Economics 2450A: Public Economics and Fiscal Policy I

Section 11: New Dynamic Public Finance

Michael Droste

Fall 2022

Outline

1. Final Exam Logistics
2. New Dynamic Public Finance (Dynamic Mirrlees)
   - Motivation
   - Setup
   - First-Best Allocation
   - Second-Best Allocation and Feasible Variations
   - Positive Savings Wedge

Final Exam

- This is the last section! The end of the semester snuck up on us.
- The final exam will in-class (75 minutes) on November 30th.
- As Stefanie mentioned in her announcement, the lectures on surveys, New Dynamic Public Finance, and capital taxation (last 3 weeks of class approximately) will not be tested on the final exam. Everything else is testable; see Stefanie's announcement for general advice.
- Make sure to pay attention to the practice problems - you'll learn more reading those. There is also a decent amount of 'fact-based' content; I personally like flash cards, your mileage may vary.

Final Exam Review Session

- I will be holding a review session, date TBA, likely a few days before the final exam.
- The session will be held live in Zoom and recorded so that you can follow offline.
- My current plan is to walk through the review problems that Stefanie and I posted. I'll hang around to take questions as well.

Motivation

- New Dynamic Public Finance (NDPF) is a fancy name for a very modern branch of optimal tax theory that extends the Mirrlees optimal tax model into a dynamic (multi-period) framework.
- Dynamic (multi-period) setup allows us to think about a tax's impact on intertemporal choices (e.g. household consumption-savings). Critical for capital taxation due to savings-investment identity in general equilibrium.
- All of the considerations from Mirrlees - and in particular the mechanism design aspect of solving Mirrlees - remains in NDPF. In addition, the model is simply richer. Households are not just trading off consumption and leisure, but also making intertemporal consumption (savings) decisions.
- So we can use this framework to analyze commodity, income, and capital taxation *all at once*.

Motivation

- The NDPF lecture notes are too complicated us to cover in a single section. And as mentioned earlier, this content will not be covered on the exam.

- With this in mind, I want to focus on the simplest possible model in the spirit of NDPF, which is essentially a two-period Mirrlees model, and focus on the 'big-picture' mechanics and intuition behind the model.

- Many of the assumptions we will make in this simple model are done only for tractability. The framework itself is extremely flexible, can be (and has been) generalized in many dimensions.

Setup and Preferences

- Two-period economy: $t = 0, 1$. Households indexed by $i \in [0, 1]$.
- Preferences: Households derive utility from consumption in both periods and disutility from working.
  Assume for simplicity that all production (and labor supply n) occurs in the first period ($t = 0$). Household's lifetime utility can be written as $U(c_0, c_1, n)$.

Skills

- Like in Mirrlees, households will be heterogeneous with respect to skill (labor productivity). Assume skill s can only take on finite set of values, $s \in \{1, 2, \dots, N_1\}$, with $p(s)$ denoting density of s across individuals.
- Timing: $t = 0$ choices made **before** s is revealed. Consumption decision in $t = 1$ pinned down by lifetime budget constraint, given $t = 0$ choices and s.
- Review question: Mike claims that if our utility function is such that $\lim_{c \to 0} u_c = \infty$, savings in the first period will always be positive, and thus so will c_1. Why?
- Interpretation: realization of s is a 'skill shock' – period 0 decisions made before it's realized.

Production and Savings

- The household produces goods with a linear production technology $y = s \cdot n$.
- Substituting this into our definition of lifetime utility, we can express lifetime utility as:

$$U(c_0, c_1(s), n) = U(c_0, c_1(s), y(s)/s)$$

- Interpretation: Household labor supply choice is equivalent to choosing how many goods are produced. Household is essentially choosing how many goods are produced and how their consumption of these goods is allocated across periods.
- Assume that households can freely save or borrow in the first period, with R denoting the gross return (1 dollar in the bank in $t = 0$ yields R dollars in the bank in $t = 1$).

Resource Constraint

- Aggregate resource constraint is:

$$c_0 + q\sum_s c_1(s)p(s) \leq q\sum_s y(s)p(s) \tag{1}$$

  where $q = 1/R$ (just for convenience).

- This says that aggregate consumption cannot exceed aggregate production - a very sensible constraint for a production economy!

- Note; we haven't explicitly solved the household's UMP or described their budget constraints: this a feature, not a bug! We're going to solve for centralized equilibria by considering what allocation a hypothetical social planner would choose - only needs to respect

First-Best Allocation

- First-best allocation: What allocation would a hypothetical social planner implement if they could directly choose any feasible allocation, knowing everyone's types?
- More precisely, suppose a policymaker is choosing an allocation $\{c_0, c_1, y/s\}$ to maximize social welfare, subject to the constraint that the allocation is feasible (satisfies resource constraint)?
- In the first-best, there are no constraints arising from incomplete information: the policymaker is assumed to directly observe type.
- First-best allocations are always a sensible benchmark, as we'll see.
- Remember from Mirrlees section/lecture: we can regard gov't choosing an allocation as corresponding to an unrestricted nonlinear tax schedule.

Solving for the First-Best Allocation

- First-best allocation can be characterized as the solution to the following social planner's problem:

$$\max_{c_0, c_1(s), y(s)} \sum_s U(c_0, c_1(s), y(s)/s) \cdot p(s) \quad \text{s.t.} \quad c_0 + q \sum_s c_1(s) p(s) \leq q \sum_s y(s) p(s) \quad \forall s$$

- In words: the social planner is choosing the allocation that maximizes (utilitarian) social welfare, subject to the constraint that the allocation satisfies the resource constraint.

- Review question: Why is this problem not written with expectations? Could it be?

Solving for the First-Best Allocation

- Setting up Lagrangian and deriving first-order conditions:

$$c_0 : \sum_s U_{c_0} \cdot p(s) = \lambda \implies \mathbb{E}U_{c_0} = \lambda$$
$$c_1(s) : U_{c_1(s)} = \lambda q$$

- Combining to eliminate λ yields optimality condition:

$$\mathbb{E}\left[U_{c_0}\right] = RU_{c_1(s)} \quad \forall s$$

- This implies full insurance: $c_1(s) = c_1$ (a constant) for all s.

- Review question: What is the intuition for why the first-best allocation is characterized by full insurance across types, i.e. c_1 does not depend on type?

Second-Best Allocation

- Next, consider the second-best allocation: What allocation would the policymaker choose if they can't observe type?

- Just like Mirrlees: imagine a hypothetical "type-reporting game" where households report their type to the policymaker.

- An allocation is incentive-compatible if a type s household would prefer the allocation they receive by truthfully reporting their type compared to any other allocation that they would receive by reporting some non-truthful type $r \neq s$.

- The policymaker would like to maximize a (utilitarian) SWF while respecting the resource constraint and incentive compatibility constraints.

Revelation Principle

- Why do incentive compatibility constraints matter? The revelation principle (RP), a powerful result from mechanism design.

- Here, the RP implies that the second-best allocation (the policymaker's welfare-maximizing choice of allocation when types are not observed) can be attained by a mechanism in which types are truthfully reported.

- This means that to characterize the second-best allocation in this model (when types are not observed by the policymaker), we only need to consider allocations characterized by incentive compatibility constraints.

- In short: RP constrains the set of allocations that can possibly be second-best. Explicitly describing the type-reporting game is often a convenient way to solve for the second-best.

Second-Best Allocation: Solving

- Second-best allocation is the solution to:

$$
\begin{aligned}
&\max_{c_0, c_1(s), y(s)} \mathbb{E}_0\left[U\big(c_0, c_1(s), y(s)/s\big)\right] \\
&\text{subject to} \quad c_0 + q\sum_s c_1(s)p(s) \leq q\sum_s y(s)p(s) \quad \forall s && (RC) \\
&\qquad U(c_0, c_1(s), y(s)/s) \geq U(c_0, c_1(r), y(r)/s) \quad \forall r, s && (IC)
\end{aligned}
$$

- As in Mirrlees: This is a really difficult problem to solve, because there are a ton of IC constraints (one for each pair of types).

- Multiple paths forward (including local IC contraints with single-crossing condition, as in Mirrlees) - one perturbation-based approach in this context is called feasible variations.

Feasible Allocations and Variations

- Very generally, a feasible allocation is an allocation that respects the constraints of the problem at hand.

- In this case, when individual types are private information (not available to the planner), an allocation $\{c_0, c_1(s), y_1(s)\}_s$ is feasible if it satisfies the resource constraint and incentive compatibility constraints.

- An important question in this class of models is whether free saving is feasible. Answering this question introduces us to the idea of feasible variations, which is a useful type of perturbation argument.

Feasible Variations

- Let F denote the set of feasible allocations and consider an arbitrary feasible allocation $\{c_0, c_1(s), y(s)\}_s \in F$.

- WE say that free savings is feasible if for *any* feasible allocation $\{c_0, c_1(s), y(s)\}$, there exists a feasible variation $\{c_0 - \Delta, c_1(s) + \Delta R, y(s)\}$ where $\Delta \in \mathbb{R}$.

- When is free savings feasible (with s private information)? Depends on income effects.

- Generally, with no income effects, free saving is feasible. With income effects (and $R > 1$), saving more leads to negative income effect that pushes labor supply down, free saving not feasible. In that case, we can still construct a feasible variation, but it's not characterized by $(\Delta, \Delta R)$.

Free Saving and Income Effects

- For remainder: let's consider separable preferences of the form:

$$U(c_0, c_1(s), y(s)/s) = u(c_0) + \beta u(c_1(s)) - h(y(s)/s)$$

  where u increasing and concave, h increasing and convex.

- Let $(c_0, c_1(s), y(s))$ denote some feasible allocation and consider variation $(c_0 - \Delta, c_1(s) + \delta(\Delta, s), y(s))$, where $\delta(\Delta, s)$ is constructed to satisfy IC constraints:

$$u(c_0 - \Delta) + \beta u(c_1(s) + \delta(\Delta, s)) = u(c_0) + \beta u(c_1(s)) + A(\Delta) \quad \forall s, \Delta$$

  for some $A(\Delta)$ and such that the variation is resource-neutral:

$$-\Delta + q \sum_s \delta(\Delta, s) p(s) = 0 \quad \forall \Delta$$

Free Saving and Income Effects

- Critical difference relative to free savings variation: now, $\delta(\Delta, s)$ depends on s - isn't simply $\Delta \cdot R$.

- In other words, with income effects, we'll be able to construct a feasible variation for some Δ by constructing $\delta(\Delta, s)$ to satisfy IC and RC – but it won't be the free saving variation, because $\delta(\Delta, s)$ will not simply be $\Delta \cdot R$ and will depend generically on type.

Inverse Euler Equation

- Suppose that an allocation $\{c_0, c_1(s), y(s)\}_s$ is the second-best allocation. Then by construction, a feasible variation cannot be welfare-improving (otherwise, it would be the second-best).

- Second-best allocation characterized by inverse euler equation:

$$\frac{1}{u'(c_0)} = \frac{1}{\beta R} \sum_s \frac{p(s)}{u'(c_1(s))} \implies u'(c_0) = \beta R \left(\mathbb{E}_0 \left[\frac{1}{u'(c_1(s))} \right] \right)^{-1}$$

- Importantly, this implies that household Euler equation characterizing optimal consumption-savings behavior - is violated by Jensen's inequality:

$$u'(c_0) < \beta R \mathbb{E}_0 \Big[u'(c_1(s)) \Big]$$

- What does this imply about the optimality of a tax on savings / capital?

Where are the Taxes?

- We haven't said anything about taxes! This parallels our discussion of Mirrlees a few weeks ago, but it is sort of an important point worth emphasizing again.
- Our planners are choosing allocations, not taxes. The only constraints on the planner are the resource constraint and the information constraints that arise from the problem (i.e. incentive compatibility). We do not restrict the form of taxes: feature, not a bug.
- But it leaves open the question of whether the second best can be attained with a given implementation (form of taxes).
- The key insight here is that when the planner is choosing allocations (i.e. a centralized equilibrium), this implicitly characterizes a tax schedule.
- Golosov, Tsyvinski, and Werning, "New Dynamic Public Finance: A User's Guide" for discussion of this point.