Policy Research Working Paper 10460

Measuring Tax Progressivity in Low-Income Countries

Alastair Thomas

WORLD BANK GROUP
Macroeconomics, Trade and Investment Global Practice
May 2023

Policy Research Working Paper 10460

Abstract

In addition to raising revenue, tax systems have a key role to play in achieving countries' equity goals. An important first step in assessing whether a particular tax system is appropriately supporting those equity goals is to assess the degree of progressivity of the tax system. However, measuring the progressivity of a tax system is challenging for several reasons. Most significantly, data limitations have a large impact on what taxes can be examined, what measurement approaches can be adopted to examine those taxes, and how reliable those approaches are likely to be. This is particularly important for low-income countries where data limitations are typically most significant. Furthermore, there is no single definition of progressivity, or single method for measuring the degree of progressivity of a tax—thereby necessitating a range of decisions and value judgements to be made in any analysis. This paper examines and assesses the potential approaches that may be adopted to measure the progressivity of tax systems, and proposes a set of metrics to assess tax progressivity in low-income International Development Association (IDA) countries.

This paper is a product of the Macroeconomics, Trade and Investment Global Practice. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The author may be contacted at athomas@worldbank.org.

Measuring Tax Progressivity in Low-Income Countries

Alastair Thomas

JEL codes: H2, H22, H23

Key words: taxes, progressivity, redistribution

Alastair Thomas is a Senior Economist in the Fiscal Policy and Sustainable Growth Unit of the Macroeconomics, Trade and Investment Global Practice at the World Bank. Helpful comments were received from Jaffar Al-Rikabi, Chadi Bou Habib, Chiara Bronchi, Marianne Caballero Parra, Matthew Collin, Gabriela Inchauste, Sebastian James, Gaute Solheim, Mariano Sosa, Charles Vellutini, and Matthew Wai-Poi. This paper was supported by the World Bank's Global Tax Program which provides support to countries in strengthening domestic resource mobilization.

Contents

Introduction

In addition to raising revenue, tax systems have a key role to play in achieving countries' equity goals. However, many tax systems in developing countries currently exhibit very little progressivity, relying heavily on indirect taxes and raising relatively little revenue from key sources of progressivity such as the personal income tax, capital income taxes and property taxes. As countries look to raise revenue following multiple crises, it is therefore important that tax reforms are implemented in ways that increase progressivity and ensure that better off households pay their fair share of tax.

In this context, the World Bank is currently undertaking a project examining how countries can increase the progressivity of their tax systems in order to both enhance domestic resource mobilization and promote inclusive growth. At the same time, for IDA20, the World Bank has committed to assessing the progressivity of tax systems in International Development Association (IDA) countries over FY2023-25.[1] This paper – which focuses on measurement of tax progressivity – contributes to both these work streams. The paper examines potential approaches to assessing the progressivity of tax systems and then, based on this analysis, proposes a set of progressivity metrics to be adopted over the next three years to assess tax progressivity in IDA countries.

Measurement of tax progressivity is a challenging issue for several reasons. Most significantly, data limitations have a large impact on what taxes can be examined, what measurement approaches can be adopted to examine those taxes, and how reliable those approaches are likely to be. This is particularly important for IDA countries where data limitations are typically most significant. Additionally, there is no single definition of progressivity, or single method for measuring the degree of progressivity of a tax, thereby necessitating a range of decisions and value judgements to be made to undertake any assessment. The appropriate definition of progressivity will depend on the underlying question being asked, and will potentially impact the way in which the degree of progressivity can be measured. Contingent upon the definition chosen, there are several approaches that can then be adopted to measure the degree of progressivity of a tax, each with benefits and drawbacks that must be weighed in determining the most appropriate approach. Finally, measurement of tax progressivity – particularly when based on microdata – can be extremely resource intensive, and this must be taken into account, particularly in the current context where the progressivity of a large number of countries needs to be assessed.

The first impact that data limitations have is on the range of taxes that can be examined. Ideally, policy makers should be interested in assessing the progressivity of the entire tax system (and indeed, more broadly, in the progressivity of the entire fiscal system). However, data limitations make it virtually impossible to undertake a detailed analysis of the progressivity of the entire tax system, yet alone the entire fiscal system – although the in-depth CEQ analyses undertaken by the CEQ Institute and the World Bank go a long way towards this by assessing the distributional impacts of a number of major components of tax, transfer and expenditure systems (see Commitment to Equity Institute Data Center on Fiscal Redistribution, 2022; Lustig, 2018).

[1] For more information on IDA, see: https://ida.worldbank.org/en/home.

A detailed analysis of progressivity (or of the broader concept of redistribution), requires a welfare ranking of individuals or households (or at least of a set of disaggregated groupings of individuals or households), and for tax liabilities of each individual or household (or group) to be ascertained. This can be achieved for certain taxes where: (a) the economic incidence of the tax can be determined with at least a reasonable degree of confidence; and (b) the necessary data is available distinguishing rich from poor among those that the economic incidence of the tax falls on.

The most notable example is the personal income tax (PIT), as tax liabilities can be linked to taxpayers at different income levels. Most simply this can be done through analysis of the underlying PIT rate schedule, or in a more complex manner through use of household income or consumption survey data (potentially augmented with tax return data), and simulation of tax liabilities. Furthermore, there is broad consensus that the incidence of the PIT falls predominantly, if not entirely, on the individuals earning the income. Consequently, the majority of empirical analyses of tax progressivity focus on the PIT. This attention is unsurprising given that the PIT is the primary tax instrument used in most countries to achieve distributional objectives. Even so, there are significant data limitations associated with PIT analyses, such as the underreporting of top incomes in survey data, tax avoidance and tax evasion, which affect the accuracy of microdata-based analyses.[2] Detailed microdata-based analyses are also highly resource intensive and often suffer from significant time lags.

For countries with comprehensive systems (where labor and capital income are taxed together under the PIT), PIT analyses can capture taxes imposed on labor and personal capital income. In schedular systems (where personal capital income is taxed separately, and often at lower rates than labor income) it is more challenging to capture personal capital income. Use of tax return data can often aid the capture of capital income taxes, but introduces other potential problems regarding the representativeness of the data.

Distributional analyses can also be undertaken for social security contributions (SSCs), although incidence becomes more questionable, especially regarding employer contributions. Furthermore, the treatment of the pension component of social contributions raises the more general question regarding the degree to which contributions reflect a pure tax as compared to compulsory savings. In some cases, information on property taxes is also obtainable from household surveys enabling distributional analysis.[3]

Meanwhile, consumption survey microdata can be used to model the progressivity of consumption taxes, such as the value-added tax (VAT) and excise taxes (including health- and energy-related taxes). In such analyses, a standard assumption is made that incidence falls entirely on consumers, although empirical evidence suggests this is not always the case (see, e.g., IHS, 2011; Benedek et al., 2019; Benzarti et al.,

[2] Underreporting of top incomes in household income survey data is a key issue being considered as part of the overall tax progressivity and inclusive growth project. Two approaches are currently being developed: (1) a methodology utilizing information from tax return data to adjust top incomes in income surveys for underreporting; and (2) a pareto distribution-based approach to adjust top incomes in income surveys when tax return data is not available. A further problem is that top incomes in tax return data will not reflect tax avoidance (or evasion). For example, income that is retained in a closely held company will typically not be realized (and taxed) at the personal level until distributed, enabling corporate structures to be used for tax avoidance. In this regard, Alstadsaeter et al. (2016) link individual and firm level data in Norway in order to attribute business income to individuals as it accrues, finding that this significantly increases the share of total income held by top income earners.

[3] There remains debate in the literature regarding the extent to which recurrent property taxes exhibit features of a benefits tax and hence are closer in nature to a payment for services (and consequently have no distributional impact) as opposed to a tax. For a summary, see Zodrow (2007).

2020). Additional challenges also arise in terms of potential differences between annual and lifetime impacts of consumption taxes (see, e.g., Thomas, 2022), and regarding combining analysis of income and consumption taxes given that their underlying bases differ.

While distributional analysis of the above taxes is challenging, the data exists to make it possible. With regard to most other taxes, there is often simply no data available to distinguish poor and rich taxpayers and allocate tax liability among them. For example, the corporate income tax (CIT) is applied to corporations rather than individuals, making it extremely challenging to attempt to allocate the CIT paid by the corporations to the individuals and households that will eventually bear the economic incidence of the tax. It is nevertheless possible to draw the broad conclusion that the CIT is progressive – based on the conventional view that capital owners bear the majority of the CIT, and given that capital income is disproportionately earned by high-income households (see, e.g., Nunns, 2012). That said, empirical evidence does suggest that at least some of the CIT may be borne by workers (in the form of lower wages), which would reduce the progressivity of the tax (see, e.g., Fuest et al., 2018).[4] While similar, high-level views of the likely distributional effects of other taxes can be made (e.g., the progressivity of a net wealth tax), detailed measurement of the degree of progressivity is not typically feasible (at least not in the absence of administrative data).

All in all, these limitations mean that detailed measurement of the *degree* of progressivity is limited to the PIT, and to a lesser extent, VAT and excise taxes. These taxes, and particularly the PIT, are consequently the main focus of this paper. However, because of the need to provide an assessment for IDA countries of the progressivity of the entire tax system, the paper also undertakes an examination of the extent to which aggregate tax revenue data – which is available for a larger number of countries in a timely manner – can be used to assess the likely degree of progressivity of tax systems in IDA countries.

It should also be borne in mind that, beyond computation of progressivity metrics, there are additional ways to help assess whether countries are taking steps to increase the progressivity and fairness of their tax systems. For example, many countries are currently implementing reforms to their international tax rules to limit the extent to which individuals and companies that operate in multiple jurisdictions can avoid or evade tax. These may include unilateral measures, or those undertaken as part of current multilateral efforts to prevent base erosion and profit shifting (BEPS) activities, address the tax challenges of digitalization (including the proposed adoption of a global minimum tax), and enable the exchange of taxpayer information amongst tax administrations. While beyond the scope of this paper, assessment or stocktaking of the various steps that IDA countries have taken to implement such measures would be informative.

This paper is structured as follows: Section 2 briefly introduces various concepts of tax progressivity and progressivity measurement. Sections 3-5 then consider the three main approaches to measuring tax

[4] There indeed remains considerable academic debate regarding the incidence of the CIT. See also, e.g., Gale and Thorpe (2022). Harris (2009) finds that even if labor bears the majority of the CIT, it remains progressive. While empirical attempts can be made to allocate the CIT to taxpayers based, for example, on tax return microdata specifying capital income (see, e.g., Nunns, 2012; Rohaly et al., 2005), this uncertainty regarding incidence, and hence the appropriate amount of CIT to allocate to capital income earners, makes it an extremely challenging exercise. A further complication arises regarding how foreign ownership should impact progressivity measurement. For example, it is arguable that any CIT borne by a foreign owner could be considered to have a neutral distributional impact as it is not paid by poor or rich domestic households. This may be a particular issue in developing countries where many large companies may be predominantly foreign owned.

progressivity, detailing the pros and cons of each approach, and illustrating their use drawing on results from the literature, where possible for IDA countries. Section 6 then considers the use of aggregate revenue data to help assess tax progressivity. Drawing on the preceding sections, Section 7 then proposes a set of indicators to use to assess the progressivity of tax systems in IDA countries.

1 Defining and measuring progressivity

The commonly accepted definition of a progressive tax is one where the average tax rate increases as income (or other relevant welfare metric[5]) increases. Meanwhile, for a regressive tax the average tax rate decreases as income increases, and for a proportional tax the average tax rate is constant as income increases. Progressivity can equivalently be specified as requiring the marginal tax rate to be higher than the average tax rate (see, e.g., Musgrave and Thin, 1948).

This concept of progressivity derives from concepts of vertical equity, in particular the well-known "ability-to-pay" principle, as well as "equal sacrifice" principles. The ability-to-pay principle, which has its origins in Smith's (1776) *Wealth of Nations*, states that someone with greater ability to pay tax should pay more tax (although how much more is not articulated).[6] Ability-to-pay is typically measured in relation to income. Equal sacrifice, which originates from Mill (1848), instead is linked more closely with utility. The principles of equal *absolute* sacrifice and equal *proportional* sacrifice, respectively, require that taxpayers should face the same absolute and proportional reduction, respectively, in utility from the imposition of a tax. Assuming that the marginal utility of income decreases with income, both concepts imply that richer households should pay more in tax than poorer households, although they do not necessarily imply that they should pay more as a proportion of their income, as this will depend on the underlying utility functions (see, e.g., Young, 1987).[7] It is important to note that the "increasing average tax rate" definition of progressivity requires not just that the total tax liability increases with income, but that it increases more than proportionally with income.

When it comes to *measuring* progressivity, this basic definition becomes problematic. In particular, progressivity is defined relative to a point in the income (or other relevant welfare metric) distribution. This makes it difficult to assess (and compare) the overall degree of progressivity of a tax (and is especially problematic where the average tax rate varies non-monotonically over the relevant distribution, as may occur, e.g., with the VAT or excises). A second problem with the basic definition is that it is vague regarding exactly what is being measured: the progressivity of the tax rate structure,[8] or the progressivity of the tax

[5] It is necessary to choose a welfare metric to rank households from poor to rich. While typically for an income tax this will be income or equivalized income, for the VAT it makes more sense to measure the tax burden relative to (equivalized) expenditure (see Thomas, 2022; IFS, 2011).

[6] Adam Smith proposes both the benefits principle and ability-to-pay principle. As noted by Musgrave and Musgrave (1989), the benefits principle derives also from earlier work by, amongst others, Locke and Hobbes.

[7] Theoretically, progressivity metrics can be based on the change in utility induced by a tax (e.g., compensating or equivalent variations), but, in practice, metrics are generally based on tax paid as per the basic increasing average tax rate definition.

[8] In this paper, we use the term "rate structure" in a broad sense, referring to the combined impact of all tax rules in determining the tax liability including not just the progressive rate schedule, but also any applicable allowances, deductions and tax credits.

actually paid by households? This distinction is particularly relevant for the PIT, where there may be considerable interest in studying just the underlying tax structure.

Historically, there have been two main approaches to measuring income tax progressivity: examining the progressivity of the tax structure; and examining the progressivity of the taxes paid. Musgrave and Thin (1948) distinguished these as "structural progressivity" and "effective progressivity" measures.[9] A third, more recent, concept is "progressive capacity" which captures elements of both structural and effective progressivity.

"Structural progressivity" focuses purely on the tax rate structure, and typically produces "local" progressivity metrics (i.e., measures of progressivity at different points along the income distribution) based on the basic "increasing average tax rate" definition. A strength of this approach is that it is easy to compute and requires no information beyond the underlying tax rules. However, this is a cost also. In addition to producing only local measures of progressivity, it entirely ignores the underlying income distribution. Effective progressivity measures address this concern, but at the cost of greater data needs.

"Effective progressivity" considers the interaction of both the tax schedule and the underlying income distribution, and typically produces "global" progressivity metrics (i.e., a single figure assessing the progressivity of the tax in question). The most commonly used global progressivity metric is the Kakwani (1977) index which, for an income tax, defines the overall progressivity of a tax based on the degree of deviation in the distribution of tax payments away from the distribution of pre-tax income. (Full detail on the calculation of the Kakwani index is provided in section 3). This global definition of progressivity is effectively an extrapolation from the basic average tax rate definition, requiring now a global rather than local deviation away from proportionality.

Effective progressivity arguably provides a far more complete picture of the progressive impact of a tax because it is based on the actual tax liabilities of different households. For example, it is inconsequential to apply a very high marginal tax rate on a particular income range if there are no taxpayers actually in that income range. It also enables the overall redistributive impact of the tax to be examined, and the contribution of progressivity towards this. The key drawback of the effective progressivity approach is that it is data intensive. It requires information on both the tax system and on the underlying distribution of income (ideally microdata, although disaggregated data can also be used). In contrast, the structural progressivity approach only requires information on the tax system itself.

"Progressive capacity" measures are a type of hybrid approach, producing global metrics of progressivity, but attempting to abstract away from the underlying pre-tax income distribution to focus on the tax structure. This may be particularly useful for cross-country comparison if the analyst is interested in comparing tax rate structures, because differences in income distributions across countries would mask to some extent the comparative progressivity of the different tax structures when using an effective progressivity metric. The approach essentially looks to model the different tax schedules of countries using a common income distribution. However, a drawback of this approach is that the progressive capacity metrics, and consequent country rankings, are not necessarily invariant to the choice of 'base' income distribution. Recent papers have looked to address this base sensitivity problem, though they rely on strong assumptions.

[9] Sometimes also distinguished as progression (structural progressivity) vs progressivity (effective progressivity).

Each of these three progressivity concepts, and the metrics that can be adopted under each, are examined in more detail in the next three sections of the paper. Effective progressivity is considered first as it is the most commonly adopted approach.

2 Measuring effective tax progressivity

The most common and widely applicable type of progressivity measurement is effective progressivity measurement. Effective progressivity metrics consider the actual tax liabilities of individuals or households rather than just the tax schedules. Subject to data availability, effective progressivity metrics can be calculated for the PIT, SSC, VAT as well as various excise taxes (such as those on fuels used for heating and transport, electricity, and on alcohol, tobacco and sugar-sweetened beverages).

This section first discusses data requirements and key methodological issues before presenting the most commonly used global measures of structural progressivity. It then presents examples of their use, focusing where possible on IDA countries, and finally summarizes the strengths and weaknesses of effective progressivity metrics.

2.1 Data, welfare metric, and unit of analysis

To calculate effective progressivity metrics, data is needed that provides an appropriate welfare metric to rank individuals or households (or at least of a set of disaggregated groupings of individuals or households) from poorest to richest, and enables the tax liabilities of each individual or household (or group) to be ascertained. Ideally, microdata is used, although (simplified) calculations can also be made with semi-aggregated data, such as average income levels across percentiles or deciles.

To examine the progressivity of the PIT, income survey data will typically be used, which can be weighted from the representative sample to reflect the population. Income and household characteristics are included to enable the actual tax liabilities of taxpayers to be modelled, including non-standard reliefs where the required information is included in the survey. To examine the VAT or excise taxes, household budget survey data can be used as these surveys typically provide a breakdown of total consumption across a large number of consumption categories, enabling differentiation between consumption subject to different VAT rates and different excise tax rates.

Before calculating progressivity or other metrics, some adjustments (in addition to basic data cleaning) may be made to the microdata to adjust for limitations associated with that data. For example, a common problem with income survey data is that incomes at the top end are often underreported, and various methods can be applied to adjust for this (see, e.g., Blanchet et al., 2022; Ruiz and Woloszko, 2016). Meanwhile, in the absence of income survey data, if consumption survey data is being used to model the PIT, adjustments based on savings patterns can be made to proxy better for income data. A parallel paper examines such techniques (including greater use of tax return microdata) in detail. Further discussion of these techniques is left to that paper.

Additional adjustments may also be made to specify the welfare metric and the unit of analysis. For PIT analysis, where income data is available it will be the preferred starting point in calculating a welfare

metric.[10] For VAT, although some studies have used income as the welfare metric, this can produce misleading results, and the use of consumption should be preferred (see, e.g., Thomas, 2022; IFS, 2011).

Even once income (or consumption) is chosen as the base for the welfare metric, the question then arises whether it should be household income or individual income. Studies have taken both approaches, however there is a broad consensus that individual income should be preferred. Choosing the individual rather than the household as the unit of analysis ensures that equal weighting is given in the analysis to the welfare of each individual. In contrast, use of the household as the unit of analysis would mean that the welfare of a one-person household would be given as much weight as that of a large household with many individuals. That approach, while often adopted for ease of computation, is difficult to justify for welfare analysis.

The key difficulty in applying the individual as the unit of analysis is the fact that income (and consumption) is often only provided in surveys on a household basis. To adjust the unit of analysis to the individual it is therefore necessary to multiply household survey weights by household size. Implicit in this approach is the assumption of equal sharing of resources within a family so that the measured welfare of each household member (whether income or expenditure) is identical.

Prior to adjusting the welfare metric to an individual basis, there is also a strong case to make adjustments to ensure that households with different compositions are comparable. This is referred to as equivalization and involves dividing household income by an "equivalence scale". The simplest option is to just divide household income by the number of individuals in the household. However, this adjustment does not account for differing degrees of need within a household (e.g., a child typically needs less food than an adult to maintain the same welfare level). Additionally, households can be expected to benefit from economies of scale in relation to heating and certain other costs. Two common equivalence scales used are the square root scale and the OECD-modified scale. The square root scale simply takes the square root of household size. The "OECD modified" scale gives a fixed weighting of 1 to the first adult household member, 0.5 to the second and additional household members aged 14 and over, and 0.3 to each child under 14.[11]

2.2 Global indicators of effective progressivity[12]

Musgrave and Thin (1948) first proposed a global measure of effective progressivity as the ratio of (one minus) the Gini coefficient of post-tax income to (one minus) the Gini coefficient of pre-tax income. However, more recent literature has distinguished the concepts of progressivity and redistribution more distinctly. Indeed, Kakwani (1977) pointed out that the Musgrave and Thin (1948) metric was actually a

[10] Although there are strong arguments that consumption provides a better reflection of the wellbeing of poor households, the fact that income is the base of the tax to be modelled creates a natural preference for income data as this will maximize the accuracy of the tax modelling.

[11] A more elaborate option is the parametric scale introduced by Cutler and Katz (1992). The scale is calculated as: $m_i = \left(n_{a,i} + \theta n_{c,i}\right)^{\alpha}$, where m_i is the equivalent size of household i, θ measures the degree of need of children relative to adults; α specifies economies of scale in consumption; $n_{a,i}$ is the number of adults in household i and $n_{c,i}$ is the number of children. This scale explicitly allows for adjustment of need between adults and children, and for economies of scale with increases in need-adjusted household size. A further benefit of the approach is that its explicit nature easily enables sensitivity analysis. See Creedy and Sleeman (2005) for a detailed comparison of different equivalence scales.

[12] This subsection draws heavily on Thomas (2020).

measure of the redistributive effect of a tax, rather than progressivity, where the redistributive effect is defined in terms of a change in a measure of inequality. Progressivity, on the other hand, to be consistent with the increasing average tax rate concept, is defined in terms of a divergence in the tax burden from proportionality.

The most commonly used global indicator of effective progressivity of a tax (or expenditure program) is that proposed by Kakwani (1977).[13] For PIT, the Kakwani index can be calculated as the difference between the tax concentration coefficient calculated with individuals ranked by (equivalized) pre-tax income and the Gini coefficient for (equivalized) pre-tax income.[14] Gini and concentration coefficients are both measures of dispersion from equality across a cumulative frequency distribution.[15] As such, the Kakwani index measures how much further from equality is the distribution of tax paid than the distribution of pre-tax income (without changing the ranking of individuals). This can be seen graphically in Figure 1 as the area between the Lorenz and concentration curves. It can range from -1 to 1; with a positive figure reflecting progressivity and a negative figure reflecting regressivity.

Figure 1: Lorenz and concentration curves for calculating the Kakwani index

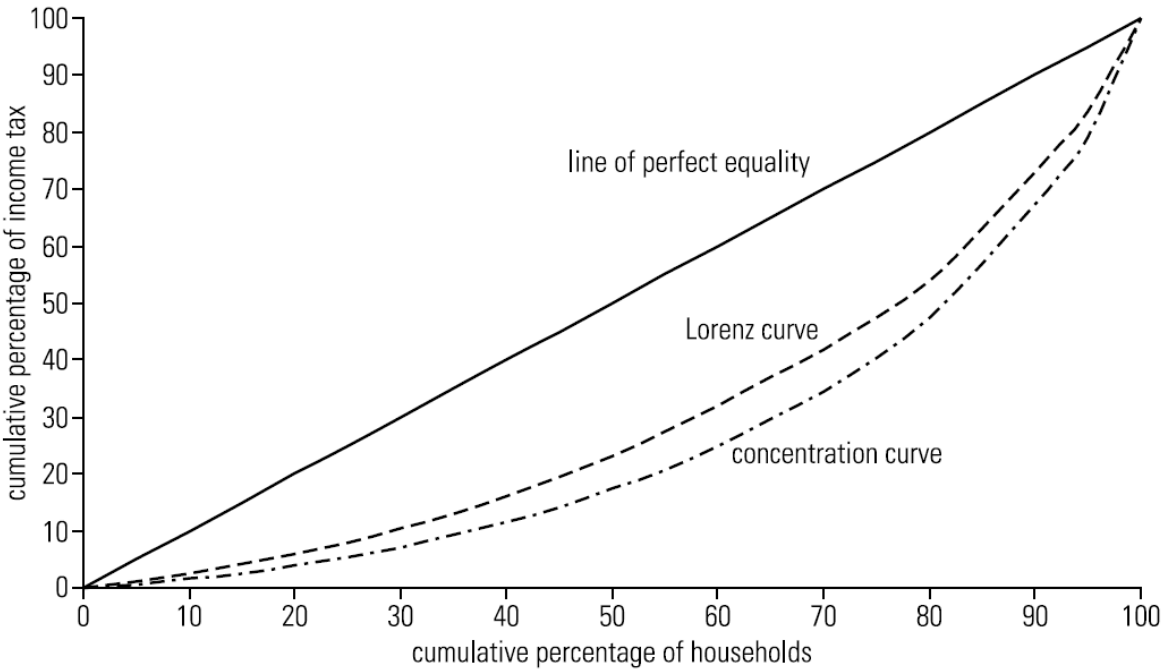

Source: Haughton and Khandker (2009), *Handbook on Poverty and Inequality*, World Bank.

[13] Another popular – and very similar – index is the Suits index. See Suits (1977).

[14] While the Kakwani index was developed to assess the income tax (as well as expenditure programs), it can also be adapted to examine the VAT (see Thomas, 2022).

[15] The Gini coefficient is based on the Lorenz curve, plotted in Figure 1, which is a cumulative frequency curve of income across households. In Figure 1, the Gini coefficient can be defined as two times the area between the Lorenz curve and the line of perfect equality. Likewise, a concentration coefficient can be defined as two times the area between the concentration curve (a cumulative frequency curve of a particular variable, in this case tax paid) and the line of perfect equality. See Haughton and Khandker (2009), p104 for further detail.

Mathematically, the Kakwani index (π^K) can be expressed as follows:

$$\pi^K = \ C_T^Y \ - G_Y$$

or:

$$\pi^K = 2\int_0^1 [L_Y(p) - L_T^Y(p)]\, dp$$

where C_T^Y is the concentration coefficient for tax liability (with individuals ranked by pre-tax income); G_Y is the Gini coefficient for pre-tax income; $L_Y(p)$ is the Lorenz curve for pre-tax income and $L_T^Y(p)$ is the concentration curve for tax liability (with individuals ranked by pre-tax income).

A related indicator commonly used (and close to that initially proposed by Musgrave and Thin, 1948) is the Reynolds-Smolensky (1977) index, which provides a measure of the overall redistributive effect of a tax. For PIT, it is calculated as the difference between the Gini coefficient for (equivalized) pre-tax income and the concentration coefficient for (equivalized) post-tax income, ranked by pre-tax income. As such, the Reynolds-Smolensky index measures how much closer to equality is post-tax income than pre-tax income (without changing the ranking of individuals) – i.e., the reduction in inequality due to the tax. It can be expressed as follows:

$$\pi^{RS} = G_Y - \ C_{Y-T}^Y$$

or:

$$\pi^{RS} = 2\int_0^1 [L_{Y-T}^Y(p) - L_Y(p)]\, dp$$

where C_{Y-T}^Y is the concentration coefficient for after-tax income (with individuals ranked by pre-tax income) and $L_{Y-T}^Y(p)$ is the concentration curve for after-tax income (with individuals ranked by pre-tax income). While the ranking is unlikely to differ for income tax calculations, it may differ for transfers which may be targeted in a different manner to that reflected in the equivalization scale, or for VAT where tax liability depends on the particular consumption patterns of households rather than on total income or consumption.[16]

The Kakwani and Reynolds-Smolensky indices are linked. The overall redistributive effect measured by the Reynolds-Smolensky index can be broken down into two components, a progressivity component measured by the Kakwani index, and an average tax rate component, as follows:

$$\pi^{RS} = \frac{t}{1-t}\pi^K$$

where t is the aggregate average tax rate. This relationship highlights the fact that redistribution can be achieved even by a tax system with only a small degree of progressivity if the average tax paid is high.

[16] An alternative measure of the redistributive effect simply compares the Gini coefficient on pre-tax income with the Gini coefficient on after-tax income. Unlike the Reynolds-Smolensky (RS) index, ranking of individuals may now differ between the two component indices. The difference between this redistributive effect (RE) index and the RS index reflects a ("Atkinson-Plotnick") re-ranking effect. The RE index can be expressed as: $\pi^{RE} = G_Y - \ G_{Y-T}$, or: $\pi^{RE} = 2\int_0^1 [L_{Y-T}(p) - L_Y(p)]\, dp$, where G_{Y-T} is the Gini coefficient for after-tax income and $L_{Y-T}(p)$ is the Lorenz curve for after-tax income.

Equally, a tax system with low tax rates requires a highly progressive system to achieve the same degree of redistribution.[17]

2.3 Examples

This section presents examples of effective progressivity measures from the recent empirical literature. It focuses on PIT and VAT.

2.3.1 Income taxes

There is a significant literature examining the distributional impacts of the PIT (see, e.g., for OECD countries, Joumard et al., 2012; Causa and Hermansen, 2017). However, given the significant data requirements for calculating effective progressivity metrics, results for IDA countries are sparce. For example, results are available from detailed CEQ distributional analyses for more than 40 developing countries, but for only six IDA countries. The CEQ analyses emphasize the importance of examining the distributional impact of tax and expenditure systems together, particularly given the interactions between different fiscal instruments.[18] Nevertheless, results focusing on direct taxes can be drawn out from the CEQ's database (Commitment to Equity Institute Data Center on Fiscal Redistribution, 2022). Results for the six available IDA countries are presented in Table 1. Direct taxes include the PIT together with SSC but excluding pension contributions (which are considered compulsory savings).[19] The results for IDA countries can be expected to predominantly reflect the PIT due to the minimal use of social contributions in IDA countries.

In addition to global metrics, it is quite common to present a visual representation of the distribution of tax burdens by plotting average tax rates across deciles or quintiles.[20] Table 1 presents both global indicators and graphical results for each country. To provide context to the IDA country results, Table 1 also presents results available for the United States. Results all relate to the particular year of the underlying microdata used in the analysis. Note that income survey microdata was available in only two of the six IDA countries (Honduras and Nicaragua).[21] In the other four countries, consumption data was used to proxy for income.[22]

[17] A characteristic of the Gini-based dispersion measures on which the Kakwani and Reynolds-Smolensky indices are based is that they are less responsive to changes at the tails as compared to the middle of the distribution being examined. An alternative approach to measure redistribution is to calculate the difference in the Atkinson inequality index before and after tax. The Atkinson index enables greater emphasis to be applied to the bottom of the income distribution by specifying the degree of inequality aversion in its underlying social welfare function.

[18] For example, Lambert's conundrum illustrates that, due to path dependency, the combined impact of a regressive tax and a progressive transfer can lead to lower inequality than without the regressive tax (Lambert, 2001).

[19] See Lustig (2018) for more detail on the CEQ methodology.

[20] These remain *effective* progressivity measures as they are the averages of the modelled tax burdens of the households in each decile or quintile.

[21] Income survey data is now also available for Ethiopia and will be used in distributional analysis to be undertaken by the World Bank's Fiscal Policy and Sustainable Growth Unit in the second half of 2022.

[22] Because richer households tend to save more than poorer households, using consumption data can be expected to underestimate the income levels of high-income households and the resulting metrics may therefore underestimate the total progressivity achieved in the country. To a smaller extent this can also occur for the countries based on income survey data due to the underreporting of top earners in income survey data.

Table 1: Effective direct tax progressivity and redistribution in IDA countries

	RS	Kakwani	ATR	Income deciles
Ethiopia (2010)	**0.017**	**0.285**	**0.056**	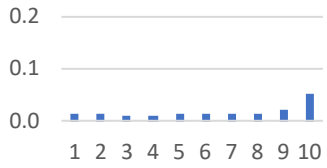
Ghana (2012)	**0.013**	**0.294**	**0.043**	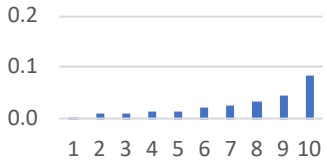
Honduras (2011)	**0.006**	**0.137**	**0.044**	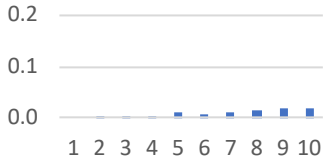
Nicaragua (2009)	**0.020**	**0.311**	**0.059**	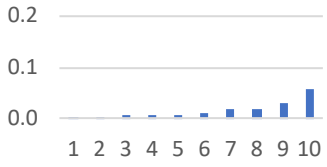
Tanzania (2011)	**0.028**	**0.520**	**0.050**	
Uganda (2012)	**0.012**	**0.471**	**0.025**	
United States (2016)	**0.065**	**0.136**	**0.325**	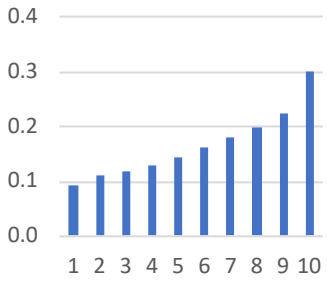

Source: authors calculations based on CEQ Data Center on Fiscal Redistribution (2022)

Looking first at the graphical results, direct taxes appear progressive in each country as average tax rates increase at higher income deciles, though to differing degrees.[23] That said, Honduras appears arguably only slightly progressive as even the top deciles pay very little tax. More generally, the graphs highlight clearly two points about the PIT in many developing countries: the PIT is paid predominantly by high income households, and the overall levels of taxation are low. For example, in Uganda, only the top three income deciles pay any significant amount of tax. The overall low levels of taxation are evident when compared to the United States, where average tax rates are significantly higher right across the income distribution, and even lower decile household pay significant levels of tax.

While the graphs provide a broad indication that direct taxes in these countries are progressive, the Kakwani index enables the degree of progressivity in each country to be assessed and compared. Among the IDA countries, Tanzania and Uganda are shown to have the highest degrees of progressivity, driven by the substantially higher tax burdens on the top decile and minimal tax burdens elsewhere. The lowest degree of progressivity amongst IDA countries is in Honduras where the top six deciles pay very low levels of tax, and the bottom four deciles pay virtually no tax. Interestingly, though, the degree of progressivity in Honduras is still slightly higher than in the United States.

The Reynolds-Smolensky index results provide further insight. While direct taxes in the United States exhibit less progressivity than in the IDA countries, they achieve significantly more redistribution. The reason for this is shown clearly by the average tax rate in the United States which is vastly greater than in any IDA country. In contrast, Uganda, which has the second most progressive tax system, achieves the lowest degree of redistribution because the average tax rate is extremely low.

More generally, the decomposition of the redistributive effect into its progressivity and average tax rate components shows that a highly progressive system on its own is not sufficient to achieve a significant reduction in inequality. It also illustrates that countries can take different approaches to achieve the same degree of redistribution. Consider, for example, Ghana and Uganda which both achieve almost the same degree of redistribution. Uganda achieves this through a highly progressive system that targets only the highest income households, whereas Ghana achieves the same redistribution through a less progressive system applying positive tax rates across most of the income distribution.

2.3.2 VAT

Effective progressivity studies can also be undertaken for the VAT if detailed household budget survey (HBS) microdata is available covering a wide range of consumption items (so as to enable modelling of multi-rate VAT structures). While VAT studies have been undertaken for many high-income countries (see, e.g., Thomas, 2022 and 2020, for 27 OECD countries; IFS, 2011, for 9 European countries), as with direct taxes, the data requirements mean that studies for IDA countries are not common.

Furthermore, a range of conceptual modelling difficulties are associated with modelling VAT burdens. In particular, many studies of the distributional impact of the VAT have examined the VAT burdens across

[23] A close examination of the results shows some local regressivity in Ethiopia (1st-2nd decile) and Honduras (5th-6th decile). This could be due to the provision of direct transfers incorporated in the CEQ analysis, but potentially also due to measurement or sampling error. While the overall trend in these countries still appears progressive, this also highlights the usefulness of global metrics in providing a framework to weight the regressivity in one party of the distribution against progressivity elsewhere to generate an overall measure. This becomes particularly relevant for analysis of the VAT where effective tax burdens can fluctuate non-monotonically across the income and expenditure distribution.

the income distribution and conclude that the VAT is highly regressive. However, this ignores the impact of savings behavior and the consequent need to consider the VAT in a lifetime or at least multi-period context. If instead VAT burdens are measured relative to their actual base – consumption – then they are typically found to be roughly proportional or even slightly progressive (see Annex for further detail). Additionally, recent evidence suggests that informality is likely to result in the VAT being slightly progressive in some developing countries (Bachas et al., 2020).

Table 2 presents expenditure-based effective VAT progressivity results for 26 OECD countries from Thomas (2020). As in Table 1, Kakwani, Reynolds-Smolensky and average tax rate results are presented. Note, though, that average tax rates are presented as *tax exclusive*, rather than *tax inclusive*, rates – thereby enabling clearer comparison with statutory standard VAT rates which are typically specified on a *tax exclusive* basis.

Kakwani index results in Table 2 are all very low, with even the highest result (Italy at 0.0822) significantly lower than the smallest direct tax result in Table 1 (United States at 0.136). In 22 of 26 countries the results are positive, showing the VAT to have a very small progressive effect. This slight progressivity is driven by the presence of reduced or zero VAT rates on products that typically make up a greater proportion of the expenditure of poorer households than richer households, particularly food products purchased for home consumption.

In four countries (Chile, Hungary, Latvia and New Zealand) the Kakwani index results are slightly negative, showing the VAT to have a very small regressive effect. These slightly regressive results provide two interesting insights: first, low spending households in these countries do not benefit significantly from reduced VAT rates. On close consideration, this is not surprising: Chile and New Zealand both have very few, if any, reduced rates in comparison to the majority of countries covered in the study. Hungary and Latvia also have relatively few reduced rates, and importantly the vast majority of food products are subject to the standard rate.[24] Second, higher spending households in these – and presumably other – countries spend a greater proportion of their total expenditure on items that are either untaxed or exempt from tax (for example, financial services, international air travel).

Turning to the Reynolds-Smolensky results, it follows logically that the 22 countries with a positive Kakwani index, also have a positive Reynolds-Smolensky redistribution index – showing that inequality has fallen in these countries. Conversely, the four countries with a negative Kakwani index, also have a negative Reynolds-Smolensky index – indicating an increase in inequality. However, in all cases the redistributive effect is very low, ranging from 0.0082 in Belgium to -0.0025 in New Zealand. Redistribution is low despite average tax rates typically being above 11%, further emphasizing the very low degree of progressivity (or regressivity) in the VAT. This highlights that even in countries with the most extensive use of reduced VAT rates and exemptions, the VAT can be expected to have very little impact on redistribution.

[24] Note that while Latvia applied the standard VAT rate to all food products except baby food in 2010 (the year modelled), it now also applies a reduced VAT rate to some fresh fruit and vegetables. Hungary now also applies a reduced VAT rate to a wider range of basic food products than in 2010.

Table 2: Effective VAT progressivity and redistribution in OECD countries

	Reynolds-Smolensky	Kakwani	$\frac{t}{1-t}$
AUT	0.0056	0.0428	0.1300
BEL	0.0082	0.0738	0.1117
CHE	0.0020	0.0419	0.0469
CHL	-0.0009	-0.0080	0.1099
CZE	0.0018	0.0114	0.1572
DEU	0.0049	0.0477	0.1017
DNK	0.0064	0.0382	0.1665
ESP	0.0032	0.0401	0.0799
EST	0.0024	0.0153	0.1558
FIN	0.0043	0.0346	0.1237
FRA	0.0040	0.0457	0.0867
GBR	0.0029	0.0277	0.1059
GRC	0.0040	0.0394	0.1010
HUN	-0.0010	-0.0059	0.1651
IRL	0.0019	0.0215	0.0876
ITA	0.0063	0.0822	0.0766
LUX	0.0031	0.0464	0.0667
LVA	-0.0001	-0.0007	0.1550
NLD	0.0054	0.0561	0.0961
NZL	-0.0025	-0.0188	0.1301
POL	0.0069	0.0562	0.1228
PRT	0.0039	0.0378	0.1026
SLV	0.0053	0.0478	0.1103
SVK	0.0010	0.0067	0.1520
SWE	0.0074	0.0631	0.1175
TUR	0.0010	0.0106	0.0988
Average	**0.0034**	**0.0328**	**0.1138**

Source: Thomas (2020)

2.3.3 Other taxes

Effective progressivity studies can also be undertaken for excise taxes (including health- and energy-related taxes) where microdata is available on consumption by different households across products subject to excise taxation. Excise taxes are often found to be regressive, although this is not always the case (e.g., Flues and Thomas, 2015, find transport fuels in some OECD countries to have some progressive impact). Furthermore, beyond the pure fiscal impact, health- and energy-related excise taxes, through eliciting behavioral changes, can have broader progressive impacts (see, e.g., Fuchs et al., 2019); Heine and Black, 2019; Sassi et al., 2018).

2.4 Pros and cons of effective progressivity metrics

Effective progressivity metrics have significant advantages. Most obviously, effective progressivity considers the interaction of both the tax schedule and the underlying income distribution. Effective progressivity measures can be used to examine not just the PIT and social contributions, but also the VAT and excise taxes. The key drawback of the effective progressivity approach is that it is data intensive. It requires information on both the tax system and on the underlying distribution of income. The difficulty in obtaining data for IDA countries was illustrated by the fact that even the wide ranging CEQ analyses cover only six IDA countries, and only two based on income survey data. Furthermore, any microdata-based analysis is resource intensive.

In addition to the difficulty obtaining data, there are also a wide range of data limitations that impact effective progressivity analysis. In particular, top incomes are typically underreported in income surveys. Evasion, particularly by high-income taxpayers, will also not be captured in income survey data. If consumption data is used for PIT modelling, it will provide only an imperfect proxy for income data due to savings patterns. Additional problems arise when modelling indirect taxes (beyond the conceptual difficulties noted above). For example, expenditure is typically underreported in household budget surveys, due to a range of factors including imperfect recall and intentional underreporting of some expenditure categories. VAT exemptions are difficult to model, and ideally input-output data is needed to estimate the impact of the denial of input tax credits. The impact of informality can also pose modelling challenges in developing countries, particularly if a place of purchase variable is not available. There are also challenges associated with modelling *ad quantum* excise taxes with expenditure data.

3 Measuring structural tax progressivity

Structural progressivity focuses purely on the income tax rate schedule itself, ignoring the underlying distribution of income, and produces local measures of progressivity. This section first presents a range of local measures of structural progressivity before presenting examples of their use, and finally summarizing their strengths and weaknesses.

3.1 Structural measures

A simple measure of structural progressivity linked to the basic increasing average tax rate definition, is "average rate progression", which measures the "rate of change in the average rate of tax" (Musgrave and Thin, 1948, p499). This can be defined with respect to a discrete income range as:

$$\text{Average rate progression} = \frac{\frac{T_1}{Y_1} - \frac{T_0}{Y_0}}{Y_1 - Y_0} = \frac{A_1 - A_0}{Y_1 - Y_0}$$

where T_1 is tax paid on income Y_1, and T_0 is tax paid on income Y_0 (with $Y_1 > Y_0$), while A_1 and A_0 are the average tax rates for incomes Y_1 and Y_0. While the name derives from Musgrave and Thin (1948), this structural progressivity measure has been proposed by various authors including Pigou (1928) and Slitor (1948). For a progressive tax, the average rate progression coefficient will be positive, while it will be equal to zero for a proportional tax, and negative for a regressive tax. Musgrave and Thin (1948) also show that average rate progression can be written as:

$$\text{Average rate progression} = \frac{1}{Y_1}(M_{1-0} - A_0)$$

where M_{1-0} is the tax rate on marginal income, $Y_1 - Y_0$. This shows that for a tax to be locally progressive, the marginal rate must be higher than the average tax rate.

In addition to average rate progression, Musgrave and Musgrave (1989) propose two further measures of structural progressivity (originally proposed by Musgrave and Thin, 1948[25]): liability progression and residual income progression. Liability progression is defined as the elasticity of the tax liability with respect to pre-tax income (i.e., the ratio of the percentage change in tax liability to the percentage change in pre-tax income):

$$\text{Liability progression} = \frac{T_1 - T_0}{T_0} \cdot \frac{Y_0}{Y_1 - Y_0}$$

Musgrave and Thin (1948) also show that average rate progression can be written as the ratio of marginal to average tax rates:

$$\text{Liability progression} = \frac{M_{1-0}}{A_0}$$

A progressive tax, with marginal rate higher than average rate, will have a liability progression coefficient greater than 1, while a proportional tax will have a coefficient equal to 1, and a regressive tax will have a coefficient less than 1.

Residual income progression is defined as the elasticity of after-tax income with respect to pre-tax income (i.e., the ratio of the percentage change in after-tax income to the percentage change in pre-tax income):

$$\text{Residual income progression} = \frac{(Y_1 - T_1) - (Y_0 - T_0)}{Y_0 - T_0} \cdot \frac{Y_0}{Y_1 - Y_0}$$

Musgrave and Thin (1948) show that residual income progression can be written as the ratio of the marginal residual income to average residual income:

$$\text{Residual income progression} = \frac{1 - M_{1-0}}{1 - A_0}$$

Unlike liability progression, a progressive tax will have a residual income progression coefficient less than 1, a proportional tax will have a coefficient again equal to 1, while a regressive tax will have a coefficient greater than 1.

3.2 Examples

While examples of structural progressivity metrics are not readily available for IDA countries, OECD (2014) presents a relatively recent example of a structural progressivity indicator for OECD countries. OECD (2014) presents an average rate progression metric for both the PIT and for the tax wedge. The tax wedge is a broader measure of the labor tax burden incorporating both PIT and SSC. It is defined as: (PIT + employee SSC + employer SSC) / (gross income plus employer SSC). Figure 2 presents both these metrics for four different family types. In each case, the metric is calculated over five different income intervals

[25] A further indicator proposed by Musgrave and Thin (1948), marginal rate progression, is not detailed here.

based on percentages of the average wage in each country: 50%-67%; 67%-100%; 100%-133%; 133%-167%; and 167%-200% of the average wage. The only information necessary for the calculations is the average wage and the underlying PIT and SSC rules in each country.

Figure 2 shows that the average rate progression for both the PIT and the tax wedge declines as income increases. This decreasing pattern holds for each family type. For families without children, at higher income levels structural tax wedge progressivity is below structural PIT progressivity. This is because SSC are typically levied at flat rates and in some cases apply contribution ceilings, thereby reducing overall progressivity. The pattern is quite different for families with children. Here the average tax wedge progression at the lowest two income intervals is significantly higher than the average PIT progression at those same intervals. This is due to concessionary SSC provisions being targeted at low-income families, with the average tax wedge increasing significantly as these family-based concessions are withdrawn.

Figure 2: Average rate progression, on average, across OECD countries, 2011

Source: OECD (2014)

3.3 Pros and cons of structural progressivity metrics

The key benefit of structural progressivity metrics is that they are easy to compute and do not require income distribution data to compute. The above results only required average wage figures and the underlying PIT and SSC rules in each country. For a single country analysis, the average wage would not even be necessary as calculations could simply be made across income levels defined in local currency terms.

However, the drawbacks are also obvious. Local rather than global metrics are produced, preventing an overall assessment of the degree of progressivity of a tax system. Meanwhile, as they entirely ignore the underlying income distribution, the impact of the displayed progressivity on the actual tax paid by households is unclear. They also provide no information on capital income taxes if taxed at different rates to labor income. They also cannot account for tax avoidance and evasion, although this is an equal criticism of other metrics. Finally, they are limited to analyzing PIT and SSC only.

4 Measuring progressive capacity

Progressive capacity measures are a type of hybrid of the effective progressivity and structural progressivity approaches, producing global metrics but attempting to abstract away from the underlying pre-tax income distribution to focus on the tax structure of the PIT. This section first summarizes the recent literature proposing approaches to measure progressive capacity, before presenting examples of their use, and finally summarizing the strengths and weaknesses of the progressive capacity approach.

4.1 Approaches to measuring progressive capacity of the PIT

The recent "progressive capacity" literature has been particularly motivated by the desire to undertake cross-country comparisons of the progressivity of PIT systems. Because the interest is in the structure of the PIT, rather than its actual redistributive impact, the approach aims to control for differences in pre-tax income distributions and thereby isolate the impact of the tax structure.

The first approach developed to control for differences in pre-tax income distributions was the "fixed income" approach which simply uses an identical pre-tax income distribution for each country in order to provide an identical base for the calculations. The tax rules of each country are then applied to the 'base' distribution and Kakwani and Reynolds-Smolensky indices, or similar global metrics, can then be calculated to assess progressive capacity of the PIT and to compare this across countries. Examples of this approach include Bishop et al. (1990), who compared results for four countries using each of the different distributions as the common base. Norregaard (1990) also adopts such an approach, applying the German income distribution to several other OECD countries to calculate progressivity measures. Meanwhile, Kasten et al. (1994) apply the United States' income distribution to examine progressivity over several subsequent years in the United States.

A limitation of this approach is that the application of a different income distribution may change the progressivity ranking of countries in the comparison as it may place greater weighting on different tax rates where population density is greater. To address this concern, Dardanoni and Lambert (2002) proposed a "transplant and compare" approach, which applies a transformation function to 'transplant' distributions into a common base. They show that, if an iso-elastic transformation can be found for each country, then rankings based on Kakwani and Reynolds-Smolensky indices (or any Lorenz curve-based metric) will hold for any pre-tax income distribution. Unfortunately, the requirement of an iso-elastic transformation is stringent, and will not necessarily be met in practice.

Additionally, the data requirements for the "transplant and compare" approach also remain significant as microdata is still required for each country. Vellutini and Benítez (2021) apply a simplified "transplant and compare" approach, assuming a lognormal distribution for each country, calibrated using just the Gini coefficient and average income (thereby avoiding the need for the underlying income distribution microdata). They find similar, though not identical, country rankings when comparing calculations using their simplified approach and the full microdata approach for a selection of countries where they have microdata. Nevertheless, the approach is still limited by the assumption that the underlying country income distributions are lognormal and that an isoelastic transformation can be applied.

Another simplified approach to produce a global progressivity metric without income distribution microdata has recently been proposed by Gerber et al. (2020). They focus their analysis on an income range from 1% to 500% of GDP per capita in each country and calculate average tax rates at each of these

500 points for OECD countries. They then use these 500 income points and corresponding average tax rate calculations to calculate Kakwani indices for each country. This approach consequently needs no data on the underlying income distribution beyond a figure for GDP per capita. The approach is more in the nature of the "fixed income" approaches noted above, as it implicitly applies a uniform distribution for each country, that is calibrated by GDP per capita.

While still potentially susceptible to the same ranking issues as other fixed income approaches, the Gerber et al. (2020) approach provides a simpler alternative to the "transplant and compare" approach while also limiting data requirements. It may, therefore, be an attractive option for calculating progressive capacity in developing countries and for comparing results both across time and countries. However, by capping the income distribution at 500% of GDP per capita, it risks ignoring the top marginal tax rate (or rates) in some countries. While their study was restricted to OECD countries, this may be a more significant issue in developing countries with lower GDP per capita and where PIT systems focus more on the top end of the income distribution.

An alternative hybrid approach that does not rely on standard global progressivity metrics was taken by Sabrianova Peter et al. (2010). They generate (global) marginal and average rate progression indices by first calculating marginal and average tax rates at 100 different levels of pre-tax income (between 4-400 percent, and 100-300 percent, of GDP per capita) and then regressing these marginal and average tax rates on the log of gross income. The PIT is interpreted as progressive if the slope of the regression is positive. They do these calculations for 189 countries and over the 1981-2005 period.

OECD (2014) also effectively calculate a type of global metric in their analysis (previously discussed in section 4), by presenting for each country the simple average of their average rate progression results across five different income intervals.

4.2 Examples

We focus here on results from Vellutini and Benítez (2021) and Gerber et al. (2020). Figure 3 presents results from Vellutini and Benítez (2021), where the averages for three country groups are presented over a number of years. Reynolds-Smolensky index results ("redistributive capacity") results are presented, as well as their decomposition into the Kakwani index ("progressive capacity") and average tax rate ("aggregate tax rate") components.

Figure 3 shows differing trends over time for low, middle and high-income countries. Progressive capacity is shown to have increased over time in low-income countries, to have declined in middle income countries, while staying comparatively low in high-income countries. Average tax rates, on the other hand, have fallen in low-income countries and increased in middle income countries. They have remained comparatively high in high-income countries, driving the overall significantly higher redistributive impact. The different trends for low and middle-income countries have roughly canceled each other out leading to similar redistributive capacities for low- and middle-income countries, on average, over time. The higher redistributive capacity and lower progressive capacity found here in high-income countries, broadly mirrors the higher redistributive effect and lower effective progressivity results observed for a single year for the United States when comparing with six IDA countries in section 3.

Figure 3: PIT progressive and redistributive capacities across 108 countries

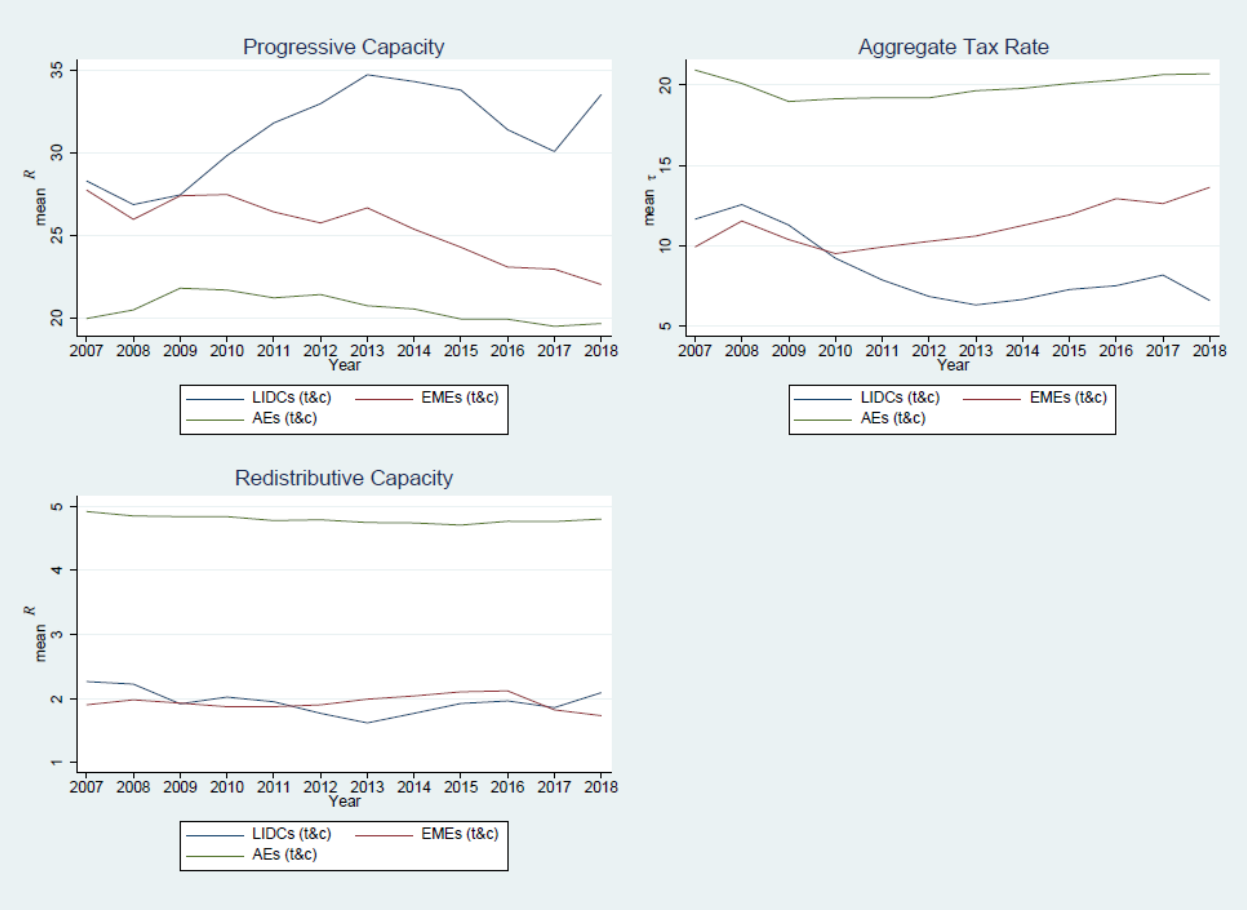

Source: Vellutini and Benítez (2021)

Figure 4 presents results from Gerber et al. (2020), who show their Kakwani index ("Average Progressive Capacity") results in addition to the global progressivity metric results of Sabrianova Peter et al. (2010) ("Average Rate Progression", and "Average Rate Progression, OECD sample"). Results are presented over a wide time span from 1981-2017 for the average progressive capacity results, which are averages across 15-35 OECD countries depending on the year. The average rate progression results cover 1981-2005, and are averages across either 161 countries, or just the OECD countries in their sample.

The two sets of results focusing on OECD countries show a strong declining trend in the progressive capacity of PIT systems in OECD countries from the early 1980s through to around the mid-2010s. Subsequently, progressive capacity in OECD countries appears to have been relatively stable at the new lower levels. The decline in progressive capacity can also be observed in the wider 161 country results from Sabrianova Peter et al. (2010), though to a far smaller degree. Gerber et al. (2020) note that this reduction in progressive capacity is consistent with the downward trend in top PIT rates across the world throughout the same period (up until around the last decade when top PIT rates in high income countries began to rise very slightly).

Figure 4: Progressive capacity across OECD countries and average rate progression metric results across 161 countries

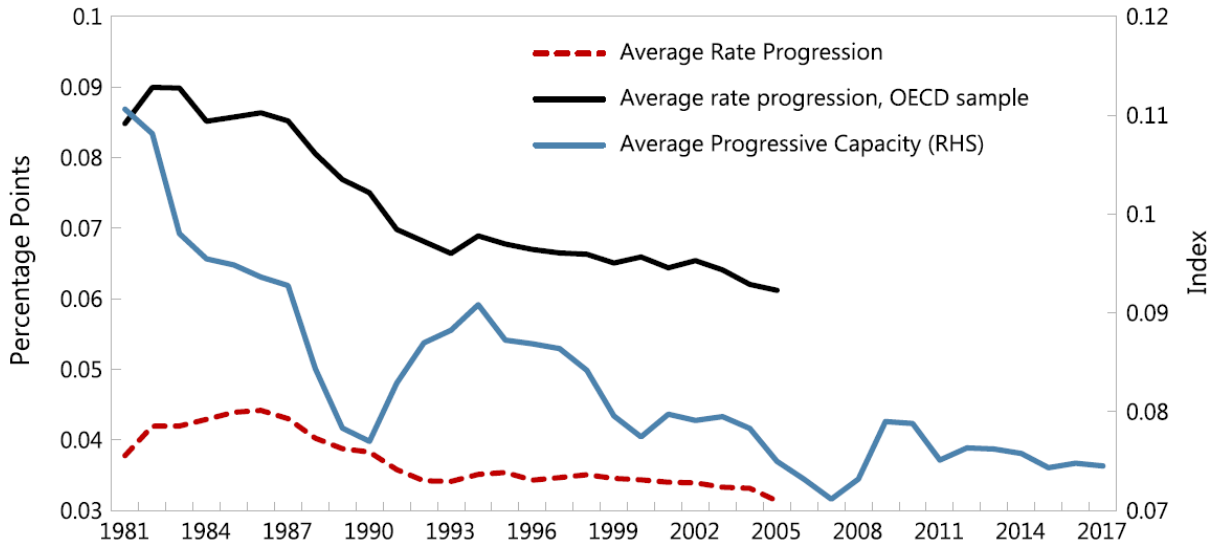

Source: Gerber et al. (2020)

4.3 Pros and cons of progressive capacity metrics

The case for using progressive capacity metrics is strongly linked to the question being asked. If the researcher is interested in the progressive impact achieved by a tax system, then it will clearly be preferable to use effective progressivity metrics. However, if the researcher is interested in measuring and comparing across countries and/or across time the degree of progressivity of the tax schedule itself, there is significant merit in adopting a progressive capacity metric. However, a drawback of this approach is that the progressive capacity metrics, and consequent country rankings, are not necessarily invariant to the choice of 'base' income distribution, and the recently developed methods to address this base sensitivity problem rely on strong assumptions.

As with other metrics, data limitations also impact the implementation of the approach, though the simplified approaches discussed above alleviate data requirements significantly. As with other approaches, the impact of tax avoidance and evasion cannot be captured, nor can the impact of capital income taxes where they are taxed at a different rate to labor income. Finally, the approach is limited to the examination of the PIT (and potentially SSC).

5 Using aggregate data

In light of the limitations of the above approaches, particularly regarding data and the difficulty undertaking analysis beyond the PIT, this section illustrates how aggregate tax revenue data can be used to help assess the progressivity of tax systems in IDA countries.

While aggregate data cannot be used to explicitly differentiate the tax paid by poorer and richer households, it can illustrate the degree of reliance a country places on taxes that typically exhibit progressivity. This may be particularly useful in countries where disaggregated data is not available, as is the case in many IDA countries. Nevertheless, while providing a useful high-level picture, aggregate tax

data remains a blunt instrument for examining progressivity. Indeed, the limitations exhibited in this section further illustrate the need for obtaining microdata to more precisely measure the progressivity of tax systems where possible.

5.1 Potential aggregate tax revenue-based metrics

The section draws on aggregate tax revenue data from the UNU-WIDER Government Revenue Dataset covering 68 of 74 countries currently eligible to receive IDA resources (59 IDA-only and 15 blend countries).[26] Three potential indicators are presented:

1. The ratio of direct tax revenue to indirect tax revenue;
2. The ratio of income tax (PIT + CIT) revenue to total tax revenue; and
3. The ratio of income tax (PIT + CIT) revenue to GDP.

The ratio of direct-to-indirect tax revenue is a commonly used aggregate metric, adopted on the basis that direct taxes can be expected to exhibit more progressivity than indirect taxes. The second and third indicators focus specifically on the PIT and CIT on the basis that these are likely to be the predominant sources of progressivity in tax systems. In the calculations, direct taxes are specified as PIT (including labor and capital income), CIT, and property taxes; while indirect taxes are specified as VAT, sales taxes, excise taxes and trade taxes.[27]

For brevity, a detailed discussion of the tax incidence literature is not provided here. However, income taxes are considered a primary source of progressivity based on the inherently progressive design of most PIT systems, and, in the case of CIT, following the conventional view that capital owners bear the majority of the CIT (and given that capital income is disproportionately earned by high-income households).[28] In contrast, while empirical results can be mixed for indirect taxes, none are typically considered to achieve a significant degree of progressivity.[29]

Property taxes, while included in the numerator for the first indicator, are excluded from the second and third for several reasons. First, because there remains significant debate in the literature regarding the extent to which recurrent property taxes exhibit features of a benefits tax (and hence are closer in nature to a payment for services thereby having no distributional impact) as opposed to a tax. Second, it avoids the inclusion of highly inefficient (but potentially progressive) property transaction taxes from having a

[26] The UNU-WIDER dataset draws on data from several sources, primarily the IMF Government Finance Statistics and OECD Global Revenue Statistics databases. See: https://www.wider.unu.edu/about-grd.

[27] The PIT and CIT categories include "taxes on income, profits and capital gains" of individuals and corporate entities, respectively.

[28] As noted earlier, empirical evidence does suggest that at least some of the CIT may be borne by workers (in the form of lower wages), which would reduce the progressivity of the tax.

[29] While the VAT was historically considered a regressive tax, more recent literature finds the VAT to be roughly proportional on a lifetime basis, and in some cases even sightly progressive (but with minimal redistributive impact) due to the application of reduced rates on consumption items that make up a greater proportion of the expenditure of poorer households such as basic food products (see, e.g., IFS 2011, Thomas, 2022), or due to informality (Bachas et al., 2020). Excise taxes are often found to be regressive, although this is not always the case (e.g., Flues and Thomas, 2015, find transport fuels in some OECD countries to have some progressive impact). Furthermore, beyond the pure fiscal impact, health and environmentally related excise taxes though eliciting behavioral changes can have broader progressive impacts.

positive impact on the indicator.[30] Ideally inheritance and wealth taxes – which are also included in the property tax category – would have been separated and included in the metrics, as they are almost certainly progressive. Unfortunately, this level of disaggregation was not available in the data. Overall, because inheritance and wealth taxes are not typically used to any significant degree in IDA countries, any benefit of including them was considered to be outweighed by the need to then also include property transaction and recurrent property taxes.[31] This exclusion is also relatively minor, as the entire property tax category raises very little revenue in IDA countries (0.12% out of 15.1% of GDP, on average in 59 IDA countries).

Social contributions are excluded entirely from the metrics for three reasons: First, SSC are typically imposed at flat rates so can be expected to have at best a broadly proportional impact, and in some cases a regressive impact (where upper contribution ceilings apply). As such, including them in the numerator of any of the indicators would not aid in assessing reliance on progressive taxes. Second, it would appear inconsistent to exclude SSC from the direct-to-indirect tax ratio (for the reason above), but to then include SSCs in the denominator of the two income tax metrics. Third, the appropriate treatment of the pension component of social contributions raises debate regarding the degree to which contributions reflect a pure tax as compared to compulsory savings. Given that SSC are applied very minimally in IDA countries, excluding SSCs from the metrics was therefore considered a reasonable means of abstracting away from this uncertainty.

The difference between the first and second indicators is relatively subtle, essentially just the exclusion of property taxes, and the tempering impact on the indicator from income taxes also appearing in the denominator. However, from a presentational perspective there may be merit in the second indicator's focus solely on taxes that are considered "good" for progressivity. This potentially contrasts with the negative connotation that may be created with the direct-to-indirect tax ratio that indirect taxes are "bad" whereas direct taxes are "good". Not only is that a simplistic interpretation from a distributional perspective, but it may cloud the fact that there are strong tax policy rationales for raising significant revenue from indirect taxes within a broader inclusive growth strategy.

The third indicator is included as a complement to the second indicator. These two indicators are intended, to an extent, to play analogous roles to the progressivity and redistribution indicators presented in section 3. That is, the income tax-to-total tax ratio gives an indication of the degree of reliance on typically progressive taxes, and so of (potential) progressivity in the tax system; while the income tax-to-GDP ratio gives an indication of the amount of revenue associated with these (potentially) progressive taxes and hence of the degree of redistribution that may be achieved with them.

5.2 Results

Figure 5 presents the three indicators for the most recent year available for 59 IDA/blend countries. For most countries data is available to 2019 or 2020. Nine countries are dropped from the originally available

[30] Transaction taxes may have a progressive impact (although they may be borne most heavily by middle and upper-middle income households rather than top-income households). However, transaction taxes are highly inefficient – distorting both housing and labor markets – and so are very unlikely to be advocated for as a means of achieving progressivity. Excluding them from the indicator avoids any implicit support of such a reform.

[31] Nevertheless, if IDA countries began to implement significant inheritance and wealth tax reforms, then inclusion of the property tax category in this indicator may need to be reconsidered.

68, because either no direct-to-indirect tax revenue breakdown was available, or no income tax revenue data was available, or because data was significantly outdated (prior to 2011).[32]

Additional data quality issues were found with data for some of the countries included. For each country, the data provides a total tax revenue figure, total direct and indirect tax revenue figures, and more detailed breakdowns across individual tax types. However, in a number of cases, the total tax revenue figure exceeds the sum of the sub-components. This may occur, for example, where some revenue was not allocated to a specific category. Perhaps of more concern, in a smaller number of countries, the sum of the components exceeds the total tax revenue figure. This includes, in several countries, the addition of the direct and/or indirect tax subcomponents not equaling the direct and indirect tax aggregates. A judgement call was therefore necessary in calculating the indicators for several countries as to whether to rely on the provided direct/indirect tax figures, or the addition of their subcomponents. To provide consistency across the three indicators, the subcomponents were used rather than the aggregates. This ensures that any differences in the country rankings in Figure 5 are due purely to the inclusion or exclusion of property taxes from the numerator.

Figure 5 clearly shows the strong reliance of IDA countries on indirect as opposed to direct taxes. Only five countries have a direct-to-indirect tax ration greater than 1. The highest of these is Papua New Guinea, where income taxes are the dominant revenue source. In contrast, Vanuatu – which employs no personal or corporate income tax – has the lowest direct-to-indirect tax ratio. While there is also no PIT in the Maldives, the CIT increases the indicator slightly. The greater reliance on indirect taxes in most IDA countries is likely driven to a significant extent by administrative capacity constraints that result in greater emphasis being placed on indirect taxes that may be easier to collect than direct taxes. Additionally, the low income levels of a significant proportion of the population can significantly limit the scope for personal income taxation.

Comparing the direct-to-indirect and income tax-to-total tax ratios, shows the tempering impact that the inclusion of income taxes in the base has. As a result the variation in indicator is lower across countries than for the direct-to-indirect tax ratio. Beyond this, the absence of property tax revenue results in some differences in ranking of countries. For example, Grenada, which raises 1.2% of GDP from property taxes, ranks 40th based on the direct-to-indirect tax ratio, but 47th equal under the income tax-to-total tax ratio.

The benefit of presenting the income tax-to-GDP ratio is also illustrated when looking, for example, at Lesotho, which ranks only moderately according the direct-to-indirect tax and income tax-to-total tax ratios. But the likely redistributive effect that the income taxes have in Lesotho is likely significantly higher than many other countries as it raises significantly more tax revenue from these taxes (as a percentage of GDP). Mozambique, meanwhile, both relies significantly on income taxes and also raises significant revenues implying both significant progressivity and redistribution. Other countries that rely heavily on direct taxes, such as Bhutan, Chad, and the Democratic Republic of Congo, may not achieve significant redistribution as a result of the comparatively low aggregate tax level.

[32] In some cases, total tax revenue was available for 2019 or 2020, but income tax revenue figures were significantly outdated: Republic of Congo (1994), Côte D'Ivoire (2005), Ethiopia (2007), Guinea (2006). In Mali, 2020 data is available for the direct-to indirect tax ratio, but income tax data is only available for 2012, so results are based on 2012 data for all three indicators (the direct-to-indirect tax ratio is very similar for 2012 and 2020).

Figure 5: Aggregate tax revenue-based progressivity indicators

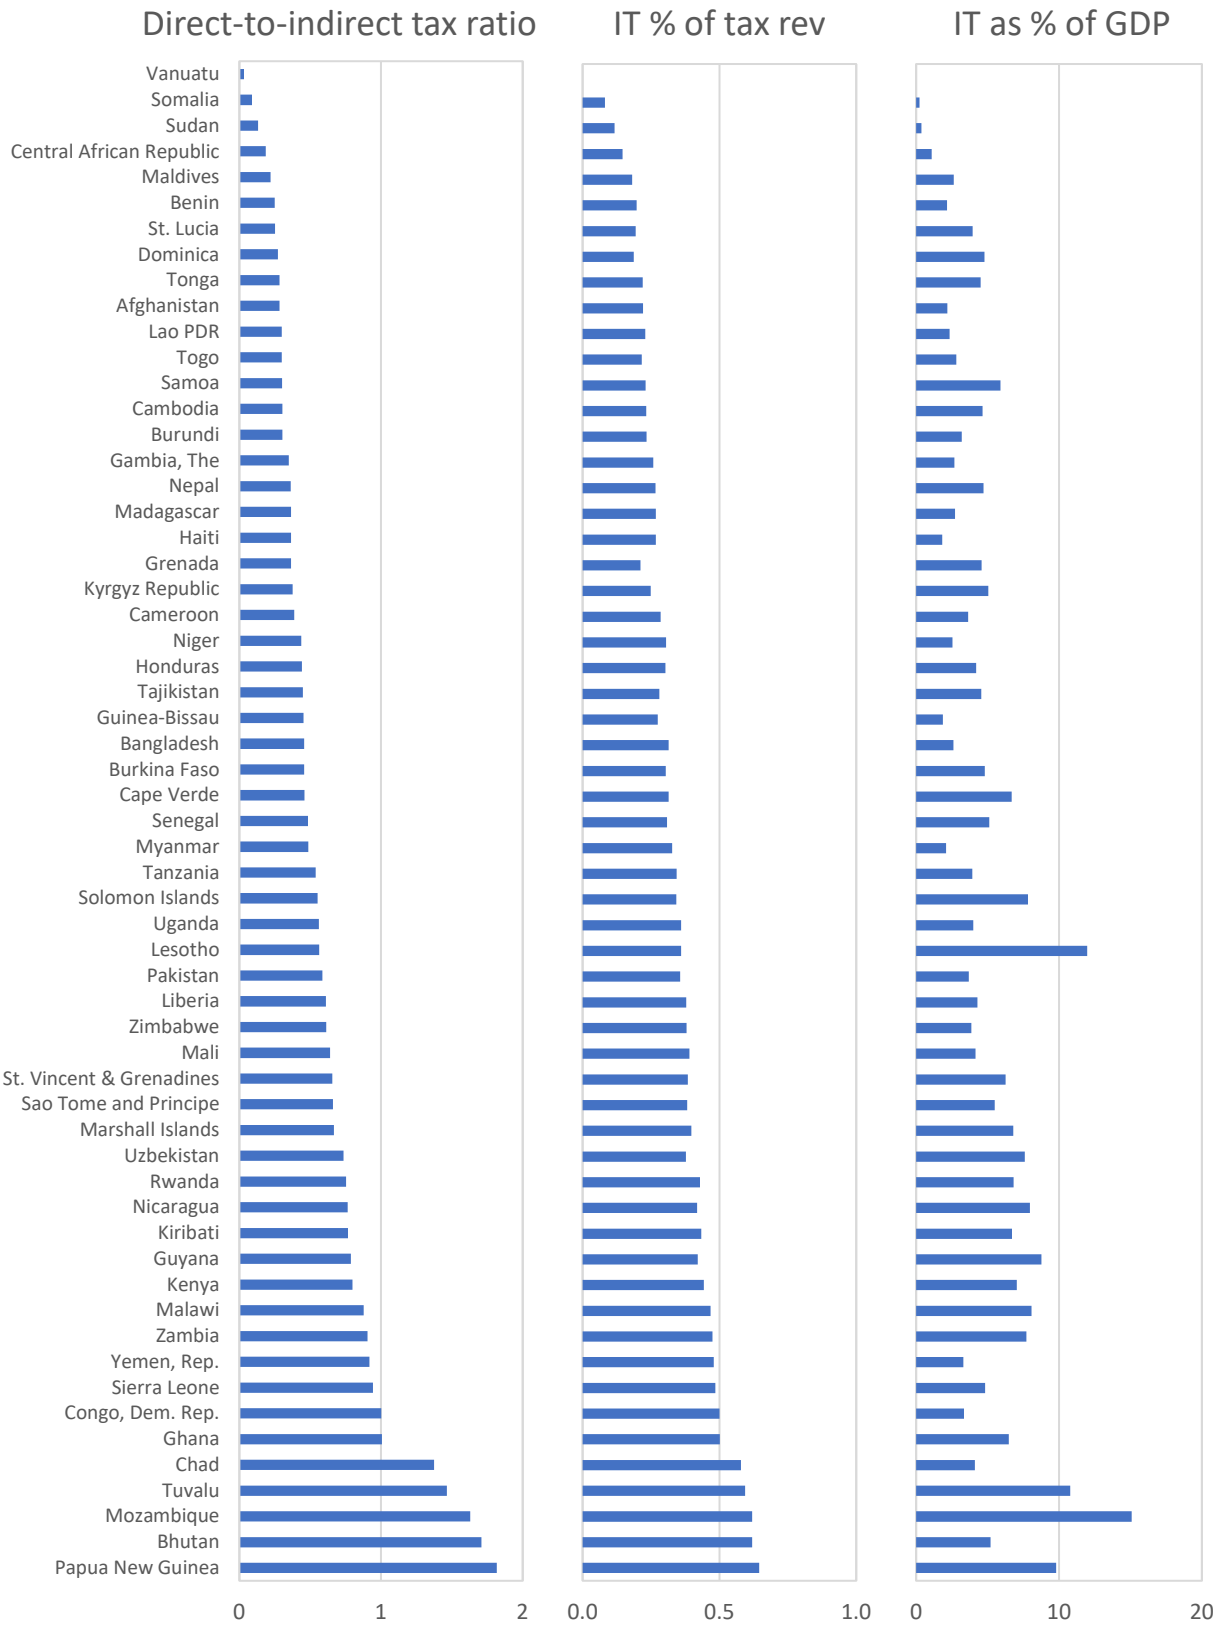

Source: authors calculations based on UNU-WIDER World Revenue Dataset

Nevertheless, the bluntness of these indicators is suggested by reference back to the microdata-based Kakwani index results for personal taxes for IDA countries in section 3. Ghana and Nicaragua appear more progressive than Tanzania and Uganda in Table 1, particularly in terms of the direct-to-indirect tax ratio. In contrast, results from section 3, albeit for different years, suggest that direct personal income taxes in both Tanzania and Uganda are significantly more progressive than in Ghana or Nicaragua (with Kakwani indices of 0.52 and 0.47 vs 0.29 and 0.31, respectively). That said, the low average tax rate effect for Uganda is consistent in both analyses, suggesting low redistributive impact. Meanwhile, Honduras exhibits comparatively low progressivity in both analyses. These comparisons emphasize the merit in presenting a range of metrics where possible for each country.

5.3 Pros and cons of aggregate tax revenue-based metrics

There appear clear benefits to the calculation of these summary metrics, which enable the impact of either all direct taxes, or of personal and corporate taxes (including tax revenue from capital income) to be incorporated. They are available for a large number of IDA countries, at low cost, and in most cases with only a time lag of around two years. However, they remain blunt indicators as they do not distinguish between taxes paid by poor and rich households and instead rely on strong assumptions regarding the incidence of different tax types. For example, if a country adopts a flat rate PIT system, the resulting lack of progressivity would be missed by these aggregate metrics.[33] Furthermore, there are data quality concerns in a number of countries, and while the time lags are generally small, for some countries the lags are large enough for them not to be useful in terms of the desired monitoring exercise for IDA20. If they are to be used, they should be complemented by additional indicators where possible, particularly of the progressivity of the PIT.

6 Measuring tax progressivity in IDA countries

Drawing on the preceding analysis, this section proposes a set of indicators for measuring progressivity in IDA countries over FY23-25. The proposal is guided by the following principles:

- The need to assess progressivity in IDA countries in FY23-25, concurrently, if possible or with minimal lag.
- The desire to assess progressivity of the entire tax system, not just the personal income tax.
- Acknowledgment of data limitations: specifically, the lack of availability of income survey microdata in most IDA countries, and difficulty in accessing consumption microdata.
- Acknowledgment of the limitations of different progressivity metrics.
- Acknowledgment of resource constraints.

As highlighted by the preceding sections, there are strengths and weaknesses with different measures of progressivity and redistribution. It is therefore proposed that a suite of metrics be employed to assess progressivity in IDA countries over FY23-25. Aggregate tax revenue metrics would be relied upon to provide a basic assessment of the progressivity of the overall tax system, while the degree of progressivity of PIT systems would be investigated in further detail utilizing both effective progressivity and progressive capacity metrics. In light of the large number of countries involved, a simplified modelling approach is

[33] Even in flat rate PIT systems, the existence of a basic allowance will still create some progressivity (i.e., the average tax rate will increase with income). However, this will be significantly less than in a typical country applying a progressive tax rate schedule.

proposed to compute effective progressivity and progressive capacity metrics, while managing data and resource constraints.

The aggregate metrics focus on direct taxes, while the more detailed metrics focus on the PIT. Detailed metrics could also have been proposed for the VAT, excise taxes, and for SSC. However, detailed analysis of these taxes is not proposed for several reasons: first, no comparable simplified approach is available to model VAT and excise taxes as there is for the PIT. As such, full microsimulation models would need to be constructed for each country (ideally also using input-output data to accurately model VAT exemptions). In light of empirical evidence that the VAT has minimal overall redistributive impact irrespective of rate structure (Thomas, 2022, 2020; see also Annex), it is considered that the additional insights obtained from detailed microdata-based analysis would not justify the resource cost involved for more than 60 countries. Additional challenges associated with modelling informality, data access, and data quality also suggest this analysis would not be warranted.[34]

While the proposed effective progressivity metrics could be extended to include SSC, this is not considered warranted for two reasons. First, SSC are not utilized to any significant degree in IDA countries. Second, as SSC are typically imposed at flat rates, they tend to have a broadly proportional impact (at least until high income levels where ceilings may be applied creating some regressivity). Third, analysis of SSC, as compared to PIT, is complicated by questions both of incidence (employee or employer) and nature (tax or compulsory savings).

Finally, broad assumptions could be made regarding the likely incidence of the CIT to enable some form of distributional analysis to be undertaken.[35] However, in light of the significant uncertainty in the literature regarding incidence of the CIT, and the at best tenuous link that could be made to microdata to enable such analysis, this is also not considered warranted. Furthermore, the influence of a progressive CIT is captured in the proposed aggregate tax revenue metrics.

The proposed metrics are discussed, in turn, below.

6.1 Aggregate tax revenue-based indicators

It is proposed that, for each IDA country, the three aggregate tax revenue-based indicators presented in section 6 are used to provide a broad measure of tax progressivity in IDA countries. The proposed indicators are:

1. The ratio of direct tax revenue to indirect tax revenue;
2. The ratio of income tax (PIT + CIT) revenue to total tax revenue; and
3. The ratio of income tax (PIT + CIT) revenue to GDP.

The first two ratios provide an indication of the degree of reliance of a country on the tax bases most likely to create progressivity in the tax system, while the third provides an indication of the magnitude of these (potentially) progressive taxes, thereby giving an indication of their potential redistributive impact.

[34] There are also challenges associated with modelling *ad quantum* excise taxes with expenditure data.

[35] For example, the microsimulation model used by the Urban-Brookings Tax Policy Center assumes that the CIT is 60% borne by capital income, and allocates projected aggregate CIT liability across the taxpayers in their tax return microdata sample in proportion to the capital income that they earn. (Nunns, 2012; Rohaly et al., 2005).

While the first two indicators are similar, there may be merit in calculating, if not necessarily presenting, both. The first indicator has the benefit that it is a commonly used and understood metric. Meanwhile, from a presentational perspective there may be merit in the second indicator's focus solely on taxes that are considered "good" for progressivity. This potentially contrasts with the negative connotation that may be created by the first indicator that indirect taxes are "bad" whereas direct taxes are "good". Not only is that a simplistic interpretation from a distributional perspective, but it may cloud the fact that there are strong tax policy rationales for raising significant revenue from indirect taxes within a broader inclusive growth strategy. Additionally, the second indicator provides a more muted response to differences in revenue source (due to direct taxes also appearing in the denominator). Given the bluntness of the indicator, this muted nature may be preferable to a more highly responsive indicator. Finally, the second indicator provides additional consistency with the third indicator as they are based on an identical numerator.

It is proposed that these indicators would be calculated for the most recent years possible for IDA countries to enable monitoring over time. In most cases, this will involve a 1-2 year lag, though for some countries it will be greater. For countries where lags are significant, an alternative approach could be to present (top) statutory tax rates for major tax types including PIT, CIT and VAT. However, these would ideally need to be complemented with information on the breadth of these tax bases (such as the C-efficiency ratio for VAT),[36] and would require similar assumptions to be made regarding the incidence of the different taxes in order to draw implications regarding progressivity.

These are obviously blunt instruments, as no detail is provided on tax burdens borne by poor vs rich taxpayers. This is particularly important in the case of the PIT where, for example, a flat tax regime could be applied that raises the same total revenue as a highly progressive tax, but this would not be differentiated using the tax revenue data. The remaining indicators therefore provide greater detail on the PIT.

6.2 Effective progressivity indicators

To examine the effective progressivity and redistribution achieved by PIT systems in IDA countries, it is proposed that Kakwani and Reynolds-Smolensky effective progressivity (and redistribution) indicators be calculated for each IDA country. These indicators could be constructed from a full microsimulation model. However, building a full microsimulation model even for a single country is highly data and resource intensive.[37] In the context of producing progressivity assessments for more than 60 IDA countries, it is proposed that a simplified modelling approach is more appropriate.

[36] The aggregate tax revenue figures effectively incorporate both the tax rate and breadth-of-base information into the one indicator.

[37] The FPSG unit has constructed PIT microsimulation models for a number of developing countries as part of its general microsimulation capacity building work. However, no models have been developed as yet for IDA countries. Constructing these models requires a significant resource investment, and so their development is more suited to country-specific projects rather than for an analysis of 60+ countries. Full microsimulation models also have a broader use than just progressivity assessment. Specifically, they are constructed to provide Finance Ministries with the ability to assess both the revenue and distributional impacts of potential tax reform options. To maximize accuracy of revenue figures, they are typically based on tax return microdata rather than income survey data (and so do not account for the full income distribution – which is necessary for constructing progressivity measures). That said, work is underway to adapt the unit's tax return data-based models to utilize income survey data, and to develop methodologies to address the limitations of different data sources.

This simplified approach would draw on the World Bank's POVCAL database, which includes income and/or consumption survey data for more than 180 countries, including 67 IDA/blend countries. Availability of the underlying microdata for each country depends on the agreements that have been made between countries and the World Bank, but, in many cases, it is anticipated that the microdata would be accessible via the respective country offices. While income data is available for many of the countries in the full database, income data is only available for five IDA/blend countries. Consumption data is available for 62 IDA/blend countries. In addition to the microdata, the POVCAL database provides, for each country, cumulative income or consumption distribution results across percentiles. Datasets are from between 2014-2019 for the vast majority of countries.

While the underlying microdata files would be required for each country to build a full microsimulation model, the publicly available percentile distributions can be used under a simplified modelling approach.[38] The simplified modelling approach would apply the PIT rules in place for single individuals to calculate the average tax liability applicable for each of the 100 percentiles. The income distribution and tax liability data would then be used to calculate Kakwani and Reynolds-Smolensky indicators of progressivity and redistribution. The modelling would account for any basic allowances or deductions, zero rate brackets, and the progressive PIT rate schedule. However, it would not account for any tax provisions that are contingent on family size – such as a child tax credit. While this simplification would be a significant concern for high-income countries, it is likely to provide less inaccuracy for IDA countries that apply far more limited use of such targeted support. A fixed adjustment would also be applied to the consumption survey percentile results to adjust for savings patterns and thereby better proxy for income.

The modelling results can be compared to the small number of countries for which full microsimulation models will be constructed as part of other projects over the period of the project. For example, a full microsimulation model will be constructed using newly available income survey microdata for Ethiopia in late 2022 as part of a project examining potential PIT reform in Ethiopia.

To enable up-to-date monitoring of progressivity in IDA countries, it is proposed that two sets of calculations be made for each country:

1. Modelling the tax rules for the year of the most recent income or consumption survey data, on that survey data.
2. Modelling the tax rules for each year from 2023-2025 on the most recent income or consumption survey data (inflation adjusted to the year of measurement 2023-25), to enable monitoring.

The second approach effectively would assume that the income distribution has remained unchanged between the survey year and the year of measurement. As income distributions can be expected to change relatively slowly, the older distribution should provide a reasonable approximation of the current distribution for most countries, with the small degree of inaccuracy a reasonable tradeoff against the benefit of providing a timely indication of effective progressivity.[39]

[38] From the cumulative income/consumption distributions it is possible to 'back-out' the average income/consumption of each percentile (with the average income and population data that is also provided with the percentile results), thereby enabling this data to be used for PIT modelling.

[39] This approach differs from a progressive capacity approach where the 'base' distribution would not necessarily have any link to the actual distribution. That is, we are using the older distribution as a proxy for the current distribution, not as a replacement for it.

6.3 Progressive capacity indicators

It is proposed also to complement the effective progressivity indicators with progressive capacity indicators. These two approaches are asking essentially two separate questions and there is consequently merit in producing both: while effective progressivity metrics assess the progressivity of the taxes actually paid by households, progressive capacity metrics assess the progressivity of the tax structure itself.

The case to also calculate progressive capacity metrics is strengthened by the fact that the proposed simplified approach to modelling effective progressivity can be leveraged to produce progressive capacity metrics at low additional resource cost. This is because: (1) the same information on underlying tax rules in IDA countries will need to be collected for either metric; and (2) a simplified tax modelling approach is adopted in each case – meaning that the percentile-based model that will be constructed for the effective progressivity metrics can easily be adapted for the progressive capacity metric.

The proposed methodology would be based on the simplified approach of Gerber et al. (2020), as outlined in section 5. This approach is based on average tax rate calculations made at 500 individual income points (from 1%-500% of GDP per capita). This provides a simple option to produce a progressive capacity metric, enabling changes in the progressivity of the tax schedules of an IDA country to be identified in isolation from their interaction with the underlying income distribution. Kakwani and Reynolds-Smolensky indices would be calculated based on this methodology for the tax rules applicable in each year from 2023-25 to enable monitoring.

As noted in section 5, Gerber et al. (2020) undertook their study for OECD countries, and it is possible that the methodology may not apply as effectively for IDA countries. This is because restricting the upper bound of the income distribution to 500% of GDP per capita could exclude significant parts of the progressive PIT schedule in some IDA countries. This will become clear as data on IDA countries PIT schedules is collected, and the methodology may be adjusted accordingly, if necessary. An alternative option that could be progressed is the simplified "transplant and compare" approach of Vellutini and Benítez (2021). While slightly more complicated to implement, this approach does not impose any ceiling on the income distribution and reduces the likelihood of country rankings being affected by the underlying 'base' distribution.

7 Conclusion

Measurement of tax progressivity is a challenging issue for several reasons. Most significantly, data limitations have a large impact on what taxes can be examined, what measurement approaches can be adopted to examine those taxes, and how reliable those approaches are likely to be. This is particularly important for IDA countries where data limitations are typically most significant. Additionally, there is no single definition of progressivity, or single method for measuring the degree of progressivity of a tax, thereby necessitating a range of decisions and value judgements to be made in any analysis. Contingent upon the definition chosen, there are several approaches that can then be adopted to measure the degree of progressivity of a tax, each with benefits and drawbacks. Finally, measurement of tax progressivity – particularly when based on microdata – can be extremely resource intensive, and this must be taken into account, particularly in the current context where the progressivity of a large number of countries needs to be assessed.

Historically, there have been two main approaches to measuring tax progressivity: examining the progressivity of a tax structure; and examining the progressivity of taxes paid. Musgrave and Thin (1948) distinguished these as "structural progressivity" and "effective progressivity" measures. A third, more recent, concept is "progressive capacity" which captures elements of both structural and effective progressivity.

"Structural progressivity" focuses purely on the rate structure of an income tax, and typically produces "local" progressivity metrics (i.e., measures of progressivity at different points along the income distribution). A strength of this approach is that metrics are easy to compute and require no information beyond the underlying tax rules. However, this is a cost also: in addition to producing only local measures of progressivity, the approach entirely ignores the underlying income distribution.

In contrast, "effective progressivity" considers the interaction of both the tax schedule and the underlying income distribution, and typically produces "global" progressivity metrics (i.e., a single figure assessing the progressivity of the tax in question). Effective progressivity arguably provides a far more complete picture of the progressive impact of a tax because it is based on the actual tax liabilities of different households. For example, it is inconsequential to apply a very high marginal tax rate on a particular income range if there are no taxpayers actually in that income range. Effective progressivity calculations can potentially be undertaken for personal income taxes (PIT), social security contributions (SSC), value-added taxes (VAT) and excise taxes (including health and environmentally related taxes). The key drawback of the effective progressivity approach is that it is highly resource and data intensive. It requires information on both the tax system and on the underlying distribution of income (ideally microdata), and presents various challenges associated with data quality.

"Progressive capacity" is a type of hybrid approach, producing global metrics of progressivity of income taxes, but attempting to abstract away from the underlying pre-tax income distribution to focus on the tax structure. This approach may be particularly useful for cross-country comparison where the analyst is interested purely in comparing tax rate structures. The approach essentially looks to model the different tax schedules of countries using a common income distribution. A drawback of the approach is that the progressive capacity metrics, and consequent country rankings, are not necessarily invariant to the chosen "base" income distribution. However, recent papers have looked to address this base sensitivity problem, though they rely on strong assumptions.

This paper has examined in detail the merits of different progressivity metrics under each of these three approaches, finding strengths and weaknesses in each approach. An overarching finding is that, due to the various data and other analytical limitations, detailed measurement of the *degree* of progressivity is limited largely to the PIT, and to a lesser extent, VAT and excise taxes. These taxes, and particularly the PIT, have consequently been the main focus of the paper. However, because of the need to provide an assessment for IDA countries of the progressivity of the entire tax system, the paper has also illustrated how aggregate tax revenue data – which is available for a larger number of countries in a timely manner – can be used to assess the likely degree of progressivity of tax systems in IDA countries.

Given the varying strengths and weaknesses of different approaches, a suite of metrics is proposed to assess tax progressivity in IDA countries over FY23-25. First, aggregate tax revenue-based metrics would be relied upon to provide a basic assessment of the progressivity of the overall tax system. These metrics assess the degree of reliance of countries on direct taxes – on the basis that these are likely to be the predominant sources of progressivity in tax systems. The degree of progressivity of the PIT – where

detailed analysis is most feasible – would then be investigated utilizing both effective progressivity and progressive capacity metrics, drawing on income and consumption distribution data for more than 60 IDA/blend countries from the World Bank's POVCAL database. In light of the large number of countries to be assessed, a simplified modelling approach is proposed to compute effective progressivity and progressive capacity metrics, while managing data and resource constraints.

References

Alstadsaeter, A., M. Jacob, W. Kopczuk and K. Telle (2016), "Accounting for Business Income in Measuring Top Income Shares: Integrated Accrual Approach Using Individual and Firm Data from Norway", NBER Working Paper No. 22888.

Bachas, P., L. Gadenne and A. Jensen (2021) "Informality, Consumption Taxes and Redistribution", NBER Working Paper 27429.

Benedek, D., De Mooij, R., Keen, M. and Wingender, P. (2019), "Varieties of VAT pass through", *International Tax and Public Finance*, 27, 890–930.

Benzarti, Y., Carloni, D., Harju, J. and Kosonen, T. (2020), "What goes up may not come down: asymmetric incidence of value-added taxes", *Journal of Political Economy*, 128, 4438–74.

Bird, R. and M. Smart (2016), "Taxing Consumption in Canada: Rates, Revenues, and Redistribution", *Canadian Tax Journal*, 64(2), 417–42.

Bishop JA, Formby JP and Thistle PD (1990), International comparisons of tax and transfer progressivity: new evidence from the Luxembourg Income Study, Luxembourg Income Study Working Paper No. 52.

Blanchet, T., I. Flores and M. Morgan (2022) "The weight of the rich: improving surveys using tax data", *Journal of Economic Inequality*, 20, pp119–150.

Caspersen, E., and G. Metcalf (1994), "Is a Value Added Tax Regressive? Annual versus Lifetime Incidence Measures", *National Tax Journal*, 47, 731–746.

Causa, O. and M. Hermansen (2017), "Income Redistribution Through Taxes and Transfers Across OECD Countries", *OECD Economics Department Working Papers*, No. 1453.

Commitment to Equity Institute Data Center on Fiscal Redistribution, 2022. Based on information from: ETHIOPIA (Hill et al., 2014); GHANA (Younger et al., 2018); HONDURAS (Espino, 2020); NICARAGUA (Cabrera and Moran, 2020); TANZANIA (Younger, Myamba, and Mdadila, 2019); UGANDA (Jellema and Renda, 2020); UGANDA (Deisy et al., 2020); UNITED STATES (Carrera et al., 2019). Available at: https://commitmentoequity.org/datacenter.

Creedy, J. (1998), "Are Consumption Taxes Regressive?", *Australian Economic Review*, 31(2), 107–116.

Creedy, J. and C. Sleeman (2005), "Adult equivalence scales, inequality and poverty", *New Zealand Economic Papers*, 39, pp51-81.

Cutler, D. and L. Katz (1992), "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's", *American Economic Review: Papers and Proceedings*, 82(2), 546–551.

Dardanoni, V. and P. Lambert (2002), "Progressivity Comparisons", *Journal of Public Economics*, 86(1), pp99–122.

Flues, F. and A. Thomas (2015) "The Distributional Effects of Energy Taxes", *OECD Taxation Working Papers*, No. 23.

Fuchs, A., P. Marquez, S. Dutta and F. Gonzalez Icaza (2019) *Is Tobacco Taxation Regressive? Evidence on Public Health, Domestic Resource Mobilization, and Equity Improvements*, World Bank.

Fuest, C., A. Peichl and S. Siegloch (2018), "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany", *American Economic Review*, 108(2), pp393-418.

Gale, W. and S. Thorpe (2022), "Rethinking the corporate income tax: the role of rent sharing", Urban-Brookings Institution Tax Policy Center, Mimeo.

Gerber, C., A. Klemm, L. Liu and V. Mylonas (2020), "Income Tax Progressivity: Trends and Implications." *Oxford Bulletin of Economics and Statistics*, 82(2), pp365–386.

Heine, D. and S. Black (2019), "Benefits beyond Climate: Environmental Tax Reform" in Pigato, M. (ed.) *Fiscal Policies for Development and Climate Action*, World Bank.

IFS (2011), "Quantitative Analysis of VAT Rate Structures" in IFS et al., *A Retrospective Evaluation of Elements of the EU VAT System*, Report prepared for the European Commission, TAXUD/2010/DE/328, Chapter 9, pp361–443.

IHS (2011), The effect of VAT on price-setting behavior. In IFS et al., *A Retrospective Evaluation of Elements of the EU VAT System*. Report prepared for the European Commission, TAXUD/2010/DE/328, Chapter 8, 285–341.

Joumard, I., M. Pisu and D. Bloch (2012), "Tackling Income Inequality: The Role of Taxes and Transfers", *OECD Journal: Economic Studies*, 2012/1, 37–70.

Kakwani, N. (1977), "Measurement of Tax Progressivity: An International Comparison." *The Economic Journal*, 87, pp71–80.

Kasten, R., F. Sammartino and E. Toder (1994), "Trends in Federal Tax Progressivity, 1980–93." In *Tax Progressivity and Income Inequality*. Cambridge University Press.

Lambert, P. (2001), *The Distribution and Redistribution of Income*, 3rd edition, Manchester University Press.

Leahy, E., S. Lyons and R. Tol (2011), "The Distributional Effects of Value Added Tax in Ireland", *The Economic and Social Review*, 42(2), 213–235.

Lustig, N. (ed.) (2018), *The Commitment to Equity Handbook*, CEQ Institute.

Metcalf, G. (1994), "Life Cycle versus Annual Perspectives on the Incidence of a Value Added Tax", *Tax Policy and the Economy*, 8, 45–64.

Meyer, B. and J. Sullivan (2003), "Measuring the Well-Being of the Poor Using Income and Consumption", *Journal of Human Resources*, 38, 1180–1220.

Mill, J.S. (1848), *Principles of Political Economy*, Longmans (1921 edition).

Musgrave, R. and T. Thin (1948), "Income Tax Progression, 1929−48," *Journal of Political Economy,* 56(6), pp. 498−514.

Musgrave, R. and P. Musgrave (1989), *Public Finance in Theory and Practice*, Fifth Edition, McGraw-Hill.

Norregaard, J. (1990), "Progressivity of Income Tax Systems." OECD Economics Studies, 15.

Nunns, J. (2012), "How TPC distributes the corporate income tax", Urban-Brookings Institution Tax Policy Center, Mimeo.

O'Donoghue, C., M. Baldini and D. Mantovani (2004), "Modelling the Redistributive Impact of Indirect Taxes in Europe: An Application of EUROMOD", *EUROMOD working paper*s, No. EM7/01.

OECD (2014), *Taxing Wages*, OECD Publishing, Paris.

Pigou, A.C. (1928), *A Study in Public Finance*, London, Macmillan.

Rohaly, J., A. Carasso and M. A. Saleem (2005), "The Urban-Brookings Tax Policy Center Microsimulation Model: Documentation and Methodology for Version 0304", Urban-Brookings, Mimeo.

Ruiz, N. and A. Trannoy (2008), "Le caractère régressif des taxes indirectes : les enseignements d'un modèle de microsimulation", *Economie et Statistique*, 413, 21–46.

Ruiz, N. and N. Woloszko (2016) "What do household surveys suggest about the top 1% incomes and inequality in OECD countries?", OECD Economics Department Working Papers, No. 1265.

Sabirianova Peter, K., S. Buttrick and D. Duncan (2010), "Global Reform of Personal Income Taxation, 1981−2005: Evidence from 189 Countries," *National Tax Journal*, Vol. 63(3), pp447−78.

Sassi, F. et al. (2018) "Equity impacts of price policies to promote healthy behaviours", *The Lancet*, Vol. 391, No. 10134, pp2059-70.

Slitor, R. (1948), "The Measurement of Progressivity and Built-in Flexibility", *Quarterly Journal of Economics*, 62, pp309-13.

Smith, A. (1776), *The Wealth of Nations*, Putnam, (1904 edition).

Suits, D. (1977), "Measurement of Tax Progressivity," *American Economic Review*, 67, pp747-752.

Thomas, A. (2022) "Reassessing the Regressivity of the VAT", *Fiscal Studies*, 43(1), pp23-38.

Thomas, A. (2020) "Reassessing the Regressivity of the VAT", *OECD Taxation Working Papers*, No. 49.

Vellutini, C. and J.C. Benítez (2021), "Measuring the Redistributive Capacity of Tax Policies", IMF Working Paper, No. WP/21/252.

World Bank (2022), “Building Back Better from the Crisis: Toward a Green, Resilient and Inclusive Future”, Report from the Executive Directors of the International Development Association to the Board of Governors.

Young, H. P. (1987), “Progressive taxation and the equal sacrifice principle”, *Journal of Public Economics*, 32, pp203-214.

Zodrow, G. (2007) “The property tax as a capital tax: a room with three views”, James A Baker III Institute for Public Policy, Rice University, Mimeo.

Annex: Distributional effects of the VAT

This annex is a condensed version of the discussion and results presented in Thomas (2022, 2020). See these papers for additional detail.

Most studies examining the distributional effects of the VAT have used cross-sectional household expenditure microdata to calculate and present average VAT rates measured in relation to either income or expenditure. This choice between income- or expenditure-based analysis has proved determinative. The common finding that the VAT is regressive has followed from the analysis of VAT burdens measured as a percentage of current income across the income distribution. Numerous European country studies (e.g., Leahy et al., 2011; Ruiz and Trannoy, 2008; O'Donoghue et al., 2004) adopt this analytical approach, and as a result conclude the VAT is a highly regressive tax. In contrast, studies that present VAT burdens as a proportion of current expenditure across either the income or expenditure distribution (e.g., Bird and Smart, 2016; IFS, 2011; Metcalf, 1994) tend to find that VAT systems are relatively proportional, or even slightly progressive.

As has been highlighted by various authors (e.g., IFS, 2011; Creedy, 1998; Metcalf, 1994; Caspersen and Metcalf, 1994), the driver of the stark difference in results between the income-based and expenditure-based approaches to measuring the distributional impact of a VAT (or other broad-based consumption tax) is the influence of savings behavior. Specifically, a single-year income-based analysis ignores the fact that income that is saved in the current year will still incur VAT when it is eventually consumed. Similarly, current expenditure, and the VAT incurred on it, may have been funded from income earned in a previous year. Because savings rates tend to increase with income, this biases income-based VAT burden results downwards at higher income levels – hence the common conclusion that the VAT is regressive.

To fully account for the impact of saving behavior, a lifetime (or at least multi-period) analysis – including calculation of both lifetime income and lifetime VAT burdens – would ideally be undertaken. Unfortunately, any attempt at estimating lifetime income and lifetime VAT burdens is highly complex, even in a single-country context, and simply impracticable for a study covering many countries. However, in the absence of such information, measuring VAT burdens relative to current expenditure can be expected to provide a more reliable estimate of the lifetime distributional impact of the VAT than measuring VAT burdens relative to current income. This is illustrated below in a simple two-period model.

Consider a taxpayer that earns income of y_1 in period 1 and y_2 in period 2. Savings, equal to s, occurs in period 1 and is fully spent in period 2. For simplicity, bequests are ignored (although they could be incorporated without altering the results).[40] Also for simplicity, a two-rate VAT system is assumed, with $x\%$ of the taxpayer's consumption subject to taxation at rate t, and $(1-x)\%$ subject to a zero tax rate. If income from savings is not taxed[41] and the return on savings equals the discount rate, savings will cancel

[40] In order to avoid the double counting of income, bequests received would be included in period 1 income, and bequests given would be subtracted from period 2 income. The analysis then remains unchanged as long as t and x remain constant for both donor and recipient. Note that, while it could be argued that donors derive some consumption value out of providing a bequest, this would lead to the double counting of at least some income.

[41] If income from savings is taxed (as is common), then the income saved may incur higher total (income plus consumption) taxation than income immediately spent. However, as Creedy (1998) emphasizes, it would be erroneous to attribute this increased tax burden to the VAT. Indeed, the NPV of the VAT paid (as opposed to income

out over the two periods, so that the net present value of lifetime income, Y_{npv}, and lifetime consumption, C_{npv}, are the same and are equal to:

$$Y_{npv} = C_{npv} = (y_1 - s) + \frac{(y_2 + s(1+r))}{1+r}$$

$$= y_1 + \frac{y_2}{1+r}$$

The net present value of tax payments, T_{npv}, is:

$$T_{npv} = tx(y_1 - s) + \frac{tx(y_2 + s(1+r))}{1+r}$$

$$= tx(y_1 + \frac{y_2}{1+r})$$

As such, the average tax rate paid over the taxpayer's lifetime is:

$$\frac{T_{npv}}{Y_{npv}} = tx$$

Knowing the lifetime average tax rate, this can then be compared with income-based and expenditure-based calculations of the average tax rate for a single period. If data are only available for period 1, the income-based average tax rate, calculated as t_1/y_1, is:

$$t_1/y_1 = \frac{tx(y_1 - s)}{y_1}$$

It is clear from this that, if any savings occurs, the income-based single period average tax rate will be an inaccurate measure of the lifetime average tax rate. Households that save will have a lower single period average tax rate than their lifetime average tax rate (the greater the amount of savings the lower the average tax rate). Meanwhile, households that dis-save will have a higher average tax rate than their lifetime average tax rate. If richer households save a greater proportion of their income than poorer households (which, as shown in Thomas (2020), is true on average in all 27 countries examined), then the average tax rate for richer households will be lower than for poorer households and the VAT will appear regressive.

In contrast, the expenditure-based average tax rate, calculated as t_1/c_1, is exactly the same as the lifetime average tax rate:

$$t_1/c_1 = \frac{tx(y_1 - s)}{(y_1 - s)}$$

$$= tx$$

Measuring VAT burdens relative to current expenditure, rather than income, removes the influence of savings behavior. The average tax rate is instead driven by the consumption pattern of the household, as

tax) will be unaffected by the taxation of income from savings as long as the taxpayer's discount rate equals the after-tax return on savings.

captured by x. The distributional impact of the VAT is therefore driven by how x varies across taxpayers. If x is constant across all taxpayers then the VAT will be proportional. However, if x is lower for poorer taxpayers – i.e., if a smaller proportion of poorer households' expenditure is subject to the standard VAT rate than of richer households' expenditure – then the VAT will be progressive. Conversely, if x is higher for poorer taxpayers then the VAT will be regressive. This is of course an empirical question, which is examined in the microsimulation modelling results summarized below.[42]

In practice, the expenditure-based current period average tax rate will still be an imperfect estimate of the lifetime average tax rate. In particular, the above model assumes that both t and x are constant over the taxpayer's lifetime. If, however, VAT rates decrease over time then the taxpayer's lifetime tax burden will be overestimated (and vice versa). Likewise, if the household's expenditure pattern shifts over time towards less heavily taxed goods, then the lifetime tax burden would also be overestimated (and vice versa).[43] This could have a regressive impact if, for example, the richest households spend a greater proportion of their savings on less-taxed items such as private education than other households do. Despite these limitations, by removing the strong influence of savings behavior, analysis based on current expenditure can be expected to provide a far more accurate picture of the distributional effect of the VAT than an analysis based on current income.

A number of additional arguments can also be made for preferring an expenditure-based analysis over an income-based analysis. These include that current expenditure may provide a better measure of an individual's welfare than current income (see, for example, Meyer and Sullivan, 2003), and that current expenditure may be a better proxy for lifetime income than current income (see, for example, Metcalf, 1994). More generally, it is arguable that any tax should be assessed relative to its base because a tax cannot redistribute something that it is not applied to – and the base of the VAT is expenditure. Importantly, as emphasized by IFS (2011), the above savings-based rationale does not rely on these additional arguments. Eliminating the misleading impact of savings behavior remains the clearest rationale for preferring an expenditure-based analysis.

Microsimulation results for 27 OECD countries broadly confirm the dichotomous results from previous smaller-scale studies: the VAT appears to be regressive when measured as a percentage of current income in all 27 countries, but appears generally either proportional or slightly progressive when measured as a percentage of current expenditure. Savings patterns are also consistent across all 27 countries, with savings rates increasing as income increases and thereby driving the regressivity of the income-based results. (Detailed country-specific results are presented in Thomas, 2020.)

Figure A1 presents the simple all-country average across 27 OECD countries of VAT burdens measured both as a percentage of income and expenditure, across both income and expenditure deciles. The left hand panel of Figure A1 shows, as expected, a strongly regressive pattern when VAT is measured as a percentage of income across income deciles. However, the other results are in strong contrast to this. When VAT burdens are instead measured as a percentage of expenditure (either across income or

[42] In practice, VAT structures are typically more complex than in this simple model, often involving multiple reduced VAT rates as well as exemptions. However, it remains the variation in consumption patterns across households that drives the distributional impact of the VAT – variation which is captured in the underlying household expenditure survey microdata used in the modelling that produces the results presented in Figures A1 and A2.

[43] As noted above, if bequests are incorporated into the model, then t and x should also be constant for both donor and recipient.

expenditure deciles), Figure A1 shows a slightly progressive pattern. Meanwhile, when VAT burdens are measured as a percentage of income, but across expenditure deciles rather than income deciles, then the results appear strongly progressive (see right hand panel). These overall trends are highly reflective of the country specific results.

Figure A1: Household average VAT burdens: 27-country simple average

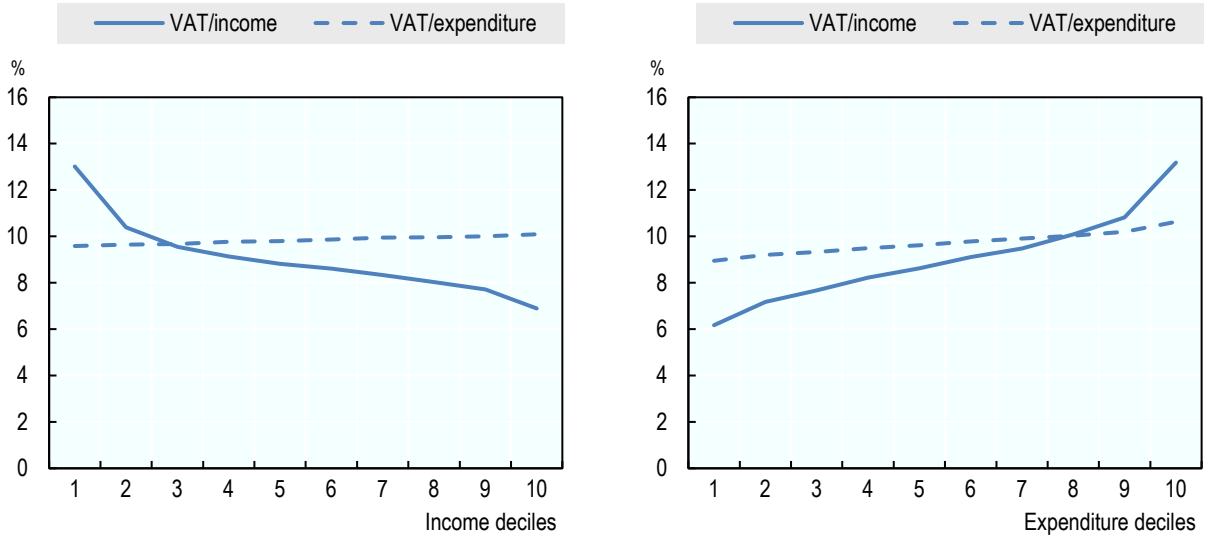

Notes. Results calculated using microsimulation models for 27 OECD countries based on household expenditure survey microdata. "VAT/income" results present the simple mean across all countries of the within-country weighted mean of the VAT paid per household as a percentage of total household income. "VAT/expenditure" results present the simple mean across all countries of the within-country weighted mean of the VAT paid per household as a percentage of total household expenditure. Results presented across equivalized income/expenditure deciles, with the individual as unit of analysis. Italy only included for VAT/expenditure across expenditure deciles. See Thomas (2022) for further detail.

Figure A2: Expenditure-to-income ratio: 27 country simple average

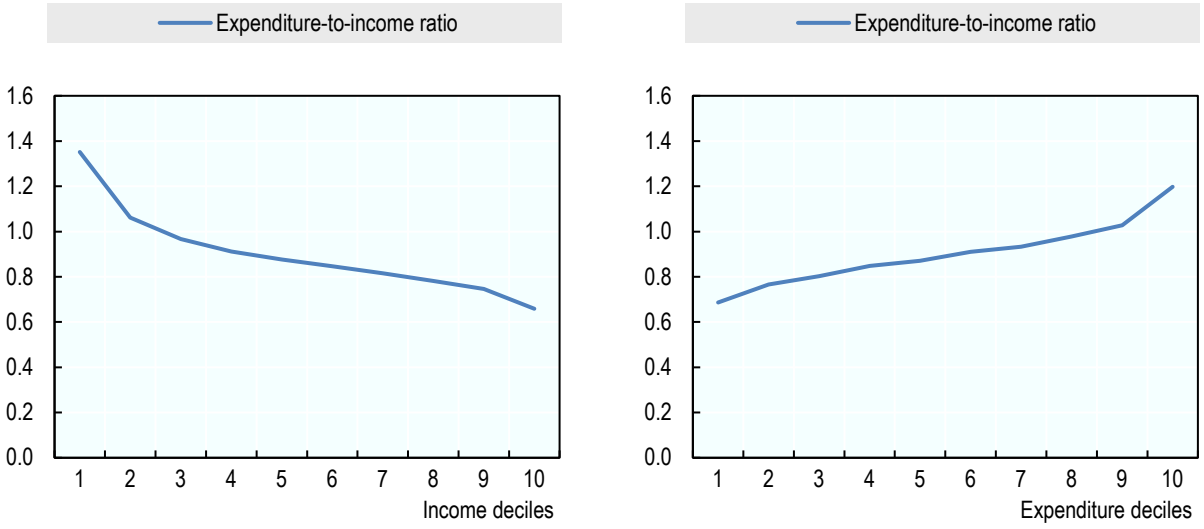

Notes. Results calculated using microsimulation models for 27 OECD countries based on household expenditure survey microdata. "Expenditure-to-income ratio" results present the simple mean across all countries of the within-country weighted mean of total expenditure divided by total income per household. Results presented across equivalized income/expenditure deciles, with the individual as unit of analysis. Italy excluded as income data not available. See Thomas (2022) for further detail.

The highly regressive pattern summarized in the left hand panel of Figure A1 when measuring VAT as a percentage of income across income deciles is consistent with previous analyses of VAT burdens as a percentage of income. The driving influence of savings behavior on these results is illustrated by the average expenditure-to-income ratios presented across income deciles in the left-hand panel of Figure A2. At low income levels, households tend to be net borrowers in all 27 countries, so average VAT burdens as a percentage of income appear relatively high. Savings rates then rise with income, lowering average VAT burdens. At high income levels, households tend to be net savers, and consequently VAT burdens as a percentage of income appear relatively low.

Equally, the highly progressive pattern shown in the right hand panel of Figure 1 when measuring VAT as a percentage of income across expenditure deciles is driven by savings behavior. This is illustrated by the expenditure-to-income ratios presented across expenditure deciles in Figure A2 (right hand panel). At low expenditure levels, households tend to be net savers, so VAT burdens as a percentage of income appear relatively low. Savings rates then fall as expenditure increases, increasing the average VAT burdens. At very high expenditure levels, households tend to be net borrowers, and so VAT burdens as a percentage of expenditure appear relatively high.

In contrast, the expenditure-based results summarized in Figure A1 remove the influence of savings behavior. As noted above, they instead identify how the presence of reduced VAT rates and exemptions move the actual VAT burden away from what would be due under a perfectly broad-based single-rate system (where all households would pay the same proportion of their expenditure in VAT). As such, they provide a far more reliable estimate of the distributional impact of the VAT.

These expenditure-based average tax rate results are confirmed by calculations of expenditure-based summary indicators of progressivity and redistribution. Kakwani progressivity index results show a low degree of progressivity of the VAT in almost all countries, often extremely close to proportionality. The exceptions are Chile, Hungary, Latvia and New Zealand, where a very small degree of regressivity is found. Reynolds-Smolensky (and Atkinson index) results show the VAT to have minimal redistributive effect, despite significant average tax rates being applied.

Overall, the VAT is found to be either roughly proportional or slightly progressive in most of the 27 OECD countries examined. Nevertheless, results for a small number of countries highlight that broad-based VAT systems that have few reduced VAT rates or exemptions can produce a small degree of regressivity. Additional results (not presented here) also show that even a roughly proportional VAT can still have significant equity implications for the poor – potentially pushing some households into poverty. This emphasizes the importance of ensuring the progressivity of the tax-benefit system as a whole, and of the merits of providing targeted support (via cash transfers where feasible) to poor households.