

Tax Equivalences and their Implications

Alan J. Auerbach
University of California, Berkeley and NBER

October 1, 2018

Executive Summary

In economic analyses of the effects of tax policies, one commonly encounters discussions of the equivalence of apparently different policies, where equivalence is defined as the policies having the same impact on fundamental economic outcomes. These related tax policies may differ in many respects, including (1) the side of a market on which they are applied; (2) the form in which they are imposed (e.g., as a unit or *ad valorem* tax, on a tax inclusive or tax exclusive basis, etc.); (3) whether they are imposed on households or firms; (4) the market in which they are directly imposed; (5) their timing; and (6) whether behavioral adjustments are involved in the equivalence. These differences give rise to conditions under which the equivalences may break down, because of several factors, including (1) differences in salience; (2) market imperfections, such as liquidity constraints, price rigidity or imperfect competition; (3) differences in information requirements and the costs of tax administration and enforcement; and (4) government accounting rules.

This paper draws out the key issues that relate to tax equivalences, using several illustrations from important instances of such equivalences that span different areas of taxation, with many of these illustrations relating to the taxation of capital income. Recognition of equivalences and the ways in which they may fail to hold is important both for positive analysis (e.g., the political reasons for choosing one approach over another) and for normative analysis (to determine which approach may be a more effective way of implementing a policy).

This paper was presented at the 2018 NBER Tax Policy and the Economy Conference, Washington, DC, September 27, 2018. I am grateful to Bill Gale, Louis Kaplow, Robert Moffitt, Emmanuel Saez, Danny Yagan, and conference participants for comments on earlier drafts.

I. Introduction

In economic analyses of the effects of government policies, one commonly encounters discussions of the equivalence of apparently different policies. Perhaps most familiar is that a tax imposed on transactions in a market should have the same impact whether it is formally levied on the consumer or the producer.




Fig. 1. A unit tax at rate t on seller or buyer.

As illustrated in Figure 1, the imposition of a unit tax at rate t on sellers raises the supply curve from S to S' , whereas the same unit tax imposed on purchasers lowers the demand curve from D to D' . In each case, purchases fall from Q_0 to Q_1 , the tax-inclusive price paid by purchasers rises to P_2 , the after-tax price received by sellers falls to P_1 , and government revenue equals $t \cdot Q_1$. The only difference between the two outcomes is that the observed

market price falls to P_1 when the tax is paid by consumers and rises to P_2 when producer pay the tax. This simple example illustrates the central elements of a tax equivalence: the budget constraints and economic incentives of agents and the government are unaffected, as is the resulting equilibrium, even though market prices may differ. It follows from the equivalence that the incidence and efficiency effects of the two tax policies are the same as well, regardless of the behavioral elasticities of individuals and firms, because these elasticities apply equally across different versions of the equivalent policy.

Implicit in this equivalence is the absence of real-world factors that might cause it to break down, such as restrictions on the movement of market prices (for example, a minimum wage or frictions that result from bargaining institutions), differences in the salience of the two taxes among market participants, and differences in the costs of taxpayer compliance or tax administration. Without any such differences, the choice between the two approaches would be arbitrary. With them, there might be an advantage to adopting one approach rather than the other, which is a good reason for identifying an equivalence in the first place. For distinct reasons, for example if one approach causes taxes to be more hidden than the other does, there could be differences in political appeal.

Equivalences among different approaches to taxation are common in the literature, but analyses that discuss them generally focus on one particular market context, such as international trade, labor markets, or saving and investment.¹ The objective of the broader review that follows is to draw out the main issues that relate to tax equivalences, using several illustrations from important instances of such equivalences that span different areas of taxation. In addition, although most of the equivalences discussed below have appeared in the

previous literature, a number of them will be unfamiliar to many readers. This lack of familiarity has contributed to erroneous arguments regarding the properties and relative attractiveness of different taxes, particularly in cases where new approaches to taxation have been the subject of discussion.

Drawing on different examples, the next section of the paper reviews the types of differences that might arise between otherwise equivalent tax systems, i.e., factors that might contribute to making an equivalence less than exact and hence relevant to the design of tax policy. Section III uses illustrations of such differences to demonstrate how they matter in the implementation of tax policy. Many of the equivalences introduced in sections II and III concern the taxation of capital income and the relationship among different taxes bases, such as wealth, capital income, consumption, and cash flow. Section IV provides some concluding comments.

While the analysis will cover many important tax equivalences, the aim here is not to provide a comprehensive catalog. In addition, we leave aside other important area of policy equivalence. One such area, covered only briefly below, is between tax-based and spending-based policies, associated with the concept of “tax expenditures.” Another is between tax-based and regulation-based policies. Familiar examples here include the classic equivalence between tariffs and quotas in international trade, those between environmental policies based on taxes and quantity restrictions (as discussed in Weitzman 1974), and indeed the recently litigated one between mandates and taxes in health insurance markets. Although the paper will not discuss these categories of policy equivalences directly, we note that many of the issues considered here carry over to these other contexts, and that many of this paper’s lessons apply.

II. Differences Relevant to Tax Equivalences

In addition to which side of a market faces the tax, there are many other differences relevant when considering tax equivalences. We discuss these differences in this section, while maintaining the assumption of perfect competition, information, etc., deferring until the next section how real-world considerations may cause the equivalences to break down. As we show, several of these differences may apply simultaneously to a particular tax equivalence.

A. Differences in Form or Characterization

Perhaps the closest equivalences involve taxes assessed in the same manner on the same individuals, but with some superficial difference in form or characterization. One such equivalence relates to whether a tax benefit is delivered as a tax credit or as tax deduction (or exclusion). For any given tax rate, t , a deduction or exclusion from the tax base of an amount is equivalent to a fractional tax credit on that amount at a rate equal to the tax rate; e.g., a taxpayer facing a 15 percent marginal tax rate would receive the same benefit from a deduction and a 15 percent tax credit. Naturally, for a progressive tax system, equivalence to a deduction would generally require a schedule of credit rates increasing with income, a point often made when comparing the distributional effects of credits versus deductions.

Another equivalence with respect to form involves the imposition of a tax on a *tax-exclusive* basis versus a *tax-inclusive* basis, i.e., having the tax apply to gross expenditures (including tax) or net expenditures (excluding tax). A tax at rate t applied on a tax-inclusive basis is equivalent to a tax at rate $\frac{t}{1-t}$ applied on a tax-exclusive basis. As an illustration, one could impose a consumption tax of 25 percent on a tax-exclusive basis or at 20 percent on a

tax-inclusive basis. It is sometimes more natural to use one approach or the other. For example, as laid out in Andrews (1974), one might tax consumption through a personal tax on income used to finance consumption expenditures (i.e., by providing a deduction under an income tax for net saving). For this, a tax-inclusive approach would be more natural: if income (Y) is exhausted by consumption (C), savings (S) and taxes (T), so that $Y = C + S + T$, then a tax on income less savings is a tax on $C + T$. It would be more natural to implement a consumption tax taking the form of a retail sales tax on a tax-exclusive basis (as sales taxes currently are), adding the tax to the net-of-tax price of consumption goods and services.

There is a one-to-one correspondence between tax-inclusive and tax-exclusive rates and hence no evident importance whether one uses a tax-inclusive or tax-exclusive approach. Unfortunately, even this equivalence can be misunderstood, as appears to have been the case during the U.S. debate over a national retail sales tax, for which proponents sometimes expressed the rate on a tax-inclusive basis, contrary to the common sales-tax approach that others might have inferred was being used (Gale 2005).

A related distinction involves whether prices are quoted inclusive or exclusive of tax. In the United States, goods subject to retail sales taxes (RSTs) have their prices quoted exclusive of tax, while excise taxes are included in listed prices, which is also the case for value added taxes (VATs) in place outside the United States. This distinction closely tracks whether the tax is assessed on sellers or buyers, but, in principle, one could have prices quoted as inclusive of tax even if the formal tax liability were on buyers.

The comparison between a proportional retail sales tax and a proportional personal consumption tax highlights another minor distinction between otherwise equivalent taxes,

namely between *indirect* and *direct* taxes, the former (in this case, the sales tax) applying to transactions and the latter (the personal consumption tax) applying to individuals (or, in other instances, businesses). As the retail sales tax is legally an obligation of individual purchasers, both taxes are proportional taxes on consumption assessed on individuals – the same tax rate, tax base, and taxpayer. Yet, the direct-indirect distinction is relevant in law, perhaps because in practice, direct and indirect taxes have usually differed in more substantive ways as well. Through treaties and other international agreements, this legal distinction between direct and indirect taxes has the potential to introduce real differences in economic impact even in cases where none would exist otherwise.

A more substantive economic difference in form occurs between unit taxes (also called specific taxes) and *ad valorem* taxes. Unit taxes are assessed per unit of quantity, q , while *ad valorem* taxes are assessed per unit of expenditure, $p \cdot q$. For any given price, p , a unit tax at rate t is equivalent to an *ad valorem* tax at rate $\frac{t}{p}$.

Finally, one tax may be equivalent to a combination of other taxes. Notably, it would be possible to implement the VAT as the combination of two taxes, one a proportional tax on wages and the other a proportional tax, at the same rate, on the non-wage components of value added. This equivalence is at the heart of *flat tax* proposals, as in Hall and Rabushka (1983). These authors' reason for breaking up the two pieces of the VAT was to facilitate a shift in the location of the wage component to the individual level from the firm – to the other side of the labor market, leaving only a cash-flow tax at the business level.² However, this second step does not inherently follow the first: one could implement the two taxes separately at the business level.

Another familiar equivalence between one tax policy and a combination of others, in the international trade context, is between a tariff on imported goods and the combination of a tax on *all* domestic sales and a domestic production subsidy at the same rate, since aggregate domestic sales equals the sum of sales by foreign sellers and those by domestic sellers.

Yet another important example of the equivalence of one policy instrument and a combination of others involves investment incentives, such as an investment tax credit or an accelerated schedule of depreciation allowances. Since these incentives lower the present value of taxes associated with *new* investment, which, as discussed below, is equivalent to reducing the tax rate on the income from such investment, one may think of them as a combination of (1) a reduction in taxes on the income from old *and* new capital; and (2) an offsetting tax increase on the income from old capital (or, alternatively, a capital levy on the value of old capital).

B. Differences in Entities

There are cases in which a given tax may be collected from more than one entity. For example, individuals may pay taxes related to their assets or associated income directly, as individual taxpayers, or indirectly through the entities or accounts they own, such as corporations or retirement accounts. In other cases, even where individuals own businesses that receive and report the income, the income is taxed directly to the individuals, through pass-through treatment.

Another example concerns the individual income tax, which is assessed on individuals but, through withholding, collected and remitted by employers. The same is true with the employee portion of the OASDHI payroll tax.

C. Differences in Markets

Consider a one-period model in which individuals earn labor income and use their income to purchase two consumption goods, so that the budget constraint in the absence of taxation is

$$p_1 c_1 + p_2 c_2 = wL, \tag{1}$$

where p_i and c_i are the price and quantity of good- i consumption, w is the wage rate and L is labor supply. Even though labor income taxes and consumption taxes are imposed in different markets (for labor and consumption goods), it is easy to see in this model that proportional taxes on labor and consumption are equivalent, with this equivalence working through the individual's budget constraint.

A tax on labor income at rate t multiplies the right-hand side of the budget constraint by $(1 - t)$, and a uniform, proportional tax on consumption at rate τ divides the left-hand side of the budget constraint by $(1 - \tau)$ (assuming that both are imposed on a tax inclusive basis.) Setting $\tau = t$ leaves the individual with the same budget constraint and provides the government with the same amount of tax revenue.

Note that the equivalence between the two taxes in this example does not require the observed price level to be the same in the two cases. For example, if the wage rate paid by firms and the prices received by them were fixed in nominal terms, a sales tax added to producer prices would raise the consumer price level, while a tax subtracted from employee wages would not. But relative prices (between the two goods and between each good and labor, and hence the real wage) are still the same in the two cases.

D. Differences in Timing

Tax policies may also differ with respect to the timing of tax payments. The equivalence between proportional taxes on labor income and consumption just discussed carries over in a life-cycle model in which individuals neither leave nor receive bequests. In a two-period version of that model, in which the individual is retired in the second period and there is one consumption good in each period, the lifetime budget constraint without taxes is

$$p_1 c_1 + \frac{1}{1+r} p_2 c_2 = wL, \quad (2)$$

where r is the interest rate, the two consumption goods from expression (1) are reinterpreted as first- and second-period consumption, and L is first-period labor supply. As in the previous example, the individual and government budget constraints are the same – in this case in present value – under equal-rate consumption and labor income taxes.

A slight modification of the budget constraint in (2) applies when the individual starts period 1 with initial assets, which may be interpreted as an inheritance, access to above-market rates of return, or rents, on certain investments of fixed size, or, in the analysis of a tax reform, as the asset accumulated by the individual in past periods. For the budget constraint,

$$p_1 c_1 + \frac{1}{1+r} p_2 c_2 = wL + A, \quad (3)$$

the tax on consumption is equivalent to the combination of two other taxes, on labor income and initial assets. This distinction between wage taxes and consumption taxes – the capital levy on initial wealth – features prominently in discussions of the efficiency and distributional properties of the adoption of a consumption tax (e.g., Auerbach and Kotlikoff 1987). Indeed,

this is the same distinction as in the case of investment incentives versus reductions in capital income taxes, given the saving-investment identity and the fact that consumption equals income less saving.

Of course, whether the component A in the budget constraint in (3) is truly exogenous from a broader perspective, for example taking account of previous decisions regarding saving or the location of investment that would be influenced by an anticipated tax on A , plays an important role in how one interprets potential efficiency effects.

E. Portfolio Adjustments

In the previous equivalence example, the individual and the government both have the same present-value budget constraints across tax systems, but translating these budget constraints into the same economic outcome requires additional adjustments to private and government saving. The individual must save more under a consumption tax than under a wage tax to achieve the same first-and second-period consumption, because some tax payments are deferred under a consumption tax. Likewise, the government must save less under the consumption tax, which provides it with some resources in the second period. There is no difference in national saving between the two cases, just the division of national saving between private and public saving. (This illustrates the limited usefulness of focusing on how tax policies affect specific components of saving.)

A more complex response is needed in cases where there is uncertainty, as the impact of tax systems on the riskiness of assets may call for changes not only in the level of saving but also with regard to portfolio composition.³ To focus specifically on this issue, suppose that in the two-period budget constraint both labor and first period consumption are fixed, so that the

individual saves a fixed amount, V , allocating a share a to a risky asset yielding a stochastic return x and a share $(1 - a)$ to a safe asset yielding a certain return r . Then, without taxes, second-period consumption equals V multiplied by $[1 + ax + (1 - a)r]$. If a proportional income tax at rate t is imposed on *excess* returns – those returns above the rate r – then the expression for second-period consumption becomes $V \cdot [1 + r + a(x - r)(1 - t)]$. It follows that the individual can achieve the same outcome as in the absence of taxation by increasing the portfolio share invested in the risky asset from a to $\frac{a}{1-t}$. From the government's perspective, the same adjustment in reverse will be needed to maintain the original level of risk in the government's revenue stream (assuming that combining different taxpayers' risky tax payments results in no additional risk-pooling⁴). That is, the government will need to shift some of its investment portfolio from risky to safe assets, as might be achieved, for example, by reducing government borrowing – which makes the government's net position in safe assets less negative – and at the same time reducing public investment, which type of investment presumably yields uncertain returns.

If a tax on excess returns has no impact on individual or government budget constraints, it follows that a tax on *all* returns, equal to a tax on the risk-free and excess return components, is equivalent to a tax only on the risk-free components. This tax on risk-free returns, in turn, is equivalent to a tax on wealth, following the equivalence between unit and *ad valorem* taxes discussed above, since a tax at rate t on returns at the risk-free rate on wealth, rV , is equivalent to a tax at rate $t \cdot r$ on wealth itself, V .

III. Tax Equivalences and the Implementation of Tax Policy

If all of the conditions for complete equivalence held, then it would be a matter of indifference through which equivalent version a particular tax policy were implemented. It is through the exceptions to equivalence that choices among otherwise equivalent policies matter.

A. Perception and Salience

Otherwise-equivalent taxes may differ in how they are perceived by market participants or by those engaged in the political process. As to the former, empirical evidence (Chetty, Looney, and Kroft 2009) suggests that taxes included in prices (e.g., excise taxes) have a larger impact on consumer responses than those that are added to seller-quoted prices (e.g., sales taxes). Lower salience can have complex effects on a policy's desirability, reducing deadweight loss because of a lower sensitivity to taxation but increasing it through a misallocation of household budget constraints. In addition, to the extent that salience differs among households, the use of less salient forms of taxation may have adverse distributional consequences (Goldin and Homonoff 2013). Rees-Jones and Taubinsky (2018) describe the various ways in which salience and other psychological biases may cause equivalences to break down, and emphasize the importance of considering the actual tax policies being implemented, rather than the underlying incentives and budget sets that may establish policies as equivalent.

As distinct from how psychological factors may influence the behavior of direct market participants, such factors may also lead the political process to treat economically equivalent taxes differently, possibly influencing the form of policies adopted. A common reference here is the distinction between the employer and employee components of the payroll tax, both of which are withheld by the employer prior to paying workers but which are officially assessed on

different sides of the labor market. Even though we lack strong evidence that these components affect labor outcomes differently simply because of the side of the market on which they are imposed (except in specific cases where wage adjustments are difficult, as where minimum wages intervene or where differential wage changes would be required)⁵, the durability of these separate components suggests that statutory incidence is politically quite relevant, perhaps because policy equivalences are not sufficiently evident or well-understood.

Standard conventions for analyzing the distributional effects of taxes may also contribute to perceived differences. For example, in earlier years it was common for U.S. government agencies not to allocate corporate taxes to individual taxpayers for purposes of analyzing the distributional effects of taxes. This meant that equivalent tax policies, such as a corporate tax reduction for dividends paid and a shareholder tax reduction for dividends received, would have been treated quite differently, being ignored in the former case and included in the latter case. More recently, the convention of viewing a standard VAT as a tax on retail sales borne by consumers, but the components of the two-part VAT (described above) as, respectively, taxes falling on asset owners and on workers, helped sow confusion over the distributional effects of a business cash-flow tax, i.e., the first component of the two-part VAT (Viard 2018).

Just as differences in the statutory incidence of taxes may affect perceptions of these taxes, so may the distinction as to whether policy changes are classified as affecting taxes or transfer payments. In 1993, for example, this labeling choice resulted in an argument in Congress about whether to characterize the pending increase in the taxation of social security benefits (from 50 percent to 85 percent of benefits subject to tax, over an exempt threshold) as

a tax increase or as a reduction in government spending, given its equivalence to a policy of reducing social security benefits. Potentially hinging on the choice of terminology was one's conclusion about the policy's impact on the "size of government," which some would see a benefit cut (but not a tax increase) as reducing, based on measuring the size of government using the level of government spending. A similar issue arises with the Earned Income Tax Credit, which for budget scoring purposes is divided between a tax reduction and a component of government spending depending on whether a recipient receives a net payment from the government (which is treated as spending), but which also could be implemented solely as an expenditure-based transfer program.

B. Market Imperfections

Tax equivalences may also break down when market imperfections are present. It may be possible to lessen the impact of such imperfections by choosing a particular version of a tax policy. We focus on three types of imperfections below: liquidity constraints, incomplete private risk-pooling, and private-sector market power.

1. Liquidity Constraints

When households or firms face liquidity constraints, shifting tax liability to periods when such constraints are relaxed may improve market outcomes. As emphasized by Hubbard and Judd (1986), this could change one's conclusions about tax policies that shift the timing of tax payments. From the perspective of equivalent taxes, it could introduce a difference between two systems that differ only with respect to timing. For example, in the pure life-cycle model with no initial assets (i.e., for the lifetime budget constraint in expression (2)), individuals facing liquidity constraints in the first period would fare worse under a labor income tax than under a

consumption tax, which shifts some of their tax liabilities to the second period, after they have accumulated assets.

Liquidity constraints may influence firms as well. Consider the case of the incentive to invest, which according to standard analysis depends on the present value of depreciation deductions, but not their timing. Under this standard analysis, immediate expensing of investment is equivalent to following a standard pattern of depreciation deductions over time, but with an additional deduction for the nominal rate of interest multiplied by the undepreciated asset basis in each year (Boadway and Bruce 1984). That is, rather than receiving an immediate deduction of the cost of investment, writing off the asset at some geometric rate, α , leaves an undepreciated basis of $e^{-\alpha t}$ and a depreciation allowance of $\alpha e^{-\alpha t}$ at date t per dollar of initial investment. Providing the taxpayer also with a deduction for the nominal interest rate, r , multiplied by the undepreciated basis yields a total deduction at date t of $(r + \alpha)e^{-\alpha t}$. The present value of all such deductions is $\int_0^{\infty} e^{-rt}(r + \alpha)e^{-\alpha t} dt = 1$, the same value as under immediate expensing.⁶

If firms are liquidity constrained, though, expensing or other investment incentives that are more “front-loaded” may be more effective in stimulating investment. There is convincing evidence that liquidity-constrained firms respond more strongly than other firms to up-front investment incentives such as expensing and “bonus” depreciation (which allows firms a fractional immediate write-off of qualifying investment expenses). This response does not occur simply through a reduction in the present value of taxes, but due to a shift in the timing of taxes as well (Zwick and Mahon 2017).

2. Incomplete Risk-Pooling

To the extent that private markets do not efficiently pool risks, taxes that shift risks from taxpayers to the government may reduce aggregate risk. In this case, the tax on all returns to investment would no longer be equivalent to a tax only on the risk-free rate of return, contrary to the result discussed earlier. Taxing all returns would reduce aggregate risk, through the pooling of the investment risks faced by taxpayers.

However, one needs to ask why markets may inefficiently pool risk. To the extent that risk-pooling would harm economic incentives (for example, in the case of entrepreneurs whose effort may influence the distribution of risky outcomes), efficiency may call for incomplete risk-pooling. In such cases, lessening an individual's exposure to risk through taxation could have the same negative effect as further private risk-pooling.

3. Market Power

When firms have market power, certain tax equivalences break down and leave one alternative preferable to the other. For example, the equivalence between unit and *ad valorem* taxes relies on the assumption that firms are price-takers. If firms are not, then they will consider the sensitivity of profits to variations in price, which *ad valorem* taxes dampen because they are imposed in relation to prices. This can lead to lower prices and more output under *ad valorem* taxes than under unit taxes (Delipalla and Keen 1992).

Another aspect of imperfect competition is the presence of non-competitive rents earned by firms. In this case, the equivalence between expensing of new investment and imposing no tax on the returns to such investment breaks down, as with expensing there is no tax in present value on the returns to marginal investment but there is a tax on investments

yielding above-market returns. For an investment with cash flows Y_i at each date i , allowance of investment expensing results in a cash-flow tax. Under a cash-flow tax, the government taxes positive flows and negative flows (i.e., investment costs) at a common tax rate t . This means that from the investor's point of view all cash flows are all multiplied by $(1 - t)$, so that after-tax cash flows at each date are $(1 - t)Y_i$. Thus, there is no change in the present value of investments with a before-tax present value of zero, for $(1 - t)$ multiplied by zero is still zero. Hence, the cash-flow tax is equivalent, except for timing, to imposing no tax on returns to marginal investments. On the other hand, positive present-value investments are reduced in value by a factor of t . This means not only that the positive returns to past investments will be taxed – in present value, there is a capital levy on returns to “old” capital, for which past investments receive no expensing deduction – but also that future rents will be taxed as well.

This ability to capture a portion of rents using a cash-flow tax holds if the tax is imposed at the firm level. Unless the firm is taxed on a pass-through basis, the timing of such rent taxation will differ for taxation at the individual level. Without rents, a cash-flow tax at the firm level and at the shareholder level would be equivalent, assuming the same rate of tax and that the value of the firm equals the value of its assets. It would not matter if cash flowing into the firm were deducted by the individual owner or by the firm itself, or whether cash flowing out of the firm were taxable to the firm or to the ultimate recipient. However, if rents are present at the firm level, this equivalence is altered, for the value of the firm should include the capitalized value of its rents.⁷ Hence, rents would face no tax at the individual level once they have been capitalized, with the tax on such capitalized rents being paid in advance by those original shareholders who owned the firm when the capitalization occurred. This difference in timing

need not break the equivalence, if tax rates remain constant over time; nor would the incidence of the tax differ, in both cases falling on the original shareholders, as future taxes imposed at the firm level would be taken into account in the determination of the value at which they could sell their shares.

C. Government Accounting and Budget Rules

Equivalences among different policies may also break down because of government accounting and budget rules and conventions. For example, having different rules that apply to transfer payments and taxes can disrupt the equivalence between a reduction in benefits and a tax on those benefits. However, perhaps the biggest impact of government accounting, for policies that differ with respect to timing, comes from the use of a limited time horizon, in particular the focus on revenues over a one-year, five-year or ten-year period. With budget measurement that applies only through a certain date, policies that shift revenue losses outside the budget window relax the constraints that legislators face.⁸ The creation and expansion of so-called “Roth” accounts for retirement saving owe much to this motive, as do tax incentives for conversion of other retirement accounts into Roth accounts.

Roth accounts follow the TEE (Taxable/Exempt/Exempt) pattern of taxation with respect to the three phases of saving: initial contributions, inside build-up during the accumulation period, and withdrawals, respectively. This is in contrast to the EET approach traditionally followed for pensions and other retirement accounts such as IRAs and 401(k) plans. The approaches are economically equivalent if the tax rate does not change over time, another instance of the equivalence between cash-flow taxation and tax exemption. The equivalence here also relies on the additional but reasonable assumptions that there are no rents earned by

investors themselves in making passive retirement-account investments and that government taxation of ex post returns from such investments under the EET approach does not improve risk-pooling.

However, when the budget window is truncated, the government revenue losses associated with the Roth accounts are understated relative to those following the traditional timing of tax payments. Relative to a full, present-value accounting of the revenue costs of adopting tax-favored saving accounts, the TEE approach understates the cost, by ignoring future revenue losses, while the EET approach *overstates* the cost, by ignoring future revenue recoupment.

A related example comes from the Tax Reform Act of 1986, which was constructed to have a balanced revenue impact over a five-year budget window. In that setting, a policy of eliminating the investment tax credit and lowering the corporate tax rate on *new* investment to maintain the same present value of taxes on such investment would have shown a revenue gain, with the tax cuts on the income from the first five years' investment falling to some extent outside the budget window and the tax increase from eliminating the investment tax credit falling fully within the budget window. This may have influenced the ultimate policy outcome, which was to reduce the corporate income tax rate (on *all* investment) while eliminating the investment tax credit, a policy that provided *windfalls* to existing capital (i.e., a reverse of the capital levy that would have resulted from increasing the investment tax credit and the corporate tax rate) while actually raising the effective tax rate on new investment (Auerbach 1987).

To the extent that distributional analysis of tax policies adheres to the limited-budget-window approach, one encounters the same problem that fundamentally equivalent policies may be perceived (or at least reported) as being quite different. For example, given that retirement saving is concentrated among higher-income individuals, the distributional impact of the EET approach will appear less progressive than it is, and less progressive than the TEE approach, because it will show tax deductions during the budget window while ignoring the eventual tax payments made on withdrawals outside the budget window.

In addition to government accounting, *private* accounting conventions may also influence how otherwise equivalent tax policies are viewed. One important example is the treatment of tax benefits based on timing, such as accelerated depreciation, which do not affect measured tax rates in public financial statements, because in such statements the time value of money is ignored in the computation of the deferred tax liabilities arising from accelerating deductions. This convention makes investment incentives provided through accelerated depreciation appear different in companies' financial statements from direct tax rate reductions, even if the present value of tax benefits is the same for the two policies. Companies seeking to report higher after-tax earnings may prefer tax rate reductions to accelerated depreciation for this reason, even if other factors (such as liquidity constraints) might point in the other direction; and this preference may have political consequences as well. As with government accounting rules and conventions, financial accounting procedures may affect both private decisions and political outcomes.

D. Tax Administration

Differences among otherwise equivalent taxes arise in many dimensions relating to the ease of administering and enforcing tax systems as well as the costs of compliance to taxpayers. We consider three such dimensions below, relating to a tax system's self-correcting features (its ability to adapt to changes in economic conditions without being explicitly modified), the costs of enforcement and compliance, and, especially, the information required by the government for its implementation.

1. Self-Correcting Features

Tax systems that respond automatically to fluctuations in the economy or a taxpayer's individual circumstances have an administrative advantage over those requiring explicit adjustment on the part of the government.

One challenge from the economic environment is inflation, which causes problems with measuring the tax base as well as with adjusting the tax schedule. Measuring the tax base is especially a problem for capital income, where nominal income must be adjusted for changes in the purchasing power of assets over time. For example, the present value of historic-cost depreciation allowances, $D(t)$ at each date t , permitted over time per unit cost of an asset purchased at date 0, equals $\int_0^{\infty} e^{-rt} D(t) dt$, where r is the nominal discount factor. As the inflation rate rises, so does r , and this reduces the present value of depreciation deductions, even if the real interest rate is unaffected. This problem motivated the proposal by Auerbach and Jorgenson (1980) for "first-year capital recovery" – permitting firms to deduct in the year of an asset's purchase the present value of depreciation allowances, calculated using an assumed

real interest rate, equal to the nominal discount rate r minus the inflation rate π , or

$$\int_0^{\infty} e^{-(r-\pi)t} D(t) dt.$$

In the absence of inflation, the two approaches – actual deduction of depreciation allowances over time, and up-front deduction of the present value of these depreciation allowances – would yield the same deduction, in present value. In the presence of inflation, it would be necessary to index depreciation allowances for inflation – to multiply them by $e^{-\pi t}$ – for the present value of allowances to equal those provided by first-year capital recovery,

$$\int_0^{\infty} e^{-rt} e^{\pi t} D(t) dt = \int_0^{\infty} e^{-(r-\pi)t} D(t) dt.$$

As for adjusting the tax schedule, this presents another difference between *ad valorem* and unit taxes. Whereas *ad valorem* taxes rise in proportion to prices, and so maintain their real value as prices rise, unit taxes do not, making continual adjustments necessary, which may be particularly burdensome in economies with high rates of inflation.

Another issue of self-correction relates to the fact that the income tax has a progressive marginal tax rate schedule applied on an annual basis. Because of this convexity in the tax schedule, for individuals with the same average income over time, those with fluctuating incomes face higher average tax rates than those with stable incomes. Here, the EET approach to pension taxation may be preferable to the TEE approach, although the ability to choose between the two may provide a further improvement in performance.

Consider, again, the two-period life-cycle model in which individuals work and save for retirement in the first period and consume in both periods. Under the EET approach, the first-period tax base equals wage income less saving, or the amount of first-period consumption expenditure, and the second-period tax base equals the amount of second-period consumption

expenditure. Thus, if the household smooths consumption fully, the tax base will be the same in both periods. Under the TEE approach, by contrast, the first-period tax base will be higher – equal to all wage income – and the second-period tax base lower – in this case equal to zero. Thus, the EET approach incorporates an automatic income-averaging feature that is absent under the TEE approach.

While the EET approach results in complete smoothing of the tax base in this simple case, it would not do so if a taxpayer's preferences led to a rising or falling consumption profile, or in cases of year-to-year fluctuations in consumption expenditures. For the general case, one could improve smoothing simply by giving the taxpayer a choice between the TEE and EET approaches each year. For example, having a choice would facilitate shifting the tax base to years with low consumer expenditures by using the TEE approach in those years. Indeed, making this choice available was suggested in U.S. Treasury (1977) for implementing a personal consumption tax, under which all saving would be tax-favored, following either the EET or the TEE approach.⁹

2. Enforcement and Compliance

Even though taxes may have the same impact in a world of costless enforcement and complete information, in a realistic setting important differences may arise.

One difference involves the identity of the person or entity with either the statutory responsibility for a tax or the obligation to remit the tax (as in the case of firms that withhold and remit income taxes on behalf of individual taxpayers who have the statutory responsibility for the taxes). In many situations, it is easier to monitor firms than individuals.

Another difference relates to the incentives for taxpayer compliance. A classic illustration is the difference between the VAT and the RST. From national income accounting, we know that final consumption equals value added less investment, government purchases, and net exports:

$$C = GDP - I - G - (X - M). \quad (4)$$

Hence, there is an equivalence between a VAT and an RST, which is imposed directly on consumption. One may express this equivalence in terms of a series of switches in tax liability, from seller to buyer, for each intermediate transaction in the value-added chain, starting from the VAT. This cascade of switches from seller to buyer at each successive stage, none of which affects the equilibrium outcome, ultimately concentrates the entire tax liability at the last point of sale, as under the RST. However, under a credit-invoice VAT, each successive seller is able to deduct purchased inputs from the tax base only by presenting evidence that tax was previously paid on the sales of such inputs, providing a third-party enforcement mechanism that is absent under the RST.

3. Information

In addition to ease of enforcement, some versions of equivalent tax systems require less information on the part of government than other versions. One may think of this, ultimately, as an issue of enforcement as well, although it is useful to understand the sources of enforcement challenges. Many of the illustrations below involve some form of cash-flow taxation, which is endowed with the natural informational advantage of being based only on observable cash flows.

Consider the taxation of consumption, which as just discussed could be implemented through a VAT. Leaving aside the tax exemption for government purchases (which, for a unified government, is of no consequence, as it amounts to a government paying taxes to itself) and the exemption for net exports – to which we return below – a VAT amounts to a tax on income net of investment. Income may be broken into wage and non-wage components, and thus one may implement the consumption tax as a cash-flow tax on non-wage income net of investment, $GDP - I - W$, plus a tax on wages, W . As mentioned earlier, one may tax these two components separately, as under a flat tax, but unless they face different rates of tax (as they would under a flat tax), this separation is unnecessary. Indeed, VATs as traditionally implemented do not distinguish between wage and non-wage components of value added.

As already discussed, a business cash-flow tax – the non-wage component of the VAT – effectively exempts the safe rate of return to investment, but not rents. To the extent that rents are not important (as in the case of passive investments), one could also exempt the safe rate of return simply by imposing a zero tax rate on returns to capital. However, imposing a zero tax rate on returns to capital and a positive tax rate on wages requires that the government be able to distinguish between capital income and labor income, whereas this is not necessary under a VAT or other forms of consumption taxation, such as a personal consumption tax, which tax expenditures rather than income. A consumption tax approach to exempting the return to capital alleviates the need to distinguish between capital and labor income, and is not subject to the argument against exempting returns to capital based on the concern that taxpayers will attempt to disguise labor income as capital income to achieve a lower tax rate. Put another way, a business cash-flow tax falls on all returns in excess of the

safe rate, including rents as well as disguised labor income, in addition to the returns to risk-taking.

Another illustration of differential information requirements in equivalent tax systems arises with respect to the taxation of capital gains. A problematic feature of capital gains taxation is the realization doctrine, under which capital gains are taxed only when assets are sold and the gains “realized.” Because delaying realization allows an investor to defer tax on gains already accrued, this causes a “lock-in” effect, an economic distortion whereby investors continue to hold assets that they otherwise would wish to sell.

A potential solution to the lock-in effect is to tax capital gains on accrual or, to deal with possible liquidity problems coming from having to pay taxes before assets are sold, to continue taxing gains upon realization but with an annual interest charge being added to the tax on gains already accrued but not yet realized. That is, as proposed originally by Vickrey (1939), the tax accrued as of the end of period t on an asset purchased at the end of period 0 would be:

$$T_t = [1 + r(1 - \tau)]T_{t-1} + \tau(A_t - A_{t-1}), \quad (5)$$

where $T_0 = 0$, r is the safe rate of return applied to unpaid taxes, τ is the income tax rate (assumed here to be the same for interest and capital gains), and A_t is the value of the asset at the end of period t . An obstacle to implementing this method of taxation (or accrual taxation) is that the government may not observe the path of A_t , but only the asset’s value at the times of initial purchase and final sale. One could assume some arbitrary path of appreciation after observing the sale price at realization, for example using a smooth geometric growth rate between the purchase and sale dates, but doing so would not fully eliminate the lock-in effect

(or its opposite, an incentive to realize losses immediately) if the assumed path differs from the actual one.¹⁰

However, as shown in Auerbach (1991), a tax liability that instead evolves according to the expression

$$T_t = \left[1 - \left(\frac{1+r(1-\tau)}{1+r} \right)^t \right] A_t \quad (6)$$

(which depends only on the holding period t and the observed value at realization) imposes the same tax rate τ at each date t on the safe component of the rate of return, rA_{t-1} , as in (5), and a tax rate of $\left[1 - \left(\frac{1+r(1-\tau)}{1+r} \right)^t \right]$ on the excess return, $(A_t - A_{t-1}) - rA_{t-1}$. This tax rate on the excess return differs from that on the safe return component, and from the tax rate under the Vickrey approach. But, as discussed above in relation to the taxation of risky assets, the investor can use a portfolio adjustment to offset the tax on the excess return, leaving the investor in the same position as under the tax system in (5), but without the government needing to observe the asset appreciation trajectory in order to eliminate the lock-in effect.

The retrospective capital gains tax system described in (6) eliminates the need for the government to observe any asset's appreciation pattern, but it still requires the observation of the holding period and sale price for each asset subject to capital gains taxation. However, even these requirements can be eliminated by adopting a generalized version of the system, based on cash-flow taxation. As Auerbach and Bradford (2004) show, taxing all of the investor's cash flows (positive and negative, from all investments) at date t at a rate equal to

$$\left[1 - a \left(\frac{1+r(1-\tau)}{1+r} \right)^{t-B} \right], \quad (7)$$

where B is some fixed base-year reference date for each taxpayer and α is an arbitrary constant, would have the same properties as the retrospective capital gains tax (a specific version of the system described in (7) for which $\alpha = 1$ and B equals the date of purchase, and hence must be different for each asset). Thus, one of the main advantages of cash-flow taxation – not needing to measure asset income – carries over even when the tax system *does* impose an income tax, effectively reducing the normal return to investment through a cash-flow tax rate that rises over time.

Returning once again to the simple, constant-rate cash-flow tax, an important equivalence first noted by the Meade Committee Report (Institute for Fiscal Studies 1978) was between different methods of implementing a cash-flow tax. As that report (Chapter 12) observed, the transactions of the business sector are related by an identity which, using our notation, may be written for a closed economy as

$$(R - I) + F + S = 0, \tag{8}$$

where R equals the cash-flow returns to real investment expenditures, I , F equals net financial inflows to the firm (excluding transactions with shareholders), including net borrowing and net interest receipts, and S equals dividend payments and net transactions in corporate shares, i.e., the net inflows to firms from their shareholders. Under the cash-flow (i.e., non-wage) component of most existing VATs, only $R - I$ is included in the tax base – financial transactions are ignored. The Meade Committee referred to this approach as the R base cash-

flow tax, and to the tax base that included the financial flows, F , as the $R + F$ base cash-flow tax.

The present value of a firm's financial flows F should equal zero to the extent that borrowing and lending are at the same interest rate – borrowing and lending activities each having zero present value – and that similar conditions hold for other financial activities, e.g., the provision of insurance is actuarially fair, etc. However, this conclusion presumes that any value added in the financial sector is measured separately from these financial flows and hence already included in the real returns, R .

In reality, much of the value added that the financial sector generates appears not as explicit payment for services, but instead takes the form of differences in rates of return (e.g., a lower rate of interest paid to depositors; a higher rate of interest received from borrowers), higher insurance premiums, or, more generally, differences between financial inflows and outflows that make F positive in present value. Thus, an R base is likely to miss a large component of the services provided by financial institutions. An information advantage of including financial flows in the tax base is that doing so obviates the need to try to distinguish the value-added components of financial flows, since real and financial flows both are subject to taxation.

Note also from expression (8) that the $R + F$ base could be implemented simply by taxing net distributions to shareholders, $(-S)$, including dividends and net share repurchases.¹¹ However, imposing a business-level cash-flow tax may be easier from the government's perspective than taxing individual shareholders on their net distributions. This is especially so in the global context, where investments by individuals and companies cross borders. To

consider this, we return to expression (4), the national income identity relating consumption to income.

Assuming, for simplicity of notation that the government imposes a tax on its own purchases, achieving a consumption tax in an open economy requires not only a tax on wages plus domestic business cash flows, $(W + R - I = GDP - I)$, but also a tax on net imports, $M - X$. This last tax component, accomplished by taxing gross imports and providing an equal subsidy for gross exports, is known as a *border adjustment*. Border adjustment is a standard component of existing VATs and is also part of the *destination-based cash-flow tax* (DBCFT), which is equivalent to the VAT combined with a deduction for wages (or a VAT combined with a payroll subsidy at an equal rate).

A familiar equivalence from international trade theory is between a tax on exports and a tax on imports, which arises from the intertemporal international constraint that a nation's exports and imports must be equal, in present value. This is the *Lerner Symmetry Theorem* (Lerner 1936) and it implies that a tax on imports and an equal-rate *subsidy* to exports is neutral. That is, border adjustment of imports and exports at the same rate is equivalent to no border adjustment at all.

However, it is the *failure* of conditions needed for the Lerner Symmetry Theorem that provides a central, information-based argument for implementing the DBCFT rather than the *origin-based* cash-flow tax (the tax on $R - I$ alone, with no border adjustment). In particular, the government's inability to observe R – the domestic returns to investment – makes the DBCFT attractive, to prevent companies from using various accounting techniques to shift profits to lower-tax jurisdictions (Auerbach 2017, Auerbach et al. 2017).

Border adjustment eliminates a company's ability to shift domestic profits to other countries using cross-border transactions with related foreign parties, because those transactions have no impact on the company's domestic tax liability with the border adjustment in place. In the computation of a company's domestic tax base, the revenues from exports and the expenses of imports are exactly offset by the border adjustments applied to these items.¹²

Understanding the relationship between the DBCFT and the taxation of shareholders requires the use of one additional identity, relating the capital and current international accounts. According to the balance of a country's current and capital accounts (and ignoring government reserve transactions and international transfer payments which also enter into the identity), the trade surplus, $X - M$, plus net foreign source income, E , equals net foreign investment; equivalently, the trade deficit equals net foreign source income less net foreign investment, or $M - X = E$. Substituting this relationship into the definition of the DBCFT tax base yields:

$$\text{DBCFT Tax Base} = (R - I) + M - X = (R - I) + E. \quad (9)$$

That is, as noted in Auerbach (1997), implementing the DBCFT is equivalent to imposing a tax on domestic cash flows $(R - I)$ as well as net cross-border cash flows (E) .¹³ This equivalence provides another perspective as to why the DBCFT eliminates the tax incentive for companies to shift earnings abroad. Because an understatement of domestic earnings and net cash flows, $(R - I)$, will lead to an equal-size increase in net foreign-source earnings and net cash flows from abroad, E , misreporting the first component of their sum will have no impact on the

DBCFT tax base, although understating $(R - I)$ in this manner may *increase* tax liabilities in other countries that still impose source-based taxes .¹⁴

From this final way of writing the DBCFT tax base, it also follows that, if one adopts the Meade Committee's $R + F$ approach and incorporates domestic financial flows F in the tax base as well, the resulting tax system amounts to the taxation of domestic resident shareholders on the distributions from their investments, whether these investments are domestic or foreign. That is, the DBCFT is equivalent to a tax on the distributions received by domestic shareholders. However, to impose the latter directly would require the government to observe its domestic shareholders' net distributions from foreign investments. Through the international accounts identity, the DBCFT accomplishes the same outcome, working through different markets (trade in goods and services versus capital and income flows), via the border adjustment.

IV. Conclusions

Recognizing tax equivalences is important to help us understand the possible economic effects of tax systems; applying what may be known about the effects of some tax systems may help us understand the effects of other systems that are equivalent.

Perhaps more important for the design and analysis of tax policy is that understanding tax equivalences is also a valuable first step in considering the ways in which equivalences may fail. Such a consideration is useful both for positive analysis (e.g., to understand the political reasons for choosing one approach over another) and for normative analysis (to determine which approach may be a more effective way of implementing a policy, given the many

deviations from the assumptions of perfect markets and information that can apply in realistic settings).

Finally, one should keep in mind that deviations from equivalence like those discussed here need not be immutable. Government budget rules could be more rational, technological innovations may increase the government's access to information about taxpayer behavior, market environments relating to competitiveness and nominal rigidities may change, and so forth. For a similar reason, specific policy prescriptions may differ across countries at any given time. Thus, perhaps the most important lesson to be taken from the discussion here relates to the process of evaluating tax policy rather than any particular results. One can learn much from understanding tax equivalences and deviations from them. Therefore, we should keep this way of framing tax policies in mind when considering the effects of new proposals as well as existing law.

Endnotes

¹ Auerbach, Frenkel and Razin (1997), a paper on which the current one relies, does consider a range of equivalences broader than the “international taxation” in its title suggests.

² Taxing wages at the individual level, in the flat tax proposal, would have facilitated the imposition of a progressive tax rate schedule on individual or household wage income, rather than the proportional one included as part of a VAT.

³ For further discussion, see Kaplow (1994).

⁴ Whether government can improve risk-pooling depends on the extent to which efficient private markets for sharing risks already exist, as well as on the extent to which individual risks are themselves correlated.

⁵ Saez et al. (2012) find that employer and employee payroll taxes have different incidence for a Greek tax reform that affected different cohorts of employees differently, for which equal incidence would have required posted wages to differ based on whether workers were subject to the tax reform.

⁶ A further equivalence here is between tax systems that provide a deduction for the nominal interest rate multiplied by the undepreciated basis, as represented here, and those that permit a deduction for interest on the share of debt used to finance the asset, d , and for the interest rate multiplied by the undepreciated basis *net of* debt, $(e^{-\alpha t} - d)$. This latter approach is known as an “allowance for corporate equity,” or ACE (Institute for Fiscal Studies 1991).

⁷ In the parlance of the investment literature, Tobin’s q would exceed 1, as the firm’s value would equal the value of its measured assets plus the present value of its above-normal returns. While expensing at the firm level would be based on the cost of asset purchases, expensing at the individual level would be based on the higher value paid for the firm.

⁸ To a lesser extent, the practice of adding up nominal effects on revenues over different years without discounting has the effect of encouraging a shift of revenue losses to earlier years within the budget window, when nominal magnitudes are smaller for any given ratio to GDP.

⁹ Taxpayers may also have an incentive to use the choice between EET- and TEE-based accounts to time their tax bases for strategic reasons, if they expect the tax system to change over time. An individual who expects marginal tax rates to rise, for example, will be more likely to opt for the TEE approach, beyond what would be needed to accomplish a smoothing of his or her tax base across periods, in order to subject a greater share of lifetime income to a lower marginal tax rate.

¹⁰ For example, if an asset appreciated substantially shortly after purchase, the use of an assumed constant annual growth rate based on the ultimate sale price would leave the investor locked in, as holding the asset longer would spread more of the large immediate gain over subsequent periods of lower rates of appreciation.

¹¹ Note that, with no net share repurchases, this amounts to a tax on dividends. This equivalence between a dividend tax and a cash-flow tax, and particularly the property that the dividend tax does not affect the after-tax rate of return to investment, is a major element of the “new view” of dividend taxation (Auerbach 1979).

¹² Note that the Lerner symmetry theorem may break down for other reasons, for example because of nominal rigidities, which may not serve as advantages for the DBCFT. See Barbiero et al. (2018). In such circumstances, the violation of other equivalences, for example leading the side of the market on which the tax is applied to matter, will be relevant for implementation.

¹³ In terms of how this cash-flow tax on foreign-held assets would be effected, the border adjustment should lead to an offsetting appreciation of the real exchange rate, through some combination of an increase in the domestic price level or an appreciation of the nominal exchange rate. This real exchange rate appreciation reduces the real value of cash flows on foreign assets. See Auerbach et al. (2017). Several factors could influence the mix of increases in domestic prices and the nominal exchange rate, including exchange rate policy and whether there were a wage deduction on the employer side or a wage subsidy on employee side, in the presence of some nominal wage rigidity.

¹⁴ As discussed in Auerbach et al. (2017), it would also be desirable to include a border adjustment for the financial flows incorporated in E , to prevent firms from using related-party cross-border lending to shift profits.

References

- Andrews, William D. (1974). "A Consumption-Type or Cash Flow Personal Income Tax," *Harvard Law Review* 87(6): 1113-1188.
- Auerbach, Alan J. (1979). "Wealth Maximization and the Cost of Capital," *Quarterly Journal of Economics* 93(3): 433-446.
- Auerbach, Alan J. (1987). "The Tax Reform Act of 1986 and the Cost of Capital," *Journal of Economic Perspectives* 1(1): 73-86.
- Auerbach, Alan J. (1991). "Retrospective Capital Gains Taxation," *American Economic Review* 81(1): 167-178.
- Auerbach, Alan J. (1997). "The Future of Fundamental Tax Reform," *American Economic Review* 87(2): 143-146.
- Auerbach, Alan J. (2017). "Demystifying the Destination-Based Cash-Flow Tax," *Brookings Papers on Economic Activity* 2017(2): 409-432.
- Auerbach, Alan J., and David F. Bradford. (2004). "Generalized Cash-Flow Taxation," *Journal of Public Economics* 88(5): 957-980.
- Auerbach, Alan J., Michael P. Devereux, Michael Keen, and John Vella. (2017). "Destination-Based Cash-Flow Taxation." Oxford Centre for Business Taxation Working Paper 17/01, January.
- Auerbach, Alan J., Jacob A. Frenkel, and Assaf Razin. (1997). "Equivalence Relations in International Taxation," in Mario I. Blejer and Teresa Ter-Minassian, eds., *Macroeconomic Dimensions of Public Finance: Essays in Honour of Vito Tanzi*: 146-163. London: Routledge.
- Auerbach, Alan J., and Dale W. Jorgenson. (1980). "Inflation-Proof Depreciation of Assets," *Harvard Business Review* 58(5): 113-118.
- Auerbach, Alan J., and Laurence J. Kotlikoff. (1987). *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press.
- Barbiero, Omar Emmanuel Farhi, Gita Gopinath, and Oleg Itskhoki. (2018). "The Macroeconomics of Border Taxes," in Martin Eichenbaum and Jonathan Parker, eds., *NBER Macroeconomics Annual* 33.
- Boadway, Robin, and Neil Bruce. (1984). "A General Proposition on the Design of a Neutral Business Tax," *Journal of Public Economics* 24(2): 231-239.
- Chetty, Raj, Adam Looney, and Kory Kroft. (2009). "Salience and Taxation: Theory and Evidence," *American Economic Review* 99(4): 1145-77.
- Delipalla, Sofia, and Michael Keen. (1992). "The Comparison between Ad Valorem and Specific Taxation under Imperfect Competition," *Journal of Public Economics* 49(3): 351-367.

- Gale, William G. (2005). "The National Retail Sales Tax: What Would the Rate Have to Be?" *Tax Notes* 107(7): 889-911.
- Goldin, Jacob, and Tatiana Homonoff. (2013). "Smoke Gets in Your Eyes: Cigarette Tax Salience and Regressivity." *American Economic Journal: Economic Policy* 5(1): 302-36.
- Hall, Robert E., and Alvin Rabushka. (1983). *Low Tax, Simple Tax, Flat Tax*. New York: McGraw-Hill.
- Hubbard, R. Glenn, and Kenneth L. Judd. (1986). "Liquidity Constraints, Fiscal Policy, and Consumption," *Brookings Papers on Economic Activity* 1986(1): 1-59.
- Institute for Fiscal Studies. (1978). *The Structure and Reform of Direct Taxation*. London: Allen and Unwin.
- Institute for Fiscal Studies. (1991). *Equity for Companies: A Corporation Tax for the 1990s*. Commentary 26.
- Kaplow, Louis. (1994). "Taxation and Risk Taking: A General Equilibrium Perspective," *National Tax Journal* 47(4): 789-798.
- Lerner, Abba P. (1936). "The Symmetry between Import and Export Taxes," *Economica*, 3(11): 306-313.
- Rees-Jones, Alex, and Dmitry Taubinsky. 2018. "Taxing Humans: Pitfalls of the Mechanism Design Approach and Potential Resolutions," in Robert Moffitt, ed., *Tax Policy and the Economy* 32: 107-133.
- Saez, Emmanuel, Manos Matsaganis, and Panos Tsakloglou. (2012). "Earnings Determination and Taxes: Evidence From a Cohort-Based Payroll Tax Reform in Greece," *Quarterly Journal of Economics* 127(1), 493-533.
- U.S. Treasury, Office of Tax Analysis. (1977). *Blueprints for Basic Tax Reform*.
- Viard, Alan D. (2018). "Too Clever by Half: The Politics and Optics of the Two-Part VAT," *National Tax Journal* 71(2): 399-416.
- Vickrey, William. (1939). "Averaging of Income for Income Tax Purposes." *Journal of Political Economy* 47(3): 379-397.
- Weitzman, Martin L. (1974). "Prices vs. Quantities," *Review of Economic Studies* 41(4): 477-491.
- Zwick, Eric, and James Mahon. (2017). "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review* 107(1): 217-248.