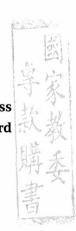




The New Dynamic Public Finance



Princeton University Press Princeton and Oxford



Copyright © 2010 by Princeton University Press Published by Princeton University Press, 41 William Street, Princeton, New Jersey 08540 In the United Kingdom: Princeton University Press, 6 Oxford Street, Woodstock, Oxfordshire OX20 1TW All Rights Reserved

Library of Congress Cataloging-in-Publication Data

Kocherlakota, Narayana Rao, 1963– The new dynamic public finance / Narayana R. Kocherlakota. p. cm. – (Toulouse lectures in economics) Includes bibliographical references and index.

ISBN 978-0-691-13915-9 (alk. paper)

- 1. Fiscal policy-Mathematical models.
- 2. Taxation-Mathematical models.
- 3. Finance. Public-Mathematical models. I. Title.

HJ192.5.K63 2010 336.001'5195-dc22

2009052297

British Library Cataloging-in-Publication Data is available

This book has been composed in LucidaBright using T_EX Typeset and copyedited by T_eX Productions Ltd, London Printed on acid-free paper. \odot press.princeton.edu

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Preface

This book is based on the Sixth Toulouse Lectures in Economics that I gave at the Toulouse School of Economics on March 17–19, 2008. Being invited to give the lectures, and presenting them, were highlights of my professional career. I thank the faculty at the Toulouse School for their thoughtful comments and lavish hospitality.

The resulting book is designed to be used in an advanced Ph.D. course (second year or beyond) in either macroeconomics or public finance. For example, I covered the material in chapters 3-6 in a special nine-hour summer course at the University of Oslo in June 2008. (The students on the course were an outstanding group of European Ph.D. students and assistant professors. All had completed at least two years Ph.D. course work in macroeconomics and/or public finance.) For those who are so inclined, it would certainly be possible to teach chapters 3, 6, and 7 in one self-contained unit. These chapters have little if any treatment of taxes, and focus instead on properties of Pareto optimal allocations in dynamic settings with private information.

I owe thanks to a lot of people for their help. I have presented this material at many seminars and conferences over the years, and I thank the attendees for their comments and feedback. Some of the best young talent in our profession has worked on the new dynamic public finance during the course of this decade,

x Preface

and my own work owes a large debt to my interactions with them. In this vein, I've learned a great deal from Stefania Albanesi, Emmanuel Farhi, Marek Kapicka, Chris Sleet, Sevin Yeltekin, and Ivan Werning. I am especially indebted to my coauthors, Marco Bassetto, Mikhail Golosov, Borys Grochulski, and Aleh Tsyvinski. I have been talking to my colleague Chris Phelan about dynamic resource allocation and private information since we were in grad school together over twenty years ago. Those conversations have been fundamental to my understanding of the new dynamic public finance—and, indeed, to my understanding of all of economics.

Kenichi Fukushima, Futoshi Narita, Machiko Narita, Ctirad Slavik, and Hakki Yazici gave close readings of the contents and helped correct many lingering errors (although probably not all!). Two anonymous reviewers also gave useful comments. Richard Baggaley and Kimberley Johnson of Princeton University Press were very helpful during the entire process. Jon Wainwright of T&T Productions Ltd provided excellent copy editing.

Finally, my wife Barbara McCutcheon has provided emotional support and a continual flow of economic insight throughout the writing of this book. I dedicate it to her.

Contents

Pro	eface		ix
1	Intro	duction	1
	1.1	Scope	4
	1.2	Lessons	5
	1.3	Structure	9
	Refe	rences	11
2	The	Ramsey Approach and Its Problems	13
	2.1	A Simple Model of Government Finance	14
	2.2	The Dynamic Ramsey Taxation Problem:	
		Setup	16
	2.3	The Dynamic Ramsey Taxation Problem:	
		Solution	19
	2.4	Problems with the Ramsey Approach	24
	2.5	Summary	28
	2.6	Technical Notes	28
	Refe	rences	33
3	Basi	cs of Dynamic Social Contracting	34
	3.1	Class of Environments	35
	3.2	Incentive-Compatibility	39

vi	Contents

vi		Con	itents	
	3.3	Remarks on Incentive-Compatibility	42	
	3.4	Pareto Optimal Allocations	45	
	3.5	The Reciprocal Euler Equation	47	
	3.6	Dynamics of Pareto Optimal Consumption	58	
	3.7	Long-Run Properties of Pareto Optima	61	
	3.8	Summary	69	
	3.9	Technical Notes	71	
	Refer	rences	79	
4	Dynamic Optimal Taxation: Lessons for			
		oeconomists	81	
	4.1	A Nonlinear Tax Problem	82	
	4.2	Any Tax Equilibrium Is		
		Incentive-Compatible	87	
	4.3	Building an Optimal Tax System	89	
	4.4	Properties of the Optimal Tax System	112	
	4.5	Remarks about the Optimal System	125	
	4.6	Summary	130	
	4.7	Technical Notes	131	
	Refe	rences	137	
5	Opti	mal Intergenerational Taxation	140	
	5.1	An Intergenerational Tax Problem	141	
	5.2	Another Reciprocal Euler Equation	150	
	5.3	Properties of Socially Optimal Allocations	154	
	5.4	Optimal Bequest Taxation	159	
	5.5	Summary	164	
	5.6	Technical Notes	165	
	Refe	rences	167	
6	Quantitative Analysis: Methods and Results		168	
	6.1	Immortal Agents in an Open Economy	169	
	6.2	A Closed Overlapping Generations		
		Economy	184	
	6.3	Summary	198	
	References		199	

Contents			vii
7	The Way Forward		201
	7.1	More Widely Applicable Tax Systems	201
	7.2	New Solution Methods	204
	7.3	Inputs from the Data	208
	7.4	Summary	210
	7.5	Technical Notes	211
	References		213
Inc	Index		

$oldsymbol{1}$ 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 vublic finance1 (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).²

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

¹I 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.

²In 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

1.2. Lessons 5

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.³ It is valid regardless of the datageneration 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

³What 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 redistributional 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.⁴

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 nontrivial function of his labor income in period (t+1).

 $^{^4}$ The proof of the result is based on the elegant approach taken by Rogerson (1985) in a two-period context.

1.2. Lessons 7

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 persons'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 1.3. Structure 9

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 longrun 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.