

BEYOND THE PRO-CONSUMPTION TAX CONSENSUS

Daniel Shaviro*

In the last two decades, the dominant norm in fundamental tax reform has shifted from income taxation to consumption taxation, among academics no less than policymakers. Few have recognized, however, that the case for a consumption tax overlaps substantially with that for lifetime income averaging, an idea that has drawn considerably less support. In particular, the case for both is strongest if one accepts the permanent income hypothesis, which holds that people's consumption decisions depend on expected lifetime income, not current-period income. Likewise, few have recognized that the grounds for unease about the case for income averaging (as an ideal system, leaving aside administrative concerns) apply equally to the case for consumption taxation.

Within a welfare economics framework, the case for both norms is close to irrefutable if one makes three key assumptions: that markets are complete, that individuals engage in consistent rational choice given their preferences, and that the only relevant information about taxpayer "ability" is that provided by an undifferentiated measure of lifetime earnings. These assumptions often fail to hold. Where they fail, (1) allowing income averaging between periods may be undesirable, (2) the case for a consumption tax becomes less clear-cut, although still compelling on administrative grounds, and (3) as revealed by the "new dynamic public finance" literature in economics, there is a strong rationale for taxing saving, as does an income tax but not a consumption tax.

INTRODUCTION.....	746
I. THE OPTIMAL INCOME TAX FRAMEWORK.....	751
A. Optimal Income Taxation and the Underlying Information Problem.....	751
B. Taxation and Inequality.....	753
C. Taxing Earnings in a Welfarist Framework.....	756
D. Extensions of the Optimal Income Tax Framework.....	758
E. Atkinson-Stiglitz and the Case for a Uniform Commodity Tax.....	759
II. THE CASE FOR (AND AGAINST) INCOME AVERAGING.....	761

* Wayne Perry Professor of Taxation, NYU Law School. For helpful comments on earlier drafts, I am grateful to Alan Auerbach, Lily Batchelder, Jason Furman, Edward Kleinbard, Yoram Margalioth, Alex Raskolnikov, Jeff Strnad, and the participants in sessions at an NYU Law School Tax Policy Colloquium and a Harvard Law School Seminar on Current Research on Taxation. I am also grateful to the D'Agostino-Greenberg Fund for financial support.

A. <i>Annual Versus Lifetime Systems</i>	761
B. <i>The Permanent Income Hypothesis</i>	763
C. <i>The Case for Income Averaging Under the Permanent Income Hypothesis</i>	766
1. <i>Distribution</i>	766
2. <i>Efficiency</i>	767
D. <i>Implementing Income Averaging</i>	767
E. <i>Problems with the Case for Income Averaging</i>	770
1. <i>Incomplete markets</i>	770
a. "Great expectations".....	770
b. Risk.....	772
c. Inability to annuitize.....	773
2. <i>Departures from consistent rational choice</i>	774
a. Hyperbolic discounting and other causes of myopia.....	774
b. Mental accounts.....	775
3. <i>Additional information</i>	776
a. Past earnings.....	777
b. Savings.....	777
c. Age-related wage distribution and labor supply elasticity.....	778
F. <i>Conclusions Regarding Income Averaging</i>	778
III. IMPLICATIONS FOR THE CHOICE BETWEEN INCOME AND CONSUMPTION	
TAXATION.....	780
A. <i>Overlap Between the Cases for Income Averaging and for Consumption Taxation</i>	780
B. <i>The Distributional Case for Consumption Taxation Under the Permanent Income Hypothesis</i>	781
C. <i>The Efficiency Case for Consumption Taxation Under the Permanent Income Hypothesis</i>	783
D. <i>Problems with the Case for Consumption Taxation</i>	783
1. <i>Incomplete markets</i>	783
2. <i>Departures from consistent rational choice</i>	784
3. <i>Additional information</i>	785
E. <i>Conclusions Regarding Consumption Taxation</i>	786
CONCLUSION.....	788

INTRODUCTION

For many decades in United States tax policy debate, fundamental tax reform was identified primarily with adopting a comprehensive income tax base.¹ In the last ten or so years, it has increasingly come to denote instead

1. See, e.g., COMM'N TO REVISE THE TAX STRUCTURE, REFORMING THE FEDERAL TAX STRUCTURE (1973); HENRY C. SIMONS, FEDERAL TAX REFORM (1950); TREASURY DEP'T, REPORT TO THE PRESIDENT: TAX REFORM FOR FAIRNESS, SIMPLICITY, AND ECONOMIC GROWTH (1984). This is not to deny that academics have often advocated consumption-based, rather than income-based, tax reform. See, e.g., NICHOLAS KALDOR, AN EXPENDITURE TAX (1959); William D. Andrews, *A Consumption-Type or Cash Flow Personal Income Tax*, 87 HARV. L. REV. 1113 (1974); David F. Bradford, *The Choice Between Income and Consumption Taxes*, 16 TAX NOTES 715 (1982). However, "[a]lthough scholars and

replacing the income tax with a consumption tax.² This shift has been as unmistakable in the academic literature as in public political debate.³ Its occurrence is important, even though in my view the prospects for fundamental reform are decidedly dim,⁴ because ideas and ideals matter.

In academic circles, the shift reflects an emerging new consensus (widespread if not universal) that an ideal consumption tax is unambiguously superior to an ideal income tax,⁵ taking into account concerns of both efficiency and distribution. This view rejects any tradeoff between the two types of ideal tax bases, such as between greater progressivity and greater efficiency,⁶ or between different kinds of efficiency.⁷ Rather, the ideal consumption tax is viewed as capable of being equally progressive but more efficient, or equally efficient but more progressive, causing it to dominate the ideal income tax in any comparison.⁸

There is something mystifying about this shift in ideals, no matter how

politicians at times have proposed switching to a consumption-based model, it was not until the last five or six years of the century that such proposals received much popular attention or support." John K. McNulty, *Flat Tax, Consumption Tax, Consumption-Type Income Tax Proposals in the United States: A Tax Policy Discussion of Fundamental Tax Reform*, 88 CAL. L. REV. 2097, 2098 (2000).

2. Prominent consumption-based reform proposals of the last twelve years or so include the Nunn-Domenici USA tax, the flat tax, and the national retail sales tax. See Chris Edwards, *Options for Tax Reform*, 106 TAX NOTES 1529 (2005); McNulty, *supra* note 1.

3. See, e.g., EDWARD J. MCCAFFERY, *FAIR NOT FLAT: HOW TO MAKE THE TAX SYSTEM BETTER AND SIMPLER* (2002); Joseph Bankman & David A. Weisbach, *The Superiority of an Ideal Consumption Tax Over an Ideal Income Tax*, 58 STAN. L. REV. 1413 (2006); Mitchell L. Engler, *A Progressive Consumption Tax for Individuals: An Alternative Hybrid Approach*, 54 ALA. L. REV. 1205 (2003); Daniel N. Shaviro, *Replacing the Income Tax with a Progressive Consumption Tax*, 103 TAX NOTES 91 (2004).

4. See, e.g., Daniel Shaviro, *Simplifying Assumptions: How Might the Politics of Consumption Tax Reform Affect (Impair) the End Product* (N.Y. Univ. Sch. of Law Ctr. for Law & Econ., Working Paper No. 06-17, 2006), available at <http://ssrn.com/abstract=896160>.

5. See, e.g., Mitchell L. Engler, *Progressive Consumption Taxes*, 57 HASTINGS L.J. 55, 55 (2005) ("After years of debate, an academic consensus has emerged that favors the consumption tax . . .").

6. See, e.g., Reuven S. Avi-Yonah, *Globalization, Tax Competition, and the Fiscal Crisis of the Welfare State*, 113 HARV. L. REV. 1573 (2000); Deborah A. Geier, *Incremental Versus Fundamental Tax Reform and the Top One Percent*, 56 SMU L. REV. 99 (2003).

7. See, e.g., JANE G. GRAVELLE, *THE ECONOMIC EFFECTS OF TAXING CAPITAL INCOME* (1994), cited in Bankman & Weisbach, *supra* note 3, at 1414 n.2. This argument about an efficiency tradeoff also appears in the first but not the second edition of a leading public finance textbook. Compare JONATHAN GRUBER, *PUBLIC FINANCE AND PUBLIC POLICY* 708 (1st ed. 2005) [hereinafter GRUBER (First Edition)] (stating that a consumption tax, while it avoids distorting savings choices, imposes greater labor supply distortions than an income tax), with JONATHAN GRUBER, *PUBLIC FINANCE AND PUBLIC POLICY* 744-45 (2d ed. 2007) (eliminating any reference to this claim). Unless otherwise indicated, all subsequent references to the Gruber text are to the second edition.

8. This precise argument is made most forcefully in Bankman & Weisbach, *supra* note 3.

intellectually persuasive one finds it. Consider a second possible shift in normative framework, from favoring an annual income or consumption tax to favoring one that is lifetime-based. Under a lifetime-based system, even if tax payments are remitted annually, lifetime rather than merely annual economic results determine how much one must pay. Thus, people with fluctuating incomes do not end up paying more tax than those with stable incomes, as happens under an annual system with graduated marginal rates. A lifetime-based system that still employs annual tax returns can be implemented through income averaging, a system that the economist William Vickrey proposed nearly seventy years ago⁹ and that the U.S. federal income tax featured in a much more limited fashion from 1964 through 1986.¹⁰ Income averaging has received extensive, but on the whole surprisingly unfavorable, attention in the tax policy literature,¹¹ even though the intellectual case for it substantially overlaps with that underlying the new pro-consumption tax consensus. At the level of ideals, leaving aside differences in administrative feasibility, it makes little sense to accept one idea while viewing the other so skeptically.

The reason for the overlap is that the case for a consumption tax, no less than that for income averaging, relies on taking a lifetime, rather than a current-year or snapshot, perspective in evaluating individuals' welfare and in predicting their behavior. A simple example can help to demonstrate this intuitively. Suppose that, in a given year, Xavier and Yolanda each spend their entire \$100,000 salaries on consumption, and thus are treated the same by a consumption tax even though Xavier has a million dollars in the bank while Yolanda has no savings. How can this be justified, when Xavier is obviously so much better-off? The answer, from a consumption tax standpoint, is that Xavier does indeed bear a much higher tax burden, because in the long run he will be taxed just as heavily as if he had spent the million dollars, on top of the \$100,000, this year.¹² This argument, however, requires taking a lifetime rather than a merely annual perspective—and indeed, perhaps, a longer than single-life perspective if Xavier leaves money to his heirs—raising, as we will see, the

9. See WILLIAM VICKREY, *AGENDA FOR PROGRESSIVE TAXATION* 164-97 (1947).

10. The U.S. income averaging rules from 1964, codified at I.R.C. §§ 1301-05 (2000), were repealed by the Tax Reform Act of 1986, Pub. L. No. 99-514, §§ 141(a), 151(a), 100 Stat. 2085.

11. See, e.g., Lily L. Batchelder, *Taxing the Poor: Income Averaging Reconsidered*, 40 HARV. J. ON LEGIS. 395 (2003); Neil H. Buchanan, *The Case Against Income Averaging*, 25 VA. TAX REV. 1151 (2006); Lee Anne Fennell & Kirk J. Stark, *Taxation Over Time*, 59 TAX L. REV. 1 (2005); Richard Schmalbeck, *Income Averaging After Twenty Years: A Failed Experiment in Horizontal Equity*, 1984 DUKE L.J. 509; Jeffrey B. Liebman, *Should Taxes Be Based on Lifetime Income? Vickrey Taxation Revisited* (Dec. 2003) (unpublished manuscript), available at <http://www.ksg.harvard.edu/jeffreyliebman/vickreydec2003.pdf>; David Weisbach, *The Optimal Accounting Period for Taxes* (2004) (unpublished manuscript, on file with the author). The articles by Schmalbeck, Weisbach, and Buchanan are especially skeptical about income averaging, and none of the articles endorse it wholeheartedly.

12. See Shaviro, *supra* note 3, at 103-04.

very same issues as the case for income averaging.

The persistence of the peculiar disjuncture between the income versus consumption tax debate and the income averaging debate reflects under-appreciation both of the affirmative case for income averaging and of its relationship to the case for a consumption tax. In addition, while consumption taxation certainly does not lack critics, its reliance on the same long-term perspective as income averaging remains under-appreciated. However, this reliance on a long-term perspective makes the cases both for income averaging and for consumption taxation depend on the accuracy of the assumptions needed to support the use of such a perspective. In particular, such reliance rests on three critical assumptions, each of which is subject to challenge.

1) *Complete markets*. Markets are complete when they cover every possible commodity and combination thereof.¹³ In illustration, with complete labor markets one can work at one's wage rate for any number of hours between zero and full-time.¹⁴ With complete capital markets, one can hold financial positions that would pay off in every possible state of the world, thus providing effective insurance against any possible contingency.¹⁵ Complete markets are necessary to the full achievement of allocative efficiency in the economy.¹⁶ In the income averaging and income versus consumption tax debates, their absence would mean that people with the same lifetime resources might actually face very different circumstances in each period, suggesting that they should not necessarily pay the same lifetime taxes.

2) *Consistent rational choice*. Under conventional economic assumptions, people have stable preferences that determine the utility they will experience in alternative states of the world and that they consult to make decisions aimed at maximizing expected utility.¹⁷ People therefore are assumed to engage in consistent rational choice, suggesting that they will make the same choice from within a given opportunity set no matter how the choices are presented or framed. Weakening this assumption, by positing that people are myopic, for example, can undermine the cases for income averaging and consumption taxation by suggesting that behavior and well-being may significantly depend on current period resources, rather than simply on the lifetime total.

3) *Within-period information*. A final assumption underlying use of the long-term perspective to support income averaging and consumption taxation relates to achieving the tax system's distributional objectives, rather than to the choice of analytic timeframe as such. Specifically, this assumption holds that,

13. HARVEY S. ROSEN, PUBLIC FINANCE 50 (5th ed. 1999).

14. Edward J. McCaffery, *Slouching Towards Equality: Gender Discrimination, Market Efficiency, and Social Change*, 103 YALE L.J. 595, 619-22 (1993).

15. Jeff Strnad, *Taxing New Financial Products: A Conceptual Framework*, 46 STAN. L. REV. 569, 578 (1994).

16. ROSEN, *supra* note 13, at 49-50.

17. See, e.g., GRUBER, *supra* note 7, at 26; Christine Jolls et al., *A Behavioral Approach to Law and Economics*, 50 STAN. L. REV. 1471, 1476 (1998).

once one has picked the relevant period for evaluating how well-off people are (such as by measuring their income or consumption for the period), information about when within the period the taxpayer acted or benefited does not provide further useful guidance for distribution policy. Thus, in an annual system, people who earn or spend a lot in January typically are treated the same as those who ended the year in the same overall position but followed a different sequence (such as earning or spending more evenly across time, or with back-loading instead of front-loading). With a lifetime perspective, this assumption becomes more controversial. The gap between, say, the ages of 21 and 75 is much bigger than that between January and December of a single year. An individual may change much more during such an extended period, and the tax system may have much more to gain informationally from looking within the period, rather than just at total results for the period as a whole. This undermines the case for income averaging and consumption taxation by suggesting that the particular sequence of the taxpayer's earnings and/or consumption, not just the taxpayer's lifetime income, should affect how she is treated both overall and at different times within her lifespan.

In this Article, I will show that none of the three assumptions that are needed to make an overwhelming case for income averaging and consumption taxation fully holds. Incomplete markets and departures from consistent rational choice can make current-period information about an individual's circumstances relatively more important than consumption taxation and income averaging effectively recognize. The effect, however, is more to muddy the analysis than affirmatively to support income taxation or a predominantly annual system.

Problems with the third assumption, pertaining to within-period information, potentially have stronger implications. As we will see, the relevance of the within-period components of lifetime information can affirmatively support taxing saving, in keeping with an income tax but not a consumption tax. In addition, in some circumstances there may be grounds for imposing higher taxes on people with declining earnings than on those with level earnings, in keeping with an annual but not a lifetime-based system if both have graduated marginal rates.

Collectively, these departures from the three assumptions refute the core conclusion of a recent leading article that "based on current understanding, ideal consumption taxes are superior to ideal income taxes."¹⁸ Reality is simply too messy for overly definite real world conclusions about the relative merits of these two systems to hold outside the contours of stylized and simplified models. Indeed, while adopting a particular tax base ideal may make sense politically—and, in this sense, I will argue that the case for shifting to a consumption tax ideal remains strong—from a pure intellectual standpoint the abstract quest for the better "ideal tax base" is misguided. To tax either income

18. Bankman & Weisbach, *supra* note 3, at 1414.

or consumption is to operate at too great a distance from the imperfectly observable attributes of individuals that we might actually want to use in allocating tax burdens for either ideal to dominate unambiguously.

For income averaging, the current consensus lies in such a different place that the identical counsel of skepticism has different implications. Despite the lack of a convincing case for lifetime income averaging as a general ideal, in various circumstances it seems likely to allocate tax burdens more equitably and efficiently than a purely annual system. Thus, there should be further exploration of how a limited income averaging system might work.

The discussion in the remainder of this Article proceeds as follows. Part I offers important background by describing the “optimal income tax” literature that has emerged in public economics over the last thirty-five years and that offers a systematic framework for thinking about tax and distribution policy. Part II discusses the case for income averaging, which depends on the merits of an important application of the complete markets and consistent rational choice assumptions. This is the permanent income hypothesis of Milton Friedman,¹⁹ under which people’s consumption decisions are based on their expected lifetime incomes, not on how much they earn in a given period. Part II also explores the significance for income averaging of modifying the assumption that distribution policy must rely exclusively on static information about overall lifetime earnings. In doing so, it makes use of an emergent branch of the economics literature, known as new dynamic public finance (NDPF), that is as yet little known to legal scholars. Part III discusses the case underlying the new pro-consumption tax consensus and shows how extensively it relies on the same assumptions as those that support income averaging and thus is subject to similar objections. Part IV offers a brief conclusion.

I. THE OPTIMAL INCOME TAX FRAMEWORK

A. *Optimal Income Taxation and the Underlying Information Problem*

To evaluate the cases for income averaging and consumption taxation, it is important to start at a foundational level. Why would we have an income or consumption tax, let alone a tax with graduated marginal rates that makes the total tax liability of individuals with fluctuating incomes depend on the choice of period? Within welfare economics, which studies the “determinants of well-being, or welfare, in a society,”²⁰ this inquiry has led over the last thirty-five years to the development of what is called the optimal income tax²¹ (OIT)

19. MILTON FRIEDMAN, A THEORY OF THE CONSUMPTION FUNCTION 20-37 (1957).

20. GRUBER, *supra* note 7, at G-11.

21. The use of the term “optimal income tax” has no bearing on the income versus consumption tax debate. As noted below, optimal income tax simulations typically concern a single period in which earnings, income, and consumption are assumed to be the same,

literature.²²

This Part briefly describes the features of this literature that are most pertinent here. To help prepare the way, however, it is worth noting up front (without attempting yet to justify) several key features of the OIT literature that may initially prove surprising or counterintuitive. Tax policy literature that predates the spread of welfare economics typically assumes that one's preferred tax base—be it income, consumption, or something else—is actually the very thing one really wants to tax.²³ OIT looks deeper, and treats these measures as merely evidence of something else that we cannot directly observe.

Going one turtle down,²⁴ the attribute of interest is earning ability, whether or not exercised. If we are concerned, as this literature is, with identifying better-off and worse-off individuals, on the view that there are strong normative grounds for redistribution from the former to the latter, the fact that someone with high earning ability prefers not to exercise it in full carries no implication that this exercise of preference made her worse-off than if she had done so. We have no reason to think, for example, that rational non-workaholics are generally worse-off than their more compulsive brethren. Another way of putting this point would be to say that, all else equal, people with greater opportunities seem likely to be better-off than those with lesser opportunities.

Even ability, however, remains at least a turtle shy of the bottom layer. In welfare economics, the sole criterion for assessing a given policy is its effect on social welfare, which "depends on how [the policy] influences individuals' well-being and on nothing else."²⁵ Thus, ability matters not for its own sake, but as evidence, in turn, of something else—people's total utility and marginal utility, the latter of which describes how much one's total utility would change if additional resources were given to or taken away from one.

This brings us to the core problem. Total and marginal utility cannot be directly observed. Even ability cannot be directly observed, since actual earnings depend on one's level of effort—itself unobservable, even if we could count people's hours of work, since time is only one of the margins at which people can vary it.²⁶ Accordingly, deciding how much tax different individuals

thereby eliminating choice of tax base issues.

22. The foundational article in this literature is J.A. Mirrlees, *An Exploration in the Theory of Optimum Income Taxation*, 38 REV. ECON. STUD. 175 (1971), which subsequently led to Mirrlees' being awarded the Nobel Prize in Economics.

23. See, e.g., Alvin Warren, *Would a Consumption Tax Be Fairer than an Income Tax*, 89 YALE L.J. 1081, 1093 (1980).

24. I refer to the old story of the woman who claimed that the earth rests on the back of a turtle and, when asked what the turtle rests on, answered that it was "turtles all the way down." Roger C. Cramton, *Demystifying Legal Scholarship*, 75 GEO. L.J. 1, 1-2 (1986). I earlier used this story in Daniel Shaviro, *Endowment and Inequality*, in TAX JUSTICE: THE ONGOING DEBATE 123, 124 (Joseph J. Thorndike & Dennis J. Ventry, Jr. eds., 2002).

25. LOUIS KAPLOW & STEVEN SHAVELL, FAIRNESS VERSUS WELFARE 16 (2002).

26. Thus, consider a law professor who could have worked the same number of hours for more pay as a transactional tax lawyer.

should pay (or what transfers they should receive),²⁷ like so much else in modern law and economics, turns into a problem of incomplete and asymmetric information. Social policy decisions ought to be based on attributes that cannot be directly observed, and as to which individuals, relative to policymakers, have private information about themselves.

Accordingly, in choosing a tax base, such as income or consumption, along with a period for applying graduated marginal rates, such as a year or a lifetime, we are blundering around several layers short of where we would really like to be. The optimal choice of tax base and tax period depend, among other factors, on how well the different alternatives would make use of any information bearing on total and marginal utility that actually is available. Moreover, once we start thinking about all of the available information, we are not necessarily limited to making simple binary choices between income and consumption taxation, or between annual and lifetime-based systems.

B. *Taxation and Inequality*

From the standpoint of efficiency, any income or consumption tax, whether with flat or varying rates, is dominated by a lump-sum tax, or one in which each taxpayer's liability is fixed without regard to any decisions that she makes.²⁸ Income and consumption taxes discourage work and market consumption, and an income tax additionally discourages saving.²⁹ A lump-sum tax could take an infinite number of different forms, including (1) a uniform head tax,³⁰ (2) a reverse lottery to assign tax liabilities randomly, and (3) a tax based purely on people's eye color. Despite its virtues from the standpoint of efficiency, such a tax is almost never seriously proposed.³¹

The reason for not having any feasible lump-sum tax is simple. Some people are better-off than others, and it is widely believed that those who are better-off should pay more tax. Thus, many would agree that Bill Gates should pay more than the average reader of this Article, who in turn should pay more than a homeless person. This sometimes is called the criterion of ability to pay, perhaps reflecting assumptions about the effect of material well-being on the disutility of paying.

In modern welfare economics, the notion of being better-off is commonly

27. See DANIEL SHAVIRO, TAXES, SPENDING, AND THE U.S. GOVERNMENT'S MARCH TOWARD BANKRUPTCY 194-214 (2006) (discussing the conceptual interchangeability of taxes and transfers).

28. See GRUBER, *supra* note 7, at G-6.

29. See Shaviro, *supra* note 3, at 97.

30. A uniform head tax simply levies the same tax on all adults who live in a jurisdiction, regardless of their individual income level.

31. *But see* Jeffrey A. Schoenblum, *Tax Fairness or Unfairness? A Consideration of the Philosophical Bases for Unequal Taxation of Individuals*, 12 AM. J. TAX POL'Y 221, 258-71 (1995) (suggesting the possible merits of a lump-sum tax such as a uniform head tax).

interpreted in terms of a budget line, reflecting the maximum combinations of available resources that an individual can acquire. Having a higher budget line suggests being better-off, all else equal, if we make two assumptions. The first is the psychological assumption of non-satiation, i.e., that more of any good is always better than less.³² Thus, if there are only two consumer goods, *A* and *B*, that are perfectly tradable for each other in complete markets, raising one's budget line means that one can get more of either or both without having to give up anything. Under non-satiation, this implies being better-off, again holding all else equal. The second assumption repeats the old maxim that there is no accounting for taste, or more precisely, that merely observing differences in taste does not immediately tell us anything about, say, who is happier or making better choices or cares more about satisfying her preferences. Drawing any such conclusions from observed differences in taste is not ruled out, but it would require supporting evidence.

Against this background, suppose that we assume two goods: (1) market consumption, comprising everything one could buy for cash, and (2) leisure, comprising not just free time but any use of one's time other than to earn as much money as possible. Everyone has twenty-four hours in a day, and without individuating information about people we might start out by assuming that everyone has an equal ability to make enjoyable use of time other than in trying to earn money. Clearly, however, people differ in wage rates, or ability to make money through economic activity. This suggests a very simple framework for evaluating relative wellbeing, in which people's budget lines vary with their wage rates. Figure 1 is an illustration, in which Andrea has a higher wage rate than Brian, and therefore is assumed to be better-off even if, owing to different preferences, she decides to earn less.

Even though Andrea earns less than Brian, we know under non-satiation that she must be better-off, all else equal, because she could have earned the same amount while having more leisure, or alternatively had the same amount of leisure while earning more cash. She is at point *Y*, rather than at a point that is equal to or better than point *X* in both dimensions, simply because she prefers *Y* to any such point.


The conclusion commonly derived from this type of analysis is that, in principle, wage rates—not actual wages—are the key determinant of relative wellbeing that we have in mind when we reject lump-sum taxation on the view that Bill Gates should pay more than we do while a homeless person should pay less. To be sure, if wages are all we can observe, as distinct from wage rate or labor supply, we will reach an erroneous conclusion in the above example, misclassifying Brian as better-off than Andrea and therefore as apparently meriting less favorable treatment.³³ Yet this error should be unusual if we have

32. See GRUBER, *supra* note 7, at 27.

33. Actually charging Andrea more tax than Brian would give rise to the "beachcomber problem" that I discuss in *Endowment and Inequality*, *supra* note 24, at 132,

no reason to think that wage rates generally are inversely correlated with preferences for market consumption. Even when we get correct relative rankings, however, it still remains to motivate responding to differences in wellbeing by imposing higher taxes on those who, by the admittedly imperfect rubric of actual earnings, appear to be better-off.

Figure 1. Individuals with Different Wage Rates and Labor Supply Decisions³⁴



It should already be clear how the analysis reflects the assumption that earnings are the only observable variable that bears on relative wellbeing. The assumptions of complete markets and consistent rational choice are important as well. For example, we might be less sure that Andrea is better-off than Brian if gaps in the labor market prevented her from optimizing her work choice given her preferences.³⁵ Thus, suppose she dislikes full-time work but would like to work twenty hours a week, and is constrained by a lack of part-time jobs

epitomized by the hypothetical case of the idle beachcomber who could have earned millions of dollars by working on Wall Street financial deals. However, the question of whether we should actually tax beachcombers on their unrealized earnings opportunities is distinct from that of whether we should regard them as better off, all else equal, than individuals with worse opportunities.

34. Figure 1 is adapted from Shaviro, *Endowment and Inequality*, *supra* note 24, at 130.

35. See *id.* at 130-31.

that would employ her for more than ten hours. This might reduce her total utility below Brian's and clearly would increase the marginal utility to her of an extra dollar, relative to the case where she was able to work more without having to choose a full-time job. As for consistent rational choice, it tells us that people's labor supply decisions meaningfully reflect their preferences and opportunities, supporting viewing earnings information as generally meaningful even if, as in the case of Andrea and Brian, it occasionally supports mistaken conclusions.

C. Taxing Earnings in a Welfarist Framework

Merely observing evidence of unequal wellbeing does not immediately motivate responding to it with differential tax treatment. For this, one needs a normative framework. Within welfare economics, with its assumption that social welfare depends purely on individuals' wellbeing, there is an easy case for applauding Pareto improvements, where someone gains and no one loses. However, the harder task of assessing tradeoffs between one person's gain and another's loss is unfortunately unavoidable in deciding how to allocate tax burdens, an exercise that inevitably has winners and losers. Within welfare economics, utilitarians assess these tradeoffs by requiring only that the net effect on welfare be positive, i.e., that gains exceed losses if one counts each individual's welfare equally.³⁶ The most common alternatives to utilitarianism give extra weight to the welfare of worse-off individuals—at the limit, by giving infinitely greater weight to the welfare of the worst-off individual than to that of anyone else.³⁷

Utilitarianism motivates redistribution from better-off to worse-off individuals through the assumption of diminishing marginal utility, which holds that unit $n + 1$ of an item yields less utility than unit n .³⁸ If people have identical utility functions characterized by declining marginal utility, then transferring resources from better-off to worse-off individuals will increase social welfare, all else equal.³⁹ Variants of welfare economics that give extra weight to the welfare of worse-off individuals provide an additional motivation for transferring resources to such individuals from those who are better-off.

36. See, e.g., GRUBER, *supra* note 7, at 53-54.

37. See ANTHONY B. ATKINSON & JOSEPH E. STIGLITZ, LECTURES ON PUBLIC ECONOMICS 339-40 (1980); Amartya Sen, *Equality of What?*, in 1 THE TANNER LECTURES ON HUMAN VALUES 195, 206 (Sterling M. McMurrin ed., 1980).

38. GRUBER, *supra* note 7, at 29-30.

39. In the case of Andrea and Brian, their different labor supply choices indicate that they do not have identical utility functions. Andrea evidently has less taste for market consumption, more taste for leisure, and/or greater work aversion. One might nonetheless treat Andrea and Brian, for purposes of the distribution decision, as if they had identical utility functions if this observed difference has no clear-cut implications for the aggregate utility effects of the transfer. See *infra* text accompanying note 63 (discussing the work of economist Abba Lerner).

Within this framework, an earnings tax provides “an insurance mechanism to mitigate undesired risk from people’s involuntary participation in the ability lottery.”⁴⁰ This provision requires departing from the complete markets assumption, since with complete insurance markets such coverage would already be available, and with consistent rational choice people would always choose to hold it if they wanted it. The gap in private markets that motivates government provision is most plausibly attributed to adverse selection, which arises when prospective purchasers of insurance have superior private information about their own prospects.⁴¹ Adverse selection can prevent insurers from being able to offer a given product without losing money, by inducing disproportionate subscription by those who expect a positive payoff given their actual odds.⁴² In the case of ability insurance, given private insurers’ lack of opportunity to offer the coverage to prospective customers from behind the veil (or before these individuals know anything about their own ability levels), adverse selection would involve disproportionate subscription by those who knew they were likely to be low-income and thus to collect on the policies. The government can address adverse selection by requiring everyone to subscribe,⁴³ but the degree of socially desirable coverage remains limited by moral hazard, or unobserved behavioral responses to the incentive effects of the insurance coverage (here, by working less in response to a redistributive tax that is based on earnings).⁴⁴

An extensive OIT literature relies on simulations, involving a hypothetical population and economy, to determine what the tax system should look like in furtherance of the ability insurance function. The main inputs to these simulations are the choice of social welfare function (utilitarian or otherwise), the distribution of earning ability in the population, the specification of people’s utility functions, and the degree of labor supply responsiveness to the tax.⁴⁵ These typically are one-period models in which members of the hypothetical population do not face the question of whether to save any of their earnings for consumption in later periods. With everything being consumed currently, earnings, income, and consumption are all the same. This eliminates any consideration either of the choice between income and consumption taxation or of income averaging, which would require multiple periods. The simulations typically focus on what an optimal tax rate structure would look like, thus helping to inform our judgment about actual tax rate structures if we are sufficiently sanguine about the relationship between the models and the much more complicated real world. The usual conclusions reached are that everyone should receive a minimum grant (regardless of income) and that

40. DANIEL SHAVIRO, MAKING SENSE OF SOCIAL SECURITY REFORM 52 (2000).

41. *Id.* at 46.

42. See GRUBER, *supra* note 7, at 324.

43. *Id.* at 328.

44. SHAVIRO, *supra* note 40, at 55.

45. See GRUBER, *supra* note 7, at 597-600.

marginal rates, as income rises, should be surprisingly non-graduated.⁴⁶

D. Extensions of the Optimal Income Tax Framework

If ability is the main attribute of distributional interest from the standpoint of welfare economics, and earnings are of interest merely as evidence of ability, there is no reason to assume that earnings are the only evidence worth observing. For example, one could in principle also try to make use of information about gender, ethnic origin, or any observable physical characteristics, such as height, determined to be statistically relevant. As yet, few practical suggestions have been made in this area apart from treating disability as evidence of low earning capacity.

One should keep in mind, however, that even ability matters only as evidence bearing on total and marginal utility. An individual in a permanent coma, while having no earning ability, would presumably lack appeal to a utilitarian as a prospective recipient of large transfers. Moreover, ability as a consumer—that is, the capacity to derive utility from resources—matters distinctly from earning ability. Differences in ability as a consumer are usually ignored, however, reflecting the difficulty of observing them along with the analytical advantages of using a simpler framework.

The static one-period approach of classic optimal income tax models likewise serves purely to simplify the analysis, concededly at the cost of accuracy and descriptive richness. Over the last twenty years, the NDPF literature has begun incorporating time, risk, and gradually unfolding information into the Mirrlees framework.⁴⁷ In particular, it emphasizes the “informational and/or enforcement frictions that limit government’s extraction power,” given that a taxpayer’s skill and effort levels may be private information, about which she knows more than anyone else.⁴⁸ Time matters in this setting both because people’s ability levels are subject to ongoing risk, and because the government can gain information about people’s past and expected current ability levels by observing their earnings sequences and their saving. Thus, rather than simply basing tax rates on current period earnings, the government can use information about past earnings and current wealth in attempting to gauge people’s current ability levels, and thus in setting their taxes. While the NDPF literature is highly technical and as yet little known to legal scholars, it has important implications for income averaging and choice of tax base, as we will see in the next two Parts.

46. *See id.*

47. For an overview, see Mikhail Golosov, Aleh Tsyvinski & Ivan Werning, *New Dynamic Public Finance: A User’s Guide*, 2007 NBER MACROECONOMIC ANN. 317 (2006).

48. Narayana Kocherlakota, *The New Dynamic Public Finance*, Slides from 2004 SED Plenary Session III (July 3, 2004), <http://www.econ.umn.edu/~nkocher/nksedtalk.pdf>.

E. Atkinson-Stiglitz and the Case for a Uniform Commodity Tax

Another important result in the broader welfare economics literature starts from the insight that we can analogize an income or consumption tax to a set of taxes on specific consumer goods or commodities, such as gasoline, DVD players, and milk. Imposing a single, comprehensive, flat-rate consumption tax is equivalent to imposing a set of separate commodity taxes, each at the same rate, on all consumer goods. This suggests that issues of tax base design amount to asking whether all commodities should be taxed at the same rate, or whether instead their tax rates should be differentiated.⁴⁹

A common intuition, based on the notion of tax neutrality, holds that all commodities should generally be taxed at the same rate. This view, however, seemingly had been rebutted in prior economics literature, concerning optimal commodity taxation.⁵⁰ That literature found that differential tax rates should indeed be employed, in a commodity tax designed purely to raise revenue at the lowest possible efficiency cost, if commodities differ in price-elasticity—that is, in their supply and demand responsiveness to changes in the tax rate. Thus, suppose that demand for orange juice is highly tax-elastic, in that people will drink a lot less of it if it is even modestly taxed, while demand for milk is highly inelastic. The central insight of optimal commodity taxation is that, in these circumstances, the tax rate on orange juice should be lower than that for milk. This follows from the idea that, in the course of raising the needed revenue, we want to distort people's decisions as little as possible. Under the elasticity assumptions, milk can bear a higher tax rate than orange juice before it starts generating the same level of tax avoidance behavior via shifts to other commodities.

An important 1976 article by Anthony Atkinson and Joseph Stiglitz⁵¹ (hereinafter "AS 1976")⁵² overturned this finding of the optimal commodity tax literature, in favor of the intuition favoring tax neutrality, for a significant set of circumstances. AS 1976 modified the assumed set-up in the optimal commodity tax literature, where the sole aim is to raise revenue at the lowest possible efficiency cost, by adding two assumptions. First, tax system design aims not just to minimize inefficiency, in which case a lump-sum tax would be preferable to any commodity tax, but also to distribute tax burdens between individuals based on ability. Second, a tax on leisure (i.e., non-market consumption) is assumed to be unfeasible even though leisure is, in effect, a

49. For this purpose, I ignore the relatively easy case of a "Pigouvian" tax that takes account of externalities, such as pollution. See GRUBER, *supra* note 7, at 134.

50. See, e.g., *id.* at 586-89.

51. A.B. Atkinson & J.E. Stiglitz, *The Design of Tax Structure: Direct Versus Indirect Taxation*, 6 J. PUB. ECON. 55 (1976). This result requires that all commodities be weakly separable from leisure, i.e., that none be leisure substitutes or leisure complements, the consumption of which was differentially affected by how much leisure one had chosen.

52. I adopt this usage from Bankman & Weisbach, *supra* note 3.

commodity that people choose like all the rest. Thus, all feasible commodity taxes, or taxes on market consumption, fall at least indirectly on labor supply, since labor is the means of generating earnings that can be spent on such consumption. In this setting, AS 1976 showed that, in general, all taxable commodities (i.e., everything other than leisure) should be taxed at the same rate, just as the neutrality concept would suggest.⁵³

The intuition underlying the result in AS 1976 is that the tax elasticity of demand for particular commodities is the wrong thing to focus on if all commodity taxes discourage work, and if underlying revenue needs mean that work is going to be tax-discouraged in one way or another. The fact that consumers will shift from orange juice to some other taxable commodity a lot faster than they will for milk offers no indication that one can mitigate the tax burden on labor supply by differentiating the two tax rates. Instead, differential commodity taxation merely takes the underlying labor supply distortion and layers a further distortion on top, by inducing taxpayers not only to work less, but also to alter their preferred commodity choices due to the tax advantage of shifting to lower-taxed goods. Needlessly adding an extra distortion without mitigating the prior one generally increases total inefficiency, and should therefore be avoided if possible.⁵⁴

In illustration, suppose that orange juice and milk were the only two commodities in the world, and that one could raise the requisite amount of revenue by either (a) taxing both at 30% or (b) raising the milk tax just slightly, to 35%, while cutting the orange juice tax all the way to 10%. Either way, given the equivalent revenue, one would be imposing about the same overall tax burden on labor supply. Option (b), however, would have the added disadvantage of inducing people to switch for tax reasons from drinking milk to drinking orange juice.

As we will see, this analysis has implications both for the choice between income and consumption taxation and for income averaging. These implications reflect that both consumption taxation and income averaging can have the effect of equalizing the marginal tax rates imposed on commodities consumed in different periods.

53. The exception to this general finding in AS 1976 was that tax rates should be higher on commodities that are leisure complements, or items which choosing leisure makes more appealing, and lower for leisure substitutes, or items that are chosen as alternatives to leisure. An example of a leisure complement might be raw food ingredients that one needs free time to turn into a meal. An example of a leisure substitute might be a restaurant meal, which becomes more appealing if one is too busy to cook. Differential tax rates for leisure complements and substitutes serve the purpose of partially undoing the distortion that results from not taxing leisure. While this exception to the general finding in AS 1976 that commodity taxes should be uniform is potentially capacious, to date few convincing illustrations of where it would apply have been offered, rendering the general case apparently of greater interest.

54. See Bankman & Weisbach, *supra* note 3, at 1422-24.

II. THE CASE FOR (AND AGAINST) INCOME AVERAGING

Having reviewed the basic optimal income tax framework for making tax policy judgments from a welfare economics perspective, we can now turn to the first of the two substantive issues that this paper examines, concerning the optimal period for determining tax liabilities. We will see that the theoretical case for a lifetime rather than an annual approach would be straightforward if markets were complete, people engaged in consistent rational choice based on their lifetime budget lines, and aggregate lifetime earnings data captured all of the available information that is relevant to making distributional judgments.

A. *Annual Versus Lifetime Systems*

The mainly annual character of the federal income tax—shared by comparable tax systems around the world, although not by Social Security and similar retirement systems—has two distinct components, apart from the frequency of filing. The first pertains to information, and the second to the timing of cash flows between the taxpayer and the government.

As to information, tax liability largely depends on measuring the taxpayer's position during the current year. For the most part, one's income for the year determines one's tax liability without regard to one's income in other years, leaving aside the carryover of various tax attributes such as net operating losses.⁵⁵ As to the timing of cash flows, taxpayers generally must settle up each year with the Internal Revenue Service. They cannot, at a market interest rate, accelerate or defer the tax payments that are due. Early payment is not rewarded with a time-value discount to keep its present value constant, and late payment (beyond permissible extensions) is penalized, rather than simply leading to the imposition of a market interest charge. Poverty programs likewise rely mainly on current-year (or even current-month) information, with exceptions such as the five-year time limit on receiving Temporary Aid to Needy Families (TANF) benefits.⁵⁶

The restriction to current-year information inevitably has important effects in the presence of non-flat marginal rates (including not only rising rates, but those that decline with income due to phase-outs (such as those of income tax credits or welfare benefits). If non-flat rates apply to annual income, then the sequence of one's earnings and other taxable income, and in particular whether they fluctuate or are relatively constant over time, affects one's lifetime tax burden. Earning fluctuations can significantly increase one's lifetime tax burden if marginal rates rise with income. In illustration, suppose the tax rate

55. However, attributes such as asset basis make possible the use of information from past years.

56. See 42 U.S.C. § 608(a)(7) (2000) (stating that generally no individual may receive welfare assistance for longer than sixty months).

were 20% on the first \$50,000 income and 50% on income above that. Under this rate structure, someone who earned \$100,000 and zero in regularly alternating years would pay \$35,000 of tax every two years, or an average of \$17,500 per year. Someone who earned \$50,000 per year, and thus exactly the same average amount, would pay only \$10,000 of tax each year.

In contrast to the effect on applying marginal rates based on annual rather than lifetime income, the requirement of annual cash flow settlement matters only contingently. In the presence of (1) complete capital markets, permitting people to borrow and lend across time however they like, and (2) rational consumers who make consistent intertemporal choices given their preferences, it would make no difference whatsoever. Having to pay federal income tax at a given time, rather than being allowed to accelerate it or defer it at a market interest rate, would have no effect on people's work or consumption activity if they could (and did) borrow and lend at will to arrange their overall cash flows as they liked. However, where markets are incomplete or people fail to exercise consistent rational choice across time, the sequence of cash flows between oneself and the government can be important.

Historically, the first of these two features of an annual tax system (restriction to current-year information) has been a lot more controversial than the second (annual cash flow settlement). In particular, concern about the former has led to calls for income averaging, under which marginal rates would apply to one's average annual income over a period longer than a year—perhaps even one's entire life. Annual cash settlement, by contrast, has prompted little dissatisfaction. While considering it important may be logically reconcilable with favoring the use of a long-term income picture,⁵⁷ the two are uncomfortable bedfellows. For annual cash settlement to matter, the current sequence of cash flows—deemed inconsequential in the long-term view—must matter.

This ambivalence is prominently on display in the best known and most comprehensive income averaging proposal to date, made by economist William Vickrey in 1939.⁵⁸ The normative criteria on which Vickrey relied included the following:

(1) The discounted value of the series of tax payments made by any taxpayer should be independent of the way in which his income is allocated to the various income years. . . .

(4) Any given tax payment should not be too large in relation to the income of the period immediately preceding.⁵⁹

57. For example, one might favor annual cash settlement on the ground that the government is ill-equipped to handle the default problems that might be raised by letting people borrow their current tax liabilities at a market interest rate.

58. I discuss the Vickrey plan *infra* Part II.D.

59. See WILLIAM VICKREY, *Averaging of Income for Income Tax Purposes*, in PUBLIC ECONOMICS: SELECTED PAPERS OF WILLIAM VICKREY 105, 107 (Richard Arnott et al. eds., 1994).

These two criteria can be (and in Vickrey's scheme are) pursued simultaneously, since the former relates to lifetime tax burden and the latter to the timing of tax payments. Yet Vickrey's Criterion (4) is a poor intellectual fit with his Criterion (1). If only the discounted present value of cash flows matters, why worry about the relationship between a given tax payment and income of the immediately preceding period? And if we do need to worry about that relationship, doesn't this suggest that not all present-value-equivalent cash flows are the same?

The reasons for this ambivalence lie in the explanatory power and limits of economic theory. Vickrey's Criterion (1) would be completely persuasive, and his Criterion (4) irrelevant, if one fully accepted the permanent income hypothesis of Milton Friedman, under which people's consumption decisions are based on their lifetime incomes, not on how much they earn in a given period.⁶⁰ However, given incomplete capital markets and time-inconsistent preferences, the permanent income hypothesis does not fully hold. As I discuss next, its descriptive validity is context-dependent in ways that the government cannot easily observe, and that general tax rules cannot easily reflect.


B. *The Permanent Income Hypothesis*

The Friedman permanent income model is mathematically sophisticated. So is the closely related life-cycle model of consumption smoothing that was pioneered by Franco Modigliani and Richard Brumberg.⁶¹ Yet the core idea behind both models is simple and intuitive, reflecting familiar economic reasoning. To illustrate it in a different setting, suppose we again posit a world with only two consumer goods, here apples and oranges, which are freely tradable for each other at a fixed ratio and at zero transaction cost. Once again, each worker has a wage rate and therefore a budget line, reflecting the largest combinations of the two consumer goods that she can earn. Here, however, I address how one chooses a particular combination of the two. One's choice depends on one's preferences, which are assumed to reflect declining marginal utility for each good, and which can be represented through indifference curves showing combinations of the two goods that one rates as equal, as shown in Figure 2.

60. See FRIEDMAN, *supra* note 19, at 20-21.

61. See Franco Modigliani & Richard Brumberg, *Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data*, in POST KEYNESIAN ECONOMICS 388 (Kenneth K. Kurihara ed., 1954); see also ANGUS DEATON, UNDERSTANDING CONSUMPTION 214 (1992) (describing the permanent income and life cycle hypotheses as "well-defined special cases of the general theory of intertemporal choice"). For convenience, given that the differences between the two models are unimportant for my purposes, I emphasize Friedman's permanent income hypothesis throughout.

Figure 2. Commodity Choice Given One's Preferences



The worker in Figure 2, who can obtain any of the apple-orange combinations on line AB , picks the combination at point C , which lies on the highest indifference curve that is tangent to AB .

Now suppose that, despite the free exchangeability of apples and oranges, some of the firms that might employ this worker pay purely in apples, while others pay purely in oranges. Under standard economic reasoning, the worker would be expected to end up at C no matter which firm ends up employing her. Given the goods' free exchangeability, and adding as well the assumption of consistent rational choice given one's preferences, she will trade her way to the favored spot even if she starts at A or B . Under these assumptions, we also would expect her choice of employer to be unaffected by whether a given firm paid in apples or in oranges.

A further possible implication requires more assumptions. Suppose that, if two workers at the same budget line chose the same commodity mix, we would assume that they had the same total utility and marginal utility for an extra unit of consumption. We might base this assumption on the view that, so far as we can tell, they have identical utility functions. This admittedly is a conclusion based on ignorance rather than knowledge—the real point being, not that we know their utility functions are the same, but rather that we lack any differentiating information. Thus, nothing that we know contradicts the

possibility that one of them might be a “utility monster”⁶² relative to the other, experiencing double the total and marginal utility under all circumstances. Given the identical information, however, picking one of them as the “utility monster” would amount to tossing a coin, with only a 50% chance of having directed the extra resources in the correct rather than the incorrect direction. As economist Abba Lerner showed, if we are entirely ignorant of whose utility is more intense but believe that people generally experience declining marginal utility, any such “coin toss” reduces the sum of the two individuals’ expected utility.⁶³

Now suppose instead that we observe two workers making different choices at the same budget line. This definitively tells us that their utility functions must be different. Still, if this observation conveys no information to us about the two workers’ total or marginal utility, the Lerner analysis continues to apply. We still lack differentiating information that is relevant to our assessment of the two individuals, suggesting that we should continue to treat them as if they were identical, at least until we gain any such information.

This little model has wide-ranging applicability. Make the goods “market consumption” and “leisure,” with tax being imposed only on the former, and we have the basic set-up for optimal income tax analysis (as discussed in Part I), with its implication that the equal taxes imposed on people who are relevantly equal should in principle reflect ability (i.e., the wage rate), rather than labor supply choice. Emphasize the assumption that the worker will head to point *C* no matter where she starts out, and essentially we have the Coase theorem.⁶⁴ Or, of greatest interest here, make the two goods “earlier consumption” and “later consumption,” and we have the permanent income hypothesis.

The permanent income hypothesis holds that people’s current consumption choices depend on their anticipated lifetime incomes, rather than their incomes for any given period. Thus, the sequence of one’s earnings is entirely independent of the sequence of one’s consumption, keeping in mind that labor effort (as distinct from earnings) in a given period may affect how much one wants to consume in that period.

While not logically necessary to the model thus described, proponents note that in practice it is likely to imply smoothing out one’s annual consumption

62. Robert Nozick proposed the notion of a utility monster who, unlike most individuals, is not subject to the principle of diminishing marginal returns as he accrues wealth, but gets “enormously greater gains in utility” from the increase of his resources. ROBERT NOZICK, *ANARCHY, STATE, AND UTOPIA* 41 (1974).

63. ABBA P. LERNER, *THE ECONOMICS OF CONTROL* 29-32 (1944).

64. The Coase Theorem holds that “the initial allocation of legal entitlements does not matter from an efficiency perspective so long as they can be freely exchanged.” Robert D. Cooter, *The Coase Theorem*, in *THE NEW PALGRAVE, ALLOCATION, INFORMATION & MARKETS* 64, 64 (John Eatwell et al. eds., 1987) (emphasis removed). See generally RONALD COASE, *THE FIRM, THE MARKET, AND THE LAW* (1988).

stream relative to one's annual earning stream.⁶⁵ In particular, to consume after retirement, people must save for it. Likewise, good years and bad years from an earnings standpoint need not have any correlation with years when one prefers high consumption as opposed to low consumption, suggesting a possible need for consumption smoothing even during one's working years.

If consumption in any one period has declining marginal utility as the amount of consumption in that period increases, people will have some tendency to prefer complete income smoothing as between periods. There is no reason to predict complete smoothing, however, with realistic utility functions. Various plausible preferences, such as for rare but long vacations, or to travel more while one is still young, or to experience constant material betterment, can lead to one's choosing uneven consumption between periods, or particular patterns such as smoothly rising or falling levels of consumption.

C. *The Case for Income Averaging Under the Permanent Income Hypothesis*

Under the permanent income hypothesis, the case for income averaging is compelling on both distributional and efficiency grounds—ignoring for now the possibility that earnings or consumption sequences, even if rationally chosen in complete markets, might offer pertinent information about ability.

1. *Distribution*

In the permanent income model, people's consumption over time depends purely on their budget lines, lifespans, and intertemporal consumption preferences. Thus, two individuals with different earnings sequences but identical lifespans, budget lines, and consumption choices would have identical total and marginal utility, so far as we know, if the differences in their earnings sequences did not convey information to the contrary.

Under a utilitarian social welfare function, the two individuals' presumed identical marginal utility would support taxing them the same. If one shifted to a social welfare function that gave extra weight to the welfare of worse-off individuals, one still would want to tax them the same, given their presumed identical total utility.

An important detail here concerns lifespan differences. The longer one lives, the greater one's lifetime consumption needs. Thus, while living longer may increase one's total utility, it also increases the marginal utility of an extra dollar. In the absence of complete markets and consistent rational choice, this would suggest that, as between two individuals with the same lifetime income, the one who lives longer should pay less in lifetime net taxes. No government response would be necessary, however, if people fully responded to life expectancy risk by annuitizing—that is, by using their wealth, minus any

65. See Fennell & Stark, *supra* note 11, at 7.

amounts they wanted to leave in bequests, to purchase life annuities that would support them for as long as they lived.

2. *Efficiency*

Whether the tax system is based on income, earnings, or consumption, the use of a lifetime measure improves efficiency, assuming non-flat marginal rates, because it causes the same marginal rate to apply to the taxpayer's choices in all years. Suppose initially that tax liability depends on income or earnings. Using a lifetime rather than an annual measure improves efficiency in two respects. First, with a shorter than lifetime measure, people have an incentive to shift their earnings to years where they face a lower rather than a higher marginal rate. A lifetime measure avoids this. Second, a lifetime measure permits tax rate smoothing. One can raise the same revenue from a given individual as under an annual system by blending what would otherwise be some years' high rates and other years' low rates into a single intermediate set of rates. This would be expected to reduce economic distortion even if earnings cannot be shifted between years, because the increase in distortion as rates increase is more than linear.⁶⁶ Thus the reduced distortion from lowering the rates in some years should exceed the increased distortion from raising them in other years.⁶⁷

Now suppose instead that tax liability depends on consumption. The use of a lifetime measure effectively results in a uniform commodity tax. One faces the same marginal rate on an extra unit of consumption no matter the year in which it occurs. If the applicable marginal rate depends instead on the taxpayer's marginal rate bracket given her other consumption in a particular year, we get a violation of the AS 1976 analysis. In effect, applicable commodity tax rates are higher in some years than in others, distorting decisions regarding when to consume, without any mitigation of the basic underlying labor supply distortion.⁶⁸

D. *Implementing Income Averaging*

William Vickrey's "cumulative averaging" scheme remains the most

66. See GRUBER, *supra* note 7, at 582-83.

67. Weisbach, *supra* note 11, emphasizes two considerations potentially raised by income averaging that I ignore here. The first is that longer accounting periods increase economic distortion if there is no time-value adjustment for when in the period one received a given dollar. In principal, however, no matter how long or short the period one uses informationally, one can make present-value adjustments within the period. Second, Weisbach notes the administrative costs of requiring more frequent filing, clearly an important consideration although distinct from the question of what information is used each time, whether in applying the rate structure or otherwise.

68. See Bankman & Weisbach, *supra* note 3, at 1422-28.

prominent implementation of the idea that lifetime tax liability should depend on permanent income.⁶⁹ Vickrey's basic idea was to have the taxpayer, each year, determine her average annual income through that year and then adjust the amount she had paid to date to equal what she would have paid to date had she earned the average amount each year.

A simplified illustration of the Vickrey proposal, assuming for simplicity an interest rate of zero, is as follows. Suppose the tax rates are 20% for income up to \$100,000 and 50% above that, and only two years are at issue. In addition, suppose that three taxpayers (*A*, *B*, and *C*) each earned \$200,000 total, but in different sequences: *A* had an even split between the two years, *B* earned everything in Year 1, and *C* earned everything in Year 2. The Vickrey system would apply as follows:

Table 1. Hypothetical Illustration of Vickrey Income Averaging

	Year 1		Year 2		Total	
	Earnings	Tax	Earnings	Tax (Refund)	Earnings	Tax
<i>A</i>	\$100,000	\$20,000	\$100,000	\$20,000	\$200,000	\$40,000
<i>B</i>	\$200,000	\$70,000	0	(\$30,000)	\$200,000	\$40,000
<i>C</i>	0	0	\$200,000	\$40,000	\$200,000	\$40,000

Note: Tax rate is 20% on income up to \$100,000 and 50% on income over \$100,000. *B* and *C* pay the amounts in Year 2 needed to equalize them with *A*.

Vickrey did not definitely commit himself to a particular averaging period, but noted the logic of extending it all the way from adulthood (reflecting that few children have significant earnings) until death.⁷⁰ This presumably would cause retirees to get annual cash refunds for as long as they lived, since, in the absence of significant current year income, the reduction in average earnings would be treated as reducing taxes due for all prior years. Vickrey's system might therefore effectively provide retirees with life annuities, albeit declining ones given that each year the arithmetic effect of adding one more year to the denominator in the average earnings computation would decline slightly. This annuity feature may be irrelevant, however, if one assumes complete markets and consistent rational choice, since under those assumptions people would arrange for the exact level of annuitization that they wanted, whether income averaging contributed to it or not.

Vickrey's system goes well beyond anything that one could realistically expect Congress to enact. However, from 1964 through 1986, the U.S. tax system had a narrower income averaging rule, under which people with rising incomes could in some circumstances apply the lower marginal rates that they had faced in preceding years to some of the growth component. In terms of Table 1, relief was offered to individuals resembling *C*, but not to those

69. See Fennell & Stark, *supra* note 11, at 4.

70. VICKREY, *supra* note 9, at 186.

resembling *B*. Income averaging was repealed in 1986, without any explanation in the legislative history of the reason for the change.⁷¹ So the main remaining form of income averaging in the U.S. system is that resulting from the allowance of net operating losses to offset positive taxable income in other years.⁷²

A second income averaging episode worth noting is that of Wisconsin, which between 1929 and 1935 based current year income tax liability on the taxpayer's average annual income for the prior three taxable years.⁷³ The rule applied automatically rather than being elective, and it benefited taxpayers with falling as well as rising income. Thus, if you earned zero in a given year after having earned \$100,000 in one of the prior two years and \$200,000 in the other, for the current year you would pay the tax due on \$100,000 of income (the moving three-year average as of year-end).

One important difference between Vickrey-style cumulative averaging and the former U.S. and Wisconsin systems concerns the effect on current year tax payments when annual income fluctuates. A rolling average system like that in Wisconsin causes annual tax liability to rise (relative to the amount that would be due under a purely annual system) when income declines.⁷⁴ Similarly, both the Wisconsin and U.S. systems caused annual tax liability to fall (relative to the outcome under a purely annual system) when annual income rose. Vickrey criticized such a relationship between income fluctuation and current year tax liability, arguing that if one "assess[es] a heavy tax based on previous high incomes in years when the income has sharply decreased . . . collection is difficult and hardship to the taxpayer results."⁷⁵ Under his system, these effects were mitigated by the fact that one was adjusting the tax payments due for prior years, in addition to using averaging to determine the amount due this year.⁷⁶ Vickrey did not notice that his quite reasonable argument for not letting the current year tax payment grow too large relative to current year income was in tension with the view, underlying the case for income averaging, that lifetime income, without regard to the exact sequence of cash flows, is all that matters.

71. See JOINT COMM. ON TAXATION, GENERAL EXPLANATION OF THE TAX REFORM ACT OF 1986, at 14-16 (1987). I worked on the 1986 repeal of income averaging as a Legislation Attorney at the Joint Committee on Taxation. My recollection is that repeal reflected the grounds (in considerable tension with each other) that (1) rate reduction made income averaging unnecessary, and (2) the revenue gain from repeal was needed to help pay for tax reform. These grounds were in tension because, if income averaging really was no longer needed, repealing it presumably would not have raised significant revenue. Staffers also were aware of Richard Schmalbeck's then-recently published work criticizing the income averaging rules, see Schmalbeck, *supra* note 11, but there was no general consensus that an annual system is best.

72. See I.R.C. § 172 (2000).

73. See VICKREY, *supra* note 9, at 183-84; Batchelder, *supra* note 11, at 415 n.68.

74. The U.S. income averaging rules provided no relief for taxpayers with falling incomes.

75. VICKREY, *supra* note 9, at 169.

76. *Id.* at 180.

E. Problems with the Case for Income Averaging

The twin assumptions of complete markets and consistent rational behavior may have been easy to accept in the context of a stylized hypothetical concerning a world with nothing but apples and oranges. The assumptions become a lot more heroic, not to mention implausible and demonstrably untrue, when we shift to real world questions of the timing of consumption across one's lifespan. In addition, with or without the twin assumptions, the case for income averaging may be weakened by the significance of additional information, whether revealed by the taxpayer's earning or consumption sequences or by our general knowledge about people's lifecycle patterns.

1. Incomplete markets

In the hypothetical with the apples and oranges, all that the "complete markets" assumption required was full tradability between these two goods. Once we turn to lifetime income, however, the complete markets that one needs are capital markets. People must be able to (1) invest current earnings for future consumption, (2) borrow against future earnings for current consumption, and (3) adjust fully and immediately to changes in current earnings whether anticipated or not. Only the first of these three requirements is relatively unproblematic. Departures from the latter two include the following:

a. "Great expectations"

Borrowing against high expected future earnings is not always feasible, even absent risk concerning a given individual's capacity to realize her "great expectations" in the future. For example, students at leading professional schools who are virtually guaranteed future high-wage employment opportunities may face significant limits on their capacity to borrow against that future capacity to fund current consumption, even if they can borrow extensively to fund their educations.

These difficulties reflect the classic insurance problems of moral hazard and adverse selection. As to moral hazard, once I have borrowed against expected future earnings but will be judgment-proof unless I actually realize those earnings, the liability functions like a tax on the earnings, diminishing my incentive to realize them. As to adverse selection, I may know more than prospective lenders about the future earnings I actually should expect. This may reflect inside information either about my ability level or about my future plans.⁷⁷

77. Education loans may be less risky than loans for general consumption purposes, notwithstanding the moral hazard and adverse selection problems, because I may be less likely to want access to funds that are being spent on the education if I do not actually plan to use it towards the goal of realizing high earnings.

To illustrate the significance of the great expectations problem for income averaging, suppose that Caleb and Diana face identical circumstances, including having the same lifetime income, except that Caleb has level earnings once he achieves adulthood, while Diana has back-loaded earnings that she cannot access during the "great expectations" stage. Suppose further that Caleb and Diana both prefer perfect consumption smoothing during their adult years. The fact that Caleb can accomplish this while Diana cannot, given available capital markets, rebuts the presumption, derived from the permanent income hypothesis, that they are equally well off. Caleb, as compared to Diana, is (a) better off on a lifetime basis, given his ability to achieve the preferred consumption pattern, (b) better off during the early adult years when Diana, unlike him, is effectively poor, and (c) worse off once Diana's high earning period has begun. At this later stage, while Diana still is worse off on a lifetime basis given the earlier deprivations that she would have avoided if possible, she now has extra cash available and, for the remaining period, is effectively richer than Caleb.

Under these circumstances, the distributional equity argument for income averaging fails to hold. In a sense, lifetime comparisons between Caleb and Diana are not even meaningful, since the relative value of a dollar to either of them depends on when it is realized (a point that the lifetime income measure ignores). Rather than ensuring that they pay the same lifetime taxes, as would a perfect income averaging system, we should want to transfer resources from Caleb to Diana before her great expectations begin to be realized, and from Diana to Caleb afterwards. This suggests treating Diana's two periods as separate rather than as subject to averaging, so that she will (rightly) appear to be a low-income taxpayer in the first period and a high-income taxpayer in the second one.

The efficiency arguments for income averaging likewise lose ground in the great expectations scenario. Level lifetime rates continue to be desirable from the standpoint of eliminating incentives to shift earnings between periods. Moreover, there is still some argument for rate smoothing insofar as (all else equal) lowering high rates tends to reduce distortion more than raising low rates increases it.⁷⁸ On the other hand, Diana is, in effect, two different individuals across time even if she has constant preferences. She is a low-earner early on and a high-earner later on. This might lead to differences in her labor supply elasticity as between the two periods, since neither the rate of trade that she faces between consumption and leisure nor her budget line is constant. Such differences might cause smooth rates to be less efficient than having a higher marginal rate in the period (whichever it is) in which her labor supply is less responsive to changes in her after-tax return. As noted in a recent contribution to the NDPF literature, when one's skill level changes, "the tradeoff between

78. See Gruber, *supra* note 7, at 582-83.

insurance and incentives then shifts and taxes should adjust accordingly."⁷⁹ Tax rate smoothing ignores this.

b. Risk

A second, less obvious missing capital market relates to risk and consequent changes in information when a risk is resolved. If I expect low career earnings and then suddenly learn that they will be high, or vice versa, I cannot (absent a time machine) suitably adjust the amounts I consumed in earlier periods to reflect what I now know. Accordingly, even if all other conditions for the model's applicability fully hold, I will fail ex post to achieve the optimal sequencing of my lifetime consumption.

The private market solution to this problem would be for me to fully insure against risks that my earning ability will change. However, moral hazard and adverse selection prevent full insurance from being available. A lesser mitigation device, precautionary saving against the risk of a decline in actual or expected earnings, increases expected smoothing but may actually end up reducing it ex post, if it turns out that I did not need to save for a rainy day after all.

The classic permanent income and consumption smoothing models exclude uncertainty,⁸⁰ which Milton Friedman conceded "blurs the sharp lines of the . . . analysis and suggests additional factors that may produce departures from the shape of the consumption function [otherwise] specified."⁸¹ While merely an additional technical challenge from the standpoint of properly specifying models to describe rational lifecycle behavior,⁸² uncertainty more definitely compromises reliance on the permanent income hypothesis to support viewing all individuals who ex post have the same lifetime incomes as having faced identical circumstances.⁸³

Uncertainty undermines both the distributional and the efficiency cases for income averaging. Suppose that two periods in a taxpayer's life differ, so far as earnings are concerned, in part predictably (e.g., one's income rises with skill or seniority) and in part unpredictably due to the resolution, at the dividing line between two periods, of a discrete risk. The predictable difference would lead a

79. Golosov et al., *supra* note 47, at 332.

80. See FRIEDMAN, *supra* note 19, at 7-14; Modigliani & Blumberg, *supra* note 61, at 392.

81. FRIEDMAN, *supra* note 19, at 15. Friedman had precautionary saving particularly in mind. *See id.* at 16.

82. *Id.* at 16 (suggesting that uncertainty would encourage precautionary saving); Modigliani & Blumberg, *supra* note 61, at 392 (arguing that "a satisfactory theory can be developed without seriously coming to grips with this rather formidable problem"); *id.* at 428-29 (noting some ways in which uncertainty may affect saving).

83. As David Kamin has pointed out to me, this problem also arises in a purely annual system, where risky income may convey different total and marginal utility than certain income even if the amounts end up being equal ex post.

rational individual to engage in consumption smoothing between the two periods, suggesting that income averaging would be appropriate unless smoothing was impeded by the “great expectations” problem. However, inability to retrofit the amount of earlier-period consumption to resolution of the risk would suggest treating the two periods as separate, and thus not permitting income averaging so far as the risky outcome is concerned.

Risk disrupts basing distributional judgments purely on lifetime income even if two individuals turn out not only to have had the same total lifetime income, but even to have earned exactly the same amount each year. Suppose Edith and Frank end up earning the same amount each year, but that Edith is subject to greater earnings risk throughout and thus rationally chooses greater precautionary saving. She thus ends up back-loading her lifetime consumption relative to Frank’s and relative to what would have been optimal for her had her actual ultimate earnings stream been certain. Frank therefore experiences greater total utility, all else equal, possibly motivating redistribution from him to Edith even though, judged through the lens of income averaging, they end up being the same.

A further risk-related distributional issue, arising from incomplete markets, concerns difficulties in promptly adjusting consumption levels in response to changes in actual and expected earnings. Suppose a high-earner who owns an expensive house suddenly and unexpectedly loses her job, indicating that she must start to live somewhere more modest. Adjusting her housing consumption may take time (e.g., since she needs to find a buyer or renter), and the disruption will be greater if she is carrying a mortgage with high monthly payments that she cannot finance out of other savings. Vickrey presumably had such scenarios in mind when he criticized Wisconsin-style income averaging for responding inappropriately to declining income. He appears, however, not to have appreciated the broader tension between such scenarios and viewing lifetime income as the canonically correct basis for determining lifetime tax burdens.

Turning to efficiency, the change in skill level likewise undermines the case for income averaging. Once again, in light of the change “the tradeoff between insurance and incentives . . . shifts and taxes should adjust accordingly.”⁸⁴

c. Inability to annuitize

Recent research suggests that life annuities generally are not available at actuarially fair prices, presumably due to adverse selection.⁸⁵ That is, people

84. Golosov et al., *supra* note 47, at 17.

85. See Amy Finkelstein & James Poterba, *Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market*, 112 J. POL. ECON 183 (2004) (finding evidence of adverse selection in U.K. lifetime annuity markets). If life annuity markets are

with private information suggesting that they will live longer than indicated by the actuarial tables sign up disproportionately, making fair returns based on those tables a losing proposition for the insurers. This gap in financial markets arguably suggests that the government should provide mandatory life annuities.⁸⁶

Under this view, the life annuity aspect that income averaging features when the averaging period continues through death may be a step in the right direction, rather than being irrelevant as suggested above. Given, however, that other instruments for mandatory annuitization exist, such as mandatory retirement programs along the lines of Social Security and Medicare, the significance of income averaging's contribution here remains uncertain. One presumably needs to coordinate it with the other programs, but it does not expand the government's opportunity set, nor does its design have anything directly to do with the question of how much mandatory annuitization is desirable.

2. Departures from consistent rational choice

A second empirical failing of the permanent income hypothesis is its assuming consistent rational choice. Behavioral economics research has revealed a number of real-world departures from this model of behavior.⁸⁷ These departures suggest that two people with the same lifetime incomes but in different sequences, and with the same preferences may make different consumption choices, implying that they will face different circumstances, even with complete markets.

a. Hyperbolic discounting and other causes of myopia

There is considerable evidence that people often fail to save adequately for retirement, reflecting myopia rather than a consistent preference for concentrating consumption in their working years.⁸⁸ Psychological

badly hampered by adverse selection, it is unclear why markets for life insurance are so robust, given that the two financial instruments reflect the same bet (whether the insured will die early or late), differing only in the side of the bet that the insured takes.

86. GRUBER, *supra* note 7, at 359-61.

87. See, e.g., BEHAVIORAL LAW AND ECONOMICS (Cass R. Sunstein ed., 2000) (collecting articles exploring the new field of behavioral law and economics); RICHARD H. THALER, QUASI-RATIONAL ECONOMICS (1991) (surveying recent research concerning departures from rational behavior). For examples of important departures from behavior predicted by the rational choice model, including declining to disregard sunk costs, self-control problems that lead to hyperbolic discounting, and over-reacting to salient or anecdotal information, see *id.*

88. See, e.g., Robert J. Shiller, *Social Security and Institutions for Intergenerational, Intragenerational, and International Risk Sharing* 42 (Nat'l Bureau of Econ. Research, Working Paper No. 6641, 1998) (noting the frequency of reduced consumption at retirement,

explanations differ, but one prominent rationale is hyperbolic discounting, or applying a much higher discount rate between the current time and any future time than between future times.⁸⁹ People who are subject to hyperbolic discounting cannot hold consistent preferences. For example, at Time 1 a hyperbolic discounter will want to apply a normal discount rate in dividing consumption between Times 2 and 3. Once Time 2 arrives, however, she will be much more inclined to concentrate her consumption in Time 2.

To the extent that people act under the sway of hyperbolic discounting, current consumption will depend on currently available resources rather than on any consistently held long-term plan. To take an extreme case, suppose people always consumed all current earnings, neither saving for the future nor borrowing against the present value of expected future earnings. Then, for any year, both total utility and the marginal utility of a dollar would depend purely on current earnings (ignoring psychological carryover effects from remembering past years or anticipating future ones), and income averaging would give current distributional weight to information from other periods that was in fact irrelevant. Even under a less extreme view, to the extent that people lean towards current consumption of currently available resources, not because they are optimizing across the lifespan but because they are behaviorally subject to “presentist” bias, current year information gains distributional relevance relative to information from other periods, thereby weakening the case for income averaging.

This point has particular implications for the design of need-based transfers such as TANF. Proponents of income averaging do not generally take the intellectually consistent step of urging that eligibility for need-based transfers depend purely on lifetime income, which would suggest denying the benefits to currently destitute individuals who had high earnings in the past. This departure from consistently favoring a lifetime standard can be defended on the ground that being destitute despite past earnings is *prima facie* evidence of poor lifetime consumption choice, rather than of having rationally preferred to concentrate one’s consumption in earlier periods.

b. Mental accounts

Even without thoroughgoing myopia, consumption behavior may diverge from that predicted by the permanent income hypothesis due to people’s use of “mental accounts” to classify dollars differently depending on their source. In effect, these are rules of thumb that people use to economize on decision costs. Thus, amounts coded as “current income” apparently are more likely to be spent than those coded as “current assets,” which in turn are more likely to be

suggesting that people act as if reaching retirement age is a surprise).

89. See, e.g., David Laibson, *Golden Eggs and Hyperbolic Discounting*, 112 Q.J. ECON. 443, 445-46 (1997).

spent than those coded as "future income."⁹⁰ Empirical research strongly supports the view that these predilections cause significant behavioral departures from the predictions of the permanent income hypothesis, in particular by causing current consumption and current income to be "much more closely linked" than one would expect if it fully applied.⁹¹ Once again, the implication is that the time sequence of earnings matters independently of their discounted lifetime value, and thus that current year information should have greater influence on current year distribution policy than information from other periods.

3. *Additional information*

The problems with income averaging that I have discussed thus far seem to suggest paradoxically that it uses too much information, by going beyond current year earnings to look at other years' earnings as well. What makes this seem paradoxical is that one would think that more information is always better than less, particularly when the underlying distribution problem is framed in informational terms as one of trying to gauge individuals' true ability levels.

The explanation for the apparent paradox is as follows. Reliance on the permanent income hypothesis involves treating information from other periods not just as currently relevant, but as equal in relevance to current period information. It flattens the lifespan, treating it as a single period and ignoring distinctions between stages, along with details concerning the exact sequences of earnings and consumption. When the internal pattern is important, a purely annual approach, while clearly suboptimal unless the other periods are totally irrelevant, can potentially be better than acting as if all information from all periods was equally relevant at all times. Better to ignore the other information, perhaps, than to over-weight it.

One can also, however, take the route of improving the lifetime measure by adding information, in lieu of subtracting it. Rather than treating the lifetime as a single period, one can instead attempt proper relative weighting and use of the information from different periods. Here the NDPF literature becomes especially pertinent, as it expressly focuses on the significance of information that unfolds over time.⁹² Examples of using information beyond current-year earnings to enrich the current year picture include the following:

90. See Hersh M. Shefrin & Richard H. Thaler, *The Behavioral Life-Cycle Hypothesis*, in *QUASI-RATIONAL ECONOMICS*, *supra* note 87, at 96-98.

91. Christopher D. Carroll & Lawrence H. Summers, *Consumption Growth Parallels Income Growth: Some New Evidence*, in *NATIONAL SAVING AND ECONOMIC PERFORMANCE* 305, 305 (B. Douglas Bernheim & John B. Shoven eds., 1991).

92. See, e.g., Kocherlakota, *supra* note 48.

a. Past earnings

Income averaging relies on the permanent income hypothesis to treat career earnings as the only relevant evidence of ability, without regard to the earnings sequence. As the NDPF literature emphasizes, however, information about ability levels plays out over time.⁹³ A high-earning year is informative about a given individual's likely ability level even if subsequent earnings are lower, since the decline, while it might reflect an adverse shock to ability, might also reflect a voluntary change in labor supply. A rational and benign government that was engaged in providing insurance against ability risk would inevitably make use of information about past earnings in deciding what taxes to levy in a given period.⁹⁴

One possible implication is what one might call reverse income averaging.⁹⁵ In Vickrey-style cumulative averaging, a taxpayer with high earnings in Year 1 and low earnings in Year 2 may pay less in Year 2 than she would have under a purely annual system because in effect she gets a refund for Year 1 as the measure of her ability is revised downward. In an NDPF framework, the implication might instead be levying a higher tax in Year 2 than we would have based purely on annual information because we have evidence suggesting reduced work effort (notwithstanding the possibility that there has instead been an adverse shock to ability).⁹⁶ This implication is all the stronger if the taxpayer has significant savings, which could have helped to motivate the decision to work less by providing an alternative source of current financial support.⁹⁷

b. Savings

The NDPF literature closely associates imposing higher taxes by reason of past high earnings with taxing saving. From the perspective of the current period, high saving may help to finance (and thus explain) an otherwise

93. See, e.g., Mikhail Golosov, Narayana Kocherlakota & Aleh Tsyvinski, *Optimal Indirect and Capital Taxation*, 70 REV. ECON. STUD. 569, 580 (2003).

94. See Golosov et al., *supra* note 47, at 20 (describing the dynamic time inconsistency problem in this framework, where a government that cannot credibly commit against subsequently exploiting early-years labor supply information about ability may overly discourage labor supply in such years); William P. Rogerson, *Repeated Moral Hazard*, 53 ECONOMETRICA 69 (1985) (considering the importance of memory in repeated principal-agent interactions with moral hazard).

95. Jeff Strnad suggested this phrase to me.

96. See Kevin Roberts, *The Theoretical Limits to Redistribution*, 51 REV. ECON. STUD. 177 (1984).

97. See Stefania Albanesi & Christopher Sleet, *Dynamic Optimal Taxation with Private Information*, 73 REV. ECON. STUD. 1 (2006); Narayana R. Kocherlakota, *Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation*, 73 ECONOMETRICA 1587 (2005).

surprisingly high decline in earnings.⁹⁸ From the perspective of past periods, since someone planning to reduce her labor supply would be expected to save, observing high savings can in effect serve as a proxy for directly observing high past earnings.⁹⁹ This association between conditioning taxes on past earnings and on savings reflects the overlap between issues of income averaging and of tax base choice. However, since the use of savings in NDPF models relates more directly and obviously to the choice of ideal tax base, I defer discussing it to Part III.

c. Age-related wage distribution and labor supply elasticity

A recent article by Michael Kremer¹⁰⁰ makes the point that an optimal tax structure might impose distinct marginal rates on people of different ages, for two reasons. The first is that labor supply elasticity may vary with age. In particular, young and old people tend to have more elastic labor supply than people in mid-career. Second, the annual earnings distribution varies with age, with young and old people being more concentrated at the low end of the scale than people in mid-career. Thus, the rates applying at low and medium income levels will apply at the margin (and thus to the next dollar of potential earnings) to a smaller proportion of mid-career than of young and old workers. Both considerations suggest that young and old people should generally face lower marginal rates than mid-career workers.

These arguments for age-based taxation are intellectually consistent with the permanent income hypothesis, since they do not rebut the claim that one allocates consumption across one's lifetime budget line without regard to the time sequence of earnings. However, they contradict the implication in support of income averaging that tax liability should depend purely on the discounted value of lifetime earnings. They imply once again that taxes should depend on when a given dollar of earnings is realized, in this case reflecting additional information about particular stages within the lifecycle.

F. Conclusions Regarding Income Averaging

The problems with the case for income averaging can be interpreted in a couple of different ways. Strictly speaking, the strong logical case for it falls short. Income averaging is not generally and universally desirable, even ignoring issues of administrative cost. One should keep in mind, however, that the presumed alternative, a purely annual system, might fare even worse if

98. See Kocherlakota, *supra* note 97.

99. See Albanesi & Sleet, *supra* note 97, at 25 (noting that their model, unlike an alternative one that observes past earnings directly, "use[s] wealth to summarize aspects of an agent's past history").

100. Michael Kremer, *Should Taxes Be Independent of Age?* (2001) (unpublished manuscript), available at <http://www.economics.harvard.edu/faculty/kremer/papers.html>.

scrutinized with comparable rigor.

For a purely annual view to hold, each year would have to be completely separate from all others, in two senses. First, people would have to be unable or unwilling to shift work or consumption between years, whether by saving or borrowing. Second, information from other years would have to be irrelevant—not merely potentially less relevant than current-period information—to the assessment of one's relevant current year attributes, such as wage rate, labor supply elasticity, and utility function. This is not compelling as a picture of the world we live in.

Given the problems with a purely annual view, there remains a significant case in support of income averaging's main effect, which is to mitigate the extra tax burdens on people who have fluctuating rather than relatively constant annual incomes. Some design details arguably supported by the preceding analysis are the following:

1) Given the "great expectations" point, the case for relief is weaker for people who newly enter high-earning stages of their careers, such as by entering a high-wage profession after graduating from school. The U.S. income averaging system that was on the books from 1964 through 1986 attempted roughly to address this problem by denying income averaging to certain individuals who had not been self-supporting throughout the averaging period (which went back four years).¹⁰¹

2) Given the existence of other tools, such as Social Security and Medicare, to address inadequate retirement saving, arguably there is no reason to let people who retire claim income averaging benefits with respect to their pre-retirement earnings. Moreover, insofar as the timing of one's retirement is a matter of discretionary personal choice, the decline in income that results from it does not indicate being worse off or having suffered a negative shock to ability. The payoff to retirement, in the form of eliminating tax liabilities with respect to the foregone earnings, may already be too high under a purely annual system, given that one presumably is not worse-off overall, despite the lost income, by reason of retiring voluntarily. Extending the benefits of income averaging in these circumstances would make the arguably excessive tax benefit of retirement greater still.¹⁰²

3) Beginning a career and ending it through retirement are only two examples of status changes across which the case for income averaging is especially weak. Income averaging might also be denied in cases where there is an identifiable change in, say, the taxpayer's health status or occupation. As for age, while denying income averaging between different life stages (such as youth and middle age) may be desirable, the problem is harder because any particular point at which one tries to draw a line between periods is likely to be

101. See I.R.C. § 1303(c)(1) (1982); Schmalbeck, *supra* note 11, at 522-23.

102. Cf. Schmalbeck, *supra* note 11, at 572 (arguing against income averaging for retirees).

arbitrary.

4) At first glance, extending income averaging to individuals with rising incomes but not to those with falling incomes (in the manner of the old U.S. rules) may appear arbitrary. It may very roughly be supported, however, by the point from the NDPF literature that high past earnings are evidence of high current ability. Even if one does not go so far as to embrace reverse income averaging, denying the positive benefits to individuals with falling incomes may be reasonable as a general rule.

5) Given the importance of current cash flows, it may be desirable to structure income averaging in the Vickrey/cumulative manner, rather than in the U.S./Wisconsin manner, so that current tax payments tend to rise and fall with current income. Given the liquidity problems that people with declining incomes may face, this is all the more important if income averaging is indeed extended to people with falling as well as rising incomes.

III. IMPLICATIONS FOR THE CHOICE BETWEEN INCOME AND CONSUMPTION TAXATION

I now turn to the application of an optimal income tax framework to the choice between the two most prominent and widely accepted tax bases, income and consumption. As we will see, the case for consumption taxation is overwhelming under the same circumstances as that for income averaging, involving in each case complete markets, consistent rational choice, and the informational completeness of aggregate lifetime earnings data.

A. Overlap Between the Cases for Income Averaging and for Consumption Taxation

A core argument of this paper is that the cases for income averaging and for consumption taxation are closely linked. This is not to deny that one could logically support one of these two approaches while opposing the other, even ignoring administrative considerations, given that either approach might be rationalized in a number of different ways. There are indeed many examples of people who have supported one while opposing the other. William Vickrey, for example, supported both income averaging and income taxation.¹⁰³ While one could rightly have criticized him for not appreciating the tension (discussed below) between income taxation and a lifetime perspective, he could equally correctly have responded that supporting a tax on saving does not pre-commit one to believing that people with fluctuating earnings should generally pay more tax than those with level earnings. After all, if two individuals both earn and save the same amounts on a lifetime basis, then presumably an income tax, no less than a consumption tax, should treat them the same, but graduated

103. See VICKREY, *supra* note 9.

annual rates might cause one to pay more than the other if they had different earnings sequences.

Likewise, one can favor consumption taxation on grounds that are distinct from adopting a lifetime perspective. Edward McCaffery, for example, believes on ethical grounds distinct from welfare economics that what he considers excessive or overly concentrated consumption should be tax-penalized. Thus, he favors using an annual rate structure to disfavor highly concentrated consumption relative to that which follows a smoother pattern.¹⁰⁴

Within a consistent welfare economics framework, however, the link between the two issues is hard to sever. Distributionally, both follow from viewing lifetime budget lines as the key attribute in discerning total and marginal utility, a stance that relies on the permanent income hypothesis to establish that all equivalent lifetime budget lines are indeed effectively the same. In terms of efficiency, both rely on the AS 1976 analysis. Objections to the underlying assumptions of the permanent income hypothesis, concerning complete markets and consistent rational choice, weigh similarly against both, as does the case for using additional information to supplement the evidence about ability that lifetime earnings provide.

B. The Distributional Case for Consumption Taxation Under the Permanent Income Hypothesis

From a welfare economics perspective, the core distributional argument for a consumption tax is that lifetime earnings determine one's budget line. Savings decisions merely reflect commodity choice within this budget line as between present and future consumption. Thus, individuals who save more and thus derive greater returns to saving (which an income tax reaches and a consumption tax does not) are not relevantly better off and should not pay more tax on a lifetime basis.

To illustrate, suppose there are two periods and that the rate of return on resources saved in Period 1 for consumption in Period 2 is 5%. George and Hilda both earn \$100 in Period 1, meaning that they can consume \$100 in Period 1, \$105 in Period 2, or some lesser combination of goods in each period. Suppose George chooses to consume \$100 in Period 1 while Hilda chooses to consume \$105 in Period 2. While they are relevantly equal, Hilda would pay more than George under an income tax. Suppose the tax rate was 40%. Both would pay \$40 in Period 1, while Hilda would pay \$2 and George would pay zero in Period 2. By contrast, under a consumption tax, they would effectively pay the same amount. Given the 5% interest rate, George's \$40 tax payment in Period 1 would have the same present value as Hilda's \$42 tax payment in Period 2. Finally, consider an earnings tax, which is equivalent to a

104. See McCaffery, *supra* note 3, at 87-91.

consumption tax since returns to saving are excluded from both.¹⁰⁵ Once again George and Hilda would be treated the same, here by reason of each paying \$40 in Period 1 and zero in Period 2.

In this hypothetical, the only required capital market is one permitting saving at a 5% rate between Periods 1 and 2. Suppose, however, that George and Hilda had each earned \$105 in Period 2, instead of earning \$100 in Period 1. The analysis would remain the same so long as available capital markets permitted George to borrow \$100 in Period 1 against his expected earnings. Now the mechanism by which an income tax, unlike a consumption tax, would treat him better than Hilda is through the allowance of a \$5 interest deduction in Year 2 when he repays the loan.¹⁰⁶

A consumption tax, by taxing George and Hilda the same, gets the correct distributional result if they are relevantly identical given that they have the same lifetime incomes. This equivalence between George and Hilda follows from the permanent income hypothesis plus the absence of other pertinent distinctions between them.

Now suppose we reject the equivalence of individuals with the same lifetime incomes, because incomplete markets or departures from consistent rational choice make the exact sequences of earning or consumption important. This does not necessarily make the case for an income tax, which appears to rest on treating income, including returns to saving that reflect commodity choice rather than different opportunity sets, as synonymous with the normative concept of ability to pay.¹⁰⁷ It does, however, move towards leveling the intellectual playing field between the two tax bases, by weakening a compelling argument for consumption taxation.

What is more, viewing periods as separate rather than as linked is generally more consistent with an income tax than a consumption tax framework. If one is looking purely at consumption opportunities (and thus ability to pay) within a given period, then all of one's wealth is relevant, and an income tax at least comes closer than a consumption tax to achieving the wealth tax ideal by taxing the return to wealth. Moreover, the argument that, rather than ignoring unspent wealth, a consumption tax merely defers payment of the tax on such wealth (without any reduction in tax burden) until consumption occurs loses force if we are focused on the current period and do not regard as relevant what might happen in future periods. The classic argument that an income tax better measures ability to pay than does a consumption tax, because it takes account of unspent wealth via the inclusion of returns to saving, is explicitly present-

105. See Shaviro, *supra* note 3, at 96.

106. The actual existing income tax disallows deductions for consumer interest. See I.R.C. § 163(h)(1)-(2) (2000). In principle, however, the argument for allowing interest deductions on all dissaving under an income tax is identical to the argument for taxing the positive interest income on net saving.

107. See, e.g., HENRY C. SIMONS, PERSONAL INCOME TAXATION (1938) (supporting income taxation on the ground that income measures ability to pay).

period focused.

C. *The Efficiency Case for Consumption Taxation Under the Permanent Income Hypothesis*

It has been known for a considerable time that an income tax distorts two choices, by discouraging both work and saving, while a consumption tax distorts only one, by discouraging work. The belief that this creates an efficiency tradeoff between the two taxes, and that an income tax might be more efficient because its imposing a tax on saving permits it to impose a lower tax burden on work, is still occasionally expressed.¹⁰⁸ However, as Joseph Bankman and David Weisbach recently showed, it is logically incorrect.¹⁰⁹ They note that an income tax is effectively a differential commodity tax, in which future consumption goods are taxed at a higher rate than current consumption goods, violating the key finding of AS 1976. Distorting decisions regarding when to consume by taxing future commodities at a higher rate than current commodities does nothing to mitigate the underlying labor supply distortion. Instead, the discouragement of saving for future consumption is simply layered on top, adding to total distortion.¹¹⁰

As noted earlier, this is equivalent to arguing for income averaging within a consumption tax, since otherwise the AS 1976 finding is violated by imposing different marginal rates on consumption in different years. It likewise is akin to the efficiency argument for income averaging that emphasizes rate smoothing as to earnings in different years, except that here the purely additive distortion relates to when one works, rather than to when one consumes. The efficiency case for consumption taxation, like that for income averaging, thereby relies on viewing the entire lifespan as a single uniform period, a view that is easiest to support if we assume complete markets, consistent rational choice, and lack of other pertinent information. As I discuss next, modifying these assumptions muddies the case for consumption taxation, just as it does for income averaging.

D. *Problems with the Case for Consumption Taxation*

1. *Incomplete markets*

As noted in Part II, when markets are incomplete, as in the “great expectations” scenario, and when earning ability is risky, lifetime comparisons are thrown off and individuals with the same lifetime incomes may differ in total utility and have varying period-specific marginal utilities even if they are

108. See, e.g., GRAVELLE, *supra* note 7; GRUBER (First Edition), *supra* note 7, at 708.

109. Bankman & Weisbach, *supra* note 3, at 1448-49.

110. See *id.* at 1427.

otherwise identical. This undermines the equity case for consumption taxation, albeit possibly without otherwise advancing the case for income taxation,¹¹¹ by contradicting reliance on lifetime income as the proper standard of comparison for distributional purposes.

2. Departures from consistent rational choice

Departures from consistent rational choice, like incomplete markets, to some extent simply disrupt the clean logical case for consumption taxation without putting anything else in its place, by indicating that we do not know enough to make confident claims about total and marginal utility. In one important respect, however, the case for income taxation may be more directly aided. This relates to the possibility that high saving is itself evidence of ability, rather than simply or even primarily evidencing a greater preference for later, as compared to current consumption.

Suppose we believe that most people, given their utility functions, should engage in substantial lifetime consumption smoothing, including saving adequately for retirement, but that this requires mental and emotional skills that the population holds very unevenly. High savers might, in this scenario, on average be more patient and farsighted than low savers, having greater self-control and capacity to restrain counter-productive impulses. In addition, they might generally do a better job of optimizing their choices, such as by modifying simple rules of thumb regarding consumption patterns where it is feasible to do better.

Under such a view, higher saving clearly would be evidence of a broader ability of some kind. The harder question is what sort of ability, and with what sorts of implications for distribution policy. Suppose initially that having these skills implied high earning ability, independent of the evidence offered by observed wages. This might, for example, reflect that savers benefit within the labor market from being more patient and farsighted than non-savers, enabling them to earn more even if their skills are otherwise the same. Taxing savers more than non-savers who had the same lifetime incomes might then be distributionally optimal, for the same reason as taxing high-earners more than low-earners.¹¹² Under this view, while an income tax would not necessarily combine the two types of information optimally, at least it would be making

111. In a truly comprehensive or Haig-Simons income tax, fluctuations in the value of human capital would have current income tax consequences, just like changes in the value of any other asset. See Louis Kaplow, *Human Capital Under an Ideal Income Tax*, 80 VA. L. REV. 1477 (1994). This would suggest taxing individuals with "great expectations" on the appreciation of their human capital even if they had no access to the value before the future labor income started being realized. This arguably would be in tension with the underlying normative basis for an income tax, if we think of ability to pay, defined in terms of income, as having something to do with currently accessible resources.

112. See Emmanuel Saez, *The Desirability of Commodity Taxation Under Non-Linear Income Taxation and Heterogeneous Tastes*, 83 J. PUB. ECON. 217, 228 (2002).

positive use of a type of information that a consumption tax, by being savings-neutral, ignores.

Now suppose instead, however, that high saving merely denotes ability as a consumer—that is, the capacity to derive more utility than others could from using the same resources, by properly deploying them to the point in one's lifetime where they improve wellbeing the most. This type of ability has mixed implications from a welfare economics perspective. On the one hand, the abler consumer presumably has the higher total utility, supporting redistribution to less able consumers if we believe that the marginal utility of a dollar generally declines with rising total utility, and/or if we give independent weight to equality in wellbeing. On the other hand, the fact that the abler consumer can derive greater utility from an extra dollar suggests redistributing wealth to her, all else being equal, rather than away from her.

3. Additional information

An income tax uses more information than does a consumption tax, since it makes lifetime net tax liability depend on savings decisions and thus on the timing of consumption relative to earning. In one important respect, this extra information may improve the performance of the income tax, relative to that of the consumption tax, as a redistributive tool. The NDPF literature strongly suggests that taxing savings, at least when it is at high levels or associated with past high earnings, may actually be optimal.¹¹³

The basic point is that, with high savings, one can more easily afford to under-utilize one's earning ability by working less, thereby in effect camouflaging oneself as a lower-ability individual than one actually is. Decisions to work less impose a negative revenue externality on the government, thus reducing its ability to provide workers with insurance against having low ability, and requiring it to meet any further revenue needs in some other way that may involve increased distortion. This negative effect of savings on labor supply potentially makes it "optimal for society to deter savings by taxing it."¹¹⁴

Articles in the NDPF literature or tradition have made this point at least since an important 1978 article by economists Peter Diamond and James Mirrlees.¹¹⁵ While some models suggest imposing no net tax on saving—for example, by having taxes on high savers or those whose earnings surprisingly

113. This point was apparently first made in Peter Diamond & James Mirrlees, *A Model of Social Insurance with Variable Retirement*, 10 J. PUB. ECON. 295 (1978). For more recent examples, see Golosov et al., *supra* note 93; Narayana Kocherlakota, *Wedges and Taxes*, 94 AM. ECON. REV. 109 (2004); Kocherlakota, *supra* note 97; Juan Carlos Conesa, Sagiri Kitao & Dirk Krueger, *Taxing Capital? Not a Bad Idea After All!* (Nat'l Bureau of Econ. Research, Working Paper No. 12880, 2007).

114. Golosov et al., *supra* note 93, at 577.

115. Diamond & Mirrlees, *supra* note 113.

decline offset by subsidies to low savers or those whose earnings surprisingly increase¹¹⁶—others agree with Diamond and Mirrlees in finding a general case for taxing saving in response to the externality. In particular, a recent article by Mikhail Golosov, Narayana Kocherlakota, and Aleh Tsyvinski finds a tax on saving to be optimal, by reason of the labor supply effect and consequent revenue externality, in a relatively general setting. These authors conclude that AS 1976, while applicable to support the desirability of uniform commodity taxation of items that are simultaneously available at a given time, does not apply (as recent legal authors such as Joseph Bankman and David Weisbach have argued)¹¹⁷ to support a zero rate of capital income taxation.¹¹⁸

E. Conclusions Regarding Consumption Taxation

Departures from the assumptions underlying the permanent income hypothesis and additional information have four effects on the otherwise compelling welfare economics case for consumption taxation. First, they introduce enough noise to make any definite conclusion about the ideal system less tenable than it would otherwise be. Second, by suggesting that current period circumstances may at times be more relevant than information about other periods, they support a more present-focused system than one focused, like a consumption tax, on lifetime income. While this may not directly transfer into support for income taxation, it may add to the distributional relevance for the current period of unconsumed wealth. An income tax, unlike a consumption tax, assigns current-period tax consequences to unconsumed wealth if there is a positive return to such wealth. Third, if higher saving, all else being equal, tends to be a signal of earning ability, then taxing returns to saving may further the core distributional aim of providing social insurance against "ability risk." (The consequences are more ambiguous, however, insofar as high savings indicate ability as a consumer, in the sense of being able to extract more utility from the same resources.) Finally, taxing saving may be socially optimal given that savings can reduce labor supply, thereby generating a negative revenue externality that adversely affects ability insurance and may require higher distortionary taxes to help finance government purchases.

These points refute the claim that an ideal consumption tax is decisively superior to an ideal income tax, as judged from the standpoint, not of a stylized economic model, but of the actual world in which we live. More broadly, searching for the best ideal tax base may be misguided as an intellectual enterprise. The issues raised by market incompleteness, inconsistent consumer

116. See Kocherlakota, *supra* note 113, at 109 (suggesting a tax on high savers and a subsidy to low savers if future skills are uncertain); Kocherlakota, *supra* note 97, at 1588 (suggesting positive wealth tax for individuals whose earnings surprisingly fall, along with a wealth subsidy for those whose earnings surprisingly rise).

117. Bankman & Weisbach, *supra* note 3.

118. Golosov et al., *supra* note 93, at 580.

choice, and within-period information are too empirical for the question of tax base choice to be resolvable at a highly abstract level. Nor should this theoretical ambiguity be a surprise once we recall the great distance, discussed at the beginning of Part I, between the things that really matter in an optimal income tax framework and the available tools. Again, if we are interested in total and marginal utility but are largely restricted to making inferences about how ability might affect them, and when we further cannot even directly observe ability, the lack of a simple, clearly dominant optimization strategy is only to be expected. One should not be startled by the lack of a uniform, one-size-fits-all answer to the question of how savings should be treated by the tax system.

While these considerations may dispose of the ideal tax base debate as a matter of pure intellectual inquiry, this is not to say that people are misguided when they argue for one or the other ideal in practice. Simplified tax base ideals may help to improve political outcomes—for example, by offering a clear framework against which to criticize and judge special interest legislation.¹¹⁹ While complicated NDPF-style inquiry may help to inform one's choice between competing ideals, it is unlikely to itself prove well-suited for holding the center stage. Accordingly, proponents of both fundamental and interstitial tax reform may reasonably conclude that they should pick one tax base ideal or the other, and mostly stick to it while engaged in public advocacy.

From this standpoint, while it comes down to personal judgment, I myself continue to favor what economist David Bradford called the “consumption strategy.”¹²⁰ I base this conclusion on two main considerations. First, while saving can have negative revenue externalities that support taxing it at a positive rate, it also may have positive externalities that economic models often ignore. For example, it may raise the living standards of future generations.¹²¹ Second, and perhaps more importantly, taxing saving requires periodically measuring the changes in value of assets held by the taxpayer. Absent realization transactions such as asset sales, this can be extremely difficult.¹²² Current income tax law responds, more or less inevitably, by providing that unrealized gains and losses, for the most part, are ignored.¹²³ This, in turn, leads not just to the exclusion of some saving from current taxation, but to a vast array of tax-motivated transactions designed to pair deferred gain against currently deductible loss. One could not easily over-estimate how much of the

119. See DANIEL SHAVIRO, WHEN RULES CHANGE: AN ECONOMIC AND POLITICAL ANALYSIS OF TRANSITION RELIEF AND RETROACTIVITY 93-98 (2000).

120. See DAVID F. BRADFORD, UNTANGLING THE INCOME TAX 312 (1986).

121. See SHAVIRO, *supra* note 40, at 81.

122. See, e.g., Alvin C. Warren, Jr., Commentary, *Financial Contract Innovation and Income Tax Policy*, 107 HARV. L. REV. 460, 462 (1993).

123. On departures under current law from requiring realization and analysis of the feasibility of extending these departures, see, for example, David A. Weisbach, *A Partial Mark-to-Market Tax System*, 53 TAX L. REV. 95 (1999).

complexity of current income tax law results from this seemingly simple problem. Hence, as William Andrews noted almost thirty years ago, a realization requirement is the Achilles' heel of any practically feasible income tax.¹²⁴ While relying on the realization problem as the main motivation for shifting to a consumption tax may seem more humdrum than deriving broad theoretical conclusions that turn out to depend on the accuracy of the permanent income hypothesis and on subtle informational issues, such reliance may ultimately prove more compelling and harder to rebut.

CONCLUSION

The permanent income hypothesis is a powerful and appealing idea, logically applying basic precepts of neoclassical economics. It also has the pleasant consequence of making the normative analysis of ideal tax policy seem a lot simpler by establishing a compelling case, within the assumptions of welfare economics, both for consumption taxation and for lifetime income averaging.

Unfortunately, from the standpoint of making life simple, the premises that the permanent income hypothesis require do not entirely hold. First, markets are not complete. For example, it may be hard to borrow against one's future expected earnings, and it is impossible to change past consumption decisions by reason of new information about available lifetime resources. Second, people do not always exhibit consistent, rational choice across time. For example, they may be myopic or prone to hyperbolic discounting, and they may use mental accounts in deciding how to use various dollars, which leads them to violate the principle that a dollar is a dollar is a dollar. These departures increase the relative importance of current period information, thereby exposing as overly flat and undifferentiated the one-period picture of an entire lifespan that the lifetime income concept offers.

Two next steps in the tax policy debate seem indicated by the analysis in this Article. The first is exploring how income averaging might work, as a technical matter, while limited to the types of circumstances where I have argued that it is most appropriate (i.e., within what is effectively a single period for the taxpayer, subject to internal consumption smoothing and without significant changes in current earning ability). The second is further exploring the case from the NDPF literature for taxing saving by reason of its negative revenue externalities. This might involve not just reassessment of the case for income taxation as a public political ideal, but also efforts to tailor the tax on saving to circumstances where the case for it is strongest.

124. See William D. Andrews, *The Achilles' Heel of the Comprehensive Income Tax*, in *NEW DIRECTIONS IN FEDERAL TAX POLICY FOR THE 1980s*, at 278, 280 (Charls E. Walker & Mark A. Bloomfield eds., 1983).

REPLY

CONSUMPTION TAXATION IS STILL SUPERIOR TO INCOME TAXATION

Joseph Bankman* & David Weisbach**

INTRODUCTION.....	789
I. OUR ANALYSIS IN <i>THE SUPERIORITY OF AN IDEAL CONSUMPTION TAX</i>	791
A. <i>Basic Argument</i>	791
B. <i>Qualifications</i>	794
II. SHAVIRO'S ARGUMENT	794
III. DISCUSSION.....	795
A. <i>Imperfections in the Credit Market</i>	795
B. <i>Inadequate Insurance Markets</i>	797
C. <i>Myopic Savings Behavior</i>	800
D. <i>Savings as Signal of Ability</i>	801
CONCLUSION.....	802

INTRODUCTION

We appreciate the opportunity to comment on Dan Shaviro's important piece on how the permanent income hypothesis relates to tax policy.¹ In this piece, Shaviro points out that the arguments in favor of a consumption tax on the one hand, and income averaging on the other, raise significantly related issues. As far as we know, Shaviro is the first person writing in the legal literature to make this connection, and his insightful work is sure to inspire further explorations of the issue.²

* Ralph M. Parsons Professor of Law & Business, Stanford Law School.

** Walter J. Blum Professor and Kearney Director, Program in Law and Economics, The University of Chicago Law School.

1. Daniel Shaviro, *Beyond the Pro-Consumption Tax Consensus*, 60 STAN. L. REV. 745 (2007).

2. See also Mitch Engler, *Tax Neutrality, Vickrey Averaging, and Income Versus Consumption Taxation* 33-36 (Apr. 5, 2007) (unpublished manuscript), available at http://www.law.nyu.edu/colloquia/taxpolicy/papers/07/Mitch_Engler_Tax_Neutrality.pdf



Shaviro's basic claim is that the permanent income hypothesis is a central underpinning of both the arguments for income averaging and consumption taxation. Market or rationality imperfections that limit the permanent income hypothesis similarly weaken the case for income averaging and consumption taxation. In particular, Shaviro raises a number of important complexities regarding the choice between an income tax and a consumption tax that are related to how capital markets work and how individuals choose their consumption patterns. He claims that these complexities make the efficiency and distributional case for a consumption tax sufficiently nuanced and dependent on empirical assumptions that the choice of tax base is better made on administrative grounds. On administrative grounds, however, Shaviro strongly prefers a consumption tax, notwithstanding the title of his article.

We agree with three of the central points of Shaviro's argument: that more complex models of behavior are likely to weaken the strong conclusions one gets from simple models; that administrative concerns may well be dominant; and that a consumption tax remains the preferred tax base. Like Professor Shaviro, we think that additional research into administering consumption taxes should be given a high priority.

The fact that we are in broad agreement on these issues may surprise readers of his article. Shaviro writes that his analysis "refute[s] the conclusion of a recent leading article," and it is our article he cites.³ The structure of his article is to cite our article for its presentation of the core argument in favor of a consumption tax, and then to critically discuss the assumptions on which that argument relies. The implication is that our article ignores these assumptions. In fact, the bulk of our article is devoted not to presenting the core argument in favor of a consumption tax, but to critically discussing the assumptions on which it relies. We discussed in detail no less than ten such assumptions and we concluded that "a complete, optimal tax analysis could produce exotic taxes that look like neither a pure consumption tax nor a pure income tax."⁴ This conclusion is similar to the conclusion reached by Professor Shaviro.

In other respects, however, our papers and analyses differ. Shaviro's primary goal is to show the connections between the permanent income hypothesis, income averaging, and the consumption tax debate. A secondary goal is to show the divergence between an ideal welfarist tax and a consumption (or income) tax. Flaws in the consumption (or income) tax matter to him even if they don't clearly alter the relative merits of the consumption tax over the income tax.

(exploring the related issues raised by income averaging and consumption taxation).

3. Shaviro, *supra* note 1, at 750 n.18, referring to Joseph Bankman & David A. Weisbach, *The Superiority of an Ideal Consumption Tax over an Ideal Income Tax*, 58 STAN. L. REV. 1413 (2006).

4. Bankman & Weisbach, *supra* note 3, at 1416.

Our paper focuses on the choice between an income tax and a consumption tax. We focus on this choice, rather than the possibility of exotic taxes suggested by some of the optimal tax literature, for two reasons. First, these are the two leading tax bases, and real world reforms will be aimed at moving the tax system toward one of those two bases. Second, relatively simple and transparent taxes reduce political economy problems, such as rent seeking, that more complex taxes create.

While we agree with Shaviro that the study of administrative and compliance costs is vital, we believe that analysis of the efficiency and distributional effects of the two ideal systems remains relevant. The reason for this is that we cannot measure the benefits of administrative and compliance cost savings without a reference point. For example, if a consumption tax is easier to implement, we need to know if an income tax is otherwise desirable to be able to determine whether the administrative savings are worth the costs in using the less preferred base.

Our two papers also differ in their evaluation of the simplifying assumptions made in the base case for a consumption tax. We conclude that, even once one examines these assumptions, including the assumptions that Shaviro finds most questionable, a consumption tax is superior to an income tax. Shaviro does not take a position on that issue.

I. OUR ANALYSIS IN *THE SUPERIORITY OF AN IDEAL CONSUMPTION TAX*

A. *Basic Argument*

Whether the tax system ought to be built around an income or a consumption tax has been a primary—some would say *the* primary—issue in tax policy for many years. Both tax bases have some attractive features. However, most scholars (including Professor Shaviro) believe that a consumption tax would be easier to implement and superior in that respect. Support for an income tax, therefore, is generally based on the belief that an income tax is fairer or (less frequently argued) more efficient than a consumption tax. Our article focuses on (and rejects) the fairness and efficiency arguments made in support of an income tax.

We begin by assuming that individuals in society vary along one or more dimensions, and in particular, in their ability to earn. The government wants to redistribute toward those with low ability to earn for some reason, whether it is declining marginal utility of income or an aversion to inequality. Unfortunately, it cannot directly observe ability, so it must rely on signals such as income or consumption. If it taxes high incomes at too high a rate, individuals will choose to work less, thereby reducing the government's ability to redistribute. The complicated balancing between redistribution and working less was analyzed in

detail in James Mirrlees's work on optimal taxation.⁵

Mirrlees imagined that individuals varied by their ability to earn—effectively their wage rate—and that the government imposed a tax only on labor income. Would a tax on savings help in such a system? The answer—given by Atkinson and Stiglitz in 1976 and further developed by Louis Kaplow in recent years—is no.⁶ The reason is that anything we can achieve with a tax on savings can be achieved better by an adjustment to the rate schedule to the tax on labor income.

Consider the following illustration, previously set forth by one of the authors:

[S]uppose that there is a 20 percent income tax. Suppose also that an individual earns \$250 before taxes, pays a \$50 tax on his labor income, and is left with \$200 to spend. Finally, suppose that the individual spends \$100 today and saves \$100 for his retirement in 25 years. The annual interest rate on savings is 5 percent.

Under the income tax, as normally conceived, the tax on capital income reduces the return on the savings by the tax rate, here from 5 percent to 4 percent. Thus, with no taxes on savings, the individual in retirement would have the future value of \$100 at 5 percent, or roughly \$340 to spend. With taxes, the individual earns only 4 percent and his future consumption is reduced to about \$270. The reduction in future consumption from \$340 to \$270 is like a 20 percent excise tax on that future consumption.

Suppose that instead of the conventional income tax, we impose a tax on labor income at 29 percent, a subsidy of 12 percent for current consumption, and a 12 percent tax on future consumption. Under this tax, the individual has after-tax wages of approximately \$180. If he spends half this amount, \$90, and receives the 12 percent subsidy, he can spend \$100 today as before. He saves the other half of his after-wage-tax earnings and must pay the additional 12 percent tax on savings. This leaves him with about \$80, which grows at the pretax interest rate to the same \$270. . . . If the individual would choose this pattern under the usual income tax, he could also choose this pattern under the [restated] tax.⁷

The wage or labor income tax in this example (without the subsidy for spending and the tax on saving) is the implicit tax on labor income imposed by the nominal rate under the income tax. We know that the actual tax imposed on labor income by an income tax is higher than the nominal rate because the tax reduces the benefit of working by more than the nominal rate. The reason is

5. See, e.g., J.A. Mirrlees, *An Exploration in the Theory of Optimum Income Taxation*, 38 REV. ECON. STUD. 175 (1971).

6. A.B. Atkinson & J.E. Stiglitz, *The Design of Tax Structure: Direct Versus Indirect Taxation*, 6 J. PUB. ECON. 55 (1976); Louis Kaplow, *On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal*, 90 J. PUB. ECON. 1235 (2006).

7. David Weisbach, Comment, in TAXING CAPITAL INCOME 143, 145 (Henry J. Aaron et al. eds., 2007) (footnote omitted). The numbers in the text reflect very crude rounding to simplify the presentation.

that it also reduces one of the benefits of working: future consumption. An hour of work brings both less immediate benefit due to the explicit labor portion of the tax and also less of one of the items one might want with one's wages: future consumption. To find the implicit rate on labor income, we find the rate where the set of taxes and subsidies on consumption choices is zero. At this point, we have captured the full effect of the tax on the benefits of working. With the numbers in the example, this rate is 29%. The restated tax simply makes this implicit tax explicit.

Compare the 29% wage tax to the income tax. By construction, they both impose the same tax on labor income. Therefore, labor effort and the labor/leisure distortion will be the same under the two taxes. This also means that both taxes raise the same revenue from wages. The income tax, however, has the offsetting tax and subsidy. This combination produces no additional revenue but distorts savings decisions.⁸ The wage tax eliminates this distortion. Leaving aside possible second-best issues discussed in our article, the wage tax is, therefore, superior for this individual. The individual has the same distortion in work effort and government revenues are the same, but the individual is able to make less distorted savings decisions.

The restated tax in this example was a tax on labor income. As is well known and illustrated in our article, however, a tax on labor income is equivalent to a consumption tax. Therefore, not only is the 29% wage tax the better choice in this example, but so is an equivalent consumption tax. A consumption tax at the right rate makes this individual better off and nobody worse off.

The analysis above concerned only a single individual and, therefore, did not take into account distributional effects. We can, however, perform the same normalization within each wage class, making the implicit tax on labor income explicit. Those with higher wage income are likely to save more. The implicit tax on labor income in an income tax, therefore, will be relatively higher for high wage classes, which means that the explicit tax rate on labor income will be correspondingly higher for these classes.

For example, the implicit tax rate in the example discussed above was 29% because of the enormously high 50% savings rate. If the individual earned less and saved less—say he saved 5% of after-tax earnings—the implicit tax on labor income would be 21%. We can find the implicit tax rate on labor within each wage or income class. If the rich tend to save more than the poor, the procedure will produce a higher explicit labor tax rate on the rich than on the poor.

We can then compare an income tax with a wage or consumption tax that uses the implicit rates on labor income created by the income tax. For the same reasons that applied to the single-individual example, for each income class,

8. The distortion of savings makes labor income less valuable (because its use to fund future consumption is hobbled), and in that way the tax/subsidy distorts labor effort as well.

individuals in that class would be better off facing a consumption tax at the implicit rate than facing the income tax. If we replace the income tax with a consumption tax using these rates, therefore, individuals in each class—poor, middle, rich, and filthy rich—would be better off. Distributional concerns would be entirely resolved, and efficiency improved. That is, we can view the choice between an income tax and a consumption tax as merely the choice of whether to impose residual, revenue- and distribution-neutral subsidies for spending and taxes on savings. Imposing this residual tax is highly unlikely to be desirable both because we tend not to want to interfere with consumption choices and because of all the choices to interfere with, encouraging spending and penalizing savings seems like a bad one. For this reason, and in addition to administrative cost concerns, an ideal consumption tax is superior to an ideal income tax.

B. *Qualifications*

Like all models, the Atkinson and Stiglitz model rests on a number of simplifying assumptions. In our article, we discuss the effect of varying those assumptions by allowing for (1) risk; (2) economic rents (supernormal profits); (3) graduated rates on capital income; (4) the presence of wealth prior to the introduction of a consumption tax; (5) inheritances; (6) savings heterogeneity; (7) myopic savings behavior; (8) imputed income from capital; (9) complementarity between leisure and some forms of consumption; and (10) savings as an indicator good that signals ability.⁹ All of these factors complicate the analysis. If tax administration were costless, some of these factors (such as the complementarity between consumption and leisure or indicator goods) would produce tax systems that look nothing like either an income or a consumption tax. We conclude, however, that given current constraints on information and administration, none of the factors substantially alters our conclusion that a pure consumption tax is superior to a pure income tax, and many make the case for a consumption tax stronger.

II. SHAVIRO'S ARGUMENT

Shaviro draws a connection between the permanent income hypothesis, income averaging, and consumption taxation. The permanent income hypothesis, in its most general form, simply states that individuals borrow and lend across time periods to maximize their utility based not on their current income but on their expected future income.¹⁰ For example, one might imagine

9. See Bankman & Weisbach, *supra* note 3.

10. MILTON FRIEDMAN, A THEORY OF THE CONSUMPTION FUNCTION (1957); Robert E. Hall, *Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence*, 86 J. POL. ECON. 971 (1978); Franco Modigliani & Richard Brumberg, *Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data*, in POST-

that to the extent people can, they would equalize marginal utility across periods.

Income averaging attempts to equalize tax rates across periods so that one's rate is based on overall lifetime (or averaging period) income or consumption. Someone with level consumption and someone with variable consumption that has the same overall present value as the level consumption would both pay the same tax under an averaging system.

A consumption tax, Shaviro argues, relies on the same arguments. Suppose Bob and Jane each earn \$100 today, and Bob consumes it all while Jane saves it until next year and consumes \$110. A consumption tax treats them the same in the sense that it imposes the same present value tax on each. To be relevantly the same, Shaviro argues, we have to believe that Bob and Jane's decisions regarding when to consume maximize their well-being, which Shaviro argues means that they follow some version of the permanent income hypothesis.

Shaviro argues that the assumptions underlying the permanent income hypothesis—and by extension, the assumptions underlying the case for income averaging or a consumption tax—are shaky. Market imperfections in the credit or insurance markets may prevent individuals from maximizing utility across periods or may otherwise make income averaging or a consumption tax less desirable. Utility maximization may be thwarted by myopia and, finally, savings may serve as a signal of utility that again makes both income averaging and a consumption tax less desirable. Two of these four arguments—that taxpayers may be myopic and that savings may signal ability—were among the ten complexities we discussed in our article. We did not discuss the implications of failures in the credit or insurance markets. We consider each of Shaviro's arguments below. Before we begin this discussion, however, we should note that the vast majority of Shaviro's paper focuses on income averaging. Neither of us has, as of yet, thought seriously about how the permanent income hypothesis affects the case for income averaging beyond reading Shaviro's paper. We do not address it here. Instead, we focus on Shaviro's arguments concerning consumption taxation.

III. DISCUSSION

A. Imperfections in the Credit Market

Shaviro begins by discussing the implications of credit market imperfections. He gives the example of Caleb and Diana with the same lifetime consumption except that Caleb has level income while Diana has back-loaded earnings. If Diana cannot borrow against future earnings when she is young because of problems with capital markets, her lifetime welfare might be lower notwithstanding equal lifetime earnings. She cannot smooth her consumption,

KEYNESIAN ECONOMICS 388 (Kenneth K. Kurihara ed., 1954).

and thus, assuming declining marginal utility of income, the decline in utility from low consumption in the first period is not offset by the increase in utility from high consumption in the later periods. We might think of Diana as a medical student or doctor just starting her practice. Diana is worse off than Caleb, whose consumption is equal across periods. A consumption tax, Shaviro argues, would treat Diana and Caleb the same when they are relevantly different.

The argument relies on credit constraints, but it is unclear how much of a social problem credit constraints are. Diana, in the example, can easily borrow to fund her medical school education and borrow to purchase her home. There are many sources of loans for starting a business. Loans to finance everyday consumption are available to essentially everyone through credit cards. Even with these many sources of credit, it is unclear if Diana would want smooth consumption. She might be happier as a student living in a dorm and eating inexpensive ethnic food, preserving more income to fund high consumption and a larger house in later years. Before using an income tax to offset market failures, we need to be sure that the market is actually failing.

Shaviro concludes that credit constraints undermine the case for a consumption tax but do not further the case for an income tax. In fact, the presence of credit constraints might be seen as a point in favor of an income tax over a consumption tax. The reason for this is that savings are a signal that an individual is not credit constrained. Savings correlates with an underlying attribute not affected by credit market imperfections. Lack of savings correlates with being affected by credit market imperfections. Since being credit constrained is undesirable, all else equal, savers as a whole might have higher utility or lower marginal utility than spenders. An income tax levies a higher tax on savers and in so doing, might redistribute from the better off (not credit constrained) to the worse off (credit constrained).

There are, however, a number of problems with this argument. First, it is not clear what portion of those who do not save are credit constrained. That is, we would have to have a sense of how good the signal is. It could be that many who do not save are myopic and wish they saved more, rather than being credit constrained and wishing they had spent more. As we note below, an income tax may exacerbate the problems caused by savings myopia.

Second, an income tax makes poor use of the signal. Even if savings is a good signal that an individual is not credit constrained, an income tax imposes a burden on each increment of savings even though the signal happens only on the first increment (or perhaps the increment beyond precautionary savings). Therefore, it will impose offsetting distortions that would have to be weighed against any benefits.

Finally, we would have to believe there were not better ways of dealing with imperfections in the credit market. For example, the best fix might be to loosen regulatory constraints or adopt policies that make borrowing cheaper. We might, for example, make it easier for individuals to get mortgages or credit

with which to purchase consumer durables. This, of course, flies in the face of current thinking (that there is too much easy credit) and supports our suggestion, above, that the problem posed by inadequate credit for the Dianas of the world may be more theoretical than real.

In sum, we believe the problem of insufficient credit is unlikely to have a meaningful impact on the debate between an income and a consumption tax.

B. *Inadequate Insurance Markets*

Shaviro argues that the new dynamic public finance (NDPF) literature significantly undermines the conclusions of hitherto conventional economical analysis of the desirability of a consumption tax and presents a strong rationale for taxing savings.¹¹

The basic idea is to extend the standard optimal tax analysis over time. The government's role in the standard analysis is to redistribute from high ability individuals to low ability individuals to maximize a social welfare function. In Mirrlees' case, we can think of the government's role as offering insurance against having low wages. In the NDPF approach, the government also insures against wage *changes* over time. That is, an individual's wage rate or ability is assumed to be risky, possibly going up or down in the future. Private insurance against the risk of wage changes is assumed to be unavailable, so the government steps in to provide it. The government uses taxes and transfers to help in this role as insurer.

Given this insurance, individuals plan their labor effort and consumption based on their wage rates and the risk of higher or lower future wages. In economic jargon, individuals dynamically plan an optimal labor and consumption path in a multi-period game. They have anything but myopia in these models.¹² The government's problem is to maximize the sum of expected discounted utility given individual planning and given that it cannot observe ability.¹³

11. See Shaviro, *supra* note 1, at 777, 785. Major works in the NDPF literature include Mikhail Golosov et al., *Optimal Indirect and Capital Taxation*, 70 REV. ECON. STUD. 569 (2003); Narayana R. Kocherlakota, *Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation*, 73 ECONOMETRICA 1587 (2005); Ivan Werning, *Optimal Unemployment Insurance with Unobservable Savings* (2002) (unpublished manuscript), available at <http://econ-www.mit.edu/files/1267>. For a summary of the literature, see Mikhail Golosov et al., *New Dynamic Public Finance: A User's Guide*, 2007 NBER MACROECONOMIC ANN. 317, available at <http://econ-www.mit.edu/files/866>.

12. A question for Shaviro's paper regards the inconsistency between arguing that individuals are myopic and arguing that we should believe the NDPF literature.

13. As with myopia, the NDPF literature makes an assumption that Shaviro elsewhere questions, which is that the government should maximize expected discounted utility. To the extent one believes Shaviro's arguments that the government should not sum up utility across periods, they undermine the NDPF arguments on which he relies.

In this setting, the private return to savings may not be the same as social return. One intuition is that the offer of insurance for someone who has low wages in a future period creates an incentive to work less to get the insurance benefits. Savings enhances the ability to shirk, ensuring consumption even with low labor effort. Therefore, there may be some role for taxes or subsidies on savings to help improve the tax/benefit scheme on labor income.

The literature is framed in terms of "wedges," which are the difference between the private and social return to savings. The private return is higher than the social return when the government offers insurance against wage changes. It takes a second and very complicated step to translate these wedges into taxes. Most often the taxes on capital are complicated, nonlinear functions of all prior period earnings and savings. Moreover, such taxes typically raise zero revenue, merely redistributing across spenders and savers. And they often are higher on those with low wages than on those with high wages because they are designed to prevent individuals from pretending that they have low ability when they are merely shirking. Except in specialized and unrealistic cases, they look nothing like income taxes.

The relevant question is whether or to what extent these arguments support an income tax. An initial observation is that they do not produce taxes that resemble income taxes. As mentioned above, the taxes are exotic and unintuitive. It is possible that with more modeling and empirical estimating of parameters, the best approximation to the resulting optimal tax is an income tax, but there is nothing in the current literature that points in that direction.

A second observation is that the results depend critically on the assumed parameters, and to get sizable wedges, the assumptions have to be made unrealistic. To illustrate, it will be useful to examine the model presented in a leading survey article cited by Shaviro and the benchmark case in the numerical simulations presented in that article.¹⁴ In *New Dynamic Public Finance: A User's Guide*, Mikhail Golosov, Aleh Tsyvinski, and Ivan Werning present a two period model in which the government tries to address a failure in the insurance markets. In period one, individuals work and save. However, in period two, each individual faces a 50% chance of suffering a 50% decline in her skills. The government would like to insure individuals against that decline. However, it has no information whatsoever to know whether an individual claiming to have suffered this decline is lying. To put the matter differently, if we characterize this severe skill shock as a form of partial disability, every individual can "fake" this disability with 100% success. To prevent everyone from faking disability, the government sets the disability payments at less than the wages available to those who do not claim disability. Some individuals will respond to the presence of uncertainty, and the availability of even partial disability insurance, by saving a great deal in the first period. These individuals will disproportionately fake disability in the second period. Their saving in the

14. Golosov et al., *supra* note 11, at 336-46.

first period ensures that they will enjoy more level consumption in the two periods, and thus optimizes utility, conditional upon claiming disability in the second period. The authors calculate that the government should make this practice less desirable by levying a wealth tax of 5%. (There is no interest in this model, so tax on savings is actually a tax on wealth.)

The wealth tax in this model is due to the fact that there is no way to detect disability. As a result, there are no private disability markets, and the government is left with the second-best scheme to provide partial insurance and to support that scheme through a series of otherwise undesirable taxes (including the wealth tax described above). In the real world, of course, there are private insurance markets. There are also ways to achieve insurance-like protection through explicit or implicit labor contracts or relationships. Examples of this include long-term contracts, union or civil service membership, and tenure or the de facto tenure found in many organizations. Where external insurance exists, in this model and other models in the literature, the tax on capital disappears.¹⁵

Absent any form of insurance, the magnitude of the capital tax in the model depends on the probability of a decline in skills. Because the government is performing an insurance function, the size of the implied capital tax falls as second period wages become more certain—whether the movement is toward a certain decline or a certain non-decline of skills. The magnitude of the implied tax on capital also declines with a decline in the magnitude of the adverse outcome. This is true even if individuals face the maximum uncertainty as to whether that outcome will be experienced. The tax on capital falls to near zero, for example, if individuals face a 50% chance of a decline of 10% in skills. In some cases, such as where individuals face a positive skill shock and risk aversion is high, the optimal tax on capital is negative—that is, the tax turns into a subsidy.

To summarize, the literature relies on government being the best and only provider of insurance against changes in income. Private markets either cannot exist or are entirely crowded out by the government. To get sizable taxes, the literature must assume a very high degree of uncertainty about future wages, both in terms of the likelihood of a change and the size of the change. Depending on the parameters, there can be a subsidy rather than a tax on capital. Moreover, even within the assumptions, the taxes that are produced are not anything like income taxes.

The NDPF literature is still too young to have produced multi-period models that incorporate realistic levels of income/skill variation and insurance protection. The literature is developing rapidly but it is simply too early to tell

15. Ken Judd makes a similar observation that the NDPF literature assumes the absence of private markets and argues that instead, “we should do as Mirrlees did, assume that private markets work, and then find the policy that best achieves the goal taking into account the presence of a private market.” Kenneth L. Judd, Comment, 2007 NBER MACROECONOMIC ANN. 381, 382, available at <http://econ-www.mit.edu/files/866>.

whether it will produce robust conclusions and what those conclusions might be.

C. *Myopic Savings Behavior*

Shaviro points out correctly that standard arguments for both income averaging and a consumption tax rely on rational choice, and writes that "[w]eakening this assumption, such as positing that people are myopic, undermines the case for both averaging and consumption tax." He discusses myopia in any detail only in connection with the case for averaging. However, the logic of his argument and the statement quoted above suggest that his discussion of myopia in the context of income averaging ought to carry over to the consumption tax context.

Briefly stated, the argument that myopic savings behavior weakens the case for a consumption tax is as follows. A consumption tax has the seemingly attractive quality of levying the same present value tax on the same present value of consumption. If the interest rate is 10%, someone who earns and spends \$100 will be treated as equally wealthy as someone who earns \$100, saves it for a year and earns 10% of interest, and then spends \$110 in the next year. In our discussion above and in our article, the choice between present and future consumption is treated similarly to the choice between vanilla and chocolate ice cream. We assume the individual chooses so as to maximize her utility but do not believe that we can say in advance which choice will leave the individual wealthier. Suppose, though, that some individuals act myopically, turning down opportunities to save even when the effective interest rate is extremely high. That would be evidence that non-savers are systemically worse off than the savers. A consumption tax ignores this effect. In contrast, an income tax recognizes the greater welfare of savers by levying a higher effective tax rate on that group.

In our article, we discuss the case for myopia and briefly review some of the empirical and experimental evidence.¹⁶ We find that the evidence is inconclusive but can be read to show some sign of myopic savings behavior, particularly among lower income individuals. (Shaviro does not discuss the literature in as much detail, so we do not know if he would agree with this conclusion.) We further conclude, however, that myopia is more likely to strengthen, rather than weaken, the case for a consumption tax. The reason for this is that myopia leads to excessive current consumption and too little savings. A conventional response to the problem of too little savings would be to reduce the tax on savings. A consumption tax does this. An income tax does the opposite. It reduces the relative cost of consumption over saving and thus encourages consumption. Adopting an income tax, which encourages spending out of solicitude for those who are made worse by excessive spending, is

16. Bankman & Weisbach, *supra* note 3, at 1444-48.

perverse. It is like noting that smokers are worse off than non-smokers and as a remedy reducing the relative price of cigarettes.¹⁷

D. Savings as Signal of Ability

The final argument¹⁸ Shaviro raises is that savings might be a signal of ability to earn. Consider two individuals with the same income but different savings. The individual with the savings might have higher ability and be working less hard. We can use the savings as a means of separating the two individuals and taxing the one with higher ability more.

This argument dates back to an article by James Mirrlees in 1976¹⁹ and was made recently by Emmanuel Saez in 2002.²⁰ We discussed it in our prior work under the rubric of indicator goods.²¹

There are a number of serious problems with the argument. The first is that there is simply no evidence as of yet that savings indicates higher ability. Shaviro cites none. Saez cites a single paper which does not support the proposition.²² It might some day be proven true, but right now it is not even a disputable proposition except through introspection.

Second, the mechanism that is claimed, through introspection, to support the claim that those who save have higher ability is that the need for savings and the sophistication to access capital markets is something that one gets through education, and that education is something that only the more able can obtain. If this is the mechanism, then education is the signal of ability, not savings.

Finally, even if savings were an indicator of ability, current models only show that a marginal tax on savings is desirable.²³ The models do not support anything like an income tax, which imposes a tax on savings at the same rate as the tax on labor income. Perhaps better models would show that the tax on savings should be large, but the current work, including everything cited by Shaviro, does not come close to supporting an income tax.

17. The adoption of an income tax would not be perverse if the behavior of spenders were fixed and unaffected by the return to savings. In that case, it would benefit a worse-off group and not exacerbate the problem it is designed to alleviate. But we are unaware of anyone who has made this empirical claim.

18. This is the final argument in our ordering, not Shaviro's.

19. J.A. Mirrlees, *Optimal Tax Theory: A Synthesis*, 6 J. PUB. ECON. 327, 341-44 (1976) (demonstrating that the presence of multi-dimensional taxpayer heterogeneity changes optimal tax selection).

20. Emmanuel Saez, *The Desirability of Commodity Taxation Under Non-Linear Income Taxation and Heterogeneous Tastes*, 83 J. PUB. ECON. 217, 225-26 (2002).

21. See Bankman & Weisbach, *supra* note 3, at 1453-55.

22. Saez, *supra* note 20, at 228 (explaining the desirability of taxing certain forms of consumption in the presence of taxpayer heterogeneity).

23. For example Saez's model only supports a marginal tax on savings even assuming that it is an indicator of ability. See *id.*

Shaviro also raises the possibility that savings might be signal for the ability to transform consumption into utility. If this were true, savers would get more utility out of a unit of consumption than spenders. Shaviro does not, however, give any reason to imagine why savers and spenders would differ in this respect. Moreover, as Shaviro notes, there is no clear-cut intuition that this difference, if it exists, would cut in support of an income or a consumption tax.

CONCLUSION

In our original paper, we framed the debate between an income tax and a consumption tax as a debate over idealized systems.²⁴ Shaviro argues that a debate over idealized systems misses many important issues, particularly tax implementation.

We agree with this point. Implementation should be central to the debate, and theoretical models that do not include implementation can lead us astray. Although more work needs to be done, it is likely, as Shaviro concludes, that implementation concerns lean heavily toward consumption taxation.

Nevertheless, the debate over idealized systems is important because it focuses research. If a consumption tax is preferable but has implementation issues, we end up with a very different research agenda than if a consumption tax is a bad idea. In the former case, research resources should be devoted to determining how best to implement a consumption tax, while in the latter case, they should not. Moreover, ideals matter because if ideals can only be reached at a cost, we must determine to what extent we should bear that cost.

In our prior paper, we argued that an ideal consumption tax is preferable to an ideal income tax. Shaviro raises additional arguments that we did not consider in our prior work. After considering these arguments, we conclude that an ideal consumption tax is still preferable.

Perhaps the lesson from Shaviro's article is that one can never be sure because models are, after all, just models. We must always continue to develop better models, to better understand the world and how the tax system can improve market outcomes. If this is his message, we agree.

24. We were not alone in framing the debate this way. See, e.g., Alvin Warren, *Would a Consumption Tax Be Fairer than an Income Tax?*, 89 *YALE L.J.* 1081 (1980).