

# 1

## Introduction

The goal of this book is to figure out at least some characteristics of the best possible tax system. This problem is a difficult one even to pose. The amount of tax that a typical citizen pays is a function of many economic variables. A far from exhaustive list includes labor earnings, interest income, dividend income, consumption, and money-holdings (via inflation). The dependence of collected taxes on these variables may be quite complicated. Moreover, taxes depend on asset incomes and asset holdings, and these represent the outcomes of decisions about how much wealth to transfer from one period to another. The problem of designing a good tax system that includes asset income taxes is intrinsically a *dynamic* one.

At the end of the 1990s, most of the research on optimal taxation in multiperiod settings was being done by macroeconomists (as opposed to specialists in public finance). Following an approach pioneered by Chamley (1986), the research made some rather strong assumptions: it restricted taxes to be linear and (generally) assumed all agents are identical. The resulting research program is extremely tractable. Unfortunately, it is also deeply flawed. Its key economic trade-off is that the government would like to make the taxes nonlinear but cannot. This basic tension is

really irrelevant in the actual design of taxes, because governments *can* (and do) use nonlinear taxes.

In response to this conceptual problem, the *new dynamic public finance*<sup>1</sup> (NDPF) thinks about how to design optimal taxes using the fundamentally different approach pioneered by Mirrlees (1971). The NDPF explicitly allows taxes to be nonlinear and allows for heterogeneity among people in the economy. The heterogeneity comes from a rather natural source. People's labor earnings depend on their choices of labor inputs (how hard or how long to work). Increasing the size of this input causes them disutility, but generates more labor income. As Mirrlees (1971) originally did, the NDPF presumes that people differ in their *skills*, that is, in how much labor input they need to generate a given level of labor income. By way of extension to Mirrlees's baseline analysis, the NDPF allows for the possibility that these skills evolve over time stochastically (so that people may gain or lose skills over time in a surprising fashion).<sup>2</sup>

In the NDPF, the government commits itself *ex ante* to a tax schedule that maximizes a (possibly weighted) average of agents' utilities. The only restriction on this

---

<sup>1</sup>I believe that I can take full credit for the name "new dynamic public finance." In July 2004, I gave a plenary talk at a conference in Florence organized by the Society for Economic Dynamics. I knew that I wanted to choose a title for the talk that would generate attendance and also signal that I wanted to discuss an agenda broader than that suggested by my own individual papers. I chose the deliberately catchy "The new dynamic public finance." The title beautifully served its main, presentation-specific, purposes. I'm more surprised that it has continued to stick as a way to describe the literature—but it has. For example, Golosov et al. (2006) refer to the literature in exactly this way.

<sup>2</sup>In an early contribution, Eaton and Rosen (1980) discuss the structure of optimal linear taxes on earnings given this kind of skill risk.

schedule is that taxes can only depend on incomes, and not directly on people's skills. This restriction immediately translates into the main trade-off that the government faces when designing its optimal tax schedule. On the one hand, the benevolent government wants to provide insurance. People can turn out to be high skilled or low skilled at the beginning of their lives or over the course of their lives. The government would like to insure them against this skill risk. This force leads the government to favor high taxes on income. On the other hand, the government would like to motivate the high-skilled people to produce more income than the low-skilled people. This force leads the government to favor low taxes. The government's problem is to figure out how to resolve this tension in various dates and states.

I have made no explicit mention of private information in describing the NDPF. However, the government's inability to condition taxes directly on skills ends up implying that it has to treat agents as being privately informed about their productivities. It follows that the optimal tax problem in the NDPF is isomorphic to a dynamic contracting problem between a risk-neutral principal and a risk-averse agent who is privately informed about productivities. There is a large literature on such dynamic principal-agent problems (including work by Rogerson (1985), Spear and Srivastava (1987), Green (1987), and Atkeson and Lucas (1995)), and the NDPF exploits its technical insights in many ways.

In the remainder of this introduction, I discuss the scope of the book. I lay out four main lessons of the new dynamic public finance. Finally, I describe the structure of the book.

## 1.1 Scope

This book is *normative*. It is interesting and important to figure out why we have the taxes that we have, but this book does not seek to answer that question. Instead, it tries to figure out what taxes we *should* have. It follows that the actual specification of taxes is irrelevant for the purposes of this book, except to indicate the range of taxation possibilities available to the government. Here's an analogy that might be helpful. The existence of agricultural subsidies and tariffs means that the government has the ability to levy these taxes. But the existence of these taxes does *not* mean that economists are wrong to recommend their elimination. In the same vein, if taxes recommended by the NDPF differ from the taxes that are actually used, there is no logical reason to conclude that there is something wrong with the NDPF.

This argument does not imply that normative economics in general or the NDPF in particular is disconnected from reality. The ultimate goal of the NDPF is to provide relatively precise recommendations as to what taxes should be. These recommendations will depend on a host of model parameters, and we will need to use data to obtain these parameters. As yet, the NDPF has not made much progress in obtaining good measures of the necessary inputs. This book reflects this weakness, but in chapter 7 I provide some ideas about how more progress can be made.

The normative focus means that I am not going to discuss two recent and related literatures. One such literature is on time-consistency. (More technically, it focuses on the structure of sequential equilibrium taxes when governments choose those taxes periodically.) The other literature is on dynamic political

economy. (It focuses on the structure of sequential equilibrium outcomes when taxes are determined by periodic voting.) These literatures examine the properties of equilibrium outcomes of particular dynamic games. Hence, they are trying to model the *actual* behavior of governments. They are not normative in nature and so lie outside the scope of this book.

## 1.2 Lessons

As the remainder of this book shows, we have learned a great deal in a short time from the NDPF. However, I think that there are four particularly important lessons that are worth emphasizing. The first three require preferences to exhibit separability between consumption and leisure. The last does not.

### 1.2.1 Lesson 1: Optimality of Asset Income Taxes

The first lesson concerns the design of optimal asset income taxes.<sup>3</sup> It is valid regardless of the data-generation process for skills. Consider a risk-averse person at date  $t$  who faces skill risk at date  $(t + 1)$ . Under an optimal tax system, the person's shadow interest rate from period  $t$  to period  $(t + 1)$  must be less

---

<sup>3</sup>What does “optimal” mean? Throughout this book, the weights in the government's objective are restricted to be independent of agents' skill realizations. Technically, this restriction means that we are focusing on properties of *ex ante* Pareto optimal tax systems, in which the government's redistributive motives are based on attributes other than skills themselves.

However, the four lessons described in this section are also valid if the weights are allowed to depend on agents' skills at birth (but not on future skill realizations). Technically, we can say that the results also characterize interim Pareto optimal tax systems, in which the government may redistribute based on agents' initial skills.

than the market interest rate. This result immediately implies that an optimal tax system must confront such a person with a nonzero asset income tax that deters him from saving.

Intuitively, when preferences are separable between consumption and leisure, leisure is a normal good. Normality of leisure means that agents with a large amount of accumulated wealth in period  $(t + 1)$  are harder to motivate in that period. Hence, on the margin, good tax systems deter wealth accumulation from period  $t$  to period  $(t + 1)$  to provide people with better incentives to work in the latter period.

This result was originally derived by Diamond and Mirrlees (1978) in the context of a model of endogenous retirement. However, Diamond and Mirrlees restricted attention to a specific data-generation process for skills (a two-point Markov chain with an absorbing state). The contribution of the NDPF (and specifically of Golosov et al. (2003)) is to show that Diamond and Mirrlees's finding applies to *all* data-generation processes for skills, and can in fact be extended to models in which skills are endogenous.<sup>4</sup>

### 1.2.2 Lesson 2: An Optimal Asset Income Tax System

The first lesson implies that any optimal tax system features nonzero asset income taxes. The second lesson is about the structure of these nonzero asset income taxes, and is best divided into two parts. The first part is that in many settings, the optimal tax on a person's asset income in period  $(t + 1)$  must be a non-trivial function of his labor income in period  $(t + 1)$ .

---

<sup>4</sup>The proof of the result is based on the elegant approach taken by Rogerson (1985) in a two-period context.

People's decisions about asset holdings in period  $t$  depend on their labor input plans in period  $(t + 1)$ , and optimal asset income taxes must take this intertemporal connection into account. (This conclusion was originally reached in work by Albanesi and Sleet (2006) and Golosov and Tsyvinski (2006).)

The second part of this lesson is that there is an optimal tax system in which taxes are linear functions of asset income in every period. In this system, given the information available at period  $t$ , period  $(t + 1)$  asset income taxes are negative for people with surprisingly high labor income in period  $(t + 1)$  and positive for people with surprisingly low labor income. The cross-sectional average asset income tax rate, and total asset income tax revenue, is always zero regardless of the aggregate state of the world. Thus, the tax system deters investment not through the level of asset income taxes, but through the positive covariance of these taxes with skill realizations. (This conclusion was originally reached in work by Kocherlakota (2005).)

### **1.2.3 Lesson 3: Optimal Bequest Taxes and Intergenerational Transmission**

Some of the most exciting work in the NDPF concerns the optimal taxation of bequests (see, in particular, Phelan 2006; Farhi and Werning 2007). There are two main results. The first has nothing to do with incentives: even if parents are altruistic, in most Pareto optimal tax systems, optimal bequest taxes are negative. The intuition is simple. In any Pareto optimum in which society puts positive weight on all people, society cares about a child in two ways: through its ancestors and directly as a person. It follows that society always puts more weight on a given child than

its ancestors do, and so society wants to subsidize parent-child transfers.

The second result is a characterization of a particular optimal bequest tax system and *is* connected to incentives and insurance. If parents are altruistic, it is optimal for a person's after-tax outcomes to depend on his/her parents' labor earnings. This dependence is a good way to motivate parents to work hard. On the other hand, society does want to insure children somewhat against their parents' outcomes. As a result, it is optimal to subsidize bequests at a higher rate for poor parents than for rich parents.

#### **1.2.4 Lesson 4: Individual Ricardian Equivalence and Social Security**

In the NDPF, a person's labor income taxes at a given date are allowed to be a function of one's full history of labor earnings. This kind of generality mimics the flexibility that governments actually enjoy. For example, in the United States, social security transfers are a function of the full history of one's labor earnings.

The fourth lesson is that, with this degree of flexibility, optimality considerations only pin down the present value of labor income taxes as a function of a person's labor earnings. Thus, if a person owes \$10,000 in taxes at age 25, the government could collect half of that at age 60 (with appropriate interest charges) without affecting individual decisions at all. This indeterminacy is essentially an individual-level version of Ricardian equivalence. (See Bassetto and Kocherlakota (2004) for a discussion.)

The government can exploit this indeterminacy to simplify the structure of labor income taxes. In particular, there is an optimal tax system in which the



government imposes a flat tax on labor earnings while people are working, and then bases post-retirement social security transfers on the full history of labor earnings. Intuitively, all that matters for incentives and insurance is the dependence of the *present value* of labor income taxes on the history of labor incomes. Any required dependence can be fully encoded into the structure of post-retirement transfers, as long as agents can borrow against these transfers. (This argument is explained more fully in Grochulski and Kocherlakota (2008).)

### 1.3 Structure

The remainder of the book is divided into six chapters. The second chapter of the book concerns the Ramsey (that is, linear tax) approach to dynamic optimal taxation. The chapter derives the classic Chamley (1986) result concerning long-run capital income taxes. The chapter also contains a discussion of the limitations of the Ramsey approach and motivates the alternative Mirrleesian approach that informs the rest of the book.

As discussed above, the NDPF is closely linked to the problem of optimal resource allocation in dynamic economies with private information. Chapter 3 provides an analysis of such problems, including a discussion of the “reciprocal” Euler equation and the long-run properties of optimal allocations. Relative to other treatments, its novelty is that it allows for general specifications of data-generation processes for individual skills. This generality rules out the recursive approaches used by, among others, Atkeson and Lucas (1992). Instead, I employ classical perturbation methods similar to Rogerson (1985). These methods are both more general and (I believe) more intuitive.

In chapter 4, I develop the implications of the NDPF for macroeconomists. I set up a canonical optimal nonlinear taxation problem in a dynamic economy with heterogeneous agents. I show how, in terms of quantities, the solution to this problem is the same as the solution to the private information allocation problem in chapter 3. I use this connection to derive general properties of optimal taxes, and discuss the properties of a particular optimal tax system.

Chapter 5 extends the analysis to bequest taxes. Mathematically, the chapter is similar to the previous one. However, the results differ in important ways, because the societal objective puts more weight on descendants than parents do. This difference affects both the sign of bequest taxes and their dependence on the income levels of parents.

The analysis in these chapters 2–5 is entirely qualitative.<sup>5</sup> In chapter 6, I set forth recursive methods that in principle allow one to find approximate solutions to the basic nonlinear taxation problem when skills follow a Markov chain. This literature is an old one (dating back at least twenty years), but progress has been slow: much remains to be done. I then solve for optimal taxes in a simple numerical example. The example is purely illustrative, but it is nonetheless suggestive.

In chapter 7, I discuss possible paths for future research. This chapter is probably the most important but it is also necessarily the most speculative.

I should add a final warning about notation. In terms of their economic lessons, the various chapters are certainly cumulative. However, the chapters use rather distinct models to derive these lessons. For

---

<sup>5</sup>The material in chapters 3–5 extends and builds upon my survey of similar material in Kocherlakota (2006).

this reason, I have made no attempt to ensure that the notation is consistent *across* chapters, although it is consistent *within* chapters.

## References

- Albanesi, S., and C. Sleet. 2006. Dynamic optimal taxation with private information. *Review of Economic Studies* 73:1–30.
- Atkeson, A., and R. E. Lucas, Jr. 1992. On efficient distribution with private information. *Review of Economic Studies* 59:427–53.
- . 1995. Efficiency and equality in a simple model of efficient unemployment insurance. *Journal of Economic Theory* 66:64–88.
- Bassetto, M., and N. Kocherlakota. 2004. On the irrelevance of government debt when taxes are distortionary. *Journal of Monetary Economics* 51:299–304.
- Chamley, C. 1986. Optimal taxation of capital income in general equilibrium with infinite lives. *Econometrica* 54:607–22.
- Diamond, P., and J. A. Mirrlees. 1978. A model of social insurance with variable retirement. *Journal of Public Economics* 10:295–336.
- Eaton, J., and H. Rosen. 1980. Taxation, human capital, and uncertainty. *American Economic Review* 70:705–15.
- Farhi, E., and I. Werning. 2007. Inequality and social discounting. *Journal of Political Economy* 115:365–402.
- Golosov, M., and A. Tsyvinski. 2006. Designing optimal disability insurance: a case for asset testing. *Journal of Political Economy* 114:257–79.
- Golosov, M., N. Kocherlakota, and A. Tsyvinski. 2003. Optimal indirect and capital taxation. *Review of Economic Studies* 70:569–87.

- Golosov, M., A. Tsyvinski, and I. Werning. 2006. New dynamic public finance: a user's guide. In *NBER Macroeconomics Annual 2006*, Volume 21 (ed. D. Acemoglu, K. Rogoff, and M. Woodford).
- Green, E. 1987. Lending and the smoothing of uninsurable income. In *Contractual Arrangements for Intertemporal Trade* (ed. E. C. Prescott and N. Wallace). University of Minnesota Press.
- Grochulski, B., and N. Kocherlakota. 2008. Nonseparable preferences and optimal social security systems. Working Paper, University of Minnesota.
- Kocherlakota, N. 2005. Zero expected wealth taxes: a Mirrlees approach to dynamic optimal taxation. *Econometrica* 73:1587-621.
- . 2006. Advances in dynamic optimal taxation. In *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Volume I (ed. R. Blundell, W. K. Newey, and T. Persson). Cambridge University Press.
- Mirrlees, J. 1971. An exploration in the theory of optimum income taxation. *Review of Economic Studies* 38:175-208.
- Rogerson, W. 1985. Repeated moral hazard. *Econometrica* 53:69-76.
- Spear, S., and S. Srivastava. 1987. On repeated moral hazard with discounting. *Review of Economic Studies* 54: 599-617.