

The Case for a Progressive Tax: From Basic Research to Policy Recommendations[†]

Peter Diamond and Emmanuel Saez

The fair distribution of the tax burden has long been a central issue in policy-making. A large academic literature has developed models of optimal tax theory to cast light on the problem of optimal tax progressivity. In this paper, we explore the path from basic research results in optimal tax theory to formulating policy recommendations.

Models in optimal tax theory typically posit that the tax system should maximize a social welfare function subject to a government budget constraint, taking into account that individuals respond to taxes and transfers. Social welfare is larger when resources are more equally distributed, but redistributive taxes and transfers can negatively affect incentives to work, save, and earn income in the first place. This creates the classical trade-off between equity and efficiency which is at the core of the optimal income tax problem. In general, optimal tax analyses maximize social welfare as a function of individual utilities—the sum of utilities in the utilitarian case. The marginal weight for a given person in the social welfare function measures the value of an additional dollar of consumption expressed in terms of public funds. Such welfare weights depend on the level of redistribution and are decreasing with income whenever society values more equality of income. Therefore, optimal income tax theory is first a normative theory that shows how a social welfare objective combines with constraints arising from limits on resources and behavioral responses to taxation in order to derive specific

■ *Peter Diamond is Professor Emeritus of Economics, Massachusetts Institute of Technology, Cambridge Massachusetts. Emmanuel Saez is Professor of Economics, University of California, Berkeley, California. Their e-mail addresses are <pdiamond@mit.edu> and <saez@econ.berkeley.edu>, respectively.*

[†] There is an Appendix at the end of this article. To access an additional online Appendix, visit <http://www.aeaweb.org/articles.php?doi=10.1257/jep.25.4.165>.
doi=10.1257/jep.25.4.165

tax policy recommendations. In addition, optimal income tax theory can be used to evaluate current policies and suggest avenues for reform. Understanding what would be good policy, if implemented, is a key step in making policy recommendations.

When done well, moving from mathematical results, theorems, or calculated examples to policy recommendations is a subtle process. The nature of a model is to be a limited picture of reality. This has two implications. First, a model may be good for one question and bad for another, depending on the robustness of the answers to the inaccuracies of the model, which will naturally vary with the question. Second, tractability concerns imply that simultaneous consideration of multiple models is appropriate since different aspects of reality can be usefully highlighted in different models; hence our reliance on trying to draw inferences simultaneously from multiple models.

In our view, a theoretical result can be fruitfully used as part of forming a policy recommendation only if three conditions are met. First, the result should be based on an economic mechanism that is empirically relevant and first order to the problem at hand. Second, the result should be reasonably robust to changes in the modeling assumptions. In particular, people have very heterogeneous tastes, and there are many departures from the rational model, especially in the realm of intertemporal choice. Therefore, we should view with suspicion results that depend critically on very strong homogeneity or rationality assumptions. Deriving optimal tax formulas as a function of a few empirically estimable “sufficient statistics” is a natural way to approach those first two conditions. Third, the tax policy prescription needs to be implementable—that is, the tax policy needs to be socially acceptable and not too complex relative to the modeling of tax administration and individual responses to tax law. By socially acceptable, we do not mean to limit the choice to currently politically plausible policy options. Rather, we mean there should not be very widely held normative views that make such policies seem implausible and inappropriate at pretty much all times. For example, a policy prescription such as taxing height (Mankiw and Weinzierl, 2010) is obviously not socially acceptable because it violates certain horizontal equity concerns that do not appear in basic models. The complexity constraint can also be an issue when optimal taxes depend in a complex way on the full history of earnings and consumption, as in some recent path-breaking papers on optimal dynamic taxation.

We obtain three policy recommendations from basic research that we believe can satisfy these three criteria reasonably well. First, very high earners should be subject to high and rising marginal tax rates on earnings. In particular, we discuss why the famous zero marginal tax rate at the top of the earnings distribution is not policy relevant. Second, the earnings of low-income families should be subsidized, and those subsidies should then be phased out with high implicit marginal tax rates. This result follows because labor supply responses of low earners are concentrated along the margin of whether to participate in labor markets at all (the extensive as opposed to the intensive margin). These two results combined imply that the optimal profile of transfers and taxes is highly nonlinear and cannot be well approximated by a flat tax along with lump sum “demogrants.” Third, we argue that capital income should be taxed. We will review certain theoretical results—in particular,

those of Atkinson and Stiglitz (1976), Chamley (1986), and Judd (1985)—implying no capital income taxes and argue that these findings are not robust enough to be policy relevant. In the end, persuasive arguments for taxing capital income are that there are difficulties in practice in distinguishing between capital and labor incomes, that borrowing constraints make full reliance on labor taxes less efficient, and that savings rates are heterogeneous.

The remainder of the paper is organized as follows: First, we consider the taxation of very high earners, second, the taxation of low earners, and third, the taxation of capital income. We conclude with a discussion of methodology, contrasting optimal tax and mechanism design (“new dynamic public finance”) approaches. In an appendix, we contrast our lessons from optimal tax theory with those of Mankiw, Weinzierl, and Yagan (2009), recently published in this journal.

Recommendation 1: Very high earnings should be subject to rising marginal rates and higher rates than current U.S. policy for top earners.

The share of total income going to the top 1 percent of income earners (those with annual income above about \$400,000 in 2007) has increased dramatically from 9 percent in 1970 to 23.5 percent in 2007, the highest level on record since 1928 and much higher than in European countries or Japan today (Piketty and Saez, 2003; Atkinson, Piketty, and Saez, 2011). Although the average federal individual income tax rate of top percentile tax filers was 22.4 percent, the top percentile paid 40.4 percent of total federal individual income taxes in 2007 (IRS, 2009a). Therefore, the taxation of very high earners is a central aspect of the tax policy debate not only for equity reasons but also for revenue raising. For example, setting aside behavioral responses for a moment, increasing the average federal income tax rate on the top percentile from 22.4 percent (as of 2007) to 29.4 percent would raise revenue by 1 percentage point of GDP.¹ Indeed, even increasing the average federal income tax rate of the top percentile to 43.5 percent, which would be sufficient to raise revenue by 3 percentage points of GDP, would still leave the after-tax income share of the top percentile more than twice as high as in 1970.² Of course, increasing upper income tax rates can discourage economic activity through behavioral responses, and hence

¹ In 2007, the top percentile of income earners paid \$450 billion in federal individual taxes (IRS, 2009a), or 3.2 percent of the \$14,078 billion in GDP for 2007. Hence, increasing the average tax rate on the top percentile from 22.4 to 29.4 percent would raise \$141 billion or 1 percent of GDP.

² The average federal individual tax rate paid by the top percentile was 25.7 percent in 1970 (Piketty and Saez, 2007) and 22.4 percent in 2007 (IRS, 2009a). The overall average federal individual tax rate was 12.5 percent in 1970 and 12.7 percent in 2007. The pre-tax income share for the top percentile of tax filers was 9 percent in 1970 and 23.5 percent in 2007. Hence, the top 1 percent after-tax income share in 1970 was 7.6 percent = $9\% \times (1 - .257)/(1 - .125)$, and in 2007 it was 20.9 percent = $23.5\% \times (1 - .224)/(1 - .127)$ and, with a tax rate of 43.5 percent on the top percentile (which would increase the average tax rate to 17.7 percent), would have been 16.1 percent = $23.5\% \times (1 - .435)/(1 - .177)$.

potentially reduce tax collections, creating the standard equity-efficiency trade-off discussed in the introduction.

The Optimal Top Marginal Tax Rate

For the U.S. economy, the current top income marginal tax rate on earnings is about 42.5 percent,³ combining the top federal marginal income tax bracket of 35 percent with the Medicare tax and average state taxes on income and sales.⁴ As shown in Saez (2001), the optimal top marginal tax rate is straightforward to derive. Denote the tax rate in the top bracket by τ . Figure 1 shows how the optimal tax rate is derived. The horizontal axis of the figure shows pre-tax income, while the vertical axis shows disposable income. The original top tax bracket is shown by the solid line. As depicted, consider a tax reform which increases τ by $\Delta\tau$ above the income level z^* . To evaluate this change we need to consider the effects on revenue and social welfare. Ignoring behavioral responses at first, this reform mechanically raises additional revenue by an amount equal to the change in the tax rate ($\Delta\tau$) multiplied by the number of people to whom the higher rate applies (N^*) multiplied by the amount by which the average income of this group (z_m) is above the cut-off income level (z^*) so that the additional revenue is $\Delta\tau N^*[z_m - z^*]$. As we shall see, the top tail of the income distribution is closely approximated by a Pareto distribution characterized by a power law density of the form C/z^{1+a} where $a > 1$ is the Pareto parameter. Such distributions have the key property that the ratio z_m/z^* is the same for all z^* in the top tail and equal to $a/(a-1)$. For the U.S. economy, the cutoff for the top percentile of tax filers is approximately \$400,000, and the average income for this group is approximately \$1.2 million, so that $z_m/z^* = 3$ and hence $a = 1.5$.


Raising the tax rate on the top percentile obviously reduces the utility of high-income tax filers. If we denote by g the social marginal value of \$1 of consumption for top income earners (measured relative to government revenue), the direct welfare cost is g multiplied by the change in tax revenue collected.⁵ Because the government values redistribution, the social marginal value of consumption for top-bracket tax filers is small relative to that of the average person in the economy, and so g is small and as a first approximation can be ignored. A utilitarian social welfare criterion with marginal utility of consumption declining to zero, the most commonly

³This top marginal tax rate is much higher than the current average tax rate among top 1 percent earners mentioned above because of deductions and especially lower tax rates that apply to realized capital gains.

⁴The top tax rate τ is 42.5 percent for ordinary labor income when combining the top federal individual tax rate of 35 percent, uncapped Medicare taxes of 2.9 percent, and an average combined state top income tax rate of 5.86 percent and average sales tax rate of 2.32 percent. The average across states is computed using state weights equal to the fraction of filers with adjusted gross income above \$200,000 that reside in the state as of 2007 (IRS, 2009a). The 2.32 percent average sales tax rate is estimated as 40 percent of the average nominal sales tax rate across states (as the average sales tax base is about 40 percent of total personal consumption.) As the 1.45 percent employer Medicare tax is deductible for both federal and state income taxes, and state income taxes are deductible for federal income taxes, we have $((1 - .35) \times (1 - .0586) - .0145) / (1.0145 \times 1.0232) = .575$, and hence $\tau = 42.5$ percent.

⁵Formally, g is the weighted average of social marginal weights on top earners, with weights proportional to income in the top bracket.

Figure 1
Optimal Top Tax Rate Derivation



Source: The authors.

Notes: The figure depicts the derivation of the optimal top tax rate $\tau^* = 1 / (1 + ae)$ by considering a small reform around the optimum which increases the top marginal tax rate τ by $\Delta\tau$ above z^* . A taxpayer with income z mechanically pays $\Delta\tau[z - z^*]$ extra taxes but, by definition of the elasticity e of earnings with respect to the net-of-tax rate $1 - \tau$, also reduces his income by $\Delta z = ez\Delta\tau / (1 - \tau)$ leading to a loss in tax revenue equal to $\Delta\tau ez\tau / (1 - \tau)$. Summing across all top bracket taxpayers and denoting by z_m the average income above z^* and $a = z_m / (z_m - z^*)$, we obtain the revenue maximizing tax rate $\tau^* = 1 / (1 + ae)$. This is the optimum tax rate when the government sets zero marginal welfare weights on top income earners.

used specification in optimal tax models, has this implication. For example, if the social value of utility is logarithmic in consumption, then social marginal welfare weights are inversely proportional to consumption. In that case, the social marginal utility at the \$1,364,000 average income of the top 1 percent in 2007 (Piketty and Saez, 2003) is only 3.9 percent of the social marginal utility of the median family, with income \$52,700 (U.S. Census Bureau, 2009).

Behavioral responses can be captured by the elasticity e of reported income with respect to the net-of-tax rate $1 - \tau$. By definition, e measures the percent increase in average reported income z_m when the net-of-tax rate increases by 1 percent.⁶ At the optimum, the marginal gain from increasing tax revenue with no behavioral response and the marginal loss from the behavioral reaction must be equal to each

⁶ Formally, this elasticity is an income-weighted average of the individual elasticities across the N^* top bracket tax filers. It is also a mix of income and substitution effects as the reform creates both income and substitution effects in the top bracket. Saez (2001) provides an exact decomposition.

other. Ignoring the social value of marginal consumption of top earners, the optimal top tax rate τ^* is given by the formula

$$\tau^* = 1/(1 + ae).$$

The optimal top tax rate τ^* is the tax rate that maximizes tax revenue from top bracket taxpayers.⁷ Since the goal of the marginal rates on very high incomes is to get revenue in order to hold down taxes on lower earners, this equation does not depend on the total revenue needs of the government. Any top tax rate above τ^* would be (second-best) Pareto inefficient as reducing tax rates at the top would both increase tax revenue and the welfare of top earners.

An increase in the marginal tax rate only at a single income level in the upper tail increases the deadweight burden (decreases revenue because of reduced earnings) at that income level but raises revenue from all those with higher earnings without altering their marginal tax rates. The optimal tax rate balances these two effects—the increased deadweight burden at the income level and the increased revenue from all higher levels. τ^* is decreasing with the elasticity e (which affects the deadweight burden) and the Pareto parameter a , which measures the thinness of the top of the income distribution and so the ratio of those above a tax level to the income of those at the tax level.

The solid line in Figure 2 depicts the empirical ratio $a = z_m/(z_m - z^*)$ with z^* ranging from \$0 to \$1,000,000 in annual income using U.S. tax return micro-data for 2005. We use “adjusted gross income” from tax returns as our income definition. The central finding is that a is extremely stable for z^* above \$300,000 (and around 1.5). The excellent Pareto fit of the top tail of the distribution has been well known for over a century since the pioneering work of Pareto (1896) and verified in many countries and many periods, as summarized in Atkinson, Piketty, and Saez (2011).


If we assume that the elasticity e is roughly constant across earners at the top of the distribution, the formula $\tau = 1/(1 + ae)$ shows that the optimal top tax rate is independent of z^* within the top tail (and is also the asymptotic optimal marginal tax rate coming out of the standard nonlinear optimal tax model of Mirrlees, 1971). That is, the optimal marginal tax rate is approximately the same over the range of very high incomes where the distribution is Pareto and the marginal social weight on consumption is small.⁸ This makes the optimal tax formula quite general and useful.

⁷ If a positive social weight $g > 0$ is set on top earners’ marginal consumption, then the optimal rate is $\tau = (1 - g)/(1 - g + ae) < \tau^*$. With plausible weights that are small relative to the weight on an average earner, the optimal tax does not change much.

⁸ If the elasticity e does not vary by income level, then the Pareto parameter a does not vary with τ . If the elasticity varies by income, the Pareto parameter a might depend on the top tax rate τ . The formula $\tau^* = 1/(1 + ae)$ is still valid in that case, but determining τ^* would require knowing how a varies with τ .

Figure 2

Empirical Pareto Coefficients in the United States, 2005



Source: The authors using public use tax return data.

Notes: The figure depicts in solid line the ratio $a = z_m / (z_m - z^*)$ with z^* ranging from \$0 to \$1,000,000 annual income and z_m the average income above z^* using U.S. tax return micro data for 2005. Income is defined as Adjusted Gross Income reported on tax returns and is expressed in current 2005 dollars. Vertical lines depict the 90th percentile (\$99,200) and 99th percentile (\$350,500) nominal thresholds as of 2005. The ratio a is equal to one at $z^* = 0$, and is almost constant above the 99th percentile and slightly below 1.5, showing that the top of the distribution is extremely well approximated by a Pareto distribution for purposes of implementing the optimal top tax rate formula $\tau^* = 1 / (1 + ae)$. Denoting by $h(z)$ the density and by $H(z)$ the cumulative distribution function of the income distribution, the figure also displays in dotted line the ratio $\alpha(z^*) = z^* h(z^*) / (1 - H(z^*))$, which is also approximately constant, around 1.5, above the top percentile. A decreasing (or constant) $\alpha(z)$ combined with a decreasing $G(z)$ and a constant $e(z)$ implies that the optimal marginal tax rate $T'(z) = [1 - G(z)] / [1 - G(z) + \alpha(z) e(z)]$ increases with z .

The Tax Elasticity of Top Incomes

The key remaining empirical ingredient to implement the formula for the optimal tax rate is the elasticity e of top incomes with respect to the net-of-tax rate. With the Pareto parameter $a = 1.5$ if $e = .25$, a mid-range estimate from the empirical literature, then $\tau^* = 1 / (1 + 1.5 \times .25) = 73$ percent, substantially higher than the current 42.5 percent top U.S. marginal tax rate (combining all taxes).⁹

⁹ Using g^* of .04, the optimal tax rate decreases by about 1 percentage point.

The current rate, $\tau = 42.5$ percent, would be optimal only if the elasticity e were extremely high, equal to 0.9.¹⁰

Before turning to empirical estimates, we review some of the interpretation issues that arise when moving beyond the simplest version of the Mirrlees (1971) model. In the Mirrlees model, there is a single tax on each individual. With many taxes, for example, in many periods, the key measure is the response of the present discounted value of all taxes, not the response of revenue in a single year. This observation matters given significant control by some people over the timing of taxes and over the forms in which income might be received. Also, because the basic Mirrlees model has no tax-deductible charitable giving, a tax-induced change in taxable income involves only distortions from reduced earnings. However, when an increase in marginal tax rates leads to an increase in charitable giving, the gain to the recipients needs to be incorporated in the efficiency measure (Saez, 2004). Other tax deductions are more difficult to consider. In the Mirrlees model, compensation equals the marginal product. In bargaining settings or with asymmetric information, people may not receive their marginal products. Thus, effort is responding to a price that is higher or lower than marginal product, and the tax rate itself may affect the gap between compensation and marginal product.

The large literature using tax reforms to estimate the elasticity relevant for the optimal tax formula has focused primarily on the response of reported income, either “adjusted gross income” or “taxable income,” to net-of-tax rates. Saez, Slemrod, and Giertz (forthcoming) offer a recent survey, while Slemrod (2000) looks at studies focusing on the rich. The behavioral elasticity is due to real economic responses such as labor supply, business creation, or savings decisions, but also tax avoidance and evasion responses. A number of studies have shown large and quick responses of reported incomes along the tax avoidance margin at the top of the distribution, but no compelling study to date has shown substantial responses along the real economic responses margin among top earners. For example, in the United States, realized capital gains surged in 1986 in anticipation of the increase in the capital gains tax rate after the Tax Reform Act of 1986 (Auerbach, 1988). Similarly, exercises of stock options surged in 1992 before the 1993 top rate increase took place (Goolsbee, 2000). The Tax Reform Act of 1986 also led to a shift from corporate to individual income as it became more advantageous to be organized as a business taxed solely at the individual level rather than as a corporation taxed first at the corporate level (Slemrod, 1996; Gordon and Slemrod, 2000). The paper Gruber and Saez (2002) is often cited for its substantial taxable income elasticity estimate ($e = 0.57$) at the top of the distribution. However, its authors also found a small elasticity ($e = 0.17$) for income before any deductions, even at the top of the distribution (Table 9, p. 24).

When a tax system offers tax avoidance or evasion opportunities, the tax base in a given year is quite sensitive to tax rates, so the elasticity e is large, and the optimal top tax rate is correspondingly low. Two important qualifications must be made.

¹⁰ Alternatively, if the elasticity is $e = .25$, then $\tau = 42.5$ percent is optimal only if the marginal consumption of very high-income earners is highly valued, with $g = .72$.

First, as mentioned above, many of the tax avoidance channels such as retiming or income shifting produce changes in tax revenue in other periods or other tax bases—called “tax externalities”—and hence do not decrease the optimal tax rate. Saez, Slemrod, and Giertz (forthcoming) provide formulas showing how the optimal top tax rate should be modified in such cases. Second, and most important, the tax avoidance or evasion component of the elasticity e is not an immutable parameter and can be reduced through base broadening and tax enforcement (Slemrod and Kopczuk, 2002; Kopczuk, 2005). Thus, the distinction between real responses and tax avoidance responses is critical for tax policy. As an illustration using the different elasticity estimates of Gruber and Saez (2002) for high-income earners mentioned above, the optimal top tax rate using the current taxable income base (and ignoring tax externalities) would be $\tau^* = 1/(1 + 1.5 \times 0.57) = 54$ percent, while the optimal top tax rate using a broader income base with no deductions would be $\tau^* = 1/(1 + 1.5 \times 0.17) = 80$ percent. Taking as fixed state and payroll tax rates, such rates correspond to top federal income tax rates equal to 48 and 76 percent, respectively. Although considerable uncertainty remains in the estimation of the long-run behavioral responses to top tax rates (Saez, Slemrod, and Giertz, forthcoming), the elasticity $e = 0.57$ is a conservative upper bound estimate of the distortion of top U.S. tax rates. Therefore, the case for higher rates at the top appears robust in the context of this model.

Link with the Zero Top Rate Result

Formally, z_m/z^* reaches 1 when z^* reaches the level of income of the single highest income earner, in which case $a = z_m/(z_m - z^*)$ is infinite, and indeed $\tau^* = 1/(1 + ae) = 0$, which is the famous zero top rate result first demonstrated by Sadka (1976) and Seade (1977). However, notice that this result applies only to the very top income earner; its lack of wider applicability can be verified empirically using tax data.¹¹ If one makes the reasonable assumption that the level of top earnings is not known in advance, and instead consider having potential earnings drawn randomly from an underlying Pareto distribution then (as we show in the Appendix available online with this paper at <http://e-jep.org>), with the budget constraint satisfied in expectation, the formula, $\tau^* = 1/(1 + ae)$, remains the natural optimum tax rate. This finding implies that the zero top rate result and its corollary that marginal tax rates should decline at the top have no policy relevance, a view that we believe is widely shared among public finance economists.¹²

¹¹ If, for example, the second-highest income is only one-half of the highest earner then $z_m/z^* = 2$ (and hence $a = 2$) when z^* is just above the second-highest earner, so that convergence of z_m/z^* to one really happens only between the top and second-highest earner. The IRS publishes statistics on the top 400 taxpayers (IRS, 2009b). In 2007, the threshold to be a top 400 taxpayer was \$138.8m and the average income of top 400 taxpayers was \$344.8m so that $a = 1.67$ at $z^* = \$138.8m$, very close to the value of 1.5 at the top percentile threshold, and still very far from the infinite value it takes at the very top income.

¹² With a known finite distribution, the marginal tax rate at the top is zero, but the average tax rate between the highest and second-highest earners is so large that highest earner gets no additional utility from being more productive than the next-highest earner.

Should Marginal Tax Rates Rise with Income?

Assuming away income effects on labor supply, the optimal marginal tax rate formula at any income level (applying to the combination of all taxes) takes a form that can be expressed directly as a function of the income distribution as follows (Diamond, 1998):

$$T'(z) = [1 - G(z)]/[1 - G(z) + \alpha(z) e(z)]$$

where $e(z)$ is the elasticity of incomes with respect to the net-of-tax rate at income level z , $G(z)$ is the average social marginal welfare weight across individuals with income above z , and $\alpha(z) = (zh(z))/(1 - H(z))$ with $h(z)$ the density of taxpayers at income level z and $H(z)$ the fraction of individuals with income below z .¹³ The expression $\alpha(z)$ reflects the ratio of the total income of those affected by the marginal tax rate at z relative to the numbers of people at higher income levels. A derivation of the optimal formula is presented in an appendix available with this paper at <http://e-jep.org>.

For Pareto distributions, $\alpha(z)$ is constant and equal to the Pareto parameter. However, the empirical U.S. income distribution is not a Pareto distribution at lower income levels. The $\alpha(z)$ term is depicted in dotted line on Figure 2 for the empirical 2005 U.S. income distribution. It is inversely U-shaped, reaching a maximum of 2.17 at $z = \$135,000$, then decreasing and staying approximately constant around 1.5 above $z = \$400,000$. Because social welfare weights are lower for higher incomes, $G(z)$ decreases with z . Therefore, assuming a constant elasticity e across income groups, the formula implies that the optimal marginal tax rates should increase with income in the upper part of the distribution. This result was theoretically established by Diamond (1998) and confirmed by all subsequent simulations that use a Pareto distribution at the top as in Saez (2001) or Mankiw, Weinzierl, and Yagan (2009). Quantitatively, this increase is substantial. For example, assuming again an elasticity $e = .25$ and that $G(z) = 0.5$ at $z = \$100,000$, corresponding to the top decile threshold where $\alpha = 2.05$, we would have $T' = 49$ percent at this income, well below the value of 73 percent for the top percentile as calculated above.

In the current tax system with many tax avoidance opportunities at the higher end, as discussed above, the elasticity e is likely to be higher for top earners than for middle incomes, possibly leading to decreasing marginal tax rates at the top (Gruber and Saez, 2002). However, the natural policy response should be to close tax avoidance opportunities, in which case the assumption of constant elasticities might be a reasonable benchmark.

¹³ Technically, Saez (2001) shows that $h(z)$ is the density of incomes when the nonlinear tax system is linearized at z . Saez (2001) also shows that a similar but more complex formula can be obtained with income effects that is quantitatively close to the equation above.

Additional Considerations

To some readers, proposing marginal income tax rates on the top percentile of earners, along with a broadened tax base, in a range from 48 to 76 percent may seem implausibly high. One way to judge how seriously to take such numbers is to consider whether elements left out in the derivation push for a significantly different answer. Two key omitted elements are the presence of capital income and a longer-run dynamic perspective.

Does the presence of capital income mean that earnings should be taxed significantly differently? When we discuss taxation of capital income in a later section, we note that the ability to convert some labor income into capital income is a reason for limiting the difference between tax rates on the two types of income—that is, an argument for taxing capital income. Plausibly, it is also an argument for a somewhat lower labor income tax, assuming that labor income should be taxed more heavily than capital income.

Perhaps most critically, does an estimate based on a single period model still apply when recognizing that people earn and pay income taxes year after year? First, earlier decisions such as education and career choices affect later earnings opportunities. It is conceivable that a more progressive tax system could reduce incentives to accumulate human capital in the first place. The logic of the equity–efficiency trade-off would still carry through, but the elasticity ϵ should reflect not only short-run labor supply responses but also long-run responses through education and career choices. While there is a sizable multiperiod optimal tax literature using life-cycle models and generating insights, we unfortunately have little compelling empirical evidence to assess whether taxes affect earnings through those long-run channels.

Second, there is significant uncertainty in future earnings. Such uncertainty gives an insurance role for earnings taxation and, as we shall see, also has consequences for the taxation of savings.¹⁴ However, the applicability of results for policy seems unclear to us.

Recommendation 2: Tax (and transfer) policy toward low earners should include subsidization of earnings and should phase out the subsidization at a relatively high rate.

Transfers are naturally integrated with taxes in an optimal tax problem. Such transfers often take the general form of a maximum benefit for those with no income, which is phased out at high rates as earnings increase. For example, in the United States, TANF (Temporary Aid to Needy Families) and SNAP (Supplemental Nutrition Assistance Program, formerly known as Food Stamps) operate in this way. A growing fraction of means-tested transfers is now administered through refundable tax credits

¹⁴ The “new dynamic public finance” analyzes such settings using mechanism design. The new dynamic public finance has made recent progress on the optimal labor income taxation in the dynamic context. See Farhi and Werning (2011).

such as the EITC (Earned Income Tax Credit) or the Child Tax Credit. Such programs are typically first phased in and then phased out with earnings so that benefits are concentrated on low-income working families instead of those with no earnings. Many studies have found compelling evidence of substantial labor supply responses to transfers along the extensive margin of whether or not to work. For example, the EITC expansions have encouraged labor force participation of U.S. single mothers (Meyer, 2010). However, there is much less compelling evidence of behavioral responses along the intensive margin—that is, hours of work on the job—for lower-income earners. As we shall see, these facts play a critical role in the optimal profile of transfers.

Intensive Elasticities

In the Mirrlees (1971) model, behavioral responses take place only through the intensive margin of the number of hours worked. In that context, it is optimal to provide income to those with no earnings, which is then phased out with earnings, possibly at a high rate—which acts as an implicit tax (see the online appendix with this paper at <http://e-jep.org>) for a derivation). The intuition is that a high phase-out rate allows the government to target transfers to the most disadvantaged families. A high phase-out rate does reduce earnings for low-income families, because they reduce hours worked. However, because earnings of those in the phase-out are small to start with, this elasticity applies to a low income base. Therefore, increasing the maximum benefit (to those with no earnings) and increasing the phase-out rate is desirable for redistribution, and the behavioral responses create modest fiscal costs relative to the redistributive gains, as long as the phase-out rate is not too high. Hence, the Mirrlees model of optimal income taxation generates a traditional welfare program where benefits are concentrated on non-earners with high phase-out rates on low-income workers.

Extensive Elasticities

However, the optimality of traditional welfare with a high phase-out rate depends critically on the absence of labor supply responses along the extensive margin, that is, whether or not to work. If labor supply responses are concentrated along the extensive margin, then it is optimal to give higher transfers to low-income workers than nonworkers, which amounts to a negative phase-out rate, as with the current Earned Income Tax Credit (Diamond, 1980; Saez, 2002a).

To see this, suppose the government starts from a transfer scheme with a positive phase-out rate—that is, the transfer is gradually reduced as earned income rises—and introduces a small additional in-work benefit for low-income workers. Ignoring behavioral responses, such a reform is desirable if the government values redistribution to low-income earners. If behavioral responses are solely along the extensive margin, this reform induces some nonworkers to start working to take advantage of the in-work benefit. However, because we start from a situation with a positive phase-out rate, this behavioral response increases tax revenue as low-income workers still end up receiving a smaller transfer than nonworkers. Hence, with the availability of a desirable redistribution and a gain in revenue from the

behavioral response, a positive phase-out rate is not optimal (we provide a more detailed graphical derivation in the online appendix.)

In practice, both extensive and intensive elasticities are present. An intensive margin response would induce slightly higher earners to reduce labor supply to take advantage of the in-work benefit, reducing tax revenue. Therefore, the government has to trade-off the two effects. If, as empirical studies show, the extensive elasticity of choosing whether to participate in the labor market is large for those with low incomes relative to the intensive elasticity of choosing how many hours to work, initially low (or even negative) phase-out rates combined with high positive phase-out rates further up the distribution would be the optimal profile.

In recent decades in most high-income countries, a concern arose that traditional welfare programs overly discouraged work, and there has been a marked shift toward lowering the marginal tax rate at the bottom through a combination of: a) introduction and then expansion of in-work benefits such as the Earned Income Tax Credit in the United States; b) reduction of the statutory phase-out rates in transfer programs for earned income, as under the U.S. welfare reform; and c) reduction of payroll taxes for low-income earners, as in the recent U.S. Making Work Pay credit. Those reforms are consistent with the logic of optimal taxation we have outlined, as they both encourage labor force participation and provide transfers to low-income workers, seen as a deserving group.

Recommendation 3: Capital income should be taxed.

With the standard model for static labor supply decisions, the simplicity of a one-period model and the extensive empirical literature on labor supply elasticities, it is possible to provide useful quantitative analysis of optimal marginal tax rates. In contrast, the literature on saving behavior sees a wide variety of basic behaviors, more widely varying elasticity estimates, and a complexity that comes from the importance of the future for decisions affected by capital income taxation. Thus, we limit our discussion to a single qualitative recommendation: capital income should be subject to significant taxation. This conclusion is important in light of repeated calls for not taxing capital income.

Academic arguments against capital income taxation typically draw on one or both of two theoretical analyses: (1) the theorem that the optimum has no asymptotic long-run taxation of capital income in Chamley (1986) and Judd (1985); and (2) the theorem that the optimum has no taxation of capital income in Atkinson and Stiglitz (1976).¹⁵ For lengthier discussion of these arguments, see Banks and

¹⁵ The aggregate efficiency theorem in Diamond and Mirrlees (1971) is sometimes cited as support for not taxing capital income. Taxes on transactions between households and firms (that do not vary with the particular firm) do not interfere with production efficiency. While taxing all capital income of households will generally change the level of savings, and so investment, it does not move the economy inside the production possibility frontier. Thus, the aggregate efficiency theorem, that the optimum is on the production frontier, has *no* direct implications relative to taxing the capital income of households.

Diamond (2010). We address each of these in turn. We then address four arguments for positive taxation of capital income: the difficulty of distinguishing between capital and labor incomes; the positive correlation between earnings opportunities and savings propensities; the role of capital income taxes in easing the tax burden on those who are borrowing constrained; and the role of discouraging savings in encouraging later labor supply in the presence of uncertain future wage rates.

Chamley and Judd

In the models analyzed in Chamley (1986) and Judd (1985), with infinitely-lived agents, an asymptotically zero tax on capital income is optimal. In order to evaluate the relevance of this result for policy purposes, one needs to understand the logic of the result, and particularly its robustness to key assumptions. As pointed out in Judd (1999), the logic for the result is straightforward. A constant capital income tax rate creates a growing tax wedge between current consumption and future consumption as the horizon grows. With interest rate r and no capital income taxes, a dollar today is worth $(1 + r)^T$ after T years. If an investor is subject to an annual tax at rate τ on capital income, then the investor can convert one unit of consumption today into only $(1 + (1 - \tau)r)^T$ units after T years. Hence, the tax wedge $1 - (1 + (1 - \tau)r)^T / (1 + r)^T$ grows with T .¹⁶ For example, with $r = .05$ and $\tau = 30$ percent, the tax wedge is a modest 13.4 percent when $T = 10$, but is a substantial 43.8 percent when $T = 40$. In order to avoid tax compounding that grows without limit as the horizon extends, the optimal average rate must go to zero, although no individual tax rate needs to be zero.

Therefore, the result relies critically on the assumption that individuals make consistent rational decisions about savings behavior across very long horizons, as in the standard intertemporal model. When agents have long horizons, modeling their current decision making using an infinite horizon model can be mathematically more tractable while doing little violence to conclusions that relate to current behavior. In contrast, substituting an infinite-horizon decisionmaker for a sequence of finite-horizon decisionmakers can make a large difference in the asymptotic position of the economy. In an overlapping generations model with no bequests and so no dynastic linkage, the optimal capital income tax is generally not zero, even in the long-run (Diamond, 1973; Atkinson and Sandmo, 1980). Thus, the strong asymptotic zero tax result of Chamley (1986) and Judd (1985) requires that rational intertemporal decision making not only holds for entire lifetimes, but extends across dynasties. Both assumptions have been heavily challenged in the empirical literature.

First, the recent behavioral economics literature has cast much doubt on the standard model of intertemporal decision making for a significant fraction of the population. A growing body of empirical work shows that savings decisions are

¹⁶ While interest income and dividends are taxed in this compounding way, the same is not true for capital gains that are taxed on a realization basis. Nor is it true for tax-favored retirement saving, such as IRA or 401(k) accounts.

heavily influenced by psychological elements (such as self-control) or minor transaction costs (like the default effects in employer-sponsored 401(k) plans).

Second, empirical analyses of gifts and bequests, while clearly showing concerns about heirs, are not supportive of the rigorous version of the dynasty model required for the Chamley–Judd result. People leave bequests for many reasons: unintended bequests due to lack of annuitization or to love of wealth accumulation *per se*; intended bequests arising out of bargaining with heirs, “warm glow” preferences, or altruism. The optimal tax treatment of bequests depends heavily on the mechanism behind bequests (Cremer and Pestieau, 2006, provide a survey). For example, unintended bequests should be taxed heavily because they do not affect donors and inheritances induce donees to work less through income effects. In contrast, if bequests are altruistic and the social planner takes into account both parents’ and kids’ welfare (as opposed to parents’ only in the traditional dynastic model), then it can be desirable to subsidize bequests, especially among the poor (Farhi and Werning, 2010). The dynastic model reflects special forms of both altruism and the social welfare function and hence likely captures only one aspect of bequest behavior.

Rejecting the policy relevance of the zero taxation result does not remove the relevance of the compounding of capital income taxes noted above. This concern adds support to the case for tax-favored retirement savings accounts coming from concern of inadequate savings by some (because saving for retirement involves long horizons). Conversely, the presence of such accounts supports higher taxation of capital income than without such a savings option.

Another straightforward conclusion coming out of the Chamley–Judd model is that it is better to tax existing wealth rather than future capital income because a tax on current wealth is lump-sum, while a tax on future capital income distorts intertemporal choices. While the asymptotic zero capital income tax result has drawn great attention, the initial result is largely ignored for policy purposes, although the same perspective, clearly stated in the literature, lies behind arguments for switching from income taxation to consumption taxation in overlapping generation models as a way to transfer wealth away from older cohorts at the time of tax implementation with little in the way of distorting incentives (Auerbach, Kotlikoff, and Skinner, 1983). However, taxing initial wealth as much as the available tax tools allow (whether as a wealth tax or a capital income tax) strains the relevance of the assumption that the government is committed to a policy that this taxation of wealth will not be repeated. Without a credible commitment (which may not be possible), confiscatory wealth taxation would adversely affect saving behavior and have serious efficiency costs because of concerns that such taxation will return. In short, we do not believe that the modeling assumptions behind the Chamley and Judd results are strong enough to support drawing policy lessons about the appropriate taxation of capital.

Atkinson and Stiglitz

In a two-period model with one period of work, the Atkinson–Stiglitz theorem (1976) states that when the available tax tools include nonlinear earnings taxes, differential taxation of first- and second-period consumption is not optimal if two

key conditions are satisfied: 1) all consumers have preferences that are separable between consumption and labor; and 2) all consumers have the same subutility function of consumption. The underlying logic behind the result starts with the observation that the incentive to earn comes from the utility achievable from consumption purchases with after-tax earnings. With separable preferences and the same subutilities for everyone, differential consumption taxation cannot accomplish any distinction among those with different earnings abilities beyond what is already accomplishable by the earnings tax, but would have an added efficiency cost from distorting spending choices. Thus the use of distorting taxes on consumption is a more costly way of providing the incentives for the “optimal” earnings pattern in equilibrium.¹⁷ In this two-period model, differential consumption taxation is the same thing as capital taxation.

While the Atkinson–Stiglitz theorem requires an absence of a systematic pattern between earnings abilities and savings propensities, there appears to be a positive correlation between labor skill level (wage rate) and savings propensities. With this plausible assumption, implying that those with higher earnings abilities save more out of any given income, then taxation of saving helps with the equity–efficiency tradeoff by being a source of indirect evidence about who has higher earnings abilities and thus contributes to more efficient redistributive taxation (Saez, 2002b).¹⁸

The dimensionality of worker types (relative to tax tools) matters in models of capital income taxation. This point can be brought out by contrasting the analysis of the taxation of capital income in a model with two types of workers in Diamond (2003) with that in a model with four types of workers in Diamond and Spinnewijn (forthcoming). Both papers use two-period models and assume additive preferences, with workers varying in both skill and discount factor. With two types, in the optimum, the high earner has no marginal taxes. In contrast, with more types of workers and diverse discount rates at each earnings level, the optimum has taxation of savings of high earners and subsidization of savings of low earners. The underlying logic comes from the incentive compatibility constraints, since high-discount types are more willing to work than low-discount types given the same skill and savings taxes. Recognizing the relevance and importance of heterogeneity in

¹⁷ Laroque (2005) and Kaplow (2006) provide an elegant and straightforward proof of this point. They show that one can always move to a system of nondistorting consumer taxes coupled with an appropriate modification of the earned income tax and generate more government revenue while leaving every consumer with the same utility and the same labor supply.

¹⁸ Banks and Diamond (2010) review evidence on the relationship between savings and skill levels as well as psychological evidence on discount factors. Empirical studies of savings behavior mostly find that those with higher lifetime incomes do save more, but that the full pattern of savings requires considerable complexity in the underlying model (including uncertainties about earnings and medical expenses, asset tested programs, differential availability of savings vehicles, and bequest motives) to be consistent with the different aspects of savings at different ages. Thus the higher savings rates are consistent with the preference assumption of Saez (2002b), but not, by themselves, a basis for necessarily having the discount rate pattern that Saez assumes, since these other factors are also present. Golosov, Troshkin, Tsyvinski, and Weinzierl (2009) propose a calibration exercise.

preferences within and across earnings levels, we reject the direct policy relevance of the Atkinson–Stiglitz theorem.

Distinguishing between Capital and Labor Incomes

A straightforward argument for taxing capital is that it is often difficult to distinguish between capital and labor incomes. For example, people spending time to manage their investment portfolios are converting labor time into anticipated capital income. In small businesses, profits arise both from the labor of owners and returns on assets so that, to some degree, individuals can convert labor income into capital income. For example, after the 1993 Finnish tax reform to a dual income tax with a lower rate on capital income, there were significant shifts of labor income to capital income among the self-employed (Pirttilä and Selin, 2011). In the United States, Gordon and MacKie-Mason (1995) and Gordon and Slemrod (2000) have found income shifting between the corporate tax base and the individual tax base driven by tax differentials. The existence of tax differentials between labor and capital also creates pressure to extend the most favorable tax treatment to a wider set of incomes. For example, in the United States, compensation of private equity and hedge fund managers in the form of a share of profits generated on behalf of clients is considered realized capital gains, although it is conceptually labor income.

The difficulties in telling apart labor and capital income are perhaps the strongest reason why governments would be reluctant to completely exempt capital income and tax only labor income. Christiansen and Tuomala (2008) examine a model with costly (but legal) conversion of labor income into capital income. Despite preferences that would result in an optimal zero tax on capital income in the absence of the ability to shift income, they find a positive optimal tax on capital income. Similarly, the Chamley–Judd result of zero capital income taxation does not hold in a model with an inability to distinguish between entrepreneurial labor income and capital income in the same basic model (Reis, 2007).

Borrowing Constraints

The models discussed above had perfect capital markets—including an absence of borrowing constraints.¹⁹ But borrowing constraints are relevant for tax policy, providing another reason for positive capital income taxation. Since capital income taxes fall on those who are not borrowing constrained (because they have capital), raising revenue from a capital income tax allows for a lower earned income tax, including the tax on those who are so constrained—allowing for an efficiency gain when taxes are collected. For example, Aiyagari (1995) and Chamley (2001) consider borrowing-constrained agents in an uncertainty setting in an infinitely-lived

¹⁹ Zeldes (1989) shows that, contrary to the predictions of the consumption-smoothing model with no liquidity constraints, consumption paths track predictable changes in income for low-wealth groups.

agent model and show that capital income taxation is desirable when consumption is positively correlated with savings.²⁰

Uncertain Future Earnings

Uncertainty about future earnings opportunities is large and pervasive (Banks and Diamond, 2010). When some consumption decisions are taken before earnings uncertainties are resolved, the Atkinson and Stiglitz (1976) result does not hold and, in a two-period model, second-period consumption should be taxed at the margin relative to first-period consumption. The underlying logic of this result is that welfare is enhanced by providing insurance about future earnings opportunities through the tax system. When leisure is a normal good, more savings, *ceteris paribus*, will tend to reduce work later on. Thus, discouraging savings enhances the ability to provide insurance against future poor labor market possibilities. The advantage of discouraging savings is present in models with longer time horizons as well. The extent of insurance is limited by moral hazard concerns.

The literature making this point has two strands. First, the optimal tax strand considers optimal linear taxation of capital income along with optimal nonlinear earnings taxes. Provided an individual's plan with less future work is accompanied by more savings, introducing such taxation raises welfare (Mirrlees and Diamond, 1982; Diamond and Mirrlees, 2000). A second strand commonly called "the new dynamic public finance" (Golosov, Kocherlakota, and Tsyvinski, 2003) has made uncertain future earnings opportunities a central concern, an element largely lacking in the optimal tax approach. It uses the mechanism design approach of social welfare optimization with the government controlling individual consumption and labor, subject to incentive compatibility constraints and aggregate resources. With additive preferences, a robust finding of this literature is the "Inverse Euler Equation," which implies that in the absence of restrictions, an individual would want to save more than called for by the socially optimal plan. To implement such an allocation one needs to have a "wedge" reflecting implicit marginal taxation of future consumption relative to earlier consumption, and so an implicit marginal tax on savings or capital income. In this way, making it less attractive for someone with higher future earnings skills to imitate someone with lower earnings skills improves the equity–efficiency tradeoff.²¹

The mechanism design approach generates the allocation that is optimal, which is then supplemented by analysis of ways to implement such an optimum,

²⁰ This correlation is always positive in the Aiyagari (1995) model with independent and identically distributed labor income, but Chamley (2001) shows that the correlation can be negative theoretically.

²¹ The "Inverse Euler equation" is that the reciprocal (inverse) of the marginal utility of consumption is equal to the expectation of the reciprocal (inverse) of the future marginal utility of consumption—that is, $1/u'(c_1) = E\{1/u'(c_2)\}$. In a certainty model, the Inverse Euler equation and the familiar Euler equation are the same. However, with uncertainty, the marginal utility of present consumption is less than the expected marginal utility of future consumption when the Inverse Euler equation holds. The Inverse Euler condition comes from optimally balancing the incentives for today's work coming from additional compensation today with anticipated changes in future resources as a consequence of today's additional earnings, because the inverse of marginal utility is the resource cost of increasing utility.

sometimes using familiar tax tools. In Mirrlees and Diamond (1982) and Diamond and Mirlees (1986), the authors implement this approach through the adjustment of retirement benefits as a function of the age at retirement in a setting where the alternatives are a particular job or no work at all and there is uncertainty about the ability to hold the job. Implicitly taxing both work and savings allows for more redistribution to those who should retire early by discouraging savings done in order to take advantage of an early retirement pension. The implicit tax comes from a benefit level that grows at less than an actuarially fair rate with continued work. Golosov and Tsyvinski (2006) study optimal disability insurance and recognize a role for an asset test, as is widespread in programs for the poor. However, in a many-period model with a rich stochastic dynamic pattern of wage rates, full implementation of a mechanism design optimum calls for a complex, sophisticated tax structure. When there can be alternative ways to implement a mechanism design optimum, without further research it is not clear which approach sheds light on how to levy taxes in more realistic settings with limited tax tools.²²

The bottom line is that uncertain future earnings opportunities argue against zero taxation of capital income, as do savings preference heterogeneity, limited distinctions between capital and labor incomes, and borrowing constraints. It is true that these arguments are based on life-cycle analyses, and that the empirical literature finds that the life-cycle approach, while helpful, is limited in its success in explaining savings behavior. The belief that many people do not save enough for their own retirements often leads to policies to encourage savings, particularly retirement savings. The most widely employed method is some form of forced saving through mandatory contributions to a retirement system. This is often complemented with a combination of taxing capital income and having tax-favored retirement savings (including some subsidies) targeted to those liable to save too little.

Public Finance Methodology

If we were helping to set tax policy, we would need to reach concrete conclusions on tax bases and tax rates. In our role as part of the general discussion of taxation that may influence the tax-setting process, we look to inform thinking about taxes without necessarily getting to a concrete recommendation. In deciding what issues to promulgate and what supportive arguments to put forth, we draw on parts of the optimal tax literature. We also recognize a role for theoretical analyses in rebutting arguments that do not seem to be a good basis for making tax policy. This approach,

²² An example of a complex implementation is derived in Kocherlakota (2005). It calls for the taxes in any period to depend on the full history of earnings up to that period and has linear capital income tax rates that have a regressive relationship to contemporaneous earnings, and which collect no revenue in aggregate. This implementation discourages savings by making the return to savings stochastic even though the rate of return on investment is determinate. And the regressivity of the tax rate is designed to discourage savings by providing a higher return when marginal utility is lower. Werning (2011) proposes a tax implementation with a progressive capital income tax that can be made independent of past earnings shocks.

drawing on multiple research sources for partial insights, seems appropriate given the complexity of issues that are relevant for good tax policy, much less the even richer set of issues that would also recognize the role of arguments in a complex political process.

As a good model for addressing the many issues that matter for good tax policy, we think of the Meade Report (Meade, 1978). Chapter 2 of the Report, “The Characteristics of a Good Tax Structure,” is divided into six sections: Incentives and economic efficiency; Distributional effects; International aspects; Simplicity and costs of administration and compliance; Flexibility and stability; and Transitional problems. To consider direct taxation in the United Kingdom, the Meade Committee examined each of these issues and then combined the insights into a policy recommendation. It seems to us that economic analysis needs to proceed in a similar fashion.

Optimality analyses of taxation have flourished in two (mostly) separate research communities. The public economics community has been actively doing optimal tax analyses since the mid 1960s, while the macro community, under the banner of “new dynamic public finance,” has been active since the mid 1980s. The standard optimal tax analysis begins with a set of allowable tax structures and optimizes the tax rates and/or tax bases in the allowable structure. In contrast, the macro analysts use a mechanism design approach, which begins by deriving each individual’s marginal rates of substitution consistent with the individual’s optimized consumption and labor allocation—the best possible allocation that is consistent with agents revealing their underlying “types.” The next step is to find a taxation mechanism that can implement this allocation.

This approach rules out taxes that are modeled as requiring information that the government is assumed not to have. A drawback of the mechanism design approach is that it allows—indeed, it often prescribes—complex tax structures that are quite unlike any existing public policies. For example, the literature typically proceeds on the assumption that individuals choose from the allowable set of complete lifetime consumption and earnings plans. It then derives optimal tax mechanisms that make taxes contingent on every observable variable that may be correlated with the key unobservable variable (as in optimal contracting theory, Holmstrom, 1979). The tax faced by a person under this approach might typically depend on the complete earnings and consumption history of that person, even without recognition of other observable variables. Just as the recognition of complexity should limit allowable tools, there is a similar role for public perceptions of tax fairness.²³

Also, analysts using the two approaches sometimes differ in how they approach policy implications. While the public economics community looks for lessons for

²³ A model (for example a game-theoretic equilibrium) that may be perfectly sensible with a small number of sophisticated agents may not be helpful for a large population with limitations in attention to long-term consequences, limited information about tax structures, and limited payoff possibilities. More concretely, legislators, tax administrators, and taxpayers have limited abilities to design, enforce, and comply with complex tax structures. We think that model tractability makes it appropriate to assume rather than derive plausible conditions when one thinks the two approaches would lead to the same central conclusion, even though, of course, some other conclusions would not carry over.

diverse settings, in Kocherlakota (2010, p. 1), which provides a comprehensive treatment of the new dynamic public finance, the author states: “The goal of this book is to figure out at least some characteristics of the best possible tax system.” While the public economics community draws on multiple models, seeking insights, not precise answers, Kocherlakota (2010, p. 4) says: “The ultimate goal of the NDPF [new dynamic public finance] is to provide relatively precise recommendations as to what taxes should be.”

While lessons from mechanism design have added to our understanding of taxation, this methodological narrowness rejects analyses that might be on point for a government considering a limited tax reform. A limited approach to tax reform may reflect political feasibility, a value placed on historical continuity, or limits in acceptable complexity and record-keeping requirements. Therefore, in our view, limited tax reform analysis can inform relevant policy questions and hence should not be rejected on methodological principles.

In our view, the models available for analysis, like much of the underlying theory, remain limited and still too far from reality to proceed in any other fashion than that followed by the Meade Committee. Thus, we have identified basic research findings that we find relevant in thinking about practical tax setting, and also basic research findings that others may find relevant but we do not. In the latter category, we include arguments for high implicit marginal tax rates on low earners in models with only an intensive margin (because the extensive margin is so important for low earners); a zero optimal tax rate at a known top of the earnings distribution (because the top is not known); the low and decreasing marginal tax rate on very high earners that comes from simulations using the lognormal distribution of skills (because the Pareto distribution is well documented to be a better fit); zero taxation of capital income based on the aggregate efficiency result (because the theorem does not have that implication); zero taxation of capital income asymptotically (because bequest behavior does not conform with what is needed for this description of the asymptotic position of the economy); and zero taxation of capital income based on the Atkinson–Stiglitz theorem (because savings rates are not uniform in the population).

■ *We are grateful to Henry Aaron, Alan Auerbach, James Poterba, Ivan Werning, Joel Yellin, and the editors for helpful comments and discussions.*

References

- Aiyagari, S. Rao. 1995. “Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting.” *Journal of Political Economy*, 103(6): 1158–75.
- Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2011. “Top Incomes in the Long Run of History.” *Journal of Economic Literature*, 49(1): 3–71.
- Atkinson, Anthony B., and Agnar Sandmo. 1980. “Welfare Implications of the Taxation of Savings.” *Economic Journal*, 90 (359): 529–49.
- Atkinson, Anthony B., and Joseph E. Stiglitz.

1976. "The Design of Tax Structure: Direct versus Indirect Taxation." *Journal of Public Economics*, 6(1–2): 55–75.

Auerbach, Alan J. 1988. "Capital Gains Taxation in the United States." *Brookings Papers on Economic Activity*, no. 2, pp. 595–631.

Auerbach, Alan J., Laurence J. Kotlikoff, and Jonathan Skinner. 1983. "The Efficiency Gains from Dynamic Tax Reform." *International Economic Review*, 24(1): 81–100.

Banks, James, and Peter Diamond. 2010. "The Base for Direct Taxation." Chap. 6 in *Dimensions of Tax Design: The Mirrlees Review*, ed. by the Institute for Fiscal Studies, 548–648. Oxford: Oxford University Press for the Institute for Fiscal Studies.

Chamley, Christophe. 1986. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives." *Econometrica*, 54(3): 607–22.

Chamley, Christophe. 2001. "Capital Income Taxation, Wealth Distribution and Borrowing Constraints." *Journal of Public Economics*, 79(1): 55–69.

Christiansen, Vidar, and Matti Tuomala. 2008. "On Taxing Capital Income with Income Shifting." *International Tax and Public Finance*, 15(4): 527–45.

Cremer, Helmuth, and Pierre Pestieau. 2006. "Wealth Transfer Taxation: A Survey of the Theoretical Literature." In *Handbook of the Economics of Altruism, Giving and Reciprocity*, Vol. 2: Applications, ed. Serge-Christophe Kolm and Jean Mercier Ythier, 1107–34. Amsterdam: Elsevier, North Holland.

Diamond, Peter A. 1973. "Taxation and Public Production in a Growth Setting." In *Models of Economic Growth*, ed. J. A. Mirrlees and N. H. Stern, 215–34. London: MacMillan.

Diamond, Peter. 1980. "Income Taxation with Fixed Hours of Work." *Journal of Public Economics*, 13(1): 101–110.

Diamond, Peter A. 1998. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review*, 88(1): 83–95.

Diamond, Peter A. 2003. *Taxation, Incomplete Markets and Social Security: The 2000 Munich Lectures*. Cambridge: MIT Press.

Diamond, Peter A., and James A. Mirrlees. 1971. "Optimal Taxation and Public Production I: Production Efficiency." *American Economic Review*, 61(1): 8–27.

Diamond, Peter A., and James A. Mirrlees. 1986. "Payroll-Tax Financed Social Insurance with Variable Retirement." *Scandinavian Journal of Economics*, 88(1): 25–50.

Diamond, Peter A., and James A. Mirrlees. 2000. "Adjusting One's Standard of Living: Two

Period Models." In *Incentives, Organization, and Public Economics: Papers in Honour of Sir James Mirrlees*, ed. P. J. Hammond and G. D. Myles, 107–122. Oxford: Oxford University Press.

Diamond, Peter, and Johannes Spinnewijn. Forthcoming. "Capital Income Taxes with Heterogeneous Discount Rates." *American Economic Journal: Economic Policy*.

Farhi, Emmanuel, and Iván Werning. 2010. "Progressive Estate Taxation." *Quarterly Journal of Economics*, 125(2): 635–73.

Farhi, Emmanuel, and Iván Werning. 2011. "Insurance and Taxation over the Life Cycle." NBER Working Paper 16749.

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski. 2003. "Optimal Indirect and Capital Taxation." *Review of Economic Studies*, 70(3): 569–87.

Golosov, Mikhail, Maxim Troshkin, Aleh Tsyvinski, and Matthew Weinzierl. 2009. "Preference Heterogeneity and Optimal Capital Income Taxation." NBER Working Paper 16619.

Golosov, Mikhail, and Aleh Tsyvinski. 2006. "Designing Optimal Disability Insurance: A Case for Asset Testing." *Journal of Political Economy*, 114(2): 257–69.

Goolsbee, Austan. 2000. "What Happens When You Tax the Rich? Evidence from Executive Compensation." *Journal of Political Economy*, 108(1): 352–78.

Gordon, Roger H., and Jeffrey K. MacKie-Mason. 1995. "Why is There Corporate Taxation in a Small Open Economy?" In *The Effects of Taxation on Multinational Corporations*, eds. Martin Feldstein, James R. Hines, Jr., and R. Glenn Hubbard, 67–94. Chicago: University of Chicago Press.

Gordon, Roger H., and Joel B. Slemrod. 2000. "Are 'Real' Responses to Taxes Simply Income Shifting Between Corporate and Personal Tax Bases?" In *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. Joel B. Slemrod. Harvard University Press: Cambridge.

Gruber, Jonathan, and Emmanuel Saez. 2002. "The Elasticity of Taxable Income: Evidence and Implications." *Journal of Public Economics*, 84(1): 1–32.

Holmström, Bengt. 1979. "Moral Hazard and Observability." *The Bell Journal of Economics*, 10(1): 74–91.

Internal Revenue Service (IRS). 2009a. "SOI [Statistics of Income] Tax Stats—Individual Income Tax Returns Publication 1304." Individual Income Tax Returns for 2007. <http://www.irs.gov/taxstats/indtaxstats/article/0,,id=134951,00.html>.

Internal Revenue Service (IRS). 2009b. "The 400 Individual Income Tax Returns Reporting the Highest Adjusted Gross Incomes Each

- Year, 1992–2007.” <http://www.irs.gov/pub/irs-soi/07intop400.pdf>.
- Judd, Kenneth L.** 1985. “Redistributive Taxation in a Simple Perfect Foresight Model.” *Journal of Public Economics*, 28(1): 59–83.
- Judd, Kenneth L.** 1999. “Optimal Taxation and Spending in General Competitive Growth Models.” *Journal of Public Economics*, 71(1): 1–26.
- Kaplow, Louis.** 2006. “On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal.” *Journal of Public Economics*, 90(6–7): 1235–50.
- Kocherlakota, Narayana R.** 2005. “Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation.” *Econometrica*, 73(5): 1587–1621.
- Kocherlakota, Narayana R.** 2010. *The New Dynamic Public Finance*. Princeton, NJ: Princeton University Press.
- Kopczuk, Wojciech.** 2005. “Tax Bases, Tax Rates and the Elasticity of Reported Income.” *Journal of Public Economics*, 89(11–12): 2093–2119.
- Laroque, Guy R.** 2005. “Indirect Taxation is Superfluous under Separability and Taste Homogeneity: A Simple Proof.” *Economics Letters*, 87(1): 141–44.
- Mankiw, N. Gregory, and Matthew Weinzierl.** 2010. “The Optimal Taxation of Height: A Case Study of Utilitarian Income Redistribution.” *American Economic Journal: Economic Policy*, 2(1): 155–76.
- Mankiw, N. Gregory, Matthew C. Weinzierl, and Danny Yagan.** 2009. “Optimal Taxation in Theory and Practice.” *Journal of Economic Perspectives*, 23(4): 147–74.
- Meade, James Edward.** 1978. *The Structure and Reform of Direct Taxation*. Report of a Committee chaired by Professor J. E. Meade. London: George Allen & Unwin.
- Meyer, Bruce D.** 2010. “The Effects of the EITC and Recent Reforms.” In *Tax Policy and the Economy*, vol. 24, ed. Jeffrey R. Brown, 153–180. Cambridge: MIT Press.
- Mirrlees, James A.** 1971. “An Exploration in the Theory of Optimal Income Taxation.” *Review of Economic Studies*, 38(2): 175–208.
- Mirrlees, James A., and Peter A.** 1982. “Social Insurance with Variable Retirement and Private Saving.” MIT Working Paper 296.
- Pareto, Vilfredo.** 1896 [1965]. “La courbe de la répartition de la richesse.” In *Écrits sur la courbe de la répartition de la richesse*. Writings by Pareto collected by G. Busino, Librairie Droz, 1965, pp. 1–15.
- Piketty, Thomas, and Emmanuel Saez.** 2003. “Income Inequality in the United States, 1913–1998.” *Quarterly Journal of Economics*, 118(1): 1–39 (Associated data series updated to 2007 in August 2009.)
- Piketty, Thomas, and Emmanuel Saez.** 2007. “How Progressive is the U.S. Federal Tax System? A Historical and International Perspective.” *Journal of Economic Perspectives*, 21(1): 3–24.
- Pirttilä, Jukka, and Håkan Selin.** 2011. “Income Shifting with a Dual Income Tax System: Evidence from the Finnish Tax Reform of 1993.” *Scandinavian Journal of Economics*, 113(1): 120–144.
- Reis, Catarina.** 2011. “Entrepreneurial Labour and Capital Taxation.” *Macroeconomic Dynamics*, 15(3): 326–35.
- Sadka, Efraim.** 1976. “On Income Distribution, Incentive Effects and Optimal Income Taxation.” *Review of Economic Studies*, 43(1): 261–68.
- Saez, Emmanuel.** 2001. “Using Elasticities to Derive Optimal Income Tax Rates.” *Review of Economic Studies*, 68(1): 205–29.
- Saez, Emmanuel.** 2002a. “Optimal Income Transfer Programs: Intensive versus Extensive Labour Supply Responses.” *Quarterly Journal of Economics*, 117(2): 1039–73.
- Saez, Emmanuel.** 2002b. “The Desirability of Commodity Taxation under Non-linear Income Taxation and Heterogeneous Tastes.” *Journal of Public Economics*, 83(2): 217–30.
- Saez, Emmanuel.** 2004. “The Optimal Treatment of Tax Expenditures.” *Journal of Public Economics*, 88(12): 2657–84.
- Saez, Emmanuel, Joel Slemrod, and Seth Gietz.** Forthcoming. “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review.” *Journal of Economic Literature*.
- Seade, Jesus K.** 1977. “On the Shape of Optimal Tax Schedules.” *Journal of Public Economics*, 7(1): 203–36.
- Slemrod, Joel B.** 1996. “High Income Families and the Tax Changes of the 1980s: The Anatomy of Behavioral Response.” In *Empirical Foundations of Household Taxation*, eds. Martin Feldstein and James Poterba, 169–92. Chicago: University of Chicago Press.
- Slemrod, Joel B., ed.** 2000. *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*. Harvard University Press: Cambridge.
- Slemrod, Joel, and Wojciech Kopczuk.** 2002. “The Optimal Elasticity of Taxable Income.” *Journal of Public Economics*, 84(1): 91–112.
- U.S. Census Bureau.** 2009. “Median Household Income for States: 2007 and 2008: American Community Surveys.” American Community Survey Report, September 2009.
- Werning, Ivan.** 2011. “Nonlinear Capital Taxation.” Unpublished working paper. <http://dl.dropbox.com/u/125966/implementation.pdf>
- Zeldes Stephen P.** 1989. “Consumption and Liquidity Constraints: An Empirical Investigation.” *Journal of Political Economy*, 97(2): 305–346.

Appendix

Comparison with Mankiw, Weinzierl, and Yagan (2009)

In this journal, Mankiw, Weinzierl, and Yagan (2009)¹ present eight lessons that they draw from the optimal tax literature. Our paper agrees with some of their lessons but also draws some very different conclusions. In this appendix, we discuss some of the discrepancies between our interpretations, following the order of the eight lessons they present.

Lesson 1: Optimal Marginal Tax Rate Schedules Depend on the Distribution of Ability: We agree.

Lesson 2: The Optimal Marginal Tax Schedule Could Decline at High Incomes: Major disagreement.

Mankiw, Weinzierl, and Yagan base this lesson on two arguments: First, they present the zero top marginal tax rate result, which, combined with positive marginal tax rates below the top, implies that the tax rate should decline as it approaches the top. Second, they discuss numerical simulations using log-normal skill distributions that show modest rates that sometimes decrease in the upper part of the distribution. They dismiss the results that use Pareto distributions and which obtain high tax rates on upper incomes on two grounds: First, they claim that one cannot infer the ability distribution without making unduly strong assumptions. Second, they examine the right tail of the wage density distribution using Current Population Survey (CPS) data (Mankiw, Weinzierl, and Yagan, Figure 1, p. 154) and conclude that it is not possible to distinguish Pareto versus log-normal distributions from such data. We find both of those arguments invalid—the zero at the top is not relevant for policy, as discussed above and in the online appendix, and the evidence is strongly supportive of the Pareto distribution.

For any distribution with a thinner top tail than the Pareto distribution, such as a log-normal distribution, the parameter $a = z_m / (z_m - z^*)$ diverges to infinity. The test of Pareto versus lognormal right tails presented by Mankiw, Weinzierl and Yagan (2009) in their Figure 1 lacks power for two reasons. First, it uses CPS data that is thin and top coded in the upper part of the distribution. Indeed, their graph covers a range of earnings from \$80,000 to \$150,000 (for full-time and full-year individuals working 2,000 hours per year). Second, it plots density fits that are inherently imprecise. As we made clear in the optimal top tax rate derivation, the statistic of central interest is $a = z_m / (z_m - z^*)$ as depicted on Figure 2 in our paper. This statistic is a much more precise way to estimate the relevant shape than a density fit. Using individual tax return data which have high sampling at the top and no top coding, we show that the

¹ Mankiw, N. Gregory, Matthew C. Weinzierl, and Danny Yagan. 2009. "Optimal Taxation in Theory and Practice." *Journal of Economic Perspectives*, 23(4): 147–74.

statistic a remains extremely stable over a very large range of incomes, which is much broader than the range considered by Mankiw, Weinzierl, and Yagan.

Furthermore, as our derivation has made clear, the optimal top rate derivation we have proposed does not require unusually strong assumptions in terms of homogeneity of preferences in the population or functional form assumptions. Mankiw, Weinzierl, and Yagan cite the simulation results of Saez (2001), which by necessity require making functional form assumptions, but fail to note that the general theoretical tax rate formula $\tau = 1/(1 + ae)$ is much more general than the numerical illustration. The virtue of the formula $\tau = 1/(1 + ae)$ is precisely that it depends only on estimable sufficient statistics.

Lesson 3: A Flat Tax, with a Universal Lump-Sum Transfer, Could Be Close to Optimal: Major disagreement.

The analysis we presented showed that, at the bottom, transfers initially increase with earnings to preserve incentives to participate in the labor force and then are phased out with income at a high rate. Mankiw, Weinzierl, and Yagan do not discuss the participation margin. After transfers are phased out, marginal tax rates should be lower for the broad middle class and then rise in the upper income groups due to declining marginal welfare weights and the Pareto shape of the income distribution toward the top but not lower down. Therefore, the optimal system appears quite different from the flat tax with a universal lump-sum transfer that they advocate.

Lesson 4: The Optimal Extent of Redistribution Rises with Wage Inequality: We agree.

Lesson 5: Taxes Should Depend on Personal Characteristics as Well as Income: Some disagreement.

While a model ignoring both issues of complexity and social acceptability would reach this conclusion for many observable characteristics, we think that these two issues should get due respect. In practice, taxes and transfers depend significantly on only few characteristics (besides income), and those characteristics, such as family structure or disability status, are related to need.

Lesson 6: Only Final Goods Ought to be Taxed, and Typically They Ought to Be Taxed Uniformly: Some disagreement.

Limiting variation in commodity (or VAT) taxes is appropriate, but some variation seems well justified, although too much variation seems to be present in some systems. We disagree with their inference that this line of argument supports not taxing capital income; in fact, this disagreement is a central part of our presentation.

Lesson 7: Capital Income Ought to Be Untaxed, at Least in Expectation: Major disagreement.

Mankiw, Weinzierl, and Yagan invoke three arguments for zero capital income taxes, including two that we have addressed in the text and found not

policy relevant. They claim that the Diamond–Mirrlees aggregate efficiency result implies that capital income should not be taxed. We explained that the theorem does not have this implication in footnote 15.

Lesson 8: In Stochastic Dynamic Economies, Optimal Tax Policy Requires Increased Sophistication: Some disagreement.

We agree that stochastic elements call for more sophisticated analysis and justify more sophisticated structures. However, we disagree with their emphasis on the optimality of a regressive interaction between capital income taxation and labor income, and where capital income taxation raises no revenue in expectation (referred to in the title of Lesson 7). As discussed in footnote 21, they focus on one implementation of the mechanism design optimum with this property but ignore the presence of a different implementation of the same optimum that has positive taxation of capital (Werning, 2010). It is not clear how to draw inferences from different implementations of a full mechanism design optimum when limited complexity implies that the full optimum is not being implemented.