

IFS

The Base for Direct Taxation: Commentary

Robert E. Hall

Prepared for the Report of a Commission on
Reforming the Tax System for the 21st Century,
Chaired by Sir James Mirrlees

www.ifs.org.uk/mirrleesreview

The Institute for Fiscal Studies

Full Report to be published by
Oxford University Press



Commentary on “The Base for Direct Taxation” by James Banks and Peter Diamond

Robert E. Hall

Stanford University

March 2008, revised August 2008, September 2008

James Banks and Peter Diamond have thought long and deep about the complex issues of designing a tax system. The chapter spans a huge area of modern research. Its main point is to question the view that a tax on consumption at a single rate is the optimal way for the government to raise revenue.

The authors spurn the framework implicit in the title of the chapter in two important ways. First, they reject, quite explicitly, the notion that the relevant question is the definition of a tax base. They frame much of their discussion in terms of the tax rate on capital income, a question whose answer is a real number (contingent on potentially a huge number of variables), not the question of the inclusion or exclusion of capital income from the tax base, a question whose answer is a binary yes or no. In this respect the authors follow the Meade report, which took the same sensible approach to tax design.

The second departure from the framework suggested by the chapter’s title is that the chapter encompasses indirect taxes as well as direct taxes imposed on the incomes or earnings of individuals. This chapter reaches conclusions about a comprehensive tax system, not just the piece labeled direct. Banks and Diamond hardly comment on this departure from the framework of the report. They have my wholehearted support—I do not believe that the distinction between direct and indirect taxation is meaningful or that

one should design a direct tax without coordinating with the design of the indirect taxes that are part of the overall tax system. Later I will describe my own approach to tax design, which defies classification into direct or indirect.

The chapter considers an ‘Arrow-Debreu’ setup where decisions about behaviour under all possible future contingencies are made prior to the onset of economic activity. The authors are clear that the issue is the variation in tax rates across types of consumption, which may vary by taxpayer characteristics, type of good, time of consumption, and state of the world. Taxation of capital income is a way to achieve higher tax rates for forward consumption. Although it’s only a matter of taste, it strikes me that tax rates on different times or types of consumption is the more satisfactory framework for discussing the issues of the chapter. Taxing capital income is a tool for implementing a tax system. Positive tax rates on capital income tax forward consumption at higher rates than current consumption. A related point is that a tax with a uniform positive rate on earnings and a zero rate on capital income achieves equality of implied rates on current and future consumption. The chapter derives the relation between the earnings tax rate and the equivalent consumption tax rate. Consumption-tax systems don’t necessarily measure and tax each taxpayer’s consumption of goods and services. The value-added tax is the leading practical example; a sales tax is another. Toward the end of this comment I will discuss my work with Alvin Rabushka on a simple, progressive way to implement a consumption tax without needing to keep track of each person’s consumption. The basic idea is to start with a value-added tax and make it progressive through a rebate built into the tax that workers pay.

Many generations of economists have come down in favour of consumption taxation. The traditional line of argument reached that conclusion from the observation that consumption is the best measure of economic well-being. In the days when tax theory was framed in terms of, first, a measure of taxable capacity and, second, equitable distribution of the burden of government across taxpayers with different capacities, a progressive consumption tax seemed to be the answer. After James Mirrlees created the modern theory of mechanism design and applied it to taxation, rationalization of tax proposals took a more sophisticated form. The chapter is firm in rejecting the traditional line of argument in favour of the modern paradigm: Define a class of permissible tax functions, determine information limitations and preferences, and choose the tax function that maximizes social welfare within the permissible class.

Two lines of thought within the modern paradigm, thoroughly discussed in the chapter, rationalize the taxation of current and future consumption at the same rate. First is the Atkinson-Stiglitz theorem. Under special but not totally unreasonable assumptions about preferences, as the chapter explains, the optimal tax when the government observes earnings but not work effort is a tax on earnings alone, with zero tax on income from savings. As I noted earlier, a tax on earnings implies equal tax rates on current consumption and future consumption. Thus another interpretation of Atkinson-Stiglitz is that a tax imposed on consumption should have equal current and future rates, under the assumptions of their theorem.

The second line supporting zero capital-income taxation is Chamley and Judd's observation that any positive tax on capital income compounds into a high implicit tax rate on forward consumption. The chapter's Table 1 illustrates the resulting distortionary

wedges. The distortion becomes overwhelming in an economy with integrated dynastic decision making, but the authors are skeptical on the relevance of that case (I've never been able to make up my mind on this point). The chapter emphasizes the incompleteness of economists' understanding of economic relations between parents and offspring.

The chapter makes it clear that taking consumption as the base for taxation is not a full statement of an optimal tax system. If the consumption tax is progressive and consumption rises over the life cycle, forward consumption incurs a higher tax and the tax system implies an implicit tax on capital income. Only a linear consumption tax with the same marginal tax rate under all circumstances generates the flat profile of forward rates that is the hallmark of zero tax on capital income. The chapter argues that such a flat profile should not be a focal point of tax design. The pure consumption tax is not the ideal. Rather, the many factors considered in the chapter imply that tax design is a much more complex task, involving many research topics as yet unresolved and in some cases as yet unexplored. Economists should not blindly advocate any tax reform that moves the system toward the flat profile of taxation of forward consumption, in the authors' view. At this stage, they argue, there is reasonable evidence supporting ideas that imply that a rising forward profile of consumption tax rates is probably preferred to the flat profile. Some implicit or explicit taxation of capital income is likely to be part of the optimal tax.

The bias in the chapter toward taxing capital income would make an outsider guess either that Britain had zero tax rates on capital income today or that there was a dangerous probability that a reform was likely that would result in zero rates. In fact, the chapter on the British tax system suggests that the tax rate on capital income is fairly

high. My reading of the Banks-Diamond chapter does not convince me that lowering of the British tax on capital income would be an obvious mistake.

The chapter gives the impression that the authors lean in the direction of convincing the reader that taxation of capital income is a better idea than economists generally think. I'm not sure its coverage of ideas pointing in the opposite direction is complete. It omits discussion of Judd's (2002) work on capital taxation with market power, where the optimal rate on capital income is negative.

I have to admit that I'm in favour of work like the chapter of Brewer, Saez, and Shephard, which takes modern theory and modern empirical results, finds the optimal policy design, and then makes a practical reform proposal based on that design. Banks and Diamond take a rather different approach, encompassing a wider variety of issues but not winding up in a policy proposal. How different this chapter would be if the authors had taken the approach of William Nordhaus's (2008) new book on global warming economics. Nordhaus tackles a problem of similar complexity, where research is inconclusive and incomplete, makes assumptions on every point needed, and reaches a definite conclusion about optimal policy (I'm not sure I agree with his conclusion, but I certainly admire the approach). A similar attack on optimal taxation would have wound up with a definite proposal for an integrated income tax system with rates differentiated by age, earnings, and capital income. The Banks-Diamond chapter leaves the reader yearning for some informed guesses about the optimal age profiles, earnings tax schedules, and capital income tax schedules.

There is an interesting point of contact—and divergence—between this chapter and Crawford, Keen, and Smith's chapter on the indirect taxes. These authors discuss a

paper of theirs that asks the completely parallel question about equality of tax rates across commodities. The paper estimates a demand system and compares the optimal, differentiated tax rates, to equal rates and concludes that little is lost by equality. One could imagine taking exactly the same approach to equality of rates over time, given that Arrow and Debreu taught us that time is just another way of indexing commodities. Maybe the intertemporal version of Crawford-Keen-Smith would conclude that there is no advantage to differentiating rates by time or age. Or maybe, contrary to my belief, it would conclude that there is no real disadvantage to the steeply inclined profile of forward consumption tax rates implied by the heavy taxation of some forms of capital income in the UK.

By contrast, the Banks-Diamond approach would challenge equality of rates across commodities in the VAT both for the considerations of preferences emphasized by Crawford-Keen-Smith and for more complicated and interesting reasons of the type emphasized in the new public finance based on dynamic mechanism design. In that theory, every commodity conveys information useful to the tax designer trying to make taxes progressive when key information about the taxpayer is hidden.

In my youth, I advocated a consumption tax that went pretty far in removing discrimination in tax rates in every dimension—time, commodity, and state of the world (Hall and Rabushka, 1995). Many people, critics and friends, pointed out that, because the tax exempted lower-consumption families from any taxation, it had two rates, and thus discriminated between consumption at times and in states of the world when a person was in the zero rate bracket relative to the positive rate bracket. I saw that discrimination as an acceptable departure from zero discrimination.

Now that I'm older, I favour greater discrimination, in the form of at least one more bracket to preserve the desirable feature of current taxes that they are paid mainly by the prosperous. The widening of the consumption distribution in Britain, the U.S., and most other high-income countries seems to call for further discrimination. Still, I find persuasive the Chamley-Judd critique of discrimination by time that reaches extreme rates after a few decades—I'm not ready to endorse any systematic taxation of income from saving. I'm in favour of taxing business income—a feature of my tax reform proposals from the start—because parts of that income are earnings and rents.

The structure of taxation in my original proposal with Rabushka still strikes me as appropriate. It stands halfway between a direct tax and an indirect one. The easiest way to understand it is to start with a comprehensive Value Added Tax covering all of GDP with a deduction for investment and exports and no deduction for imports—the VAT setup the European Union mandates. Let the rate be the same for all goods and services. Now make a change that is purely administrative—let the part of the VAT previously paid by the employer on labour compensation be paid by the employee instead. Give the business a deduction for compensation and tax the worker on the compensation at the VAT rate. Apart from enforcement issues, the substance of the tax remains unchanged. It is a uniform tax on all consumption. Finally, give each worker an exemption from the tax up to a designated level of, say, £8,000.

This setup achieves a two-rate consumption tax. Workers whose earnings are below the exemption level face no wedge between consumption and work effort, while those above the exemption level face a wedge defined by the VAT rate. The tax is progressive in the most important way—it removes any burden of taxation from the

lowest earners. Its approach to progressivity is way more satisfactory than the method currently in use in the EU of zero-rating products that account for large shares of low-income budgets and small shares of high-income ones.

Because the modified VAT has two rates, it does not achieve complete flattening of the profile of forward consumption-tax rates. Rather, workers who begin their careers with earnings below the exemption level face rising expected consumption tax rates. As we learn from Banks and Diamond, this rising profile may be desirable.

Today, I am inclined to advocate at least one more bracket in the individual compensation tax. Instead of the single positive rate of 19 percent that Rabushka and I proposed in 1981, rates of say 15 and 30 percent would come closer to matching the distribution of the burden of the personal and corporate income taxes in the United States. This three-rate system would make the profile of forward consumption-tax rates rather steeper.

The recent literature on tax design has explored linkages of tax rates to personal characteristics. The chapter describes this research in detail. The desirability of linking tax rates to age is virtually a consensus in the literature and receives a strong endorsement from Banks and Diamond. In a few more years, we may have a better idea of the age profile of the optimal consumption tax. Height is another matter. The chapter's discussion of this topic is illuminating—in a full mechanism design framework, it's really hard to avoid the conclusion that taxes should depend positively on height, almost certainly adjusted for sex and for heights of parents. For all their enthusiasm for differentiated tax rates, the authors don't make the leap to endorsing taxes based on height. Neither they

nor I have a totally coherent framework for explaining why we oppose taxation of height. We talk about simplicity, but even that, as the authors point out, is an elusive concept.

I've started work on the design of a fiscal system that considers the implications of the growth in health spending that will surely occur during the rest of this century. Serious fiscal involvement in health finance is the rule among all high-income countries and it seems likely to grow in the United States, where the involvement is currently limited to people over age 65 and to the poor. Because a large fraction of health spending occurs in retirement, because an important fraction of the population seems not to save enough for backloaded spending, because unregulated private insurance markets are not viable, and because the public appears to support more even distribution of health care than of wealth, the governments of high-income countries seem fated to taking in a growing share of output as taxes and distributing the proceeds as insurance payoffs to mostly elderly citizens.

My thinking remains that the revenue needed to accomplish the redistribution toward people with expensive disorders and toward those with low command over resources will and should come from an earnings tax or VAT. As a prediction, this seems to be on reasonably firm ground. Financing of retirement and health programs from earnings taxes is the rule around the world. As a prescription based on principles of tax design, it is a statement that the profile of forward consumption taxes has approximately the right slope from the progressivity of the earnings tax and does not need to be boosted with any tax on the return to saving.

One of the ways to evaluate the material in this chapter is to ask why I, the most open-minded, middle-of-the-road member of the economics profession, find the arguments in this chapter favouring a moderate tax on capital income unpersuasive.

Notice first that I am not the victim of taking the flat profile of forward consumption taxes as the focal point of tax design. More than ever, I'm in favour of a progressive consumption tax, which implies quite a bit of upward slope to the profile.

One obvious reason that the chapter is unpersuasive is that it makes no attempt to provide a number. Did the United States move in the right or wrong direction in 2003 when it cut dividend and capital gains rates? Should Britain have a 10-percent or a 60-percent top rate on capital income?

My biggest concern is the Chamley-Judd argument that extreme rates of forward taxation of consumption are surely not the right profile. This argument has its teeth with respect to people who actually hold non-human wealth. It says that we can do a better job of extracting revenue from the wealthy by flattening the forward profile of consumption taxation. The main substantive argument in the chapter against this view is that the wealthy accumulate wealth for its own sake and not just because wealth permits consumption.

I'm hoping that this most interesting chapter will stimulate work that finds the optimal tax design based on empirical research and recognizing the constraints that govern tax design in practice. That's definitely a goal that the authors and I share.

References:

Robert E. Hall and Alvin Rabushka, *The Flat Tax*, 2nd edition, Hoover Institution Press, 1995.

William D. Nordhaus, *A Question of Balance: Economic Modeling of Global Warming*, Yale University Press, 2008 (forthcoming)

Kenneth Judd, "Capital Income Taxation with Imperfect Competition", *American Economic Review Papers and Proceedings* 92, May 2002