Notes for Econ202A:
The Ramsey-Cass-Koopmans Model

Pierre-Olivier Gourinchas
UC Berkeley

Fall 2015

1 Introductory Remarks

- In the last few weeks, you studied the Solow model in great details. The Solow model is a very important tool to understand the determinants of long term growth. It's main conclusion – that long term growth in standards of living across countries and over time cannot be accounted for simply by the accumulation of physical capital– is a key message.
- But the Solow model makes some important simplifying assumptions along the way. One of the most striking simplification is that aggregate consumption is simply a linear function of aggregate output, so that the fraction of output devoted to investment (=saving in a closed economy) is also constant.
- This is really a strong assumption. An entire body of literature on consumption behavior (which we will study in more details in the second part of this course) emphasizes that household consumption is much more complex than simply a fixed proportion of income. In fact, Milton Friedman in 1976 and Franco Modigliani in 1985 both won the Nobel Prize in Economics in part for their analysis of aggregate consumption and saving behavior.
- So today, we add one layer to the model we've been working with: we allow households to make **optimal consumption/saving decisions** at the microeconomic level, given the environment they are facing. As a result, the evolution of the capital stock will reflect the interactions between utility-maximizing households (supplying savings) and profit-maximizing firms (demanding investment). In this model, the saving rate may not be constant anymore.
- This model was originally developed by Frank P. Ramsey, a precocious mathematician and economist who died at age 26! (1903-1930).[1] The original Ramsey problem was a planning problem (i.e. the allocation of resources chosen optimally by a planner that tries to maximize the utility of households). The model was later extended by David Cass and Tjalling Koopmans in 1965 (in separate contributions) to a decentralized environment where households supply labor, hold capital and consume optimally, given prices and wages, while firms rent capital, hire labor to maximize profits, given prices and wages; and markets clear.[2] The two approaches are identical, because there are no market imperfections, so the first welfare theorem holds: the competitive, decentralized equilibrium is a solution to the planner problem. Historically the model is often referred to as the **Ramsey-Cass-Koopmans** model.
- Today we will look at the competitive equilibrium. You will see the corresponding planner's problem later with David Romer.
- Extending the model in this direction achieves three purposes:
    1. (least important). it will provide a check on the Solow model. We want to know if the insights from that model will survive once we allow for more complex and endogenous saving behavior. If they did not, then we would have to conclude that the assumption of a constant saving rate is quite important. As it turns out, we will see that in some important way, Solow's insights survive.

[1] Ramsey wrote only three papers in economic theory, all of which are foundational papers, one on subjective utility, on on optimal taxation and one on applying calculus of variation to the question of consumption and saving. Had he lived longer, he would have easily won the Nobel prize in Economics himself.

[2] Koopmans also received a Nobel prize in 1975 for his contributions to the "theory of optimal allocation of resources."

2. it allows us to address welfare issues. This is a major benefit of having a fully microfounded approach: the utility of the household is well specified and can be evaluated along a number of alternative scenarios. In the Solow model, we can only look at aggregate variables (output, consumption etc...) but we cannot specify what is desirable from the point of view of aggregate welfare.
3. from a methodological point of view, this allows us to introduce important new tools. Namely, in today's lecture we will introduce the techniques for *dynamic optimization in continuous time*. These are valuable and very powerful tools. Infinite horizon (and later Overlapping Generation) models come up in all kinds of models all over economics, so they should be part of your standard toolkit.

2 The Ramsey-Cass-Koopmans Model

2.1 Firms

- There is a large number of identical firms, with access to a production function $Y(t) = F(K(t), A(t)L(t))$ with the same properties as in the Solow model (i.e. constant returns to scale, $\partial F/\partial K > 0$ and $\partial^2 F/\partial K^2 \leq 0$, $\lim_{K\to\infty} \partial F/\partial K = 0$ and $\lim_{K\to 0} \partial F/\partial K = \infty$, and $F(0, AL) = 0, \forall AL$).
- Technology is the same as in the Solow model and grows exogenously at rate $\dot{A}(t)/A(t) = g$
- To simplify things a little, we assume there is no depreciation: $\delta = 0$.
- Firms hire workers at real wage $W(t)$ at time t and rent capital at rate $r(t)$ to maximize profits
- Firms rebate profits (if any) to households (i.e. the owners of the firm)
- $K(0)$, $A(0)$ and $L(0)$ all given and all > 0.[3]

2.2 Households

- There is a large number of identical, infinitely lived, households. The size of each household grows at rate n. Denote H the number of households. Population $L(t)$ also grows at rate n and the size of each household is $L(t)/H$.
- Because each household is small, they take wages and interest rates as given.
- Each member of the household supplies 1 unit of labor inelastically (so # workers = #people $= L(t)$)
- Each household initially holds $K(0)/H$ units of capital.
- Each household receives income from the following sources:
  - labor income (wages of the household members),
  - capital income (from renting out capital to the firms)

[3] Note that here, unlike in the Solow model, we have to be concerned with what happens if the households decide to consume everything, leaving no capital behind. From that point on, the assumptions on the production function would imply that there is no output ever again since $C < 0$ cannot happen. As we will see these paths are obviously not optimal.

- profits from the firms (rebated to the households), if any.

- Each household has to decide how much to consume and how much to save (in the form of capital accumulation).

- How do they choose between various consumption sequences? By maximizing lifetime utility:

$$U = \int_{t=0}^{\infty} e^{-\rho t} u(C(t)) \frac{L(t)}{H} dt \tag{1}$$

  Note the various elements in this expression:

  - integral of some flow utility defined over consumption per worker $u(C(t))$. Here, $C(t)$ denotes *consumption per worker*, so aggregate consumption is $C(t)L(t)$.
  - multiplied by the number of people in the household $L(t)/H$
  - discounted at rate ρ.

- We specialize the flow utility $u(C(t))$ to:

$$u(C(t)) = \frac{C(t)^{1-\theta}}{1-\theta} \quad ; \quad \theta > 0 \quad ; \quad \rho - n - (1-\theta)g > 0 \tag{2}$$

  - θ is the *coefficient of relative risk aversion*, defined as $-Cu''(C)/u'(C)$. The coefficient or relative risk aversion tells us how the household is going to rank different lotteries.
  - θ plays another role: it is also the inverse of the instantaneous elasticity of intertemporal substitution (IES). The intertemporal elasticity of substitution between two points in time t and $s > t$ is defined as $\sigma_{t,s} = -d\ln(C_s/C_t)/d\ln(u(C_s)/u(C_t))$. If you take the limit of this expression as $s \to t$ you get $\sigma = -u'(C)/Cu''(C) = 1/\theta$. The IES tells us how willing the household is to shift consumption from one period to another.
  - Here risk considerations are irrelevant since the environment is deterministic, so what matters is the IES. A low θ means a high IES: marginal utility fall more slowly and the household is more willing to substitute consumption over time. When θ is high, the IES is low: marginal utility falls rapidly and the household is less willing to substitute.
  - Because the coefficient of relative risk aversion is constant, these preferences are called Constant Relative Risk Aversion, or CRRA.
  - The assumption $\rho - n - (1-\theta)g > 0$ ensures that lifetime utility is well defined. To see this, note that along a balanced-growth-path where consumption per capita grows at rate g the integrand term in U grows at rate $-\rho + n + (1-\theta)g$. We want that term to be negative so that the integral converges. If this is not satisfied, the maximization problem is not well-defined.

2.3 The Behavior of Firms

- Firms take input prices as given and choose how much labor to hire and how much capital to rent to maximize profits $\pi(t)$:

$$\pi(t) = Y(t) - W(t)L(t) - r(t)K(t)$$

  where $W(t)$ is the real wage and $r(t)$ is the rental rate of capital.

- **Capital**: Since firms pay a price $r(t)$ for renting a unit of capital and there is no depreciation, they will equate the return to capital and the marginal product of capital: $r(t) = \partial F(K, AL)/\partial K$. Since $\partial F(K, AL)/\partial K = f'(k)$ where $f(k) \equiv F(k, 1)$ and $k = K/AL$, firms will rent capital up to the point where:

$$f'(k(t)) = r(t) \tag{3}$$

- **Labor**. Similar reasoning tells you that firms will hire workers up to the point where the real wage $W(t)$ equals the marginal product of labor $\partial F(K, AL)/\partial L$. To express the marginal product of labor in terms of $f(.)$, observe that

$$\begin{aligned}\frac{\partial F(K, AL)}{\partial L} &= \frac{\partial}{\partial L} ALf(K/AL) \\ &= Af(k) - ALf'(k)\left(\frac{K}{AL^2}\right) \\ &= A\left[f(k) - kf'(k)\right]\end{aligned}$$

This expression tells us that the wage per effective unit of labor $w(t) \equiv W(t)/A(t)$ satisfies:[4]

$$w(t) = f(k) - kf'(k). \tag{4}$$

$$f(k) = kf'(k) + F_L(K, AL)/A = kf'(k) + w$$

- This expression also tells us that the firm does not generate any pure profits. This was to be expected given that the production function exhibits constant returns to scale.

2.4 The Behavior of Households

Until now, everything was more or less straightforward. We are now gearing up to the big challenge: how to characterize optimal consumption paths? To simplify things, let's first normalize H to 1 (this is without consequence).

2.4.1 The Problem that Households Face

The problem that the household solves is therefore:

$$\max_{\{C(t)\}} U = \int_{t=0}^{\infty} e^{-\rho t} u(C(t)) L(t) dt$$

subject to the following dynamic budget constraint:

$$\dot{B}(t) = r(t)B(t) + W(t)L(t) - C(t)L(t) \tag{5}$$

where we define $B(t)$ as family wealth (i.e. the stock of saving of the household at time t). One way to think about this is as follows: suppose there are financial intermediaries in this economy. These financial intermediaries operate costlessly and competitively (these are strong assumptions!). Households deposit their wealth with these financial intermediaries. The intermediaries can then

[4] Note that we could have obtained this result more directly by observing that F has constant returns to scale, so that by Euler's Theorem:

$$F(K, AL) = K\partial F(K, AL)/\partial K + AL\partial F(K, AL)/\partial(AL)$$

dividing by AL and rearranging gives: (where we use the fact that $\partial F(K, AL)/\partial L = A\partial F(K, AL)/\partial(AL)$)

purchase capital that is rented out to firms, or -conceivably- make loans to other households. Renting capital to firms or lending to household must generate the same return $r(t)$ since there is no uncertainty (otherwise arbitrage opportunities would arise), and this return to capital is paid back to the households in the form of capital income.

Of course, in equilibrium, we will need to have $B(t) = K(t)$, since in a closed economy, the stock of savings by households must be equal to the stock of physical capital, but conceptually, it is clearer if we think of $B(t)$ as the result of saving decisions by the household and $K(t)$ as the result of rental decisions by the firms. The reason that this separation makes sense here, is that it allows us to think about what happens if the household wants to borrow (i.e. $B(t) \leq 0$).

The household takes $B(0) = K(0) > 0$ as given, and also takes as given the path of rental rates $\{r(t)\}$ and the path of wages $\{W(t)\}$.

Now, if the household could borrow unlimited amounts, then there would be an easy solution to its maximization problem: just borrow an infinite amount and never repay (i.e. simply roll over the debt, principal and interest, permanently into the future). Such a path amounts to a 'Ponzi scheme', where you keep borrowing in order to fund current consumption and roll over existing debt. Clearly we want to rule out such paths: it is unlikely that anyone would be willing to take the other side of that trade. So we are going to impose a *No-Ponzi Condition* (NPC) on the problem:

$$\lim_{t\to\infty} e^{-R(t)}B(t) \geq 0 \tag{6}$$

where $R(t) = \int_0^t r(u)du$ is the *continuously compounded* interest rate between time 0 and time t. $R(t)$ tells you how much you would earn if you invest \$1 at time 0 and keep re-investing the interest between 0 and t.

Condition (6) simply says that it is not possible for the household's debt to grow faster asymptotically than the real interest rate. You should convince yourself that it does not prevent the household from borrowing (even asymptotically).

2.4.2 The Intertemporal Budget Constraint

From the dynamic budget constraint (5), we can derive an *intertemporal budget constraint*. To see this, let's take the dynamic budget constraint (5) and pre-multiply each side by $e^{-R(t)}$, then integrate between time t and time $T > t$:

$$\int_t^T e^{-R(s)}\dot{B}(s)ds = \int_t^T e^{-R(s)}r(s)B(s)ds + \int_t^T e^{-R(s)}(W(s) - C(s))L(s)ds$$

We can use the Integration by Part formula to replace the left hand side with

$$\int_t^T e^{-R(s)}\dot{B}(s)ds = B(T)e^{-R(T)} - B(t)e^{-R(t)} + \int_t^T e^{-R(s)}r(s)B(s)ds$$

where we use the fact that $dR(s)/ds = d(\int_0^s r(u)du)/ds = r(s)$.

Substituting this expression into the budget constraint and canceling terms, we obtain:

$$B(T)e^{-R(T)} - B(t)e^{-R(t)} = \int_t^T e^{-R(s)}(W(s) - C(s))L(s)ds$$

Now, take the limit as $T \to \infty$ to obtain:

$$\lim_{T\to\infty} B(T)e^{-R(T)} = B(t)e^{-R(t)} + \int_t^\infty e^{-R(s)}(W(s) - C(s))L(s)ds$$

The NPC condition (6) implies that the term on the left hand side is positive, so the right hand side must also be positive:

$$B(t) + \int_t^\infty e^{-R(t,s)}(W(s) - C(s))L(s)ds \geq 0 \quad (7)$$

where we define $R(t,s) = \int_t^s r(u)du$ as the continuously compounded return between time t and time $s > t$.[5]

This is the **intertemporal budget constraint** faced by the household. It has a simple interpretation: the present value of consumption has to be smaller than total wealth, defined as the present value of future labor income plus current assets.

2.4.3 An Alternative Formulation of the Household Problem

It is easier to work with variables expressed per units of effective labor (the lower case variables) rather than the variables themselves. So we now transform the problem from one of choosing $\{C(t)\}$ into one of choosing $\{c(t)\}$, where $c(t) = C(t)/A(t)$.

Observe that we can write

$$\begin{aligned} u(C(t)) &= \frac{(A(t)c(t))^{(1-\theta)}}{1-\theta} \\ &= A(0)^{1-\theta}e^{(1-\theta)gt}u(c(t)) \end{aligned}$$

where we use the fact that $A(t) = A(0)e^{gt}$.

Then we can write the welfare of the household as:

$$\begin{aligned} U &= \int_{t=0}^\infty e^{-\rho t}u(C(t))L(t)dt \\ &= A(0)^{1-\theta}L(0)\int_{t=0}^\infty e^{-(\rho-n-(1-\theta)g)t}u(c(t))dt \\ &\propto \int_{t=0}^\infty e^{-\beta t}u(c(t))dt \end{aligned}$$

where $\beta = \rho - n - (1-\theta)g > 0$ by assumption. It follows that maximizing U over $\{C(t)\}$ is the same thing as:

$$\max_{\{c(t)\}} \tilde{U} = \int_{t=0}^\infty e^{-\beta t}u(c(t))dt$$

Similarly, one can show that the dynamic budget constraint can be rewritten as:[6]

$$\dot{b}(t) = (r(t) - g - n)b(t) + w(t) - c(t)$$

[5] Observe that $R(t,s) = R(s) - R(t)$.

[6] To see this, observe that $\dot{B}(t)/A(t)L(t) = \dot{b}(t) + (g+n)b(t)$.

The initial condition $B(0) = K(0) > 0$ translates directly into an initial condition for $b(0) = k(0) > 0$.

Lastly, the No-Ponzi condition (6) can be expressed as:

$$\lim_{t\to\infty} e^{-R(t)+(n+g)t} b(t) \geq 0$$

So to recap, the statement of the problem we want to solve is:

Definition 1 (The Household Problem) *The household solves the following problem*

$$\max_{\{c(t)\}} \tilde{U} = \int_{t=0}^{\infty} e^{-\beta t} u(c(t)) dt \tag{8}$$

subject to:

$$\dot{b}(t) = (r(t) - g - n)b(t) + w(t) - c(t) \tag{9a}$$

$$c(t) \geq 0 \tag{9b}$$

$$b(0) > 0, \{r(t), w(t)\} \text{ given} \tag{9c}$$

$$\lim_{t\to\infty} e^{-R(t)+(n+g)t} b(t) \geq 0 \tag{9d}$$

In this problem, $c(t)$ is the *control* variable, while $b(t)$ is the *state* variable.

2.5 The Maximum Principle

To begin with, let's just state the main theoretical result. Then, we will work backwards and look at the intuition behind that result in a simple and tractable case.[7]

Proposition 1 (Maximum Principle) *Consider the Household Problem in definition 1 above. Let $c^*(t)$ be a consumption sequence that solves this problem. Then there exists a* co-state variable *$\lambda(t) \geq 0$ such that the Hamiltonian*

$$\mathcal{H}(c(t), b(t), \lambda(t)) = u(c(t)) + \lambda(t)[(r(t) - g - n)b(t) + w(t) - c(t)] \tag{10}$$

is maximized at $c^*(t)$ given $\lambda(t)$ and $b(t)$:

$$\frac{\partial \mathcal{H}}{\partial c}(c^*(t), b(t), \lambda(t)) = 0 \tag{11}$$

at all times. Furthermore, the co-state variable satisfies the following differential equation:

$$\dot{\lambda}(t) = \beta\lambda(t) - \frac{\partial \mathcal{H}}{\partial b}(c^*(t), b(t), \lambda(t)) \tag{12}$$

Finally, the co-state variable $\lambda(t)$ satisfies the Transversality Condition (TC):

$$\lim_{t\to\infty} b(t)\lambda(t)e^{-\beta t} \leq 0 \tag{13}$$

[7] This is a heuristic presentation of the Maximum Principle. For a more detailed treatment, please refer to Atle Sierstad & Knut Sydsaeter, **Optimal Control Theory with Economic Applications**, North Holland, Amsterdam, 1987, or a similar textbook.

The Maximum Principle gives a set of *necessary conditions* that any solution must satisfy. Note that it does not establish existence (it assumes that a solution exists). The conditions stated above are also *not sufficient*. Some additional conditions also need to hold (akin to second order conditions) for sufficiency.

The Maximum Principle can look a bit daunting, but if you look more closely, you will see that there are a number of familiar elements here. First observe that the co-state variable $\lambda(t)$ is similar in many ways to a *Lagrange multiplier* associated with the budget constraint of the household (9b), while the Hamiltonian in (10) itself looks like a static Lagrangian.

This is an important insight: the Maximum Principle turns a complex dynamic problem into a static one. The Hamiltonian compares the gain from additional consumption today ($u(c(t))$) versus the loss in terms of capital accumulation $\dot{b}(t)$. It assigns a 'weight' $\lambda(t)$ to that loss. In that sense $\lambda(t)$ can be interpreted as the *shadow value* of wealth $b(t)$ (in utility terms).

If we know what the shadow value of wealth $\lambda(t)$ is, then condition (11) tells us that we should simply choose $c(t)$ so as to solve this static problem, trading-off a higher utility now and a lower wealth in the future. At the optimum, we should be indifferent. Of course the dynamic dimension of the problem has not really disappeared. Rather, it is summarized by condition (12) that tells us how the co-state variable needs to evolve. We will see in the following section where this equation comes from, but we can already provide some intuition for it. Let's rewrite (12) as:

$$\beta = \frac{\dot{\lambda}(t) + \frac{\partial \mathcal{H}}{\partial b}(c^*(t), b(t), \lambda(t))}{\lambda(t)} \tag{14}$$

Now remember that $\lambda(t)$ can be interpreted as the shadow price of wealth. If we think of $\lambda(t)$ as an asset price, then the right hand side of (14) can be interpreted as the *return* on that asset. It has two components: a yield, coming from the increased value of the Hamiltonian $\partial \mathcal{H}/\partial b$, and a *'capital gain'* coming from the change in the value of the co-state variable itself. Equation (14) tells us that this return should equal β, the discount rate of the household. So the Maximum Principle states that the dynamic evolution of the co-state variable needs to ensure that the shadow value of wealth evolves in line with other asset prices in the economy.

Finally, the Transversality Condition (13) simply says that it does not make sense to accumulate large amounts of wealth. $b(t)\lambda(t)$ can be interpreted as the quantity of wealth $b(t)$ times its price (in utility terms) $\lambda(t)$. So it represents the shadow value of total wealth. If the limit term in (13) was strictly positive, it would mean that this wealth (measured in utility terms) would grow faster than the discount rate. This is not optimal: the household would be better off consuming more today and holding less wealth in the long run.

2.6 Applying the Maximum Principle to the Household Problem

We start by writing the necessary first-order condition associated with (11):

$$u'(c(t)) = c(t)^{-\theta} = \lambda(t)$$

This equation has an important interpretation. It tells us that the co-state variable $\lambda(t)$ is simply the marginal utility of consumption, along the optimal plan. So the shadow value of wealth is simply the marginal utility.

Next, we substitute $\lambda(t)$ in the differential equation (12), observing from (10) that $\partial\mathcal{H}/\partial b = \lambda(t)(r(t)-g-n)$:

$$\dot{\lambda}(t) = \beta\lambda(t) - \lambda(t)(r(t)-g-n)$$

from which we conclude that:

$$\frac{\dot{\lambda}(t)}{\lambda(t)} = -\theta\frac{\dot{c}(t)}{c(t)} = \beta - (r(t)-g-n)$$

Re-arranging terms and substituting for β, we obtain:

$$\begin{aligned} \frac{\dot{c}(t)}{c(t)} &= \frac{r(t)-g-n-\beta}{\theta} \\ &= \frac{r(t)-\rho-\theta g}{\theta} \end{aligned}$$

One can then obtain the growth rate of consumption per worker:

$$\frac{\dot{C}(t)}{C(t)} = \frac{\dot{c}(t)}{c(t)} + g = \frac{r(t)-\rho}{\theta} \tag{15}$$

Equation (15) is a key equilibrium condition. It characterizes the slope of the optimal consumption profile $\dot{C}(t)/C(t)$ at every instant. It is called the **Euler equation**.

Interpretation.

- Consumption per worker $C(t)$ grows or decline depending on the difference between the interest rate $r(t)$ and the discount rate ρ.
- This reflects the effect of the two forces that shape consumption/saving decisions. On the one hand the household is impatient and wants to consume now. On the other hand, postponing consumption today means more consumption tomorrow.
- A higher interest rate makes savings more desirable. Therefore consumption tomorrow will be higher than today's consumption and consumption growth increases.
- When the interest rate $r(t)$ equals the discount factor ρ, the two forces balance out and the household will choose to leave consumption unchanged.
- Observe the role of θ: as θ increases the IES is smaller and consumption growth is less responsive to the interest rate. This is because if the intertemporal elasticity of substitution is lower, the household is less willing to adjust consumption in response to a change in the interest rate.
- The Euler equation only pins down the 'slope' of consumption. The level of consumption will be determined by the intertemporal budget constraint (7).
- The population growth rate n does not enter into these expressions. Can you explain why?[8]

[8] Think about the effect of a higher population growth rate on preferences and on wealth accumulation.

Consider now the Transversality Condition (13). We can rewrite the differential equation for $\lambda(t)$ (12) and integrate it between 0 and t to obtain:

$$\lambda(t) = \lambda(0)e^{-R(t)+(\rho+\theta g)t}$$

If we substitute this expression in the Transversality Condition, we obtain:

$$\lim_{t\to\infty} b(t)\lambda(t)e^{-\beta t} = \lim_{t\to\infty} b(t)e^{-R(t)+(n+g)t} \leq 0 \tag{16}$$

This expression has a very simple interpretation. As we will see a bit later, the term in the exponent is asymptotically negative. In other words, the interest rate $r(t)$ will settle to a value higher than $n + g$.[9] Condition (16) states that it is optimal for the household not to accumulate extremely large amounts of wealth. To see this, observe that if the limit was strictly positive, aggregate wealth holdings $B(t)$ would be growing faster than the compound interest rate $R(t)$. Households would have to keep saving more than the interest income they earn on their wealth. In the limit, this cannot be optimal and the household would prefer to consume more.

Taken together, the No-Ponzi condition (6)) and the Transversality Condition (16) imply that the optimal consumption plan satisfies:

$$\lim_{t\to\infty} e^{-R(t)+(n+g)t}b(t) = 0 \tag{17}$$

A direct consequence is that the Intertemporal Budget Constraint (7) will hold with equality in equilibrium:

$$B(t) + \int_t^\infty e^{-R(t,s)}(W(s) - C(s))L(s)ds = 0$$

2.7 Closing the Ramsey-Cass-Koopmans Model

Applying the Maximum Principle to the Household Problem, we have shown that the optimal consumption/saving problem satisfies the following two equations:

$$\begin{aligned}
\frac{\dot{c}(t)}{c(t)} &= \frac{r(t) - \rho - \theta g}{\theta} \\
\dot{b}(t) &= (r(t) - g - n)b(t) + w(t) - c(t)
\end{aligned}$$

From the household's perspective, the rental rate $r(t)$ and the wage $w(t)$ are exogenous. But we know that in equilibrium, they must satisfy the firms' first order condition, equations (3) and (4). Finally, remember that equilibrium on the asset market requires that family wealth equals the stock of capital: $b(t) = k(t)$. Substituting these expressions we obtain the following dynamical system:

$$\begin{aligned}
\frac{\dot{c}(t)}{c(t)} &= \frac{f'(k(t)) - \rho - \theta g}{\theta} \\
\dot{k}(t) &= f(k(t)) - c(t) - (g + n)k(t)
\end{aligned}$$

This is a dynamical system in $(c(t), k(t))$, with initial condition $k(0) > 0$ and terminal condition:

$$\lim_{t\to\infty} e^{-\int_0^t f'(k(s))ds+(n+g)t}k(t) = 0$$

[9] This will be the case precisely when the condition we imposed so that utility would be finite, i.e. $\beta > 0$ holds

3 Some Intuition for the Maximum Principle

In this section, we consider a slightly simpler set-up where we can relatively easily gain some intuition for the Maximum Principle. These derivations follow closely Obstfeld's 'guide for the perplexed.' The simpler problem is as follows

Definition 2 (A Stationary Household Problem (SHP)) *Consider a household solving the following stationary problem*

$$\max_{\{c(t)\}} \tilde{U} = \int_{t=0}^{\infty} e^{-\beta t} u(c(t)) dt$$

subject to:

$$\dot{b}(t) = (r - g - n) b(t) + w - c(t) \tag{18a}$$
$$c(t) \geq 0 \tag{18b}$$
$$b(0) > 0, \text{ given} \tag{18c}$$
$$\lim_{t\to\infty} e^{-(r-n-g)t} b(t) \geq 0 \tag{18d}$$

This problem assumes that the rental rate $r(t)$ and the wage $w(t)$ in equations (18) are constant, instead of time-varying.[10]

3.1 A recursive representation

The first thing to observe is that the stationary problem has a lot of structure. In particular, it is stationary: nothing changes if we write the problem starting from some time t, with initial condition $b(t)$ instead of time 0 with initial condition $b(0)$. This implies that the problem is also nicely *recursive*. This will allow us to break it up in smaller problems that are easier to solve.

To see this more formally, define $J(b(0))$ the maximand of the stationary household problem. This is called the **value function** of the problem. It is a very important object since it encodes the maximal utility that can be achieved by the household.[11]

Associated with $J(b(0))$, there is a *feasible* consumption sequence $\{c^*(t)\}$ (possibly not unique) that delivers that utility, and an associated path for $\{b^*(t)\}$. By definition of the optimum:

$$J(b(0)) = \int_0^{\infty} e^{-\beta t} u(c^*(t)) dt$$

Now suppose that the optimal plan has been followed from time 0 to some time $T > 0$. The associated wealth at that time is $b^*(T)$. Suppose a new household head is appointed at time T and decides to solve the problem *as of time* T, that is to maximize

$$\max_{\{c(t)\}} \tilde{U} = \int_T^{\infty} e^{-\beta(t-T)} u(c(t)) dt$$

[10] Note that this implies that $R(t) = rt$.

[11] The general problem is not stationary since $r(t)$ and $w(t)$ are exogenous to the household and evolve over time. So in full generality we would want to write the value function as $J(b(t), t)$. This complicates matters somewhat and since this is not essential to the intuitions here, this is why I present the derivations for a stationary problem instead.

given $b^*(T)$, and constraints (18) from T onwards . It should be quite obvious that this new head of household will decide to carry on with the initial optimal plan. That is, it's optimal consumption will be $c^*(t)$ from T onwards. [Can you think of a proof of that result?]

The fact that the household wants to follow with the initial plan is at the heart of the notion of *time-consistency* (a time-inconsistent plan is a plan devised by the household at time 0 that the household at time T would rather not follow.)

The implication is that we can rewrite the value function as

$$J(b(0)) = \int_0^T e^{-\beta t} u(c^*(t))dt + e^{-\beta T} J(b^*(T))$$

This implies that we can think of the original problem as a *finite horizon problem* where we have to choose consumption $c(t)$ between 0 and T so as to maximize:

$$J(b(0)) = \max_{\{c(t)\}} \int_0^T e^{-\beta t} u(c(t))dt + e^{-\beta T} J(b(T))$$

where $b(T)$ is the level of family wealth that results from the consumption decisions, and subject equation (18).

This is *Bellman's principle of dynamic programming*. It means that the problem can be thought of as choosing consumption over some interval of time (here $[0, T]$), then we can rely on the continuation of the program and the fact that it follows the same structure.

3.2 A discrete time approximation and the Hamilton-Jacobi-Bellman equation

Now to make some progress on the optimal consumption-saving choice, we are going to make the time-interval $[0, T]$ very small, so small in fact that we will only have to choose consumption at one instant of time. To do this, it is best to proceed by way of a discrete-time approximation of the problem, then take the continuous-time limit as the discretization step becomes infinitesimal.

Let's carve time into tiny slices of length h. We assume that over these tiny intervals the control variables and the state variables (as well as the exogenous variables) are *constant*. We can then consider at discrete times $t = 0, h, 2h, ...$

The discretized stationary household problem consists in maximizing:

$$\max_{\{c(t)\}} \sum_{t=0,h,2h,...} e^{-\beta t} u(c(t))h$$

subject to:

$$b(t + h) - b(t) = [(r - g - n)b(t) + w - c(t)]h$$

with $b(0)$ given as before. In the above expressions, t increases in increments of h, and $c(t)$ is the (constant) value of consumption over the time interval $[t, t + h)$. Observe that the integral has been replaced by a summation, since variables are constant over intervals of length h.

The idea is now to write the Bellman equation between time t and time $t + h$:

$$J(b(t)) = \max_{c(t)} u(c(t))h + e^{-\beta h} J(b(t + h))$$

Now, we are going to expand this expression when h becomes infinitesimal. To do this, observe that we can write the following expressions (all resulting from first-order Taylor expansions):

$$\begin{aligned} e^{\beta h} &= 1 - \beta h + o(h) \\ J(b(t + h)) &= J(b(t)) + J'(b(t))(b(t + h) - b(t)) + o(h) \end{aligned}$$

Substituting these expressions in the Bellman equation we obtain:

$$\beta J(b(t)) = \max_{c(t)} u(c(t)) + J'(b(t))[(r - g - n)b(t) + w - c(t)] \quad (19)$$

Equation (19) is called the **Hamilton-Jacobi-Bellman** equation. This equation has a couple of very useful interpretations.

- A Hamiltonian interpretation. Recall the definition of the Hamiltonian in the Maximum Principle. One can see that the term on the right is simply the Hamiltonian, evaluated at $\lambda(t) = J'(b(t))$:

  $$\mathcal{H}(c, b, J'(b)) = u(c) + J'(b)[(r - g - n)b + w - c]$$

  The household needs to choose whether to consume today (which generates utility $u'(c)$) or to save and therefore increase b. Not surprisingly, the shadow price of increasing wealth is $J'(b)$, i.e. the increment in overall utility. The Hamilton-Jacobi-Bellman equation can then be rewritten as:

  $$\beta J(b(t)) = \max_{c(t)} \mathcal{H}(c(t), b(t), J'(b(t)))$$

- An asset return interpretation. We can rewrite the Hamilton-Jacobi-Bellman equation as:

  $$\beta = \frac{\max_{c(t)} u(c(t)) + J'(b(t))\dot{b}(t)}{J(b(t))}$$

  the term on the left hand side is the discount rate for the stationary household problem. The term on the right can be interpreted as the return on the optimal plan. It has two components. The first component is the current utility from consumption. This is like the yield or dividend on an asset. The second component, $J'(b)\dot{b}$, can be thought of as a capital gain. It reflects the change in the state variable and how it affects the value function. The total return to the optimal plan (the sum of the dividend and capital gain divided by the value of the asset $J(b(t))$) must equal the discount rate.

3.3 Optimality Conditions

The Hamilton-Jacobi-Bellman equation (19) gives us (11), the first condition from the Maximum Principle: given $b(t)$ and $\lambda(t) = J'(b(t))$, the optimal consumption plan satisfies:

$$\partial \mathcal{H}(c(t), b(t), \lambda(t))/\partial c(t) = 0.$$

For the Stationary Household Problem, the associated (necessary) optimality condition is:

$$u'(c(t)) = J'(b(t))$$

This equation states that the marginal value wealth (also equal to the shadow price of wealth) is equal to the marginal utility of consumption. This is an intuitive result: an increase in consumption yields a bit more utility today. But it also means less asset accumulation and therefore a decline in the value of the program by $J'(b)$

It may look as if this equation is not terribly useful since we don't know what $J'(b)$ is. To make progress, let's use the Envelope Theorem.[12] That theorem tells us that we can take a derivative of the Hamilton-Jacobi-Bellman equation (19) with respect to b to obtain:

$$\beta J'(b(t)) = J''(b(t))\dot{b}(t) + J'(b(t))(r - g - n)$$

Finally, recall that $\lambda(t) = J'(b(t))$ so that by the chain-rule theorem, $\dot{\lambda}(t) = J''(b(t))\dot{b}(t)$. Substituting into the above expression yields

$$\beta\lambda(t) = \dot{\lambda}(t) + \lambda(t)(r - g - n)$$

This is precisely the differential equation (12) for the co-state variable derived in the Maximum Principle, since

$$\frac{\partial \mathcal{H}}{\partial b} = \lambda(t)(r - g - n)$$

Finally, one can check also that $J'(b(t)) = \lambda(t)$ satisfies the Transversality Condition (13).

[12]According to the Envelope Theorem we don't need to worry about the effect of changes in $b(t)$ on $c(t)$ at the optimum.