Economics 1011B
Section 3

Spring 2023

Today's Outline

- Neoclassical Growth Model
    - Motivation
    - Household and Firm Problems
    - Competitive Equilibrium
    - Steady State
    - Transition Dynamics
    - Phase Diagram
    - Planner's Problem / Centralized Equilibrium
- Bonus Slides
    - Pareto Efficiency
    - First Fundamental Theorem of Welfare Economics

From Solow to the Neoclassical Growth Model

- The goal of this week is to develop a growth model that builds on the Solow model we developed in week 1, using the consumption/savings framework we developed in week 2.
- Remember that in Solow, we simply assumed that the savings rate, s, was fixed.
- But we know from micro data (on individual households) and from macro data (on aggregate savings/consumption) that savings rates are *not* fixed (in aggregate or at the individual level).
- By embedding the consumption-savings framework inside Solow, we'll be 'endogenizing' the savings decision in the context of the model.

From Solow to the Neoclassical Growth Model

- Going forward, models in this course will generally involve the same elements:
    1. Optimization: Households choose things (i.e. consumption, savings, labor supply) to maximize utility subject to constraints; firms choose things (i.e. input demands) to maximize profits subject to constraints.
    2. Equilibrium conditions: equations describing how supply and demand interact; for instance, prices are set such that supply equals demand (e.g. for capital, for consumption goods, for labor).
    3. Resource constraints: Economy-wide resource constraints hold. For instance, capital evolves according to a law of motion.

- Modern macro also includes an 'aggregation' step. In our models so far (and most places elsewhere in this course), aggregation will be trivial. This is only because allowing for heterogeneous households complicates the math quite a bit: but there are many places in macro where this is really important (e.g. studying inequality).

- This framework allows for us to incorporate arbitrary 'frictions': borrowing constraints, behavioral agents, monopolistic competition, etc.

Neoclassical Growth Model: Household Problem

- Time is discrete, indexed $t = 0, 1, 2, \dots$. A representative household chooses consumption c_t and savings a_{t+1} each period to maximize the present discounted value of lifetime utility, subject to their budget constraint.

- The household receives labor income w_t (from inelastically supplying one unit of labor) and interest income $(1 + r_t)a_t$ from last period's savings a_t.

- The household's problem can be written as (with period budget constraints):

$$\max_{\{c_t, a_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t) \quad \text{subject to} \quad c_t + a_{t+1} = (1 + r_t)a_t + w_t$$

$$a_0 \text{ given}$$

where the household takes initial savings a_0, the interest rate r_t, and the wage w_t as given.

Neoclassical Growth Model: Household FOCs

- The (present value) Lagrangian associated with this problem is:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left[u(c_t) + \lambda_t \Big((1 + r_t) a_t + w_t - c_t - a_{t+1} \Big) \right]$$

- Notice that the Lagrangian multiplier is indexed by t - one constraint each period!

- When we differentiate $\mathcal{L}$ with respect to a_{t+1}, it shows up in λ_t and λ_{t+1} (why?)

- The first-order conditions for this problem:

$$\begin{aligned} c_t : &\qquad u'(c_t) = \lambda_t \\ a_{t+1} : &\qquad \beta^t \lambda_t = \beta^{t+1} \lambda_{t+1} (1 + r_{t+1}) \end{aligned}$$

- As in the consumption-savings model, we combine them to derive the Euler equation:

$$u'(c_t) = \beta (1 + r_{t+1}) u'(c_{t+1})$$

Neoclassical Growth Model: Firm Problem

- The firm's problem is to choose capital k_t and labor ℓ_t to maximize profits each period t. The firm produces output using a CRS production function $F(k_t, l_t)$, with the price of the output good being normalized to 1. The firm borrows at an interest rate r_t to purchase capital, and they sell off undepreciated capital at the end of each period. Lastly, they pay a wage rate w_t for each unit of labor they hire.

- The firm's problem can be written:

$$\max_{k_t, \ell_t} F(k_t, \ell_t) + (1-\delta)k_t - (1+r_t)k_t - w_t\ell_t$$

- FOCs are really easy: no constraint above, so no need for Lagrangian:

$$\begin{aligned} k_t : \qquad & F_k(k_t, \ell_t) = r_t + \delta \\ \ell_t : \qquad & F_\ell(k_t, \ell_t) = w_t \end{aligned}$$

- Interpretation: The firm will hire labor/capital up until the point where the marginal product of labor/capital is smaller than the (constant) cost of labor/capital.

Neoclassical Growth Model: Competitive Equilibrium

- The competitive equilibrium of our economy consists of an allocation $\{c_t, a_t, k_t, \ell_t\}_{t=0}^{\infty}$ and prices $\{w_t, r_t\}_{t=0}^{\infty}$ such that:
  1. Household optimality: $\{c_t, a_t\}_{t=0}^{\infty}$ solves the household's problem, taking prices as given.
  2. Firm optimality: $\{k_t, \ell_t\}_{t=0}^{\infty}$ solve the firm's problem, taking prices as given.
  3. Financial market clearing: $a_t = k_t$ for all t
  4. Labor market clearing: $\ell_t = 1$ for all t
  5. Resource constraint holds: $k_{t+1} = (1-\delta)k_t + F(k_t, \ell_t) - c_t$ for all t

- Technical note: goods market clearing is implied by (3-5). At a deep level, this follows from Walras' Law in micro: $N-1$ markets clearing implies the last market clears in general equilibrium..

- Notice that this is a road-map to solve the model! 6 unknowns (prices, quantities) and 6 equations (firm optimality has 2 FOCs).

Neoclassical Growth Model: Competitive Equilibrium

- Intuitively, competitive equilibrium is the outcome that prevails when agents in our model (households, firms) are acting in their own self-interest (maximizing profit/utility), markets clear (prices are set such that supply equals demand), and any economy-wide resource constraints hold.

- The definition of competitive equilibrium pins down exactly what values our endogenous variables will take *in equilibrium*. It also provides a road-map to solving the model, by reminding us what equations we need to solve for the unknown endogenous variables (market clearing, FOCs).

- Because agents are acting in their own self-interest, it is always worth asking: in a given model, could we do better than the competitive allocation? We might if, for instance, there are externalities that agents in our model are not internalizing. More later!

Neoclassical Growth Model: Solving for Steady State

- Goal: derive difference equations for c and k expressing their $t+1$ values in terms of t values and constants. In particular, we want to substitute out prices (e.g. r and w) and other endogenous variables (e.g. ℓ).

- Start with Euler equation:

$$\begin{aligned}
u'(c_t) &= \beta(1+r_{t+1})u'(c_{t+1}) && \text{(Euler equation)} \\
&= \beta(1+F_{k,t+1}(k_{t+1},\ell_{t+1})-\delta)u'(c_{t+1}) && \text{(Firm FOC for } k\text{)} \\
&= \beta(1+F_{k,t+1}(k_{t+1},1)-\delta)u'(c_{t+1}) && \text{(Labor market clearing)}
\end{aligned}$$

- This gives an equation that implicitly characterizes c_{t+1} as a function of c_t, k_{t+1}, and constants.

Neoclassical Growth Model: Solving for Steady State

- Recall the resource constraint (law of motion for capital):

$$\begin{aligned} k_{t+1} &= (1-\delta)k_t + F(k_t, \ell_t) - c_t && \text{(Resource constraint)} \\ &= (1-\delta)k_t + F(k_t, 1) - c_t && \text{(Labor market clearing)} \end{aligned}$$

- That was easy! Given this period's capital and consumption, this difference equation tells us what capital will be next period.

- Just to recap, we have the following two difference equations:

$$u'(c_t) = \beta(1 + F_k(k_{t+1}, 1) - \delta)u'(c_{t+1}) \tag{1}$$

$$k_{t+1} = (1-\delta)k_t + F(k_t, 1) - c_t \tag{2}$$

- These play the same role that the difference equation for k played in the Solow model.

Neoclassical Growth Model: Key Equations

- To solve for the steady state, impose $c_t = c_{t+1} = c^*$ and $k_t = k_{t+1} = k^*$. The equations from the previous slide become:

$$u'(c^*) = \beta(1 + F_k(k^*, 1) - \delta)u'(c^*) \implies F_k(k^*, 1) = \beta^{-1} - 1 + \delta \quad (3)$$

$$k^* = (1 - \delta)k^* + F(k^*, 1) - c^* \implies c^* = F(k^*, 1) - \delta k^* \quad (4)$$

- Notice that (3) alone is sufficient to solve for k^*: If the production function is Cobb-Douglas, i.e. $F(K, L) = K^\alpha L^{1-\alpha}$, can use algebra to solve for k^* in terms of parameters using only (3).

- Meanwhile, equation (4) has c^* as a function of k^*.

- If we plot these equations in (c, k) space (c as a function of k), the steady state of the neoclassical growth model is where the two lines intersect.

Neoclassical Growth Model: Dynamics

- So, the steady state is easy enough to characterize: equations (3) and (4) will do. We can use those equations to think about how the steady state/long-run equilibrium will change when e.g. a parameter changes.

- But how do we think about changes in the short-run, or the transition path to the steady state?

- Multiple difference equations can be complicated to interpret on their own. One useful graphical tool to think about how c_{t+1} and K_t move toward the balanced-growth path from any initial condition (i.e. initial capital stock, or initial savings) is a phase diagram, which work especially well for two-equation systems.

- Our model involves c and K, so those will be the axes. In this space, we plot our difference equations, imposing $c_{t+1} = c_t$ and $K_{t+1} = K_t$.

Neoclassical Growth Model: Dynamics

- Before I go through these plots, I want to emphasize that the elements of the figure I'm about to draw are not exam-eligible per se. Ludwig did not cover them in lecture.

- I want to work through the "full" phase diagram for the neoclassical growth model for two reasons:
    1. It is a very common part of the canonical neoclassical growth model, and may help you understand these parts of the textbook readings better
    2. I believe it may help you understand where the optimal consumption path comes from in this model (and in Ludwig's phase diagram from lecture 5 slides).

Neoclassical Growth Model: Phase Diagram (Steady State)

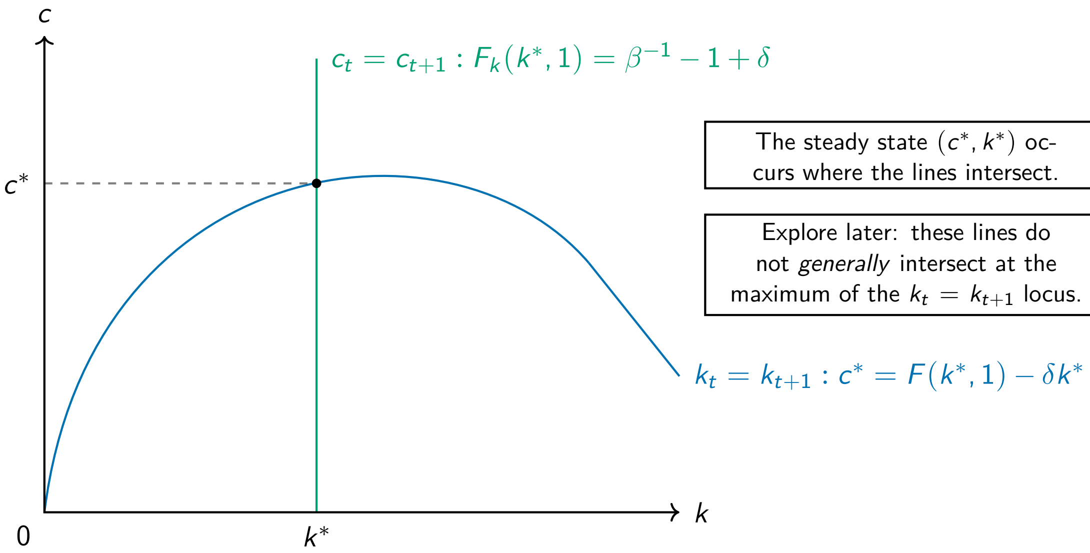

The steady state (c^*, k^*) occurs where the lines intersect.

Explore later: these lines do not *generally* intersect at the maximum of the $k_t = k_{t+1}$ locus.

Neoclassical Growth Model: Phase Diagram (Transition Dynamics)

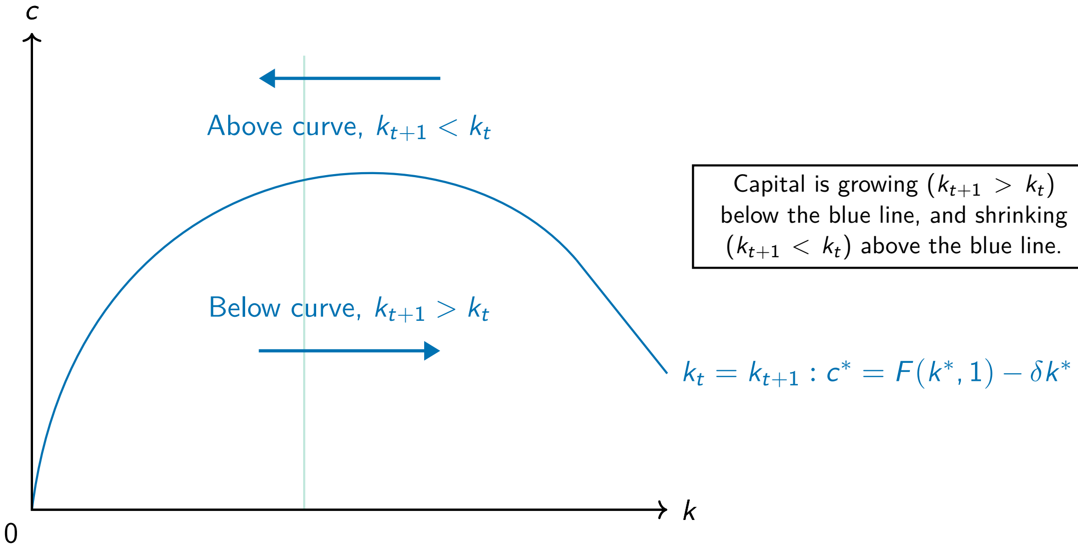

Capital is growing ($k_{t+1} > k_t$) below the blue line, and shrinking ($k_{t+1} < k_t$) above the blue line.

Neoclassical Growth Model: Phase Diagram (Transition Dynamics)

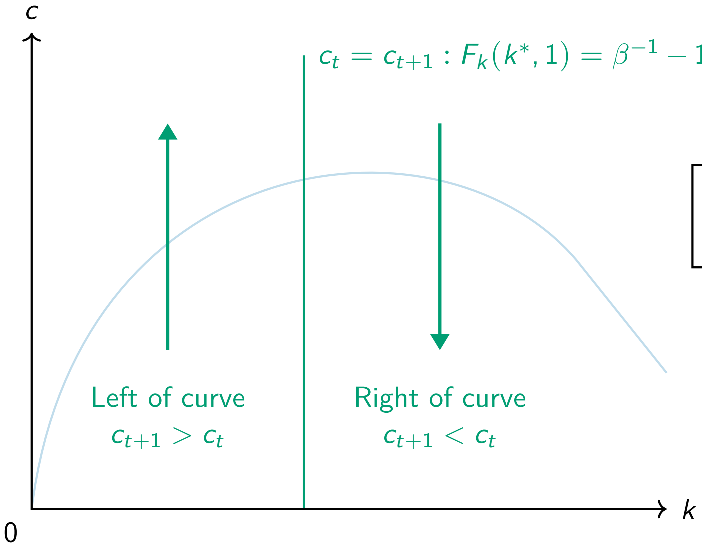

Consumption is growing ($c_{t+1} > c_t$) left of the green line, and shrinking ($c_{t+1} < c_t$) right of the green line.

Neoclassical Growth Model: Phase Diagram (Transition Dynamics)

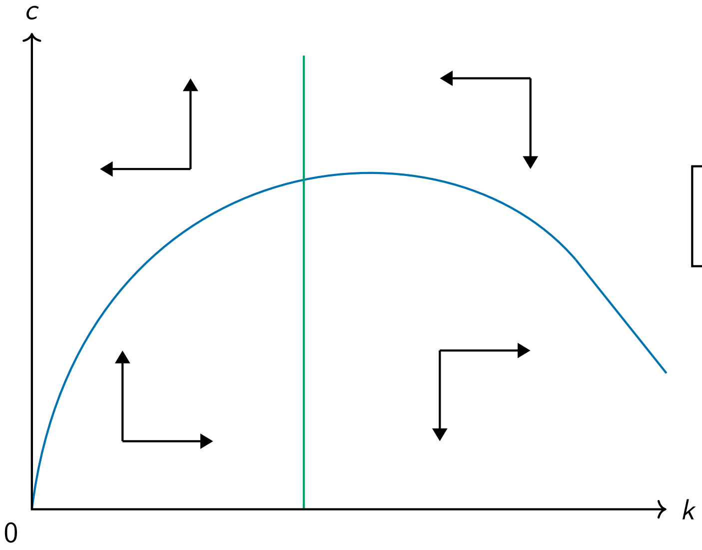

The "traffic arrows" we've drawn indicate the paths that consumption and capital must follow.

Neoclassical Growth Model: Phase Diagram (Transition Dynamics)

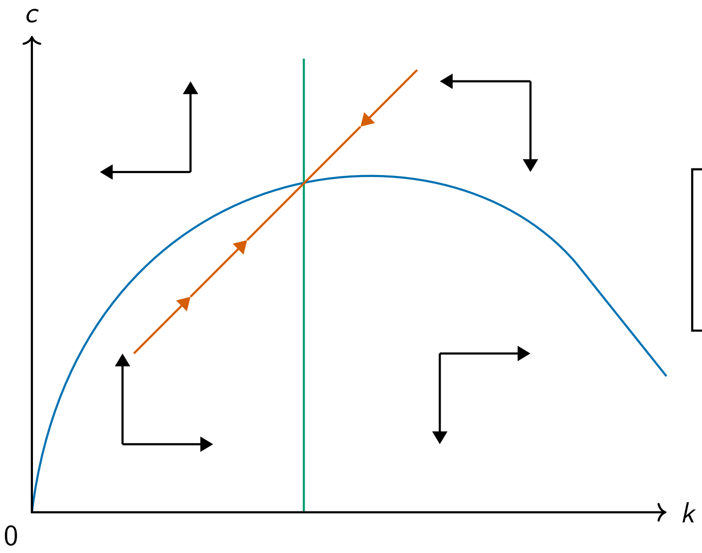

The only values that consumption and capital will take in equilibrium are on the red line – the saddle path. This is the unique path that returns to the steady state.

Neoclassical Growth Model: Phase Diagram (Transition Dynamics)

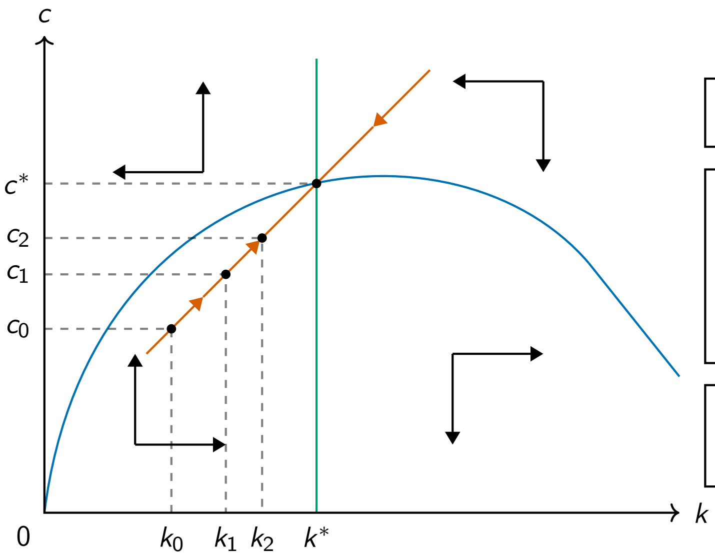

Consider an initial capital level $k_0 < k^*$ as drawn.

The c_0 corresponding to k_0 must lie on the red line: otherwise, consumption and capital would not converge to the steady state, but would instead shoot off (following the traffic signals) toward 0 or ∞.

After period 0, c_t and k_t follow their difference equations (staying on the red line) back to the steady state.

Neoclassical Growth Model: Phase Diagram (Transition Dynamics)

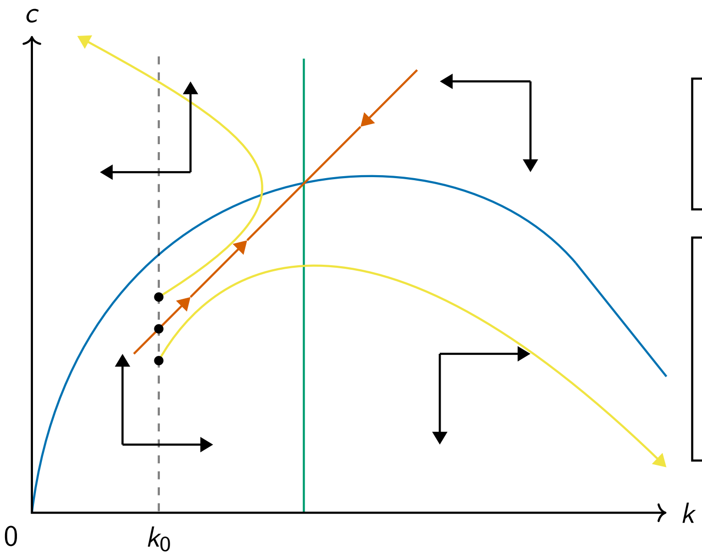

Why must c lie on the red line? What does it mean for a consumption path to not converge to steady state?

The yellow lines depict what would happen if c_0 was not on the optimal consumption path, given $k_0 < k^*$. For each value of k, there is exactly one point on the red line; the red line is sometimes called a 'saddle path'.

Neoclassical Growth Model: Dynamics

- Model exhibits what is sometimes called saddle path stability: for any given k_0, there is precisely one value of consumption c_0 that returns to the steady state. Any other value (even a little bit higher or lower!) would 'explode off', given the traffic arrows we drew, toward 0 capital or 0 consumption.

- As you can see: analyzing steady state dynamics (how does the steady state change if a parameter changes?) is much easier than analyzing transition dynamics.

- To answer what happens to variables in the steady state, we can use algebra (with our key difference equations, imposing steady state), or we can use the graphs and figure out how a parameter changing shifts our lines. The steady state always occurs at the intersection of the blue and green lines, but only values on the red line are attained in equilibrium.

Neoclassical Growth Model: Steady State and Transitions

- We can ask what happens (1) to the steady state and (2) on the transition path (to the new steady state) when a parameter in this model changes: for instance, β, α, or δ.

- What happens to steady-state capital and consumption if β falls? If we were at the steady state prior to β falling, how does consumption respond immediately after the steady state?

- What about if δ falls? Make sure to ask what happens to both lines!

Neoclassical Growth Model: Phase Diagram (Example)

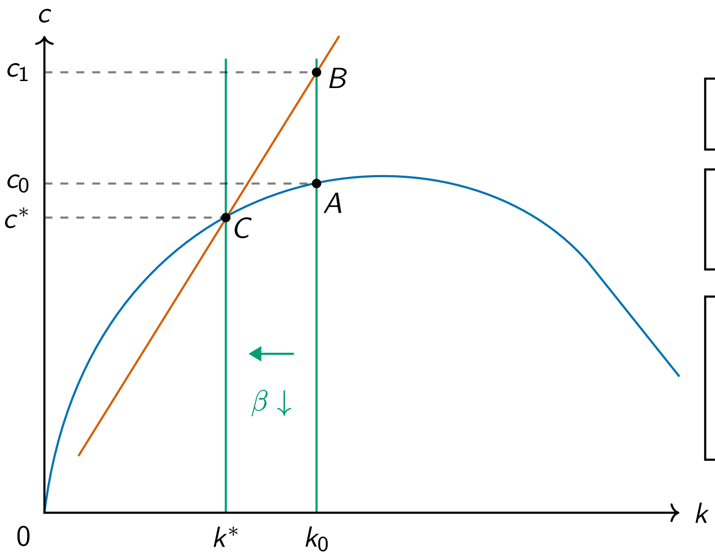

Starting at a steady state (c_0, k_0), what happens if $\beta \downarrow$?

$\beta \downarrow$ shifts the green locus inwards (why?). The new steady state has lower c and k.

On the transition path, c is higher! (less savings). Shoot up to point B. Then follow new optimal consumpton path down to new steady state $(c*, k^*)$.

Neoclassical Growth Model: Phase Diagram vs. Lecture

- The preceding figures are the canonical phase diagrams for neoclassical growth, exactly as they appear in our textbooks (Kurlat Fig 9.3.1; Romer Figs. 2.2, 2.3, 2.4)

- The phase diagram in the Lecture 5 slides is a somewhat simplified version of these graphs, intended to convey just the key intuition.

- The red line on the phase diagram (known as the saddle path/optimal consumption path/policy function) is **not** given by either of the difference equations in the model per se: the difference equations give us the 'traffic arrows'. The orange line represents the only set of points in (c, k) space for which this model converges back to the steady state when we apply the difference equations.

- In fact, there isn't even a closed-form solution for the red line generally! To solve for the red line (which is generally not linear as drawn), you need to use numerical approximation techniques - beyond the scope of 1011B.

Neoclassical Growth Model: Recap

- We have nice, closed-form solutions for the steady state values c^* and k^* in this model. We can obtain them from setting $c_t = c_{t+1} = c^*$ and $k_t = k_{t+1} = k^*$ in our key difference equations and using algebra to rearrange.

- The red line in our phase diagram indicates the path of values that consumption and capital can take in this model. **We do not actually have a closed form equation characterizing the red line**. Solving for this (sometimes called a policy function) is beyond on the scope of this course, requiring some numerical approximation tools. So for transition dynamics in this model (assessing what happens to c, k, w, r etc if some parameter changes), the graphical approach is best. We have waved our hands a bit on this out of necessity.

- The neoclassical growth model is therefore a little delicate in the transition. You are not expected to know about saddle path stability - Ludwig intentionally did not raise these points in lecture. I'm presenting them just to provide additional context.

Neoclassical Growth Model: Centralized Equilibrium / Planner's Problem

- The competitive equilibrium we have solved for is sometimes called a decentralized equilibrium, because individual economic actors (households and firms) make their decisions independently in their own self-interest.
- It is worth asking: If we could directly assign allocations to people (i.e. force households and firms to adopt the consumption/savings and input demand profiles we want, in order to maximize the household's utility), could we do better than the competitive equilibrium?
- This perspective is sometimes called the planner's problem, or the centralized equilibrium. It's a hypothetical benchmark that is easier to solve, and doesn't involve prices (because there are no markets - allocations are assigned, not bought and sold). We only have to solve one problem rather than every agents' separately!
- This is 'bonus content' in Ludwig's lectures – not essential, but included here because we have time.

Neoclassical Growth Model: Centralized Equilibrium / Planner's Problem

- The planner's problem in this case is to choose consumption and capital to maximize utility:

$$\max_{\{c_t, K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t) \qquad \text{subject to} \quad K_{t+1} = (1-\delta)K_t + F(K_t, 1) - c_t$$

$$K_0 \text{ given}$$

- The planner's only constraint is the economy-wide resource constraint (e.g. the law of motion for capital). We can solve this by setting up a Lagrangian, as before.

- (Practice problem) Can show that the first-order condition is exactly the same as the household's problem - leads to the same difference equations governing the model as the competitive equilibrium.

- Therefore, the competitive equilibrium and the planner's equilibrium coincide. The competitive equilibrium, in this environment, is optimal (are you surprised? why or why not?)

Bonus Slides

Bonus Slides: First Fundamental Theorem of Welfare Economics

- Can we ever improve on the outcomes delivered by private markets? Sometimes - even if the markets are perfectly competitive! But when, and how? What does it mean for an allocation (of goods, productive inputs, etc) to be efficient? Under what conditions is a competitive equilibrium efficient?
- We've now developed theoretical tools that we can use to answer these questions for any economic model.
- With a little more mathematical sophistication, we could extend this framework to a case with many heterogeneous consumers and a notion of *social welfare*.
- This is an important topic in the field of public economics, with substantial intersection with macro.

Bonus Slides: Pareto Efficiency

- First, we need to carefully define what we mean for an outcome (allocation) to be *efficient*. Many definitions exist.

- Strongest, most common, notion of allocative efficiency: Pareto efficiency. In words: an allocation is Pareto efficient if you cannot change the allocation to make everyone (weakly) better off without making at least one person worse off.

- Key virtue of Pareto efficiency: does not rely on interpersonal comparisons ("we both want this good, but Mike wants it more!"). In our models so far, not an issue, we've assumed a single (representative) consumer/household. But if we have two or more distinct (heterogeneous) households, this matters a great deal!

- Key drawback of Pareto efficiency: Very strong, hard-to-satisfy criterion. In the real world, Pareto efficient allocations or Pareto-improving policies are hard to come by. Policies and interventions almost always generate both winners and losers.

Bonus Slides: First Fundamental Theorem of Welfare Economics

First fundamental theorem of welfare economics

If an economy is characterized by:

1. Perfectly competitive markets (individuals, firms are price-takers)
2. Perfectly rational agents who realize their optimal choices
3. Complete markets and information
4. No externalities

Then the competitive (decentralized) equilibrium is Pareto efficient.

- This is a theorem – a rigorously-proven statement that provides sufficient conditions for an equilibrium allocation in a given model to be Pareto efficient.

- This is one of the most powerful and important results in economic theory. It captures the notion of Adam Smith's invisible hand. In this environment, price signals allow markets to allocate goods and productive inputs into the hands of those who value them most. It relies on **very strong** assumptions, **all** of which are suspect in the real world.

Bonus Slides: First Fundamental Theorem of Welfare Economics

- Strictly speaking, the first welfare theorem provides sufficient conditions for the competitive equilibrium to be efficient. But deviations from the (very strong) premises in the welfare theorem often indicate the competitive equilibrium is not Pareto efficient.
- For any given model, we can always solve for either a competitive equilibrium or a hypothetical centralized equilibrium in order to address whether they coincide. In they neoclassical growth model, they do: this model satisfies the premises of the first welfare theorem.
- When does this matter in the real world? All of the premises are suspect – for instance, I can't name any real-world market that is truly perfectly competitive.
- Real-world policy might be aimed at restoring the premises of the theorem, e.g. by correcting externalities, alleviating frictions that prevent agents from realizing their optimal choices (e.g. borrowing constraints), making markets more competitive.