

Pure Exhaustible Resources and Optimal Control

Lecture 5 of Economic Dynamics

Stephen J. Turnbull

Division of Policy and Planning Science
University of Tsukuba

November 4, 2022

In the first part of this lecture, we look at the theory of pure exhaustible resources.

Then we look at the problem of optimization in a dynamic context.

Oil: A Pure *Exhaustible* Resource

- Oil exists as a stock, and is used up, *i.e.*, “exhausted.” Once the stock is exhausted, there will never be any more.
 - This is an approximation: after 20 million years or so, new stocks will form in the ocean floors. But we use up oil *much* faster than that.
- Oil is *storable*, so the rental price (price of consumption) must be equal to the asset price. True of any inventory, but in theory of the firm we equate assets to (generic) capital. (*E.g.*, rental car.)
- Resources like oil that must be used up are called *pure exhaustible resources*.
- Since oil cannot be recovered once used up, several conceptual issues arise:
 - The model should have an infinite horizon.
 - Steady states are impossible.
 - Although this lecture abstracts from them, intergenerational considerations are important.

Pure Exhaustible vs. Renewable Resources

- Our analysis of the fishery, like our analysis of macroeconomic growth, focused on *steady states*.
- Both are supply-side models, lacking consumer choices. But the existence of steady states has important *economic* implications:
 - In growth theory, stable levels of *per capita* income and consumption can be maintained forever, but there are both practical maxima (the k^* for a particular socially chosen saving rate s) and an absolute sustainable maximum (for the *golden rule* s^*).
 - *Optimistic* because a Malthusian crisis need not occur. *Pessimistic* in that we can't grow out of conflicts caused by inequality.
 - The fishery is incompatible with Solow-style growth: a fixed *total* for everybody, so population growth implies *less for each* over time.
- Steady state analysis is mathematically simple.
- We have some tricks left, but the mathematics necessarily gets harder as we introduce more consumer choice. Pure exhaustible resources force the choices of *when* or *how fast* on us.

Steady States are Impossible with Pure Exhaustible Resources

- Actually, there is an infinite set of possible steady states, but they are all uninteresting, except for one, and that one is impractical.
- Steady states where the state variable is the stock imply zero consumption. Uninteresting and clearly not economic equilibrium.
- As in growth theory, consider per capita stock as the state variable. Positive consumption requires decreasing stock, so maintaining steady state with positive consumption implies decreasing population.
 - Matching willingness to accept restrictions on family size and the preferred level of resource consumption seems unlikely.
- This implies we need a utility function (criterion for trading consumption in one period against consumption in another).

Intergenerational Considerations

- Since people live a finite time, it is nearly certain that their preferences for consumption during life will differ from their preferences for consumption later, even if they care about their children and future generations.
 - They may value them less, or
 - Care about their (differing) utility rather than consumption, or
 - Be imperfectly informed about the future.
- But the current generation's actual choices change the *constraints* for future generations in a way the future generations cannot affect the current generation.
- This is different from growth theory where you can always return to the “golden rule” steady state.
- For simplicity, we ignore issues of intergenerational equity, but they are extremely important in practice (consider Japan right now!)

Exhaustible Resources Compared to Growth Theory: Ideas

- We consider resources which exist in finite amount, and necessarily are depleted by consumption.
- Contrast growth theory based on factors which are produced and accumulated without bound: capital, technology.
- It's not enough to just change the sign. There is a lower bound of zero for exhaustible resources; they cannot keep decreasing at a given rate forever.

Exhaustible Resources Compared to Growth Theory: Modeling

- Because of depletion, we cannot use steady state analysis.
- Steady state analysis, although based on a *dynamic model*, allows us to think about the economics in the same way as we do with equilibrium.
- Without a “good” steady state to aim at (e.g., Solow’s Golden Rule), we *must* make explicit intertemporal tradeoffs, *i.e.*, dynamic economic analysis. It’s not possible to reduce the (pure) exhaustible resource problem to a pure dynamic model that is economically interesting.

Pure Exhaustible Resources

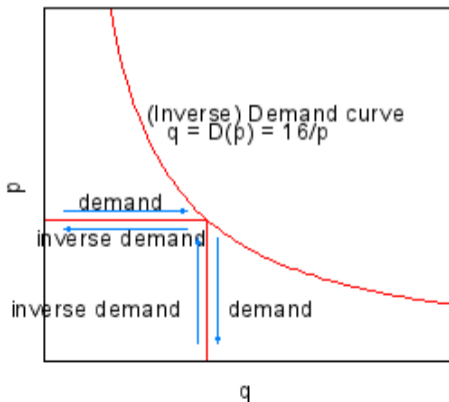
- Abbreviate *pure exhaustible resource* to *exhaustible resource*.
- Exhaustible resources are *rival* goods.
- They may be non-excludable (ocean fishing), partially excludable (large oil fields), or completely excludable (small oil fields).
 - Socially optimal usage patterns don't vary; they depend only on the stock.
 - Degree of excludability helps determine market structure, and the equilibrium usage patterns are different.
- The basic ideas of dynamic optimization can be seen with minimum technical difficulty in a model of a monopoly business which owns the whole stock of a resource.

- Single-decision-maker, completely excludable exhaustible resource.
 - Social planner or monopoly; monopoly is simpler.
 - Unlike intellectual property, there are no spillovers from consumption.
 - Monopoly (or government) is special: other producers are excluded.
- Production subtracts from the *stock* of the resource, but provides a *flow* of benefits to consumers. Like saving in growth theory, where production (Y) can be used as consumption $((1 - s)Y)$ or savings (sY) , and when saved is *added* to the capital stock (\dot{K}).
 - We suppose production is costless for simplicity; marginal analysis is possible as usual.
- Critical point: price of the good sold to consumer must equal price of the good saved as asset: they are perfect substitutes.

- Consumer behavior: market demand curve.
 - Consumers do not store the good.
 - The demand curve will be the same for all market structures.
 - Examples: *linear* (constant-slope) and *constant-elasticity* demand curves.
- Assume the market demand function for the resource is constant over time. At each instant of time, the relation between price and total quantity demanded of the resource is the same.
 - After covering the basic theory, we will consider the adjustments that must take place in a growing economy.

Demand

- We denote the instantaneous or one-period *demand curve* by $q = D(p)$. We also use the *inverse demand curve* $p = D^{-1}(q)$, which has the same graph. The inverse demand curve is also called the *marginal willingness to pay curve*.



Dynamics and Demand

- With an exhaustible resource, price must rise to choke off demand.
- Suppose $p_t \leq \bar{p}$ for all t . Then $q_t = D(p_t) > D(\bar{p}) \equiv \bar{q}$ since demand is downward sloping. For example, in the figure above, you could take $\bar{p} = 4$, implying $\bar{q} \equiv D(4) = 4$.
- If current stock is S_0 , stock is exhausted no later than time $T \equiv S_0/\bar{q}$, when price must rise to p where $D(p) = 0$, to maintain equilibrium.
- Setting price $= \bar{p}$ until stock runs out, then jumping to choke price, is not equilibrium. The marginal customer at time T has much lower value than the marginal customer at time $T + 1$. Under a plan to exhaust the resource in time T , both a profit-making firm and a social planner want to reduce consumption in time T and increase it (to greater than zero) in time $T + 1$.
- This applies to any pair of periods. Price rises gradually, forever.
 - “Gradually” is relative here: it’s *exponential*.

Asset Pricing as Portfolio Choice

- The exhaustible resource is an asset, and is priced by comparing it to other assets, in particular, bonds. We consider whether the firm's future profitability is increased by selling more of the resource and buying more bonds, or by selling less of the resource and buying less bonds (*N.B.* marginal analysis).
- Suppose that at date 0 the market price of the resource is P_0 . This is both the price as a commodity (sold to consumers) and as an asset (held for its future value).
- Suppose that the firm can buy or sell bonds with an interest rate of r . (That is, paying 1 yen for a bond today will yield a return of $1 + r$ yen tomorrow.)

Asset Pricing as Portfolio Choice, cont.

- The firm decides quantity by comparing two plans for the marginal unit:
 - a. Invest p_0 in the resource now, returning p_1 tomorrow.
 - b. Invest p_0 in bonds now, returning $(1 + r)p_0$ tomorrow.

The investments are the same, so we just compare the future values:

$$(1 + r)p_0 > p_1.$$

- Then the firm faces one of three situations about P_1 , the next period price:
 - $p_1 < (1 + r)p_0$. Bonds increase in value faster. The resource is *overpriced* today.
 - $p_1 = (1 + r)p_0$. neither increases in value faster. the resource is *correctly priced* today.
 - $p_1 > (1 + r)p_0$. Resource value increases faster. The resource is *underpriced* today.

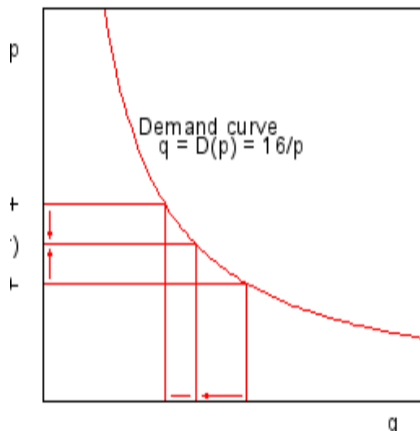
In present value terms, $p_0 > \frac{1}{1+r}p_1$.

The preceding argument is called an *arbitrage argument*. If there is a sure way to make money by changing your portfolio of assets in the financial markets, *then the market is not in equilibrium*. Taking advantage of such a situation is called *arbitrage*. The activity of arbitrage tends to have adverse effects on the *terms of trade*, i.e., increasing demand, and price, in the lower-priced market, and decreasing price by increasing supply in the higher-priced market. Eventually the price differential is eliminated, and the markets are in equilibrium.

This gives us an equation, the two markets must have the same price, that characterizes equilibrium. So, we have learned how to evaluate the price path by looking at the *financial market equilibrium*. This is common in dynamic studies.

Interaction of Investment Decision and Market Price

The firm is a monopolist, facing downward sloping demand. Its “investment decision” moves the terms of trade against it.



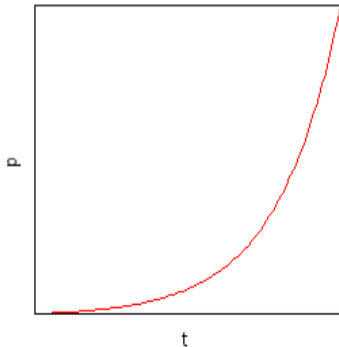
Dynamic Price Path

The firm must expect the price to rise over time according to the rule $p_{t+1} = (1+r)p_t$, i.e. $p_t = (1+r)^t p_0$.

Otherwise the firm will want to “play the market,” but this also alters the market price of the exhaustible resource, returning the price to this path.

For given p_0 , at time t the depletion is $D((1+r)^t p_0)$, and the stock at future time T is

$$S_T = S_0 - \sum_{t=0}^{T-1} D((1+r)^t p_0).$$



Initial Price p_0

- For given p_0 , at time t the depletion is $D((1+r)^t p_0)$, and the stock is

$$S_T = S_0 - \sum_{t=0}^{T-1} D((1+r)^t p_0).$$

- Price p_0 cannot be too low, or $S_0 < \sum_{t=0}^{\infty} D((1+r)^t p_0)$, and the stock is used up in finite time. This would drive price higher than the path assumed ($(1+r)^t p_0$ is finite).
- What if price p_0 is too high? Then $S_0 - \sum_{t=0}^{\infty} D((1+r)^t p_0) > 0$, and some of the stock is never used.
 - Social planners (maximizing utility) would always want to supply the excess.
 - A monopolist might not sell all with sufficiently inelastic demand. Optimal initial price depends on elasticity of demand.
 - The situation for a competitive industry is complicated.

- If the choke price is not infinite (e.g., with linear demand), in equilibrium the stock will be used up on the date when the price hits the choke price.
 - Interpretation: there is a perfect substitute whose price is the choke price.
 - So when price hits that level switch to the substitute.
- The social planner will surely use up all of the stock at infinity; this results in maximum benefit to society.
- If stocks are excludable, a competitive market in stocks will achieve the social optimum.
- If stocks are not excludable, the tragedy of the commons probably results in overexploitation, and possibly exhaustion in finite time.

The Hotelling Rule

- It is convenient to analyze this model in continuous time.
- By a usual kind of limiting argument applied the length of each period of time in the discrete arbitrage equation:

$$p(t + \delta) = (1 + r)^\delta p(t),$$

we can compute a continuous price path according to the differential equation

$$\dot{p} = rp.$$

- This is the *Hotelling Rule* (after the economist Harold Hotelling).

Initial Condition and Equilibrium

- We have characterized the price path in terms of the no-arbitrage condition (the Hotelling Rule). But we haven't given the level of the price.
- First, integrate Hotelling's Rule to get $p_t = p_0 e^{rt}$ (again assume r constant over time for computation).
- Let's try the consumption path where quantity supplied is equal to quantity demanded, at a price determined by inverse demand.
- Long-run resource balance condition requires $\int_0^\infty D_t(p_0 e^{rt}) dt = S_0$, which is complicated but can be solved to get p_0^* .
- So we propose a price path of $p^*(t) = p_0^* e^{rt}$.

Verifying the Equilibrium

- Of course “supply = demand” strongly suggests market equilibrium, but we need to confirm that all agents are making optimal decisions.
- Is $p^*(t) = p_0^* e^{rt}$ an equilibrium? Yes:
 - The condition of consumers' optimum in each period is implied by the assumption that price is at the inverse demand of the quantity sold. (Note: we call such consumers are *myopic* (near-sighted): they don't consider the future.)
 - No firm can profit by selling more now: they run out in finite time because of resource balance.
 - The argument against selling less now is not the “mirror image”! Resource balance is not a constraint against consuming less.

Uniqueness of Equilibrium

- Have we shown that there is *no* equilibrium that leaves some resource “left over at the end of time”?
- No! Because Hotelling’s Rule makes them indifferent, the “other firms” *may* optimally choose to fill up the gap left by the leader now, and “give way” to the leader when it comes back in the future, keeping price the same. Then, they are indifferent between earning now and later (when the *leader*, the firm that sells less now, sells its “excess reserve”) because of Hotelling’s Rule.
- However, there may be other equilibria. It may be that all of the firms decide to “follow the leader” and maintain a higher price. As long as forever after, the price follows Hotelling’s Rule from the current price, no firm has an incentive to change its behavior and sell its reserve early, as long as it believes the other firms will *not* “give way” (and therefore the price will drop).

Is There a Proof of Multiple Equilibria?

- Can this argument be proved? The answer is that it depends on the elasticity of demand and the interest rate. If demand is very inelastic and the interest rate high, the profit to reducing supply now (and forcing price up) may be great enough to support a restriction on supply for ever.
- But if demand is elastic, profit from selling a little more now at (almost) no reduction in price is very great. A firm forecasting it has excess stock forever wants to sell it at the current price, returning price to the level at which the stock is just exhausted “at the end of time.”
- Then that equilibrium is unique.
- This argument corresponds to a mathematical condition called *transversality*.

Rational Price Bubbles

- A “bubble” is a price path that is not supported by fundamental value of the good or security. Normally bubbles are *irrational*, and eventually “burst” when it becomes clear to some agents that excess prices cannot be sustained by actual value. However, in this market, when the equilibrium is not unique, the non-competitive (high price) equilibria are what is called “rational bubbles”.

Why infinite horizon?

- If we set a finite planning horizon T , then a rate of consumption of $1/T$ is an obvious plan, and usually a pretty good approximation to the best plan.
- But this naturally uses up all of the stock.
- This is not a problem in inventory management: you just order more for the next planning period. But the definition of *exhaustible* is that once used up, there will never be any more. For any reasonable finite horizon, the question of “what do we do after the resource is used up?” becomes critical—we, and the need that the resource satisfied, will still exist “after.”
- Since it's hard to imagine that if you manage to survive to day N , there is *zero* chance of surviving to day $N + 1$, it becomes natural to consider an infinite horizon (this is the *principle of mathematical induction*).

Infinite Horizon Models

- We prefer models with an infinite horizon because they correspond to an autonomous recursive model. That is, we know that tomorrow is mostly like today.
- In particular, just as tomorrow follows today, the day after tomorrow follows tomorrow. Since tomorrow can never be yesterday, time is an infinite sequence of days. There's always a next day and it is never a day we've already experienced.
- Also, with finite horizon there is the technical problem of “end-point effects”: if the calculation ends at time T , then it will give the highest value to consuming all remaining goods at time T . Such behavior is not only unrealistic, the expectation of an “end-time potlatch” alters incentives in the near future.

Preferences in Infinite Horizon Models

- Since human needs change little from time to time, and must be satisfied at each point in time, we use a (additively) *separable representation* of preferences: $U(c_0, c_1, \dots) = \sum_{t=0}^n u_t(c_t)$.
- It seems natural to simplify by assuming symmetry, $u_t(c_t) \equiv u(c_t)$ for all t . But there are two difficulties:
 - In steady state, $\sum \bar{u} = \infty$ for every steady state $\bar{u} > 0$, so maximization is impossible because we can't compare infinities in a useful way! We can compare steady states (as we did in growth theory), but we can't be sure that there is no non-steady state preferred to any steady state (as happens in the “collapsing competitive fishery”).
 - With exhaustible resources, we run out so $U(c_0, c_1, \dots) < \infty$, but still

$$U(1, 0, \dots) = U(0, 1, 0, \dots) = U(0, 0, 1, 0, \dots) = \dots$$

We can't decide when to consume the last one!

Preferences with Infinite Horizon (2)

- Both problems can be solved with *discounting*:

$$U(c_0, c_1, \dots) = \sum_{t=0}^{\infty} \delta_t u(c_t),$$

where $\sum_{t=0}^{\infty} \delta_t < \infty$. Usually we put $\delta_t \equiv \delta^t$ (so that $\delta_0 = 1$).

- Other possibilities sometimes used in advanced theory include
 - long run average utility*:
$$U(c_0, c_1, \dots) = \lim_{m \rightarrow \infty} \sup_{n > m} \sum_{t=0}^n \frac{u(c_t)}{n},$$
 - and the *overtaking criterion*:
 $(c_0, c_1, \dots) \succ (c'_0, c'_1, \dots)$ if and only if there exists n such that $c_t > c'_t$ for all $t > n$.
 - These emphasize the very long run, and so are used for evaluating government policy. They're not appropriate for consumer or business optimizations.

Constraints in Infinite Horizon Models

- Constraints are easier than preferences.
- A bound on total consumption of the resource over time:
$$\sum_{t=0}^{\infty} c_t \leq X_0, \text{ or}$$
- a recursive constraint as in Solow's growth model:
$$K_{t+1} \leq K_t - D_t + F(K_t, L_t) \text{ (Solow used a differential equation)}$$
- or period by period constraint $c_t \leq \bar{c}_t$ or $c_t \geq \underline{c}_t$.

Backward Induction: An Example

- A classic example of dynamic optimization from game theory is the stage game of the Chain Store Paradox. There is a large chain of stores (the *monopolist*) facing a new *entrant* with one store.
- After the entrant comes in, the monopolist has a choice of cutting price to chase him out, or sharing the market (here, at the monopoly price). See the table. Of course the monopolist will threaten to fight.
- There are *two* Nash equilibria in pure strategies (and none in mixed strategies). But is the deter/out equilibrium plausible?
- Not if the monopolist decides on quantity *after* seeing entry behavior, and the *entrant* plans ahead, thinking about the monopolist.

	deter	share
in	-1, -1	1, 1
out	2, 0	2, 0

Backward Induction: Adding Dynamics

- deter/out is implausible if the *entrant* plans ahead, thinking about the monopolist.
- This structure can be represented as a *tree*.

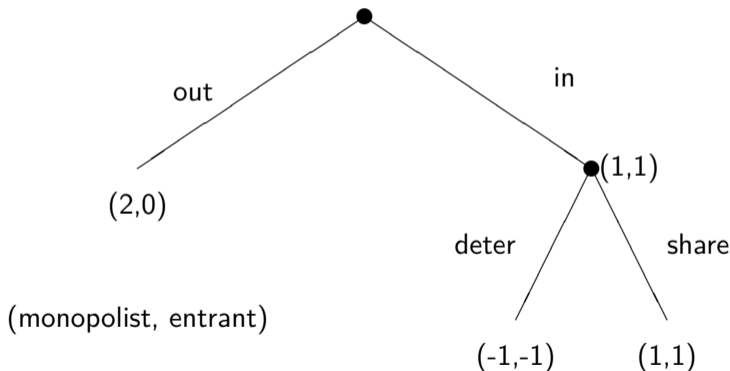


Figure: The Tree Form of the Chain Store Game

Backward Induction: “Perfect” Equilibrium

- The entrant thinks, if I enter, M's $1 > -1$, so he will *share*.
- Therefore I should choose *in*.
- The monopolist, observing entry, does indeed want to *share*.

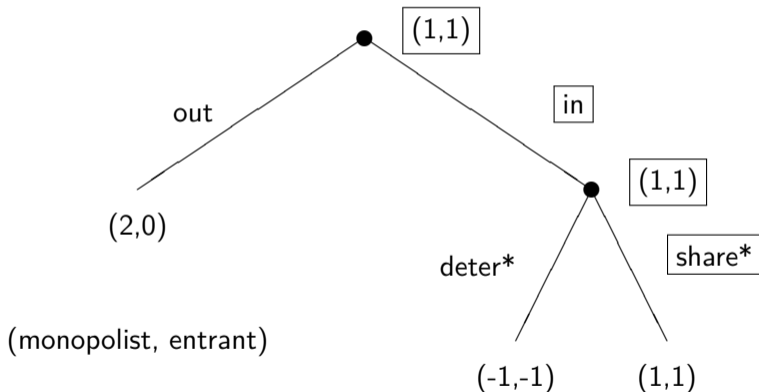


Figure: The Solution to the Chain Store Game

Backward Induction

- Backward induction is a method of solving dynamic problems.
- First, construct a *tree* to represent the dynamic structure.
 - Branches represent alternative behaviors.
 - The “chain-store paradox” game involves only agents’ decisions, but we can also represent exogenous uncertainty as a “choice” by “nature” or “a god”.
- Add payoffs to represent agent preferences.
- Start at the *end* (leaves) of the tree, solving each small subtree, then work “backward” to the root.

- Present discounted value
 - “Price” of profit in period t is $\left(\frac{1}{1+r}\right)^t$
 - Add up for all periods $\sum_{t=0}^T \left(\frac{1}{1+r}\right)^t \pi_t$
- Discounted utility
 - $\sum_{t=0}^T \delta_t U_t$
 - Usually $\delta_s < \delta_t$ for $s > t$
 - Often $\delta_t = \delta^t$
- If we propose a *plan* for the first T_0 periods, we can define a *value function* for the plan $x_0, x_1, \dots, x_{T_0-1}$ (x_t is the decision at time t) as $\sum_{t=0}^{T_0} \delta_t U_t$.

Suppose you are planning for an expense far in advance. In 20 years you will need to have savings of \$100,000 (about 10,000,000 yen).

There are many plans:

- Put \$100,000 in the bank today ($r = 0$).
- Put \$100,000 in the bank in 20 years.
- Put \$5,000 in the bank every year.
- put \$10,000 in the bank every other year.
- Put \$5,000 in the bank every other year, and in year 11 put in \$50,000.

In fact, there are an infinite number of ways. And the decision is very complex; for example, all of the above are probably impossible.

How can you make a plan for this?

Backward Induction

- It's hopeless to start out by picking a target savings this year. What you want to do this year will be affected by what you intend to do next year. And what happens if next year you realize that you really should have saved more this year? Won't you regret that?
 - So you have to make a whole plan for each of the next 20 years, and if you have any weak points, you must start over again.
- There is an important simplification that we can make because of the irreversibility of time. **Nothing you do in year 20 affects feasibility of actions in years 1–19.**
 - So year 20 actions depend only on your preferences and what you've done to that point. And you will *know* what you've done, it's history!

Backward Induction Example

We consider a smaller example, of only 4 years, needing \$20,000.

- For example, after year 3, you will have X_3 in savings, and you'll know exactly what X_3 is—it's in your account passbook. So your decision in year 4 is trivial: you must save $20000 - X_3$ to meet your goal.
- And after year 2, you'll have X_2 in savings. Your decision looks complicated: you must decide for year 3 and also year 4. But wait: decide your savings S_3 , and that determines $X_3 = X_2 + S_3$. But then you automatically know what S_4 must be.
- Of course, the balance between year 3 and year 4 will be determined by the relative utility cost of saving (not consuming) in each year. But you can make this decision with only a limited amount of knowledge: X_2 . That decision then forces the optimal decisions for all later periods (of course, there's only one).

Backward Induction Optimal Saving

- So, we can write
 - the optimal decision for savings in year 4: $S_4^*(X_3)$;
 - and for year 3: $S_3^*(X_2)$;
 - but wait, $X_3 = X_2 + S_3$, so we can write $S_4^*(X_2 + S_3^*(X_2))$.
 - Now both S_4^* and S_3^* depend on X_2 .
- This process can be continued for period 2:
 - for year 2: $S_2^*(X_1)$;
 - for year 3: $S_3^*(X_1 + S_2^*(X_1))$;
 - for year 4: $S_4^*(X_1 + S_2^*(X_1) + S_3^*(X_1 + S_2^*(X_1)))$.

Backward Induction: Principle

- Written in algebra, it is quite complicated.
- However, when you solve for each period, you *assume* you have a solution for all later periods; there is actually only one “active” dependent variable (S_t) to solve for in each period, and one independent variable (the X_t). This is *relatively* simple.
- We can write a *formula* for each S_t , which we call $S_t^*(X_{t-1})$.
- This proves it's possible to solve but normally we don't solve by algebra. Instead, we use a *recursive* solution based on *value functions*.

First, let's finish the solution.

Backward Induction: Initial Conditions

- This could go forever, except for one thing: you must have an initial condition for X_0 . For example, most college students can assume $X_0 = 0$. Since all the S_t^* are actually formulæ, $S_1^*(X_0)$ is a formula, and you can put the known value of X_0 in and compute it.
- Now you know $X_1 = X_0 + S_1$, and you can substitute into the formula $S_2^*(X_1)$, and so on. Working all the way back up the tree to $S_4^*(X_3)$, and you're done.
- This is very similar to the process of “Gaussian elimination” used in solving linear equations. The only difference is that the S_t^* functions are usually more complex.

Backward Induction: Utility and Value Functions

- The main problem left is deciding the balance between one year and the (whole) future. This is done using a *value function*.
- With no discounting, an income of \$10,000 each year, and utility of consumption in each period is $\ln C$, total utility is $\sum_{t=1}^4 \ln C_t$.
- As before, $S_4^*(X_3) = 20000 - X_3$. It's forced.
- It turns out to be useful to figure out the maximum utility in terms of the previous amount of savings:

$$V_4^*(X_3) = \ln(10000 - S_4^*(X_3)) = \ln(X_3 - 10000)$$

- This is the **value function** of the savings after 3 years. We can assume a value function for “period 5” like $V_5(X_4) = -\infty$ for $X_4 < 20000$ and $V_5(X_4) = 0$ for $X_4 \geq 20000$, which is just a mathematical way of saying “forced to have 20000.” This notation is useful because we can generalize to problems where there is *scrap value*, or there is a non-infinite penalty if $X_4 < 20000$.)

Backward Induction: Example and Value Functions

- Note that the value function gives a *future* value V_{t+1} in terms of a *current* variable X_t .
- Now when we come to year 3, for $S_3^*(X_2)$, we need to maximize $\sum_{t=1}^4 \ln C_t$, but in fact this is $\sum_{t=1}^4 \ln(10000 - S_t)$. Furthermore, we can no longer change our decisions about S_1, \dots, S_2 : they happened in the past.
- So we want to maximize $\ln(10000 - S_3) + \ln(10000 - S_4)$ subject to $X_2 + S_3 + S_4 = 20000$, which has *two* variables.
- However, we can take account of these by using the value function for the period 4 utility:
 $u(C_3) + V_4^*(X_3) = u(10000 - S_3) + V_4^*(S_3 + X_2)$ and now everything is in terms of predetermined variables (X_2) and a single decision S_3 .

Backward Induction: Example and Value Functions, cont.

- So

$$V_3(X_2) = \max_{S_3} \ln(10000 - S_3) + \ln(X_2 + S_3 - 10000)$$

- The first order condition is $\frac{1}{10000 - S_3}(-1) + \frac{1}{X_2 + S_3 - 10000} = 0$ which simplifies to $S_3 = 10000 - \frac{X_2}{2}$ and implies $S_4 = 10000 - \frac{X_2}{2}$.
- The *value function* is simply the maximized value of the remaining utility, and is given by substituting in the optimal savings:

$$V_3(X_2) = 2 \ln\left(\frac{X_2}{2}\right).$$

- This is the **utility value of the savings** after 2 years.

Backward Induction: Example and Value Functions, cont.

- It is the next step that is the key. Instead of savings for periods 3 and 4 as tradeoffs to savings for period 2, use the value function!
- To determine S_2^* , we maximize $\ln(10000 - S_2) + V_3(X_2)$ subject to $X_1 + S_2 = X_2$, or $\ln(10000 - S_2) + V_3(X_1 + S_2)$, which is

$$\ln(10000 - S_2) + 2 \ln\left(\frac{X_1 + S_2}{2}\right).$$

- The first order condition is $0 = \frac{1}{10000 - S_2}(-1) + 2 \frac{2}{X_1 + S_2}(\frac{1}{2})$ or $X_1 + S_2 = 2(10000 - S_2)$.
- The solution to the first order condition is $S_2 = \frac{20000 - X_1}{3}$, and the value function is $V_2(X_1) = 3 \ln\left(\frac{X_1 + 10000}{3}\right)$.

- For period 1, we use its value function. To determine S_1^* , we maximize $\ln(10000 - S_1) + V_2(X_1)$ subject to $X_0 + S_1 = X_1$, or more simply $\ln(10000 - S_1) + V_2(X_0 + S_1)$, which is

$$\ln(10000 - S_1) + 3 \ln\left(\frac{X_0 + 10000}{3}\right).$$

- The first order condition is $0 = \frac{1}{10000 - S_1}(-1) + 3 \frac{3}{X_0 + S_1 + 10000}(\frac{1}{3})$
- Restated $X_0 + S_1 + 10000 = 3(10000 - S_1)$
- The solution to the first order condition is $S_1 = \frac{20000 - X_0}{4}$, and the value function is $V_1(X_0) = 4 \ln\left(\frac{X_0 + 20000}{4}\right)$.

Using the Initial Conditions

- The initial condition is $X_0 = 0$.
- This gives $S_1 = 5000$ and $X_1 = 5000$.
- Then $S_2 = \frac{20000 - X_1}{3} = 5000$ and $X_2 = 10000$.
- Then $S_3 = \frac{20000 - X_2}{2} = 5000$ and $X_3 = 15000$.
- Finally $S_4 = 20000 - X_3 = 5000$ and $X_4 = 20000$, of course.
- Note that $V_1(X_0) = \sum_{t=1}^4 U(y_t - S_t)$. This is a general property: the optimal value of the value function for the whole problem is the optimal value. It's constructed that way.

These require symmetry of y_t and U_t . They can be generalized to discounted utility easily.

- $S_t^*(X_{t-1}) = \frac{X_T - X_{t-1}}{T - (t-1)}$
- $V_t^*(X_{t-1}) = (T - (t-1))U(y - \frac{X_T - X_{t-1}}{T - (t-1)})$

Optimal control theory

- There are two basic general approaches to dynamic optimization. The first is *optimal control theory*.
- Optimal control theory states the optimization problem mathematically as choice of a function of *time* from a set of such functions satisfying some constraints. The dynamic constraints are typically defined using a system of differential equations (or difference equations, when the problem is formulated in discrete time).
- The solution to the optimal control problem was first characterized by the Russian mathematician L. S. Pontryagin.
- Pontryagin's *maximum principle* is defined in terms of the derivatives of a function called the *Hamiltonian function* of the problem.

Dynamic programming

- The second approach is *dynamic programming*.
- Dynamic programming states the optimization problem mathematically as choice of a function of some *state* from a set of such functions. The special property of this function is that it gives the optimal value of solving the problem if we restart now.
 - This is very similar to use of phase diagrams in growth theory.
- The solution to the dynamic programming problem was first characterized by the American mathematician Richard Bellman, using an equation called the *Bellman equation*.
 - In stochastic problems such as those encountered in finance (e.g., the Black-Scholes model of option pricing) a more general equation called the *Bellman-Hamilton-Jacobi* equation is used.
- Bellman's *principle of optimality* is defined in terms of the function of the state, which is called the *value function* of the problem.

Comparison of dynamic optimization methods

- Any given problem can be formulated as either an optimal control problem or a dynamic programming problem and solved. *Both methods must give the same optimal value, and usually the same solution.*
- The optimal control formulation is most useful for *open-loop* or “fire-and-forget” *solutions*, where the control is given as a function of time. These are common in physical situations (such as ballistics).
 - These require a very accurate formulation of the model. You have to be willing to commit to a specific plan for the control variable, and the solution method is not very helpful in understanding how to change the control if conditions change.
- Economics generally prefers to use dynamic programming (also called “recursive methods”), which gives *closed-loop solutions* where the control is given as a function of state, and thus can be adapted to unexpected changes in conditions.

Dynamic Optimization

- Simple backward induction based on comparison to next period.
- If we have discrete time and a terminal period, the value function for the last period is constant.
- Then we use the current utility plus (discounted) future value, optimize and get next-to-last period value function.
- This process continues until first period, and we're done.
- But what if we have infinite time or continuous time, or want a steady state?

Dynamic Programming in a Steady State

- “Backward induction” doesn’t really make sense in continuous time or steady state, so we use more general term “dynamic programming”.
- In discrete time the value function was

$$V_t(X_{t-1}) = \max_{y_t} \{U(y_t) + \delta V_{t+1}(X_t)\},$$

and in steady state this would be

$$V(X_{t-1}) = \max_{y_t} \{U(y_t) + \delta V(X_t)\}.$$

- We can’t find V without knowing V itself, it seems! This is a *fixed point problem*.
 - With an ordinary variable this is not hard. We have some function $f(x)$, and the fixed point problem is “solve $f(x) = x$ for x .”
 - But for a *function*, this is hard.

Calculus of Variations

- General problem is “nonlinear programming in infinite dimensions.”
- “Classical” approach is “calculus of variations”.
 - Original use was computing things like shape of soap bubbles.
 - It uses methods like “Gateaux derivatives” and “Frechet derivatives” to derive first order conditions for the V , as we do in ordinary calculus.
- However, these methods are difficult and less general than modern *optimal control theory*.
- Advantage is direct relation to value function.

Derivatives with Respect to Functions

- A real-valued function of a function is a *functional*.
 - The definite integral of a function is a functional.
 - The value of a solution to a dynamic problem is a functional (typically expressed as a sum or integral).
- In finite dimensions, say optimizing the consumption of apples (q_a) and bananas (q_b), we take derivatives with respect to q_a , q_b , and λ of the Lagrangean

$$L(q_a, q_b, \lambda) = U(q_a, q_b) + \lambda(y - p_a q_a - p_b q_b)$$

E.g. for apples: $\frac{\partial L(q_a, q_b, \lambda)}{\partial q_a}$ and so on.

- But a function may have special constraint such as continuity.

Interpretation of the Value Function

- The value function is what is says, a way to evaluate the value of the state variable, e.g. a stock of an exhaustible resource.
 - We can compare it to the *market value* to evaluate the market performance: is $V(X) = PX$, where X is a state variable representing a stock of some good or productive input, and P is the market price of that good? If not, there is a *market failure* of some kind.

Interpretation of the Value Function's Derivative

- The derivative of the value function is comparable to the price of the state variable, which can be compared to *market prices*. Consider the derivative of V with respect to X , which is given in *units of value* (e.g., money) per *unit of the variable*. Compare to the derivative of *market value*:

$$\frac{d}{dX}PX = P.$$

We saw a similar idea with the comparison of rate of interest to the derivative of H in the theory of renewable resources.

Optimal Control Theory

- The modern approach is *optimal control theory*.
- Developed from (literal) rocket science used to control missiles and warplanes as well as launching artificial satellites.
- More general than calculus of variations.
- Pontryagin Maximum Principle

Example: Optimal Growth

- We can easily transform the Solow model to an optimal control problem.
- We start from the characteristic differential equation

$$\dot{k} = sf(k) - (n + d)k$$

by recognizing that saving is no longer determined by a constant proportion, but by choosing consumption in each period c_t . So we substitute $f(k) - c_t$ for $sf(k)$

$$\dot{k} = f(k) - c_t - (n + d)k.$$

- The system is no longer autonomous, so let's rewrite to make time explicit

$$\dot{k}(t) = f(k(t)) - c(t) - (n + d)k(t),$$

giving the law of motion.

The Objective

- As usual, we want to maximize a per-person quantity.
- To trade off present versus future consumption, we use a concave utility function $u(c)$ with $u' > 0$ and $u'' < 0$, and a discounted sum (or really an integral, in continuous time)

$$\int_0^T u(c(t))e^{-\delta t}dt.$$

The Problem Stated

$$\begin{aligned} \max_{c(t)} \quad & \int_0^T u(c(t)) e^{-\delta t} dt \\ \text{s.t.} \quad & \dot{k}(t) = f(k(t)) - c(t) - (n + d)k(t) \\ & k(0) = k_0, k(t) \geq 0, c(t) \geq 0 \end{aligned}$$

The Hamiltonian Function

- The Hamiltonian formula according to theory is

$$H(k(t), c(t), t, p(t)) = q(t)u(c(t))e^{-\delta t} + p(t)(f(k(t)) - (n+d)k(t) - c(t)).$$

- But the necessary condition requires $q(t) = \bar{q}$ for all t , and we can assume $\bar{q} > 0$ (otherwise utility has no effect on the solution!). Since as usual only relative “prices” matter, we can set $\bar{q} = 1$ as the numeraire, giving

$$H(k(t), c(t), t, p(t)) = u(c(t))e^{-\delta t} + p(t)(f(k(t)) - (n+d)k(t) - c(t)).$$

- Note that the Hamiltonian is one equation at each instant of time. While that is a lot of equations, note that we have reduced this (in principle) to a set of static problems.

The Necessary Conditions

- The Hamiltonian formula is

$$H(k(t), c(t), t, p(t)) = u(c(t))e^{-\delta t} + p(t)(f(k(t)) - (n+d)k(t) - c(t))$$

(multipliers on the nonnegativity conditions for k and c are omitted).

- The necessary conditions for an optimum are

$$\dot{k}(t) = \frac{\partial H}{\partial p} = f(k(t)) - (n+d)k(t) - c(t)$$

$$\dot{p}(t) = -\frac{\partial H}{\partial k} = -p(t)(f'(k(t)) - (n+d))$$

and $c^*(t)$ maximizes $H(k(t), c, t, p(t))$ over c .

- The solution is

$$c^*(t) = -\frac{u'}{u''}(f'(k(t)) - (n+d+\delta)).$$

Interpretation of the Solution and Conditions

- By construction, the condition $\dot{k}(t) = \frac{\partial H}{\partial p}$ simply reproduces the law of motion. This is similar to the way $\frac{\partial \mathcal{L}}{\partial \lambda} = 0$ reproduces the budget constraint in the Lagrangean formulation of the consumer's optimization: $\max_{x,\lambda} \mathcal{L} = u(x) + \lambda(y - px)$.
- The Hamiltonian is related to the value function. For this reason, it should be unsurprising that the condition $\dot{p}(t) = -\frac{\partial H}{\partial k}$ establishes a relationship between the “shadow price” p and the derivative of H with respect to k . (The negative sign is due to the fact that the Pontryagin Maximum Principle is from physics, and there naturally formulated as a *Minimum Principle*; but *maximizing* U is the same as *minimizing* $-U$, giving the opposite sign.)
- There is a *duality* relationship between the conditions for \dot{p} and \dot{k} : each is characterized by the derivative with respect to the *other*.

- The negative sign in $-\frac{u'}{u''}$ is due to $u' > 0$ and $u'' < 0$ as usual.
 - Note that u' and u'' are *time-varying*; they depend on $c^*(t)$.
 - We often use a *constant absolute risk aversion* (CARA) utility function in computing examples (even though there's no risk involved!), because then $-\frac{u'}{u''}$ is a constant.
- In a steady state, we get Solow's Golden Rule because

$$0 = -p(t)(f'(k(t)) - (n + d)).$$

- Note that the optimal path deviates only slightly from Solow's Golden Rule, by the introduction of δ in $n + d + \delta$.

Turnpike Theorems

- In many contexts (ballistics, various production processes, analysis of derivative securities, short-run growth planning), a specific state should be achieved at the end of the planning period.
 - In growth theory, the “target” at the end of time is left free.
- Often such problems can be characterized in three steps:
 - 1 Consider the set of optimal “free” solutions, parameterized by the initial state.
 - 2 Find the optimal “free” solution whose path “goes through” the target state at the end of the planning period.
 - 3 Starting from the actual initial state, adjust the controls to go to the optimal free solution as fast as possible, then switch to the controls for that path.
- Such results are called “turnpike theorems” (“turnpike” is an old word for “highway”).
- Turnpike theorems are *why* rockets burn all their fuel as quickly as possible, then coast to the destination.