

INTERNATIONAL MONETARY FUND

Efficiency Aspects of the Value Added Tax

Ruud de Mooij, Shafik Hebous, and Michael Keen

WP/25/165

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2025
AUG**



WORKING PAPER

IMF Working Paper
Fiscal Affairs Department

Efficiency Aspects of the Value Added Tax

Prepared by Ruud de Mooij*, Shafik Hebous* and Michael Keen**

Authorized for distribution by Ruud de Mooij
June 2025

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: This paper examines the efficiency of the Value Added Tax (VAT), focusing on its role as a revenue-raising tool and its use to achieve non-revenue objectives. The analysis highlights the VAT's potential ability to generate revenue with minimal distortions, emphasizing its advantages over alternative taxes, such as turnover taxes and tariffs, particularly in minimizing the cascading effects of input taxation. The paper also explores the VAT as a macroeconomic policy tool, especially in counter-cyclical fiscal policy, and its potential to address environmental and health objectives. It concludes that a well-designed and implemented VAT is a highly efficient revenue-raising tool, surpassing other forms of consumption taxation, but cautions against its misuse in industrial policy and other contexts for which it is ill-suited.

RECOMMENDED CITATION: De Mooij, R., Hebous, S., and Keen, M. (2025). Efficiency Aspects of the Value Added Tax. IMF Working Paper WP/25/165.

JEL Classification Numbers:	H21
Keywords:	Value Added Tax; Efficiency; Welfare
Author's E-Mail Address:	rdemooij@imf.org , shebous@imf.org , mickjkeen@gmail.com

* International Monetary Fund, Fiscal Affairs Department.

** Tokyo College, University of Tokyo; Institute for Fiscal Studies, London; Centre d'Études et de Recherches sur le Développement International; Centre for Business Taxation, Oxford.

* This is the draft of a chapter for the forthcoming IMF book *The Changing World of VAT*, edited by Andrew Okello, Artur Swistak, and Christophe Waerzeggers. We are grateful to the editors and to seminar participants at the IMF for valuable comments and suggestions on earlier draft.

Contents

1. Introduction	2
2. The VAT as an Efficient Revenue-Raiser	2
2.1 Why Tax Consumption?	3
2.2 Efficiency in the taxation of final consumption	11
2.3 Why a VAT?	14
2.4 How Large are the Welfare Gains from the VAT?	18
2.5 Who Really Pays the VAT?.....	19
3. Beyond Revenue	22
3.1 The VAT and Short Run Macroeconomic Management	23
3.2 The VAT and Long Run Economic Performance	27
3.3 The VAT in Pursuit of Other Objectives	30
4. Conclusion.....	32
References	34

BOXES

1. Individual-based Consumption Taxation	3
2. Distortions from input taxation.....	7
3. Gains from Shifting from a Tariff to a Consumption Tax	9

FIGURES

1. Deadweight loss from taxing inputs	6
2. Searching for a Correlation between Inflation and VAT Rate Changes	27

1. Introduction

This paper is concerned with the question of whether, and, if so, to how significant an advantage, the Value Added Tax (VAT) is suited to the pursuit of ‘efficiency’ in relation to a range of objectives that governments might set for it. The most fundamental purpose that the VAT serves, of course, is raising revenue: Efficiency in this dimension, and the related question of who it is that really pays the VAT, are the subject of the first part of this paper (Section 2). As the VAT has become an increasingly large part of overall tax systems, however, and as new or reinvigorated policy objectives have emerged, so the VAT has become increasingly looked on to serve non-revenue objectives. These are the subject of the second part (Section 3).

To begin, one needs some clarity on what is meant by ‘efficiency,’ a term that is often used with little precision. Here we take it to mean minimizing the resources used in pursuing whatever the objective to hand may be. This is intended as a more accessible formulation of the fundamental notion of ‘Pareto efficiency’: a situation in which it is impossible to make one person better off without making someone else worse off. In some contexts, we follow the conventional practice of supposing there to be a single representative consumer, so that efficiency boils down to maximizing their well-being subject to some kind of constraint.¹ Importantly, the pursuit of efficiency as we use the term here is not necessarily a matter of maximizing output, or GDP. Reducing harmful emissions, for instance, may require reducing output; efficiency, in that context, would mean minimizing the fall in output required to achieve a given level of emissions reduction. It is not always easy to operationalize this very general idea of efficiency, but it provides a basis for some analytical rigor and, with added assumptions, in some cases enables rough quantification of the potential efficiency effects from the VAT.

In focusing on efficiency concerns, this paper sets aside, as far as possible, issues of equity and fairness around the VAT: see e.g. De la Feria and Swistak (2024). They must, however, be touched on here too, as the two sets of issues are not completely separable. Equity objectives may require compromising efficiency in the pursuit of other objectives, for example, but that compromise can be achieved in more or less efficient ways. While they thus need to be recognized, equity aspects of the VAT are addressed in this paper only in so far as they affect what the pursuit of efficiency requires.

2. The VAT as an Efficient Revenue-Raiser

The primary merit widely claimed for the VAT is that it is, or can be, a particularly efficient way of raising substantial amounts of tax revenue. The factual question of whether or not the VAT has indeed generated significant amounts of revenue—its ‘effectiveness’—is subject of analysis in e.g. Baer et al. (2025). Here the focus is on the other element of that claim: that the way in which the VAT raises revenue is, or might be, particularly efficient; more efficient, that is, than practicable alternatives.

Since the VAT is simply one form of consumption tax, the analysis needs to proceed in steps. Why tax consumption at all? What would an efficient consumption tax look like? And then: Why do those considerations point to the use of the VAT, and how do they suggest it be designed? It then turns, speculatively, to the critical bottom line: How large might the efficiency gains from using the VAT to raise revenue plausibly be? The section

¹ In other contexts, it is harder to make a direct link between policy and efficiency, so that we make broad but seemingly reasonable assumptions on this (such as, for instance, that reductions in unemployment are efficiency-improving).

concludes by asking whether, as is often quietly presumed, the VAT is truly a tax ‘on’ consumption: That is, who really pays the VAT?

2.1 Why Tax Consumption?

Consumption—whether in aggregate or of particular commodities—is just one of many possible tax bases: others include production, trade, income, and wealth. What then does consumption taxation have to offer when there are many other tax instruments that might be deployed instead, or in addition?

Three ideas suggest distinct though not necessarily unique appeal to consumption as a tax base: that consumption is relatively easy to observe and so to tax; that a person’s consumption is a good indicator of their ‘ability to pay,’ and so can provide a fair tax base; and that taxing consumption avoids distorting production, which is inherently desirable. It is the third of these ideas that is most critical for efficiency concerns, but there are also lessons from the other two.

The first serves as reminder of the need for clarity on what ‘consumption’ is. From the economic perspective, the core idea of ‘consumption’ is that it captures the enjoyment (or possibly harm) that individuals realize by their direct use of real resources. So one key element is that consumption is something that only real people, not businesses, do: talk of ‘business consumption’ in the VAT or other context can be a source of confusion. Seen in that light, what is often observable and therefore taxable is not consumption in this wider enjoyment sense but the act of purchase: when someone buys a raincoat their actual enjoyment is spread over the months or years over which they can wear it. Purchase is thus taken as a proxy for consumption, and taxing purchases of durables is in effect taxing future consumption, which, given the importance of some durables—cars, housing, and secondhand items—raises distinct issues for the VAT. Worth bearing in mind too is that some purchases by people in their role as final consumers may not yield enjoyment in themselves but in effect provide services that do: this too turns out to have some significance for the VAT, for example in deciding the appropriate tax treatment of financial services.

The second idea is essentially a matter of equity. Importantly, however, the debate in this important strand of the literature revolves around the case for an individual-based and likely progressive tax on total consumption, not that for a commodity-based consumption tax such as the VAT: Box 1. elaborates. The focus in this paper is of course on the latter, though the distinction is not entirely clear cut: a commodity-based consumption tax could be an integral part of even a progressive individual-based one. Moreover, proposals for individual-based consumption taxation also raise issues of efficiency that will need to be taken up in considering below the respective roles and effects of the VAT and the personal income tax.

Box 1. Individual-based Consumption Taxation

The idea that taxing consumption has an inherent element of fairness dates back at least to Thomas Hobbes (1651) “... the Equality of Imposition consist rather in the Equality of that which is consumed, than of the riches of the persons that consume the same.” Modern variants would stress that, when individuals have some ability to save or borrow, the level of consumption they choose at each point in their life is a better indicator of the level of well-being that they can sustain than is their current income: Comparing an 80 year old and a 30 year old on the same income, the former, having no great need to save for the future, is likely to be meaningfully better off, as reflected in their higher level of consumption.

This perspective of consumption as a better proxy for economic well-being than is current income—combined with the view that there may be efficiency merits in excluding capital income from the tax base, a point taken up below—has led to powerful advocacy for replacing the income tax by an individual-based progressive tax on aggregate consumption (an ‘expenditure tax’).¹ Doing so on a progressive basis, however, would require implementation by means quite different from the usual mechanisms of the VAT, which treats all commodity purchases the same, and so is anonymous in the sense of taking no account of the wider circumstances of the purchaser.² For this reason the paper does not pursue arguments for expenditure taxation in this sense, though, as will be seen below, the view taken as to the relative merits of expenditure and current income in proxying welfare significantly affect views as to the progressivity or otherwise of the VAT.

¹ The classic treatments are in Bradford et al. (1977) and Meade Report (1978), with a recent variant in Mirrlees Review (Adam et al., 2011). Many personal income tax systems do allow for an element of consumption taxation in the form of providing a deduction for pension contributions and then taxing the funds (cumulated with interest) on withdrawal, the result being that only spending is taxed. One form of progressive consumption tax would simply extend this treatment to all savings.

² There are more limited, novel ways to improve progressivity within a commodity-based VAT as discussed in De la Feria and Swistak (2024).

The efficiency argument for not taxing production

The third idea, more subtle, raises principles of taxation that are key to understanding efficiency considerations in the pursuit of tax revenue.

By changing the relative prices that taxpayers face, taxes on consumption and on production both lead to inefficiency in the precise sense of imposing private costs that exceed the amount of tax paid.² On the consumption side, suppose for instance that a taxpayer initially goes to the hairdresser once a month, paying \$10 each time. If a 20 percent tax on hairdressing were levied on each such trip, however, they might decide to go only once every two months. The result would be that they not only pay tax of \$1 per month but are also worse off because their hair is less well cared for. This additional loss is referred to as an ‘excess burden’ (or ‘deadweight loss’): it is the amount by which the loss to the consumer exceeds the loss that is unavoidable if resources (in this case \$1 per month) are to be taken from the consumer and transferred to the government.³ All else equal, good policy will want to make this excess burden as small as possible. Precisely analogous arguments apply to the taxation of business inputs: imposing a tax on machinery, for example, may lead to a shift to less capital-intensive production—methods that, if it were not for the tax, would be more expensive—with the result that the increase in costs to business exceeds the amount of tax they end up paying on the fewer machines they now use.

With both consumption and production taxes generating costly distortions—which tend to increase rapidly with the tax rate applied—one might have expected the minimization of excess burden to require using both types of tax to some moderate degree. Remarkably, however, if the government is able to tax final consumption at

² As used here, ‘paid’ refers to the amount ultimately received by government, irrespective of who literally remits it and of who ultimately bears the real burden of the tax.

³ To see that it is relative prices that matter here, suppose that all prices faced by this consumer had gone up by 20 percent, or the government had simply taken \$1 from their purse. They would still have been worse off by that \$1 but would have had no reason to change their consumption pattern other than because their real income had fallen. Excess burden would be zero.

whatever rates it likes (and some other conditions are met),⁴ this is not so: it is then desirable to tax only consumption and leave production undistorted. The reason for this—the desirability of retaining production efficiency—is quite intuitive. Production inefficiency means that there is less final output available for final consumption than it would be technologically possible to produce; but that cannot be a sensible outcome if, as seems very likely, that additional output could be put to some good use.

This result on the desirability of production efficiency is pivotal for policy analysis because it creates a presumption that production efficiency is desirable unless it is shown specifically how it can be improved upon. It is a powerful tool in comparing the merits of consumption as a tax base relative to alternatives, to which we now turn.

Production efficiency and the choice of tax base

To see the power and limits of the production efficiency result in understanding the role and weaknesses of the VAT, it helps to start by comparing the taxation of consumption with some of its main alternatives. It will be assumed—for now—that each tax is perfectly enforced and complied with. The realities of imperfect implementation are critical, not just for the operational challenges they raise but because they frame the context in which the efficiency or otherwise of alternative tax arrangements must ultimately be assessed.⁵ For the moment, however, we leave those realities aside.

Taxes on turnover

By a ‘turnover’ (or ‘gross receipts’) tax is meant one charged on sales, whether to businesses (as inputs to production) or to final consumers. This approach is still sometimes used as a simple form of taxation for small taxpayers, or as a minimum tax on businesses, but turnover taxes are rarely now a major source of revenue—and are generally not seen as a realistic alternative to the VAT. They do nonetheless retain some lingering appeal, most evidently in recurrent support for some form of financial transactions tax. More to the immediate point, however, is that turnover taxation provides an essential conceptual benchmark in assessing the efficiency, potential and actual, of the VAT, not only because they are what in many cases the VAT replaced but also because—as will be seen later—imperfections in the design and/or implementation of the VAT can lend it properties equivalent to those of a turnover tax.

Such appeal as turnover taxation retains lies in the common but often misplaced idea that the best approach to taxation is to set a low rate on a broad base. It is indeed possible⁶ that, since its tax base would be much broader, a general turnover tax could raise, at a much lower rate, the same revenue as a consumption tax. Attractive though that may sound—the point is often made by advocates of taxing financial transactions—it does not, in fact, imply that a turnover tax is superior to a consumption tax. Account must also be taken of what


⁴ The result (due to Diamond and Mirrlees, 1971) presumes conditions to be competitive, and that the government can tax rents—profits in excess of the minimum required by investors—at whatever rate it chooses (otherwise, there may be a case for taxing or subsidizing inputs on distributional grounds, with related considerations also arising in an international context (Keen and Wildasin, 2004). It also requires that the government be unconstrained in its ability to tax consumption—an important point to which we return in Section 2.3.

⁵ While practical problems of implementation will frame much of that analysis that follows, for brevity we set aside costs of administration and compliance even though, to the extent that these involve real resource use, they matter for efficiency.

⁶ The ‘possible’ is because, exactly for the reasons given in this paragraph, the maximum revenue that could be attained under a turnover tax is less than that possible under a consumption tax.

that base is. And the superiority of taxing consumption follows immediately from the production efficiency result above: by taxing production, even at a low nominal rate, a turnover tax reduces output in a way that a consumption tax does not.

Quite how large this consequent excess burden from turnover taxation is—and so how large the welfare gains from instead using a tax on consumption—depends, beyond the tax rate itself, on three factors. One is the number of stages of production involved, and the consequent ‘cascading’ (or ‘pyramiding’) effect as tax accumulates across them: a 5 percent tax levied on some input becomes a tax of (slightly more than)⁷ 10 percent if also levied on whatever it is used to produce. The second factor is the ease with which producers can substitute away from the taxed inputs. A tax on software, for instance, might lead an accounts department to use less sophisticated software and more labor than it otherwise would have. The greater is the scope for such a change in production practices, the greater is the excess burden: it is zero if, and only if, the amount of the input required per unit of output is a technological constant. (Importantly—a point often missed—this means that cascading of the kind just described is not necessarily a source of inefficiency.)⁸ The third is the proportion of inputs to which the tax applies. When this is small, the efficiency cost, all else equal, is also likely to be small. Less obviously, the inefficiency may also be small when the tax is levied across a very wide range of inputs, the reason being that if, in the extreme case, all input prices rise by the same proportion then their relative prices are unchanged and so decisions as to the mix of inputs to use is unaffected.



⁷ The effective rate in this example is 10.25 percent (because $1.05 \times 1.05 = 1.1025$).

⁸ Even when no production inefficiency arises, cascading has the problematic consequence that the effective tax rate on any commodity becomes hard to gauge, since it depends on the tax rates on all the various inputs used in its production across potentially multiple stages of production. (On the calculation and uses of these effective rates, see Appendix II of Ebrill et al. (2001) and Ahmad and Stern (1986)). This makes it hard to understand exactly what effects the tax system is having, a loss of transparency that can only impede coherent policy making.

Note: Authors' calculation. DWL is deadweight loss from the tax. Based on Keen (2014): $DWL (\% \text{ of cost}) = \frac{1}{2} \alpha(1-\alpha)\sigma t^2$, where α is the share of taxed inputs, σ is the degree of substitution between taxed inputs, and t is the ad valorem tax rate, assumed here to be 5%. The DWL as a share in *tax revenue* is: $\frac{1}{2}(\sigma(1-\alpha) \alpha t)/(1 + \alpha(1+t))$. For example, when $\sigma=2$ and $\alpha=0.5$, the DWL is 0.06% of the cost and 0.82% of tax revenue.

The deadweight loss from a consumption tax is generally a function of the share of taxed input, the elasticity of substitution between inputs, and the tax rate. Keen (2014) derives a specific formula, akin to the famous Harberger's approximation, depicted in Figure 1. As the share of taxed inputs rises, the deadweight loss increases too. The share of taxed inputs tends to be higher under a turnover tax than under a VAT. Moreover, the higher the elasticity of substitution between inputs (given a tax rate and the share of taxed input) the higher deadweight loss.

One way in which firms may be able to avoid the cascading effect is by vertical integration, to conduct the same succession of operations within what is for tax purposes a single entity. While doing so may save the business some tax, it does not eliminate excess burden. It simply changes its form, since there is presumably some cost of integrating (perhaps transaction costs of changing legal form, or the additional cost in now needing to coordinate activities) otherwise it would have been undertaken already. Excess burden can arise in this way even if there is no substitution between the taxed and other inputs.⁹

There is indeed compelling empirical evidence that turnover taxes can cause significant distortions and reductions in activity levels (Box 2). These difficulties with turnover taxation have commonly become very apparent to policy makers, whose initial reaction is often ad hoc. In several African countries, for example, 'ring' systems were developed that exempted certain commodities transacted between traders registered for this purpose, with exemptions also provided for presumptively intermediate goods, such as raw materials or capital equipment. In Uganda and elsewhere, these piecemeal and complex arrangements ultimately proved a steppingstone to adoption of a VAT.¹⁰

Box 2. Distortions from Input Taxation

Hansen, Miller, and Weber (2021) find that the removal in 2015 of a 25 percent turnover tax on sales of cannabis by growers to manufacturers in Washington state substantially reduced vertical integration, and increased production by 23 percent. Smart and Bird (2009) find that the transition from a retail sales tax (largely applying to business inputs) to a VAT led to a 17 percent increase in investment in machinery and equipment.¹¹ The inefficiencies from taxing turnover appear to be large in lower income countries too. For P.R. China, Xing, Bilicka and Hou (2022) find that the replacement of a turnover tax by VAT treatment for services in 2012 increased the sales of affected firms by 12 percent; their employment and R&D spending also rose significantly. Chen et al. (2023) find that allowing firms to deduct input VAT in P.R. China increased investment. Agrawal and Zimmerman (2024) estimate that the shift in India from a sales tax marked by significant input taxation towards a VAT increased sales in the medium term by 57 percent.

The simplicity of turnover taxes nonetheless remains. In the smallest island states, with little domestic production activity, this may more than outweigh possible inefficiencies. Moreover, and especially in developing

⁹ Inefficiencies from turnover taxation even arise in the absence of distortion to input choices or artificial vertical integration, because of the difficulty of establishing trade neutrality by appropriately remitting tax on exports and charging it on imports. This, indeed, was a major consideration in the decision to adopt the VAT as the common form of commodity taxation in the European Union

¹⁰ See Chapter 5 of Ebrill et al. (2001).

¹¹ These reforms were broadly revenue neutral, so that the effects reflect changes in the structure rather than the level of taxation.

countries, turnover taxes may have a useful role in ensuring that at least some tax is paid by traders below the VAT threshold. While the potential revenue from such taxes is often very limited, there may be potential benefits in encouraging a wide citizenry to hold government accountable¹² and, more directly relevant to the concerns of this paper, they can ease distortions of competition relative to larger and/or more compliant enterprises. Indeed, some form of simple taxation on small traders may be warranted as improving production efficiency even if the revenue it directly raises is less than the cost of collecting it.¹³ This use, however, is essentially as a backstop to the core VAT system, not—our primary concern here—as general commodity taxes.

The taxation of consumption avoids these inefficiencies because, in principle, it simply does not tax inputs. In practice, however, imperfections of design and implementation, including of the VAT, commonly imply an element of input taxation, in which case the production inefficiencies just described arise once more. We come back to this in Section 2.3.

Taxes on trade

The ability to intercept physical commodities crossing jurisdictional boundaries has long made the taxation of imports—and, to a lesser extent these days, exports—one of the easiest of taxes to implement. And so it remains (notwithstanding the growth of digitally-deliverable services) in many developing countries, which often raise 25 percent or more of their revenue from tariffs.


A general tax on domestic consumption, in contrast, would be levied not only on imports but, at the same rate, on domestic production; and it would entirely exclude exports: it would be charged, that is, on a ‘destination basis.’ As a result, it would be a more efficient source of revenue than taxes on trade,¹⁴ at least for economies that are small in world markets (in the sense that the prices at which they trade are unaffected by their own actions). This again follows from the result above on the desirability of production efficiency. For a small economy, the social cost of its purchases (or sales) on world markets is given by international prices, since it is these that indicate the prices which the nation ultimately pays or receives. Trade taxes drive a wedge between these social costs and the prices that producers actually pay or receive and so guide their decisions. They are thus akin to taxes on production and so, when the conditions for the desirability of production efficiency hold, are inferior to taxes on consumption.

¹² Some such idea recurs in the literature on taxation and state building.

¹³ Suppose, for instance, that direct revenue gain from a tax on small businesses (initially paying no tax) is viewed as offset by the welfare loss they consequently suffer. Then the relevant comparison in gauging the overall impact on welfare is between the administrative costs incurred and the additional tax paid by the larger businesses that *are* initially paying tax (since their competitive position is affected), with the latter signaling a reduction in the production inefficiency arising from the differential treatment of small and large businesses; see Keen (2013).

¹⁴ We set aside here non-revenue motives for tariffs, such as the protection of favored producers from foreign competition or the restriction of trade so as to secure advantageous movements in prices on world markets.

Box 3: Gains from Shifting from a Tariff to a Consumption Tax



The chart shows a small open economy facing a fixed world price P for some tradable good, with domestic demand DD and domestic supply SS . In the absence of any intervention, domestic demand is at point a , and domestic production at b , so that imports are given by the length ab .

Imposing a tariff of T raises the price faced by both consumers and domestic suppliers to $P + T$. Demand consequently falls to c ; domestic supply increases to d , as producers are able to expand production even though their product is more costly to produce than is that of their foreign competitors. Imports fall to cd , with revenue from the tariff given by the area $cdef$.

Replacing the tariff by a consumption tax at the same rate T has no effect on the consumer price but restores that faced by producers to the world price P , and hence their production to b . Revenue from taxing consumption is $cfhg$.

In welfare terms, the shift from the tariff to the consumption tax leaves consumers indifferent (since they face the same price), reduces producer surplus (because of the withdrawal of protection) by $dbgh$ and increases tax revenue by $degf$. Overall, there is a net gain of deb . This comes from elimination of the production inefficiency induced by the tariff, being exactly the amount by which the increase in domestic production under the protective shield of the tariff was more costly than the imports it replaced.

¹ The argument here is developed in Keen and Ligthart (2002). It assumes competitive conditions; the implications of imperfect competition being explored in Keen and Ligthart (2005).

The point is most easily seen with the example of a tariff on an imported final product that is also produced domestically. Box 3. spells out the argument but, in short: replacing a tariff by a consumption tax at the same rate leaves the price to consumers unchanged and leads to increase in tax revenue (because it widens the tax base by including domestic production as well as imports) that more than offsets the loss to producers. Overall, the replacement of the tariff by a consumption tax is, for the economy as a whole, a clear win—at least in principle.

In practice, however, many developing countries have found the replacement from domestic sources of trade tax revenue lost from trade liberalization—achieving what is sometimes referred to as the ‘tax transition’—difficult. One estimate is that, between 1970 and 2006, 40 percent of developing countries were still suffering a reduction of total revenue 10 years after a liberalization episode (Cagé and Gadenne, 2018). The issue is of significance beyond the VAT itself, of course, since other domestic taxes are also available to replace lost revenue. Nonetheless, the arguments of principle above do suggest that the VAT—or, more precisely, consumption taxes, which would also include many excises—can play a key role. On this, Baunsgaard and Keen (2010), who also find recovery to have been very problematic in lower income countries over a similar period, conclude that the presence of a VAT did not significantly affect the degree of revenue recovery. Somewhat more positively, and perhaps reflecting a focus on a more recent period (1985-2013), Adandohoin (2021) finds that in developing Asia and Africa increased VAT revenue has replaced 33-37 cents of every dollar of lost trade tax revenue.

Quite why taxes on consumption in general, and the VAT in particular, have not played a more evidently decisive role in supporting the tax transition remains unclear. Certainly, the average effects reported above conceal considerable diversity in countries’ experiences, with some signs that,¹⁵ as one might suspect, imperfections in design and implementation of the VAT play a role. But that does not provide a complete explanation: even a VAT that succeeded in collecting revenue only on imports—which should be straightforward for any revenue administration to achieve—should at least match the revenue performance of a tariff. Quite why the VAT has not enabled a smoother tax transition remains something of a puzzle—and one worth more (and updated) attention than it has received.

Leaving aside the comparison with trade taxes, a general consumption tax—one levied at a uniform proportionate rate on all commodities—is expected to have no impact on trade, so long as prices and/or the exchange rate are sufficiently flexible. This result and empirical evidence on whether the VAT has indeed proved to be trade neutral are explored in Section 3.2.

The personal income tax (and transfers)

The personal income tax (PIT) differs from consumption taxation in two fundamental respects. The first is the base, which for the PIT is essentially wages (payment for labor services), capital income (the return to savings) and business income (in turn commonly a mixture of wages and capital income).¹⁶ There are close similarities between wage taxation and consumption taxation: For an individual who receives only wage income over their lifetime and spends it all, and is able to freely borrow and lend, a flat tax applied at the same rate over time to all wages is equivalent to a tax on all consumption applied at a constant rate over time: one is a tax on all ultimate sources of income, the other is a tax on all final uses. They will thus create the same excess burden, though the possibility of taxing different commodities at different rates may be able to mitigate the effect—an issue for the next subsection.

The more fundamental difference of base is in relation to capital income. When only consumption is taxed, the ‘normal’ return to capital—meaning for present purposes the minimum return required by the investor—is

¹⁵ As in Adandohoin (2021) and IMF (2005).

¹⁶ These may be taxed at different rates, but this is inessential for what follows.

excluded from tax;¹⁷ under an income tax, it is not. Some have argued that the excess burden from taxing the normal return can be extremely large, and this has been another reason for advocacy of a broad-based progressive individual consumption tax of the kind discussed in Box 1. The claim of large excess burden, however, is not as theoretically clear-cut as was once thought.¹⁸ While views continue to differ, it is now at least widely felt that some taxation of the normal return may well be justified on efficiency grounds, which implies that consumption taxation alone is not enough—though quite how high the rate should be remains contentious.

The second fundamental difference is in terms of rate structure: income taxation can be made progressive in a way that commodity-based consumption taxation cannot, because it can be shaped to overall circumstances of the taxpayer whereas taxation based only on the fact of purchase cannot. As a result, the best design of such a commodity-based consumption tax has to be seen in connection with the design of the PIT and social transfers, in ways to which we now turn.

2.2 Efficiency in the taxation of final consumption

A central and often controversial choice that arises with any commodity-based consumption tax is the rate structure: whether to charge all commodities at the same proportional rate (a ‘uniform’ structure) or at different rates.¹⁹

To begin, we set aside distributional issues and assume that all items of final consumption can be taxed at whatever rate may be desired. The latter means that production efficiency will be desirable, and the design problem is that of setting tax rates on the various items of consumption that are efficient in the sense of minimizing the excess burden of the kind seen in Section 2.1 to arise from the taxation of final consumption.

This collateral damage from taxing consumption comes, potentially, from two sets of distortions in relative prices: between different items of consumption, to the extent that these are taxed at different proportionate rates (so that their relative post-tax prices differ from their relative pre-tax prices) and between consumption and ‘leisure’ (meaning time spent doing something other than earning money). The consumer-leisure distortion means, for example, that taxes on consumption make it more attractive, at the margin, to undertake home repairs oneself, rather than work more in order to hire a contractor whose services will be taxed, or to spend

¹⁷ To see this, note that, when only final consumption is taxed, the present value of the reduction in current tax liability from consuming less in order to save more now and the increased future liability when consuming the proceeds is zero if the saving just earns the return required by the investor.

¹⁸ See Straub and Werning (2020). The production efficiency result does not apply here since it is a receipt of income by final consumers that is at issue. Beyond the present efficiency concerns, however, there may be a distinct argument case for taxing the normal return if, for instance, the better off have high capital income not simply because they can afford to save more but because they tend to have a higher propensity to save (Saez, 2002).

¹⁹ The account that follows focuses on intuition more than precision; for the underlying formalities, see for instance Hindriks and Myles (2013) and the references in Crawford, Keen, and Smith (2010).

more time cooking food at home rather than working to earn enough to eat out.²⁰

This distortion against 'leisure' can be eased by charging a somewhat lower tax on items that make it more attractive to undertake paid work—such as childcare services, or commuting charges—and somewhat higher ones on those associated with leisure—such as downloaded movies or, possibly, alcohol.²¹ Such differentiation induces a distortion of the first kind just mentioned, in relative prices across commodities, but by mitigating the distortion against paid work can in principle reduce excess burden while raising the same amount of revenue. This then is a general principle of commodity-based consumption taxation: Items that are relatively strong complements with 'leisure' (so that an increase in their price reduces the amount of leisure chosen) should be taxed at higher rates than those that are relatively strong substitutes for leisure.

If demands for each commodity are independent of the prices of all other commodities, this principle reduces to an 'inverse elasticity' rule: the tax rate should be higher on items in more inelastic demand (meaning that demand for them is less sensitive to price). Though sometimes invoked as a fundamental principle in itself, this rule is a special case of the general principle of the preceding paragraph: inelastic demand for some item means that an increase in its price leads to a relatively small reduction in spending on it, and if commodity demands are independent, so that there is no fall in spending on other items, the only way to pay for that is by spending more time in paid work. That is, a relatively low elasticity means relatively high complementarity with leisure, and so points to a relatively high tax rate.

Allowing for distributional concerns, another principle comes into play. This is that the tax structure should be such as to reduce the total consumption of some commodity by a greater proportion the more strongly it is concentrated among the well-off²² and the more that policymakers care about alleviating inequality. The intuition is straightforward: all else equal, equity calls for lower consumption tax rates (or even subsidies) on items that are important to the poor. Such distributional concerns then need to be considered alongside the efficiency considerations set out above. There are, however, limits to the effective redistribution that can be achieved through differential rates of consumption taxation alone, since, for example, the rich also benefit from low tax rates on, say, foods. What matters for the distributional power of rate differentiation is the extent of dissimilarity in the proportions of their total consumption that rich and poor spend on particular items. The more similar they are, the less rate differentiation can achieve; and indeed simple calculations often suggest that surprisingly little can be achieved through differential consumption taxation.²³

Governments, however, commonly have instruments other than consumption taxes by which they can address their distributional concerns. These include not only the personal income tax, but other measures of income support; whether income-related (meaning not just the personal income tax but such measures as earned

²⁰ This does not mean that the overall effect is to decrease paid work. To the extent that the tax makes individuals worse off, they can afford less leisure and so may well work more; but this (income) effect is an unavoidable consequence of any form of tax levied on them, so does not give rise to excess burden.

²¹ Piggott and Whalley (2001), for instance, show how bringing into tax services that are a close substitute for self-supply can reduce welfare. Crawford, Keen and Smith (2010) provide a brief review of other work exploring in more detail the tax implications of how 'leisure' time is really spent.

²² This is in contrast to another property of the tax structure when there are no equity concerns, which is that—leaving aside the income effect mentioned in the footnote above, it reduces the demands for all commodities by roughly the same proportion.

²³ See for instance Sah (1983) and Box 7.4 of Ebrill et al. (2001).

income tax credits²⁴ of the kind provided in the US and several other advanced economies) or not (such as categorical benefits to particular types of households, or even poll subsidies payable to all). Since such instruments can take fuller account of individuals' wider circumstances than can a tax on some item of consumption (which treats everybody the same), they are generally better suited to addressing fairness concerns. As a general principle, the point can be simply put: the richer is the available set of instruments that can be tailored to equity concerns, the weaker is the case for differentiating consumption tax rates on distributional grounds. In the limit, when there is full flexibility in taxing or subsidizing earnings the distributional argument for rate differentiation vanishes: The only case for departing from uniform taxation is then the efficiency one, hinging on differing degrees of substitutability with leisure.²⁵

There can thus be good reasons to differentiate tax rates across consumption items, on efficiency and—in the absence of instruments by which to pursue them more effectively—on distributional grounds too. That said, however, while the efficiency argument for structuring the tax system to bear less heavily on relative substitutes for leisure is clear, the practical challenges of inferring from estimated demand systems the precise pattern of rates needed to achieve this are extremely challenging, even in the advanced countries for which demand systems have been most extensively studied.²⁶ One is then left with little more than speculation as a basis for differentiating rates on efficiency grounds. Given, too, the administrative and compliance costs implied by varying tax rates across commodities, the cases in which departing from uniformity on efficiency grounds is clearly warranted are likely to be rare. Even for the commodities often asserted as being suitable for heavy taxation on elasticity grounds—such as tobacco and alcohol—the empirical evidence is by no means always compelling, and rationales in terms of externalities and/or internalities may be more compelling.

Equity arguments for rate differentiation, moreover, are evidently weaker the more fully developed are other instruments for pursuing distributional policies. And in advanced economies they can be very weak. In the UK, for instance, it is widely recognized that food could be brought fully into tax with the poor fully protected against the consequences while leaving a substantial net increase in government revenue.²⁷ The case for equity-based rate differentiation may be stronger in lower income countries, with less sophisticated instruments for supporting the poorest: while it may be, for the reasons above, that relatively little improvement in the circumstances of the least well-off can be achieved through differential commodity taxation, even a small improvement may have considerable social value. Nonetheless, technological developments are rapidly making other instruments for social support increasingly effective in lower income countries, creating options that—through either better targeted transfers or direct VAT rebates to the poorest²⁸—may further limit the appeal of generalized rate differentiation.

There is, in short, no grand theoretical result establishing the universal optimality of uniformity in consumption tax rates, or even that it minimizes excess burden. What theory does suggest, however, is the importance of providing cogent reasons to depart from it. All the above, moreover, abstracts from practical issues of

²⁴ This provides a payment to, rather than by, the taxpayer, if their income is sufficiently low.

²⁵ See for example the characterization of Pareto efficient tax structures in Edwards et al. (1994).

²⁶ Crawford, Keen and Smith (2010) elaborate on this.

²⁷ See for example Crawford, Keen and Smith (2010).

²⁸ The 'personalized VAT' of Barreix et al. (2022) for instance, envisages a transfer to the poorest calibrated to the amount of VAT payable on a typical bundle of purchases; the 'progressive VAT' of de la Feria and Swistak (2024) provides a real time and transaction-related rebate at the time of purchase.

implementation and the distinctive structural features of the VAT as a particular form of consumption tax. It is to these that we now turn.

2.3 Why a VAT?

So far, we have barely mentioned the VAT, but talked rather of commodity-based consumption taxation. This is in order to emphasize that the properties and results highlighted above do not rely on the crediting and refund mechanism that is the defining feature of the VAT. They apply to any form of consumption tax, not least to a Retail Sales Tax (RST), under which tax is levied only on final sales to domestic consumers—which remains the leading and perhaps only alternative to a VAT.²⁹ Understanding what may lend particular appeal to the VAT requires relaxing the assumption above that the design and implementation of the various taxes discussed is perfect and complete.

To this end we look first at two aspects of the distinctive design of the VAT that are intended to counter noncompliance (meaning by that simply the tendency of some taxpayers to evade or perhaps be unaware of their tax obligations). One is that, by its fractional nature, the VAT secures some revenue even if there is noncompliance later in the production chain, notably at retail level—filling, at least partly, gaps in collection that go unfilled under an RST. The second is that it can directly counter noncompliance by propagating incentives to comply along chains of production. We then focus on an imperfection of the VAT that can arise even with perfect implementation: exemption, and its problematic consequences.

Filling gaps

With full and complete implementation, a VAT and RST would be equivalent. The imperfections of implementation that are the sole source of difference between a VAT and—in practice what is most often seen as its main competitor—an RST are of several kinds. From the efficiency perspective, one key difference arises in respect of non-compliance in the form of failure to remit any tax dues on sales, at the final, retail stage.

One common concern with an RST, much emphasized over the years, is that being collected only at the retail stage, it is entirely reliant on the ability to ensure payment by businesses that have long been regarded as hard-to-tax, being in many cases small and—still the case in many developing countries—dealing largely in hard to trace cash. If, for some reason, final sales escape taxation, revenue is zero. An RST puts all the tax administration's eggs in what may be a very flimsy basket. Though this is the aspect generally highlighted, there is another: when some businesses are compliant et al. among their competitors are not, a production inefficiency arises to the extent that the former, even if their pre-tax costs are higher, are able to expand their output at the expense of the latter.

Against this, the argument goes, the VAT has the advantage of not being solely reliant on capturing revenue at the last stage; so one can hope that at least some revenue is raised at earlier stages (with even noncompliant traders perhaps charged on their purchases from larger firms, or on imports). The VAT also eases the

²⁹ The RST is still found not only in the US, at state and local levels, but in 26 other countries (Annacondia, 2018). Malaysia, one of the few countries that has removed a VAT without (yet) reintroducing it, replaced it with an RST.

production inefficiency between the compliant and noncompliant, since by taxing them on their inputs it reduces the competitive advantage that the noncompliant would otherwise enjoy.

Precisely because it taxes the inputs of the noncompliant, however, the VAT potentially imposes an additional source of production inefficiency, of exactly the kind discussed in connection with turnover taxation above—an inefficiency that does not arise with an RST.³⁰ This is not an inherently fatal flaw. Among the conditions needed for the desirability of production efficiency noted above is that the government be unrestricted in its ability to tax all final sales: when it cannot—precisely the situation at issue here—some production inefficiency is generally appropriate.³¹ The RST, that is, can be improved upon by introducing some production inefficiency. This does not mean, however, that any form of production inefficiency is then desirable, so that it is not entirely clear cut that this imperfectly implemented VAT is more efficient than the imperfectly implemented RST. Simple models suggest, however, that—unless substitutability across inputs in the production decisions of the noncompliant is very high, so that the costs of production inefficiency are very large—the VAT will, in such contexts, commonly dominate an RST,

Complicating these arguments, however, is another feature of RSTs often observed in practice: that while they may thus under-tax the right thing (final sales) they may also over-tax the wrong thing: business inputs. In both the US and Canada, about 40 percent of the revenue from state/provincial RSTs has been estimated as coming from business purchases.³² This is a potential source of production inefficiency, just as with the VAT above. But there is a difference. The input taxation that arises under the RST is essentially haphazard in that it is not conditional on noncompliance at a later stage. That under the VAT is systematic form: input tax will remain exactly and only when output tax is not being properly remitted. The latter seems more likely to be helpful.

The balance of advantage between an RST and VAT, in efficiency as in other terms, will of course depend on the precise design and implementation of each. An old rule of thumb, drawn from experience, has it that the difficulties discussed above make the RST highly problematic at rates in the order of 10 percent (and most US rates, for example, are below this). It may be, as Slemrod and Velayudhan (2022) suggest, that technological advances over the last decade or two (some of them touched on below) make it easier to monitor retailers. That though also applies to the VAT, so that the arguments above suggest strong advantages to the VAT when, as in most countries, the ambition is to tax consumption at substantially higher rates than that.

Building chains

The second aspect of the VAT much highlighted as combating noncompliance is that it builds in an element of self-enforcement, by establishing a conflict of interest between seller and buyer: the seller wants to minimize reported sales in order to minimize output VAT while the buyer wants to maximize the reported purchases from that seller in order to maximize their credit or refund for input VAT. This conflict manifests itself in two, related

³⁰ The argument in this paragraph is spelled out in Keen (forthcoming).

³¹ As explored in Newbery (1986). Intuitively, the production efficiency result fails in this case because when the government cannot control consumer prices it cannot necessarily make good use of the additional output that eliminating production inefficiency would imply.

³² The estimates, from Ring (1999) and Phillips and Ibaid (2019) for the US and Smart and Bird (2009 for Canada, are remarkably close, at 41, 42 and 43 percent respectively.

but distinct ways. One is through the trail of evidence—now increasingly electronic rather than on paper—provided by the invoices that pass between them, with the information provided by one side effectively serving as third party information by which the government can check the declaration of the other. The availability of that information gives each side an enhanced incentive to report truthfully, so long as there is enough chance that audit (or, in the more advanced administrations, automatic cross-checking) will detect falsehood and lead to meaningful penalties. The other is through the financial interests involved: the VAT in effect acts as a creditable withholding tax, so that, for example, a noncompliant operator, faced with otherwise unrecoverable input tax, has an incentive to register, declare its purchases and recover the tax paid on its input.

These pressures create the possibility of generating chains of VAT compliance. A business whose customers are VAT-compliant has an incentive to become VAT-compliant too, to recover its own input VAT and so be able to sell more cheaply (the output VAT that it must then charge being recovered by its client). And a business whose suppliers are VAT-compliant has an incentive to become VAT compliant too in order to recover the VAT they will be charged.

There are, however, inherent limits to the effectiveness of these mechanisms. Retailers, most obviously, face no pressures countervailing the incentive to understate sales (the ‘last mile’ problem); and the same is true for sales to exempt businesses. Perhaps more fundamentally, what can be self-enforcing can also prove self-defeating. That is, the same logic which led above to the formation of ‘good’ VAT chains can also lead to the formation of ‘bad’ ones. If a firm’s customers are not VAT-compliant, for instance, then it has an incentive not to be compliant itself, so as to avoid charging VAT that their customers will not recover; and if their suppliers are not charging VAT, there is no input VAT to be recovered by registering.

It is only in the last fifteen years or so that hard evidence, often based on detailed tax records, has begun to emerge on the formation and operation of VAT chains. In both Brazil and West Bengal, there are indeed signs of ‘good’ and ‘bad’ chains co-existing, with traders (not) registered for the VAT systematically more likely to trade with others who are also (not) registered.³³ The element of third party information alone appears to have some effect, with an experiment in Chile finding that firms threatened with audit adjusted their behavior less where the paper trail was strongest, suggesting that awareness that such a trail existed had already induced them to be more truthful.³⁴ And there is evidence of compliance effects propagating through VAT chains. When energy supplies were subjected to a positive VAT rate in Pakistan, for example, manufacturers substantially increased their own reported sales (enabling them to make full use of the increased input tax credit), and some previously unregistered firms appear to have registered for the VAT.³⁵ In Chile, an increased threat of audit delivered to some firms increased the net VAT paid by their suppliers: as the threatened firms moved towards compliance, so it became mutually advantageous for suppliers to remit more creditable output tax.³⁶

It is all too obvious, from the estimates of the compliance gaps (often large, especially in lower income countries) that the VAT is in practice not fully self-enforcing. Nonetheless, elements of self-enforcement are evidently at work, and part of the task of tax administration is to exploit them more fully and effectively. One

³³ See, respectively, de Paula and Scheinkman (2010) and Gadenne, Nandi and Rathelot (2021).

³⁴ Pomeranz (2015).

³⁵ Waseem (2022).

³⁶ The compliance effect also seems to have passed forward, since the net VAT of the customers of alerted firms remained essentially unchanged even though their input VAT presumably increased.

way to do so is by adopting mechanisms intended to generate fuller and firmer third-party information. A wide range of possible measures have been used to this end, including: mandating the use of electronic invoices (these being shared with the tax authorities), electronic fiscal devices of various kinds (storing and making available to the authorities information on sales), lotteries that create an incentive for final consumers to request receipts³⁷ and tax reductions for purchases made by credit card.³⁸ Each of these measures has its limitations.³⁹

A second way to strengthen the elements of self-enforcement is by reinforcing the operation of withholding mechanism. For this, there is some logic in focusing on the retail end of the production chain, so addressing the last mile problem directly and propagating compliance backwards. This though can be problematic in practice given the sheer number and often small scale of retail enterprises. Some of the measures just mentioned can help address this, but ultimately effective discovery and prompt recovery through audit is likely to remain key, and in this area the capacity of many developing countries remains severely hampered. Large firms (and government agencies), relatively easy to monitor, may then have a pivotal role to play. Typically upstream, they are well-positioned to propagate compliance forward, and indeed there is some evidence that effects through this route can be particularly powerful.⁴⁰ Such enterprises can also be leveraged to apply withholding tools that go beyond the standard VAT structure, collecting VAT on their purchases and/or charging an additional creditable withholding tax on their sales; these devices, again not without their difficulties.

The trouble with exemptions

VAT exemptions—under which no tax is charged on sales but, more to the immediate point, tax on inputs is not credited or refunded—may arise *de facto* when chains of compliance are incomplete, or input VAT is not fully refunded on exports (a widespread problem with the VAT). It also seems to be the case, at least in developing countries, that a significant number of firms do not claim the input VAT to which they are entitled.⁴¹ They may also arise *de jure*, whether as a deliberate alternative to full taxation, perhaps in pursuit of equity or other objectives, as a response to practical problems of implementation, as in the taxation of financial services or in order to save administration and compliance costs by excluding firms below some threshold level of turnover from the obligation to remit output VAT and the right to receive credit for input VAT.

Whatever their source exemptions are anathema to the logic of the VAT. This is because they imply the taxation of inputs and so create all the potential production inefficiencies discussed above in relation to turnover taxation.⁴² The analogy to vertical integration as a means of avoiding the taxation of inputs discussed there, for

³⁷ On e-invoicing and electronic fiscal devices, see Hardy et al. (2025); on lotteries, see for example Naritomi (2019).

³⁸ As for example in Korea: see Sung, Awasthi and Lee (2017).

³⁹ Another common theme is that improved information along one dimensions may result in more concealment along others. In Ecuador, for instance, Carillo, Pomeranz and (2017) find (in the context of the corporate tax) that while improved third party information on sales led to higher declared sales the effect was largely offset by increases in reported levels of hard-to-monitor expenses; Slemrod et al. (2017) find a similar effect from the use of third party information on sales from credit card companies in the US.

⁴⁰ As in the results of Waseem (2022) above, and the signs in West Bengal that the tax status of suppliers has a larger effect on registration decisions than that of customers (Gadenne, Nandi and Rathelot (2021), Table 5).

⁴¹ Brockmeyer et al. (2024).

⁴² The Australian usage of 'input-taxed' as an alternative to 'exempt' is a reminder of the point.

instance, is ‘self-supply’; provision of commodities in-house rather than from a third party so as to avoid unrecovered input tax, again with excess burden arising from the additional pre-tax costs this likely implies.⁴³ In open economies, production inefficiency can arise even in the absence of substitution in input use because imports, likely zero-rated (and so freed from input tax) abroad, will be favored over domestic products even if their pre-tax cost is higher. Little is known about the scale of production inefficiencies associated with lingering input taxation under the VAT.⁴⁴ They may in some cases be a price worth paying in pursuit of other objectives, as was argued above may often be the case in filling gaps from noncompliance at retail level. And the inefficiencies from exempting small traders may be outweighed by the saving of implementation costs (and, perhaps, distributional advantages). That they inherently undermine the efficiency properties of the VAT implies, nonetheless, that exemptions be used with extreme caution.

Exemptions carry one further risk. Namely, the problems caused by unrecovered VAT lead to pressures to eliminate (or reduce) it by exempting (or applying a lower rate to) products used as inputs into the exempted activity. But such ‘exemption creep’ is, in effect, an unravelling of the VAT.

2.4 How Large are the Welfare Gains from the VAT?

The arguments above pointing to potential efficiency gains from the VAT are ones of principle. They thus leave open to questions of how large those gains might be, and whether they are realized in practice. It is hard to answer these questions with great precision, not least because they will evidently depend on specifics of the conditions under which the VAT is introduced (what, if anything does it replace?) and design (how extensive are exemptions and the consequent undermining of production efficiency?). It is possible, however, to give a broad sense of the likely orders of magnitude.

For this, it is useful to focus on effects operating through the ‘marginal cost of public funds’ (MCPF), which is the social cost of raising one additional dollar of revenue. This is generally greater than one to an extent that reflects the excess burden of taxation (because that generates a welfare cost without contributing to revenue).⁴⁵

One would thus expect that adoption of the VAT, if indeed more efficient than the instruments previously being used, would reduce the MCPF. And a lower marginal welfare cost of raising revenue would be expected to increase the total amount of tax revenue raised. The converse is also true, which generates a test of whether the VAT has generated an efficiency gain: this will be so if and only if its adoption increases the tax ratio. Empirical findings by Baer et al (2025) suggest that this test has commonly been passed, indicating that, consistent with the underlying theory set out above, the VAT has indeed generally proved efficiency-enhancing.

⁴³ ‘Grouping rules’ may facilitate this by disregarding for VAT purposes transactions between sufficiently closely connected businesses (though commonly not across national borders),

⁴⁴ As with turnover taxation, VAT exemptions also create complexity and non-transparency in effective rate of commodity taxation: see Appendix II of Ebrill et al. (2001).

⁴⁵ Our purpose being very broad brush, the discussion here abstracts from many subtleties in the concept of the MCPF (including the role of income effects and distributional concerns: see for instance Bastani (2024).

Quite how large these gains are, however, depends on quite how much adoption of the VAT reduces the MCPF. To get some handle on this, it can be shown, in an admittedly simple context, that the welfare gain from adoption of the VAT can be approximated as the reduction in the MCPF multiplied by the initial level of tax revenue.⁴⁶ Applying this convenient formula is not straightforward, since the MCPF is not directly observed. But there are some clues to be found.

One comes from the work of Auriol and Warlters (2012) who used computable general modelling to estimate the MCPF for 38 African countries. Among those with a VAT, their median estimated MCPF is 1.17; among those without, it is 1.23.⁴⁷ Supposing, heroically, that the difference of 0.06 is entirely due to the presence of the VAT, the implication is that, at a tax ratio of say 12 percent, adoption of the VAT leads to a welfare gain of 0.72 percent of GDP; at a tax ratio of 15 percent, the gain is 0.9 percent of GDP. These are sizable gains, especially bearing in mind that they are in principle not one-off but arise on a continuing basis.

One last approach is to infer the change in the MCPF from the change in total revenue associated with introduction of a VAT. This requires further strong assumptions, including on the shape of the underlying welfare function—about which little is known.⁴⁸ Suppose, for instance, that the initial value of the MCPF is the 1.21 noted above for non-VAT countries, that the initial revenue ratio is 12 percent, and that adoption of the VAT is associated with an increase in total revenue amounting to 2 percent of GDP. Then, across fairly extreme assumptions on that shape, the reduction in the MCPF might be as low as 0.03 or as high as 0.18, and the welfare gain consequently between 0.32 and 2.1 percent of GDP. Even the lower of these figures, as a recurring gain, is sizable.

2.5 Who Really Pays the VAT?

In both public debate and formal analyses, it is commonplace to assume that changes in VAT rates are fully reflected in the prices paid by final consumers. But this may not be the case: it may be that, given the market condition they face, those setting prices find it in their best interest to pass some (or even all) of that additional tax to their suppliers, including employees, in the form of reduced prices of wages, and/or to swallow a reduction in their own profits. Thus, while the VAT is thought and spoken of as a tax on consumption, and even though it may indeed be the intention to tax only consumers, the real burden—the ‘incidence’ of the tax—may ultimately not be on final consumers. It may be, wholly or in part, on those working for, supplying or owning the firm that remits the tax.

Understanding how its incidence plays out is critical to almost all aspects of VAT policy. It evidently matters from a distributional perspective: a reduced rate will do little to help the poorest consumers if it does not result in a decrease in the price they face. It matters also for assessing its potential as a tool for macroeconomic policy:

⁴⁶ Suppose welfare W to be of the form $U(C) + V(R)$, where C denotes private consumption and R the level of tax-financed public spending. Consumption is $C = Y - \theta R$, where Y denotes total output, taken as fixed, while $\theta \geq 1$, the MCPF in this setting, captures the idea that, as a result of inefficiencies, taxation may reduce consumption more than one-for-one. With R chosen to maximize $W(R, \theta) = U(Y - \theta R) + V(R)$, an exogenous change in R implies (normalizing the derivative U' to unity) that $dW = -R d\theta$.

⁴⁷ From their Tables 2 and 8.

⁴⁸ In the same setting as the earlier footnote: maximizing W with respect to R defines $R(\theta)$; solving $dR = R'(\theta)d\theta$ for $d\theta$ (and assuming $V'' = 0$) allows the change in the MCPF to be inferred from the change in revenue as $d\theta = -\left(\frac{\theta dR}{R + \frac{1}{\theta A}}\right)$, where $A \equiv -U''/U' \geq 0$. θdR .

cutting the VAT rate will fail to directly stimulate demand if the cut is not passed on as a reduction in consumer prices. Less often noted, its incidence also matters for efficiency in the sense that, as will be seen, alternative outcomes in terms of incidence can have quite different efficiency implications.

The general principle shaping the incidence of any tax is that a larger part of the real burden will be borne by whichever side of the taxed transaction has fewer alternatives to it. In the case of a single competitive market, this translates into the implication that the ‘pass through’ of a tax on some commodity—the extent to which a tax increase is reflected in the price paid by the consumer—will be greater the less elastic is demand for it (reflecting limited possibilities for substituting towards other products) and the more elastic is its supply (reflecting greater ease of shifting productive resources into other activities). There is a hint here, of particular importance for the VAT, that pass through will be greater for broad-based rate increases, which leave consumers with few options for substituting towards more lightly taxed goods, than for narrower ones: it may be greater, in particular, for changes in the standard rate of VAT than for changes in a more narrowly applicable rate. This is not theoretically assured,⁴⁹ and so is ultimately an empirical question. Further considerations arise when one looks beyond standard competitive conditions. Under imperfect competition, incidence depends not simply on elasticities but on how those elasticities vary with the price level: as a result, pass through may—but need not be—more than 100 percent.⁵⁰ There are also cases, not all implausible, in which pass through may even be negative.⁵¹

The last few years have seen massive growth in empirical work on the incidence of the VAT,⁵² comparing price changes across similar or comparable products that are and are not subject to some rate changes. Results remain mixed, but—at some risk of putting more structure on the results than their diversity calls for—suggest some broad tentative conclusions:

- *Changes in the standard rate of VAT seem to be about fully passed on.* This emerges from analyses of experience across a wide range of commodity groups in the eurozone over 1999-2013 and of scanner prices for durable goods in Germany over 2002-13.⁵³ In both cases, notably, most of the adjustments comes not at the time of the rate change but between its announcement and implementation and, especially, after it. Examination of the pricing of around 300 commodity groups in Japan around a 2014 increase in the standard VAT rate also points to something like full pass through on average, but with significant variation across products and in this case very largely around the time of the reform.⁵⁴ Again for

⁴⁹ One source of ambiguity arises from feedback effects of substitution away from the taxed goods. (For instance, less substitution of this kind implies a smaller increase in the demand for and hence prices of alternative untaxed goods; that means a greater reduction in demand for the taxed goods; and hence—through this route—lower pass through. On this, see Benedek et al. (2020).

⁵⁰ The intuition is that under imperfect competition firms’ markups—the excess of price over marginal cost—depends on the elasticity of demand; how that price responds to taxation thus depends on how that elasticity varies with price and how marginal cost varies with output; see for example Delipalla and Keen (1992) and Weyl and Fabinger (2013).

⁵¹ This ‘Edgeworth paradox’ can arise when firms produce many products (Armstrong and Vickers, 2023) and in ‘two-sided’ markets (Kind et al., 2008). As an example of the latter, a publisher may react to a tax on newspapers by reducing its final price in order to increase circulation and thereby attract more paid advertising.

⁵² There is work too on the incidence of sales taxation in the US and of excises, which for brevity is not reviewed here: a recent overview is in Lyssiotou and Savva (2021).

⁵³ Respectively, Benedek et al. (2020) for the eurozone and Buettner and Madzharova (2021) for Germany.

⁵⁴ Shiraishi (2020). The impression of full pass through comes from combining the estimates for the months before and after reform in Table 2 there.

Germany, and looking at over one million products, pass through of the 2020 reduction in the standard rate in response to COVID has been put at around 70 percent: somewhat less than complete, perhaps reflecting its explicitly temporary nature.

- *Pass through appears to be noticeably less than full for VAT rates applying to a narrower set of commodities.* Studies of rate cuts on car sales in France, haircuts in Finland and sit-down restaurant meals in France and Sweden, for example, find pass through of ultimately only (respectively) about 55, 50 and 14-25 percent respectively.⁵⁵ For the eurozone between 1999 and 2013, average pass through of changes in the reduced rate has been put at around 30 percent.⁵⁶ Not all studies reach conclusions of this kind: in Cyprus, an increase in the VAT on basic foods appears to have been close to fully passed through, as was a temporary reduction (and subsequent reversal) in the VAT rate on food in Portugal.⁵⁷ Policy makers clearly need to look to the specifics of their own circumstances. In the broadest terms, nonetheless, the hint above as to the likely difference between broad and narrow VAT changes finds some support in the data.
- The point is of obvious importance to the distributional case for setting reduced rates on items important to the less well-off: this is clearly weaker (even leaving aside other limitations of this approach mentioned in Section 2.4 above) if it is others than consumers who receive much of the benefit.
- *Responses may be asymmetric and heterogenous.* Again for the EU, there is evidence that pass through—generally presumed to be the same for rate increases and decreases—is substantially larger for increases than for decreases, though the explanation for this is unclear.⁵⁸ Also striking is evidence that the average pass through may reflect quite different responses by different types of firms and quite different circumstances. A rate cut on restaurants in Finland, for instance, produced an ‘all-or-nothing’ pattern, with most independent restaurants keeping their prices unchanged while chains were more likely to fully pass on the rate cut; a difference that, again somewhat puzzlingly, persisted into the long term.⁵⁹ In Japan, pass through appears to have been systematically lower for goods in more elastic demand, consistent with theoretical predictions. In Germany, the same temporary VAT rate cut in 2020 implied different degrees of pass-through of the same product in different regions with different market structures.⁶⁰

Some very broad empirical regularities are thus beginning to emerge,⁶¹ but accompanied by puzzles and with a diversity of estimates which indicates that incidence is likely to be highly context specific. Importantly, moreover, these results are all for developed countries. One can speculate on how matters might differ for developing countries: high dependence on imports whose prices are effectively fixed on world markets, for

⁵⁵ Respectively Carbonnier (2007), Kosonen (2015), Benzarti and Carloni (2019), and Harju et al. (2018). Carbonnier (2007) also finds, noticeably above the other estimates, 77 percent pass through for home repair services in France.

⁵⁶ Benedek et al. (2020).

⁵⁷ See, respectively, Lyssiotou and Savva (2021) and Bernardino et al. (2024).

⁵⁸ Benzarti et al. (2020).

⁵⁹ Harju et al. (2018).

⁶⁰ Barczay et al. (forthcoming).

⁶¹ Including also that there seem to be no instances of pass-through exceeding 100 percent for VAT rate changes, though there are such for some excises (as in Delipalla and O'Donnell, 2001).

example, may suggest high pass through, as perhaps would significant concentration in the retail sector; ease of moving into and entering informality, on the other hand, might point to low pass through. But paucity of evidence means that little is known about the incidence of the VAT outside developed countries; this is one of the greatest gaps in current understanding of the VAT.

Assessing the extent of pass through into consumer prices is, however, only part of understanding the incidence of the VAT: If pass through is not complete, the question arises as to who it is that bears its real burden. Value added, being the difference between all receipts and all input costs, including wages and investment spending, is the sum of wages and profits defined on a cash flow basis, with the latter corresponding to ‘rents’: payments in excess of the minimum required by investors.⁶² The two leading candidates to bear some burden of the VAT when pass through is incomplete are thus wages and rent. And that brings us back to a point noted but not explained above: that incidence can matter for efficiency.

To see this, note first that to the extent that the VAT is passed backwards in the form of reduced wages this will have much the same qualitative effects, including in distorting labor supply decisions, as an increase in consumer prices (though with some difference across the generations, with a lesser impact on those living on their savings). On the other hand, to the extent that the burden falls on rents there will be no distortion: being an excess over the minimum required by investors, rents can in principle be taxed at any rate without affecting those decisions.

Again, however, very little is known about where the incidence of the tax falls when pass through to consumer prices is incomplete. The one study that focuses on this, looking at changes in the rate on a narrowly defined product, finds that much of the benefit from a rate reduction does indeed go to owners, as increased profits, rather than to workers as higher wages or increased employment.⁶³

One aspect of VAT incidence that has attracted particular attention is the role of the VAT threshold. If traders below the threshold retain their sale prices unchanged when the VAT rate increases, then their customers suffer no burden from the rate increase. Since the poor may indeed be more likely to buy from small retailers, this effect may make the VAT significantly less harmful for the poor than might otherwise be supposed.⁶⁴ If, however, small retailers raise their prices when the VAT rate increases, even though not required themselves to charge that higher rate, then they enjoy an increase in real income from an increase in the VAT: for them, the real burden is then in effect negative. We know, however, of no evidence on which of these outcomes has been observed in practice.

3. Beyond Revenue

So great is the revenue raised by the VAT that its implications for—and possible use to pursue—policy objectives beyond the raising of tax revenue add further to its importance. Turning to these, this section

⁶² With the full cost of investment deducted, tax is levied, in present value (and from the perspective of the investor) only to the extent that the return on the investment exceeds the investors’ required return. See for instance Chapter 2 of Devereux et al. (2021).

⁶³ Benzarti and Carloni (2019) and Harju et al. (2018).

⁶⁴ See Jenkins and Jenkins (2006) and Bachas et al. (2023). By the same token, of course, they also benefit less from rate cuts than might otherwise be supposed.

considers the use of the VAT as a tool for macroeconomic stabilization, its impact on long-term growth and trade performance, and its potential to address environmental, health and industrial policy concerns.

3.1 The VAT and Short Run Macroeconomic Management

The VAT can be used to serve macroeconomic objectives in several ways.

Keynesian use of the VAT

The VAT can be used as part of a standard Keynesian counter-cyclical fiscal policy: raising the VAT during a boom to cool the economy down or reducing it during a bust so as to stimulate aggregate demand and hence the overall level of activity. As a broad-based tax, a VAT rate cut, for example, would be expected, to the extent that it is passed through to consumers, to generate an immediate and sizable increase in real incomes that allows them to buy more goods and services. The strength of this stimulus will depend on the proportion of the tax saving that households spend rather than save. In this respect, a VAT cut might be particularly effective since it directly reaches all households, including those with low and middle incomes who typically have a relatively higher marginal propensity to consume. It may be more effective, than reducing personal income tax, for instance, since many of those on the lowest incomes may in any case have little liability—indeed, especially in developing countries, they may (legally or otherwise) have none. Also important for its effectiveness of a VAT reduction in stimulating activity is the extent to which households spend on imports or services performed abroad (such as tourism), since lower net exports would offset increases in domestic production, thus reducing the stimulus impact of a VAT cut on GDP. The latter concern is particularly relevant for developing countries, which often rely heavily on imports of consumer goods.

The overall effect of a VAT change on aggregate demand can be measured by the ‘tax multiplier’: the dollar increase in GDP resulting from a dollar reduction in revenue. This critically depends on the design and timing of the stimulus. For instance, tax cuts are in general more effective during recessionary times when there are significant slack resources in the economy. Estimates suggest that tax multipliers can significantly exceed 1, especially during recessions; and that—reflecting differences in marginal propensities to consume—they tend to be larger for tax cuts targeted at low- and middle-income earners, rather than high-income earners. One estimate puts the VAT multiplier during fiscal consolidation episodes in advanced economies at 3: that is, a VAT cut amounting to 1 percent of GDP raises GDP in the first year by 3 percent. The VAT multiplier has also been found to be larger in countries where the VAT rate is initially higher—notably for advanced European countries. For the UK, there are signs that VAT multipliers are smaller than income tax multipliers and that shifting the burden of taxation from income to consumption would be expansionary.⁶⁵

It may be tempting to stimulate aggregate demand through selective VAT rate reductions or exemptions. For instance, VAT cuts might focus on commodities that are heavily consumed by the poor, given their higher marginal propensity to consume. But a selective approach can also bring challenges associated with

⁶⁵ The studies referred to in this paragraph are, respectively, Riera-Crichton et al. (2016), Gunter et al. (2018) and Nguyen et al. (2021) and Dabla-Norris and Lima (2023).

complexities in the VAT and distortions in consumer choices—along with the possibility highlighted above that selective rate cuts will be less than fully passed on. The experience of the EU from the early 2000s, where the VAT rate on certain labor intensive services was reduced by several countries to boost demand and create jobs for low-skilled workers, was largely disappointing: the reduced rates were mainly used to increase the margins of service providers, and even where the reduction was passed on to consumers, it was only temporary, as prices subsequently increased.⁶⁶ VAT cuts during COVID on heavily affected sectors, such as accommodation and restaurants, were largely ineffective in boosting output due to supply constraints (Blundell et al. 2020). A tax cut in such circumstances becomes to some degree a de facto form of income support to providers in the favored sectors, which might itself have served distributional concerns and/or prevented further reductions in activity.

Somewhat more subtly, governments may also use measures related to the administration of the VAT as, in effect, a form of stimulus during severe adverse economic shocks. To relieve liquidity problems for firms during COVID-19, for instance, many governments extended filing and payment deadlines, relaxed the terms and conditions for paying tax arrears, and accelerated VAT refund payments. Indeed, even in the absence of any policy measures, compliance with the VAT may operate as an automatic stabilizer. For instance, during recessions traders may be cash constrained and seek to reduce VAT payments by underreporting sales and delaying or even failing to make payments. One estimate, for a panel of 27 European countries, is that in the short run, a 1 percentage point reduction of actual output below potential immediately increases the VAT compliance gap by 0.4 percent.⁶⁷ The concern, of course, is for the government to ensure that compliance is quickly restored as the economy recovers.

Inducing intertemporal substitution

An explicitly time-bound VAT reduction may also induce an intertemporal substitution effect by creating incentives for households to bring their consumption forward, drawing on their savings, saving less or even borrowing in order to buy durable and/or storable goods and services and so benefit from the temporarily lower tax rate. When several countries reduced their VAT in the aftermath of the global financial crisis and during COVID-19, many of them did so explicitly temporarily in order to have such an effect. These VAT cuts had some effects. For instance, the 2.5 percentage point VAT rate reduction in the UK (over 13 months during 2008-09) generated a 0.4 percent increase in total consumption largely on account of a substitution rather than an income effect.⁶⁸ In Germany, the temporary VAT cut by 3 percentage points in the second half of 2020 significantly raised aggregate consumption.⁶⁹ While a temporary VAT rate cut increases consumption in the short run, by the same token retail sales are likely to fall after the cut ends.⁷⁰

An intertemporal substitution effect can also be achieved by announcing a future VAT rate increase, rather than temporarily reducing the current level of VAT. This may be an unintended consequence of pre-announced rate increases, or it might be deliberate policy: A continuing process of small increases in the VAT rate (combined

⁶⁶ EC (2003).

⁶⁷ Keen (2016).

⁶⁸ Crossley et al. (2014).

⁶⁹ Bachmann et al. (2021).

⁷⁰ Crossley et al. (2014) and Misch and Seymen (2012).

with reductions in the rate of income tax) was proposed, for example, as one way to increase activity and strengthen inflationary expectations in Japan during the 2000s.⁷¹

There is substantial evidence of such intertemporal substitution at work. Many European VAT reforms in the early 2010s triggered increases in the sales of durable goods in the period immediately before the VAT increase. Buettner and Madzharova (2021) find that a VAT increase of one percentage point in several European countries triggered an average increase in sales in the month before the VAT increase of around 2.5 percent. In Germany, the 3-percentage point VAT increase in 2007 announced 13 months in advance led to a sizable increase in consumption, especially immediately before the rate increase: consumption in the last quarter of 2006 increased by 2.2 percent year-over-year.⁷² Similarly, in Japan, each announcement of a VAT rate increase (from 3 to 5 percent in 1997, and from 5 to 8 percent in 2014) led to an increase in private consumption in the period leading up to the rate increase, especially the purchases of durables. This intertemporal shifting again implies a contraction of consumption after the rate increase occurs. For example, the Japanese VAT increase from 5 to 8 percent in 2014 was associated with a decline in consumption in the first quarter after the reform by 5.8 percent.⁷³

Fiscal devaluation

The VAT can also be used as part of a revenue-neutral package to mimic the effects of a devaluation, at least for some time. This possibility attracted particular attention during the early 2010s in several members of the eurozone, including Greece and Portugal, who faced the combination of severe fiscal constraints and painful downturns in activity that they were unable to counteract by a nominal devaluation.

This idea was to accompany an increase in the VAT with such a reduction in employers' social security contributions as to leave revenue broadly unchanged. With the nominal wage rates paid to workers (net of those contributions) fixed in the short term, the reduction in social contributions would be expected to reduce domestic producer prices, including those of exported goods and services. In contrast, the higher VAT rate bears on domestic consumption, but not on exports. Hence, it offsets the impact on the consumer prices of domestically produced commodities of the reduction in domestic wage costs while also increasing the consumer price of imports. Foreign demand for exports increases and domestic demand for imports falls; consequently, the current account improves—as it would with a depreciation of the real exchange rate. Crucially, moreover, if labor is initially in excess supply, the increase in the demand for labor generated by the reduction in employers' wage costs leads to a reduction in unemployment.

The effectiveness of this strategy requires rigidity in both the exchange rate and the nominal wage exclusive of the employers' social contribution. With a flexible exchange rate, the increased demand for exports and reduced demand for imports would, as above, cause an appreciation of the nominal exchange rate that undoes the competitiveness impact of the tax shift. Even if the exchange rate is fixed, a fiscal devaluation will also have no real effect if—or when—domestic wages adjust. As workers find their real wage reduced by the increase in the domestic price of imports, they (or their trade unions) will aim to increase their nominal wages, moving the

⁷¹ Feldstein (2002).

⁷² Carare and Danninger (2008).

⁷³ On the Japanese experience, see Cashin and Unayama (2016), Lubik and Rhodes (2019) and IMF (2018).

wage back towards its pre-reform equilibrium, with no long run impact on trade or labor markets. This adjustment of nominal wages, however, may be expected to take some while: in the meantime, employment was expected to increase and the current account to improve.

Even though temporary, the effects of a fiscal devaluation may be important, especially when the economy is initially in a disequilibrium position with an overvalued real exchange rate and involuntary unemployment. In these conditions, a fiscal devaluation can speed up the adjustment in the labor market, which may otherwise take a long time to implement. This faster speed of adjustment is critical in countries where doubts may otherwise arise on the sustainability of the adjustment process under a pegged exchange rate.

De Mooij and Keen (2013) explore the impact of employer social contributions and VAT on net exports in OECD countries. They find that the balance between the two does indeed matter in the Eurozone countries (with a fixed exchange rate) and that a fiscal devaluation has significant short-run implications for trade: a tax shift from employer's social contributions to VAT equivalent to 1 percent of GDP is estimated to generate a short-run increase in net exports of about 0.9 percent of GDP. As the arguments above suggest would be the case, this effect disappears in the long run (though, as theory also predicts, less rapidly than in countries without a fixed exchange rate): the case for the fiscal devaluation thus rests on its potential to accelerate adjustment to deeper underlying problems, not as a permanent resolution of them.⁷⁴

Inflation

The concern is sometimes heard that an increase in the rate of VAT will lead to inflation: a continuing increase in the level of consumer prices.⁷⁵ In principle, however, a one-off increase in the VAT rate should lead simply to a one-off increase in the level of consumer prices (possibly with less than full pass through, as discussed above), with no impact on the continuing rate of structural inflation. It might perhaps be feared, nonetheless, that even a one-off increase would somehow create expectations—which might, through wage and price adjustments, prove self-fulfilling—of a continuing increase in the price level.


In practice, it is hard to find signs of this happening. Case studies commonly find no such lasting effect on inflation.⁷⁶ More systematically, a scatterplot of VAT rates changes in year t (compared to $t-1$) and the change in inflation in years before and after the VAT rate change (year $t-1$ and $t+1$) fails to mark a significant correlation (Figure 2). While statistically insignificant, the downward slope of the regression line is consistent with some

⁷⁴ A somewhat similar issue is raised by the concern, expressed most notably by some US observers that the remission of VAT on exports and inclusion of VAT on imports by European countries places them at a competitive advantage in world markets relative to countries that are more reliant on the corporate income tax (CIT; which, like social contributions, is not remitted on exports). Keen and Syed (2006) explore the trade implications of both CIT and VAT, their empirical finding being that while the balance between them matters significantly for the trade balance in the short run this effect too vanishes in the long run.

⁷⁵ A different concern that sometime arises at the time of VAT increases (or introduction) is that of 'gouging': prices rising by more than the additional tax. Governments sometimes establish oversight bodies to monitor this (or at least serve as a deterrent and signal of political concern, it being impossible to monitor all prices), as at the time of the 2014 increase in Japan. Ultimately, however, market forces can be expected to dominate, and as seen above over-shifting may indeed arise for some products in circumstances of imperfect competition.

⁷⁶ One example being Bolton and Dollery (2005).

anticipation of the rate change at t being reflected in prices at $t-1$, along lines encountered in Section 2.5.⁷⁷



3.2 The VAT and Long Run Economic Performance

Beyond an essentially counter-cyclical role, the VAT may have lasting effects on international trade and long-term growth.

⁷⁷ For a panel of 10 advanced economies, Gonçalves et al. (2025) find a persistent effect on inflation of unanticipated VAT shocks in the rather different sense of unanticipated increases in VAT revenues, an effect that can last for between 2 and 3 years after the VAT reform. Such VAT shocks, measured as the "intended VAT revenues in t over GDP in $t-1$ ", where revenues are forecasts produced by country authorities when the VAT rate or base measures were announced, can arise from sources other than a rate change, and will exclude effects from anticipated rate changes.

Trade

How the VAT might affect trade is of natural interest to policy makers. On this, the benchmark result is one of neutrality: once all prices have adjusted, a VAT (or consumption tax more generally) that is levied at a uniform rate and perfectly enforced has no effect on the level or composition of trade.

To see this at its simplest, suppose that labor is perfectly inelastic in supply and fully employed and imagine that the rate of VAT is increased. Taking first the case in which the exchange rate and the tax-exclusive prices received by domestic producers are all fixed in nominal terms, the effect is then to increase the level of consumer prices. Real wages consequently fall; but the assumptions on the labor market mean that employment and output do not change. Returning the revenue raised by the VAT to consumers as a lump sum enables them to nevertheless maintain the same level of consumption. And, since import and export prices are unchanged in both domestic and foreign currencies, trade is also unaffected. The same trade-neutrality also follows if instead the nominal exchange rate and producer prices are flexible while consumer prices are fixed. The increased VAT then causes producer prices, including the wage to fall by the same proportion while consumer prices remain unchanged. This again has no effect in the labor market and the return of revenue enables the same level of consumption to be attained. Now, however, the fall in producer prices means that exports have a lower foreign currency price. The increased attractiveness of the country's exports in world markets that results from this will lead to an appreciation of the currency, reducing consumer prices in domestic currency so as to restore the real wage to its initial level. Once again, trade and all real variables will ultimately be unchanged.⁷⁸

This property of trade neutrality rests of the presumption that the VAT is applied at a uniform rate (otherwise, for example, changes affecting relative consumer prices cannot be undone by changing the single instrument of the exchange rate). In practice, however, rate differentiation is commonplace. More specifically, Feldstein and Krugman (1990) note that tradables are often taxed more heavily under the VAT than are non-tradables: for example, several EU Member States currently make use of special provisions enabling the application of reduced rates to specific non-tradable labor-intensive services. In these circumstances, a higher standard VAT rate will reduce the relative consumer price of non-tradables, encouraging substitution of domestic production out of tradables and lead a reduced volume of trade.

A second common departure from the required condition is that, especially in developing countries, VAT refunds on inputs used by exporters are paid, if at all, with a significant delay: Pessoa et al. (2021), for instance, show that refunds in low-income countries are on average less than 10 percent of gross VAT receipts, while this ratio exceeds 30 percent in high-income countries (a finding in line with those of Harrison and Krellove (2005) and Keen (2007)). In such cases the VAT is levied not only on consumption but also acts in part as an export tax and so would be expected to reduce the size of the tradables sector and to reduce export intensity.

⁷⁸ This account abstracts from a variety of considerations addressed in the literature. When labor supply is not fixed, for example, since the reduction in the real wage that arises in both cases will affect the labor market equilibrium; trade neutrality then requires the VAT increase to be accompanied by an equi-proportion reduction in labor taxes. When adjustment is partly through the exchange rate, neutrality also requires that citizens hold no assets denominated in a currency other than their own. Recent statements of trade neutrality propositions are in Benzarti and Tazhitdinova (2021) and Schneider, Stähler, and Thuncke (2023) (both for fixed labor supply) and, perhaps most generally, Farhi, Gopinath, and Itskhoki (2014).

Whether such imperfections and other failures of the neutrality conditions mean that the VAT has an impact on trade volumes—most considerations suggest, if anything, a negative one—is an empirical question.

Several studies do suggest that trade impacts of the VAT might indeed be sizable. For a panel of over 100 countries, Sharma (2020) finds that VAT adoption has tended to reduce exports from developing countries by an amount that is larger the greater is the use of intermediate inputs, with no such effect in advanced economies—results that are consistent with an impact from less effective refunding there. An earlier study by Desia and Hines (2002) also found a negative effect, which was noticeably more pronounced in lower income countries. Focusing on experience in P. R. China, where incomplete refunding has been used as a matter of policy, Chandra and Long (2013) find that every USD 1 of rebate that is not refunded to exporters reduces exports by USD 4.70; this effect is more pronounced for Chinese firms facing financial constraints (Lu and Ma, 2023). Studies focused on advanced economies generally find little if any effect. For the EU, Benzarti and Tazhitdinova (2021) find essentially no impact.⁷⁹ Running counter to an emerging conventional wisdom, however, Schneider et al. (2023), also looking at the EU experience but using a very different methodology find a strong negative effect: a one percentage point VAT increase is estimated to reduce aggregate imports and internal trade by 3.1 percent. Whether trade neutrality of the VAT is a reasonable approximation where refund procedures are adequate thus remains, for now, unclear.

Long-run growth

The effect of the VAT on long-term growth cannot be assessed in isolation but depends on how the revenue that it raises is spent or/and on which alternative sources of revenue it replaces. Pursuing the latter perspective, there is evidence that consumption taxation (generally comprising VAT and excises) is relatively growth-friendly in the sense that increasing reliance upon it while reducing that on the personal (and, even more, than corporate) income tax has been associated with faster growth.⁸⁰ This finding, however, is for middle- and high-income countries: the relative growth-friendliness of the VAT is less clear for low-income countries.⁸¹

It is very likely, however, that it is not just the revenue that these taxes raise but their design that is critical for its impact on growth. On this there is some evidence that a rise in the VAT, financed by a fall in income taxes, promotes growth only when the additional revenue is raised not by increasing the standard rate of the VAT but by expanding its base.⁸² The implication is that a policy of base broadening and rate reduction in the VAT could spur economic growth. This result needs, however, to be treated with caution. It is not easy to explain as a matter of theory, and as a matter of practice some ways of apparently expanding the base (such as delaying or

⁷⁹ For samples of OECD countries, both Keen and Syed (2006) and de Mooij and Keen (2013) also find no lasting trade effect from the VAT.

⁸⁰ Kneller, Bleaney, and Gemmell (1999) and Arnold et al. (2011).

⁸¹ Acosta-Ormaechea, Sola, and Yoo (2019).

⁸² Acosta-Ormaechea and Morozumia (2019). By an expansion of the base is meant here an increase in 'C-efficiency (the ratio of revenue from the VAT to the product of the standard rate and private consumption).

denying VAT refunds to exporters) seem unlikely to be growth-supporting. It does, nonetheless, lend support to arguments favoring VATs with extensive coverage and limited rate differentiation.

3.3 The VAT in Pursuit of Other Objectives

The VAT is sometimes also called on to serve other purposes of social and economic policy—to which it is in many cases not well-suited.

Environmental and health objectives

A case can be made for a differentially high ‘corrective’ tax on the consumption of a commodity whenever it generates adverse ‘externalities’—damage to others not party to the consumption decision—or is associated with ‘internalities’—problems of self-control. Externalities are often associated with environmental damage of some kind. Examples include the generalized harm that the burning of fossil fuels causes by aggravating global warming, and damage through passive smoking. Internalities are often associated with health issues: the difficulty, for example, that heavy smokers or drinkers are likely to have, when the time comes, in meeting their perfectly honest intention to cut back in the future. The range of social concerns generating calls for some tax-based correction has become increasingly large, extending beyond the traditional areas of smoking, drinking, and burning fuels to include, for example, the consumption of sugars and meat.

What these considerations generally call for, however, is not a differentially high rate of VAT but a specific tax—one specified, that is, as a monetary amount rather than as a proportion of price—on the commodity concerned (or, where the activity is beneficial rather than harmful, a specific subsidy).

Conceptually, there are two reasons for this. First, a specific tax is appropriate where the damage is linked to the amount of the commodity used, not how much is paid for it. The climate harm caused by burning one liter of oil, for example, does not depend on the price paid for it. Most externalities, and environmental harm in particular, do indeed have this feature of being rooted directly in the use of some commodity. For internalities, this is perhaps less clear cut.

The second conceptual point is that when damage arises from the use of the commodity itself it is immaterial whether that use is for final consumption or as an input to production. Burning fossil fuels generates the same CO₂ whether it is being done to power a passenger car or to keep a factory operating. A specific tax, applied on all purchases, will affect all uses of the product, but the crediting mechanism of the VAT (if it is well-functioning) means that a differential rate of VAT will affect only its use for final consumption.⁸³

The VAT should of course continue to apply in addition to any specific tax (or subsidy) levied for corrective purposes, since the revenue-raising motive continues to apply.⁸⁴ To keep these roles distinct, it is generally recognized as good practice that the VAT be applied to the price inclusive of any specific corrective tax: in this

⁸³ Practical considerations also likely favor the use of specific taxation to serve corrective ends, given the difficulties associated with implementing differential VAT rates. That said, establishing specific taxes on the traditional triumvirate (tobacco, alcohol, fuel) can create its own challenges, as does a proliferation of specific taxes; and specific subsidies may lead to further difficulty.

⁸⁴ On precisely how the corrective and revenue concerns interact to shape the appropriate overall commodity tax, see Sandmo (1975) and Pirttilä and Tuomala (1997).

way,⁸⁵ the corrective tax is targeted at the relative price of the concerned commodity while changes in the VAT affect only the general level of consumer prices.

The VAT and industrial policy

There is one use to which the VAT is very poorly suited, but for which it is nonetheless sometimes used: as a tool of industrial policy, to promote the domestic production of particular commodities or sectors by either charging their sales at a reduced (possibly zero) rate or exempting them. This is a problematic policy choice, for somewhat different reasons, whether that product is for final consumption or intermediate use; and it is also the case, again for somewhat different reasons, whether the chosen instrument is a reduced rate or exemption.

Take first the case in which the commodity whose domestic production it is intended to remote is an item of final consumption: electric cars, for example. Then by leading to a lower consumer price, a reduced rate or exemption at the final stage can indeed be expected to increase domestic demand. But while that will tend to increase domestic production of electric cars it will also tend to increase imports. Both VAT instruments are thus poorly targeted: the policy objective would be better achieved by subsidizing domestic production directly, without extending the benefit to foreign producers. Exemption, moreover, induces a potential production inefficiency of the kind discussed above (because of the unrecovered tax on inputs); a reduced rate does not, so long as input tax is fully credited or refunded.

Use of the VAT to pursue industrial policy is even more problematic if the commodity at issue is an intermediate good. If purchasers are VAT-compliant, a reduced rate will simply be ineffective since their input VAT, whatever the rate at which it is charged, will be fully credited or refunded to them. And exemption will be worse than ineffective: by leading to unrecovered input tax, it may actually increase the price that the exempted firms charge their customers—and so achieve the exact opposite of the policy intention. This reduction in demand for the exempted item, moreover, will be amplified to the extent that competing imported products are zero-rated abroad and so experience no such increase in their input costs. The likely adverse effect on the price charged to the purchaser will be lessened if production can be shifted away from the taxed inputs; but that signals a production inefficiency of the kind discussed above.

Reduced rates or exemption can be expected to increase demand for and production of intermediate goods, however, if their purchasers are not VAT-complaint: the situation is then as if the sale were to final consumers. But these measures are then counterproductive in an even deeper sense, since they then have an effect only because they further undermine an already compromised VAT.

That the VAT is not well-suited to manipulate production patterns is no great surprise given that its fundamental attraction stems from taking consumption as its base. And these arguments illustrate another general truth: that

⁸⁵ Assuming here that the incidence of all taxes is entirely on consumers.

reduced rates are, in principle—and practical considerations may point in a different direction⁸⁶—generally preferable to exemption in avoiding production inefficiency.

4. Conclusion

The primary purpose of the VAT is to raise revenue. So effective has it been at this that it is inevitably important in relation to macroeconomic policy and long-term performance. For these, however, the distinctive features of the VAT as a particular form of consumption are largely essential. It is in relation to the revenue objective and more micro-level policies that the architecture of the VAT is critical.

On the latter, this paper ultimately arrives at what might seem a well-known conclusion: that the VAT is, or at least can be, a particularly efficient revenue-raiser. While that is indeed often said to be the case, the reasoning is rarely made precise—and is more subtle than generally recognized.

The precision is in the meaning to be attached to the word ‘efficiency.’ In the context of revenue raising, this can be interpreted, fairly precisely, as the minimization of excess burden. And that, as an important benchmark, requires production efficiency, meaning an absence of distortions to production methods. This, importantly, has nothing to do with rates imposed on final consumption. The application of a uniform rate on final consumption, for example, does not have any particular efficiency merit: minimizing excess burden may well require, in principle, some rate differentiation. It is the treatment of inputs that matters.

The subtlety is that the VAT ceases to guarantee production efficiency whenever the only structural feature which distinguishes it from a retail sales tax—fractional collection—comes into play, since inputs are then taxed. One cannot simultaneously laud the VAT as assuring production efficiency and praise the fractional collection which is its distinctive feature as securing revenue when compliance is incomplete. There is reason to hope that the production inefficiencies associated with an imperfectly implemented VAT are reasonably benevolent and a price worth paying to secure revenue. But there is no guarantee that this will be the case.

That, in turn, is a reminder of the dangers of damage through VAT exemptions, whether de jure as a conscious policy choice or de facto as a result of noncompliance, including in the form of ‘bad’ VAT chains. While in some cases perhaps technically unavoidable, in all cases exemptions undermine the smooth functioning of the VAT. Conversely, the VAT is poorly suited as an instrument to affect production decisions precisely because that is exactly what a well-functioning VAT does not do.

Empirical understanding of how the VAT operates in practice has advanced enormously in recent years, but how much all this ultimately matters for economic performance and consumer welfare remains to some degree an open question. Indeed, this will clearly vary depending on the deism of and circumstances in which any particular VAT operates. The tentative results presented here suggest, nonetheless, that potential efficiency gains from the VAT do indeed exist, have been quite widely realized, and can be sizable.

⁸⁶ Exemption, for instance, avoids the need to pay refunds that may arise from setting a reduced rate; but it also creates a need to apportion input VAT when suppliers also sell taxed items.

It is natural, finally, to ask what this likely means for the future of the VAT. One aspect of the VAT touched on above, is that it is in part a tax on rents. This led Ebrill et al (2001) to suggest that the role of the VAT as an implicit form of corporate tax would attract increasing attention. This has to some extent happened, most clearly with advocacy of a Destination Based Cash Flow Tax—in effect combining a VAT with a wage subsidy—to replace traditional corporate taxes.⁸⁷ Whether the continuing pressure on the international tax system will press reform in this direction remains to be seen. Perhaps more fundamentally, the fractional nature of the VAT has its roots in the weaknesses of and desire to strengthen third party information. As the ability to collect and process this information by other routes grows, so the apparatus built around the idea of fractional collection may come to seem increasingly cumbersome and dispensable. It is worth remembering that if all transactions were observed with complete accuracy, and final sales identifiable, there would be no need for a VAT.

⁸⁷ Devereux et al. (2021).

References

- Acosta-Ormaechea, Santiago and Atsuyoshi Morozumia (2019). “The Value-Added Tax and Growth: Design Matters.” *International Tax and Public Finance*, Vol. 28, p. 1211–1241.
- Acosta-Ormaechea, Santiago, Sergio Sola, and Yiae Yoo (2018). “Tax Composition and Growth: A Broad Cross-country Perspective.” *German Economic Review*, Vol. 20, pp. e70-e106.
- Adandohoin, Kodjo (2021), “Tax Transition in Developing Countries: Do Value Added Tax and Excises Really Work?” *International Economics and Economic Policy*. Vol. 18, pp. 379–424.
- Adam, Stuart, Tim Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, James Mirrlees and Gareth Myles (2011). *Tax by Design* (“The Mirrlees Review”). (Oxford University Press: Oxford).
- Agrawal, David R and Laura Zimmerman (2024). “The Effects of Adopting a Value Added Tax on Firms.” Mimeo, University of Kentucky.
- Ahmad, Ehtisham and Nicholas Stern (1986). “Tax Reform for Pakistan: Overview and Effective Taxes for 1975-76.” *Pakistan Development Review*, Vol. XXV, pp. 43-72.
- Annacondia, Fabiola. 2018. “Overview of General Turnover Taxes and Tax Rates.” *International VAT Monitor*, Vol. 29, pp. 81–91.
- Armstrong, Mark, and John Vickers (2023). “Multiproduct Cost Pass-Through: Edgeworth’s Paradox Revisited,” *Journal of Political Economy*, Vol. 131, pp. 2645 – 2665.
- Arnold, Jens Matthias, Bert Brys, Christopher Heady, Åsa Johansson, Cyrille Schwellnus, Laura Vartia (2011). “Tax Policy for Economic Recovery and Growth, *Economic Journal*, Vol. 121, pp. F59–F80
- Aslam, Aqib, and Maria Coelho (2021). “A Firm Lower Bound: Characteristics and Impact of Corporate Minimum Taxation.” IMF Working Paper 21/161.
- Auriol, Emanuelle, and Michael Warlters (2012). “The Marginal Cost of Public Funds and Tax Reform in Africa,” *Journal of Development Economics*, Vol. 97, pp. 58-72.
- Bachas, Pierre, Lucie Gadenne and Anders Jensen “Informality, Consumption Taxes, and Redistribution,” *Review of Economic Studies*, Vol. 91, pp. [2604-2634](#).
- Baer Katherine, Li Liu and Miguel Almunia (2026). “VAT Effectiveness in Raising Revenue,” Chapter IV in: Artur Swistak, Christophe Waerzeggers and Andrew. Okello (eds). *The Changing World of VAT*. International Monetary Fund, Washington D.C. (forthcoming in 2026)

Barczay, Michael, Shafik Hebous and Tom Zimmermann (forthcoming). The Heterogeneous Incidence of a VAT Cut: Evidence from Germany.

Barreix, Alberto, Martín Bès, Oscar Fonseca, Maria Belén Fontañez, Dalmiro Morán, Emilio Pineda and Jerónimo Roca (2022). “Revisiting Personalized VAT: A Tool for Fiscal Consolidation with Equity,” Inter-American Development Bank

Bastani, Spencer (2024). “The marginal value of public funds: A brief guide and application to tax policy,” *International Tax and Public Finance*, <https://doi.org/10.1007/s10797-024-09860-x>

Baunsgaard, Thomas and Michael Keen (2010). “Tax Revenue and (or?) Trade Liberalization,” *Journal of Public Economics*, Vol. 64, pp. 563–577.

Benedek, Dora, Ruud de Mooij, Michael Keen and Philippe Wingender (2020). “Varieties of VAT Pass Through,” *International Tax and Public Finance*, Vol. 27, pp. 890–930.

Benzarti, Youssef and Dorian Carloni (2019). “Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France.” *American Economic Journal: Economic Policy*, Vol. 11, pp. 38–63.

Benzarti, Youssef and Alisa Tazhitdinova (2021). “Do Value-Added Taxes Affect International Trade Flows? Evidence from 30 Years of Tax Reforms.” *American Economic Journal: Economic Policy*, Vol. 13, pp. 469–489.

Benzarti, Youssef, Dorain Carloni, Jarkko Harju, and Tuomas Kosonen (2020). “What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes.” *Journal of Political Economy*, Vol. 128, pp. 4438–4474.

Bernardino, Tiago, Ricardo Duque Gabriel, João Qeulhas and Márcia Silva-Pereira (2024). “The Full, Persistent, and Symmetric Pass-Through of a Temporary VAT Cut,” Mimeo, IIEES, Stockholm University.

Best, Michael, Anne Brockmeyer, Henrik Kleven, Johannes Spinnewijn and Mazhar Waseem (2015). “Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan.” *Journal of Political Economy*, Vol. 123, pp. 1311–1355.

Boadway, Robin, Maurice Marchand, and Pierre Pestieau (1994), “Towards a Theory of the Direct–Indirect Tax Mix,” *Journal of Public Economics*, Vol.55, pp. 71–88.

Bolton, Tom, and Brian Dollery (2005). “An empirical Note on the Comparative Macroeconomic Effects of the GST in Australia, Canada and New Zealand.” *Economic Papers*, Vol. 24, pp. 50–60.

Bradford, David F., and US Treasury Tax Policy Staff (1977). *Blueprints for Basic Tax Reform* (US Treasury: Washington DC).

Brockmeyer, Anne, Giulia Mascagni, Vedanth Nair, Mazhar Waseem and Miguel Almunia (2024). “Does the Value-Added Tax add value? Lessons using administrative data from a diverse set of countries.” *Journal of Economic Perspectives*, Vol. 38, pp. 107–132.

Buettner, Thiess and Boryana Madzharova (2021). "Unit Sales and Price Effects of Preannounced Consumption Tax Reforms: Micro-level Evidence from European VAT." *American Economic Journal: Economic Policy*, Vol. 13, pp.103–134.

Cagé, Julia and Lucie Gadenne (2018). "Tax Revenues and the Fiscal Cost of Trade Liberalization, 1792–2006." *Explorations in Economic History*, Vol. 70, pp.1-24.

Carare, Alina and Stephan Danninger (2008). "Smoothing and the Modest Effect of VAT in Germany." IMF Working Paper 08/175.

Carbonnier, Clément (2007). "Who Pays Sales Taxes? Evidence from French VAT Reforms, 1987–1999." *Journal of Public Economics*, Vol. 91, pp. 1219–1229

Carillo, Paul, Dina Pomeranz and Monica Singhal (2017). "Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement." *American Economic Journal: Applied Economics*, Vol. 9, pp. 144–64.

Cashin, David, and Takashi Unayama (2016). "Measuring Intertemporal Substitution in Consumption: Evidence from a VAT Increase in Japan." *Review of Economics and Statistics*, Vol. 98, pp. 285–97.

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

Chandra, Piyush and Cheryl Long (2013). "VAT Rebates and Export Performance in China: Firm-level Evidence." *Journal of Public Economics*, Vol. 102, pp. 13-22.

Chen, Zhao, Xian Jiang, Zhikuo Liu, Juan Carlos Suárez Serrato, Daniel Yi Xu (2023), Tax Policy and Lumpy Investment Behaviour: Evidence from China's VAT Reform," *The Review of Economic Studies*, Vol. 90, pp. 634–674.

Crawford, Ian, Michael Keen and Stephen Smith (2010). "VAT and Excises," pp. 275-362 in James Mirrlees et al., *Dimensions of Tax Design: The Mirrlees Review* (Oxford: Oxford University Press for Institute for Fiscal Studies).

Crossley, Thomas, Hamish Low and Cath Sleeman (2014). "Using a Temporary Indirect Tax Cut as a Fiscal Stimulus: Evidence from the UK." Institute for Fiscal Studies Working Paper 14/1g.

Dabla-Norris, Era, and Frederico Lima (2023). "Macroeconomic effects of tax rate and base changes: Evidence from fiscal consolidations," *European Economic Review*, Vol 153.

de Mooij, Ruud and Michael Keen (2013). " 'Fiscal Devaluation' and Fiscal Consolidation: The VAT in Troubled Times," pp. 443-485 in Alberto Alesina and Francesco Giavazzi (eds) , *Fiscal Policy After the Crisis* (Chicago: University of Chicago Press), 2013.

de Feria, Rita, and Artur Swistak (2024). "Designing a Progressive VAT," IMF Working Paper 2024/078.

de Paula, Áureo and José Scheinkman (2010). "Value-Added Taxes, Chain Effects, and Informality," *American Economic Journal: Macroeconomics*, Vol. 2, pp. 195-221.

Delipalla, Sofia and Michael Keen (1992). "The Comparison between Ad Valorem and Specific Taxation under Imperfect Competition." *Journal of Public Economics*, Vol. 49, pp. 351-367.

Delipalla, Sofia and Owen O'Donnell (2001). "Estimating Tax Incidence, Market Power and Market Conduct: The European Cigarette Industry." *International Journal of Industrial Organization*, Vol. 19, pp.885–908.

Desai, Mihir and James R. Hines Jr. (2002). "Value Added Taxes and International Trade: The Evidence," Mimeo, University of Michigan.

Devereux, Michael, Alan Auerbach, Michael Keen, Paul Oosterhuis, Wolfgang Schön and John Vella (2021). *Taxing Profit in a Global Economy*. (Oxford University Press: Oxford).

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

Ebrill, Liam, Michael Keen, Jean-Paul Bodin and Victoria Summers (2001). *The Modern VAT*. (Washington, DC: International Monetary Fund).

Edwards, Jeremy, Matti Toumala and Michael Keen (1994). "Income tax, commodity taxes and public good provision: A brief guide," *FinanzArchiv*, Vol.51, pp. 472-487.

Emran, Shahe and Joseph Stiglitz (2005). "On Selective Indirect Tax Reform in Developing Countries," *Journal of Public Economics* Vol. 89, pp. 599–623.

European Commission (2003). "Experimental application of a reduced rate of VAT to certain labour-intensive services," Report from the Commission to the Council and the European Parliament, SEC(2003)622.

Farhi, Emmanuel, Gita Gopinath, and Oleg Itskhoki (2014). "Fiscal Devaluations," *Review of Economic Studies*, Vol. 81, pp. 725-760

Feldstein, Martin (2002). "The Role for Discretionary Fiscal Policy in a Low Interest Rate Environment." NBER Working Paper 9203.

Fuest, Clemens, Florian Neumeier, and Daniel Stöhlker (2021). "The Pass-through of Temporary VAT Rate Cuts: Evidence from German Supermarket Retail, CESifo Working Paper 9149.

Goncalves, Carlos, Ahn Nguyen and Era. Dabla-Norris, 2025, "Tax-based Fiscal Consolidations and Inflation," mimeo, International Monetary Fund.

Gunter, Samara, Daniel Riera-Crichton, Carlos Vegh, and Guillermo Vuletin (2019), "Non-linear effects of tax changes on output: the role of the initial level of taxation," NBER Working Paper 26570.

Harju, J., Kosonen, T., and Skans, O. N. (2018). "Firm Types, Price-Setting Strategies, and Consumption Tax Incidence." *Journal of Public Economics*, 165(100), 48–72.

Lucie Gadenne, Lucie, Tushar Nandi and Roland Rathelot (2021). "Taxation and Supplier Networks: Evidence from India" Institute for Fiscal Studies WP 19/21.

Hansen, Benjamin, Keaton S. Miller, and Caroline Weber (2021). "Vertical Integration and Production Inefficiency in the Presence of a Gross Receipts Tax," National Bureau of Economic Research Working Paper 28478.

Harrison, Graham, and Russell Krelove (2005). "VAT Refunds: A Review of Country Experience." IMF Working Paper 05/218.

Hindriks, Jean and Gareth Myles (2013). *Intermediate Public Economics* (2nd edition) (MIT Press: Cambridge MA).

Hobbes, Thomas (orig. 1651). *Leviathan*. (Penguin: London).

International Monetary Fund (2005). *Dealing with the Revenue Consequences of Trade Reform*. (IMF: Washington DC).

International Monetary Fund (2018). "The Impact of Consumption Tax Increases and their Policy Implications for Japan," pp. 3-15 in *Japan: Selected Issues*.

Jenkins, Glenn P., Hatice Jenkins, and Chun-Yan Kuo (2006). "Is the Value Added Tax Naturally Progressive?" Working Paper 1059, Department of Economics, Kingston University.

Keen, Michael (2008). "VAT, Tariffs, and Withholding: Border Taxes and Informality in Developing Countries," *Journal of Public Economics*, Vol. 92, pp.1892-1906.

Keen, Michael (2013). "Taxation and Development—Again," pp.13-41 in Clemens Fuest and George Zodrow (eds), *Studies of Critical Issues in Taxation and Development* (Cambridge: MIT Press),

Keen, Michael (2014). "Targeting, Cascading and Indirect Tax Design," *Indian Growth and Development Review*, Vol. 7, pp. 181-201.

Keen, Michael (2016). "Some Empirics of Compliance Gaps," pp. 55-67 in Masamitsu Muchizuki (ed.) *Fiscal and Tax Reform in Abenomics*. (Yuhikaku: Tokyo).

Keen, Michael (forthcoming). "Only Connect: Aspects of the Tax Systems Approach," *National Tax Journal*.

Keen, Michael and Ravi Kanbur (2015). "Reducing Informality," *Finance and Development*, March, pp.52-54.

Keen, Michael and Jenny Ligthart (2002). "Coordinating Tariff Reduction and Domestic Tax Reform" *Journal of International Economics*, Vol 56, pp. 489– 507.

Keen, Michael and Jenny Ligthart (2005). "Coordinating Tariff Reduction and Domestic Tax Reform under Imperfect Competition," *Review of International Economics*, Vol. 13, pp. 385-390.

Keen, Michael and Murtaza Syed (2006). "Domestic Taxes and International Trade: Some Evidence," IMF Working Paper WP/06/47.

Keen, Michael and David Wildasin (2004) "Pareto Efficient International Taxation, *American Economic Review*, Vol. 94, pp. 259– 275.

Kind, Hans Jarle, Marko Koethenbuerger, and Guttorm Schjelderup (2008). "Efficiency Enhancing Taxation in Two-sided Markets," CESifo Working Paper, No. 2262,

Kneller, Richard, Michael Bleaney, and Norman Gemmell (1999). "Fiscal Policy and Growth: Evidence from OECD Countries." *Journal of Public Economics*, Vol. 74, pp. 171-190

Kosonen, Tuomas (2015). "More and Cheaper Haircuts after VAT Cut? On the Efficiency and Incidence of Service Sector Consumption Taxes," *Journal of Public Economics* 131, 87–100.

Lu, Bing & Ma, Hong, 2023. "The "Matthew Effect" in Rebates: How Does VAT Rebates Allocation affect Firm Export Performance," *Journal of Public Economics*, Vol. 225.

Lubik, Thomas and Karl Rhodes (2019). "A closer look at Japan's rising consumption tax." Richmond Fed Economic Brief, At <https://api.semanticscholar.org/CorpusID:211379051>

Lyssiotou, Panayiota and Elena Savva (2021). "Who Pays Taxes on Basic Foodstuffs? Evidence from Broadening the VAT Base," *International Tax and Public Finance*, Vol. 28, pp. 212–247

Meade, James (1978). *The Structure and Reform of Direct Taxation: Report of a Committee chaired by Professor J. E. Meade for the Institute for Fiscal Studies*. (London: George Allen & Unwin).

Misch, Florian and Atilim Seymen (2012). "The effects of Countercyclical Fiscal Policy: Firm-level Evidence from Temporary Consumption Tax Cuts in Turkey." ZEW-Centre for European Economic Research Discussion Paper No. 12– 082

Naritomi, Joana (2019). "Consumers as Tax Auditors," *American Economic Review*, Vol. 109, pp. 3031–3072.

Newbery, David (1986). "On the Desirability of Input Taxes," *Economics Letters*, Vol. 20, pp. 267–70.

Nguyen, Anh D. M., Luisanna Onnis, and Raffaele Rossi (2021). "The Macroeconomic Effects of Income and Consumption Tax Changes." *American Economic Journal: Economic Policy*, Vol. 13(2), 2021, pp. 439–66.

Pessoa, Mario, Mario, Andrew Kazora Okello, Artur Swistak, Muyangwa Muyangwa, Virginia Alonso-Albarran, and Vincent Koukpaizan (2021). How to Manage Value-added Tax Refunds. IMF How to Note 21/03

Phillips, Andrew, and Muath Ibaid (2019). "The Impact of Imposing Sales Taxes on Business Inputs." Ernst & Young LLP.

Piggott, John and Whalley, John (2001), "VAT Base Broadening, Self Supply, and the Informal Sector", *American Economic Review*, Vol. 91, 1084–94.

Pirttilä, Jukka and Matti Tuomala (1997). "Income Tax, Commodity Tax and Environmental Policy." *International Tax and Public Finance*, Vol. 4, pp.379-393.

Pomeranz, Dina (2015): "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax," *American Economic Review*, Vol. 105, pp.2539–2569.

Riera-Crichton, Daniel, Vegh, Carlos and Vuletin, Guillermo, (2016), Tax Multipliers: Pitfalls in Measurement and Identification, *Journal of Monetary Economics*, Vol. 79, pp. 30–48.

Ring, Raymond J., Jr. (1999), "Consumers' Share and Producers' Share of the General Sales Tax", *National Tax Journal*, Vol. 52, pp. 79-90.

Saez, Emmanuel (2002). "The Desirability of Commodity Taxation under Non-linear Income Taxation and Heterogeneous Tastes," *Journal of Public Economics*, Vol. 83, pp. 217–230.

Sah, Raaj K. (1983). "How Much Redistribution Is Possible through Commodity Taxes?" *Journal of Public Economics*, Vol. 20, pp. 89–101.

Sandmo, Agnar (1975). "Optimal Taxation in the Presence of Externalities," *Swedish Journal of Economics*, Vol. 77, pp. 86–98.

Schneider, Georg, Frank Stähler and Georg U. Thuncke (2023). "The (Non-)Neutrality of Value-Added Taxation." Max Planck Institute for Tax Law and Public Finance Working Paper 2023 – 20.

Sharma, Rishi R. (2019). "Does the VAT Tax Exports?" *Economic Inquiry*, Vol. 58, pp. 225–240

Shiraishi, Kosuke (2020). "Determinants of VAT Pass-through under Imperfect Competition: Evidence from Japan," *Japan and the World Economy* 61, 101120.

Slemrod, Joel, Brett Collins, Jeffrey L. Hoopes, Daniel Reck and Michael Sebastiani (2017). "Does Credit-Card Information Reporting Improve Small-Business Tax Compliance?" *Journal of Public Economics*, Vol. 149, pp. 1–19. <https://doi.org/10.1016/j.jpubeco.2017.02.010>

Slemrod, Joel and Tejaswi Velayudhan (2022). "The VAT at 100: A Retrospective Survey and Agenda for Future Research", *Public Finance Review*, Vol. 50, pp. 4–32

Smart, Michael and Richard Bird (2009). "The Impact on Investment of Replacing a Retail Sales Tax with a Value-Added Tax: Evidence from Canadian Experience," *National Tax Journal*, Vol. LXII, pp. 591-609.

Sung, Myung Jae , Rajul Awasthi and Hyung Chul Lee (2017). "Can Tax Incentives for Electronic Payments Reduce the Shadow Economy? Korea's Attempt to Reduce Underreporting in Retail Businesses." Policy Research Working Paper 7936. World Bank, Washington, DC.

Waseem, Mazhar (2022). "The Role of Withholding in the Self-Enforcement of a Value-Added Tax: Evidence from Pakistan," *Review of Economics and Statistics*, Vol. 104, pp. 336–354..

Weyl, E. Glen, and Michal Fabinger (2013). "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, Vol. 121, pp. 528-583

Xing, Jing, Katarzyna Bilicka and Xipei Hou (2022). "How Distortive are Turnover Taxes? Evidence from Replacing Turnover Tax with VAT," CESifo Working Paper 9511.



PUBLICATIONS

Efficiency Aspects of the Value Added Tax
Working Paper No. WP/2025/165