

ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

A UNIFIED PERSPECTIVE ON EFFICIENCY, REDISTRIBUTION, AND PUBLIC POLICY

Louis Kaplow

Discussion Paper No. 1027

01/2020

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center

The Social Science Research Network Electronic Paper Collection:
<https://ssrn.com/abstract=3518787>

***A UNIFIED PERSPECTIVE ON
EFFICIENCY, REDISTRIBUTION, AND PUBLIC POLICY***

Louis Kaplow*

January 13, 2020

Abstract

Specialized theoretical and empirical research should in principle be embedded in a unified framework that identifies the relevant interactions among different phenomena, enables an appropriate matching of policy instruments to objectives, and grounds normative analysis in individuals' utilities and a social welfare function. This article advances an approach that both provides integration across many dimensions and contexts and also identifies which tasks may be undertaken separately and how such analysis should be conducted so as to be consistent with the underlying framework. It employs the distribution-neutral methodology and welfare analysis developed in Kaplow (2008a) and related work, offering applications to income taxation, commodity taxation, tax expenditures, externalities, public goods, capital income and wealth taxation, social security and retirement savings, estate and gift taxation, and transfer programs. It also explores welfare criteria and examines how their consideration enables the normative analysis of the taxation of families, heterogeneous preferences, and tax administration and enforcement.

Keywords: income taxation, optimal taxation, inequality, efficiency, redistribution, commodity taxation, public goods, externalities, tax expenditures, capital taxation, wealth taxation, estate and gift taxation, social insurance, transfer programs, welfarism, taxation of families

JEL Codes: D61, D62, D63, H20, H21, H23, H24, H26, H41, H43, H53, H55

*Harvard University and NBER. This article originated in my Holland Medal address delivered at the 2019 annual meeting of the National Tax Association. I am grateful to the John M. Olin Center for Law, Economics, and Business at Harvard University for financial support. Disclaimer: I occasionally consult on antitrust cases, and my spouse is in the legal department of a financial services firm.

I. INTRODUCTION

A unified view of a field is complementary to specialized theoretical and empirical work that abstracts from even basic matters. Most research inevitably must be specialized, but it is useful at times to consider how various strands connect to the whole. Unless the relevant linkages are periodically identified and assessed, there is the risk that efforts may be misdirected and that results may provide misleading guidance to policymakers. Moreover, in setting research agendas, it is important to appreciate what, in principle, needs to be known to provide prescriptions.

This article, drawing primarily on my own writing, offers a unified conceptual perspective on efficiency, redistribution, and public policy. In doing so, it seeks to identify connections within and across a number of dimensions:

- *Policy instruments*: the effects and optimal design of the income tax and transfer system, differential commodity taxes, Pigouvian corrections (taxation and regulation), public goods, capital and wealth taxation, wealth transfer taxation, and social security.
- *Policy objectives*: in particular, the relationship between efficiency and redistribution.
- *Utility and welfare functions*: the need to trace effects to individuals' utility, implications of notions of social welfare that do not depend entirely on individuals' utility, the choice of social welfare function, and implications for thinking about equality, mobility, poverty, horizontal equity, taxation of the family, heterogeneous preferences, and tax administration and enforcement.

Accomplishing this mission is a tall order. This investigation elaborates key concepts and connections that provide a rough framework that is useful for a variety of more concrete tasks.

To further motivate the inquiry, consider empirical research that estimates the magnitude of the elasticity of charitable giving with respect to the net-of-tax cost. A common benchmark for the “efficiency” of the deduction for charitable contributions is whether the magnitude of this elasticity exceeds 1.0, for then the deduction (at the margin) induces more than a dollar of giving for each dollar of revenue cost. See, for example, Boskin and Feldstein (1977).¹ When seeing subsequent papers on the subject, I asked myself whether I could think of a plausible model in which this policy prescription made sense. (In today's phraseology, this is akin to asking whether this elasticity is a sufficient statistic for welfare analysis, with a further particular claim regarding the implications of various magnitudes.) I could not.

In my subsequent research (largely not on this particular subject), the core normative question—how society should determine the optimal level of the charitable deduction (or other such instrument)—is analyzed by breaking it down in the following way:

- Charitable giving, the target of the subsidy, is a species of giving. It differs from the other main category of giving (between individuals, primarily from parents to children) in involving the support of various public goods.
- Giving of the ordinary sort, in turn, is a species of consumption. It differs in three principal respects from ordinary consumption: (1) there is a positive externality to the donee that even an altruistic donor does not fully credit; (2) there is a negative externality on the fisc (in today's parlance, a fiscal externality) to the extent that the income or wealth effect induces the donee to work less and thus pay less in taxes;

¹ Boskin and Feldstein (1977, p. 351) refers to whether the deduction is “fully efficient in this sense,” without later elaboration. Although not explicitly advanced as a sufficient statistic for welfare analysis, a motivation for this line of research seems to be that this elasticity is a central determinant of optimal policy.

- (3) gifts involve voluntary income redistribution that changes the marginal utilities of the donor and donee and thus interacts with the optimal redistributive income tax problem.
- Various types of ordinary consumption themselves are objects of policy, ordinarily analyzed under the subject of differential commodity taxation.
 - Consumption as a whole is the use of labor income that is subject to income taxation; commodity taxation and income taxation interact, so income taxation must be addressed as part of the analysis.
 - Optimal income taxation, in turn, depends on a number of empirical parameters and on the social welfare function.

Perhaps the reader is wondering: What about the elasticity of charitable giving? In a simple setting in which a charitable deduction is modeled as a corrective subsidy designed to internalize a positive externality, this elasticity not only fails to be a sufficient statistic, but in basic cases does not even appear in the formula—which calls for the standard Pigouvian correction, that is, a subsidy equal to the marginal external benefit. Relaxing one of the assumptions (weak separability of labor effort in the utility function), this elasticity becomes relevant in light of a fiscal externality correction. Even then, the elasticity can enter with either sign (depending on the direction of the interaction)—so that a higher elasticity could favor a smaller subsidy rather than a larger one—and there is no particular significance to the magnitude 1.0.

The point of this exercise is not that researchers should avoid specialization and, in particular, the bracketing of some phenomena in order to illuminate others. Nor should empirical research limit itself to the estimation of sufficient statistics for welfare analysis, eschewing all manner of study that illuminates behavior and various effects of policies. Rather, the central claim is that we cannot understand relevant behavior or properly inform policy without an appreciation of how various pieces fit together. Moreover, we cannot readily leverage what is learned about some subjects to enhance our understanding of others and of the system as a whole.

This article sketches key features of a unified perspective that follows my book, Kaplow (2008a), and related work. Section II develops the core distribution-neutral framework that enables a modularized treatment of the efficiency and distributive effects of reforms.² Specifically, it is explained how, generically, various policy reforms (whether of commodity taxes, public goods, regulation, or much more) can be decomposed into two components. First is a distribution-neutral module that consists of the core reform (say, an increase in a public good) combined with an adjustment to the income tax and transfer schedule that renders the combined package distribution neutral. This package can be analyzed using standard efficiency tests (such as the Samuelson Rule for public goods). Second is a purely redistributive module that converts the aforementioned tax adjustment into whatever actual income tax modification was imagined to accompany the core reform. This component, constituting a purely redistributive change to the income tax system, can be analyzed as such.

As section II explains, there are a number of benefits of employing this two-step decomposition. To begin, analysis is clarified and results are more readily communicated. When core effects and redistribution are entangled—and in different ways in different investigations of the same policy instrument—it is difficult to compare results, and both

² Modularity here takes the meaning common in computer programming and complexity theory: analysis within each module can be undertaken independently even if an output of one is an input to the other or a higher level module combines the output of separate, lower-level modules. For an interesting essay on aspects of modularity by the Nobel-Prize-winning economist and polymath Herbert Simon, see Simon (1962).

researchers and policymakers can readily misinterpret the conclusions. The two-step decomposition avoids these problems. Moreover, the decomposition facilitates specialization, something most researchers appropriately do in any event, but enabling it to proceed in a legitimately grounded and more reliable manner. Analysts, say, of environmental policy need not simultaneously model and simulate income tax policy, which requires additional effort as well taking a stand on empirical matters like elasticities and the distribution of abilities and on the social welfare function. Third and related, difficult political economy assumptions are often embedded (at least implicitly) in analyses that eschew this decomposition—assumptions that differ across studies, are not empirically grounded, and may well be implausible.

Last, section II illustrates the analysis with the familiar differential commodity tax problem. It is explained how a variety of reforms can be assessed without assuming either that the income tax is optimal or confining attention to reforms in the neighborhood of the commodity tax optimum. In a basis case with weak separability of labor effort in the utility function, the optimality of uniformity obtains; moreover, proportional reforms in the direction of uniformity can be implemented so as to generate a strict Pareto improvement, as can other conventionally efficiency-enhancing reforms. The role of separability is examined, emphasizing that relaxing the assumption does adjust the efficiency test, on account of a fiscal externality, but does not undermine the proposed modularity and thus the benefits of the two-step decomposition that begins with distribution-neutral analysis. It is also explained how, despite Atkinson and Stiglitz's (1976) article, other, inconsistent methods largely dominated public economics through the 1990s and to an extent beyond, practices that are now increasingly being displaced by work that builds on or often is inspired by a newer view that is reflected in the distribution-neutral approach.

Section III examines a wide range of applications. It begins with the familiar subject of tax expenditures, which remarkably had not been analyzed as an application of differential commodity taxation; when it is, important corrections to conventional wisdom become apparent. Next examined are externality-correcting taxes and public goods, two areas that have received substantial attention in the second-best public economics literature but that, until more recently, had failed to disentangle core efficiency effects from (often implicit) redistribution. Additional applications include capital income and wealth taxation, social security and retirement savings, estate and gift taxation (reprising the above illustration involving gifts), and the rather different but quite important problem of the optimal design of transfer programs. As will be seen, the core framework illuminates how best to identify the distinctive effects of different policy instruments and sharpens our understanding of the targeting principle regarding how they should optimally be used—in many cases, by following the dictates of domain-specific efficiency tests, leaving distribution to the income tax and transfer system.

Section IV extends the analysis to consider individuals' utility functions and the social welfare function, which underlie the welfare economic approach to policy analysis. Following Kaplow and Shavell (2001, 2002), it is explained how the Pareto principle substantially restricts normative analysis to a much greater degree than is widely appreciated. These lessons (and some of my own prior work) are then used to clarify the appropriate (and more limited) role of notions and measures of horizontal equity, mobility, equality, and poverty. This framing is then used to illuminate difficult normative issues regarding taxation of the family, heterogeneous preferences, and problems of tax administration, compliance, and enforcement—all realms in which it is challenging but critical to trace policies' effects to individuals' utilities and thus ultimately to measures of social welfare. One cannot identify sufficient statistics for welfare analysis without articulating and analyzing the underlying determinants of welfare.

This article addresses a wide range of topics and maps many of their interrelationships. Although the framing is largely normative, much of the analysis—and the core virtues of the two-step decomposition used throughout sections II and III—is positive. It enhances clarity and rigor, and it aids in setting empirical research agendas and interpreting estimates of the effects of policy experiments. The exposition throughout is cursory in a number of ways, with many qualifications and omissions left unstated.³ The presentation is also largely informal (with hints throughout of the formal arguments) and, in light of the article’s origins, highly self-referential (see Kaplow 2007c, 2008a for fairly comprehensive references to those dates). The purpose here is not to give a comprehensive account but rather to articulate a framework that is useful in a range of settings and that, due to its elemental character, makes it easier to criticize, qualify, and improve as research progresses.

II. EFFICIENCY AND REDISTRIBUTION: A FRAMEWORK FOR ANALYSIS

This section presents a framework for analyzing a wide variety of policy instruments and relating them to social objectives. Section II.A articulates a two-step decomposition that separates the core features of (non-purely-redistributive) policies, which can be assessed using efficiency tests, from the redistributive dimension, which is subject to analysis following Mirrlees (1971). Some readers may recognize this proposed modularity of the two types of analysis from Musgrave’s (1959) suggestive distinction between what he termed the Allocation and Distribution Branches of government.

Section II.B elaborates the advantages of employing this decomposition: conceptual clarity, specialization, and detaching considerations of political economy. The framework developed here also helps to formalize the idea that instruments should be properly matched to objectives and to indicate how specialized research can more effectively be conducted. Section III.C illustrates the methodology by analyzing commodity tax reforms and relating the method advanced here to that of Atkinson and Stiglitz (1976) and also to prior literatures that adopt neither the present approach nor that of Atkinson and Stiglitz.

A. Two-Step Decomposition

Suppose that we wish to evaluate some change in policy—perhaps in the level of a public good, a commodity tax, or a regulation—which we will denote as ΔP . Assume further that this is imagined to be accompanied by some change to the income tax and transfer *schedule*, ΔT^P . The only restriction on ΔT^P is that the package as a whole, $\Delta P + \Delta T^P$, is *budget neutral*. (Note, for example, that if ΔP is a regulatory reform that has no effect on the government’s budget, ΔT^P could be the null policy, but it need not be.) Our task is to assess the aggregate policy, $\Delta P + \Delta T^P$.

The analysis will be conducted in a standard Mirrlees (1971) setting. Individuals differ in their (unobservable) earning ability and choose labor effort (and make any other domain-relevant choices, depending on the policy being considered) so as to maximize their (common) utility functions, taking policies and the nonlinear income tax schedule as given.⁴ There is some

³ The article sets to the aside (in some cases entirely) behavioral economic considerations, problems of administration and enforcement, empirical work and related practicalities of matching evidence to the theory, political economy concerns, and macroeconomic considerations.

⁴ Extensions for heterogeneous utility functions and different family types are considered in section IV.

individualistic social welfare function (SWF)—that is a welfare function that depends (only) on individuals’ utilities—which, as we will see, does not need to be specified further for much of the analysis because the methodology developed here enables Pareto comparisons. Moreover, the analysis is largely independent of the initial income tax and transfer schedule T , including whether it is set optimally according to some SWF. Indeed, T will not appear in our analysis.

The fairly general methodology elaborated here requires that we introduce only one further element. Define ΔT^{DN} to be the adjustment to the income tax and transfer schedule such that the aggregate policy, $\Delta P + \Delta T^{DN}$, is distribution neutral. Specifically, ΔT^{DN} is set at every level of income such that $\Delta P + \Delta T^{DN}$ provides individuals who initially earned that income the same level of utility as before. That is, ΔT^{DN} is the compensating variation for policy change ΔP at every level of income. Note that, for a marginal change one can determine this compensation in a straightforward way because any adjustments to an individual’s behavior have no effect on the individual’s utility by the envelope theorem. (For example, for the increase in a public good examined in section III.C, the income tax adjustment would be an increase equal to the marginal rate of substitution at each income level.)

Note that this package as a whole is actually more than just distribution neutral. It keeps each individual at the same utility level, which also means that we need not choose some definition of “distribution neutral” that allows comparisons across distributions with different levels of income or of utility. As a further note on terminology, when we consider a tax schedule adjustment that creates a package that is, as a whole, distribution neutral, it is helpful to describe this tax adjustment in isolation as distributively offsetting, or an offsetting income tax adjustment.

The main lesson of this section can now be stated immediately using the above notation and definitions:

$$\Delta P + \Delta T^P = (\Delta P + \Delta T^{DN}) + (\Delta T^P - \Delta T^{DN}).$$

That is, we can rewrite (decompose) the policy we seek to evaluate, $\Delta P + \Delta T^P$, into two components, $\Delta P + \Delta T^{DN}$ and $\Delta T^P - \Delta T^{DN}$. These components, in turn, can be viewed as steps, wherein we first imagine implementing the policy comprising the first parenthetical term, instantaneously followed by implementing the policy comprising the second:

$$\Delta P + \Delta T^P = \underbrace{(\Delta P + \Delta T^{DN})}_{\text{Step 1}} + \underbrace{(\Delta T^P - \Delta T^{DN})}_{\text{Step 2}}.$$

It is also useful to restate this expression verbally:

- Step 1. Combine the policy in question (sans finance) with a distributively offsetting adjustment to the income tax and transfer schedule.
 - Efficiency test.
- Step 2. Transform the foregoing (hypothetical) income tax schedule into the actually proposed income tax schedule.
 - Redistribution assessment.

This simple two-step decomposition has remarkably useful properties. Step 1 can be assessed *entirely on efficiency grounds*. No matter what the SWF, because Step 1 as a package is distribution-neutral, there is no redistribution to account for. Relatedly, there is no change in

labor supply distortion associated with any change in redistribution to account for either. The basic tradeoff embodied in the Mirrlees problem is entirely absent. By contrast, Step 2 is a *purely redistributive adjustment* to the income tax and transfer system. It *is* the Mirrlees problem and should be assessed accordingly. Let us now elaborate each of these features.

Step 1: As mentioned, Step 1 is a pure efficiency test. In particular, it yields a simple way of assessing policies (within Step 1): Focus on whether there is a budget surplus or deficit.⁵ If there is a surplus, one could rebate it, say, pro rata and generate a strict Pareto improvement. If there is a deficit, the same exercise would make everyone worse off, but in that case reversing the policy (implementing $-\Delta P$ and likewise reversing the sign of the distribution-neutral income tax adjustment, creating $-\Delta T^{DN}$) would, for marginal policies, enable a Pareto improvement—and, for nonmarginal policies, a reversal, or some partial version, would enable a Pareto improvement.

Moreover, it is often possible to understand fairly directly whether Step 1 involves a budget surplus or deficit. Recall that the income tax adjustment, ΔT^{DN} , is the compensating variation at each level of income. Integrating over the population, the net change in income tax and transfer revenue is the total compensating variation in the population. As will be explained below, this yields fairly simple (Economics 101, or perhaps 201) policy rules in basic settings. For a public good, if the sum of individuals' willingness to pay exceeds the cost of providing the good, there will be a surplus. For a commodity tax reform, if the move is in the direction of efficiency (in basic cases, toward uniformity), the net of changes in commodity tax revenue and in the revenue from ΔT^{DN} will be positive. And so forth. (Qualifications are discussed in section II.C.)

Step 2: To see that Step 2 is, indeed, a purely redistributive adjustment to the income tax and transfer system, simply examine the term $\Delta T^P - \Delta T^{DN}$. It consists of the difference between the actually imagined income tax adjustment, ΔT^P , and our hypothetical distribution-neutral income tax adjustment, ΔT^{DN} . Obviously, this is purely a change in the income tax schedule. Moreover, the particular change is indicated by the manner in which the imagined tax adjustment, ΔT^P , itself departs from distribution neutrality (viewing the policy package as a whole). If there is, say, a simple increase in the overall degree of redistribution, we would have whatever social welfare gain is associated with that and a concomitant distortionary cost, and conversely for a simple decrease in overall redistribution. More broadly, whatever is the deviation from distribution neutrality, evaluation of Step 2 is a pure version of the Mirrlees exercise. And when the policy under examination and its imagined means of finance involve marginal changes, Step 2 is a simple perturbation of the income tax schedule.

As a final note on Step 2, recall the point that the initial income tax and transfer schedule, T , need not be specified (only ΔT 's are in our expression). Relatedly, it was not assumed that the initial T was optimal. Of course, if T was optimal and, moreover, if the perturbation in Step 2 was marginal, then Step 2 would have no effect on the overall welfare evaluation of the policy package, $\Delta P + \Delta T^P$. In all other cases, as with the pure Mirrlees problem, we would need to choose an SWF and determine how much change in distortion was associated with the change in redistribution entailed by $\Delta T^P - \Delta T^{DN}$ in order to assess Step 2.

Finally, let us reflect on this two-step decomposition. The most important thing to note is that the claimed modularity is a tautology (although the term is often understood as derogatory, it

⁵ Put another way, the change in the government's budget is a sufficient statistic for welfare analysis. To be sure, no simple observable tells us what this change will be, but in many applications elaborated below this may be easier to predict than might have been thought because the underlying determinants are fairly straightforward.

is better to be in the position of defending a tautology's validity than to be attacking it). To be sure, the construction in which this decomposition is embedded does entail certain assumptions, notably, common utility functions.⁶ But it is still quite general, particularly regarding the policies it encompasses. Also, this decomposition does *not* assume that individuals' utility functions are weakly separable in the disutility of labor: as will be elaborated in section II.C, that assumption can be useful for tractability and clarity but relaxing it calls for a simple adjustment to the efficiency analysis in Step 1 rather than any qualification to the modularity embodied in the two-step decomposition.

B. Advantages of the Framework

Conceptual clarity: Use of the two-step decomposition is advantageous to researchers and policymakers in a number of ways. First and foremost, conceptual clarity is enhanced by the two-step decomposition and its concomitant separation of the distinctive efficiency features of a variety of policies from their (in many respects generic) redistributive effects. Economists frequently decompose complex formulas into more elemental components to ease communication and understanding. The distinction between efficiency and distribution has a long history, rooted in the fundamental theorems of welfare economics and featured in Musgrave's proposed taxonomy of government functions.

The flipside is equally important to consider: when different types of effects are entangled, it can be difficult to understand any of them, and our impressions of the whole may readily be confused, often without our realizing it. This point can be illustrated by the challenge of comparing studies of a common instrument, say, an increase in a tax on fuel.

Suppose that study A finds that an increase in the fuel tax reduces distortion while study B finds that it raises distortion. If the two studies use different ways of balancing the government's budget, with different distributive incidences, it can be very difficult to determine what we have learned. Perhaps study A's reduction in distortion arose despite its findings of adverse efficiency effects from higher fuel taxes as such because the policy package under examination used the higher fuel tax revenue to reduce the progressivity of income taxes, causing distortion from the latter to fall by an even greater amount. Perhaps study B found that higher fuel taxes themselves improved efficiency, but the revenue was used in a highly redistributive way that raised total distortion by even more. Furthermore, if both studies, in the main text, emphasized modeling and data relating to the fuel tax, rendering details on the income tax, labor supply, and so forth to an appendix, and presented complex simulations of policies with only bottom-line results, it would be challenging even for an expert reader to untangle these differences, and most of them would be missed by other consumers of the research. Over time, as more studies of fuel tax policies accumulated, there may be ever increasing latent knowledge, but little of it may be absorbed or reconciled by anyone. And policymakers might easily be led astray by what seems to have been learned from the research.

By contrast, if each study used distribution-neutral finance—focusing on Step 1 in the two-step decomposition—all differences would be due to the modeling of and data relating to fuel tax policies themselves. Significant complexity and disagreement may remain, but the far greater comparability of methods and conclusions would enhance understanding and facilitate

⁶ Specifically, the construction requires that ΔT^{DN} exists, and for the precise Pareto statements in particular, this assumes either commonality of utility functions (so that, at each observed level of income, the compensating variation is the same for all individuals) or that any differences can be observed (for example, they may be a function of age, family composition, or an observable disability). Heterogeneity is explored further in section IV.D.

progress. This conclusion holds even if particular researchers' beliefs or specific contemporary proposals in the political sphere do not envision distribution-neutral implementation. That is, focusing on Step 1 does not in any way reflect a belief in the reality of distribution neutrality; rather, the virtue is in the method of analysis and the conceptual clarity it brings.

Comparability yields huge advantages, and in order for a wide range of researchers to generate comparable results it is necessary for there to be a benchmark, or focal point, for adjustments to the tax and transfer system. Moreover, as the foregoing discussion indicates—and the next segment on specialization reinforces—there are substantial, distinctive advantages to choosing distribution neutrality (ΔT^{DN}) as *the* benchmark from among the infinity of ways that the income tax could be adjusted to achieve budget neutrality. Indeed, this choice uniquely enables the separation of efficiency and redistribution. As will be seen in section II.C and the applications in section III, the method greatly facilitates the analysis of a broad range of policies.

Specialization: Use of the two-step decomposition also greatly facilitates specialization. Some economists can focus on fuel taxation or other energy and environmental issues, others on education, others on health care, others on the taxation of capital, . . . , and others on redistribution through the tax and transfer system. This is a somewhat paradoxical lesson of the view advanced here and in much of my prior work. On one hand, it emphasizes relating everything to foundations, and to everything else, but on the other hand, it enables specialization. The reconciliation is that we can better know how to specialize when we appreciate how different system components interact and, in particular, when we can determine what sorts of modularization are appropriate. Specialization is often useful even when different domains are not fully modular, or even close to it, but specialization is especially effective when there is substantial modularity, as exists here when the two-step decomposition is employed.

From the perspective of researchers who do not primarily study redistribution through the income tax and transfer system, the benefits are clear. In addition to achieving a sharper understanding through comparability, there is also much less work to do. If one eschews distribution-neutral analysis, then one must undertake all that is involved in the redistributive income tax problem every time: taking a stand on labor supply elasticities, the distribution of abilities, and the choice of a social welfare function. Moreover, as will be elaborated momentarily, there is a third, even more challenging field that the researcher must engage when departing from an analytical benchmark like distribution neutrality: choosing the particular tax adjustment, ΔT^P , to pair with the policy under analysis, ΔP , which entails an exercise in political economy. That is, alternative approaches require more work to produce results that, as per the prior point, are more difficult to interpret both for subsequent researchers and for policymakers.

Those who study optimal income taxation and transfer programs already tend to specialize, on Step 2. At that point, an analyst or policymaker, in the true spirit of modularity, can take the outputs from different, specialized research modules and combine the results to assess a particular policy, $\Delta P + \Delta T^P$, when $\Delta T^P \neq \Delta T^{DN}$, by summing the results from Step 1 and Step 2.

Political economy: We can see that there are significant sacrifices to conceptual clarity, specialization, and resulting progress in research and communication to policymakers when investigators eschew the two-step decomposition and analyze policy packages in which $\Delta T^P \neq \Delta T^{DN}$. The most plausible justification for ever doing so is political reality: if it is believed that in fact the reform to be implemented, if any, will be $\Delta P + \Delta T^P$, then it may seem appealing to analyze that package as a whole.

Actually, not so: with the two-step decomposition, one does fully analyze just that package, and more easily, with the added benefit of providing a clearer picture to policymakers.

Indeed, perhaps the policymaker who is inclined to like the package as a whole will realize that the $\Delta P + \Delta T^{DN}$ component is undesirable, causing a shift to just the redistributive portion, $\Delta T^P - \Delta T^{DN}$. Or it may learn that $\Delta P + \Delta T^{DN}$ is highly desirable, so it may be superior to implement it alone, foregoing the unfavorable or unpopular redistributive component of $\Delta T^P - \Delta T^{DN}$ (or otherwise modifying ΔT^P to better meet distributive objectives).

Furthermore, just how is it that the researcher knows what ΔT^P should be associated with a particular ΔP ? Many policies considered are entirely hypothetical, variations on those policies are often examined as well, and even policies under active consideration by the government are often obscure regarding ΔT^P . And when they are explicit, that choice may change as the debate proceeds, the tax change may be one that would have been implemented regardless, and the income tax and transfer schedule may well be reformed more broadly in short order, sweeping away the interim effects of the particular ΔT^P associated with a particular ΔP . These reasons relate to the preceding comment that a third specialty, political economy, is required if one is to analyze particular ΔT^P 's because they are real. In practice, of course, this work is not generally done. The many studies that do not employ the two-step decomposition and instead analyze only a combined package, $\Delta P + \Delta T^P$, do not purport to defend their choice of ΔT^P using empirical evidence on prevailing redistributive politics. Relatedly, the fact that different researchers embed different ΔT^P 's in their analyses of common ΔP 's suggests that at least some of these predictions, if they are taken as such, are mistaken.

If one did have to advance a generic political economy perspective in this domain, perhaps $\Delta T^P = \Delta T^{DN}$ is as good a conjecture as any, particularly in the long run. At any given time, a political system gravitates toward some sort of equilibrium regarding redistribution policy, and the system's view of redistribution tends to respond to aggregates. Hence, over time, as various reforms are enacted, there will be a tendency to maintain some such equilibrium. Of course, that equilibrium is likely to evolve over time, but a similar point about balances of redistributive forces can be made about this trajectory.

Note further that there is another appeal to ΔT^{DN} in this realm: implementation of only Step 1's $\Delta P + \Delta T^{DN}$ entails a Pareto improvement when the reform is an efficient one, so it is a package with potential political appeal. Consider, for example, the 1986 U.S. tax reform that purported to be distribution as well as revenue neutral, aiming to achieve efficiency gains through myriad changes in particular provisions that broadened the income tax base and lowered marginal rates.⁷ And when Step 1 is inefficient, this means that there is not enough gain to buy off losers, which makes implementation more difficult.

The foregoing is simplistic and in some respects Panglossian. It is well known that politics is highly imperfect even in the best of times and that there are systematic illusions about policies' effects, distorting influences of special interest groups, and other factors that muck up the works.⁸ Moreover, heeding my own lesson, I am entirely inexperienced in political economy. Hence, readers should take these conjectures not as good approximations of reality but rather as spurs to thinking. The main takeaway here is that, for researchers who are not experts in political prognostication, there is yet another reason to employ the two-step decomposition with

⁷ As will be discussed in section III.A, however, analyses of broad-based reforms that reduce tax expenditures are often misleading in ways that are illuminated by the two-step decomposition.

⁸ In this regard, the discussion earlier in this section about how the failure to employ the two-step decomposition is confusing and can readily mislead policymakers is apt. Perhaps policy illusions, particularly regarding matters involving distribution, would not be as severe if analysis was more digestible.

Step 1's distribution-neutral approach. As noted, this also means that those not engaged in optimal income taxation or the politics of redistribution can skip these subjects altogether.

C. Illustration: Commodity Tax Reforms

To make the two-step decomposition more concrete, this section considers its application to commodity tax reforms. The focus will be on Step 1 because, as already explained, Step 2 is simply the familiar Mirrlees problem (which, although challenging, presents nothing new). The discussion first outlines how the literature evolved and then sketches the application itself.

1. *History of Thought*

The analysis of differential commodity taxation was long associated with Ramsey (1927) and, in the simplest case, the inverse-elasticity rule. The standard formulation used a representative individual and hence entailed no concern for distribution. Relatedly, income taxation was not one of the admissible instruments.⁹ Analysis of this model and various extensions are still featured in textbooks and surveys. See Myles (1995), Auerbach and Hines (2002), and Salanié (2011).

In parallel with the increase in theoretical study of optimal taxation in the 1970s following Mirrlees (1971), it is notable that much work outside the income tax as such—even by Mirrlees and other prominent scholars associated with the development of optimal income taxation—stuck mostly to variations of the representative-individual framework that did not allow an income tax. See, for example, Diamond and Mirrlees (1971), Stiglitz and Dasgupta (1971), and Atkinson and Stern (1974).

This changed—but not really—with Atkinson and Stiglitz (1976), who showed that if there was an income tax and it was set optimally, then optimal commodity taxes entailed no differentiation if labor was weakly separable in the utility function.¹⁰ Their proof made use of the first-order conditions for the two related problems.

Interestingly, it took decades for Atkinson and Stiglitz (1976) to begin to have a substantial influence on work in public economics. For example, at the start of the 1990s, substantial literatures on the second-best analysis of the optimal provision of public goods and of optimal environmental taxation and regulation stuck with the representative-individual framework. Many papers departed from classic Ramsey models in allowing an income tax

⁹ Standard formulations do not allow for a nonzero intercept—which, in a representative-individual model, would have rendered further analysis moot because a lump-sum tax would be optimal. (Statements such as that in Salanié (2011) that introduce treatments of the Ramsey approach by observing that a wage or income tax is allowed but must be linear are misleading because of omission of the further restriction that there be a zero intercept. Actually, that restriction, not linearity, is necessary to generate the core results in that literature.) Note that a uniform commodity tax (or an extracted average level of commodity taxes) is equivalent to a linear income tax with a zero intercept. For elaboration, see the discussion of literature on environmental taxation in section III.B and, more broadly on the relationship of Ramsey models to current literature that admits an income tax, see Atkinson and Stiglitz (1976), Stiglitz (1987), Mirrlees (1994), and Kaplow (2008a, ch. 6D).

¹⁰ They also remarked on the implications of nonseparability, discussed in subsection 2, and surprisingly stated the implications backwards (that leisure complements should be subsidized rather than taxed), an error variously replicated in subsequent texts (Myles 1995 and Salanié 2003) but that (fortunately) did not seem to influence most economists' understanding of the problem. For a formal treatment, tracing the misunderstanding to a misinterpretation of the sign of the costate variable in the Hamiltonian for the optimal income tax problem, see Kaplow (2010b).

(often a linear one) that was adjusted in various ways to balance the budget. But these tax adjustments differed within and across papers and produced myriad results that were not generally in the spirit of what might have been extrapolated from Atkinson and Stiglitz (1976). That article usually was not cited and its methods were not employed. It does not appear that the fundamental shift in analysis and implications from the Ramsey framework was appreciated at least through the 1990s.¹¹ Perhaps one reason for this widespread omission was the commonly held belief expressed in Boadway and Pestieau's (2003, p. 400) essay—specifically on the Atkinson and Stiglitz article, in a festschrift for Stiglitz—which closed with the statement: “without optimal income taxation there is no A–S theorem.”¹² That is, the result was thought to depend on the optimality of the income tax, and since few believed that this assumption held in reality, all bets were off.

This implication, however, was fundamentally mistaken. After all, the intuition for the optimality of no differentiation (elaborated below) is fairly robust. Moreover, it is not generally sound to suppose that relaxing an assumption throws out all that one has learned from a foundational model. Instead, we understand that the results may change but usually suppose that the manner in which they do can be traced to the role of the assumption and how it is relaxed.

Meanwhile, there emerged another line of work (in what was initially a parallel universe) that began with an underappreciated paper by Hylland and Zeckhauser (1979) that argued against using distributive weights in cost-benefit analysis. Their article, which did not cite Atkinson and Stiglitz, employed the sort of distributively offsetting income tax adjustment presented above. With some exceptions, this alternative approach lay largely dormant into the 1990s, particularly in public economics.¹³

Starting in the 1990s, I extended this approach and applied it directly to a range of questions in the field: my 1993 NBER working paper (published as Kaplow 1996c) examined public goods (and briefly externalities); Kaplow (1998b, 2001a) investigated the taxation of gifts and bequests relating the problem to differential commodity taxation; my 2004 NBER working paper (published as Kaplow 2006a) and my *Journal of Economic Perspectives* paper (Kaplow 2004) presented and discussed the generalizations of Atkinson and Stiglitz (1976) sketched in the next subsection. Since then, much of my work on public economics (including my book, Kaplow 2008a) and a growing portion of the literature by many researchers directly embeds analysis in a framework that combines Mirrlees (1971) and Atkinson and Stiglitz (1976), sometimes making explicit use of the distribution-neutral approach (Step 1) that does not require the income tax to be optimal. In addition, complementary work, such as Slemrod and Yitzhaki

¹¹ For example, the text by Myles (1995, p. 100) opens its extensive treatment of commodity taxation, which employs the Ramsey framework, by stating that relaxing the restriction disallowing a uniform lump-sum tax (which “is assumed [to be unavailable] for simplicity” would “not significantly modify the conclusions,” and the discussion of Atkinson and Stiglitz (1976) in a later chapter is brief and does not mention any implications for the results on commodity taxation. For references discussing the implications, see note 9.

¹² In addition, a number of leading scholars reacted to some of my papers in the 1990s and 2000s by questioning whether they could be right or even insisting that they were wrong, purporting to offer counterexamples. It would seem that these economists were either unaware of or failed to appreciate Atkinson and Stiglitz (1976), a further reflection of the lack of penetration of their approach during this period.

¹³ The method was applied to legal rules by Shavell (1981), who drew on Hylland and Zeckhauser and likewise did not cite Atkinson and Stiglitz. As best I can tell, for some time no one was aware of both lines of work, at least not sufficiently to identify and exploit the synergy.

(2001) and Hendren (2014, 2016), employs integrated approaches that do not use the two-step decomposition.¹⁴

2. Analysis

The exposition here is an abbreviated form of Kaplow (2006a) and Kaplow (2008a, ch. 6). To make the setting explicit (even though no formal derivations will be offered), suppose that individuals differ in their ability (wage rate) w , and that they choose their labor effort l and levels of commodities x_1, \dots, x_n to maximize the utility function

$$u(x_1, \dots, x_n, l)$$

subject to the budget constraint

$$\sum_{i=1}^n (p_i + \tau_i)x_i = wl - T(wl),$$

where p_i is the price and τ_i is the tax (or, if negative, subsidy) on commodity x_i , wl is the income earned by an individual of type w who exerts labor effort l , and $T(\cdot)$ is the initially given (and not necessarily optimal) nonlinear income tax and transfer schedule.

Let us now examine a policy reform. We will start with arbitrary but nonuniform commodity taxes. For concreteness, it might be helpful to suppose that there are higher tax rates on some luxury goods and lower rates on some necessities. Our reform, ΔP , will be to eliminate all commodity tax differentials, and to ease the exposition suppose that this is implemented by eliminating all commodity taxation.

The offsetting income tax adjustments, the schedule ΔT^{DN} , that we need to construct will involve higher taxes at every income level. This income tax increase has two components. First, imagine that individuals continued to earn the same (before-tax) income and consume the same consumption bundles as before. Because they no longer pay commodity taxes, they can afford to buy more of everything, so our income tax adjustment will need to tax away all of this savings. Second, even when that is done, all individuals will in fact be better off, so their income taxes will have to be raised on this account as well since our overall tax adjustment is designed to hold them to the same utility level. Why is everyone indeed better off before this second income tax adjustment? The reason is that we now have different price ratios from those prevailing before because differential commodity taxation has been eliminated. Hence, starting at the initial consumption bundle (which, after the first component of our income tax adjustment, individuals can still just afford), all will choose different bundles. By revealed preference, this raises their utility, so indeed they must pay more income tax to be held to the same utility level.

Let us take stock of where we are so far. Individuals are all at the same utility level. From the first component of the income tax adjustment, the government breaks even (it raises more in income tax revenue precisely the amount that it no longer collects from commodity taxes), and from the second component of the income tax adjustment, the government generates

¹⁴ Most alternative approaches that seek to address both efficiency and distribution combine in some fashion measures of the marginal cost of public funds (MCPF) and the marginal value of public funds (MVPF), which requires explicit use of distributive weights. For a recent application, see Hendren and Sprung-Keyser (2019).

a surplus.¹⁵ This surplus can now be rebated (say, pro rata) so as to make everyone better off, a strict Pareto improvement.

Note further that section II.A's characterization of this analysis in Step 1 as an Economics 101 "efficiency test" is entirely apt. Consider that the magnitude of this surplus is precisely equal to the second component of the income tax adjustment, which in turn is just the amount necessary to tax away individuals' utility increases due to the restoration of the economy's price ratios to their undistorted levels. That is, individuals, but for this part of the income tax adjustment, would all have had a utility gain just equal to the efficiency gain. The total utility gain thus equals the total efficiency gain in the economy, and this is what (translated into dollars) constitutes the budget surplus available for the rebate.

The careful reader may have noticed an important gap in the foregoing argument: specifically, what if individuals, as a consequence of the policy experiment, $\Delta P + \Delta T^{DN}$, are induced to change their labor effort? In that case, there would be another, third source of change in income tax revenue that must be accounted for. To analyze this potential fiscal externality, it is helpful to proceed in two steps: first, with weak separability, and second, with that assumption relaxed.

First, assume for simplicity that the utility function is weakly separable in the disutility of labor effort. That is, the utility function can be written as $u(v(x_1, \dots, x_n), l)$: individuals have a common subutility function $v(\cdot)$ that indicates the utility they derive from consumption. The core implication of this separability assumption is that the disutility of labor may depend on the overall level of (sub)utility from consumption but not on the particular combination of commodities that generates that level of (sub)utility. Conversely, the level of labor effort required to generate the disposable income necessary to fund consumption does not affect the (sub)utility derived from any given consumption bundle.

With weak separability, it can be demonstrated that there is no change in any individual's labor effort, in which event the above analysis is the complete story. To see why labor effort does not change, note that, if labor effort is indeed constant, then the tax adjustment ΔT^{DN} that keeps an individual of any type, and thus any level of income earned, at the same utility level is one that generates the same level of subutility, $v(\cdot)$, as prevailed before the reform. (Indeed, when the argument is developed formally, the tax adjustment is defined such that this is so.) Now, consider the labor effort choice of any individual: to each l and hence each level of disposable (after-income-tax) income—the right side of the above budget constraint—there is associated the same subutility $v(\cdot)$, as before. Therefore, the same l will continue to maximize that individual's utility.

Second, consider how this very simple (Economics 101) efficiency test needs to be modified if we relax the weak separability assumption, returning to our original utility function $u(x_1, \dots, x_n, l)$. Now, labor effort may rise or fall. Note, however, that the same Step 1 income tax adjustment, ΔT^{DN} , would be correct for a marginal reform, by the envelope theorem.¹⁶ In any case, the change in labor effort itself will generate a fiscal externality to the extent that the marginal tax rate on labor income is positive, as is generally the case. If labor effort rises, therefore, there will be more income tax revenue than otherwise, suggesting that the reform is

¹⁵ Note that individuals' consumption reallocations do not affect commodity tax revenue because our experiment sets all such taxes to zero. It can be shown that if taxes were made uniform at some nonzero level, the same bottom-line results would obtain.

¹⁶ For nonmarginal reforms, a revealed preference argument like that employed above would be applicable, suggesting that ΔT^{DN} would raise more revenue on that account.

more efficient than otherwise, whereas if labor effort falls, there will be less revenue and lower efficiency.

Which will be the case depends on the familiar intuition associated with Corlett and Hague (1953). Suppose, for example, that the reform raises the relative tax rate on basic foods, like rice or flour. Assume further that home-prepared meals are, as suggested by the evidence in Ab Iorwerth and Whalley (2002), leisure complements. Then eliminating this preference for these necessities would make leisure less attractive, which increases labor effort. That effect further boosts income tax revenue, making the hypothesized reform more efficient than otherwise. Indeed, under this assumption, it would be optimal to tax basic foods at a higher, not lower rate than that applied to other commodities.

Relaxing weak separability thus adds another component to the efficiency test, but the analysis thereof also reinforces basic intuitions and is wholly consistent with the modularity suggested by the two-step decomposition. Specifically, this analysis can be conducted entirely as part of Step 1 using a pure efficiency test, with potential Pareto improvements indicated by the change in the government's budget as a consequence of undertaking the policy experiment $\Delta P + \Delta T^{DN}$. This is true regardless of the distributive incidence of ΔP or of ΔT^{DN} standing alone. If one wishes instead to consider a policy experiment consisting of $\Delta P + \Delta T^P$, one can, as argued previously, employ the two-step decomposition, confining the (pure) analysis of redistribution to Step 2, which examines $\Delta T^P - \Delta T^{DN}$. To emphasize this point: nothing in the two-step decomposition itself depends on weak separability. That simplifying assumption matters only for how Step 1 is conducted: whether the efficiency test includes a supplemental fiscal externality term.

Second, separability clarifies the intuition behind the determination of optimal commodity taxes in the presence of a distorting labor income tax. A loose but misleading statement of the general theorem of the second best is that, once there are other distortions in the economy, no policy prescriptions can be offered. A more moderate and more accurate, but not always very helpful, version is that once there are other distortions, we must simply analyze everything, turn the crank, and see what pops out.

It is best to focus directly on the underlying logic of the welfare-analytic challenge generated by the existence of multiple distortions. A second distortion *might* be optimal when there is another distortion, but this *will* be so only when it helps to offset it. Accordingly, additional distortions will be optimal in the presence of a distorting income tax and transfer system when they act so as to reduce the labor-leisure distortion. Because the weak separability assumption breaks the connection between consumption choices and the labor-leisure choice, there is no room for improving (or worsening) the labor-leisure distortion through differential commodity taxation. Therefore, we are back to basics, that is, our first-best, Economics 101 efficiency prescription that it is not optimal to distort consumption choices.¹⁷

Relatedly, when the weak separability assumption is relaxed, both the direction and magnitude of the optimal secondary distortion will depend on the interaction with the labor-leisure distortion. Notably, it has nothing to do with the distributive incidence of the ΔP one is considering. This is well illustrated by the foregoing case of commodity tax preferences for basic foods: no preference is optimal with weak separability, and when separability is relaxed in the manner suggested by the evidence in Ab Iorwerth and Whalley (2002), it becomes optimal to

¹⁷ There are, of course, other qualifications. See Mirrlees (1976) and Kaplow (2008a, ch. 6C). Many reinforce the core ideas in the text. For example, just as nonseparability allows the possibility that expenditures on a commodity may be correlated with unobserved labor effort, so it is that preferences that depend on ability allow expenditures on a commodity to be correlated with unobserved ability.

(relatively) tax rather than subsidize basic foods—keeping in mind that, either way, distribution is being held constant because ΔT^{DN} is defined by the ΔP under consideration.¹⁸

More broadly, as will become clear in Section III's broader set of applications, the direction and magnitude of the nonseparability adjustment relates entirely to the fiscal externality determined by the interaction with the labor-leisure choice and has nothing to do with anything that may be regarded as special about the target of the adjustment. Here, it does not matter for the direction or magnitude of the optimal correction whether the good is a luxury or a necessity, and this point about the irrelevance of the distributive incidence of the contemplated ΔP standing alone is entirely general. It likewise does not matter whether the commodity under consideration is a so-called dirty good or one that generates positive externalities. As discussed in section III.B, relative to the otherwise-optimal first-best Pigouvian tax or subsidy, the correction depends, again, only on the nature of the nonseparability. Qualitatively (and in respects quantitatively), the nonseparability correction is the same as what would be optimal if the good involves no externality, as implicitly assumed here.

Finally, before concluding this illustration of the two-step decomposition, it is worth noting that the method can be used to analyze all manner of commodity tax reforms, not just the complete elimination of any differentiation. Conventionally, analysts following Atkinson and Stiglitz (1976) focus on local reforms, using first-order conditions, so they are examining the neighborhood of uniformity. But the two-step approach outlined here is not nearly so limited: indeed, no first-order conditions characterizing the optimum—whether of the nonlinear income tax problem or of the commodity tax problem—are needed. The above discussion considered a nonlocal reform that removed all commodity taxes and subsidies. One can (reverting for ease of exposition to the case of weak separability) likewise show that any partial reform that moves proportionally in the direction of uniformity enables a Pareto improvement. For example, cutting all commodity taxes and subsidies by half, or by any other factor, has this property.

More broadly, it can be demonstrated that any change in commodity taxes generates a Pareto improvement in the Step 1 experiment, consisting of $\Delta P + \Delta T^{DN}$, if and only if that commodity tax reform passes a pure efficiency test. See Kaplow (2006a; 2008a, ch. 6B). Moreover, my other work shows that this sort of encompassing result has analogues in a range of settings, from externality correction (Kaplow 2012) to competition policy (Kaplow 2019).¹⁹ Therefore, the sense in which the analysis of Step 1 can be regarded as a pure efficiency test is quite general. There is indeed substantial traction to the modularity advanced throughout this section, wherein we can decompose a range of policy problems into two steps, one concerned solely with efficiency and one solely with redistribution.

¹⁸ The magnitude of the optimal second-best adjustment will depend on the elasticity of food expenditures and on the strength of the interaction with the disutility of labor effort. Contrary to the Ramsey prescription, however, the magnitude of the elasticity does not bear on the sign of the optimal relative tax rate; a higher elasticity favors less of a deviation from neutrality, in whichever direction the optimal deviation may be.

¹⁹ The key insight was previously stated: the ability to give a rebate that generates a Pareto improvement depends on whether the change in the government's net revenue is positive, and the change in income tax revenue (with weak separability) is given by the integral of individuals' compensating variations, which capture the efficiency effects of the policy. For example, in Kaplow (2019), an intermediate step in the proof shows that the change in the government's budget is given by the sum of the changes in consumer and producer surplus (the firms being owned by individuals, perhaps mostly by those with high incomes). See also the discussion of public goods in section III.C.

III. APPLICATIONS

This section further illustrates the power of the two-step decomposition by sketching a wide array of applications. Most of these reside in Step 1. Myriad policy instruments can be understood more clearly—regarding both their distinctive effects and their optimal design—when distribution-neutral implementation is contemplated. Consistent with section II’s analysis, efficiency tests and the targeting principle hold over a broad domain. A final application, to income transfer policies, elaborates Step 2 of the framework.

A. Tax Expenditures

The subject of tax expenditures is one of many that had long been treated as a subject unto itself. Yet it should be apparent that tax expenditures can readily be analyzed using the two-step decomposition. See Kaplow (2017). Indeed, tax expenditures are a form of differential commodity taxes and subsidies—typically the latter—wherein a tax credit is a subsidy at the credit rate and a tax deduction is a subsidy at the taxpayer’s marginal income tax rate. Hence, in our benchmark case with weak separability of labor, it is optimal to eliminate tax expenditures.

Exceptions would be for the usual reasons, such as when a positive externality is associated with the expenditure (see section III.B), an often advanced justification for the charitable contribution deduction, energy conservation provisions, and some others. Another broad class of exceptions concerns the tax expenditure concept itself. Deductions that are part of defining income are not regarded as tax expenditures. For example, a sole proprietor whose store generates \$1,000,000 in gross receipts but incurs costs (of goods sold, rent, salaries, utilities, etc.) of \$950,000 is (and should be) permitted to deduct these expenditures, yielding taxable income of \$50,000. Tax expenditures are taken to be defined as deviations from this benchmark.

Even though many of these ideas are well understood, it is important to embed tax expenditure analysis in the comprehensive framework that employs the two-step decomposition, something not usually done. For example, recurring proposals to substantially reduce tax expenditures are often advanced and analyzed because they are regarded to reduce labor supply distortion (by lowering marginal tax rates as part of the package), enhance redistribution without the usually associated efficiency trade-off, and more. See Feldstein (2015) and Burman et al. (2017). Yet such is an illusion. To illustrate some of the problems, suppose that there was a flat-rate income tax of 40% and that everyone undertook deductible tax expenditures equal to 5% of their income. These deductions could be eliminated, enabling a reduction of the tax rate to 38%. But labor supply distortion would not drop in any direct manner because the previous marginal *effective* tax rate was only 38% to begin with. And if rates were reduced, say, only to 39%, using the revenue from the other 1% to fund a larger lump-sum grant, redistribution would increase but so would the concomitant labor supply distortion ordinarily associated with greater redistribution. The only direct gain, distinctive to the elimination of otherwise unwarranted tax expenditures, would be the efficiency enhancement from avoiding the distortion of expenditure choices—that is, precisely the efficiency benefit in Step 1 that is associated with the elimination of differential commodity taxation.

B. Externalities

The core intuition for why differential commodity taxation is inefficient in basic settings is that the price ratios faced by consumers should equal the ratios of the marginal resource costs of producing the commodities. In the presence of externalities, the logic is that the price ratios should equal the ratios of the marginal social (not private) resource costs. This favors first-best Pigouvian taxes and subsidies equal to the marginal external cost or benefit associated with the commodity. One might also adduce associated principles for the correction of “internalities” (Gruber and Koszegi 2001).

As outlined in Kaplow (1996c, 2004) and demonstrated formally in Kaplow (2012), this simple logic extends to second-best settings with a redistributive income tax, which may be demonstrated using the two-step decomposition. Because the logic should now be familiar, the exercise will not be repeated.

Nevertheless, much second-best environmental economics literature, starting in the 1990s, focuses on the question whether there might be a “double dividend” from corrective environmental taxation, wherein the government could collect extra revenue while correcting rather than causing further distortion.²⁰ Subsequent work explored various special cases, often finding that the interaction with the income tax was adverse, calling for optimal corrections below the first-best prescription. For a collection and a survey, see respectively Goulder (2002) and Bovenberg and Goulder (2002).

As the present framework suggests, however, this literature almost had to be misleading in some fashion. In fact, it often employed representative-individual models, so that there could be no social welfare gains or losses associated with changes in redistribution itself, yet the analysis also featured an income tax that was adjusted in various ways in connection with different reforms. Hence, in basic cases in which further distortion was identified, we might expect that the posited change in the income tax made the system as a whole (the environmental policy reform combined with the income tax adjustment, $\Delta P + \Delta T^P$) *implicitly* more redistributive (ΔT^P was more redistributive than ΔT^{DN}).

Indeed, the literature often examined precisely such packages, as discussed in Kaplow (2012). A key lesson is that, if one is using a representative-individual model and embedding an income tax to more realistically account for society’s (unmodeled) concerns for distribution, it is essential to conduct a Step 1 distribution-neutral analysis (i.e., as if there were individuals of different earnings abilities). And if a non-distribution-neutral analysis (in this sense) is undertaken instead, one really does have to append Step 2, associating any change in the distortion of labor supply with the implicit change in redistribution and then accounting for the latter with an SWF. It is misleading to report a finding that “welfare” falls due to an increase in labor supply distortion from a more redistributive system and, moreover, to regard aggregate changes in distortion as measuring “social” welfare, without even identifying the implicit increase in redistribution. (This example well illustrates many of the points in section II.B about the advantages of employing the two-step decomposition.)

As a final way to drive this point home, consider the following example (which is similar to that underlying some of the environmental taxation literature as well as, *mutatis mutandis*, that on public goods, considered in the next section). The utility function is additively separable in utility from the consumption of commodities (which include a clean and a dirty good), harm

²⁰ Note the similarities to the aforementioned free lunches often thought to be associated with the elimination of tax expenditures.

from an atmospheric externality, and the disutility of labor effort. For curvature parameters associated with an upward-sloping labor supply curve, raising the tax on the dirty good while lowering a linear labor income tax rate so as to maintain budget balance has the following properties: labor supply falls, the benefit of controlling the externality (although identical for everyone in utils) rises less than proportionately with disposable income when the benefit is measured in dollars, and the magnitude of the utility cost from altered consumption is rising with disposable income. That is, the posited policy experiment of further controlling the externality reduces labor supply but also unambiguously increases redistribution. Not surprisingly, in Step 2 reduced labor supply and increased redistribution go hand in hand. When using the two-step decomposition, wherein this coupling is isolated in Step 2 and thus made explicit, it is hard to lose sight of this core principle in the manner that some of the otherwise-sophisticated second-best literature over the past half century sometimes has.

C. Public Goods

Public goods are subject to a similar analysis that, again, the two-step decomposition makes simple. See Kaplow (1996c, 2004, 2006b). As explained in section II.A, the distribution-neutral income tax adjustment, ΔT^{DN} , is the compensating variation for the underlying policy change, ΔP . Hence, if utility is weakly separable in labor, the tax revenue raised from ΔT^{DN} when a public good is marginally increased is the sum (integral) of individuals' marginal rates of substitution. Therefore, there will be a Step 1 budget surplus if and only if the Samuelson Rule (1954) is satisfied. Adjustments for nonseparability and redistributive income tax adjustments in Step 2 (for $\Delta T^P \neq \Delta T^{DN}$) are as before. As with the analysis of externalities, prior work—often employing Ramsey models or supplementing them with income tax adjustments that implicitly changed the extent of redistribution—neither confined itself to Step 1's distribution-neutral approach nor used the two-step decomposition to disentangle redistribution and hence (at least appeared) to produce different conclusions that were, in these respects, misleading. See, for example, Atkinson and Stern (1974), Diamond and Mirrlees (1971), and Stiglitz and Dasgupta (1971), and the subsequent elaboration and survey in Ballard and Fullerton (1992).

The juxtaposition of public goods and externalities is itself instructive. Subtle second-best literatures developed independently—without cross-citations—on each through the 1990s and beyond, without recognizing that in most relevant respects they are the same problem. Application of the two-step decomposition, however, clarifies the overlap: obviously, the purely redistributive Step 2's are the same. And the Step 1 rules are not merely both simple, Economics 101 efficiency tests (in basic cases), but in a sense they are the same tests. After all, most of the environmental economics literature models atmospheric externalities, and clean air (etc.) is a public good, the value of which is the sum of individuals' marginal rates of substitution. Relatedly, if one thinks of the Samuelson Rule (1954) in terms of Lindahl (1919) pricing, then deviations can be viewed as employing (implicit) price ratios that do not reflect the proper social valuations, just as with corrective taxes and subsidies that deviate from the first-best Pigouvian prescription.²¹ Moreover, both public goods' and externalities' interactions with labor supply, once separability is relaxed, are much the same. For example, nature preserves are public goods that may be leisure complements, and these preserves, in turn, may be polluted by emissions.

²¹ Hence, it is no surprise that the example in section III.B with additive separability is precisely the same whether the (separable) benefit is regarded as a public good or an externality (often taken to be negative).

Therefore, policies that improve the value of these resources—whether by improving access or by reducing pollution—should be analyzed along all dimensions using the same methods.

D. Capital Income and Wealth Taxation

The optimal taxation of savings also has a long lineage. Interestingly, Atkinson and Stiglitz's (1976) original paper on commodity taxation in the presence of an income tax explained the direct implication for the optimal taxation of savings: with weak separability, the optimal rate is zero. This follows immediately from applying their model to a world with two commodities, x_1 and x_2 , where the subscripts refer to two time periods: period 1 is the present (one's working life), in which labor is also supplied, and period 2 is the future (retirement), in which the individual only consumes. The optimality of no differentiation implies that individuals should face no wedge on their allocation of after-labor-income-tax income between present and future consumption.

Remarkably, it was decades before this important lesson substantially penetrated. For example, the models of optimal capital income taxation associated with Chamley (1986) and Judd (1985) are not embedded in a Mirrlees model and do not incorporate the differential commodity taxation insight.²² See Golosov, Tsyvinski, and Werning (2007) and Kaplow (2008a, ch. 9). In parallel with the increasing invocations of Mirrlees (1971), Atkinson and Stiglitz (1976), and work on the two-step decomposition, this situation had changed substantially by the 2000s. There is now a dynamic Mirrlees literature (Golosov, Tsyvinski, and Werning 2007) that examines labor supply over multiple periods in the presence of earnings uncertainty (which may generate precautionary savings that is excessive in light of the negative fiscal externality in later periods) and some other extensions as well—all in models of capital income taxation that admit the income tax as an instrument.²³ As always, the core point is not whether the simple benchmark (here, an optimal tax rate of zero) is correct or even approximately so, but rather the use of a framework that makes analysis of the problem both more transparent and ultimately more rigorous at the same time.

Another value of this depiction of the problem is that it facilitates the analysis of a wide range of related policy instruments. Most obvious is a wealth tax, which is understood to be equivalent to a capital income tax in simple settings. For example, if the real return on capital is 4%, a 50% capital income tax and a 2% annual wealth tax are equivalent. In basic cases, this equivalence also holds in models extending Domar and Musgrave's (1944) treatment of systematic risk that allows for portfolio adjustments. See, for example, Bulow and Summers (1984) and Gordon (1985). Kaplow (1994) shows how, in a simple general equilibrium model, a proportional capital income tax with full loss offsets is equivalent to an ex ante wealth tax or to an ex post wealth tax, with these in turn equivalent to a tax on only the riskless return to capital.

²² At a presentation of one of these papers, I raised a question to the effect "What about Atkinson-Stiglitz?" The query did not register with the author or, as best I could tell, with others in the audience.

²³ For example, some attention has been devoted to the qualification mentioned in note 17 regarding correlations between expenditures on particular commodities and ability; here, the idea is that higher capital income may signal a higher labor-income-earning ability. See, for example, Saez (2002). It is worth reflecting, however, on the empirical predicate. All but the highest income and ability taxpayers tend to confine investments to mutual funds, whereas the highest income individuals (and perhaps, as the argument goes, the highest ability types within an income level) actively manage their investments, with the aid of expensive advisors, producing net (fee- and risk-adjusted) returns that may well underperform the market, which suggests a negative correlation. Also important is that what is sometimes classified as capital income of top earners—the business income of founders—may better be understood as labor income, a point discussed later in this section.

Moreover, this equivalence is not merely *ex ante* but also entails equal after-tax income for taxpayers and equal government tax receipts in every state of the world.²⁴

The actual taxation of capital income is highly complex and differs substantially across jurisdictions and over time.²⁵ Different types of entities may be subject to different tax regimes, an entity's debt and equity may be taxed asymmetrically, rules regarding depreciation and expensing may differentially tax different types of investments, capital gains may be taxed at different rates from other capital income and subject to realization requirements, and more. A helpful way to address these and other issues is to employ a modest extension of Step 1's protocol: in addition to considering distribution-neutral income tax adjustments, ΔT^{DN} , it is also helpful to consider reforms of various elements of capital taxation that also keep the wedge on savings constant. See Kaplow (2008a, ch. 9). In that way, one can set aside questions of the optimal level of capital income taxation and of redistribution, focusing on the most efficient way to tax capital income at any particular overall (effective) rate.²⁶

To further illustrate the power of the two-step decomposition in this realm, consider incomes (disproportionately at the top of the income distribution) that may derive from investments in firms that earn large markups due to imperfect competition. To the extent that such prospective profits are capitalized and accounted for in *ex ante* investment decisions, it is not evident that the analysis changes. In any event, Kaplow (2019) examines a model in which there may be rents due to markups in various sectors and determines the implications for both competition policy and redistributive income taxation. The conclusion in a basic model (with weak separability of labor) is that competition policy should maximize total surplus—the sum of consumer and producer surplus—regardless of how skewed the ownership shares of the latter may be: again, a standard efficiency test. Moreover, the optimal income tax problem is qualitatively similar (it can be mapped directly into the standard Mirrlees problem), although the optimal schedule itself differs (perhaps substantially) depending on features of the rent-generating process and the distribution of ownership shares.

As a final note, it seems that one important impetus for higher capital income or wealth taxation is that much that appears in these categories (on income tax returns or Forbes measures) is in essence labor income that may never have been taxed. This can arise from income generated in the underground economy, but in developed countries this is most associated with entrepreneurial income, particularly of highly successful founders. Smith et al. (2019) find that much of the income of the top 1% and a large portion of the increase in their share since 2000 is in the form of business income of pass-through entities that flows through to active owners. Until the 2017 U.S. tax reform, such income was taxed at the (here, top) labor income tax rate, but there are many settings (often involving different legal structures) in which this may not in fact be the case. This problem requires further research given the inevitable entanglement of what might conventionally be viewed as labor and capital income. For example, what is the optimal taxation of the rise or fall in the value of ownership shares that are retained by active managers as incentive pay and because their sale is deterred by asymmetric information?²⁷

²⁴ It also extends to the schemes of retrospective capital gains taxation developed by Auerbach and Bradford (2004).

²⁵ Significant issues involved with international taxation are omitted here.

²⁶ On another dimension, it is also important to distinguish optimal long-run or steady-state capital income taxation from the capital levies or windfalls that may be associated with transitions from one regime to another, which some important work in the field does not fully disentangle. For elaboration, see Kaplow (2008a, ch. 9) and Kaplow (2009).

²⁷ Accounting for the many losses (which receive far less publicity than the spectacular gains of a few) and making appropriate risk adjustments are first-order considerations. Hall and Woodward (2010) find that, in their preferred

E. Social Security (and Retirement Savings Policies)

Many aspects of social insurance schemes—here, the focus is on provisions for retirement funded by payroll taxes—can be analyzed using a similar approach. See Kaplow (2008a, ch. 11). First, consider the two-step decomposition itself, which strips out, placing in Step 2, any redistributive aspects that are present. This dimension can be assessed in the usual way, to a first approximation.²⁸

This leaves, in Step 1, what we will take to be an actuarially fair program that is tantamount to forced savings. Looking at any given level of income (and keeping in mind our distribution-neutral framing, so that each income level may be considered separately), the minimum savings constraint may or may not be binding. If it is not, the policy and marginal adjustments thereto have no effect. If it is binding, then we can analyze the constraint as equivalent to a capital income subsidy, set at a level that is just sufficient to produce the forced level of savings. Tightening the constraint is, therefore, equivalent to raising the subsidy. As discussed in the preceding section, in a simple benchmark case this policy would be inefficient, and moving toward uniform taxation (a zero wedge on capital, and thus no binding forced savings constraint) would improve individuals' welfare. Note also that, by linking the analysis to that of savings subsidies, the optimal design of various retirement savings policies is illuminated as well.²⁹

Other benchmarks would, of course, be appropriate under other assumptions pertaining to optimal capital income taxation. It is interesting in particular to incorporate the primary justifications for mandatory social insurance schemes and to consider their implications within this framework (Kaplow 2011, 2015a, 2015b). One is the existence of the Samaritan's dilemma (Buchanan 1975), wherein other individuals (relatives in particular) or the government will be unable to resist aiding those who have provided inadequately for their retirement, the anticipation of which induces rational individuals to save too little *ex ante*. This is a form of externality that is subject to standard corrective analysis, as developed in section III.B.

Another justification involves self-control problems, such as individuals' myopia leading to savings that are suboptimal from the perspective of their "true" utility function (Bernheim and Rangel 2007). Here, we have an internality, the analysis of which (again) is similar to that of an externality, in which case promoting savings may raise welfare. If everyone is homogeneous, a capital subsidy and forced savings would have the same effect. If, instead, only some are subject to myopia (or the degree of myopia varies), then it may be optimal to force some savings rather than employ a savings subsidy, the latter of which distorts those who would already have saved adequately.

Once one departs from the neoclassical framework in this manner, it is also necessary to account for how forced savings or capital subsidies may influence individuals' labor supply

scenarios, founders of Silicon Valley firms supported by leading venture capitalists approximately break even on an *ex ante*, risk-adjusted basis.

²⁸ One qualification relates to the arguments that follow: once one introduces, say, behavioral considerations, it is possible that they will amend the redistribution problem as well. Another relates to intergenerational redistribution and risk sharing.

²⁹ Retirement savings subsidies are both often a form of differential capital income taxation and also a form of tax expenditure (if one takes as a baseline a Haig-Simons income tax rather than a cash flow consumption tax, a point of controversy in stating tax expenditure budgets). Making such connections helps to provide an integrated analysis and to leverage what is already understood about closely related subjects.

differently from in the standard treatment. If individuals' savings decisions are guided by a behavioral utility function different from their normative ("true") utility function, it may not be obvious what utility function governs their labor supply decisions. This question could be of great consequence. After all, payroll taxes that fund social insurance are of a similar order of magnitude (for all but wealthy taxpayers) to income taxes. Moreover, myopia's implied underweighting of the future might be imagined to lead individuals to treat currently incurred social insurance taxes as taxes but to heavily discount the distant future benefits, which might be imagined to produce a much larger labor wedge than otherwise. This and other possibilities are modeled in the aforementioned literature, where it is shown that, depending on the assumptions (on which there is little empirical evidence), this is indeed possible, but other seemingly plausible cases cover a wide range of outcomes, including no labor supply effect and even a boost in labor supply (a positive fiscal externality from stronger forced savings).

F. Estate and Gift (Wealth Transfer) Taxation

Estate and gift taxation—wealth transfer taxation as distinguished from wealth taxation—also had a tradition of being treated as a subject unto itself. As indicated in the Introduction, however, it is best to view gifts of all sorts as a species of consumption and thereby integrate their analysis into the broader framework. Although estate and gift taxation is typically confined to high-wealth taxpayers and hence is strongly associated with redistribution, when we employ the two-step decomposition we can see the taxation of estates and gifts as a distinct policy instrument that, at a given level of (donor) income, differentially taxes own-consumption and transfers that fund the consumption of others. Accordingly, the optimal setting of the instrument depends on the differences between these two forms of consumption. This perspective on the subject was initiated by Kaplow (1998b, 2001a) and has been followed in many of the leading treatments since then. See Fahri and Werning (2010) and Kopczuk (2013).

As the Introduction mentioned briefly, gifts (typically to donors' children) are distinctive from own-consumption in three respects. First, there is a positive externality on donees; in informal terms, two individuals benefit from a single round of consumption of the same real resources. Even an altruistic donor does not place the weight on the donee that would an SWF that registers the utility of both the donor (which includes any altruistic or warm glow utility) and the donee. Some have regarded this as double counting (Diamond 2006), but in that event it must be either that we do not have an individualistic SWF (elaborated in section IV)—here, ignoring the utility parents obtain from their children (on reflection, a remarkable departure from respecting preferences)—or that children (even adult children) are not in the SWF, also a radical view and one at odds with many social policies that aim to help children.

Second, gifts ordinarily have an income (wealth) effect on donees that suppresses their labor effort. (Note that a marginal adjustment in donees' labor effort has no effect on their utility due to the envelope theorem.) Annual gifts could have this feature, as could a trust fund received early in life or the anticipation of a later inheritance. It is also possible in some cases that liquidity constraints would be relaxed, enabling entrepreneurship that generates higher incomes (Cox 1990, Holtz-Eakin, Joulfaian, and Rosen 1994a, 1994b), but presumably that is not the dominant effect across the population. In a world with an income tax, this effect generates a negative fiscal externality.

Third, Step 1 of our two-stage decomposition can be thought of here as designed to be distribution neutral in the parent (donor) generation, but the transfers, a form of voluntary redistribution, affect the marginal utility and utility levels of donees, which will be relevant

under typical SWFs to how they should be treated when examining optimal redistribution. In this instance, we have a spillover from our Step 1 analysis to that in Step 2. To some extent, gifts and bequests are probably equalizing among the children of given parents. But across society as a whole, greater amounts are typically received by individuals who are already better off or who are made so by the transfers themselves.

Combining these three considerations is not an easy task, in part because it can be important to disentangle different donor transfer motives, but the two-step decomposition makes it fairly clear in principle how to proceed in the standard welfare economic framework.³⁰ Moreover, we can see that this roadmap differs widely from how estate and gift taxes are discussed in society at large, although there are some loose analogues. For example, those concerned about equal starting points might be seen as advancing a version of the third point about the distributive impact in the donees' generation.³¹ And some advance arguments for estate and gift taxation (as well as for wealth taxation, examined in section III.D) on the grounds of a negative externality on the political process.³²

G. Income Transfers

Most of the foregoing analysis has focused on Step 1: determining the effects of policies in a distribution-neutral setting that focuses attention on any given policies' distinctive features, which usually pertain only to efficiency. This section takes up a particular and important aspect of Step 2, regarding redistribution, examining transfer policies aimed at the lower end of the income distribution. As emphasized in Kaplow (2007a; 2008a, ch. 7), viewing them in this manner is helpful, and Step 1 analysis is also valuable for analyzing some key issues.

The first and most important point is that, to a surprising degree, analyses of transfer programs are not integrated into the standard optimal income tax framework. Yet Mirrlees (1971) specifically commented on this application. If one wants to know how transfers should be designed—notably, their optimal levels and phaseouts—one can simply look to the left of a diagram of an optimal income tax and transfer schedule and take note of the intercept and slope. In Mirrlees (1971) and much subsequent work, the optimal grant was substantial and optimal marginal tax rates were high in the lower range of incomes. There are two primary explanations for the latter. First, marginal tax rates are inframarginal for all individuals with higher incomes. For marginal income tax rates near the bottom of the distribution, a large portion of the population is inframarginal, so substantial revenue is raised that can fund a more generous grant

³⁰ As noted in Kaplow (1998b; 2001a; 2008a, ch. 10), different transfer motives can be relevant to determining the effects of wealth transfer taxation on giving behavior and on donors' and donees' utilities.

³¹ However, as a practical matter most inheritances are received when children are in their 50s, and most estate, inheritance, and gift tax regimes exempt all support of children early in life, including parents' expenditures on good neighborhoods and on education, so the empirical connection between the notion of equal starting points and these policies is quite attenuated.

³² However, taxing high levels of wealth or wealth transfers encourages more own-consumption, here by the rich, so the result may be greater expenditures by the wealthy on politics. Also, it is familiar that autocratic regimes may have the most to fear from concentrations of wealth that can provide independent bases of power, which explains why some of them either attack or bring under state control potential sources of countervailing economic power. A proper assessment of such externalities, as well as optimal deployment of other instruments (such as campaign finance limitations and media deconcentration), is an important subject outside the scope of this article. The point is that, from this perspective, estate and wealth taxes are corrective taxes and hence, as a first cut (and ignoring the other factors identified in the text), they should be set equal to the otherwise uncorrected marginal harms associated with wealth transfers or wealth, whatever they may be.

without causing distortion (and the income effect generates a further positive fiscal externality). Second, the lost productivity from low-ability individuals whose work effort is distorted downward by high marginal rates is relatively small, by definition.³³

Some literature on the design of transfer programs, following Diamond (1980), has suggested the optimality of lower (or even negative) marginal tax rates at the bottom in models in which behavioral responses are only at the extensive margin. However, it does not seem realistic to rule out part time work, especially at the lower end of the income distribution. And such literature further assumes that individuals are unable to mimic lower types (for example, a doctor may not be a good construction worker), but it seems unlikely that most low-skilled individuals, who have negligible occupation-specific human capital, are incapable of performing jobs at a slightly lower level of skill. Yet other (empirical) research assesses a qualitatively different type of program wherein transfers are limited to those working, say, at least 30 hours per week (for example, Michalopoulos, Robins, and Card 2005), but the presumed observability of hours (and thus, implicitly, the wage as well) eliminates the information constraint on the Mirrlees problem altogether. Hence, if such an assumption held and one considered the unconstrained optimal design of transfer programs, even higher social welfare could be achieved through other means.

Taking a simple, integrated view of the design of optimal transfer policy yields some additional insights. One is that transfers may optimally be targeted (tagged; see Akerlof 1978), for example, by providing more generous support for the disabled or for single parents with young children. If one oversimplifies by taking these traits to be exogenous, the optimal income tax problem can be restated as determining a separate tax and transfer schedule—with its own grant and marginal rate structure—for each group, with the schedules implicitly linked by a common shadow price of government revenue.³⁴ One important implication is that thinking in terms of “phaseouts”—wherein any grant increase must be phased out via higher marginal tax rates—is inapt, among other reasons because more generous grants for some groups may be financed by higher tax rates on other groups.

Another important feature of the Mirrlees (1971) formulation is that the income tax and transfer schedule is understood to be comprehensive: it should be taken to include all taxes (income taxes, payroll taxes, a VAT) and all transfers (social insurance, welfare, EITC-like provisions in the income tax code, child credits). For any group, what matters for behavior and utility is the aggregate grant and the combined marginal tax rate at any income level, inclusive of both explicit taxes and phaseouts. In this setting, it is incoherent to think in a vacuum in terms of the optimal design of, say, the EITC or some particular transfer program. (Likewise, viewing the familiar “EITC trapezoid” in isolation is misleading.) For example, taking an existing transfer program (say, food stamps, now SNAP, and supposing that the food expenditure requirement is nonbinding), imagine that we decrease the phaseout rate at very low incomes, which extends the

³³ In a standard Mirrlees framework, optimal marginal income tax rates are always under 100%. Simulations suggest the optimality of rates as high as, say, 60%. This naturally involves significant work disincentives, and many favor provisions such as the EITC (see just below) for this reason. Consider, however, the cost of reducing the marginal tax rate by 10% (that is, 10 percentage points) between \$0 and \$10,000. Focusing just on the inframarginal effect, \$1000 in revenue would be forgone from each person who earns above \$10,000. Suppose that there are 100 million such individuals; then, the annual cost of this improved incentive is \$100 billion. There may be positive externalities from the enhanced work effort of those in the \$0 to \$10,000 income range (perhaps the positive role model effects significantly outweigh any losses from reduced parental supervision), but they would have to be quite large in light of this revenue cost to justify substantially lower marginal rates than typical simulations suggest to be optimal.

³⁴ This construction will be discussed further below, in section IV.C on taxation of the family.

eligibility range, calling for an additional phaseout at somewhat higher incomes. This reform is identical to an analogous expansion of the EITC. Likewise, enacting such a reform along with a corresponding reduction of the EITC would entail no change whatsoever for the pertinent eligible groups.³⁵

There are reasons for varied, specialized transfer programs, but it is helpful to analyze differences among them, in light of the foregoing, in a distribution-neutral fashion, that is, using Step 1. We can ask, for a particular group of individuals, whether it is better to provide a given level of transfer at a given level of income via cash or in kind. For example, should food stamps that are binding be offered in lieu of cash of the same amount? Should some assistance be in the form of housing or medical care? Should there be free or subsidized pre-school programs?

Arguments in favor (other than optics and politics) usually involve paternalism (behavioral concerns) and externalities (including, as in the discussion of gifts, that even parents who care strongly about their children care less than society does). In addition to direct, current benefits to children (the beneficiaries of many in-kind provisions), there are also future benefits to them, which parents may not adequately perceive, and spillovers, such as through reduced crime and positive fiscal externalities generated by greater earnings (Heckman 2006, Hendren and Sprung-Keyser 2019). Some of these benefits might be generated, as well or better, by unrestricted cash transfers, but others may not.

IV. UTILITY AND SOCIAL WELFARE

The analysis to this point, despite its normative focus on the optimal setting of various policy instruments, is for the most part positive as well. The two-step decomposition is simply a way of analyzing all manner of policy experiments in a way that brings clarity and facilitates specialization. Utility functions were important among other reasons because they determine individuals' behavior, which must be understood to predict the effects of policies. Explicit normative analysis, however, must take a stand on additional issues in order to reach a bottom line.

This section presents a perspective on these subjects and the core rationales behind them, drawing on Kaplow (2008a, chs. 3, 13-15) and my prior work. Naturally, normative analysis is pertinent to the choice of the SWF, examined first, but it is also relevant to such notions as horizontal equity, mobility, and equality. Furthermore, it is necessary to attend explicitly to the normative framework in order to address a range of important subjects, from the treatment of different family groupings to the analysis of tax administration and enforcement. It is impossible to identify sufficient statistics for welfare analysis when the underlying basis for evaluation is unspecified—and worse, when implied specifications cannot be grounded in any plausible SWF.

A. Welfarism and the Pareto Principle

It is conventional in welfare economics to employ an SWF of the following form (expressed here for a finite population):

$$SW(z) \equiv W(u_1(z), \dots, u_I(z)),$$

³⁵ The optics, however, are different: expanding phaseouts is regarded as an increase in the generosity of welfare whereas the corresponding expansion of the EITC is viewed as a tax cut.

where u_i is the utility of individual i in a society of I individuals and z is a complete description of the state of the world. This is called an individualistic SWF because it depends on, and only on, each individual's utility. Put another way, no aspect of z affects the level of social welfare other than through its effects on individuals' utilities.

This type of SWF is often employed in the analysis of optimal income taxation and more broadly, usually without explicit justification. It seems normatively plausible that a social evaluation of any state z should indeed depend on how the state affects each individual's utility, and moreover it is not clear why aspects of the state should matter to social welfare without regard to such effects. Indeed, it can be demonstrated that any SWF that is not an individualistic SWF—in particular, one that depends on anything other than individuals' utilities—violates the Pareto principle. See Kaplow and Shavell (2001, 2002). The logic is related to the foregoing intuition. More precisely, if an SWF did regard welfare to be higher in one state than another for a nonutility reason, and if the SWF was continuous, we could imagine a state in which that nonutility reason was inoperative but each individual's utility was higher by ε . For ε sufficiently small, the hypothesized nonindividualistic SWF would regard social welfare to be lower in this modified state even though every individual's utility was higher.³⁶

This restriction on the range of normatively relevant SWFs has many implications, some of which will be elaborated below.³⁷ For general distributive purposes, Kaplow (2007b) shows that proposed substitutes for utility—famously, Rawls's (1971, 1982) primary goods and Sen's (1985a, 1985b, 1997) capabilities and functionings—entail violations of the Pareto principle and related conundrums. Similar analysis can be applied to Musgrave's (1959) notion of "merit" goods, unless their in-kind provision is fully justified by the sorts of efficiency considerations elucidated in section III.G.

More subtly, analysts sometimes use social welfare weights as a shortcut, where what is weighted is income rather than utility. This practice can likewise prescribe policies that violate the Pareto principle although it often would not in the domains in which these constructs are employed. Another problem with welfare weights is that income (really, after-tax disposable income, i.e., consumption) is endogenous, and what really matters (marginal utilities and utility levels) is endogenous as well. However, for analysis in the neighborhood of an optimum or that seeks to map a Pareto frontier, this endogeneity may not cause problems. In any case, the deeper lesson is to exercise caution: unless one literally means that income or consumption rather than utility is what matters socially—and one is willing to violate the Pareto principle—one should keep in mind the underlying SWF from which any shortcuts are derived.³⁸

³⁶ Perhaps surprisingly, the Pareto principle also has implications for the form of the SWF within the individualistic class. Specifically, any SWF that is not utilitarian violates the Pareto principle. See Kaplow (1995; 2008a, ch. 14) for analysis, related ideas, and comments on the literature. Some economists will already be familiar with Harsanyi (1953, 1955), who offers two demonstrations that the SWF should be utilitarian, one using a veil of ignorance (first stated to imply utilitarianism in Vickrey 1945) and another extending the rationality axioms from expected utility theory to the SWF as well.

³⁷ For a survey emphasizing contrasting views, see Fleurbaey and Maniquet (2018).

³⁸ There has also been confusion about the welfare gains from redistribution associated with the concavity of utility functions and of the SWF (on which, see also the preceding footnote). These forces are disentangled in Kaplow (2008a, ch. 3; 2010a).

B. Horizontal Equity, Mobility, Equality, and Poverty

Beginning at least with Musgrave's (1959) discussion of horizontal equity, economists examining income taxation and other distributive policies have considered an array of desiderata, including various indices for mobility, equality, and poverty. It is evident that, if any of these measures are to be interpreted not merely as proxies for welfare-related factors but as social ends in themselves, they would be inconsistent with an individualistic SWF and thus in conflict with the Pareto principle.³⁹ It seems that such confusion often arises, and that economists advancing these sorts of criteria have been largely unaware of these and other difficulties. See Kaplow (1989, 1995, 2005).⁴⁰

The problem is probably most familiar with concerns about inequality, wherein it is understood that one way to improve equality scores is to tax the rich beyond the top of the Laffer curve (and by more than can be justified by externalities). Likewise, many concerned with reducing poverty appreciate the imperfections of poverty measures and that they only capture part of the population (ignoring, for example, the near-poor). But perhaps there would be less—or at least better directed—fuss with measures of equality and poverty if there was a greater recognition that these notions are indeed proxies. Moreover, because one often needs much of the directly welfare-relevant information to construct the better proxies (and because one often is discarding some of that information in the process), it usually makes little sense to take such indices as policy targets.

Notions of horizontal equity and mobility have generated greater confusion when they have been employed. For example, King's (1983) formal indexes of each of the two are actually identical except for a minus sign: that is, two purportedly desirable objectives have been defined such that they are opposites (Kaplow 1989). Stepping back, both of these criteria do seem best understood as loose proxies of varying application. Horizontal equity—usually defined in a question-begging fashion as calling for the equal treatment of equals, without any independent normative criterion for who are “equals” (see Westen 1990)—is a principle that is ordinarily satisfied by maximizing an objective function: if the inputs are identical, so (typically) should be the outputs. Moreover, unequal treatment is historically associated with invidious discrimination and corruption. Similarly, obstacles that inhibit mobility have often been imposed for similar reasons and tend to be inefficient. Hence it makes sense to be aware of seeming violations of horizontal equity and impediments to mobility, but this important point does not transform them into welfare-relevant policy objectives in and of themselves.⁴¹

³⁹ To take a subtle example related to the discussion of income-based welfare weights in section IV.A, Auerbach and Hassett's (2002) proposed measure of horizontal equity depends only on incomes, but because incomes are not utility and, moreover, incomes in an identified status quo are given independent weight, their measure violates the Pareto principle. See Kaplow (2001b).

⁴⁰ The titles of two of these articles suggest some of the issues: “Horizontal Equity: Measures in Search of a Principle” and “Why Measure Inequality?” For an overview and also for applications to sacrifice theories, the benefit principle, notions of ability to pay, and the use of notions such as the Haig-Simons definition of income as norms, see Kaplow (2008a, ch. 15).

⁴¹ Treating them as such can readily lead policy astray. For example, random inspections and related enforcement policies directly violate horizontal equity, but banning them would eliminate most law enforcement. And forcing random assignment of jobs and all rewards would maximize mobility by many measures but eviscerate incentives.

C. Taxation of the Family

Income tax systems and many transfer programs, as well as social insurance, often impose differential treatment that reflects family composition, depending possibly on whether an adult is single or married and often on the number of children. There are two frequent deficiencies in the analysis of policies along these dimensions, corresponding to the central thrusts in this article. The analysis here follows Kaplow (1996b; 2008a, ch. 12).

First, by failing to employ something akin to the two-step decomposition, and in particular to engage in Step 1's distribution-neutral analysis, researchers often conflate distributive issues (which are sometimes termed "vertical") with the question of how, if at all, different family configurations at a given standard of living should be distinguished. For example, it is often suggested that some forms of support for children should be phased out as income rises because needs are falling with income, or that preferential status, say, for households with a single adult along with children, are regressive because they deliver greater dollar benefits to higher-income households. But in the spirit of Step 1, it is helpful to compare households of similar income and ask, essentially at each level of income, what differentiation is optimal. For some of the reasons elaborated below, it often will be optimal to make larger dollar transfers on account of children at higher levels of income. Note that, in the framework of the two-step decomposition, we can think of these larger transfers as being funded by similarly high-income households with fewer dependents. In this sense, one can maintain distribution neutrality.

The foregoing comments, however, are suggestive and incomplete because of our second issue, relating to utility functions and the SWF. Specifically, it is not obvious just which households (say, that differ only in the number of children) are to be compared to which—that is, which income levels for different family configurations should be regarded as comparable. This problem has often been dealt with using so-called family equivalence scales, but that formulation uses proxies that are not directly grounded in the pertinent utility functions, so the discussion in section IV.A warns us that we are likely to go astray.

To launch the proper analysis, it is useful to have in mind another piece of apparatus that was introduced in section III.G on transfer programs. We can extend the Mirrlees framework to contemplate different income tax schedules—each with its own intercept and schedule of marginal tax rates—for each group. (In the analysis that follows, we will continue to set aside the important point that family composition is endogenous, so the schedules will affect such choices as marriage and divorce and decisions to have children.) For each of these schedules, there will be two sets of concerns. One (which has received significant attention in the literature) is the resulting distortion of labor effort. For example, if second earners in married couples have more elastic labor supply, optimal marginal tax rates on them may be lower (Boskin and Sheshinski 1983).

Less attention has been devoted to the focus here, which involves determination of the appropriate utility and welfare effects for those subject to each of these schedules. Echoing a point in section III.F on estate and gift taxation, both parents and children are individuals, and each of their utilities should ordinarily be taken to enter into the SWF. (If one is to use a "family" or "household" utility function for purposes of assessing social welfare, rather than for predicting behavior, it ought to be something like the sum of family members' utilities, which would be precisely correct under a utilitarian SWF.) When the problem is broken down to its fundamentals, a number of implications can be derived.

First, for a given level of resources, larger families tend to have a higher marginal utility of disposable income—which is to say, for a given allocation mechanism within the household, the marginal dollar will variously flow to individuals whose consumption is lower and hence whose marginal utility is higher when the household is larger.⁴² This point favors more generous provision to larger families. Moreover, brief reflection suggests that, *ceteris paribus*, the dollar transfers that would equalize marginal utilities across different-sized households would be larger at the upper end of the income distribution.

Second, there are often regarded to be economies of scale in household size, due to the (partially) public good nature of some forms of expenditure, such as on housing. It is widely believed that this factor unambiguously favors less generous treatment of larger households, but this point overlooks the fact that disposable income is for that very reason a more efficient utility generator in larger households, which cuts the other way. It can be shown that the curvatures of the scale economies function, utility functions, and the SWF determine which force dominates and to what extent.

Third, there is a further question of distribution within the household, say between two adults. It is commonly thought that unequal sharing—say, because in traditional households, the husband controls the allocation of resources and takes a disproportionate share of the marginal dollar—favors less generous treatment. (This argument is also advanced to favor individual taxation rather than treating married couples as a single unit, as is done in the United States.) This view, however, is at best misleading, also because of a failure to derive conclusions from first principles. Consider that, on one hand, a marginal dollar will mostly go to the favored spouse (who has a lower marginal utility of income), but, on the other hand, the portion going to the disfavored spouse delivers higher marginal utility because the baseline is so much lower. Again, depending on the curvatures of utility functions and of the SWF, either force could dominate.

The optimal design of the tax and transfer system with regard to the treatment of different family configurations is a complex issue. This is true because the foregoing utility- and welfare-related factors are many and subtle and because incentives to supply labor and to form, dissolve, and augment family units are endogenous to the tax and transfer system. Nevertheless, this highly challenging and contentious problem is brought into sharper focus both by use of the two-step decomposition and by tracing effects directly to individuals' utility functions rather than relying on loose and potentially misleading notions.

D. Heterogeneous Preferences

Much policy analysis assumes that all individuals have the same utility function.⁴³ If sources of heterogeneity can be observed, different tax schedules may be employed. When it is

⁴² It might also be the case that utility *levels* per person are lower, which is relevant under an SWF that is concave in utilities but not, for example, under a utilitarian SWF. It is sometimes said that the larger “family’s” (total) utility is lower, but this suggestion is misleading. For example, if in one family, two adults voluntarily and successfully give birth to or adopt a child, and moreover the child is healthy and well treated, it is hard to see in what sense that “family” is worse off than another with two otherwise identical adults who were unsuccessful in adding a child.

⁴³ Of particular relevance here, Step 1, in constructing T^{DN} , assumes that individuals' utility functions are the same (or differ in ways that are observable). When this is not so, the resulting distribution neutrality can only be approximate, say, within income levels, and precise Pareto comparisons are impossible. Ng (1984) discusses “quasi-Pareto” comparisons in this regard.

only possible to observe signals of underlying differences or if no pertinent traits are observable, the problem is more challenging.⁴⁴

Exploration of heterogeneity is difficult for another, qualitatively distinct reason: it is necessary to engage in interpersonal comparisons of utility, and, moreover, ones that cannot (as is usually done) be rendered tractable by assuming that everyone's preferences are the same. See Boadway et al. (2002) and Sandmo (1993). Another key point is that some literature introduces heterogeneity in one of many possible ways and derives characterizations that are attributed to heterogeneity as such. Yet, as developed in Kaplow (2008b), heterogeneity is itself a heterogeneous phenomenon. Even with simple functional forms, one individual can derive more utility from consumption than another in qualitatively different ways that, among other things, can have opposite implications for the direction of the difference in marginal utilities and hence in the direction of optimal (utilitarian) redistribution.

In some settings, policy moderation may be optimal when heterogeneity cannot be observed and hence directly compensated for. For example, Stern (1982) showed that error in income measurement favors less redistributive income taxation. The idea is that when, say, an observed signal of income may over- or understate true income, the welfare loss from taxing the former too much exceeds the welfare loss from taxing the latter too little. On the other hand, similar considerations may favor a more generous grant or provisions for those who are probably (but not certainly) disabled because errors of underprovision have a higher welfare cost than errors of overprovision, *ceteris paribus*.⁴⁵

E. Administration, Compliance, and Enforcement

Tax administration, compliance, and enforcement are extremely important features of tax system design, yet they have not received nearly the attention they deserve. Under one common approach, these concerns are simply included in a list of objectives—along with, say, efficiency, equality, and perhaps some other notions of fairness—that taxation aims to advance. The lack of either a common denominator or a link to individuals' utility and thus to a proper measure of social welfare renders this approach obscure.

A more determinate but also unsatisfying method is to ask how to set tax administration and enforcement instruments so as to maximize tax revenue. Whatever may be the virtue of such a command as part of the solution to a second-best delegation problem, net revenue is an unsatisfying core objective. Instead, administrative, compliance, and enforcement concerns need to be embedded in the sort of social welfare framework presented here.

Slemrod and Yitzhaki (1987), Kaplow (1990), and Mayshar (1991) emphasized the need to focus on welfare rather than revenue maximization. And some of the earliest work that builds on Allingham and Sandmo (1972) embeds optimal tax enforcement in a welfare-based framework. A more recent line of work that analyzes the elasticity of taxable income likewise takes this tack (Feldstein 1999, Slemrod and Kopczuk 2002). Kaplow (1996a, 1998a) analyzes accuracy, complexity, and enforcement, and the book by Slemrod and Gillitzer (2014) covers many of the modern developments. Nevertheless, much work remains to be done that traces the

⁴⁴ One of the arguments for in-kind provision of welfare is that it may help with screening when relevant differences are unobservable. See Blackorby and Donaldson (1988).

⁴⁵ All may not be equal, however, because increasing generosity causes more individuals to feign disability, so overprovision may be provided to a greater number of individuals (and to ones for whom the social welfare gain is not as high as, say, for those truly disabled but not by quite enough to properly qualify).

effects of both government actions and taxpayers' responses to individuals' utility and embeds the analysis in the broader optimal tax framework.

V. CONCLUSION

Most work in public economics is inevitably and appropriately specialized, focusing on particular policy instruments and analyzing one or a few margins of behavior. But it is essential to occasionally step back and consider how all the pieces fit together in order to synthesize what we know, to identify what we need to understand but do not, and to better guide policymakers. This article offers a unifying framework that aims to advance this integrative mission and enable greater progress through specialization that is grounded in fundamentals.

The primary tool of analysis is a two-step decomposition of policy reforms into a distribution-neutral component that can be analyzed entirely on efficiency grounds and a purely redistributive component that embodies the familiar tradeoff of redistribution and labor supply distortion. This decomposition greatly enhances conceptual clarity and the comparability of different research findings, indicates the appropriate contours of specialized research by identifying respects in which analysis is modular, and allows for separate analyses of both the optimality and political economy of redistribution.

These lessons sharpen positive analysis and provide simple benchmarks for a wide range of policies, including commodity taxation, tax expenditures, externality correction, public goods, capital income and wealth taxation, social security and retirement savings, and estate and gift taxation. Simple Economics 101 efficiency tests (supplemented by an additional term for any fiscal externality due to interactions with labor supply) support the general targeting principle relating instruments and objectives that are often associated with particular margins of behavior. The framework also illuminates the optimal design of transfer policies, which constitute an important aspect of the second, redistributive component of the decomposition.

Finally, normative analysis—including the identification of sufficient statistics for welfare analysis—requires the explicit tracing of policies' effects to individuals' utilities and the aggregation of utilities through a social welfare function. Although not widely appreciated, adherence to the Pareto principle requires that social welfare be a function of only of individuals' utilities, without any deviations such as by giving independent weight to horizontal equity, mobility, equality, and poverty (even though these concepts may in respects serve as proxies for underlying welfare-relevant considerations). Tracing effects directly to individuals' utilities is also essential for the assessment of policies involving the taxation of families, heterogeneous preferences, and tax administration and enforcement.

The simple yet encompassing framework advanced here takes up the daunting task of relating everything to everything else, and in particular to fundamentals, but delivers the good news that a great deal of specialized research may be conducted in ways that are both easier than often imagined and more rigorously grounded at the same time. These encouraging messages, to be sure, are overly simple and optimistic. Many subtleties were noted, and other qualifications should be easier to identify and explore due to the transparency of the analysis. Although the gap between what we in principle need to know and what we do know or can plausibly learn anytime soon will sometimes be large, such limitations need to be acknowledged and addressed forthrightly.

REFERENCES

- Ab Iorwerth, Aled, and John Whalley, 2002. "Efficiency Considerations and the Exemption of Food from Sales and Value Added Taxes." *Canadian Journal of Economics* 35, 166–182.
- Akerlof, George A., 1978. "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning." *American Economic Review* 68, 8–19.
- Allingham, Michael G., and Agnar Sandmo, 1972. "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics* 1, 323–338.
- Atkinson, Anthony B., and Nicholas H. Stern, 1974. "Pigou, Taxation and Public Goods." *Review of Economic Studies* 41, 119–128.
- Atkinson, Anthony B., and Joseph E. Stiglitz, 1976. "The Design of Tax Structure: Direct versus Indirect Taxation." *Journal of Public Economics* 6, 55–75.
- Auerbach, Alan J., and Bradford, David F., 2004. "Generalized Cash-Flow Taxation." *Journal of Public Economics* 88, 957–980.
- Auerbach, Alan J., and Kevin A. Hassett, 2002. "A New Measure of Horizontal Equity." *American Economic Review* 92, 1116–1125.
- Auerbach, Alan J., and James R. Hines, Jr., 2002. "Taxation and Economic Efficiency." In Auerbach, Alan J., and Martin Feldstein (eds.), *Handbook of Public Economics*, vol. 3, 1347–1421. Elsevier, Amsterdam.
- Ballard, Charles L., and Don Fullerton, 1992. "Distortionary Taxes and the Provision of Public Goods." *Journal of Economic Perspectives* 6 (3), 117–131.
- Bernheim, B. Douglas, and Antonio Rangel, 2007. "Behavioral Public Economics: Welfare and Policy Analysis with Nonstandard Decision-Makers." In Diamond, Peter, and Hannu Vartiainen (eds.), *Behavioral Economics and Its Applications*, 7–77. Princeton University Press, Princeton.
- Blackorby, Charles, and David Donaldson, 1988. "Cash versus Kind, Self-Selection, and Efficient Transfers." *American Economic Review* 78, 691–700.
- Boadway, Robin, Maurice Marchand, Pierre Pestieau, and Maria del Mar Racionero, 2002. "Optimal Redistribution with Heterogeneous Preferences for Leisure." *Journal of Public Economic Theory* 4, 475–498.
- Boadway, Robin, and Pierre Pestieau, 2003. "Indirect Taxation and Redistribution: The Scope of the Atkinson–Stiglitz Theorem." In Arnott, Richard, Bruce Greenwald, Ravi Kanbur, and Barry Nalebuff (eds.), *Economics for an Imperfect World: Essays in Honor of Joseph E. Stiglitz*, 387–403. MIT Press, Cambridge, MA.
- Boskin, Michael J., and Martin Feldstein, 1977. "Effects of the Charitable Deduction on Contributions by Low Income and Middle Income Households: Evidence From the National Survey of Philanthropy." *Review of Economics and Statistics* 59 (3), 351–354.
- Boskin, Michael J., and Eytan Sheshinski, 1983. "Optimal Tax Treatment of the Family: Married Couples." *Journal of Public Economics* 20, 281–297.
- Bovenberg, A. Lans, and Lawrence H. Goulder, 2002. "Environmental Taxation and Regulation." In Auerbach, Alan J., and Martin Feldstein (eds.), *Handbook of Public Economics*, vol. 3, 1471–1545. Elsevier, Amsterdam.
- Buchanan, James M., 1975. "The Samaritan's Dilemma." In Phelps, Edmund S. (eds.), *Altruism, Morality, and Economic Theory*, 71–85. Russell Sage Foundation, New York.
- Bulow, Jeremy I., and Lawrence H. Summers, 1984. "The Taxation of Risky Assets." *Journal of Political Economy* 92, 20–39.

- Burman, Leonard E., Eric J. Toder, Daniel Berger, and Jeffrey Rohaly, 2017. “Economic and Distributional Effects of Tax Expenditure Limits.” In Auerbach, Alan J., and Kent Smetters (eds.), *The Economics of Tax Policy*, 109–144. Oxford University Press, Oxford.
- Chamley, Christophe, 1986. “Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives.” *Econometrica* 54, 607–622.
- Corlett, W. J., and D. C. Hague, 1953. “Complementarity and the Excess Burden of Taxation.” *Review of Economic Studies* 21, 21–30.
- Cox, Donald, 1990. “Intergenerational Transfers and Liquidity Constraints.” *Quarterly Journal of Economics* 105, 187–217.
- Diamond, Peter A., 1980. “Income Taxation with Fixed Hours of Work.” *Journal of Public Economics* 13, 101–110.
- Diamond, Peter A., 2006. “Optimal Tax Treatment of Private Contributions for Public Goods with and without Warm Glow Preferences.” *Journal of Public Economics* 90, 897–919.
- Diamond, Peter A., and James A. Mirrlees, 1971. “Optimal Taxation and Public Production II: Tax Rules.” *American Economic Review* 61, 261–278.
- Domar, Evsey D., and Richard A. Musgrave, 1944. “Proportional Income Taxation and Risk-Taking.” *Quarterly Journal of Economics* 58, 388–423.
- Farhi, Emmanuel, and Ivan Werning, 2010. “Progressive Estate Taxation.” *Quarterly Journal of Economics* 125 (2), 635–673.
- Feldstein, Martin S., 1999. “Tax Avoidance and the Deadweight Loss of the Income Tax.” *Review of Economics and Statistics* 81, 674–680.
- Feldstein, Martin S., 2015. “Raising Revenue by Limiting Tax Expenditures.” In Brown, Jeffrey R. (ed.), *Tax Policy and the Economy*, vol. 29, 1–11. University of Chicago Press, Chicago.
- Fleurbaey, Marc, and François Maniquet, 2018. “Optimal Income Taxation Theory and Principles of Fairness.” *Journal of Economic Literature* 56 (3), 1029–1079.
- Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning, 2007. “New Dynamic Public Finance: A User’s Guide.” In Acemoglu, Daron, Kenneth Rogoff, and Michael Woodford (eds.), *NBER Macroeconomics Annual 2006*, vol. 21, 317–388. MIT Press. Cambridge, MA.
- Gordon, Roger H., 1985. “Taxation of Corporate Capital Income: Tax Revenues versus Tax Distortions.” *Quarterly Journal of Economics* 100, 1–27.
- Goulder, Lawrence H. (ed.), 2002. *Environmental Policy Making in Economies with Prior Tax Distortions*. Edward Elgar, Cheltenham, UK.
- Gruber, Jonathan, and Botond Koszegi, 2001. “Is Addiction ‘Rational?’ Theory and Evidence.” *Quarterly Journal of Economics* 116 (4), 1261–1305.
- Hall, Robert E., and Susan E. Woodward, 2010. “The Burden of the Nondiversifiable Risk of Entrepreneurship.” *American Economic Review* 100, 1163–1194.
- Harsanyi, John C., 1953. “Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking.” *Journal of Political Economy* 61, 434–435.
- Harsanyi, John C., 1955. “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility.” *Journal of Political Economy* 63, 309–321.
- Heckman, James J., 2006. “Skill Formation and the Economics of Investing in Disadvantaged Children.” *Science* 312 (5782), 1900–1902.
- Hendren, Nathaniel, 2014. “Efficient Welfare Weights.” NBER Working Paper No. 20351. National Bureau of Economic Research, Cambridge, MA.
- Hendren, Nathaniel, 2016. “The Policy Elasticity.” In Brown, Jeffrey R. (ed.), *Tax Policy and the Economy*, vol. 30, 51–89. NBER and University of Chicago Press, Chicago.

- Hendren, Nathaniel, and Ben Sprung-Keyser, 2019. "A Unified Welfare Analysis of Government Policies." NBER Working Paper No. 26144. National Bureau of Economic Research, Cambridge, MA.
- Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen, 1994a. "Entrepreneurial Decisions and Liquidity Constraints." *Rand Journal of Economics* 25, 334–347.
- Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen, 1994b. "Sticking It Out: Entrepreneurial Survival and Liquidity Constraints." *Journal of Political Economy* 102, 53–75.
- Hylland, Aanund, and Richard Zeckhauser, 1979. "Distributional Objectives Should Affect Taxes But Not Program Choice or Design." *Scandinavian Journal of Economics* 81 (2), 264–284.
- Judd, Kenneth L., 1985. "Redistributive Taxation in a Simple Perfect Foresight Model." *Journal of Public Economics* 28, 59–83.
- Kaplow, Louis, 1989. "Horizontal Equity: Measures in Search of a Principle." *National Tax Journal* 42, 139–154.
- Kaplow, Louis, 1990. "Optimal Taxation with Costly Enforcement and Evasion." *Journal of Public Economics* 43, 221–236.
- Kaplow, Louis, 1994. "Taxation and Risk Taking: A General Equilibrium Perspective." *National Tax Journal* 47, 789–798.
- Kaplow, Louis, 1995. "A Fundamental Objection to Tax Equity Norms: A Call for Utilitarianism." *National Tax Journal* 48, 497–514.
- Kaplow, Louis, 1996a. "How Tax Complexity and Enforcement Affect the Equity and Efficiency of the Income Tax." *National Tax Journal* 49, 135–150.
- Kaplow, Louis, 1996b. "Optimal Distribution and the Family." *Scandinavian Journal of Economics* 98, 75–92.
- Kaplow, Louis, 1996c. "The Optimal Supply of Public Goods and the Distortionary Cost of Taxation." *National Tax Journal* 49 (4), 513–533.
- Kaplow, Louis, 1998a. "Accuracy, Complexity, and the Income Tax." *Journal of Law, Economics, & Organization* 14, 61–83.
- Kaplow, Louis, 1998b. "Tax Policy and Gifts." *American Economic Review Papers and Proceedings* 88 (2), 283–288.
- Kaplow, Louis, 2001a. "A Framework for Assessing Estate and Gift Taxation." In Gale, William G., James R. Hines Jr., and Joel Slemrod (eds.), *Rethinking Estate and Gift Taxation*, 164–215. Brookings Institution Press, Washington, D.C.
- Kaplow, Louis, 2001b. "Horizontal Equity: New Measures, Unclear Principles." In Hassett, Kevin A., and R. Glenn Hubbard (eds.), *Inequality and Tax Policy*, 75–97. American Enterprise Institute Press, Washington, D.C.
- Kaplow, Louis, 2004. "On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy." *Journal of Economic Perspectives* 18 (4), 159–175.
- Kaplow, Louis, 2005. "Why Measure Inequality?" *Journal of Economic Inequality* 3, 65–79.
- Kaplow, Louis, 2006a. "On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal." *Journal of Public Economics* 90 (6–7), 1235–1250.
- Kaplow, Louis, 2006b. "Public Goods and the Distribution of Income." *European Economic Review* 50, 1627–1660.
- Kaplow, Louis, 2007a. "Optimal Income Transfers." *International Tax and Public Finance* 14, 295–325.

- Kaplow, Louis, 2007b. “Primary Goods, Capabilities, or . . . Well-Being?” *Philosophical Review* 116, 603–632.
- Kaplow, Louis, 2007c. “Taxation.” In Polinsky, A. Mitchell, and Steven Shavell (eds.), *Handbook of Law and Economics*, vol. 1, 647–755. Elsevier, Amsterdam.
- Kaplow, Louis, 2008a. *The Theory of Taxation and Public Economics*. Princeton University Press, Princeton.
- Kaplow, Louis, 2008b. “Optimal Policy with Heterogeneous Preferences.” *B.E. Journal of Economic Analysis and Policy: Advances* 8 (1: 40), 1–28.
- Kaplow, Louis, 2009. “Capital Levies and Transition to a Consumption Tax.” In Auerbach, Alan J., and Daniel N. Shaviro (eds.), *Institutional Foundations of Public Finance: Economic and Legal Perspectives*, 112–146. Harvard University Press, Cambridge, MA.
- Kaplow, Louis, 2010a. “Concavity of Utility, Concavity of Welfare, and Redistribution of Income.” *International Tax and Public Finance* 17, 25–42.
- Kaplow, Louis, 2010b. “Taxing Leisure Complements.” *Economic Inquiry* 48, 1065–1071.
- Kaplow, Louis, 2011. “Targeted Savings and Labor Supply.” *International Tax and Public Finance* 18 (5), 507–518.
- Kaplow, Louis, 2012. “Optimal Control of Externalities in the Presence of an Income Tax.” *International Economic Review* 53 (2), 487–509.
- Kaplow, Louis, 2015a. “Government Policy and Labor Supply with Myopic or Targeted Savings Decisions.” In Brown, Jeffrey R. (ed.), *Tax Policy and the Economy*, vol. 29, 159–193. NBER and University of Chicago Press, Chicago.
- Kaplow, Louis, 2015b. “Myopia and the Effects of Social Security and Capital Taxation on Labor Supply.” *National Tax Journal* 68, 7–32.
- Kaplow, Louis, 2017. “A Distribution-Neutral Perspective on Tax Expenditure Limitations.” In Moffitt, Robert A. (ed.), *Tax Policy and the Economy*, vol. 31, 161–188. NBER and University of Chicago Press, Chicago.
- Kaplow, Louis, 2019. “Market Power and Income Taxation.” NBER Working Paper No. 25578. National Bureau of Economic Research, Cambridge, MA.
- Kaplow, Louis, and Steven Shavell, 2001. “Any Non-welfarist Method of Policy Assessment Violates the Pareto Principle.” *Journal of Political Economy* 109, 281–286.
- Kaplow, Louis, and Steven Shavell, 2002. *Fairness versus Welfare*. Harvard University Press, Cambridge, MA.
- King, Mervyn A., 1983. “An Index of Inequality: With Applications to Horizontal Equity and Social Mobility.” *Econometrica* 51, 99–115.
- Kopczuk, Wojciech, 2013. “Incentive Effects of Inheritances and Optimal Estate Taxation.” *American Economic Review* 103 (3), 472–477.
- Lindahl, Erik, 1919. “Die Gerechtigkeit der Besteuerung: Eine Analyse der Steuerprinzipien auf der Grundlage der Grenznutzentheorie. Lund: Gleerup and H. Ohlsson.” In Musgrave, Richard A., and Alan Peacock (eds.), *Classics in the Theory of Public Finance*, 168–176 (Chapter 4, “Positive Lösung,” translated as “Just Taxation—A Positive Solution”). Macmillan: London (1958).
- Mayshar, Joram, 1991. “Taxation with Costly Administration.” *Scandinavian Journal of Economics* 93, 75–88.
- Michalopoulos, Charles, Philip K. Robins, and David Card, 2005. “When Financial Work Incentives Pay for Themselves: Evidence from a Randomized Social Experiment for Welfare Recipients.” *Journal of Public Economics* 89, 5–29.

- Mirrlees, James A., 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38(2), 175–208.
- Mirrlees, James A., 1976. "Optimal Tax Theory: A Synthesis." *Journal of Public Economics* 6, 327–358.
- Mirrlees, James A., 1994. "Optimal Taxation and Government Finance." In Quigley, John M., and Eugene Smolensky (eds.), *Modern Public Finance*, 213–231. Harvard University Press, Cambridge, MA.
- Musgrave, Richard A., 1959. *The Theory of Public Finance*. McGraw-Hill, New York.
- Myles, Gareth D., 1995. *Public Economics*. Cambridge University Press, Cambridge, UK.
- Ng, Yew-Kwang, 1984. "Quasi-Pareto Social Improvements." *American Economic Review* 74, 1033–1050.
- Ramsey, Frank P., 1927. "A Contribution to the Theory of Taxation." *Economic Journal* 37, 47–61.
- Rawls, John, 1971. *A Theory of Justice*. Harvard University Press, Cambridge, MA.
- Rawls, John, 1982. "Social Unity and Primary Goods." In Sen, Amartya, and Bernard Williams (eds.), *Utilitarianism and Beyond*, 159–185. Cambridge University Press, New York.
- Saez, Emmanuel, 2002. "The Desirability of Commodity Taxation under Non-Linear Income Taxation and Heterogeneous Tastes." *Journal of Public Economics* 83, 217–230.
- Salanié, Bernard, 2003. *The Economics of Taxation*. MIT Press, Cambridge, MA.
- Salanié, Bernard, 2011. *The Economics of Taxation*, 2nd ed. MIT Press, Cambridge, MA.
- Samuelson, Paul A., 1954. "The Pure Theory of Public Expenditure." *Review of Economics and Statistics* 36, 387–389.
- Sandmo, Agnar, 1993. "Optimal Redistribution When Tastes Differ." *Finanzarchiv* 50, 149–163.
- Sen, Amartya, 1985a. *Commodities and Capabilities*. North-Holland, Amsterdam.
- Sen, Amartya, 1985b. "Well-being, Agency and Freedom: The Dewey Lectures 1984." *Journal of Philosophy* 82, 169–221.
- Sen, Amartya, 1997. *On Economic Inequality*, enlarged ed. Clarendon Press, Oxford.
- Shavell, Steven, 1981. "A Note on Efficiency vs. Distributional Equity in Legal Rulemaking: Should Distributional Equity Matter Given Optimal Income Taxation?" *American Economic Review Papers and Proceedings* 71 (2), 414–418.
- Simon, Herbert A., 1962. "The Architecture of Complexity." *Proceedings of the American Philosophical Society* 106 (6), 467–482.
- Slemrod, Joel, and Christian Gillitzer, 2014. *Tax Systems*. MIT Press, Cambridge, MA.
- Slemrod, Joel, and Wojciech Kopczuk, 2002. "The Optimal Elasticity of Taxable Income." *Journal of Public Economics* 84, 91–112.
- Slemrod, Joel, and Shlomo Yitzhaki, 1987. "The Optimal Size of a Tax Collection Agency." *Scandinavian Journal of Economics* 89, 183–192.
- Slemrod, Joel, and Shlomo Yitzhaki, 2001. "Integrating Expenditure and Tax Decisions: The Marginal Costs of Funds and the Marginal Benefit of Projects." *National Tax Journal* 54, 189–201.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick, 2019. "Capitalists in the Twenty-First Century." NBER Working Paper No. 25442. National Bureau of Economic Research, Cambridge, MA.
- Stern, Nicholas H., 1982. "Optimum Taxation with Errors in Administration." *Journal of Public Economics* 17, 181–211.

- Stiglitz, Joseph E., 1987. "Pareto Efficient and Optimal Taxation and the New New Welfare Economics." In Auerbach, Alan J., and Martin Feldstein (eds.), *Handbook of Public Economics*, vol. 2, 991–1042. North-Holland, Amsterdam.
- Stiglitz, Joseph E., and Partha Dasgupta, 1971. "Differential Taxation, Public Goods, and Economic Efficiency." *Review of Economic Studies* 38, 151–174.
- Vickrey, William, 1945. "Measuring Marginal Utility by Reactions to Risk." *Econometrica* 13, 319–333.
- Westen, Peter, 1990. *Speaking of Equality: An Analysis of the Rhetorical Force of "Equality" in Moral and Legal Discourse*. Princeton University Press, Princeton.