Measuring Income and Income Inequality[1]

Conor Clarke[2]
Washington University in St. Louis
School of Law

Wojciech Kopczuk[3]
Columbia University
Department of Economics

April 3, 2025

[1]For helpful comments and discussions we thank Alan Auerbach, Jerry Auten, John R. Brooks, Jim Hines, Russell Osgood, Jonathan Parker, Nina Pavcnik, Robert Pollak, Emmanuel Saez, Peter Siminski, David Splinter, Timothy Taylor, Heidi Williams, and Eric Zwick.

[2]conor.clarke@wustl.edu

[3]wojciech.kopczuk@columbia.edu

Abstract

Income inequality is important, but attempts to measure it arrive at strikingly different conclusions. Why? We use recent disputes over measuring United States income inequality to return to first principles about both the income concept and inequality measurement. We emphasize two broad points. First, no measure of the income distribution is truly comprehensive, or could attempt to be comprehensive without making controversial choices. We document the practical and conceptual problems that the standard ideal—comprehensive Haig-Simons income—raises. Second, much of the controversy in this area turns on the many tradeoffs between starting with individual tax data versus attempting to match more expansive income concepts. Individual tax data reflects only a shrinking subset of a more comprehensive income concept—but it is individual data. Broader alternatives, on the other hand, are harder to allocate to individuals. We document some of the most important and contestable assumptions that allocating national income requires.

Introduction

The distribution of income is an alluring topic for researchers, policymakers, journalists, and the public. The economics literature on measuring income inequality extends back at least to Kuznets (1953). But the contrast between two recent and prominent contributions by Piketty, Saez and Zucman (2018) and Auten and Splinter (2024*a*) raises basic methodological issues that go to the heart of this literature. The two papers pursue the same goal—the measurement of United States top income inequality—but arrive at strikingly different conclusions: Piketty, Saez, and Zucman show a large increase in recent inequality, while Auten and Splinter show a pattern that is much more flat. In this paper, we use the controversy to return to first principles of measuring income and income inequality. We begin with a basic discussion of the income concept—starting with so-called comprehensive Haig-Simons income—and what it can and cannot capture. We then critically evaluate the income metrics in recent literature, the decisions that go into constructing those metrics, and the problems with underlying data sources.

Our goal is not to adjudicate the recent debate. Instead, we emphasize two themes about income and inequality measurement. First, no current measures of the income distribution use (or could use) the comprehensive Haig-Simons measure of income—the leading income concept that attempts to capture all changes in savings and consumption. Indeed, no measure could attempt to be truly comprehensive without making controversial methodological choices, and it is far from obvious that a comprehensive income measure would capture the most important distributional details that society should rightly value. Haig-Simons income is a common and influential ideal—and can be defended as a kind of necessary evil—but its virtues should not be overstated and its vices should not be forgotten.

Second, the efforts to quantify the elusive income concept raise tradeoffs. Papers that stick closely to what is observable in administrative tax data to measure income have many virtues, such as plentiful individual-level data, and data that speaks to the very top of the income distribution.[1] But the income that appears on annual tax returns reflects only a shrinking subset of a more comprehensive Haig-Simons concept. On the other hand, papers that attempt to use a more comprehensive (if still incomplete) income definition derived from the national accounts must extend beyond what can be observed on tax returns—and, often, any administrative data sources. Allocating these broader income concepts without individual administrative data requires contestable assumptions—which, in turn, helps to explain the recent controversy over the different findings of Piketty, Saez and Zucman (2018) and Auten and Splinter (2024*a*).

In short, the ideal income concept is flawed, and attempts to approach this flawed ideal must come at a steep methodological cost. We conclude by sketching a few additional implications for the contemporary debate over income inequality and future literature.

What measure of income do we want, and why?

There is a bewildering profusion of income concepts: national, personal, labor, capital, factor, fiscal, cash, market, expanded, taxable, gross, net, pre-tax, post-tax, post-transfer, disposable—and many more. The sheer profusion of income concepts can give the impression that "income" results are driven as much by definitional choices as real economic changes. It is also not clear that *any* of these concepts align well with what non-academics intuitively consider income. For example, Dahl and Ransom (2002) surveyed Mormons about their understanding of income for religious tithing purposes and concluded that it did "not coincide with current tax laws or economists' views of comprehensive income."

[1] Because we focus on attempts to measure top income shares, tax data rather than survey data is our default for discussion.

Figure 1 provides a sense of the relationship between various income measures and concepts. The bottom two categories show labor and capital income that appears on individual tax returns as a share of national income, as defined by US national accounts. Such income is a declining share of national income over time. The next level shows untaxed labor income—things like employer-provided health insurance and other fringe benefits—while the following level shows capital income that is not subject to individual taxation (and, hence, not directly linked to particular individuals). This category of capital income includes a variety of gains that the US tax system does not target, such as the rental value of owner-occupied housing, the interest from certain tax-advantaged retirement accounts. It also includes corporate retained earnings and corporate tax liability itself. These forms of income are included in national income as measured by our national accounting system—every component of which is included under the horizontal line at 100 percent.

But national income isn't everything. The categories above the national income line offer a sense of forms of individual economic gain that US national accounts do not include. The first category shows realized capital gains, which are a familiar part of US tax law—you pay tax on an appreciated asset when you sell it—but are not part of "national income" as officially defined.[2] The next category shows another familiar form of economic benefits: payments from Social Security, disability insurance, and unemployment insurance. The category above that shows all other government transfers, including those for health care spending. The sum of those categories, plus national income, can be viewed as a 'narrow' version of the comprehensive Haig-Simons definition. The next category above shows the value of unpaid household production—which is based on surveys of time-use and included in official 'satellite' national accounts. Adding household production to the categories below it creates a total that can be viewed as closer to a 'broad' Haig-Simons definition. But adding other categories is possible, too. The top category includes a rough and illustrative value of leisure—based on time-use survey averages and multiplied by an indexed version of the federal minimum wage.[3] One could also imagine adding yet more hard-to-value components of welfare on top of that—sleep, happiness, health, and the like. This is intended to be provocative more than precise. But, hopefully, it illustrates both the fuzzy contours of the income concept and the large differences between aggregate definitions.

The Haig-Simons baseline

Where does our modern, baseline concept of "income" come from? The so-called "Haig-Simons" comprehensive definition of income dates to the writings of economists Robert Haig (1921) and Henry Simons (1938). Haig defined income as "the money value of the net accretion to one's economic power between two points of time," and Simons defined it as "the algebraic sum of (1) the market value of rights exercised in consumption and (2) the change in the value of the store of property rights between the beginning and end of the period in question." Broadly and somewhat informally, the Haig-Simons concept—sometimes shorthanded as "economic" income—defines the income of an individual as the market value of all changes in the ability of that individual to save or consume over a given period.

Haig and Simons were writing in the wake of the ratification of the Sixteenth Amendment (1913), which empowered the US Congress "to lay and collect taxes on incomes, from whatever source derived."[4] The

[2]The precision of this figure should be not be overstated: some components of asset appreciation—such as corporate retained earnings—are included in the national accounts, so that adding realized capital gains may double-count some gains. In addition, the measure on the graph does not directly include tax-exempt housing capital gains and it does not account for inflation.

[3]The short appendix describes the construction of this figure in more detail.

[4]More specifically, the Sixteenth Amendment allowed Congress to tax incomes without apportioning such a tax on the basis of state population—i.e., such that a state with 5% of the national population would bear 5% of the total tax burden. In 1895, in the landmark case of *Pollock v. Farmers' Loan and Trust Company*, the United States Supreme Court held that an income tax needed to be apportioned by population. Because an apportioned income tax cannot guarantee the equal treatment of individuals with the same income in different states—or progressive rates—*Pollock* was a fundamental obstacle to US income taxation that led to the eventual ratification of the Sixteenth Amendment.

Figure 1: Components of different income concepts (national aggregates)

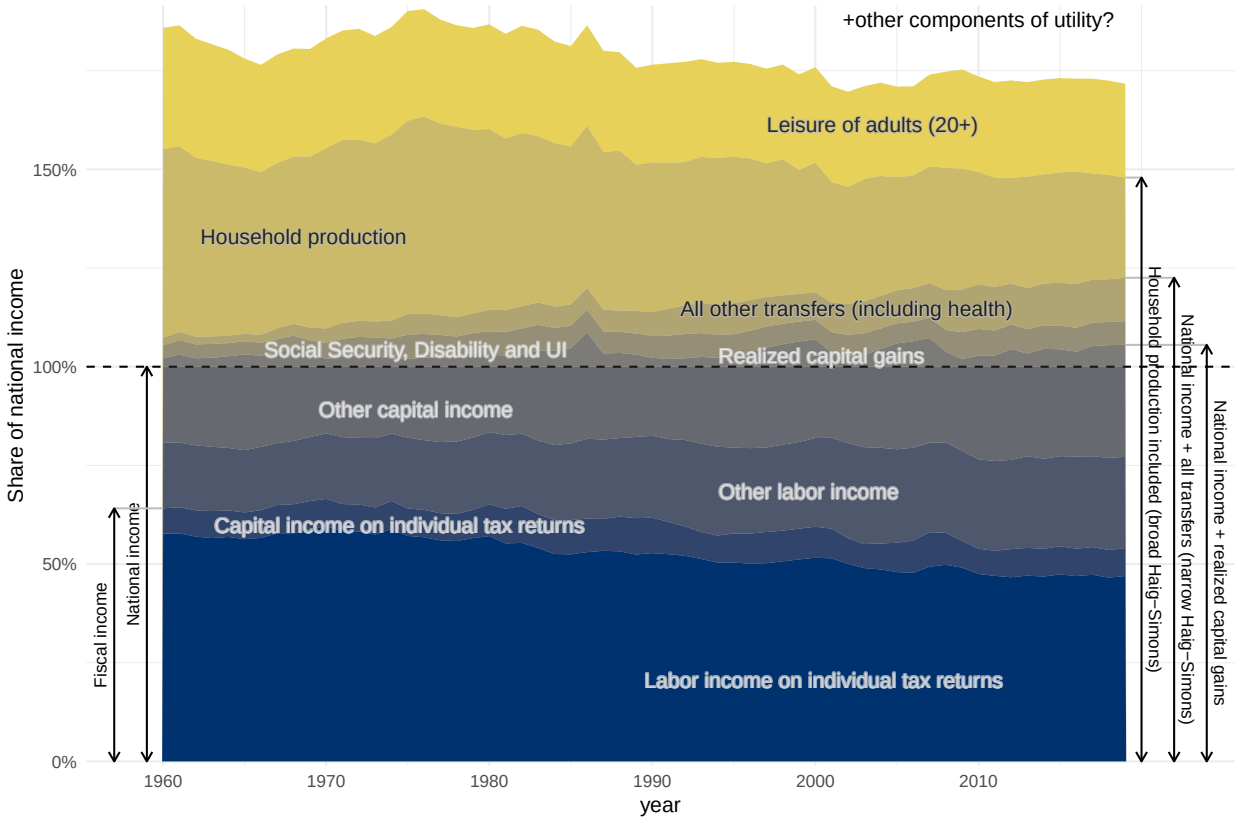

backdrop for their pioneering work was thus an early administrative income tax system for which the basic concepts—like the appropriate scope of "taxes on incomes"—were still being worked out. The concept of "income" existed before the invention of an income tax, of course. But, prior to its use as a tax base, "personal income" was (to our knowledge) not a central concept in scholarly work. In Adam Smith's *The Wealth of Nations* (1776), for example, the concept of "income" makes only a vanishingly small number of appearances. Smith's discussion of taxation acknowledges that individuals should contribute "in proportion to the revenue which they respectively enjoy under the protection of the state." But he considered such taxes impossible. For Smith, "[t]he impossibility of taxing the people, in proportion to their revenue" is what led to "the invention of taxes upon consumable commodities"—which were more easily measurable and, Smith argued, will be "nearly in proportion to [one's] revenue." Early and temporary US experimentation with an income tax during the Civil War encountered a similar issue: "The government had no scientific way to measure personal income" (Brownlee, 2016).

Haig and Simons were thus concerned with fashioning a new concept of income that would be relevant and administrable for both lawyers and policymakers. They had only a limited practical and philosophical inheritance upon which to draw—though their discussions did mirror some earlier German-language discussions of the income concept that flowed from European experimentation with income taxation (for example, von Kleinwächter, 1896).

Haig (1921), for example—leaning on a tradition of early utilitarian thinkers—emphasized that an abstract and ideal economic analysis might recognize income as "a flow of satisfactions, of intangible psychological experiences." But he quickly conceded that practical analysis required "something more definite and

more homogeneous—less diaphanous and elusive than these psychic satisfactions." For this reason, Haig argued that a usable concept of income must look to gains that could be readily converted into market prices. "The basis of comparison," he argued, "is, of course, that of the common, universally-acceptable unit of value—money." Simons (1938), likewise, argued that a workable definition of personal income "has to do not with sensations, services, or goods but rather with rights which command prices (or to which prices may be imputed)."

Over time, the Haig-Simons concept has obtained implicit dominance in tax law and tax policy. Many introductory legal and economics textbooks introduce students to the concept of Haig-Simons income as a baseline for thinking about the difference between "economic" and "legal" income, and for assessing individual provisions of tax law (for example, Graetz and Alstott, 2022 and Gruber, 2022, p. 538). National governments rely on the Haig-Simons concept too. Each year, for example, the federal government estimates the size and costs of various "tax expenditures" (the name for quasi-spending programs that run through the tax code) and clarifies that such estimates use the baseline of a "comprehensive income tax, which defines income as the sum of consumption and the change in net wealth in a given period of time"—that is, Haig-Simons income. The Haig-Simons framework also underlies the guidelines of the United Nations Statistics Division, which provides standards for how statistical agencies around the world define income (see Larrimore et al., 2021).

To be sure, there are excellent reasons for an income tax system to express some aspirational interest in capturing *all* changes in an individual's ability to save or consume. A more naïve concept of income—say, 'cash received'—would quickly fall victim to an even more sophisticated world of fringe benefits: tax-free meals, housing, insurance, and so forth. One defense of the broad Haig-Simons concept is therefore that it expresses the ambition to capture all forms of real economic gain—such as benefits paid in-kind—rather than mere salaries paid in cash or cash gains realized upon the sale of an asset.

But there are also serious drawbacks to the Haig-Simons concept, many of which have been known since the beginning and have been emphasized in recent research (as in Brooks, 2018). In the rest of this section, we discuss five interrelated problems that have particular relevance for the measurement of income and income inequality. First, Haig-Simons income is perhaps best conceptualized as an aspirational concept for administrative tax law, not a concept developed with justice or fairness in mind—and its proper role in debates over justice is debatable. Second, there are serious conceptual difficulties about what the comprehensive Haig-Simons concept can and should include, which raises stubborn questions about its implementability. Third, even where there is conceptual consensus about what to include, many forms of savings and consumption are inherently difficult to value. Fourth, the Haig-Simons concept does not tell us what the appropriate unit of observation is (individuals? households?) for measuring income and inequality. And, fifth, the Haig-Simons concept does not tell us how to attribute government spending to personal income (especially given that government spending was already income for someone else before it was taxed and spent).

Haig-Simons and fairness. Haig and Simons were not writing about personal income in the context of fairness and equality; they were writing about income in the context of a nascent tax system attempting to find its footing. This meant the income concept that Haig and Simons developed already started with concessions to reality—such as Haig's comment that a more ideal utilitarian concept would concern itself with the "flow of satisfactions" and "intangible psychological experiences," rather than gains that could be readily translated into market prices. Other measures beyond "utility"—like some measure of welfare, flourishing, happiness, capability, opportunity, ability to pay, and so forth—would have a similarly imperfect overlap with the Haig-Simons concept.

One such concession to reality is that the measurement of Haig-Simons income limits itself to a single administrative time period (in the case of the US tax system, normally one year). Again, there is a good reason for this: Much of collective human life is conducted on an annual basis. But justice and fairness should care about more than a year. Income changes over the life cycle. (For some discussion of inequality in this context, see Auerbach, Kotlikoff and Koehler, 2023.) It is not necessarily a social fairness problem if the 20 year-old college student has a low Haig-Simons income in a given year. Income can also be lumpy. It is not necessarily a distributional problem if that college student grows up and, in one year, earns a one-time bonus or inherits a house.[5]

Of course, the concept of Haig-Simons income does have *some* relationship to justice and fairness. Ideally, the tax system cares about measuring income because it cares about measuring ability to pay, and this particular concept of income is constructed with a defensible notion of ability to pay in mind. And, as anticipated centuries ago by Adam Smith, the tax system should care about measuring ability to pay because it reflects a fundamentally fair and reasonable way to allocate the cost of government. But "income"—as defined by Haig and Simons—is at best a rough indicator of the other measures of justice listed above, and it is an open question whether first principles of justice should care about one's time-restricted saving and consumption divorced from its impact on those other measures. Accordingly, the standard concept of "justice" used in economics (and based on a utilitarian welfare function) does not intersect directly with Haig-Simons income.

Perhaps the best defense of using Haig-Simons as a baseline for thinking about distributive issues goes like this. First, it makes sense to use income data to study distributional issues because we *have* income data—both from the work of the administrative tax system and from our system of national accounts. Second, while income is an imperfect way of measuring fairness, we have some reason to believe that it is connected to the measures we care about, because it relates to the ability to pay (as discussed above and by, for example, Musgrave, 1967; Thuronyi, 1990). Finally, every measure of social welfare or distribution has its faults (see, for example, Deaton, 2020 and Glogower, 2023). Still, the particular faults of income as a rubric for thinking about justice and fairness—it was only partially intended to capture those things, and it only captures them imperfectly—should not be forgotten.

Haig-Simons and implementability. How feasible is it to implement a broad and comprehensive concept of income? A fundamental problem is that that the concepts of "consumption" and "savings" (or changes in wealth) are not obvious or self-defining. Indeed, even in cases where there *is* rough consensus about the concepts, there can be uncertainty in how the concepts can be measured in common units of exchange that a tax system (and the Haig-Simons concept) require. The issues here are vast—indeed, defining and policing the boundaries of those concepts is a primary object of tax law and policy—and we provide here only a selective overview of some of the more fundamental issues.

One basic question is how to think about the frontier between personal consumption and expenses that represent costs of producing income. It is crucial to an income tax base—as opposed to a tax on the volume of activity or receipts—that it nets out the costs of producing income. If two business owners each bring in $100,000 of gross receipts in a given year, but one has $90,000 in expenses and the other has $10,000 in expenses, an income tax should not treat those owners as equivalent. If it did, the tax would be one on gross receipts, not income. Indeed, many taxes (including many American taxes) were, and sometimes still are, imposed on gross receipts rather than on net income—presumably because gross receipts are easier to measure, even if they do less to capture ability to pay.

[5]Similar issues arise concerning whether the nation-state is the best level of generality in which to study income. Again, much collective life is conducted through nations. But it might be that local or global inequality is what matters more. For some discussion of this issue, see Hayashi (2023).

But figuring out what is a cost of business—as opposed to an item of personal consumption—turns out to be surprisingly tricky. (For some evidence of businesses blurring the line between costs and personal consumption, see Alstadsæter, Kopczuk and Telle, 2014 and Leite, 2024.) As Simons observed almost 100 years ago, something can obviously be either, depending on the context. A professional artist who buys paint can treat it as a business expense, while a hobbyist who buys the same paint treats it as consumption. For Simons, this dual quality suggested an unsettling conclusion: "there is something quite arbitrary about the distinction between consumption and accumulation."

This issue is not a minor one. Just as the same activity may be viewed reasonably as representing consumption or production—depending on the context—so too may the same good or activity be shot through with both consumption and production. The activities most fundamental to the human experience—food, sleep, shelter—are both an obviously major component of consumption and a necessary condition of earning income. Indeed, the US tax system is somewhat schizophrenic in whether it treats such expenses as consumption or as a cost of production. For example, the deductibility of certain personal medical expenses may be viewed as either a departure from the normal income tax baseline (i.e., it is a form of consumption we decline to tax) or as a reasonable part of the normal income tax baseline, because it does not tax a cost of labor production (i.e., we need to be alive to produce income).

Or consider the issue from the other direction: Many jobs are enjoyable, and in that sense can be (partially) conceptualized as substantial consumption opportunities. A young economics or law student may decide to pursue the academic life—under the impression that it will be fun—over a career in finance or at a law firm, even if the latter may yield a far higher salary. Again, the issues here are conceptual, practical, and moral. How should we think about taxation of those who choose the "fun job," although their talents could have earned higher market returns in different professions? Should the tax system create incentives for individuals to pursue a "calling" that pays a lower income, or should it attempt to reach talents and abilities directly? (For some discussion of these issues in both economics and in law, useful starting points include Lockwood, Nathanson and Weyl, 2017 and Zelenak, 2006.)

Imputations and the problem of valuation. Both Haig and Simons imagined a concept of income that would add up "the money value" (in Haig's terms) or "the market value" (in Simons's) of changes in consumption and saving. But, even when we can agree on the right conceptual bucket for an activity (see above), many large components of consumption and savings do not have readily ascertainable market values. Again, this raises a mix of issues that are both practical and philosophical. For example, the change in value of corporate equities constitutes a large fraction of annual Haig-Simons income, but there is (by definition) no public market for closely held corporations. Or consider: Leisure can be thought of as a large component of consumption—and indeed working hours and vacation time are often in employment contracts—but there is no agreed-upon method for valuing moments of idle reverie in which one does nothing. (But see Figure 1 for one rough attempt.)

The so-called "problem of valuation" is closely related to a recurring issue in the income and income inequality literature over what "imputations" of income can and should be done (where "imputation" in this context refers to an assignment of value that must be done indirectly). Most observers would agree, for example, that in thinking about the distribution of Haig-Simons income we should impute rental income to owners who occupy their own houses, on the theory that occupying a home increases one's ability to save or consume—just as having an employer cover your rent, or owning a home and renting it out as business, can increase one's ability to save or consume. But practice diverges: US tax law does not impute (and thereby attempt to tax) rental income to owner-occupiers, while US national accounts do attempt to include this as part of "national income."

What imputations of value are possible and where should one stop? Perhaps, for example, we could and should impute the income of businesses as it is earned to the owners of those businesses—as our tax system does in the case of some (but not all) businesses, and as we discuss in Clarke and Kopczuk (2017) and further below. This discrepancy is broadly related to another major deviation between the Haig-Simons concept and the tax system in the United States: the timing of when capital gains income is recognized by the tax system, or deemed "realized" (which generally happens when an asset is sold, rather than when it appreciates in value). As a legal matter, the status of realization remains contested.[6] As a policy matter, one traditional argument in favor of a realization requirement is that some assets are too difficult to value on an accrual basis.

And then there is household production. Many tasks could in principle be outsourced to the labor market—cooking, cleaning, driving, childcare, and so on. Just as one might think about imputing rental income for one's own house, we might think about imputing labor income for one's own labor. This is not just a toy philosophical issue. Household labor consumes a lot of time—hundreds of billions of hours each year in the United States (Bridgman, Craig and Kanal, 2022)—and correlates with other elements of Haig-Simons income. Household labor can also change dramatically over time—as we have seen in both recent decades (as the composition of the workforce has changed) and in recent years (as the pandemic temporarily changed labor markets). There is no consensus on how to account for household production—and our tax system and standard national accounts do not attempt to account for it. However, satellite estimates of gross domestic product suggest that including household production would add more than 25 percent to standard measures of GDP in 2020 (Bridgman, Craig and Kanal, 2022).[7] The income inequality literature is only starting to address this issue, and Figure 1 provides some sense of its potential scale. Frazis and Stewart (2011), and the more recent Gautham and Folbre (2024), suggest that incorporating household production into inequality measures can have a large effect on measured levels and trends.

The proper unit of observation. In their initial discussions, Haig and Simons were concerned with "personal" income of individuals. But it is not obvious what the best unit of analysis is, and tax systems vary in how they approach this issue. For example, for the first few decades, the American income tax system only looked to individual income and did not, as now, allow varieties of joint (household) filing to occur. (Indeed, at certain junctures this created an incentive for a household to attempt to reduce its taxes by formally splitting a husband's income with a nonworking wife, on the theory that they would each pay a separate individual tax at a lower marginal rate.)

For normative and distributional analysis, there is certainly a case for looking to individuals. Much of life concerns individuals. We *are* individuals. It seems obvious and reasonable to consider the gains and

[6] In the high-profile 2024 Supreme Court case of *Moore v. United States*, 602 US 572 (2024), four of the nine justices indicated their view that "realization" was a constitutional requirement—in other words, that income must be "realized" in some way before the tax system can reach it under the authority of the Sixteenth Amendment. For present purposes, the opinion illustrates that there may be important and ongoing differences in how economists and the legal system think about "income."

[7] Whether household work should be included in national production and national income has been an issue since the very beginning. Here is the rather conclusory and unsatisfying discussion of the issue in The Department of Commerce (1934), which was the government's original report on the national income concept developed under the guidance of Simon Kuznets: "The volume of services rendered by housewives and other members of the household toward the satisfaction of wants must be imposing indeed, when totaled for the 30 million families comprising the population of this country; and the item is thus large enough to affect materially any estimate of national income. But the organization of these services render them an integral part of family life at large, rather than of the specifically business life of the nation. Such services are, therefore, quite removed from those which gainfully occupied groups undertake to perform in return for wages, salaries, or profits. It was considered best to omit this large group of services from national income, especially since no reliable basis is available for estimating their value. This omission, unavoidable though it is, lowers the value of national income measurements as indexes of the nation's productivity in conditions of recent years when the contraction of the market economy was accompanied by an expansion of activity within the family."

capacities of individuals. But other facts of life seem obvious and reasonable, too. Individuals organize themselves into households, extended families and even communities to pool resources and pursue shared plans. It therefore seems reasonable to think about both the "income" of individuals and of households—and indeed potentially of larger units, such as extended families, communities, and so forth. On the other hand, households form, change, and dissolve over a lifetime, which may push in favor of using individuals as the right unit of measurement. The naked concept of "income" generates no consensus on the *right* way to think about these issues—although any administrative tax system and any discussion of inequality must necessarily take a stance on them.

Government spending and income. From the Haig-Simons perspective, some government spending may seem to translate easily into income. Someone who receives a Social Security check, for example, has money to save or consume. (And Social Security benefits generally contribute to taxable income above a certain level.) But even this seemingly simple case disguises great complexity. After all, Social Security benefits correspond to taxes that somebody else pays. Including both taxes paid and benefits received in a measure of income would involve double-counting. We might be tempted to simply deduct Social Security taxes from income. But that would diverge from one administrative purpose of the underlying income concept—namely, measuring one's ability to pay for the costs of government.

How other government spending translates into ability-to-pay poses even harder challenges. Paved roads paid for by a government also increase one's ability to save or consume, but it is harder to attribute the resulting income to individuals or households. Similarly, national defense supports the presence of relatively secure property rights and the absence of pillaging invaders improves savings and consumption. But it is difficult to say how such forms of spending should be allocated to households and individuals across the income spectrum, if at all. The size of such benefits—and how fast the benefits rise in proportion to other forms of income—have been debated by economists, philosophers, and lawyers for decades (for example, Blum and Kalven, 1952, Murphy and Nagel, 2004, Piketty, Saez and Zucman, 2018).

The question begins to approximate: What is government worth to each individual? But most tax systems are concerned with estimating one's ability to pay *for* the costs of government, not with estimating how the costs of government improve one's ability to save or consume. Moreover, the issues of how government taxing and spending are interrelated with income are not easily resolved by appealing to *both* a pre- and a post-tax perspective, as is the case in recent work allocating national income by both Piketty, Saez and Zucman (2018) and Auten and Splinter (2024*a*). This recent work roughly allocates *current* government spending—already a formidable and contestable imputation task, as discussed more below. But it does not attempt to impute income—which would be akin to a kind of imputed rental income—from the myriad forms of capital (both physical and institutional) that the government owns, represents, or controls.

The five issues described above are neither trivial nor new—much less resolved. Many of them were reflected in the early German-language discussions of the income concept, were repeated again by Haig and Simons, and were raised again in the debates in the 1960s and 70s over the concept of "tax expenditures" and the development of a national tax expenditure budget (for example, Bittker, 1967). Similar concerns have been raised in recent legal scholarship (as in Brooks, 2018). But they may have been forgotten in some of the recent, public debates over the distribution of income—which have treated income as both an obvious and obviously meaningful unit of measurement and dimension of comparison.

Recent controversies over US income inequality

In recent studies of US top income inequality, two broad families of "income" measure are in wide use: income measures that start with tax data and income measures that start with national accounts data. But the measures of income that reach beyond what can be directly observed in administrative tax data have to rely on extensive imputations based on other data sources. These other data sources can include aggregates, surveys, smaller-sample estimates, and other administrative information that cannot be directly linked to observed income. Sometimes, evidence may only be available for a select number of years, so researchers must project to other periods. Sometimes, there is not much information at all.

These data limitations require many assumptions that leave room for subjective judgment calls. It is therefore unsurprising that different researchers pursuing nominally the same exercise will reach different answers. At the same time, the different answers put consumers of such research in the unenviable position of trying to evaluate the underlying judgment calls—often with no consensus methodology on which to rely. Sometimes, the best a casual reader can do is rely on the authors' own summaries of the differences.

Figure 2 provides an overview of how different choices and definitions produce different distributional results. It displays the top 1 percent income shares over time drawn from some prominent sources in both government and the academic literature. Most of these series are pre-tax, although a number of the measures include some government transfers. The series from Larrimore et al. (2021) nets out taxes but are included because they are unique for their detailed accounting for capital gains. Given the ambiguity and flexibility of the income concept, the variations in the levels and trends are, unsurprisingly, large. In recent years, for example, the measured top 1 percent share of income varies between almost 25 percent and less than 9 percent, depending on which measure of income is selected. The figure may be a little messy and disorienting to look at. But that's the point: the lack of consensus produces disorienting results.

There are several main sources here. Three estimates are from the Congressional Budget Office, based on household data: one for market income, one adding just social insurance transfers (like Social Security), and one adding all transfers. Several of the results are drawn from studies by Piketty and Saez, and sometimes also Gabriel Zucman. For example, Piketty and Saez (2003) uses tax units as data in one line, and then adds realized capital gains in a second line. Piketty, Saez and Zucman (2018) uses individuals as the unit of observation and use national income as the basis for analysis. Another set of results is drawn from Larrimore et al. (2021), who start with a series that accounts for cash transfers and in-kind private and public benefits, then nets out taxes and adds both realized and accrued capital gains. Finally, we show two measures from Auten and Splinter (2024*a*): their measure of pre-tax national income and a second that includes government transfers.

We do not think that any of these approaches is "right" or "wrong." And we do not think that there is a general-purpose series that one should pick above the rest. The estimates correspond to different concepts of income and different definitions of the population, and they vary along other dimensions too. For example, some authors vary in whether they include payroll taxes or Social Security benefits in the pre-tax concept of income, or in how they account for tax evasion, or in how they distribute business incomes that are not easily attributable to individuals or households using administrative tax data. Given the varying concepts of income, the varying notions of the income-earning unit, and the varying data limitations, there are important methodological choices to make.

Nevertheless, some general patterns are discernible. How one accounts for capital gains matters. They are a large and volatile component of income—even more so when one considers them as they accrue rather than as they are realized. Whether and how one includes social insurance and transfers matters quite a bit too, as one might expect. Finally, there appears to be less disagreement before the mid-1980s than

Figure 2: Different estimates of the top 1% share in the literature

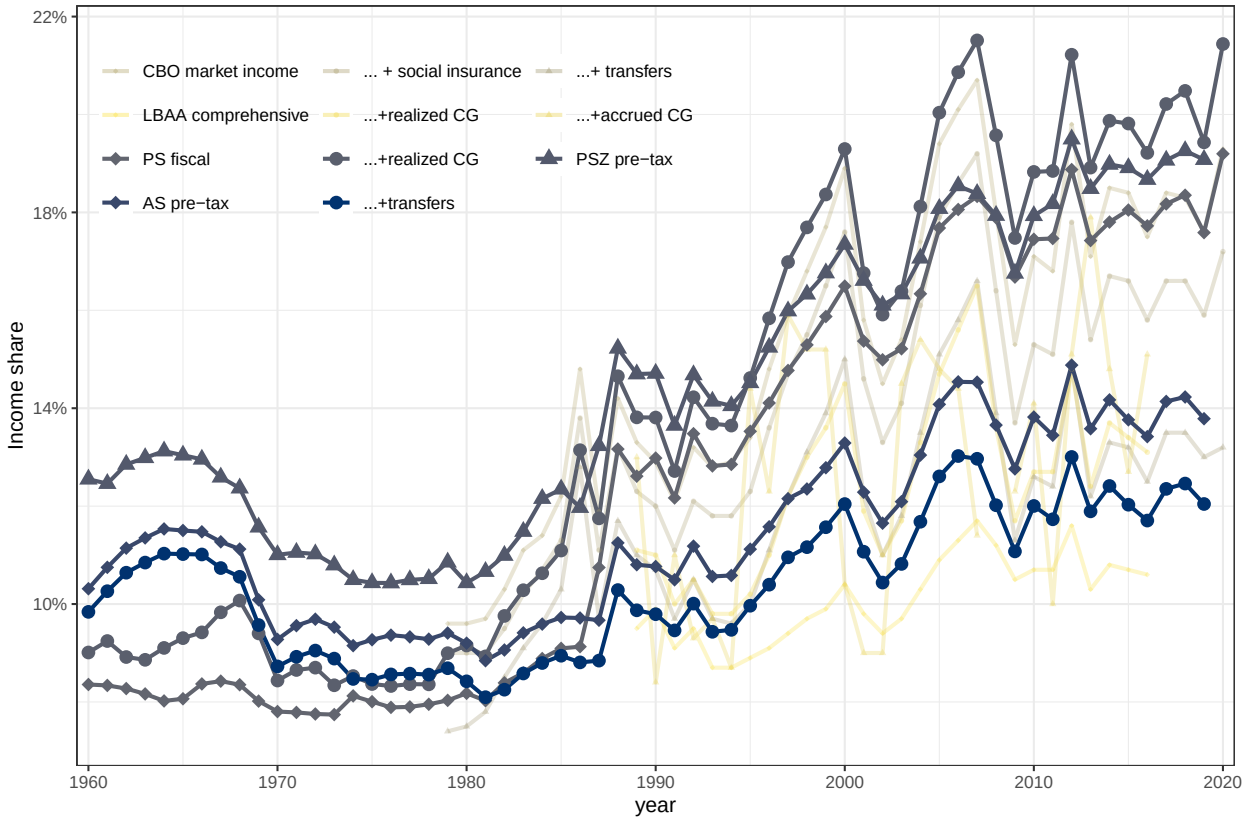

Estimates from Larrimore et al. (2021) (LBAA), Congressional Budget Office (2023) (CBO), Piketty and Saez (2003) (PS, updated series), Piketty, Saez and Zucman (2018) (PSZ, updated series) and Auten and Splinter (2024*a*) (AS)

after—or, in other words, the trends vary. One striking aspect of these trends is that the pre-1980s series are smoother and less cyclical, which suggests to us that the earlier series may not be accounting fully for business incomes—something that we will come back to and which we explore in other work.

The tax-law income baseline

The influential efforts by Piketty and Saez (2003) to measure US income inequality relied on data generated by the individual income tax, a concept they described as "fiscal income." This approach had, and continues to have, several strong practical justifications. First, prior literature that relied on survey data had incomplete coverage at the very top of the income distribution. But, at least in principle, the modern tax system sees virtually all taxpayers. Second, the tax data enables (with some assumptions) both the observation of individual units and the construction of national aggregates. Third, US income tax data goes back a long time—in some cases, close to the ratification of the Sixteenth Amendment in 1913. In principle, therefore, tax data allows for the construction of long-run series on "income" that is rich and granular along many dimensions.[8]

[8]Not everyone has to file a tax return. But, nowadays, tax authorities also have some information about the income of non-filers from "information tax returns," which are required from some to provide information about taxes that are owed by others (say, a firm that hires an independent contractor). This information allows for estimating the aggregate income of non-filers and—when coupled with the assumption that non-filers are at the bottom of the income distribution—still allows for estimating top income shares. But this issue becomes thornier the further back in time one goes, both because the extent of information reporting has changed over time, and because the income tax system was far from universal until after the Second World War.

But relying on tax data results in large deviations from the baseline Haig-Simons income concept. As Atkinson, Piketty and Saez (2011) acknowledged, all the income estimates in this literature "follow the tax law, rather than a 'preferred' definition of income, such as the Haig-Simons comprehensive definition." As a result, researchers who use tax data to measure "income" inequality are not actually measuring what the "income" terminology often suggests. Moreover, we should expect the size and nature of the difference between a 'preferred' income measure and one derived from tax law to evolve over time—as both tax law and the real economy evolve.

One glimpse of these points is visible back in in Figure 1, which illustrates that taxable labor and capital income is an incomplete share of national income and other income concepts, and that this share changes over time. As Piketty, Saez and Zucman (2018) note, "the fraction of national income that is reported in individual income tax data has declined from 70% in the late 1970s to about 60% today." Why is that? Here, we recount a few of the particularly large divergences between national income and tax-law income.

First, the United States taxes many forms of economic income only when they are "realized," rather than when they accrue. For example, a family may own real property, equities, and other important assets that increase in value over the course of a year—thereby increasing that family's ability to save or consume. But the US tax system does not see taxable income until (for example) an asset is exchanged or sold—a so-called "realization event." Realization is a tax-law term of art that has no conceptual analog in finance or economics. It is not just equivalent to the sale or disposition of an asset, or receipt of cash. And it is not always intuitive. (For example, there are cases where a person can exchange one asset for another of a similar type without triggering income tax liability.) The realization requirement represents an enormous deviation from an ideal economic concept of income and has been famously characterized as the "Achilles heel" of the US income tax (Andrews, 1983).

Second, the United States separately taxes individual income (via an individual income tax) and certain corporate income (via a corporate income tax). Income earned by corporations that pay the corporate income tax is not taxed to individuals until it is distributed to them, or until individuals sell their equities and 'realize' income. On the other hand, the income of many "pass-through" entities is immediately attributed to (and thus individually taxable to) their owners. The separately taxed corporate sector—and the shifting legal line between the individual and corporate sectors for tax purposes—raises many complex issues for distributional work that relies on tax data (for an overview, see Clarke and Kopczuk, 2017).

Third, and related, personal income on individual tax returns does not account for business and employer-side taxes (like the corporate income tax and an employer's share of payroll taxes). This means that the income appearing on individual tax returns cannot be used to construct an airtight measure of "pre-tax" income. Individual pre-tax income is already net of business-side taxes.

Fourth, through the "tax expenditure" policies mentioned earlier, the United States excludes from taxable income many large items that uncontroversially increase the ability of individuals to save or consume (for an overview, see Joint Committee on Taxation, 2022). For example, the value of employer-provided healthcare plans is generally excluded from taxation. And the basis of property acquired at death—say, inherited stock—is stepped up to its "fair-market value," meaning that untaxed appreciation will escape taxation altogether. These tax expenditure policies are numerous and large.

Fifth, income that is subject to tax evasion and tax avoidance is not fully accounted for in a tax-based measure of income. (The income of non-filers is also not on tax returns, although Piketty and Saez (2003) do estimate and add this as a component of income in the aggregate.)

Despite these issues, the use of tax data to study distributional questions has a strong rationale. After all, tax law reflects a particular society's answer (at a particular time) to the difficult and inevitable questions about ability-to-pay that have been known since the time of Haig and Simons. Moreover, especially at the

top of the income distribution, this is the best individual data on income that we have.

But this justification must come with caveats. Any tax law-derived income measure will be missing important forms of economic gain—which will, in turn, vary systemically across the income distribution. For example, we should expect taxpayers to differ systematically across income levels in the degree to which they benefit from the realization doctrine, own corporations, and receive large employer-provided health-care plans. Legal definitions also change over time. Large legal changes—like the major statutory tax reform that happened in 1986—cast doubt on the ability of law-derived measures to produce a consistent measure of income. And, finally, tax law-derived concepts of income will vary across countries—casting doubt on the power of international comparisons, as well as historical ones.

National accounts and opposing views

The scale and nature of the deviations between tax-law-based income measures and a more comprehensive definition of income led to a shift in the research agenda on income inequality: a turn from nearly exclusive reliance on tax data to instead starting analysis with national accounts data. Thus, Piketty, Saez and Zucman (2018) attempted to address some of the drawbacks of the earlier Piketty and Saez (2003) approach by allocating national income as defined by the national income and product accounts from the US Bureau of Economic Analysis. But this approach also has drawbacks.

For one, the national-accounts definition of income—while broader than tax-derived measures of income—is not the same as comprehensive Haig-Simons income. This point is typically not emphasized (or even mentioned) in the economics literature, but it is important. Of course, national-accounts income does include many classic forms of economic income that our tax system does not attempt to reach—such as the imputed rent of owner-occupied housing (Mayerhauser and Reinsdorf, 2007). But those imputations go only so far. For example, the national accounts do not attempt to impute household labor production, and they do not account for unrealized increases in asset values. Once again, a sense of the magnitude of this issue is provided by Figure 1.

There is also an important practical tradeoff between measures that start with tax data and national accounts data. The attempt to approach—if never quite arrive at—a comprehensive definition of income comes at the expense of many of the practical virtues described above. It is intrinsically difficult to allocate to individuals income that is not reported on individuals' tax returns. Researchers must instead resort to imputations based on imperfect information and debatable assumptions, which inevitably leads to disagreements. Mirroring our discussion above, we offer an illustrative (and by no means exhaustive) list of six categories in the national accounts without reliable individual data.

A first example: undistributed corporate earnings are included in national income; realized capital gains are not. Undistributed corporate income is part of the value of corporate equities owned by individuals, but it is not synonymous with accrued capital gains and certainly not with realized ones. Without good data to link firms and individuals—and such data does not, to our knowledge, exist in the United States—researchers need some method for imputing undistributed corporate income to individuals. When starting with individual tax data, researchers can rely on proxies for ownership, like dividends and capital-gains realizations. But firms that pay dividends are not necessarily the same firms that retain earnings (after all, paying dividends *reduces* retained earnings). Indeed, the relationship between retained income and payouts is likely heterogeneous over time and across firms. Legal changes like the 1986 Tax Reform Act—which massively expanded the number of "pass-through" corporations and moved whole categories of firms off the corporate income tax—surely did not have a neutral effect on the relationship between dividends and corporate retained earnings. Relying on dividends to impute corporate retained earnings does poorly in Norway, where more data is available to

link firms and owners and thus study this connection (Alstadsæter et al., 2025). Similarly, evidence from Honduras suggests most income at the top of the distribution consists of undistributed corporate profits, even though most individuals at the top do not receive dividends (Carmen et al., 2023).[9]

Second, savings accumulate in pension funds and retirement accounts, but researchers do not, at present, observe the ownership of these accounts. What should be imputed to individuals, and when? Do retirement savings generate Haig-Simons income at the time of contribution and accrual, or the time of withdrawal, after retirement? Most approaches treat retirement savings as income when they are distributed from retirement accounts. But this approach raises a practical problem: national accounts measure retirement savings on an accrual basis, and—if the researcher is committed to allocating the full annual aggregates that appear in the national accounts—the difference between accrual and distribution must be allocated. Retirement income is also behind one of the mistakes that crept up when dealing with imperfect data—as Auten and Splinter document, the approach of Piketty, Saez and Zucman (2018) misinterpreted some rollovers of retirement funds to different IRA and 401(k) accounts as new income—because these mere transfers of assets from one account to another do appear on tax returns. Because large rollovers are concentrated at the top of the income distribution, this had the effect of increasing top income shares.

Third, reliance on national accounts requires choices about the treatment of Social Security—and, more generally, blurs the line between the nominally pre- and after-tax approaches to measuring income. Social Security transfers are funded out of payroll taxes, but they also constitute a component of income that is partially subject to taxation. As noted above, including both the taxes and benefits would create double-counting. In the face of this difficulty, Piketty, Saez, and Zucman choose to exclude payroll taxes and include Social Security benefits in their "pre-tax" measure of income—a departure from national accounts that requires, for the aggregates to reconcile, imputing the difference between payroll taxes and benefits to individuals. Auten and Splinter do the opposite—including payroll taxes and excluding Social Security benefits from their pre-tax income measure. They also, separately, construct a pre-tax-plus-transfers concept of income that includes government transfer benefits. They justify this latter measure as providing a more complete picture of the total resources available for saving, consuming, and paying taxes—and as consistent with a long tradition of government measurement for those purposes, even if it does not conform with the national accounts' notion of income.

This Social Security issue is connected to a fourth, and wider, problem: approaches that use national accounts data—and that appeal to an after-tax perspective on inequality—must allocate *all* current government spending to individuals (including national defense, infrastructure, education, and so forth), because it is part of national accounts. Auten and Splinter assign half of this spending on an equal per-capita basis, while Piketty, Saez, and Zucman assign it proportionally to other income. While one can debate the merits of both approaches (does Elon Musk benefit more from national-defense spending than you do?), there is no consensus on the right approach. (Although, at a minimum, we suspect many Americans would be surprised to learn that—at least according to some researchers—defense spending primarily benefits the top of the income distribution.)

Fifth, national income approaches must deal with difficulties that arise from the treatment of business investment expenses. Tax law has its own (often-changing) rules for how capital investments should be depreciated—that is, allocated over the life of the asset. These rules often do not track the real economic decline of an asset. National accounts, by contrast, attempt to estimate actual economic depreciation. Reconciling national accounts with tax data therefore requires allocating the difference in depreciation. This issue becomes especially important after the 2017 tax reform expanded the "expensing" of certain

[9]Post-publication revisions by Piketty, Saez and Zucman (2018) allocate retained earnings based on their own estimates of equity wealth, but they do not fundamentally change this point, because those estimates rely on a capitalization method that leverages observed capital income.

investments—that is, treating 100 percent of certain investment costs as costs in the year they were made (allowing for an immediate deduction from income), rather than depreciating (and therefore deducting) the cost of the investment over time. Piketty, Saez and Zucman (2024) argue that Auten and Splinter (2024*a*) allocate to partnership owners too little of the "excess depreciation"—that is, the more generous IRS tax treatment of depreciation, relative to the national accounts. In particular, the trio argues that Auten and Splinter assign too much excess depreciation to sole-proprietors. The core of the problem is the data: most excess depreciation is not directly observed on the individual tax returns. In the face of this uncertainty, Piketty, Saez, and Zucman allocate the aggregate excess depreciation of partnerships proportionally to observed partnership income. This approach is appealing in its simplicity. But it might not be correct to assume that the ownership of investment- and capital-intensive partnerships—the ones that generate the greatest excess depreciation deductions, such as utilities and real estate companies—is well-approximated by taxpayers who have the most partnership income (predominantly finance). In their response, Auten and Splinter (2024*b*) discuss and defend their original allocation in detail and report the results from a new analysis based on linking the tax returns of pass-throughs (S-corporations and partnerships) and individuals (based on the work of Love, 2021) to estimate the total (rather than excess) depreciation share of top 1% tax returns. They then make adjustments to approximate the *excess* depreciation that should be allocated to the top 1% of individuals. These results are similar to their original method.

Fundamentally, one approach falls back on the assumption that unobserved income is distributed similarly to some observed quantity. The other approach brings in additional microdata that provides some new distributional information—but that does not match precisely the underlying concepts. Understandably, neither approach is perfect. But the implications are surprisingly large: according to Piketty, Saez and Zucman (2024), this one issue accounts for a 1.2pp difference in the change in the top 1% income share between 1979 and 2019—or 30% of the total difference in trends between the two published papers. This issue also offers a nice illustration of how much is effectively hidden from a casual reader. A quick search for "net operating loss," "depreciation," "partnership," and "capital consumption adjustment" in the published versions of the two papers does not warn a reader of this issue. Auten and Splinter note: "We also account for other differences, such as faster depreciation in tax data than in national accounts due primarily to expensing on tax returns. See the online appendix for details." They discuss their assumptions in the appendix, asserting that they have little impact on results. Piketty, Saez and Zucman (2018) do not comment on this particular issue at all.

A final but related issue concerns the potential re-ranking of individuals in the income distribution. When starting with income tax data, individuals are simply ranked in the income distribution according to their observed reported income. But bringing in other data requires adjustments—most importantly in this context, adjustments to individuals with low reported taxable income who are deemed to be high income when their income is "corrected." Such adjustments result in changes to the composition of the higher-income groups. This generates issues with how income that's responsible for moving individuals to top groups is treated. As a concrete example, how should we distribute the aggregate amount of income subject to tax evasion to individuals who otherwise report losses? Should we move all of them up in the income distribution uniformly, or should we move some of them up by a lot? If the latter, what assumptions should be used? Some partial microdata may be available in these and other cases—but by definition such data is not going to be comprehensive, and one needs auxiliary assumptions.

Again, this list of issues concerning the allocation of national accounts across the income distribution is not exhaustive. But it illustrates the difficulties that arise when no precise data is available. (And if income with an unknown distribution is simply allocated proportionally to observed income, then the shift to using national accounts data would be pointless: it would just replicate the same income distribution we observe

in the administrative tax data.)

The turn to national income can still be defended as representing a kind of middle-ground—a middle-ground between the under-inclusive concept of income derived from tax law and the quixotic, unreachable Haig-Simons ideal. But these approaches come with tradeoffs. Narrower income concepts have better data and are easier to measure. But they are narrow, and what is included in income categories can evolve over time. Broader income concepts are, by definition, less narrow. But they are harder to measure. The harder one searches for the Haig-Simons ideal, the farther one strays from reliable data.

Practically, once one leaves the comfortable confines of microdata—and attempts to change the concept of income or the unit of observation—a number of problems arise. First, one needs to measure the magnitude of the additional income. Reliance on national accounts helps address this, but comes with other limitations. Second, there is the question of how the additional income is distributed relative to observed income. An attractive assumption that researchers often use is to assume that what we do not know is proportional to what we know. Taken to its extreme—if we assume that all unobserved income is distributed proportionally to observed income—this assumption adds no additional information. Much of the controversy over measuring income inequality stems from the choice between using this kind of easy-to-understand assumption, and approaches that use imperfect auxiliary information to construct an alternative allocation. Based on various auxiliary data sources, Auten and Splinter (2024*a*) effectively move farther away from a proportionality assumption than Piketty, Saez and Zucman (2018) do. Third, allowing for heterogeneity conditional on observed income changes the ordering of individuals in the distribution.

Many of the difficulties that arise in accounting for tax evasion, excess depreciation, and undistributed corporate earnings stem from these three issues. Even if one settles on the total amount of income to be allocated, variation in who this income belongs to produces uncertainty about where it falls in the observed income distribution and how the allocation of this income modifies the observed ranking.

Why it matters, and the next frontier

We conclude by offering a few thoughts on the implications of the issues discussed above for income measurement and inequality measurement.

First, and perhaps most obviously, there is a need for greater sensitivity to the limitations of existing data, and a need for better data. Particularly valuable is data that can link information on aggregate income with individuals. Researchers in some other countries, for example, have been constructing increasingly robust data linking firm-level income with owners: for Canada, see Wolfson et al. (2016); for France, see Bach et al. (2024); for the Netherlands, see Lejour (2024); for Chile, see Fairfield and Jorratt De Luis (2016); for Norway, see Alstadsæter et al. (2025); and in Honduras, see Carmen et al. (2023). This international evidence suggests that correctly allocating business income to owners can make a large difference in top income shares relative to standard imputation methods.

Some recent work on the US economy has made improvements on this front. In particular, Smith et al. (2019) link the universe of pass-through firms (firms where the entity-level income is passed along to owners as taxable income, and that do not pay a separate corporate tax) to their owners for the years 2001–2014. Love (2021) refines and extends this further for partnerships. But such work remains incomplete for some years of greatest interest—such as the 1980s, in which there were both major changes in tax law and a sudden observed increase in top income shares. Such work also remains incomplete for other kinds of firms. For example, we do not know of any data that links the universe of C-corporations (firms that pay the separate corporate tax and that retain more income) to their individual owners.

Second, we need more attention to legal changes that affect the measurement of inequality. For example, the line between individual and firm income has shifted in the United States, stemming from legal changes in how businesses are taxed (Clarke and Kopczuk, 2017, 2025). Before the Tax Reform Act of 1986, for example, one of the most common (and widely publicized) tax-planning strategies in the United States was to use C-corporations as a tax-deferral device. By retaining earnings inside a corporation, taxpayers—especially taxpayers with high marginal rates—could defer individual taxation and pay only the lower, entity-level corporate tax. But the 1986 tax reform reduced individual marginal rates and changed entity-level rules, and the incentive for individuals to use C-corporations to shelter income declined. Indeed, several hundred thousand C-corporations disappeared from the US economy in the years after 1986—replaced by pass-through entities, in the form of S-corporations and, starting in the 1990s, partnerships. When firms switch in this fashion, a greater share of business income will show up on individual tax returns, and firms will change their retention behavior. How should we think about these legal changes? To what extent do changes in US income inequality reflect real economic changes versus legal changes that merely affect what we observe? Similarly, to what extent do cross-country differences reflect real differences, versus differences in what national laws make measurable?

Third, we see an opportunity to appreciate the tradeoffs between different approaches to income inequality—rather than preferring a tax-based approach or a distributional national accounts approach as strictly better or worse. With either approach, we are using an imperfect measure of comprehensive income or welfare. Tax data provides a limited view of comprehensive income. But tax law defines a concept of income that can be defended as representing *society's* answer—through the democratic lawmaking process—to a hard question: "what should be counted for determining one's ability to pay for the cost of government?" The answer to this question will naturally evolve over time. We are still observing this definitional process in action today. Recent debates over taxing accrued (rather than realized) capital gains, and recent litigation over the scope of the Sixteenth Amendment, reveal that there are ongoing disagreements over exactly what constitutes, and should constitute, "income."

Fourth, we urge modesty and skepticism in the face of the profusion of income measures and inequality results. Some studies in this area produce a monolithic result with no consumer warning label: 'this is the income share of the top 1 percent.' Our Figure 2 reports a number of those estimates. But it would be better to think about inequality estimates as representing bounds that emerge under different assumptions. Alvaredo et al. (2024) adopt something like this approach for the case of Latin America. They note that, since no single method for measuring inequality "is fully convincing at present, we are left with (often wide) ranges, or bands, of inequality as our best summaries of inequality levels." Thinking of individual inequality results as falling within a wide band of reasonable measures would restore a sense of perspective to this literature.

Finally, we stress that the income concept itself is nebulous and evolving. Measures that researchers work with are influenced by what "market" transactions they can observe. Likewise, governments tax what they can measure. But these things change over time. As the line between household production and market transactions shifts, "income" changes. As the retirement landscape changes—and changes the balance of various public, private, and tax-deferred savings mechanisms—when and how "income" is recognized evolves. And on and on.

One can choose a single income concept and run with it. But no single concept is perfect.

References

Alstadsæter, Annette, Martin Jacob, Wojciech Kopczuk, and Kjetil Telle. 2025. "Acccounting

for Business Income in Measuring Top Income Shares: Integrated Accrual Approach Using Individual and Firm Data from Norway." *Journal of the European Economic Association*. Conditionally accepted.

Alstadsæter, Annette, Wojciech Kopczuk, and Kjetil Telle. 2014. "Are Closely-Held Firms Tax Shelters?" *Tax Policy and the Economy*, 28: 1–32.

Alvaredo, Facundo, François Bourguignon, Francisco Ferreira, and Nora Lustig. 2024. "Inequality bands: Seventy-five years of measuring income inequality in Latin America." World Inequality Lab Working Papers 08.

Andrews, William D. 1983. "The Achilles' Heel of the Comprehensive Income Tax." In *New Directions in Federal Tax Policy for the 1980s*. Ed. Charles E. Walker and Mark A. Bloomfield, 278–280. Ballinger Publishing Company.

Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2011. "Top Incomes in the Long Run of History." *Journal of Economic Literature*, 49(1): 3–71.

Auerbach, Alan J., Laurence J. Kotlikoff, and Darryl Koehler. 2023. "US Inequality and Fiscal Progressivity: An Intragenerational Accounting." *Journal of Political Economy*, 131(5): 000–000.

Auten, Gerald, and David Splinter. 2024*a*. "Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends." *Journal of Political Economy*, 132(7).

Auten, Gerald, and David Splinter. 2024*b*. "Reply to Piketty, Saez, and Zucman (2024): Income Inequality in the United States." Office of Tax Analysis, U.S. Treasury and Joint Committee on Taxation, mimeo.

Bach, Laurent, Antoine Bozio, Arthur Guillouzouic, and Clément Malgouyres. 2024. "Do Billionaires Pay Taxes?" Paris School of Economics, mimeo.

Bittker, Boris I. 1967. "A "Comprehensive Tax Base" as a Goal of Income Tax Reform." *Harvard Law Review*, 80(5): 925–985.

Blum, Walter J., and Harry Jr. Kalven. 1952. "The Uneasy Case for Progressive Taxation." *University of Chicago Law Review*, 19(3): 417–429.

Bridgman, Benjamin, Andrew Craig, and Danit Kanal. 2022. "Accounting for Household Production in the National Accounts: An Update 1965–2020." *Survey of Current Business*, 102(2).

Brooks, John R. 2018. "The Definitions of Income." *Tax Law Review*, 71: 253.

Brownlee, W. Elliot. 2016. *Federal Taxation in America*. 3rd ed., Cambridge University Press.

Carmen, Giselle Del, Santiago Garriga, Wilman Nuñez, and Thiago Scot. 2023. "Two Decades of Top Income Shares in Honduras." World Bank Policy Research Working Paper 10722.

Clarke, Conor, and Wojciech Kopczuk. 2017. "Business Income and Business Taxation in the United States since the 1950s." *Tax Policy and the Economy*, 31: 121–59.

Clarke, Conor, and Wojciech Kopczuk. 2025. "Income Inequality and the Corporate Sector." Washington University and Columbia University, mimeo.

Congressional Budget Office. 2023. "The Distribution of Household Income in 2020." https://www.cbo.gov/publication/59509.

Dahl, Gordon B., and Michael R Ransom. 2002. "The 10% Flat Tax: Tithing and the Definition of Income." *Economic Inquiry*, 40(1): 120–137.

Deaton, Angus. 2020. "GDP and Beyond: Summaries from the 2020 Annual Meeting of the American Economic Association." *Volume 100, Number 6*.

Fairfield, Tasha, and Michel Jorratt De Luis. 2016. "Top Income Shares, Business Profits, and Effective Tax Rates in Contemporary Chile." *Review of Income and Wealth*, 62: S120–S144.

Frazis, Harley, and Jay Stewart. 2011. "How Does Household Production Affect Measured Income Inequality?" *Journal of Population Economics*, 24(1): 3–22.

Gautham, Leila, and Nancy Folbre. 2024. "Household Production and Inequality in Living Standards in the U.S., 1965–2018." Paper prepared for the 38th IARIW General Conference, August 26-30, 2024.

Glogower, Ari. 2023. "A Basic Needs Baseline for Distributional Analysis." *Brigham Young University Law Review*, 48(6): 1697–1768.

Graetz, Michael, and Anne Alstott. 2022. *Federal Income Taxation, Principles and Policies.* 9th ed., Foundation Press.

Gruber, Jonathan. 2022. *Public Finance and Public Policy.* 7th ed., Macmillan Learning.

Haig, Robert M. 1921. "The Concept of Income—Economic and Legal Aspects." In *The Federal Income Tax*. 1–28. New York:Columbia University Press.

Hayashi, Andrew T. 2023. "The Federal Architecture of Income Inequality." *Virginia Public Law and Legal Theory Research Paper 2023-80*.

Joint Committee on Taxation. 2022. "Estimates of Federal Tax Expenditures for Fiscal Years 2022-2026." Report prepared by the Staff of the Joint Commitee on Taxation Report JCX-22-22.

Kuznets, Simon. 1953. *Shares of Upper Income Groups in Income and Savings.* New York:National Bureau of Economic Research.

Larrimore, Jeff, Richard V. Burkhauser, Gerald Auten, and Philip Armour. 2021. "Recent Trends in US Income Distributions in Tax Record Data Using More Comprehensive Measures of Income Including Real Accrued Capital Gains." *Journal of Political Economy*, 129(5): 1319–1360.

Leite, David. 2024. "The Firm as Tax Shelter: Micro Evidence and Aggregate Implications of Consumption Through the Firm." Paris School of Economics, mimeo.

Lejour, Arjan. 2024. "The economic position of the wealthy in the Netherlands since 2006." Tilburg University.

Lockwood, Benjamin B., Charles Nathanson, and E. Glen Weyl. 2017. "Taxation and the Allocation of Talent." *Journal of Political Economy*, 125(5): 1635–1682.

Love, Michael. 2021. "Where in the World Does Partnership Income Go? Evidence of a Growing Use of Tax Havens." SSRN Working Paper. Available at https://ssrn.com/abstract=3985535.

Mayerhauser, Nicole, and Marshall Reinsdorf. 2007. "Housing Services in the National Economic Accounts." *Bureau of Economic Analysis*, Updated as of September 11, 2007, `https://www.bea.gov/sites/default/files/methodologies/RIPfactsheet.pdf`.

Murphy, Liam, and Thomas Nagel. 2004. *The Myth of Ownership: Taxes and Justice.* Oxford:Oxford University Press.

Musgrave, R. A. 1967. "In Defense of an Income Concept." *Harvard Law Review*, 81: 44.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913-1998." *Quarterly Journal of Economics*, 118: 1–39.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *The Quarterly Journal of Economics*, 133(2): 553–609.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2024. "Income Inequality in the United States: A Comment." Paris School of Economics and UC Berkeley, mimeo.

Simons, Henry. 1938. *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy.* Chicago:University of Chicago Press.

Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations.* London:W. Strahan and T. Cadell.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick. 2019. "Capitalists in the Twenty-First Century." *Quarterly Journal of Economics*, 134(4): 1675–1745.

The Department of Commerce. 1934. "National Income, 1929-32." *U.S. Government Printing Office*, Government Report.

Thuronyi, Victor. 1990. "The concept of income." *Tax Law Review*, 46(1): 45–105.

von Kleinwächter, Friedrich Ludwig. 1896. "Das Einkommen und seine Verteilung [Income and its Distribution]." *Hand- and Textbook of Political Sciences*, 5.

Wolfson, Michael, Michael Veall, Neil Brooks, and Brian B. Murphy. 2016. "Piercing the Veil – Private Corporations and the Income of the Affluent." *Canadian Tax Journal/Revue Fiscale Canadienne*, 64(1).

Zelenak, Lawrence. 2006. "Taxing Endowment." *Duke Law Journal*, 55(4): 1145.

Appendix

This short appendix describes the basic construction of figures in the document. Full details are in the online appendix.

Figure 1: Layering income concepts. This figure begins with the breakdown of national income that appears in Piketty, Saez and Zucman (2018). That gives us the share of national income, both labor and capital, that appears on individual tax returns. We also take their measures of labor income and capital income that do not appear on individual tax returns. We then add in aggregate government transfers and other aggregate government spending (also from Piketty, Saez and Zucman (2018)) as well as realized capital

gains drawn from IRS and CBO data. We do not make adjustments for accrual in including capital gains, so there is some potential here for double-counting. Our estimate of the value of household production is drawn from the BEA satellite accounts, which cover all years in the figure. Our rough estimate of the value of leisure—which is offered to provide only a general sense of how one might think about the order of magnitude—starts with the American Time Use Survey measure of the hours each day spent in leisure, multiplied by the annual adult population and days in the year. (Before 2003, we assume that individual daily leisure is the same as the average of measured years after 2003.) For the wage rate, we use the most recent federal minimum wage, bench-marked to 2010 (the first full year it was law), indexed to other years using the Social Security Administration's measure of wage growth. The purpose here is to provide only a rough and illustrative sense of the magnitude, not anything approaching a precise treatment of the issue, which we suspect is not possible—there is no consensus on how to think about the value of leisure. (One of the many nice puzzles one might consider here—and which we do not attempt to address—is: How to think about the leisure of children?)

Figure 2: **Different measures of the top 1%**. Here we pull directly from the appendix tables of several works: The CBO "Additional Data for Researchers" available through their ongoing series on household income distribution; Larrimore et al. (2021) we draw directly from table A2 in their appendix supplement; Piketty and Saez (2003) we draw directly from their updated distribution tables (A1 and B1); Piketty, Saez and Zucman (2018), we draw from TA1, TB1, and TC1 in their distribution tables; Auten and Splinter (2024*a*) we draw from T-A1 in their appendix spreadsheet.