frac{\partial v}{\partial C} / \frac{\partial v}{\partial E} = \left(\frac{E}{C}\right)^{1+\alpha}$$

The right-hand side can be rewritten it as follows:

$$\frac{\Delta E}{E} = \left(\frac{E}{C}\right)^{\alpha} \frac{\Delta C}{C}$$

In other words, to achieve an increase of 1% in the environmental good, the representative individual is willing to give up $\left(\frac{C}{E}\right)^{\alpha}$ % of the consumer good.

- if $-1 < \alpha < 0$, the willingness to pay for the environmental good decreases as E/C decreases, that is, as it becomes relatively scarcer. This is the case when the environmental good and the consumption good are *substitutes.*
- if $\alpha > 0$, the willingness to pay for the environmental good increases as it becomes relatively scarcer. This is the case when the two goods are *complements.*

As in the standard model, we assume that there is a production function $Y = F(K, AL)$, which is positively homogeneous of degree one, and where the labour force $L(t) = L_0 \exp nt$ and the technological progress $A(t) = A_0 \exp gt$ are exogeneously given. The relevant measure of consumption then is the consumption per household, which is given by:

$$c(t) = \frac{C(t)}{L_0 A(t)} = \frac{C(t)}{L_0 C_0 e^{gt}}$$

We will assume that the environmental good cannot be produced (or destroyed), and that it is available in a fixed quantity $\bar{E}$. We will also consider that it is a *public good*, so that it is non-exclusive: every household can enjoy the full quantity available available to everybody.

The utility of the representative household is then given by:

$$\begin{aligned}
u(C) &= \frac{1}{1-\theta}\left(\frac{1}{C^{\alpha}}+\frac{1}{\bar{E}^{\alpha}}\right)^{-(1-\theta)/\alpha} \\
&= \frac{1}{1-\theta}\left(\frac{1}{L_0^{\alpha}C_0^{\alpha}e^{\alpha g t}c^{\alpha}}+\frac{1}{\bar{E}^{\alpha}}\right)^{-(1-\theta)/\alpha} \\
&= \frac{1}{1-\theta}\bar{E}^{1-\theta}\left(1+\frac{\bar{E}^{\alpha}}{L_0^{\alpha}C_0^{\alpha}}\frac{1}{c^{\alpha}}e^{-\alpha g t}\right)^{-(1-\theta)/\alpha}
\end{aligned}$$

The representative household's optimisation problem then becomes:

$$\max \int_0^{\infty} \tilde{u}\left(1+\frac{\bar{E}^{\alpha}}{L_0^{\alpha}C_0^{\alpha}}\frac{1}{c^{\alpha}}e^{-\alpha g t}\right)e^{-\rho t}dt$$

$$\frac{dk}{dt} = f(k) - (n+g)\,k - c$$

where $\tilde{u}(x) := \bar{E}^{1-\theta}\frac{1}{1-\theta}x^{-(1-\theta)/\alpha}$. Note that it is a function of one variable only.

A.2.1 The case $\alpha > 0$

In that case, we find that, for large t, the utility function can be approximated by:

$$\begin{aligned}
\tilde{u}\left(1+\frac{\bar{E}^{\alpha}}{L_0^{\alpha}C_0^{\alpha}}\frac{1}{c^{\alpha}}e^{-\alpha g t}\right) &= \frac{1}{1-\theta}\bar{E}^{1-\theta}\left(1+\frac{\bar{E}^{\alpha}}{L_0^{\alpha}C_0^{\alpha}}\frac{1}{c^{\alpha}}e^{-\alpha g t}\right)^{-(1-\theta)/\alpha} \\
&\simeq \frac{1}{1-\theta}\bar{E}^{1-\theta}\left(1-\frac{1-\theta}{\alpha}\frac{\bar{E}^{\alpha}}{L_0^{\alpha}C_0^{\alpha}}\frac{1}{c^{\alpha}}e^{-\alpha g t}\right) \\
&= \frac{\bar{E}^{1-\theta}}{1-\theta} - \frac{1}{\alpha}\frac{\bar{E}^{\alpha+1-\theta}}{L_0^{\alpha}C_0^{\alpha}}c^{-\alpha}e^{-g\alpha t}
\end{aligned}$$

The constant plays no role in intertemporal optimisation, and we are left with the criterion:

$$\max \int_0^{\infty} -\frac{1}{\alpha}c(t)^{a}\, e^{-(\rho+\alpha g)t}dt$$

This is precisely the standard problem with $\alpha = \theta - 1$. The interest rate on consumption then is:

$$r_C = \rho + (1+\alpha)\,g$$

Note that the parameter θ does not appear in the formula - in the utility function of the representative household does not come into play ! This is known in the lterature as "ecological stunting": as the technological progress drives up the production of consumer goods, the environment becomes comparatively more valuable, and long–term interest rates are determined only by ρ, the pure rate of time preference, g, the technological grownt rate, and α - the larger α, the less and increase in consumption can compensate for a decrease in environment quality, and the higher the interest rate on consumption.

We can now find the *interest rate on environmental investments* by asking how much consumption the representative household would be willing to forgo today to increase the quantity of environmental goods by 1% at time t. The benefit of such an increase, evaluated in the consumer good, is equal to:

$$\left(\frac{C\left(t\right)}{\bar{E}}\right)^{\alpha} = \text{Const } e^{\alpha g t}$$

which leads to an ecological interest rate of:

$$r_E = r_C - \alpha g = \rho + g < r_C$$

A.2.2 The case $-1 < \alpha < 0$

As $t \longrightarrow \infty$, we find that:

$$\begin{aligned}
\tilde{u}\left(1 + \frac{\bar{E}^{\alpha}}{L_0^{\alpha} C_0^{\alpha}} \frac{1}{c^{\alpha}} e^{-\alpha g t}\right) &= \frac{1}{1-\theta} \bar{E}^{1-\theta} \left(1 + \frac{\bar{E}^{\alpha}}{L_0^{\alpha} C_0^{\alpha}} \frac{1}{c^{\alpha}} e^{-\alpha g t}\right)^{-(1-\theta)/\alpha} \\
&\simeq \frac{1}{1-\theta} \bar{E}^{1-\theta} \left(\frac{\bar{E}^{\alpha}}{L_0^{\alpha} C_0^{\alpha}}\right)^{-(1-\theta)/\alpha} c^{(1-\theta)} e^{(1-\theta) g t}
\end{aligned}$$

and we are back into the standard model, with the interest rate on consumption given by:

$$r_C = \rho + \theta g$$

References

[1] Guesnerie, Roger (2004) "Calcul economique et developpement durable", Revue Economique, 55 (3), p. 363-382

[2] Nordhaus, William D. (1991) "To slow or not to slow: the economics of the greenhouse effect" Economic Journal 101 p. 920-937

[3] Nordhaus, William D. (1992) "Lethal model 2: the limits to growth revisited" Brookings papers on economic activity (2) P. 1-43

[4] Romer, David (1996), "Advanced macroeconomics", McGraw and Hill

.1 Modifications 2: Uncertainty on the growth rates

.1.1 A classical argument

Let us start from Ramsey's formula:

$$r = \rho + \theta g$$

Assume now that *we believe in the model*, but are uncertain about the growth rate g :

$$g \sim \mathcal{N}\left(\bar{g}, \sigma^2\right)$$

so that $C(t) = C(0)\, e^{gt}$ is lognormal.

Assume moreover that we are utility maximizers, and handle uncertainties à la von Neumann-Morgenstern. We ask, as always, how much consumption $\Delta C(0)$ we are willing to forgo today to increase by $\Delta c(t)$ our consumption at time t

$$u'(C(0))\, \Delta C(0) = E\left[u'(C(T))\, \Delta C(T)\right] e^{-\rho t}$$

with $u(C) = \frac{c^{1-\theta}}{1-\theta}$ so that $u'(c) = c^{-\theta}$. This gives:

$$\frac{\Delta C(0)}{\Delta C(t)} = \frac{e^{-\rho t}}{u'(C(0))} E\left[u'(C(t))\right]$$

The computation gives:

$$\frac{e^{-\rho t}}{C(0)^{-\theta}} E\left[C(0)^{-\theta} e^{-g\theta t}\right] = e^{-rt}$$

$$r = \rho + \theta \bar{g} - \frac{1}{2}\theta^2 \sigma^2 \qquad \text{(Interest 2)}$$

Uncertainty *lowers* the interest rate. This corresponds to the standard fact that individuals are risk-averse. Note that this runs counter to an argument that politicians and businessmen have been making for many years, namely that we should do nothing about climate change, because it is not certain and it may turn out to be all right after all. From what we know, people are risk-averse for themselves, at least when the stakes (magnitude of potential losses) is large, meaning that the downside is more important to them than the upside, and it is difficult to understand why society should behave differently

.1.2 Pooling opinions of experts

Weitzman made the following, very general, observation. Suppose you consult two experts, whom you equally trust, about which interest rate to choose, and that they come up with two different opinions, namely r_1 and r_2, with $r_1 < r_2$ What value should you take ? As you trust them equally, it seems reasonable to pick the mean value, namely $\bar{r} = (r_1 + r_2)/2$. As Weitzman points out, this is wrong: what these experts are really saying is that one dollar today is worth respectively $e^{-r_1 t}$ and $e^{-r_2 t}$ at time $t > 0$. So if a mean is to be taken, it should be the mean of those values, leading to a interest rate $\tilde{r}$ given by:

$$\tilde{r}(t) = \frac{1}{t} \ln\left(\frac{1}{2} e^{-r_1 t} + \frac{1}{2} e^{-r_2 t}\right)$$

Note that this interest rate is *not constant*. It is approximately equal to $\bar{r} = (r_1 + r_2)/2$ for the short term, but for the long term is equal to the lowest rate r_1. This is the Weitzman lesson: for the long term, the lowest rate should prevail

Weitzman put his idea into practice. He pooled $I = 1,800$ economists and asked them for an assessment of interest rates to be applied for investment projects. Economist i answered with a constant rate r_i. leading to a discount rate $R_i(t) = e^{-r_i t}$. He found that the r_i were distributed according to a Gamma distribution with parameters (α, β):

$$f(r) \tilde{} \frac{\beta^\alpha}{\Gamma(\alpha)} r^{\alpha-1} e^{-\beta r}$$

Averaging the discount rates, he then derives the aggregate interest rate:

$$R(t) := \int_0^\infty f(r) A(r) \, dr = \left(\frac{\beta}{\beta+t}\right)^\alpha = \frac{1}{(1+t\sigma^2/\mu)^{\mu^2/\sigma^2}}$$

In terms of the mean μ and the variance σ^2 of the Gamma distribution (α, β). The corresponding interest rate then is:

$$r(t) = -\frac{1}{R}\frac{dR}{dt} = \frac{\alpha}{\alpha+\beta} = \frac{\mu}{1+t\sigma^2/\mu}$$

Note that very long-term interest rates are 0. The question is, how far out is the very long term ? Within the time horizon of the Stern review, from 50 to 200 years, Weitzman finds an interest rate of 1.75%, very much in line with the value 1.4% chosen by Stern himself.

References

[1] Weitzman, Martin (2001) "Gamma discounting", American Economic Review 91 (1) p. 260-271

.1 Modifications 3: Uncertainty on the model.

Up to now, the modelling does not capture one of the main features of very-long term decisions, namely the possibility of major catastrophes with unknown probabilities. The fact that these probabilities are unknown is an added ingredient to risk, which is not captured by simply assigning a priori probabilities, as in classical economic theory. Indeed, a classicial experiment by Ellsberg indicates that people have a specific aversion to ambiguity, that is to facing unknown probabilities. This is not captured by the von Neumann-Morgenstern approach to decision under uncertaintly, and the paradigm has to be changed. There is at present an active and promising literature on that.

Weitzman has pointed out another problem: whatever probability distribution our model works with, this will not be the one we work with. Indeed, we do not observe the distribution, all we can do is to infer it from a *finite* (and, in the case of climate change, pitifully small) amount of data. This means that, even if our model specifies Gaussian or Poisson distributions, which is usually the case, and which are nice because they have "thin tails" (large deviations have small probabilities) the ones we will end up working with may well have "fat tails", meaning that all long-term calculations break down.

To take a specific example, go back to the formula (Interest 2):

$$r = \rho + \theta \bar{g} - \frac{1}{2}\theta^2\sigma^2$$

which is based on the modelling assumption that $g \tilde{} \mathcal{N}\left(\bar{g}, \sigma^2\right)$. Weitzman'point is that, even if we agree with that specification, we know neither $\bar{g}$ nor σ. We will have to estimate them, and for this we need not only the data but an a priori distribution.

A standard way (Jeffreys prior) to choose such a distribution is to suppose that $\ln \sigma$ is Gaussian. If there are N experimental values available, we are led to a classical problem in statistics (find the variance of a Gaussian variable given N experimental values), the answer to which is a Student distribution with N degrees of freedom. It is well known that this distribution has fat tails. More precisely, if we have observed we find that:

$$\frac{1}{u'\left(c\left(0\right)\right)} E\left[u'\left(c\left(t\right)\right) \mid c\left(t_1\right), ..., c\left(t_N\right)\right] = +\infty$$

with the specification $u\left(c\right) = c^{1-\theta} / \left(1-\theta\right)$. In other words, given a finite number of observations, society should be willing to give up an unlimited amount of consumption today to gain any certain amount of consumption in the future. This corresponds to an interest rate of $r = -\infty$!

References

[1] Heal, Geoffrey (2007) "Climate change economics: a meta-review and some suggestions"

[2] Weitzman, Martin (2008) "On modeling and interpreting the economics of catastrophic climate change", American Economic Review 91 (1) p. 260-271

.1 Modifications 4: Equity and redistribution

.1.1 The problems

Consider again the standard model: there is an infinite-lived representative consumer, who strives to maximize

$$\max \int_0^\infty u\left(C\left(t\right)\right) e^{-\rho t} dt \qquad \text{(Ramsey Growth Model)}$$

Problem 1: there is no such thing as a representative consumer People are different - in their tastes (utility function u), in their expectations (probability p). More importantly, some are rich, but most are poor. The first question is dealt with by aggregation theory (see the lectures by Jouini in this summer school). The second question, to my knowledge, has not attracted academic attention - except from Ramsey himself ! He devotes the last section of his seminal paper (1928) to this problem and concludes : "In such a case, therefore, equilibrium would be attained by a division of society into two classes, the thrifty enjoying bliss and the improvident at the subsistence level". It would of course be politically quite incorrect to mention the poor nowadays, and this is why academics gladly adhere to the fiction of the representative consumer.

Problem 2: no one lives for ever This means that the coefficient ρ (pure rate of time preference) will apply to different persons in the short to middle term (where the present generation is alive) and in the middle to long term (when we are all dead, and our descendants rule or are ruled). This means that this parameter is put to two different uses:

- for weighing consequences to me of my own actions
- for weighing consequences to others of my own actions

In a seminal paper, Sumaila and Walters (2005) separate the (psychological) impatience from the (ethical) concern for future generations. Their model combines three parameters:

- the population grows at the rate γ
- each generation has a pure rate of time preference ρ
- each generation discounts at the rate $\delta < \rho$ the utility of future generations

For an event which is to happen at time t, we find that the discount factor to apply is:

$$\begin{aligned} R(t) &= e^{-\rho t} + \gamma \int_0^t e^{-\delta s} e^{-\rho(t-s)} ds \\ &= (1-\lambda) e^{-\rho t} + \lambda e^{-\delta t} \end{aligned}$$

with:

$$\lambda = \frac{\gamma}{\rho - \delta}$$

. Note that this corresponds to a non-constant rate of time preference:

$$r(t) = -\ln\left((1-\lambda) e^{-\rho t} + \lambda e^{-\delta t}\right)$$

- $r(t) \simeq -\delta$ in the long term
- $r(t) \simeq \lambda\rho + (1-\lambda)\delta$ in the short term

Problem 3: time-inconsistency An individual alive today then faces the problem (in the deterministic case; we set $n = g = 0$ for simplicity):

$$
\begin{aligned}
& \max \int_0^\infty R(t)\, u(c(t))\, dt, \\
\frac{dk}{dt} &= f(k(t)) - c(t) \text{ and } k(0) = k_0
\end{aligned}
$$

This means that the present generation is dependent upon others (namely future generations) to carry out the policies that it designs today. However, the future generations may not agree with decisions taken on their behalf many years before they were around, and decide to carry out different ones. This is the *non-commitment* problem, which can be avoided only if whatever seemed optimal when certain policies were decided still seems optimal when the time has come to carry them out. Unfortunately, with a non-constant rate of time preference, this will not happen.

Take two scenarios $c_1(\cdot)$ and $c_2(\cdot)$, both of which kick in at time T. In other words, $c_1(\cdot)$ and $c_2(\cdot)$ are defined for $s \geq T$. Say that we compare them, at some time $t_1 < T$, and we find $c_1(\cdot)$ is superior to $c_2(\cdot)$:

$$\int_T^\infty R(t-t_1) u(c_1(t))\, dt \geq \int_T^\infty R(t-t_1) u(c_2(t))\, dt \tag{1}$$

Let some time elapse, and do the comparison again at some later instant $t_2 < T$. Is it still true that we will find $c_1(\cdot)$ superior to $c_2(\cdot)$? This would mean that:

$$\int_T^\infty R(t-t_2) u(c_1(t))\, dt \geq \int_T^\infty R(t-t_2) u(c_2(t))\, dt \tag{2}$$

In the case when the discount rate is constant, so that $R(t) = \exp^{rt}$, the first inequality implies the second because of the special properties of the exponential function. We have:

$$
\begin{aligned}
\int_T^\infty R(t-t_2) u(c_1(t))\, dt &= \int_T^\infty e^{r(t-t_2)} u(c_1(t))\, dt \\
&= e^{r(t_1-t_2)} \int_T^\infty e^{r(t-t_1)} u(c_1(t))\, dt \\
&= e^{r(t_1-t_2)} \int_T^\infty R(t-t_1) u(c_1(t))\, dt
\end{aligned}
$$

so that 2 is derived from 1 by multiplying both sides by a constant.

In the case of non-constant discount rates, 1 no longer implies 2 ! In fact, a policy which is optimal for the decision-maker at time t_1, no longer is optimal for the decision-maker at a later time t_2 (even though the utility function $u(c)$ is unchanged). There is no control that will be simultaneously optimal for all those who will have to implement it. We need a new concept, and we will introduce it in the next lecture. Please refer to the bibliography and to the presentation "Lecture 5 - Time inconsistency" on the PIMS website:

http://www.pims.math.ca/scientific/summer-school/summer-school-perceiving-measuring-

and-managing-risk-illiquidity-long-term-

References

[1] Ekeland, Ivar and Lazrak, Ali (2005) "Being serious about non-commitment" http://arxiv.org/abs/math/0604264

[2] Sumaila, Ussif and Walters, Carl (2005) "Integenerational discounting: a new intuitive approach" Ecological Economics 52 p. 135-142