Working Paper

INTERNATIONAL MONETARY FUND

IMF Working Paper

The Anatomy of the VAT

Michael Keen

IMF Working Paper

Fiscal Affairs Department

The Anatomy of the VAT

Prepared by Michael Keen*

Authorized for distribution by Michael Keen

May 2013

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

Abstract

This paper sets out some tools for understanding the performance of the value added tax (VAT). Applying a decomposition of VAT revenues (as a share of GDP) to the universe of VATs over the last twenty years, it emerges that developments have been driven much less by changes in standard rates than by changes in ‘C-efficiency’ (an indicator of the departure of the VAT from a perfectly enforced tax levied at a uniform rate on all consumption). Decomposing C-efficiency into a ‘policy gap’ (in turn divided into effects of rate differentiation and exemption) and a ‘compliance’ gap (reflecting imperfect implementation), results pieced together for EU members suggest that the former are in almost all cases far larger than the latter, with rate differentiation and exemptions playing roles that differ quite widely across countries.

JEL Classification Numbers: H21, H25, H83

Keywords: Value added tax, tax gaps, tax compliance

Author’s E-Mail Address: mkeen@imf.org.

*Forthcoming in the *National Tax Journal*. The author has benefited greatly from comments from and many discussions with Richard Bird, Ruud de Mooij, John Norregaard, Victor Thuronyi, Christophe Waerzeggers, and colleagues at the IMF working on VAT gap measurement, especially Eric Hutton, Kentaro Ogata, Mick Thackray, and Juan Toro, and from Kelsey Moser’s excellent research assistance. Nonetheless, the views expressed here are the author’s alone, and should not be attributed to the IMF, its Executive Board, or management.

Contents	Page
I. Introduction	3
II. Understanding trends in VAT revenue.....	4
A. Trends in VAT Revenues and Rates	4
B. Decomposing VAT Revenue	6
III. Understanding C-efficiency	8
A. Concept and Measurement.....	10
B. ‘VAT Gaps’ and the Decomposition of C-efficiency	14
IV. Concluding.....	21
Table 1. Decomposing C-efficiency	200
Figures	
1. Average VAT Revenue by Income Group.....	4
2. Average VAT Revenue by Region	5
3. Average Standard Rate of VAT by Income Group.....	5
4. Average Standard Rate of VAT by Region	6
5. Decomposing Changes in VAT Revenue, 1993–95, by Income Group	8
6. C-efficiency by Income Group	9
7. C-efficiency by Region	9
Appendixes	
1. Data	23
2. Welfare and C-efficiency	25
Table 1. Summary Table, 1993–2012.....	23
Figure 1. Number of Countries with the VAT	24
References.....	26

I. INTRODUCTION

This paper is about some tools for understanding developments and differences in the performance of the value added tax (VAT). These tools have been emerging over the last few years and are now receiving considerable practical attention and conceptual development as their needs for fiscal consolidation focus policymakers' thoughts ever more closely on the potential for increasing the revenue yield of the VAT without unduly compromising other policy objectives. Increases in the standard rate of the VAT have been one of the main ways in which heightened revenue needs have been addressed since the onset of crisis in 2008: in the two years before that, only one European Union (EU) country increased its standard rate; in the two years after, 13 (of the 27) did so. And as standard rates edge higher, so the roles of reduced rates of VAT and non-compliance become greater concerns. It is this nexus of issues that the tools addressed here aim to illuminate. They are, it should be stressed, essentially descriptive, intended to identify problems and opportunities, not to spell out precise responses: they are anatomy—at best, diagnosis—and not medicine.

The general approach is simply that of decomposing VAT-related quantities of interest into elements that meaningfully, if inevitably somewhat loosely (everything being inter-related), capture distinct forces at work. Section II first does this for VAT revenue itself, relating it to the standard rate, the average propensity to consume, and 'C-efficiency': an indicator of the departure of the VAT from a perfectly enforced tax levied at a uniform rate on all consumption. Applying this decomposition to the universe of VATs over the last twenty years it emerges, perhaps surprisingly, that changes in VAT revenues have been driven much less by changes in the standard rate than by changes in C-efficiency. Section III then explores C-efficiency itself, breaking it into a 'policy gap' (in turn divided into effects of rate differentiation and exemptions) and—a particular and increasing focus of much work in many tax administrations in recent years—a 'compliance' gap reflecting imperfect implementation. Here the empirical application must be narrowed in country coverage and become more speculative, but results pieced together for a subset of EU countries suggest policy gaps that are in almost all cases far larger than compliance gaps; reflecting, to degrees that differ across countries, the impact of both rate differentiation and exemptions. Section IV concludes.

II. UNDERSTANDING TRENDS IN VAT REVENUE

It is natural to begin with the most basic indicator of the performance of any VAT: the revenue it raises.

A. Trends in VAT Revenues and Rates¹

Figure 1 plots VAT revenues, in percent of GDP, for essentially all VATs in place in each of the last twenty years (1993–2012)—a total of 150 by the end of the sample period—averaged over countries within four income groups. In the high income group, there has been only a modest increase: from a little under 7 percent to a little over. Elsewhere, however, the increase has been marked: upper middle income countries now raise about as much from the VAT as do high income, and in low income countries VAT revenue has about doubled since the mid-1990s. While it remains the case that, broadly speaking, VAT revenue relative to GDP increases with the level of per capita income, the difference has become very much less marked.

Figure 1. Average VAT Revenue by Income Group (in percent of GDP)

Figure 2 shows the same developments by region. It points to quite diverse experiences. In the Western Hemisphere (the Americas and Caribbean), the Middle East and Central Asia (over the last ten years or so) and (less steadily) sub-Saharan Africa, VAT revenues have increased quite substantially.

¹¹ Appendix 1 describes the data and sources used in this paper. All averages, of course, are over only countries with a VAT at that time.

Figure 2. Average VAT Revenue by Region (in percent of GDP)

An obvious first thought in trying to make sense of these developments is that increased revenue from the VAT has perhaps been brought about by increases in the standard rates at which it is applied (the rate, that is, which applies most widely). Figures 3 and 4, however, suggest that this will not do as an explanation. Again dividing countries by income group, Figure 3 shows that in all cases the trend since the mid-1990s has if anything been toward lower standard rates—the notable exception being the recent increases in the high income countries of Europe. In the lowest income countries, the average standard rate has fallen by around 2 percentage points over the last twenty years. Viewing countries by region, Figure 4 shows, as well as the recent uptick in Europe, trend decreases in sub-Saharan Africa and Middle East and Central Asia. These patterns do not map straightforwardly into revenue developments: in particular, the trend increase in average VAT revenues in sub-Saharan Africa—and even more strikingly for low income countries as a group—has gone along with a trend reduction in the average standard rate.

Figure 3. Average Standard Rate of VAT by Income Group (in percent of GDP)

Figure 4. Average Standard Rate of VAT by Region (in percent of GDP)

This of course suggests further hypotheses—maybe, for instance, lower standard rates have more than paid for themselves? But the more general implication is a need to identify and explore more systematically a fuller range of drivers of VAT revenue.

B. Decomposing VAT Revenue

One way of approaching this is by writing VAT revenue (denoted V) in percent of GDP (Y) as

$$\frac{V}{Y} = \tau_s E^C \left(\frac{C}{Y} \right) \quad (1)$$

where τ_s denotes the standard rate of the VAT, C denotes consumption (valued at VAT-exclusive prices) and

$$E^C \equiv \frac{V}{\tau_s C} \quad (2)$$

denotes ‘C-efficiency’: the ratio of VAT revenue to the product of the standard rate and consumption.² This last has come to be a widely used tool for evaluating VATs, implicitly comparing the revenue that some VAT actually raises with that which would be raised if it were perfectly enforced and levied at a uniform rate, equal to the standard rate, on all consumption, and with no exemptions—assuming that moving to such a world had no effect

² The notion of C-efficiency originates in Ebrill et al. (2001); OECD (2008) calls it the ‘VAT revenue ratio.’

on the level or composition of consumption.³ (Some terminology: ‘exemption’—also known as ‘input taxation’—means that no tax is charged on sale, but VAT charged on inputs is not refunded or credited, and so ‘sticks’; ‘zero rating’ means no tax on sales, but full recovery of input VAT. Since they distort input choices, exemptions are anathema to the logic of the VAT in a way that zero-rating is not). Quite what C-efficiency can and cannot tell us about VATs is taken up in the next section; the focus for the moment is on its role in understanding the mechanics of VAT revenues.

The decomposition in (1) isolates three drivers of VAT revenues: the standard rate, C-efficiency, and the share of consumption in GDP. Translated into proportional changes,

$$\left(\frac{\widehat{V}}{\widehat{Y}}\right) = \widehat{\tau}_S + \widehat{E}^C + \left(\frac{\widehat{C}}{\widehat{Y}}\right) \quad (3)$$

so that the proportional difference in VAT revenues (over countries or time) is the sum of the proportionate differences in these three components. This, of course, is just mechanics, there being no reason to suppose the three components on the right of (3) to be independent; indeed there is good reason to suppose otherwise—an increase in the standard rate, for instance, might reasonably be expected to increase noncompliance and so reduce C-efficiency (and for the dataset used here there is evidence that it does), and also to lower the average propensity to consume (as Alm and El-Ganainy (2013) find to be the case for a panel of EU countries).⁴ Nonetheless, (3) is potentially informative in identifying the proximate sources of developments in VAT revenue.

Figure 5 illustrates. It shows the decomposition in (3), calculated by country and averaged within income groups, over (mostly) four-year intervals since 1993, of the overall change in VAT revenues (in percent), indicated by the bullet, as the sum of the changes in the components, indicated by the bars. Taking high income countries between 1996 and 2000, for example, VAT revenue increased on average by about 3 percent; this was the net outcome of a 1 percent increase from a higher standard rate and 2.6 increase from greater C-efficiency, balanced against a 0.6 percent reduction from a fall in private consumption relative to GDP.

³ It is safer of course to drop this last assumption and simply interpret E^C mechanically; but that is not how most convey and use the idea.

⁴ Their finding is for final household consumption expenditure; revenue effects might in principle lead to an increase in final general government consumption, usually (and here) included in the denominator of C-efficiency—as discussed below.

Figure 5. Decomposing Changes in VAT Revenue, 1993–95, by Income Group (in percent)

Two strong impressions emerge from Figure 5. The first is that the most powerful of the immediate drivers of changes in VAT revenues, by far, has been changes in C-efficiency. Changes in the standard rate, in contrast, directly account for a relatively small part of developments in VAT revenue changes, and (with some exceptions—notably in the experience of high income countries in the period including the crisis from 2008) changes in the average propensity to consumer account for even less. The second, from somewhat closer inspection, is that average C-efficiency has often moved in the opposite direction to the average standard rate. The implication is that understanding developments in VAT revenues requires understanding developments in C-efficiency.

III. UNDERSTANDING C-EFFICIENCY⁵

Figures 6 and 7 chart the course of C-efficiency by income group and region, respectively. The former shows that C-efficiency tends to be higher in higher income countries (though, it

⁵ The ideas in this section build on those set out in Ebrill et al (2001), OECD (2008), IMF (2010), and de Mooij and Keen (2012), as well as work underway on the measurement of compliance gaps at the IMF.

is important to stress, there are exceptions) and also points to trend increases in all income groups over the last twenty years or so, most modestly in Europe, and a quite marked consequent convergence. The regional breakdown in Figure 7 suggests, however, that this reflects quite complex compositional effects: in Asia-Pacific, most notably, there has been a continuing decline in average C-efficiency throughout the period. What is clear, in any event, is that average C-efficiencies have changed quite substantially over time; the challenge for policy analysts is to understand why.

Figure 6. C-efficiency by Income Group

Figure 7. C-efficiency by Region

A. Concept and Measurement

The idea of C-efficiency has proved very attractive to applied policy analysts: calculating it, and making comparisons with other countries, is a core element in much of the advice commonly given to developing countries, and is now routinely reported for its members by the OECD. One practical merit is that the data requirements for putting some number on C-efficiency are relatively modest. And these numbers can provide a vivid first assessment and expression of the potential revenue gain from base-broadening. Suppose, for instance, that C-efficiency is 60 percent and the standard rate 10 percent. Then, assuming true consumption C to be constant (implausibly, but perhaps a defensible first approximation) extending and effectively applying the standard rate to all consumption would increase revenue by two-thirds ($= (1 - E^C)/E^C$); and the same revenue could be raised by a standard rate of 6 percent ($= \tau_s E^C$) applied to all consumption.

There are, however, both the conceptual and practical issues to be faced in using C-efficiency to assess VAT systems.

Conceptual issues

There is no deep welfare basis for the use of C-efficiency in assessing a VAT. Reforms that bring a VAT closer to the benchmark of 100 percent C-efficiency do not necessarily mean a better VAT. C-efficiency can be increased, for example, by denying VAT refunds to exporters, or by introducing exemptions for intermediate goods;⁶ there may or may not be other good reasons to do so, but the effect is in each case to undermine the intended role of the VAT as a tax on domestic consumption.

Even leaving aside imperfections in the application and enforcement of the VAT as a tax on (only) final consumption, the welfare significance of C-efficiency is tenuous at best. For a single consumer economy, it is shown in Appendix 2 that C-inefficiency ($1 - E^C$) is approximately the sum of two terms: the reduction in deadweight loss in moving to a system in which the current same standard rate applies to all consumption (which can take either sign); and the associated welfare loss in moving to the latter system (expected to be positive, since rates reduced below the standard tend to apply more extensively than rates above—ignoring here the impact on public spending, which a fuller treatment would recognize).⁷ One implication is that a base-broadening reform that increases C-efficiency without changing the standard rate must reduce deadweight loss by more than it reduces welfare. But that is not a very powerful observation.

⁶ Standard procedure is to ensure that exports leave a country VAT-free by refunding exporters the VAT paid on their inputs; exemptions for intermediates lead to higher revenue both directly and indirectly (because of their likely cascading effect in increasing intermediate's sales prices and hence any VAT paid at subsequent stages).

⁷ The details are also in Appendix 2.

In welfare terms, not much can thus be read into changes in C-efficiency. The same remains true even if uniform taxation—the reference point around which the idea of C-efficiency is built—is optimal. The conditions for this are, of course, strong (and not entered into here).⁸ More to the point, even when uniformity is optimal, and even for a revenue-neutral reform, it seems that a reduction in C-efficiency does not imply an increase in welfare; and nor does a welfare-improving revenue-neutral reform necessarily increase C-efficiency. A direct welfare interpretation would follow, however, if C-efficiency were redefined to have in the denominator the revenue that would be raised by a uniform commodity tax system yielding the same welfare as the system being evaluated. Such a measure is directly proportional to the reduction in deadweight loss in moving to such a system: see Appendix 2. This though would be a substantial modeling exercise in itself.

Ultimately, the use of uniform taxation as a reference point in evaluating VAT systems reflects the pragmatic judgment that the practical case for a single rate is so strong as to make this an important benchmark. This is indeed the common view of practitioners. The case is especially persuasive in higher income countries, a central lesson of the literature being that uniform commodity taxation is more likely to be optimal the more sophisticated the range of other instruments by which equity objectives can be pursued. Whether the weaker instruments available to developing countries are nevertheless strong enough for uniformity to be the best policy, and/or the implementation difficulties associated with multiple rates large enough, is less widely agreed. In this respect, the measure may need to be used with special caution in such contexts.

Measurement issues

The relative ease with which numbers can be put together for C-efficiency risks hiding a number of issues.

Starting with the numerator, one issue is the treatment of purchases by non-residents. With the VAT intended, or at least generally viewed, as a tax on domestic consumption, C in the denominator of (2) might naturally be taken to exclude purchases by non-residents—as it generally is; and VAT revenue from sales to non-residents naturally excluded from V in the numerator—which it generally is not. In most cases the implications are likely to be modest, and one might in any event argue that it is expenditure within a country rather than consumption by residents alone that is the proper concern of policymakers. For some countries, this materially affects the interpretation of C-efficiency: it is likely one reason, for instance, why small island economies, often with strong tourism sectors, tend, all else equal, to have high C-efficiency. A further departure from the underlying concept of C-efficiency noted by OECD (2012)—which provides a good account of the measurement issues in this

⁸ See for instance Crawford, Keen, and Smith (2010).

area—is that in some cases cross-border services provided to final consumers are taxed where they are provided rather than where they are used, so that the corresponding revenue does not follow through to the location of final consumption.

Deeper issues arise in measuring the consumption variable C to which the standard rate is applied in the denominator of (1). While something with ‘consumption’ in its name can usually be found in the national accounts, this can differ significantly from what the logic of C -efficiency calls for:⁹ which is the aggregate of consumption as it would be properly defined were it to be subject to VAT.

Some issues under this heading arise from disagreement or differing practice in the treatment of particular items. The value of owner-occupied housing services, for instance, is generally included in national account measures of consumption as an imputed rental value. Taking that as a base is indeed how one ideally might wish to tax the consumption value of housing, but this is unlikely to be feasible in practice (and no country tries to do so); those countries that do aim to tax housing services do so by instead charging VAT on the initial sale of new properties (taking this as a proxy for the present value of future consumption services) or, more crudely, by taxing construction inputs.¹⁰ Even current ‘best practice’ in taxing housing thus implies a base that differs from the corresponding component of national accounts consumption.

Perhaps the most quantitatively significant issues, however, are those that arise in the treatment of the public sector.¹¹ The difficulty is that much public provision is not at anything that could be interpreted as market prices, whether as a matter of policy (basic education and health, for instance) or because their non-excludability makes it hard to conceive of market pricing (defense being the obvious if unimaginative example).¹²

The national accounts approach to such items is to treat them as consumption by the government, not households, and to proxy their value by their cost of production, which means adding to the cost of goods and services purchased (net of sales at something like market prices) the labor and capital costs incurred. This corresponds, in the jargon of the UN *System of National Accounts*, to the ‘Final consumption expenditure of government.’ It may

⁹ In fact, views may differ on this. OECD (2012) raises the possibility that measurement should follow normal (and presumably ‘good’) practice so that, for instance, housing construction, taxed in many countries, would be included in C . The risk is that this sets the bar too low, though it is also the case that—for practical reasons—the approach set out below sets it too high in implicitly including in C a valuation of pure public goods that are simply not amenable to any form of consumption tax. A useful exercise, not attempted here, would be to apply both approaches and assess the practical significance of the differences.

¹⁰ Cnossen (2011) provides an instructive account of practice, issues, and possibilities in the VAT treatment of housing.

¹¹ A very careful account of this issue, somewhat differently argued but leading to the same conclusion, is in OECD (2012).

¹² Similar but quantitatively far less important issues arise in relation to the non-profit sector.

seem odd to include salaries and wages in a notional potential VAT base in this way, since they would not themselves be directly subject to VAT. To the extent, however, that these costs reflect implicit subsidies in the provision of goods and services that could in principle be subject to VAT (such as health and education), their inclusion in C is appropriate, and simply a way of estimating the corresponding potential tax base. To the extent, on the other hand, that the government wage bill and other non-commodity input costs are largely attributable to the provision of pure public goods, which could not be taxed, one would ideally wish to exclude them from C . As a practical matter, however, this is generally not an option, especially in developing countries, as information on net commodity purchases by government—let alone such information broken by use in the producing pure public and private commodities—is rarely available. For this reason, the C-efficiency figures reported here (like those now reported by OECD) are based on final consumption expenditure, aggregated over government, households, and non-profits. They thus over-estimate the potential VAT base by including the non-commodity cost of producing public goods—by an unknown but potentially sizable amount.

Even leaving aside these issues, there can be further scope for disagreement as to what should properly be included in C . In the EU, most exemptions are mandatory for member states, so it could be argued that for the purposes of assessing the performance of any of their VATs the final consumption of these items (though not the potential unrecovered input tax in their production) should be excluded from potential revenue in the denominator—though this, by the same token, misstates their performance relative to a wider but arguably no less policy-relevant reference point. For simplicity, it is assumed in the analytics and application below that mandatory exemptions are not to be excluded.

A more precise expression for C-efficiency, reflecting these concerns, will be needed. For this, write VAT revenue in the numerator of (2) as

$$V = \sum_{i=1}^N T_i^* C_i^* \quad (4)$$

where T_i^* denotes the ‘effective’ rate of VAT on, and C_i^* the consumption of commodity $i = 1, \dots, N$, that is actually taxed; if the rate on final sale is zero, this can be taken to coincide with true consumption, C_i .¹³ More generally, non-compliance means that C_i^* may be substantially less than C_i . The structure of the T_i^* , which capture revenue collected throughout the chain of production leading to the final consumption of i (and are to be distinguished from the statutory rates T_i), depends on the structure of input-output relations (as discussed for instance in Ebrill et al., 2001), the key point for present purposes being simply that $T_i^* = T_i$ for all commodities if and only if there are no exemptions at any stage of production. For the denominator of (2) it is assumed—reflecting the considerations above—

¹³ Leaving aside here issues related to subsidized public provision of private goods, taken up later.

that the aggregate consumption measure used includes not only households' consumption of commodities $i = 1, \dots, N$, all of which, it is supposed, are properly regarded as forming part of the potential base, but also that of some commodity $N + 1$, which, reflecting the discussion of public consumption above, reflects consumption untaxed (on final sale) at household level but included in C as an item of government consumption: C_{N+1} might reflect, in particular, the labor and other non-commodity input costs of producing both pure public goods (which could not be seen as part of the potential base) and commodities that are publicly produced but for which payment by households is trivial (which, in principle, could be).¹⁴ Thus

$$C = \sum_{i=1}^{N+1} C_i . \quad (5)$$

Using (4) and (5) in (2), C-efficiency is then

$$E^C = \frac{\sum_{i=1}^N T_i^* C_i^*}{\tau_s (\sum_{i=1}^{N+1} C_i)} . \quad (6)$$

Even this formalization, it should be stressed, is not fully general. Purchases by non-residents, for example, enter V (unless some explicit adjustment is made) but not C ; the same is true, within the EU, of input taxes of exempt commodities exported to other member states; and statutory rates may differ depending on the use made of a particular commodity. Such features can be significant in particular countries, but the focus here is on only the most generic considerations.¹⁵

B. 'VAT Gaps' and the Decomposition of C-efficiency

Intuitively, it is clear that departures from C-efficiency of 100 percent arise from two broad sources: the extent to which the VAT differs in design from a uniform tax on all consumption (and only on consumption, with no exemptions and consequent possibility of unrecovered tax on inputs); and the extent to which its implementation is imperfect. The critical assumption here is that the measure of C being used includes a reasonable adjustment for unrecorded consumption—as national accounts generally try to make (though in ways that vary and will often be unclear to the user of these data—another and major reason for caution in interpretation)..

¹⁴ Conceptually, commodity inputs to the production of these publicly provided items is naturally treated as corresponding to one of the N taxable items giving rise to revenue in (4).

¹⁵ Note too that excluding mandatory exemptions would require a different approach, with potential revenue in the denominator calculated at effective rates corresponding to application of the standard rate throughout the production chain.

This can be characterized more precisely by writing

$$E^C = \left(\frac{\sum_{i=1}^N T_i^* C_i}{\tau_s (\sum_{i=1}^{N+1} C_i)} \right) \left(\frac{\sum_{i=1}^N T_i^* C_i^*}{\sum_{i=1}^N T_i^* C_i} \right) \quad (7)$$

$$= (1 - P)(1 - \Gamma) \quad (8)$$

where

$$P = 1 - \frac{\sum_{i=1}^N T_i^* C_i}{\tau_s \sum_{i=1}^{N+1} C_i} \quad (9)$$

can be thought of as a ‘policy gap’, being zero if the VAT is applied at a single rate to all (and only) consumption and, analogously,

$$\Gamma = 1 - \frac{\sum_{i=1}^N T_i^* C_i^*}{\sum_{i=1}^N T_i^* C_i} \quad (10)$$

is a ‘compliance gap,’ zero if implementation of the VAT is perfect.

The simple decomposition¹⁶ in (4) provides an elegant framework for delving deeper into the roots of C-inefficiency. It also has the convenient implication that it is not necessary to have independent measures of both gaps: given an estimate of C-efficiency (usually readily available) and just one of the gaps, the other can—and, later, will—be inferred as a residual.

Before considering each of the two gaps more closely, it should be noted that there is an important asymmetry between them, as defined here: the policy gap P is calculated assuming perfect compliance (using true consumption C_i , not taxed consumption C_i^*) but the compliance gap Γ is calculated assuming that policy is whatever it is (using effective tax rates T_i^* , not applying the standard rate τ_s to all final consumption). While one could in theory proceed differently, this approach is the most practicable in closely matching existing work on compliance gaps (which has developed much further than that on policy gaps): tax administrations naturally think of measuring implementation relative to the tax system they are asked to enforce, not some hypothetical one. This asymmetry does mean, however, that the implications of narrowing the measured gaps are quite different: for instance, assuming no behavioral impacts on true consumption, a policy change will generally affect both policy and compliance gaps, but a change in the effectiveness of implementation will affect only the compliance gap.

¹⁶ Barreix et al. (2012) discuss a similar decomposition, referring to what are here called policy and compliance gaps as ‘G-’ and ‘X-inefficiency,’ respectively.

The policy gap

The policy gap P in (9) is essentially just a normalized¹⁷ measure of tax expenditures under the VAT, measured relative to a single rate VAT applied to (only) final consumption, calculated under the assumption of full compliance. This in turn can be split multiplicatively into gaps attributable to distinct aspects of VAT design.

One partition of particular interest is that corresponding to the two key design differences between actual VATs and the reference point of a uniform tax applied only and to all consumption: the differentiation of the statutory rate across commodities and the presence of exemptions. This is captured by writing

$$1 - P = \left(\frac{\sum_{i=1}^N T_i C_i}{\tau_s (\sum_{i=1}^N C_i)} \right) \left(\frac{\sum_{i=1}^N C_i}{\sum_{i=1}^{N+1} C_i} \right) \left(\frac{\sum_{i=1}^N T_i^* C_i}{\sum_{i=1}^N T_i C_i} \right) \quad (11)$$

$$= (1 - r)(1 - x) \quad (12)$$

where the impact of differentiation in statutory rates is captured by the ‘rate gap’

$$r = \frac{\sum_{i=1}^N (\tau_s - T_i) C_i}{\tau_s \sum_{i=1}^N C_i} \quad (13)$$

and that of exemptions by the ‘exemption gap’¹⁸

$$x = 1 - (1 - \theta)(1 - \alpha) \quad (14)$$

where

$$\theta \equiv \frac{\sum_{i=1}^N (T_i - T_i^*) C_i}{\sum_{i=1}^N T_i C_i} . \quad (15)$$

and

$$\alpha \equiv \frac{C_{N+1}}{C} . \quad (16)$$

As before, there is an element of sequencing in the relationship between the two components: the exemption gap is calculated with the rate differentiation still in place, and the rate gap with exemptions removed.

¹⁷ The normalization is slightly odd, through the inclusion in the potential base of C_{N+1} ; this point is reflected in the decomposition in (11) below.

¹⁸ The decomposition here differs from that in de Mooij and Keen (2012) both in the recognition of C_{N+1} , and in the calculation of the exemptions gap here at true rather than taxed consumption. The present approach seems to capture better the essential difference between design and implementation.

The rate gap r in (13) is readily interpreted, and non-negative so long as the (true) consumption-weighted average rate is no greater than the standard rate, as seems invariably to be the case.

The exemption gap x in (14)-(16) is much more subtle, reflecting a range of effects.¹⁹ The first and most straightforward, reflected in θ , is the cumulation of unrecovered VAT on all items of private consumption. Since $T_i^* \geq T_i$ (the effective tax rate being at least as large as that of sales to final consumers, because it also includes any unrecovered VAT tax accumulated in production)²⁰ this latter component of the exemption gap, as defined, is non-positive. (This is also, of course, a reminder of the care needed in using C-efficiency as an indicator: exemption of intermediates, while fundamentally distorting the VAT away from its intended purpose as a tax on only final consumption, actually reduces the measured policy gap). Take, for instance, financial services which, when charged for in the form of a margin, are commonly exempt. Since the value of such services is included in national accounts measures of household consumption, in principle entering as one of the C_i , the impact of the zero rate on financial consumption is captured in the rate gap, while the cascading of VAT ‘stuck’ on inputs purchased by the sector is in principle captured in the exemption gap though its impact on the effective rate.

The second component of the exemption gap, α , in turn reflects two considerations that, as discussed above, are wound into C_{N+1} . One is the value of the provision of pure public goods that, ideally, would have been removed from the denominator of C-efficiency. The other is the value of commodities provided free of charge, many of which are typically VAT-exempt. For these, the impact of the zero-rate on final consumption, which would ideally be included in the rate gap, must instead be absorbed in the exemption gap solely for the practical reason that the consumption value of these items is not provided in the national accounts.

With $\alpha \geq 0$ and $\theta \leq 0$, the sign of the exemption gap x is in principle ambiguous. Loosely speaking, the exemption gap will be strictly positive so long as the accumulation of unrecovered VAT is outweighed by the impact of public provision.

The compliance gap

The compliance gap Γ in (10) is simply the difference between the amount of VAT that is payable in principle and that actually received by government, expressed as a proportion of the former. The underlying sources of imperfect compliance with the VAT are discussed, for

¹⁹ There is an element of arbitrariness here, in that one might alternatively absorb α into the rate gap term, or identify it separately. The approach here is largely driven by the practical concern that estimates of the exemption gap will later be backed out using extraneous estimates that approximate the rate gap as defined in (13). (The handling of exemptions is the trickiest and least satisfactory part of the analytical exercise here—just as they are the trickiest part of the VAT itself).

²⁰ Assuming here $T_i \geq 0$ for all i : see page 209 of Ebrill et al (2001).

instance in Baer (2012) and Keen and Smith (2006); the concern here is with its measurement.

The estimation of compliance gaps for the VAT—often referred to without qualification as ‘VAT gaps’—is now the focus of much and growing interest.²¹ The UK has produced such estimates for several years (HMRC, 2012; Thackray, 2012); Australia now does so too (Australian Taxation Office, 2012); Reckon (2009) produced estimates for all EU countries in a study commissioned by the European Commission (with an updating exercise now under way); Trigueros et al. (2012) collect estimates for Latin American countries; and the estimation of such gaps is becoming an increasing focus of the technical assistance and advice provided by the IMF and others.

Various methods can be used to estimate theoretical VAT liability and hence, by comparison with actual VAT revenue, the VAT compliance gap.

One common distinction is that between ‘top down’ and ‘bottom up’ approaches. The former—which is in practice the most widely used—effectively estimates $\sum_{i=1}^N T_i^* C_i$ directly by using information on consumption disaggregated by commodity, generally from household surveys, together with national accounts data that enables some handle on such VAT as is in principle unrecoverable on intermediate purchases and those of traders below the threshold at which registration for the VAT is compulsory. The latter approach instead grosses up operational information from audit and other activities to estimate VAT that is due but unpaid. There are other possible approaches. One currently being applied in the IMF, for example, uses national accounts source and uses tables to mimic the chain structure of the VAT, estimating unpaid VAT by sector and aggregating this to arrive at an estimate of the overall gap. Which method is preferred will depend on both data availability and the deeper objectives of the exercise. The bottom up approach, for instance, provides more insight into measures that might be taken to close the gap than does the classic top down approach; but the top down approach may be more accurate when operational information is relatively thin—or even contaminated by governance issues within the tax administration—and in any event can be constructed using only information readily available in almost all countries. Importantly, estimation methods relying on survey and national accounts data can only be as good as the adjustments for unrecorded transactions that the statisticians have already made to them; about which, often, users of these data know relatively little.

Subtle but potentially quantitatively significant questions also arise as to precisely how the compliance gap should be defined. Defining VAT receipts, in particular, is much less straightforward than it may seem. Simply defining them as cash amounts collected has obvious appeal, since revenue is of little use to the government unless it actually receives it.

²¹ There is nothing new in compliance gap analysis, of course, which dates back at least to work on the income tax using the Taxpayer Compliance Management Program of the US. But the VAT is in many respects much easier for compliance gap analysis, as the base is closer to independently-derived national accounts aggregates than that of the income tax, and is perhaps for this reason the main focus of the current resurgence of interest.

But cash collections in any year generally include collections in respect of previous years' liabilities; and part of this year's liability will be collected in the future. Moreover, some cash collections (from exporters, notably, who are invariably zero-rated) may represent input VAT that can be expected to give rise to future refunds (where, for instance, excess credits must be carried forward for some time before entitlement to refund arises).²² These considerations point toward an accrual-based measure of receipts (so that, for instance, usable excess credits would be excluded)—but then we are back to the concern that this can hide administrative imperfections if accrued revenues are never collected. In terms of theoretical liability too, there is scope for disagreement: revenue foregone as a result of avoidance activities might be excluded on the grounds that the focus should be on failure to meet legal obligations, but the UK, for example, includes them on the grounds that they should be recognized among the risks to VAT revenue.

Viewed in the context of the decomposition set out above, choices made on how to handle these issues in principle require adjustments to other components: with avoidance included in the compliance gap, for instance, the consumption variable C in the denominator of C -efficiency should be adjusted to include the corresponding base. More fundamentally, questions of the kind just raised generally arise from the deeper objective of wanting somehow to assess the performance of the tax administration. That clearly requires a much finer evaluation than any single number can provide (with gaps on both cash and accrual bases, for instance, being informative). Nonetheless, some dimensions of performance could in principle, and given adequate information, be accommodated in the framework set out here by decomposing the compliance gap. One distinction of interest, for instance, is that between VAT revenue collected at the due date and that collected subsequently, with the difference providing some indication of the effectiveness of at least some forms of intervention by the tax authorities.²³ This could be incorporated by decomposing a compliance gap defined in terms of all collections over the measurement period into parts paid on time and subsequently.

An illustration

At this point, one might hope to see a carefully executed application of the framework set out above. In fact, no one has yet provided one. But by combining elements from different sources, it is possible to at least give some impression of how such an application might look and where it might lead.

For countries belonging to both the EU and the OECD, estimates of C -efficiency for 2006 in OECD (2012), shown in the second column of Table 1, can be combined with the estimated compliance gap Γ for the same year (the latest on which it reports) in Reckon (2009), shown

²² It is also true that VAT remitted by one trader may be taken as a credit by another, but these two events are likely to take place with little if any lag, whereas the delay in obtaining refunds can be significant.

²³ Reflecting this distinction, what is called here simply the compliance gap is sometimes referred to as the 'net' gap, with the 'gross' gap being that relative to payments made by the due date.

in the third column, to arrive at the estimated policy gaps P in the fourth. Two lessons emerge. First, in all countries the policy gap is larger than the compliance gap—very much larger, except in Greece. C-inefficiency in these countries, which is for all but Luxembourg (a special case),²⁴ quite extensive, is primarily attributable to design rather than implementation. (The opposite, it should be noted, seems likely to be true in many developing and emerging economies). Second, policy gaps vary quite widely, even though all countries are constrained by the same EU rules.

Decomposing these policy gaps into rate and exemption gaps is an even riskier exercise, since information enabling this to be done does not seem to be available for 2006. The last columns of Table 1 report two approaches: using the weighted average VAT rates on final household consumption in Mathis (2004) and (in brackets) those on taxed and exempt

Table 1. Decomposing C-efficiency

Country	C-efficiency (E^c)	Compliance gap (Γ)	Policy gap (P)	Decomposing the policy gap:	
				Rate differentiation (r)	Exemptions (x)
Austria	59	14	31	18 (23)	17 (11)
Belgium	52	11	42	22 (30)	25 (17)
Denmark	64	4	33	0 (10)	33 (26)
Finland	61	5	36	12 (33)	27 (17)
France	51	7	45	26 (30)	26 (22)
Germany	57	10	37	12 (18)	28 (22)
Greece	47	30	33	30 (26)	4 (9)
Ireland	66	2	33	24 (38)	12 (-0.09)
Italy	43	22	45	26 (30)	26 (21)
Luxembourg	87	1	12	30 (34)	-26 (-32)
Netherlands	60	3	38	24 (31)	19 (11)
Portugal	53	4	45	25 (36)	27 (14)
Spain	57	2	29	33 (31)	-6 (-3)
Sweden	56	3	42	19 (22)	29 (26)
United Kingdom	48	17	42	21 (31)	27 (17)

Sources: C-efficiency (for 2006) is from OECD (2012) and the compliance gaps (for 2006) are from Reckon (2009); the policy gap is calculated from (8) as a residual. Figures for rate differentiation and, as a residual (from (12)), the exemption gap are calculated using the standard and weighted average implicit VAT rates for household final consumption (for 2000) in Mathis (2004); those in brackets use the standard and weighted average VAT rates on zero-rated, taxed, and exempt consumption (for 2011) underlying Chart 1 of Borselli, Chiri, and Romagno (2012), kindly provided by the authors.

²⁴ VAT revenues in Luxembourg likely reflect both extensive sales to non-residents and input VAT paid on the inputs of financial institutions on their exempt activities that generate final consumption in other member states.

final sales in Borselli, Chiri, and Romagno (2012) to construct (by comparison with the standard rate of that year) an estimate of the rate gap r and, hence, as the residual from (12), of the exemption gap x . The latter approach gives systematically higher rate gaps (and correspondingly lower exemption gaps) since it attributes at least some of the revenue foregone from the final sales of exempt goods to the rate structure; in terms of the notation above, C_{N+1} , and hence α , is consequently smaller.

Under either approach, the rate gap evidently varies greatly, from zero–10 percent in Denmark to 25–30 percent in countries with extensive zero and reduced rating. That, of course, was well known. Much less obvious (and complex) is the story told by the estimated exemption gaps (though these need to be interpreted with particular caution, being residuals assembled from estimates relating to different years). They are almost all positive, which is consistent (recalling the discussion of (14) above) with exemptions of public goods and subsidized items outweighing unrecovered input VAT, though with some negative entries providing a reminder that this need not be the case.²⁵ They are also, moreover, large. Some unknown part of this will relate to the inherently non-taxable proxy value of public goods included, as discussed above, in the exemption gap; nonetheless, both the levels of the two sets of exemption gap estimates, and the difference between them (giving some idea of the revenue lost from final sales of exempt items) lend some quantitative force to the view, most powerfully expressed by Cnossen (2003), that the mandatory exemptions of the EU merit, at the least, close review.

IV. CONCLUDING

The tools of the trade discussed here have many limitations. As stressed by Gemmell and Hasseldine (2012), they have no behavioral content: measures intended to make the tax structure more uniform, for instance, will affect the policy gap not only directly but indirectly too, by affecting consumption patterns, and may well also narrow the compliance gap by easing control problems. They are no substitute for efforts to identify key underlying determinants of VAT performance empirically, as a small parallel literature aims to do.²⁶ Conversely, behavioral changes unrelated to VAT policy or implementation can affect the measured gaps.²⁷

²⁵ The very negative exemption gap for Luxembourg likely reflects the special features noted above.

²⁶ Sancak, Velloso, and Xing (2010), for instance, find that both C-efficiency and the ratio of VAT revenue to GDP move strongly with the stage of the cycle, and Aizenman and Jinjirak (2008) that stronger performance is associated with more durable political regimes, consistent with the view that developing administrative capacity requires investments that are more likely to be undertaken in more politically stable environments. Preliminary analysis of the data used here also suggests a significant negative impact of the standard rate on C-efficiency.

²⁷ In Australia, for instance, VAT revenue has been powerfully affected over the last few years by a shift of expenditure toward untaxed items that largely reflects an increase in their relative price: see Box 4 of Government of Australia (2011).

Clearly too these tools need to be used with care: cross-country comparisons are problematic when methodologies differ, so that measured gaps will often be much better suited for tracking changes within countries over time than cross-sectional differences between them. Even for this, formal standard errors are generally not available,²⁸ so that assessing whether changes are significant becomes essentially a matter of judgment. And the anatomy is superficial, as the skin is superficial to the spine: deeper anatomy is needed, for instance, to understand the operation (or not) of the chain of relationships that is the essence of the VAT.²⁹ The merit of the tools set out here is in organizing thoughts and identifying issues in a way that can be conveyed to practical people. The surprise is not their profundity or even novelty, but that they are not more routinely used already.

²⁸ HMRC (2012) mentions a ± 4 percentage point margin of error for the VAT compliance gap and stresses, but does not quantify, that this is likely much smaller for changes.

²⁹ As, for instance, in the experiments reported by Pomeranz (2011).

Appendix 1. Data

Sources and classifications

VAT revenue is from the IMF Tax Policy Revenue Mobilization Database.

Standard rates of VAT are from IMF Tax Policy VAT Rates Database and the International Bureau of Fiscal Documentation.

GDP is from the World Economic Outlook, series name NGDP, measured in local currency units at current market prices.

Total final consumption expenditure is from the World Economic Outlook, series name NC, measured in local currency units at current market prices.

Countries are placed in **income groups** by the World Bank Income Classification, according to their per capita GNI at July 1st, 2012. The 26 Low Income countries have GNI per capita at \$1025 or less; 40 Lower Middle Income countries have GNI per capita from \$1026 to \$4035; the 46 Upper Middle Income countries have GNI per capita from \$4036 to \$12475; the 38 High Income countries have GNI per capita of \$12476 or more.

Countries are classified into **regions** as in the IMF area departments: Africa, Asia and Pacific, Europe, Middle Eastern and Central Asian (including North Africa), and Western Hemisphere. Details are in the Tax Policy VAT Rates database.

C-efficiency is calculated by dividing VAT revenue by the product of the standard rate and final consumption expenditure less VAT revenue. Generally, the availability of observations on C-efficiency data is constrained by the availability of VAT revenue data.

Description

Appendix Table 1 summarizes the data set used:

Appendix Table 1. Summary Table, 1993–2012

Variable	Observations	Mean	Std. Dev.	Min.	Max.
VAT revenue ¹	1192	6.04	2.52	0.02	16.02
Total final consumption ¹	2870	82.89	17.22	10.37	475.54
Standard rate	2514	15.88	4.88	1.50	35.00
C-efficiency	1136	0.52	0.19	0.002	1.23

¹ In percent of GDP.

Appendix Figure 1 shows the number of countries with a VAT at each sample date, by income group. As of 2012, 150 countries have a VAT: 26 low income, 40 lower middle income, 46 upper middle income, and 38 high income.

Appendix 2. Welfare and C-efficiency

Denoting by $m(q, u)$ the expenditure function giving the lump sum income needed to attain utility u at consumer prices q , the deadweight loss from the tax system being evaluated (following Kay (1980)) is

$$L = m(q, u) - m(p, u) - R \quad (\text{B. 1})$$

where p denotes the vector of producer prices (assumed unchanging) and R revenue; that from a system in which the current standard rate applies to all consumption can be written, in obvious notation, as

$$L^S = m(q^S, u^S) - m(p, u^S) - R^S. \quad (\text{BZ. 2})$$

Taking R^S as an approximation to the revenue that would be raised by applying the current standard rate to the current vector of consumption (that is, ignoring the impact on consumption patterns of changing the non-standard rates), C-inefficiency is $1 - E^C = (R^S - R)/R^S$. Subtracting (B.2) from (B.1), and assuming lump sum income to be unchanged, so that $m(q, u) = m(q^S, u^S)$, gives

$$(1 - E^C)R^S = \{L - L^S\} + \{m(p, u) - m(p, u^S)\} \quad (\text{B. 3})$$

the first term on the right being the reduction in deadweight loss from applying the standard rate to all consumption and the second a measure of the welfare loss also implied.

Consider instead a system that applies a single rate to all consumption, implying consumer prices q^* that, with lump sum income unchanged, generate the same welfare as that in (B.1) and yield revenue R^* . Deadweight loss is then

$$L^* = m(q^*, u) - m(p, u) - R^*. \quad (\text{B. 4})$$

Defining the alternative indicator $1 - E^{C*} = (R^* - R)/R^*$, comparing (B.1) and (B.4) gives

$$(1 - E^{C*})R^* = R^* - R = L - L^*. \quad (\text{B. 5})$$

as claimed.

References

- Aizenman, Joshua, and Yothin Jinjark, 2008, “The collection efficiency of the Value Added Tax: Theory and international evidence,” *Journal of International Trade and Development*, Vol. 17, pp. 391–410.
- Alm, James and Asmaa El-Ganainy, 2013, “Value-added taxation and consumption,” *International Tax and Public Finance*, Vol.20, pp.105–28.
- Australian Taxation Office, 2012, “Measuring compliance gaps in Australia for the GST and LCT,” available at www.ato.gov.au/content/00336240.htm
- Baer, Katherine, 2012, “What international experience can tell us about the potential challenges of administering a U.S. VAT,” mimeo, Office of Tax Analysis, U.S. Treasury Department.
- Barreix, Alberto, Martin Bès, Santiago Diaz de Sarralde and Fernando Velayos, 2012, “Value added tax: Let it be,” mimeo, Inter-American Development Bank, Washington DC.
- Borselli, Fabrizio, Salvatore Chiri, and Ettore Romagnano, 2012, “Patterns of reduced rates in the European Union,” *International VAT Monitor*, January/February, pp.13-21.
- Cnossen, Sijbren, 2003, “Is the VAT’s Sixth Directive becoming an anachronism?” *European Taxation*, pp. 434–42.
- _____, 2011, “A proposal to improve the VAT treatment of housing in the European Union,” *Fiscal Studies*, Vol. 32, pp.455–81.
- Crawford, Ian, Michael Keen, and Stephen Smith, 2010, “VAT and Excises,” pp. 275–362 in James Mirrlees and others (eds), *Dimensions of Tax Design: The Mirrlees Review* (Oxford: Oxford University Press for Institute for Fiscal Studies).
- De Mooij, Ruud, and Michael Keen, 2012, “‘Fiscal devaluation,’ fiscal consolidation: The VAT in troubled times,” forthcoming in Alberto Alesina and Francesco Giavazzi (eds), *Fiscal Policy After the Crisis* (Chicago: University of Chicago Press).
- Ebrill, Liam, Jean-Paul Bodin, Michael Keen, and Victoria Summers, 2001, *The Modern VAT* (Washington: International Monetary Fund).
- Gemmell, Norman, and John Hasseldine, 2012, “The tax gap: A methodological review,” Working Paper in Public Finance 09/2012, University of Wellington.

- Government of Australia, 2011, Budget Paper No. 1: Budget Strategy and Outlook 2011–2015, Statement 5, available at <http://www.budget.gov.au/2011-12/content/bp1/html/index.htm>
- HMRC, 2012, Methodological Annex for Measuring Tax Gaps 2012, at <http://www.hmrc.gov.uk/statistics/tax-gaps/mtg-annex2012.pdf>
- International Monetary Fund, 2010, “From stimulus to consolidation: Revenue and expenditure policies in advanced and emerging economies,” available at <http://www.imf.org/external/np/pp/eng/2010/043010a.pdf>
- Kay, John A., 1980, “The deadweight loss from a tax system,” *Journal of Public Economics*, Vol. 13, pp.111–20.
- Keen, Michael, and Stephen Smith, 2006, “VAT fraud and evasion: What do we know and what can be done?” *National Tax Journal*, LIX, pp. 861–87.
- Mathis, Alexandre, 2004, “VAT indicators,” Taxation Papers, Working Paper No. 2.
- Organization for Economic Cooperation and Development, 2008, *Consumption Tax Trends 2008* (Paris: OECD).
- _____, 2012, *Consumption Tax Trends 2012* (Paris: OECD).
- Pomeranz, Dina, 2011, “No taxation without information: Deterrence and enforcement in the Value Added Tax,” mimeo, Harvard Business School.
- Reckon LLP, 2009, *Study to Quantify and Analyze the VAT Gap in the EU-25 Member States* (Reckon LLP, London).
- Sancak, Cemile, Ricardo Velloso, and Jing Xing, 2010, “Tax Revenue Response to the Business Cycle,” IMF Working Paper 10/71 (Washington: International Monetary Fund).
- Thackray, Michael, 2012, “UK VAT gap,” presentation available at <http://www.imf.org/external/np/seminars/eng/2012/asiatax/>
- Trigueros, Miguel Pecho, Fernando Pelaez Longinott, Jorge Sánchez Vecorena, 2012, “Estimating tax noncompliance in Latin America: 2000–2010,” Centro Interamericano de Administraciones Tributarias (CIAT), Tax Studies and Research Directorate Working Paper No. 3-2012.