Microeconomic Theory
Lecture 10: Intertemporal Decisions

Marcel Fafchamps
Oxford University

Readings: Varian, Chapters 11, 19, and 20.

1. Intertemporal preferences

Let us introduce time in the standard household model. Let c_t denote a vector of consumption at time t and let t go from 1 to T. A complete description of consumption from $t = 1$ to $t = T$ is the composite vector:

$$\{c_1, ..., c_T\}$$

This is a very large vector since it describes consumption at all (relevant) points in time. Suppose that consumers have preferences that satisfy the four axioms of utility theory. Consequently, they can unambiguously rank all possible consumption vectors $\{c_1, ..., c_T\}$. Consequently, there exists a utility function:

$$U(c_1, ..., c_T) \tag{1.1}$$

that describes the consumer's preferences.

1.1. Additive separable preferences

Although there is nothing wrong with utility function (1.1), it is extremely cumbersome to work with. Except for some work in economic theory, economists nearly never work with such a general function when they discuss intertemporal phenomena. They typically assume that utility takes an additively separable form written:

$$U(c_1, ..., c_T) = \sum_{t=1}^{T} \beta^t u(c_t) \tag{1.2}$$

where function $u(c_t)$ is called the instantaneous utility function, and β is called the discount factor.

Parameter β measures how individuals discount future utility. It is assumed that $\beta \leq 1$. To see how β works, suppose there are two periods ($T = 2$) and assume that $c_t = c_{t+1} = c$. In this case, we have:

$$U(c_1, c_2) = u(c_1) + \beta u(c_2)$$

This shows that, whenever $\beta < 1$, the instantaneous satisfaction derived from immediate consumption is valued more than future satisfaction. By extension, if there are T periods and constant consumption, we see that consumption of c at m periods in the future yields β^m as much satisfaction as consumption of c now. Parameter β thus expresses the consumer's degree of impatience or time preference. Consumers with a high β are more patient; those with a low β are more impatient. A related concept is that of discount rate ρ which is defined as:

$$\begin{aligned} \rho &= \frac{1-\beta}{\beta} \text{ and thus} \\ \beta &= \frac{1}{1+\rho} \end{aligned}$$

1.2. Hyperbolic preferences and time inconsistency

Experiments have been used to test whether the assumptions underlying the above model hold in practice. In general, human agents are found to diverge from this model; they are said to have 'hyperbolic' (or quasi-hyperbolic) preferences. As it turns out, hyperbolic preferences violate not only equation (1.2) but also the axioms underlying utility theory. This is because people who have hyperbolic preferences make inconsistent choices that cannot, therefore, be expressed as a single ranking $U(c_1, ..., c_T)$. Hyperbolic preferences are thus a rejection not so much of the specific form of equation (1.2) but of the existence of a function $U(c_1, ..., c_T)$ itself.

To illustrate what the problem, consider slightly generalized preferences of the form:

$$\sum_{t=1}^{T} \beta(t) u(c_t)$$

where $\beta(t)$ is now a function of time. Assume that $T \geq 3$. Note that the discount factor between period t and $t+1$ is simply $\beta(t+1)/\beta(t)$. In the exponential case (1.2), this boils down to $\beta^{t+1}/\beta^t = \beta$: the discount factor between periods does not vary with time. Contrast this case to a situation in which $\beta(2)/\beta(1) < \beta(3)/\beta(2)$. This means that $\beta(t)$ falls more between periods 1 and 2 than between periods 2 and 3. This kind of β function is commonly referred to as hyperbolic preferences because function $\beta(t)$ could not be written in the β^t form, for any value of β. But it could be captured by another class of functions called hyperbolic. For us, this is just a name and we need not worry more about it.

Inconsistency arises if, at the beginning of every period, people change their preferences so that time is 'reset' at $t = 1$ and preferences restart with hyperbolic time preferences $\beta(t)$. In this case, people's ordering of consumption bundles change over time. To see why, again assume that $\beta(2)/\beta(1) < \beta(3)/\beta(2)$. This implies that people want to consume a lot today but are happy to save tomorrow for period 3 – because the discount rate between periods 2 and 3 is lower than then discount rate between periods 1 and 2.

Suppose that an individual with these preferences starts at period 1. He decides to consumes a lot at 1 but then to save between periods 2 and 3. When this person reaches period 2, however, his time preferences 'restart' so that again he is very impatient relative to the next period. Consequently the decision maker decides to change his initial savings decision and to consume more in period 2 than he planned to do when he decided his original consumption path in period 1. This kind of situation is called time inconsistency. Note that time inconsistency arises because preferences over consumption bundles $U(c_1, ..., c_T)$ are not stable over time – i.e., violate utility axioms of a clear and unique ranking. With time inconsistent preferences, rankings change over time: when in period 1, the decision maker wanted to save in period 2; but when period 2 is reached, the decision maker wants to consume.

Time inconsistency does not arise if the decision maker has time preferences of the form β^t. In this case, even if time is 'reset', relative discounting remains unchanged, i.e., we have:

$$\frac{\beta(2)}{\beta(1)} = \frac{\beta^2}{\beta} = \frac{\beta(3)}{\beta(2)} = \frac{\beta^3}{\beta^2} = \beta$$

This means that switching the 'point of view' from time 2 to time 1 does not change the relative discounting of time periods – and hence does not change choices. Consequently, if the decision maker decided to save at time 2, when he reaches time 2, he still decides to save the same amount even if time is 'reset' to 1.

Time inconsistencies of this kind are receiving more attention these days. They can explain many interesting phenomena, especially forced savings of various kinds: automatic savings contract with a bank; voluntary contributions to a pension scheme; life insurance; rotating savings and credit associations (ROSCAs); 'susu' collectors; etc. It also explains why we seek to limit our capacity to succumb to impulse purchases – e.g., we carry a limited amount of cash when we go to the pub... Certain forms of micro-lending could even be construed as self-control devices to help the poor save.

Another intertemporal model that generates time inconsistent decisions is the following model of impulse purchase or addiction. At EACH time period people are assumed to maximize a utility function of the form:

$$\max_{\{c_t, m_t\}} u(c_1 - m_1) + S_0 + \sum_{t=1}^{T} \beta^t u(c_t - m_t)$$

where $m_t = \{0, \overline{m}\}$ and $S_t = \overline{S}$ if $m_t = \overline{m}$ and $S_t = 0$ otherwise. The idea is that expenditure m_t generates an immediate satisfaction S_t. But the individual does not perceive any benefit from future expenditure m_t. Consequently, (for $\overline{S}$ large enough) the individual decides to set $m_1 = \overline{m}$ and incur satisfaction $\overline{S}$ now, but decides to set $m_t = 0$ for all subsequent periods. When period 2 is reached, however, the individual again solves the problem above, i.e., again perceives satisfaction level S. To understand what is going on, imagine a smoker who thinks he should stop smoking in the future but cannot help smoking 'a last cigarette' $\overline{S}$, promising himself that, come tomorrow, he will not smoke ($m_t = 0$). The problem is that when tomorrow comes, he is again tempted to consume a cigarette.

Self-control issues are now taken seriously in policy design because it is increasingly recognized that the poor find it difficult to save on their own and often call for some kind of commitment device. One could even argue that group lending is such device because it is engineered to assist people when they have a self-control problem – i.e., are tempted not to make their regular contribution to repayment of the loan. The group serves to motivate the person, a bit like an alcoholic anonymous support group. See recent experimental work on fertilizer demand in Kenya (Duflo) and on saving in the Philippines (Karlan).

2. An intertemporal household model

In Lecture 3, I introduced general notation for a household with missing markets. I now extend that model to include time. Below is an intertemporal version of that model, assuming that there are no missing markets for consumption goods c_{it} and variable inputs and outputs q_{it} (i.e., all $i \in T$):

$$\max_{\{c_t,q_t,K_t,I_t\}} \sum_{t=0}^{\infty} \beta^t U(c_{1t},...,c_{Nt}) \text{ subject to}$$

$$\begin{aligned} \sum_{i\in T} p_{it}c_{it} + p_{kt}I_t &= \sum_{i\in T} p_{it}(q_{it}+T_{it}) + S_t \\ G(q_t,K_t) &= 0 \\ K_{t+1} &= (1-\delta)K_t + I_t \qquad (2.1) \\ K_{t=0} &= \overline{K} \end{aligned}$$

where T_{it} as before denotes the endowment of good i (e.g., labor) at time t, variable K_t denotes the stock of a capital (accumulable) good at time t, I_t stands for investment in capital, S_t are exogenous transfers, β is the household's discount factor ($\beta \leq 1$), and δ stands for the depreciation rate on capital. Endowments other than capital are regarded as exogenously determined but may vary over time (e.g., through birth, aging, etc). All exogenous variables and prices are assumed to be known beforehand.

The main difference with our static model is the addition of a law of motion for capital, equation (2.1) and the starting condition $K_{t=0} = \overline{K}$. We also explicitly recognize that production depends on capital stock K_t.

2.1. Simplifying the model

At first glance, the above model appears difficult to deal with. However, it is possible to transform it into a simpler capital accumulation model (i.e., a growth model). First, we use equation (2.1) to get rid of variable I_t in the budget constraint. This eliminates one variable and one constraint. Second, we normalize prices so that $p_{kt} = 1$ for all t. This means that all prices are quoted relative to the price of capital.

Finally we note that, within each period, we can solve for decisions about which goods to consume and produce *conditional on* the value of K_t. In other words, the optimization model can be decomposed into two nested optimizations:

$$\max_{\{K_t\}} \quad \max_{\{c_t,q_t|K_t\}} \sum_{t=0}^{\infty} \beta^t U(c_{1t},...,c_{Nt}) \text{ subject to}$$

$$\begin{aligned} \sum_{i\in T} p_{it}c_{it} + K_{t+1} - (1-\delta_i)K_t &= \sum_{i\in T} p_{it}(q_{it}+T_{it}) + S_t \\ G(q_t,K_t) &= 0 \end{aligned}$$

If we ignore the investment term $K_{t+1} - (1-\delta_i)K_t$, the inner optimization is exactly the same model as what we have examined in the household model with missing markets (except that here there are no missing markets). It can be solved in the same way. This means that the solution to the inner optimization can be written in the form of an indirect utility $V(C_t, p_t)$ and a profit function $\pi(p_t, K_t)$ where C_t stands for total consumption expenditures, that is, total income minus investment:

$$C_t = \sum_{i\in T,N} p_{it}c_{it} = Y_t - I_t$$

where total income is:

$$Y_t = \pi(p_t, K_t) + S_t + \sum_{i\in T} p_{it}T_{it} = Y(p_t, K_t, X_t)$$

The first term is pure profits, the second term is transfers from the rest of the world, and the last term is payment to factors, that is, value added. To simplify the notation, I have written unearned income as $X_t \equiv S_t + \sum_{i \in T} p_{it} T_{it}$. We have:

$$C_t = Y(p_t, K_t, X_t) - K_{kt+1} + (1 - \delta_i) K_{kt}$$

With this, the outer optimization model can be rewritten more simply as:

$$\max_{\{K_t\}} \sum_{t=0}^{\infty} \beta^t V(C_t, p_t) \text{ subject to}$$
$$C_t = Y(p_t, K_t, X_t) + (1 - \delta_i) K_{kt} - K_{kt+1} \text{ for all } t$$

This resembles the standard growth model, except for the explicit recognition that the solution depends on prices.[1] This dependence on prices is usually ignored in the notation, which is equivalent to assuming either that prices do not change over time or that output is measured in constant terms.[2] The advantage of the above presentation is to make underlying assumptions explicit.

The solution to the above is discussed in the growth literature. What is typically not discussed in the growth literature is price fluctuations. There are situations in which time-vaying prices should be taken into account. Seasonal data is a case in point: relative prices are known to fluctuate systematically over the year as various crops get harvested. A seasonal analysis of intertemporal choices should therefore take price fluctuations into account. Over time, changes in household composition may also affect consumption preferences for various products – e.g., milk, nappies – and household services – e.g., childcare and elderly care. This in turn would affect $V(C_t, p_t)$ in a systematic manner. For instance, anticipated pregnancy would affect investment and labor choices of parents, especially if certain goods (milk) and services (childcare) cannot be provided by the market, possibly because of asymmetric information (milk quality, child care supervision). This, for instance, may explain why households invest less in the schooling and market-specific skills of women or why they purchase and hold cows when they have small children.

2.2. Credit markets and separability

With a little more work, we get a result similar to the separability property of household models but this time regarding savings. Let us first see what happens without credit market. Following standard practice, we rewrite the optimization model omitting the dependence on prices and on unearned exogenous income X_t. We have:

$$\max_{\{K_t\}} \sum_{t=0}^{\infty} \beta^t V(C_t) \text{ subject to}$$
$$C_t = Y(K_t) + (1 - \delta_i) K_{kt} - K_{kt+1} \text{ for all } t$$

Replacing C_t in the objective function by its value and differentiating with respect to K_t, we obtain a series of first order conditions of the form:

$$\beta^t V_t' \left(\frac{\partial Y(K_t)}{\partial K_t} + 1 - \delta \right) - \beta^{t-1} V_{t-1}' = 0$$

[1] Also I have written the production function a bit differently in the sense that I expect to recover depreciated capital at the end of the period. Many growth models 'fold' this into the production function, i.e., define $f(K) = Y(K) + K$.

[2] Atanasio and Weber (JPE 1995) examine the effect of relative price changes in an intertemporal model, assuming all prices are exogenous.

These are called Euler conditions. Rearranging we get:

$$1 + \frac{\partial Y(K_t)}{\partial K_t} - \delta = \frac{V'_{t-1}}{\beta V'_t}$$

The left-hand side is the marginal return to capital, net of depreciation; the right-hand side is the utility cost of saving. What this shows is that the choice of capital stock K_t depends on the time preference coefficient β and on the ratio of marginal utility of consumption. Put differently, the choice of K_t depends on preferences; the model is not separable. Note that, even if consumption is constant over time (as in the steady state) and $V'_{t-1} = V'_t$, we still get

$$\begin{aligned} 1 + \frac{\partial Y(K_t)}{\partial K_t} - \delta &= 1 + \rho \text{ or} \\ \frac{\partial Y(K_t)}{\partial K_t} &= \rho + \delta \end{aligned}$$

which shows that the rate of time preference ρ still enters the equation that determines the choice of capital stock K_t. In general, however, $V'_{t-1} \neq V'_t$ when accumulation is taking place and income rises over time. If $C_{t-1} < C_t$, then $V'_{t-1} > V'_t$. This is because, when consumption is low, the marginal utility of an additional \$ is worth more than when consumption is high. It follows that $\frac{V'_{t-1}}{V'_t} > 1$. Consequently, investment K_t is less than if $V'_{t-1} = V'_t$.

Now suppose that capital can be rented at cost r or, alternatively, suppose that capital can be purchased with borrowed funds, and that the interest on borrowed funds is r. In this case, the optimization problem becomes:

$$\begin{aligned} &\max_{\{K_t, W_t\}} \sum_{t=0}^{\infty} \beta^t V(C_t) \text{ subject to} \\ C_t &= Y(K_t) - \delta_i K_t - r K_t - W_{t+1} + (1+r) W_t \end{aligned}$$

where W_t is the net wealth of the household, and r is the interest rate. I have assumed that the household can save and earn an interest r. Since only K_t enters the budget constraint, not K_{t+1}, the Euler equations for K in each period is:

$$\frac{\partial Y(K_t)}{\partial K_t} = r + \delta$$

which means that the marginal return to capital equals the rental cost of capital plus depreciation. In this case, production decisions are independent from endowments (e.g., wealth W_t) and preferences (e.g., discount factor β, utility function $V(.)$). The producer adjusts instantaneously to the optimal capital level. There is no waiting time.

There is a lot of work in development economics focusing on credit constraints and the need for the poor – and poor countries in general – to accumulate before being able to make a large, lumpy profitable investments. Nurkse (Problems of Capital Formation in Underdeveloped Countries, 1953) made the point that capital formation is the key constraint of poor countries. Gerschenkron (Economic Backwardness in Historical Perspective, 1962) went further by insisting that providing funds to the country is not enough; it must reach investors. He started a large literature on the role of financial intermediation in the development process. McKinnon (1973) further argued that to speed up investment by the poor, savings instruments available to them must yield a high enough return.

Much of the literature on poverty has concerned itself with this issue – e.g., current efforts to develop micro-credit. Two observations are in order. First, the absence of credit does not, by itself, preclude social mobility; it only slows it down. In the model presented above, we see that poor people save and eventually get out of poverty. [Things would be different if there is a large return on a large and non-divisible investment but the return on a small investment/saving is small or negative – for instance because of inflation. In this case, accumulating enough to undertake the large investment may take so long that the individual is discouraged and gives up trying. In this case, we have a poverty trap, that is, there is no accumulation out of poverty.]

Second, this model does not include any risk. If we include risk, then giving credit to the poor may, if they are unlucky with their investment, push them into a debt trap which, in the absence of personal bankruptcy law, can last indefinitely. To put differently, contrary to the way many people think, there is no poverty trap in the absence of credit – just a delay. But there is a debt-driven poverty trap once credit to the poor is introduced and some risk exists. Put differently, without credit a poverty trap does not exist but with credit one appears... This is an issue I discuss in my paper with Gubert on contingent loan repayment and in my book on poverty and risk. We will discuss this issue further when I formally introduce risk.

3. Human capital

Investment in human capital can be modeled in the same way as any other accumulable asset – hence the name 'capital'. So it is possible to take the model presented so far and include human capital in the capital vector. There is an enormous literature on this.

To illustrate the kind of findings we get, let me develop a little model and illustrate how it can be used to derive interesting results. I focus here on the schooling dimension of human capital. There is a price to pay to go to school – school fees, foregone income – and a return – increased wages or earnings from self-employment. Let the cost be written $F(E_{t+1})$ and the wage/income function be written $w(E_t)$ with $F' > 0$ and $w' > 0$. To keep things simple, let us ignore other assets. We first assume there is no credit market. Furthermore, for simplicity, we consider an individual who lives two periods – i.e., works or goes to school in period 1 and only works in period 2. We also assume that $E_1 = 0$.

We have:

$$\begin{aligned} & \max_{\{E_2\}} V(C_1) + \beta V(C_2) \text{ subject to} \\ C_1 &= w(0) - F(E_2) \\ C_2 &= w(E_2) \end{aligned}$$

which can be simplified into:

$$\max_{\{E_2\}} V(w(0) - F(E)) + \beta V(w(E))$$

with first order condition:

$$\begin{aligned} -V_1' F' + \beta V_2' w' &= 0 \text{ or} \\ \beta w' &= F' \frac{V_1'}{V_2'} \end{aligned}$$

which means that the marginal return to an additional year of schooling is equated with the marginal cost of acquiring education times $\frac{V_1'}{\beta V_2'}$. The latter term is a correction for the difference in levels of consumption between the two periods. In general we assume that V' falls with consumption – the marginal satisfaction one gets from an additional unit of consumption gets smaller as income rises. If this is true, $V_1' > V_2'$ if $C_1 < C_2$. Since $w(0) < w(E)$ in general we have $V_1' > V_2'$ which implies that $w' > F'$ in most cases. Put differently, the poor underinvest in education because of the high utility cost of paying school fees when consumption is low.

3.1. Separability

As before, underinvestment disappears if we have a perfect credit market. To see this, add debt D to the model. We now have:

$$\max_{\{E_2,D\}} V(C_1) + \beta V(C_2) \text{ subject to}$$
$$\begin{aligned} C_1 &= w(0) - F(E) + D \\ C_2 &= w(E) - (1+r)D \end{aligned}$$

Solving for the two first order conditions, we obtain:

$$\begin{aligned} -V_1'F' + \beta V_2'w' &= 0 \\ V_1' - (1+r)\beta V_2' &= 0 \end{aligned}$$

from which we get:

$$\begin{aligned} \frac{V_1'}{\beta V_2'} &= 1+r \\ \frac{w'}{1+r} &= F' \end{aligned}$$

which shows that the (discounted) marginal return to education is set equal to the marginal cost.

These simple observations form the basis of many efforts to help the poor get education – e.g., by giving student loans, reducing the cost of going to school (free tuition, free meals, etc), and making it illegal for parents to send their kids to work (eliminating the opportunity cost of foregone child labor).

The model suggests that the best method is student loans, not low tuition. To see this, suppose that households differ in initial endowments, so that some people can consume some extra amount W_0. In this case, it is impossible to find a tuition grant that ensures that $\frac{w'}{1+r} = F'$ for all households. In fact, a tuition grant equal for everyone would raise education demand by the 'privileged' (lower V_1') so that, in the end, they would get more education than in the first best equilibrium – and than the poor. Needless to say, this reasoning assumes that all those who wish to study get a loan. In practice, this need not be the case if individuals who fail at their studies cannot be forced to repay the loan they received. The creates a perverse incentive issue and lenders may be reluctant to lend as a result.

Of course, starting from a situation in which education is free, the institution of a tuition fee would represent an immediate financial loss that is felt more by poor parents, given that they are less able to absorb it from their accumulated savings. Showing this formally is left as an exercise.

3.2. Human capital as bequest

The above human capital model assumes that the investment is done by the person herself. In most cases, however, education is financed by parents as a form of bequest inter vivos.[3] Schooling then is the result of an altruistic decision by parents. One way of modeling this is to write a modified human capital investment model as:

$$\max_{\{E_2\}} V_1(C_1) + \alpha V_2(C_2) \text{ subject to}$$
$$\begin{aligned} C_1 &= w(\overline{E}_1) - F(E_2) \\ C_2 &= w(E_2) \end{aligned}$$

where for simplicity I have assumed that parents live only in period one while their child studies in period 1 and works in period 2. Parameter α captures parental altruism, that is, how parents value the utility of their child $V_2(.)$. It is immediate to see that the solution to the above optimization problem is very similar:

$$w' = F' \frac{V_1'}{\alpha V_2'}$$

If parents value their child's utility at most as much as their own, and they are poorly educated with a low income, then we have the same result as before: $w' > F'$, meaning that there is underinvestment in the child's education. This means that poverty in partly inherited: poor parents invest less in the education of their children and the children remain poor. This an example of intergenerational presistence in poverty. If, however, $\alpha >> 1$, it is conceivable that parents will invest at the optimal level – or even over-invest – although they are poor.

More complex models can be investigated in which parents can choose whether to bequeath their wealth in terms of human capital or in the form of real assets such as land. Quisumbing (1994) for instance shows that rural Filipino parents educate their daughters but bequeath land to their sons.

[3] This need not be true for university education, but it is usually true for primary and secondary education.