

# Is the VAT a Money Machine?

**Abstract** - This paper considers what it might mean to describe the VAT as a "money machine," tests whether it is one, and asks if it might consequently be wise not to adopt it. We find broadly persuasive evidence, using panel data for the OECD, for a "weak form" of the money-machine hypothesis: that countries with a VAT raise more revenue than those without. But the effect may not be large. The evidence also supports a "strong form" of the hypothesis: that this association reflects not increased demand for government, but rather the greater effectiveness of the VAT in raising revenue. Models in which citizens/voters are likely to lose by entrusting politicians with a "money machine" rely on quite extreme views of their preferences and/or the effectiveness of electoral discipline.

## INTRODUCTION

"Some panelists were . . . concerned that introducing a VAT would lead to higher total tax collections over time and facilitate the development of a larger federal government—in other words, that the VAT would be a 'money machine.'"<sup>1</sup>

Over the last 50 years or so, the value-added tax (VAT) has been introduced in around 130 countries, including all OECD members—with the sole and notable exception, of course, of the United States. A central claim made by proponents of the VAT is that it is a particularly effective tax,<sup>2</sup> reducing the welfare costs of raising any given amount of revenue and so facilitating increased revenue mobilization where (as clearly remains the case, in particular, in many developing countries that have adopted the VAT) this is an object of policy. But what proponents see as a merit of the VAT is turned on its head by some of the opponents of the tax, notably in the U.S., who see it, instead, as a fundamental flaw. The VAT, on this view, is simply too easy a way of collecting revenue. Most recently and prominently, this concern weighed heavily on the minds of some members of the recent advisory panel on federal tax reform in the United States. And, indeed, it has run

---

**Michael Keen**

*Fiscal Affairs  
Department,  
International Monetary  
Fund, Washington,  
D.C. 20413*

**Ben Lockwood**

*Department of  
Economics, University  
of Warwick, CV4 7AL,  
United Kingdom*

**National Tax Journal**

Vol. LIX, No. 4  
December 2006

---

<sup>1</sup> President's Advisory Panel on Federal Tax Reform (2006, p. 192).

<sup>2</sup> We do not rehearse here the reasons why this may, or may not, be the case: Ebrill, Keen Bodin and Summers (2001) provide a fairly traditional account, and Emran and Stiglitz (2005), a more skeptical one (albeit mainly for developing countries).

through much of the debate on the VAT in the U.S., dating back to the early work of Brennan and Buchanan (1977), who specifically cite the VAT as an example of their general argument that the public well-being may ultimately be damaged by entrusting self-interested policy makers with efficient tax instruments.

The purpose of this paper is to explore and evaluate the claim—and fear—that the VAT is a “money machine.” What exactly might this irresistible but vague term mean? Is there any evidence that, somehow defined, the claim is, in fact, true? And if it is true, does that mean that adoption of the VAT is a good thing or a bad thing?

The next section sets the scene for this discussion with an overview of the key features and revenue significance of the VAT in OECD countries, which are the natural focus of interest for the U.S. Following this, we try to tease out with some precision what it might mean to describe the VAT (or any other tax) as a “money machine,” deriving testable implications that we explore using panel data for OECD members. The discussion then turns to the question of how political economy considerations might affect the desirability of entrusting policy makers with a money machine. A final section concludes.

## BACKGROUND

Table 1 shows key features of the VAT in the non-U.S. OECD countries. As shown in the first column, all OECD members other than the U.S. have adopted the VAT over the last 30 years or so, beginning

with France<sup>3</sup> continuing through adoption by Australia in 2000. The (unweighted) average standard rate of VAT is about 17 percent, but with considerable variation. Within the EU, it varies between 15 percent (the minimum permissible under the union’s rules) in Luxembourg, and 25 percent (the maximum) in Denmark and Hungary. And several non-EU countries apply far lower standard rates than this, the most striking being the five-percent rate in Japan.<sup>4</sup> Most also apply a reduced rate to some commodities, with domestic zero-rating<sup>5</sup> being quite widespread.

The fourth column shows that revenue from the VAT is also typically substantial—averaging a little over seven percent of GDP—but again with considerable variation, from a high of over 12 percent of GDP in Iceland to a low of around 2.5 percent in Japan. Importantly, this variation in revenue yield is only very imperfectly explained by differences in standard rates: the VAT in New Zealand, for example, raises nearly three points of GDP more than does that in the U.K. despite having a standard rate that is 5.5 points lower. While these differences in part reflect structural differences in the wider economy, what also evidently matters a good deal in practice is not only the standard rate but also its coverage. Some sense of this is provided by the “C-efficiency” figure in the final column of Table 1, this being the ratio of VAT revenue to the product of aggregate consumption and the standard rate. For a textbook VAT levied uniformly on all consumption, C-efficiency would be 100 percent. It is reduced below this by the application of zero or reduced rates, and by the exemp-

<sup>3</sup> Opinions differ as to precisely when France is best said to have adopted a VAT, having introduced various degrees of crediting from the late 1940s onwards.

<sup>4</sup> This rate—applied also in Netherlands Antilles, Paraguay and Singapore—appears to be the lowest in the world.

<sup>5</sup> “Zero-rating” means simply taxation at a zero rate, implying that no VAT is due on output but VAT paid on inputs is creditable in the usual way, implying such input VAT is refunded. This is quite distinct from “exemption,” which is explained below.

**TABLE 1**  
VAT RATES, REVENUES AND C-EFFICIENCY IN THE OECD, 2005

	Introduced	Standard Rate	Reduced Rates	VAT Revenue (in percent of GDP)	C-Efficiency (2005)
Australia	2000	10.0	Zero <sup>1</sup>	4.3	53.0
Austria	1973	20.0	10.0; 12.0 <sup>2</sup>	7.9	52.9
Belgium	1971	21.0	6.0; 12.0; Zero <sup>1</sup>	7.1	42.9
Canada	1991	7.0	Zero <sup>1,3</sup>	5.0	66.5
Czech Republic	1993	19.0	5.0	7.4	38.9
Denmark	1967	25.0	Zero <sup>1</sup>	9.9	51.6
Finland	1994	22.0	8.0; 17.0; Zero <sup>1</sup>	8.7	52.9
France	1968	19.6	2.0; 5.5 <sup>4,5</sup>	7.4	45.3
Germany	1968	16.0	7.0	6.2	50.5
Greece	1987	18.0	4.0; 8.0 <sup>6</sup>	8.0	51.5
Hungary	1988	25.0	5.0; 15.0	10.5	41.3
Iceland	1989	24.5	14.0; Zero <sup>1</sup>	12.1	49.2
Ireland	1972	21.0	4.8; 13.5; Zero <sup>1</sup>	7.5	55.5
Italy	1973	20.0	4.0; 10.0; Zero <sup>1</sup>	6.0	38.2
Japan <sup>7</sup>	1989	5.0	—	2.4	65.3
Korea	1977	10.0	Zero <sup>1</sup>	4.4	68.9
Luxembourg	1970	15.0	3.0; 6.0; 12.0	7.1	68.2
Mexico	1980	15.0	Zero <sup>1,8</sup>	3.7	30.4
Netherlands	1969	19.0	6.0	7.7	51.9
New Zealand	1986	12.5	Zero <sup>1</sup>	9.1	96.4
Norway	1970	25.0	7.0; 11.0; Zero <sup>1</sup>	8.5	52.5
Poland	1993	22.0	3.0; 7.0; Zero <sup>1</sup>	7.4	40.2
Portugal	1986	19.0	5.0; 12.0 <sup>9</sup>	8.5	53.7
Slovak Republic	1993	19.0	—	8.1	44.6
Spain	1986	16.0	4.0; 7.0 <sup>10,11</sup>	6.1	50.1
Sweden	1969	25.0	6.0; 12.0; Zero <sup>1</sup>	9.3	47.3
Switzerland	1995	7.6	2.4; 3.6; Zero <sup>1</sup>	4.0	71.7
Turkey	1985	18.0	1.0; 8.0	7.1	56.5
United Kingdom	1973	17.5	5.0; Zero <sup>1</sup>	7.0	46.4
Average		17.7		7.2	52.9

Source: OECD, *Consumption Tax Trends*, 2006 edition.

Notes:

<sup>1</sup>“Domestic zero rate” means tax is applied at a rate of zero to certain domestic sales.

<sup>2</sup>Applies in Jungholz and Mittelberg.

<sup>3</sup>The provinces of Newfoundland and Labrador, New Brunswick, and Nova Scotia have harmonized their provincial sales taxes with the federal Goods and Services Tax and levy a rate of 15 percent. Other Canadian provinces, with the exception of Alberta, apply a provincial tax to certain goods and services. These provincial taxes apply in addition of GST.

<sup>4</sup>Applies in Corsica.

<sup>5</sup>Applies to overseas departments excluding French Guyana.

<sup>6</sup>Applies in the regions Lesbos, Chios, Samos, Dodecanese, Cycladen, Thassos, Northern Sporades, Samothrace and Skiros.

<sup>7</sup>Central government taxes only.

<sup>8</sup>Applies in the border regions.

<sup>9</sup>Applies in Azores and Madeira.

<sup>10</sup>Applies in the Canary Islands.

<sup>11</sup>Applies in Ceuta and Melilla axes on specific goods and services.

tion<sup>6</sup> of final consumption: this largely explains the relatively low C–efficiency in the U.K., for example, which zero–rates about 13 percent of household expenditure and exempts another 30 percent. Variation in C–efficiency across the OECD is evidently also wide, from nearly 100 percent in New Zealand, whose VAT is often taken as a model of good design, to as low as around 40 percent.

Figure 1 shows the importance of the VAT as a source of revenue within the wider tax systems of the OECD, with countries ordered by the ratio of total taxation (defined throughout the paper to be inclusive of social security contributions) to GDP, and including the U.S. (and revenue from sales taxation) for comparison. The low level of sales tax revenue in the U.S., compared to the VAT elsewhere, stands out.<sup>7</sup> What is also clear, however, is that there are countries—Japan, Korea and Mexico—that have a VAT but “nevertheless” have smaller government than does the U.S. In that sense, adoption of a VAT is clearly not a sufficient condition for large government. More generally, while there is, indeed, a positive association between revenue from the VAT and total tax revenue, the cross–country differences in government size are evidently not fully explained by differences in VAT revenue: Ireland and France, for example, both collect about four percent of GDP in VAT revenue, but the overall tax ratio is over ten points lower in the former than in the latter.

Such a tabulation can reveal little, however, about the links between VAT revenue and government size. It could be that cross–sectional differences in VAT revenue substantially “explain” cross–country differences in government size if account

is also taken of other determinants of the latter (standard candidates including the levels of income per capita and openness). The difference in the size of government between Switzerland and Australia may reflect structural differences in their wider economies, for example, with the VAT affecting not the difference in government size between them but rather its level in each. And along the temporal dimension too, it could be that the change in VAT revenue in each country over time explains a good deal of the change in its total tax revenue.


All this calls for a more structured empirical analysis, and we explore this below. For background, however, Table 2 provides some basic information on temporal developments, showing changes in total and VAT revenues, and in the rate structure, between 2003 and the year in which the VAT was introduced.

The first three columns concern changes in the VAT itself. The first shows that the revenue importance of the VAT has, indeed, tended to creep up in the years following its introduction, on average by about 1.7 percent of GDP. But—as in almost all aspects of the VAT experience—there is considerable variation across countries. In several, this subsequent growth has exceeded four points of GDP; in a few others (including, somewhat surprisingly, France and Belgium), the revenue raised by the VAT, relative to GDP, has actually declined. This broad upward trend in revenue reflects a clear tendency, shown in the second column, for the standard rate of VAT to increase over time, in some cases quite substantially; but again there are exceptions, not only in formerly socialist economies (which tended

<sup>6</sup> “Exemption” means that no VAT is chargeable on sales (as under zero–rating), but (unlike zero–rating) VAT paid on inputs cannot be recovered; the Australian term “input–taxed” is more descriptive. Note the implication that while exemption of commodities purchased by final consumers reduces revenue, exemption of items used as inputs by registered taxpayers tends to increase it (because it leads to “cascading” of the VAT—tax being charged on tax) and so also tends to raise C–efficiency. This is one of the main pitfalls in using C–efficiency as an indicator of the quality of the VAT, discussed, with others, in Ebrill et al. (2001).

<sup>7</sup> It should be noted, however, that the figures for the U.S. include only federal taxes.

Figure 1. Total Tax and VAT Commodity Tax Revenue in the OECD, 2004<sup>1</sup>



Sources: OECD, *Revenue Statistics*, 2005 edition; and OECD, *Consumption Tax Trends*.

Notes:

<sup>1</sup>Central government only, and including social security. VAT revenue for all countries except the U.S., for which revenue from (federal) excises is shown.

TABLE 2  
CHANGES IN TAX REVENUES AND VAT STRUCTURE IN THE OECD<sup>1</sup>

	Change since introduction of VAT in:				Excess of change in total tax revenue <sup>3</sup> over current VAT revenue <sup>5</sup>
	VAT Revenue <sup>2</sup>	Standard Rate	Number of reduced rates	Total Tax revenue <sup>3,4</sup>	
Australia	0.4	0.0	0	0.7	-3.4
Austria	0.9	12.0	1	7.6	-0.3
Belgium	-0.4	3.0	2	11.2	4.0
Canada	0.1	0.0	0	-2.4	-7.4
Czech Republic	0.3	-4.0	0	...	...
Denmark	5.7	15.0	0	17.7	8.0
Finland	0.7	0.0	1	-0.2	-8.9
France	-0.9	6.0	-1	8.9	1.5
Germany	1.4	6.0	0	2.5	-3.7
Greece	0.2	0.0	0	5.6	-2.9
Hungary	...	0.0	1	...	...
Iceland	-0.3	0.0	0	7.4	-3.4
Ireland	3.2	4.6	0	0.8	-6.6
Italy	1.8	8.0	0	14.9	9.0
Japan	1.6	2.0	-1	-3.1	-5.6
Korea	2.0	-3.0	0	8.3	3.8
Luxembourg	3.6	7.0	1	10.7	4.7
Mexico	1.2	5.0	0	...	...
Netherlands	3.2	7.0	0	1.6	-5.7
New Zealand	5.1	2.5	0	4.5	-4.4
Norway	0.3	5.0	2	9.4	0.9
Poland	-1.2	0.0	1	-1.3	-8.8
Portugal	2.7	2.0	0	9.7	1.6
Slovak Republic	...	-4.0	0	...	...
Spain	1.0	4.0	0	8.1	2.0
Sweden	4.8	13.9	0	11.5	2.4
Switzerland	0.6	1.1	2	2.0	-2.0
Turkey	4.0	8.0	2	15.5	8.4
United Kingdom	4.4	7.5	1	2.6	-4.4
Average (unweighted)	1.7	3.7	0.4	5.7	-0.8

Source: Calculated from OECD, *Consumption Tax Trends*, 2006 edition.

Notes:

<sup>1</sup>All revenue figures in percent of GDP.

<sup>2</sup>Change between 2003 and first year of operation of the VAT.

<sup>3</sup>Including social security contributions.

<sup>4</sup>Change since year prior to VAT introduction.

<sup>5</sup>Excess of change shown in previous column over VAT revenue in 2004 shown in Table 1.

to set high standard rates when the VAT was introduced, rapidly, at the start of the transition), but also, interestingly, in both Korea and Japan, which did not begin with especially high rates. It is notable too, that several of the countries that have increased their standard rate of VAT have also increased the number of reduced rates applied, presumably in an attempt (wise or not, given the possibility that bet-

ter targeted instruments were available) to mitigate the perceived distributional impact of a higher standard rate.<sup>8</sup>

The final column of Table 2 shows that governments in OECD countries have, indeed, tended to become larger after their adoption of the VAT, in the sense that the proportion of GDP taken in taxes and social security contributions was higher in 2003 than in the year prior to

<sup>8</sup> The information available does not allow an assessment of how the coverage of reduced rates or the extent to which VAT exemptions have changed.

VAT introduction, by nearly six points. The final column of the table shows that in most cases the increase in the overall tax ratio has been less than the revenue raised by the VAT itself. Thus, the revenue raised by the VAT has been to some degree offset by reduced revenue (at least relative to GDP) from other taxes. It will be seen in the next section that the nature and extent of such offsetting is of central importance in evaluating the money machine notion, and exploring this, controlling for other potential determinants of government size, will be a key part of the later empirical analysis.

### WHAT IS A MONEY MACHINE?

What does or might it mean to say that the VAT is a money machine? The term seems to have proved too useful and evocative to define precisely. One can, however, identify (at least) two distinct hypotheses of this kind.<sup>9</sup> The first is simply that:

- Governments with a VAT raise more revenue, all else equal, than those without.

The second—recall the words “lead to” in our opening quotation from the President’s Panel—asserts causality:

- The use of the VAT has in itself been a cause of increased government size.

Since the latter implies the former, but not vice versa, we shall for brevity refer

to these as respectively the “weak” and “strong” forms of the money machine hypothesis. The question then becomes that of how one might test them empirically.

The weak money machine hypothesis is conceptually the more straightforward of the two, and there is an obvious strategy for testing it: simply add dummies (and interaction terms) for the presence of a VAT to standard “tax effort” equations that relate tax ratios (total tax revenue as a percentage of GDP) to such structural characteristics of the economy as openness and per-capita income, and check whether the apparent revenue impact of the VAT is significantly positive.<sup>10</sup> Ebrill et al. (2001) perform such an exercise for a large cross-section of countries, and Keen and Lockwood (2006) extend their analysis to a large panel. Here we undertake such an analysis for a panel of OECD countries. An important and early precursor to this work, it should be noted, is Nellor (1987). This tests what is essentially the weak form of the money machine hypothesis by modeling the tax ratio (for 11 European countries) without introducing a VAT dummy but instead testing for an increase in the mean residual pre- and post-introduction (the results being broadly supportive of the weak form hypothesis).<sup>11</sup>

Formulating and testing the strong form of the money machine hypothesis, with the element of causality, is more challenging, and does not seem to have been

---

<sup>9</sup> One version of the money machine story seems to assert that the VAT is a particularly easy way for governments to raise tax revenue because it can be “hidden” in the consumer price. This is simply wrong: there is no reason why it should not be required that VAT be separately identified in the price charged to consumers, and, indeed, several countries, including Canada and Italy, do precisely that. The discussion here focuses on “money machine” notions resting on genuinely distinctive features of the tax.

<sup>10</sup> Throughout, we conceive of government “size” in terms of total tax revenue. Non-tax revenue is relatively small for the OECD countries with which we are concerned; and revenue data is more readily available than are data on total government expenditure (the two in any event being linked, presumably, at least in present value).

<sup>11</sup> Stockfisch (1985) applies a “difference in differences” logic to address an analogous question cast in terms of the growth of government size rather than (the more natural concern) its level, asking whether the change in the growth of the tax ratio around the time of introduction of the VAT was greater in countries that adopted the tax than in those that did not (and concluding that any such effect was at best modest).

addressed in the previous literature. We shall pursue two approaches.

The first approach is essentially statistical: to ask whether revenue from the VAT Granger-causes total tax revenue (both relative to GDP), in the sense that lagged values of the former are useful in predicting future values of the latter (but not, if the case is to be clinched, vice versa).

The second approach conceives of causality not in a temporal sense, but as a comparative statics statement. More precisely, suppose that the weak form of the money machine hypothesis were empirically verified, so that there is indeed a positive association, all else equal, between VAT revenue and government size. There would be broadly two possible explanations of this. The first is that increased taste for government spending has created revenue needs which have been met by adopting a VAT. The second is that access to the VAT has in itself so increased the efficiency of the tax system that governments have found it optimal to use it to increase their total tax revenues. The latter would seem to be the key claim underlying the strong money machine hypothesis. To see how one might test for it, a simple formalization of the tax design problem is helpful.

Consider then a government that has two tax instruments at its disposal,  $A$  and  $B$ , and that chooses the revenue  $R_i$  raised by each so as to maximize an objective function

$$[1] \quad U = \lambda V(R_A + R_B) - \left(\frac{1}{2}\right)\theta_A(R_A)^2 - \left(\frac{1}{2}\right)\theta_B(R_B)^2,$$

where  $V$  denotes the private utility derived from public expenditure, with  $\lambda$  parameterizing the strength of this prefer-

ence, and the squared revenue terms capture the notion that the marginal efficiency cost of raising revenue by any tax instrument increases with the amount raised, while the  $\theta_i$  parameterize the efficiency of the two tax instruments (higher values indicating a less efficient tax). The necessary conditions on the  $R_i$  are then:

$$[2] \quad \lambda V'(R_A + R_B) = \theta_i R_i, \quad i = A, B.$$

To see the implications, suppose first that the taste for government increases, in the sense that  $\lambda$  becomes larger. Then it is readily seen<sup>12</sup> from [2] that the revenue optimally raised by each tax increases (and indeed, given the simple functional forms being used, they increase by the same proportion). Thus, an increased taste for government is optimally financed by increasing revenue from all tax instruments.<sup>13</sup>

Now suppose instead—to capture the notion that access to a VAT means access to a more effective tax—that one of these instruments, say  $A$ , becomes more efficient, in the sense that  $\theta_A$  falls. It is straightforward to show from [2] that, as one would expect, both the revenue raised by  $A$  and total revenue  $R_A + R_B$  optimally increase. Importantly, however, revenue optimally raised by  $B$ —the tax whose efficiency is unchanged—optimally falls. The reason is straightforward: the social benefit from access to a more efficient tax instrument is optimally taken partly in the form of increased public expenditure, but partly too in the form of reduced reliance on less efficient tax instruments. The degree to which increased revenue from  $A$  is offset by reduced revenue from  $B$  can be shown to be larger—and, hence, the increase in total revenue smaller—the greater in absolute value is  $V''$  and the larger is  $\theta_B$ . For the more rapidly the

<sup>12</sup> This and the other comparative statics results asserted in this section are derived in Appendix A.

<sup>13</sup> This corresponds to what Kenny and Winer (2006), in their recent analysis of the use made of different tax instruments, call the “scale effect.”



marginal valuation of public spending falls, and the more costly is the alternative revenue source, the less is the social benefit from expanding total revenue relative to that from reducing reliance on less efficient taxes.<sup>14</sup>

While the weak form of the money machine hypothesis is, thus, consistent with two possible views of the world—one in which increased taste for government generates an increase in revenue from all sources, and the other in which growth of government is driven by the greater efficiency of the VAT—the strong form of the hypothesis, which rests on the second of these views, carries the further and testable implication that the revenue that countries raise through the VAT should have been offset, to some degree, by reduced revenues from other taxes. In this case, the increase in total revenue associated with use of the VAT will be less than the revenue from the VAT itself.

EVIDENCE

Does the empirical evidence bear out either form of the money machine hypothesis? We consider each in turn.

*A Weak Money Machine?*

The natural way to test for a positive association between overall revenues and the presence of a VAT, as noted above, is simply to estimate a VAT-augmented “tax effort” equation of the general form

$$[3] \quad R_{it} = \alpha V_{it} + \beta'_v V_{it} X_{it} + \beta' X_{it} + \pi_i + \eta_t + u_{it},$$

where the dependent variable  $R_{it}$  is the ratio of tax revenue to GDP in country

$i$  at time  $t$  (taken from OECD *Revenue Statistics*),  $V_{it}$  is a dummy taking the value unity if a VAT is present and zero if not<sup>15</sup> (derived from the dates of VAT introduction given in Ebrill et al. (2001)), and the control variables in the column vector  $X_{it}$ , discussed in more detail below, are taken from standard sources (as described in Keen and Lockwood (2006)). The terms  $\pi_i, \eta_t, u_{it}$  denote, respectively, country- and year-fixed effects and a random disturbance, assumed independently and identically distributed.

Equations of the form in [3] are also estimated in Keen and Lockwood (2006), but on a much larger sample of countries—the concern there being with the impact of the VAT on the full span of countries—and using a different measure of  $R_{it}$ . That paper also addresses the potential bias arising from the endogeneity of VAT adoption. That is, there may be some unmeasured characteristic of a country that affects both the likelihood of its adopting a VAT and the likely revenue gain from doing so. This bias can be corrected, and its existence tested for, by also estimating a VAT adoption equation and then including a Heckman-type correction in the revenue equation. Keen and Lockwood (2006) find, however, no evidence of such bias, and so here we proceed by simply estimating [3] as it stands.

Table 3 shows the results of regressions along these lines. The dataset is the full set of the 30 current OECD member countries for the period 1965–2004 (covering, for each, years subsequent to membership). Throughout, the dependent variable is tax revenue including social security contributions. While the panel is unbalanced, the coverage is quite good, with at least two-thirds of all country-year

<sup>14</sup> While these formalities conceive of an increase in the efficiency of a tax instrument as a small improvement in one already adopted, it is straightforward to establish a similar result for the discrete adoption of an initially unused tax that is more efficient at the margin than that in place: total revenue can be shown to increase by an amount smaller than that optimally collected from the new tax alone.

<sup>15</sup> As the example of France noted above suggests, the date of adoption is not always entirely clear-cut.

**TABLE 3**  
DO COUNTRIES WITH A VAT RAISE MORE REVENUE?<sup>1,2</sup>

Estimation method	1 OLS	2 OLS	3 OLS	4 GMM	5 OLS	6 OLS	7 OLS
<i>R-1</i>			0.867** (41.07)	0.865** (38.00)	0.865** (39.94)	0.838 (36.91)	0.853** (37.64)
<i>Ln(YPC)</i>	2.766** (3.97)	-8.772** (8.19)	-0.379 (1.15)	-0.998 (0.25)	-0.346 (0.56)	-0.815 (1.30)	-0.065 (0.12)
<i>OPEN</i>	-0.338 (0.29)	-3.435** (3.50)	-0.607 (1.70)	-0.624 (1.66)	-1.241 (1.20)	-2.396 (2.10)	-0.222 (0.23)
<i>AGR</i>	-0.459** (7.55)	-0.558** (10.33)	-0.099** (3.93)	-0.082* (2.01)	-0.099* (2.37)	-0.084 (2.01)	-0.032 (0.72)
<i>V</i>	3.095** (9.28)	1.138** (3.57)	0.279 (1.15)	0.203 (1.00)	0.445 (0.23)	-0.015 (0.01)	4.625** (2.44)
<i>Ln(YPC)*V</i>					-0.145 (0.27)	-0.333 (0.63)	-1.368 (2.66)**
<i>AGR*V</i>					-0.026 (0.54)	-0.025 (0.52)	-0.136* (2.25)
<i>FED*V</i>					0.500 (1.58)	0.568 (1.70)	0.33 (0.99)
<i>OPEN*V</i>					0.52 (0.60)	1.84 (1.82)	-0.196 (0.23)
<i>DEPOLD</i>						17.62** (2.94)	15.439** (2.66)
<i>DEPYOUNG</i>						-2.479 (0.71)	0.874 (0.21)
<i>IMFCR</i>						0.002 (0.01)	-0.386 (1.86)
<i>Ln(POP)</i>						0.25 (0.81)	-0.035 (0.11)
Observations ( <i>N</i> )	865	865	856	826	856	856	745
R-squared	0.91	0.94	0.98	n.a.	0.98	0.98	0.98
Country dummies	yes	yes	yes	yes	yes	yes	yes
Year dummies	no	yes	no	no	no	no	no
Joint significance of VAT terms	n.a.	n.a.	n.a.	n.a.	0.0236	0.0708	0.0170
First-order serial correlation test	0.000	0.000	0.6151	n.a.	0.634	0.4801	0.9857

Notes:

<sup>1</sup>Robust t-statistics in parentheses; and \*\* indicates significance at 1 percent, \* at 5 percent.

<sup>2</sup>The diagnostic tests are: (i) an F-test for joint significance of VAT terms; (ii) a test for first serial correlation in panels, proposed by Wooldridge (2002). In each case, for ease of understanding, only the p-value of the test statistic is given. Also, n.a. indicates that the test is not applicable.

observations available for all regressions.<sup>16</sup>

Column 1 reports a regression of the form in [3], with  $X$  including variables commonly found in the tax effort literature: (the log of) income per capita ( $YPC$ ), openness to international trade ( $OPEN$ ), and the share of agriculture in GDP ( $AGR$ ). For simplicity, we initially exclude the interaction terms  $V_{it}X_{it}$ . The results<sup>17</sup> indicate that the presence of a VAT does have a significantly positive impact on the tax ratio. The implication of the point estimate is that, all else equal, tax revenue is higher by about three points of GDP when a VAT is present than when it is not. Since the mean value of  $r$  in the sample is 33.38, this corresponds to a slightly less than ten-percent increase in tax revenue. However, Wooldridge's (2002) test massively rejects the null of no first-order serial correlation.

To pick up general time variation in omitted variables, column 2 adds the year dummies  $\eta$  (the country effects  $\pi$  being included in all regressions). This reduces the effect of VAT on total revenue, but serial correlation remains, pointing to more complex dynamics. Column 3, therefore, adds a lagged dependent variable, denoted  $R - 1$ . This eliminates the serial correlation (the p-value of the test statistic is now below 0.05), but now the effect of the simple VAT dummy becomes small and insignificant.

These estimates are subject, however, to the bias that arises from estimating equation [3] by OLS in the presence of a lagged dependent variable.<sup>18</sup> This bias is of order  $1/T$ , where  $T$  is the number of time-series observations, and so may be fairly modest in a panel of the length used here. Nevertheless, as a check on this, column 4 estimates the same regression as in column 3, but using the Arellano-Bond (1991) GMM estimator. In this case, robust z-statistics, which are normally distributed, are given in parentheses. The Arellano-Bond test for second-order autocorrelation is passed easily.<sup>19</sup> The coefficients in columns 3 and 4 are reasonably similar. Moreover, when  $T$  is large, as in our case, the Arellano-Bond estimator may exhibit finite-sample bias due to overfitting.

So, we proceed by using the OLS estimator, using country dummies to pick up the fixed effects  $\pi_i$  in [3]. Regression 5, therefore, introduces interaction terms between the VAT dummy and all the standard tax effort variables, as well as an interaction with a dummy for a federal country.<sup>20</sup> Now the VAT terms are jointly strongly significant, although individually they are not. Interestingly, the sign pattern of effects from the VAT proves complex, and does not obviously point to an overall revenue gain—a point we return to shortly.

<sup>16</sup> The largest possible number of observations is  $40 \times 30 = 1,200$ , there being 30 countries and 40 years. So, the coverage for any regression can be measured as  $N/1,200$ , where  $N$  is given in Table 3.

<sup>17</sup> For brevity, we do not discuss in any detail the point estimates on the conditioning variables. It may be noted, however, that the negative (and often insignificant) effect from income per capita that recurs in these regressions is consistent with the findings of others (see for instance Rodrik (1998), and the potential explanation of this in Keen and Lockwood (2006)). That the coefficient on openness is negative and generally insignificant is more surprising, with Rodrik (1998), notably, finding a positive association even amongst higher income countries: it appears that this relationship is not present within the OECD subset of this group.

<sup>18</sup> This is the "within groups" estimator.

<sup>19</sup> The p-value for this test is 0.29, indicating that the null of no second-order autocorrelation is easily accepted. Because the Arellano-Bond estimator estimates a first difference of [3], this indicates that there is no first-order serial correlation in the  $u_{it}$ .

<sup>20</sup> The potential role of this variable in regressions of this kind is discussed in Keen and Lockwood (2006). Note that because we also have country fixed effects in all regressions, the baseline effect of a federal country on revenue is unidentified, as this dummy has no variation over time for any country in the sample.

Columns 6 and 7 test the robustness of these results. Regression 6 adds additional controls: the (log of the) of population size (*POP*), demographic variables (the proportions of the population 65 or older (*DEPOLD*) and 14 or younger (*DEPYOUNG*), and a dummy variable recording whether the country was in an International Monetary Fund (IMF) crisis program. Regression 7 addresses a possible concern that there are relatively few years of observations on the OECD dataset on some of the newer members, dropping any country from the sample for which there are less than 30 years of observations of the tax to GDP ratio.<sup>21</sup>

In column 6, the VAT terms fail the test for joint significance at five percent, although only somewhat marginally. In the final regression, the VAT terms are again jointly strongly significant. Moreover, some terms are also individually significant. First, there is a significant positive baseline effect from the simple VAT dummy—and it is large, implying an increase in the tax ratio of just under five percent of GDP. This is mitigated, however, by a significantly negative interaction effect through the income variable, implying a smaller gain in higher income countries. This is a surprising finding, the conventional wisdom (which Keen and Lockwood (2006) verify on a wider set of countries) being that the gain from the VAT is likely to be larger at higher income levels, a common argument being that in this respect income serves as a good proxy for capacity to administer and comply with the VAT. One possible interpretation of the results here is that such effects are significant only up to some basic level of capacity that is readily met by all OECD members. As one would expect given the political and technical difficulties of applying the VAT to farm-

ing, the apparent gain from the VAT tends to lower where the agricultural sector is larger. Though never individually significant, the robustly positive coefficient on the interaction of the *FED* dummy is also striking, and suggestive perhaps—this too is no more than speculation—that the technical necessity<sup>22</sup> of adopting a VAT at the central rather than lower level makes it a useful device for avoiding erosion of the tax base that may otherwise arise from allocating tax powers to lower-level governments. More puzzling is the interaction with openness. Conventional wisdom is that the VAT works best in more open economies, since there is then a large import base on which the tax can readily be levied. In these regressions, however, the coefficient on the interaction with openness not only proves to be individually insignificant but also varies in sign. There is, thus, no suggestion of such effects at work within the OECD countries.

As noted above, the pattern of sign effects means that the direction of the revenue effect associated with the VAT is in principle uncertain, depending on country-specific characteristics. A natural way to explore this is by evaluating the revenue gain from the presence of a VAT at the mean values of the controls,  $\bar{X}$  for those countries and years in which a VAT is not in place; that is, to calculate  $\Delta R = \alpha + \beta' \bar{X}$ , which is the predicted gain from the adoption of the VAT by a “typical” country in the sample without a VAT. This gain does, indeed, prove to be positive. For example, expressed as a percentage of that hypothetical country’s tax ratio, for the specification in column 7 it is 1.6 percent, which is modest but not insubstantial. For specifications 5 and 6, the gains are 2.1 and 0.5 percent, respectively.

<sup>21</sup> This removes the Czech Republic, Hungary, Iceland, Mexico, Poland, the Slovak Republic, and Turkey.

<sup>22</sup> Or perhaps not: see for instance Keen and Smith (2006). In any event, no OECD country allocates design powers to lower level governments.

Overall, then, there are signs that, within the OECD, the VAT has indeed proved to be a “money machine” in the weak sense. Though the evidence is not overpowering, and the impact of the VAT appears to be sensitive to country-specific characteristics, the presence of a VAT does seem to have a significant contemporaneous effect on the tax ratio, and, for the “typical” OECD country, it is positive—but it is also quite modest.

*A Strong Money Machine?*

With there thus being some evidence in support of the weak form of the money machine hypothesis, attention turns next to the strong form: Is there any evidence that the rise of the VAT has been a cause of increased government size or is it better seen as a consequence? As discussed above, there are broadly two ways of approaching this question empirically.

*Has the VAT Granger-caused the Growth of Government?*

The first approach is to test for causality in the statistical Granger-sense: variable *X* “Granger-causes” variable *Y*, recall, if lagged values of *X* are significant when regressed on current and lagged values of *Y*. Subject to some well-known qualifications,<sup>23</sup> which are not likely to be relevant here, Granger-causality tests in a well-defined sense for causality between economic variables.

To implements this, we run a two-variable unrestricted vector autoregression in total tax revenue (*R*) and VAT revenue (*RV*), both relative to GDP, using the panel data set described above. Generally, the regressions run were:

$$[4] \quad R_{it} = \alpha_0 + \alpha_1 R_{i,t-1} + \alpha_2 R_{i,t-2} + \beta_1 RV_{i,t-1} + \beta_2 RV_{i,t-2} + \pi_i + u_{it}, \text{ and}$$

$$[5] \quad RV_{it} = \gamma_0 + \delta_1 RV_{i,t-1} + \delta_2 RV_{i,t-2} + \phi_1 R_{i,t-1} + \phi_2 R_{i,t-2} + \theta_i + \omega_{it},$$

where  $\theta_i$  and  $\pi_i$  are country fixed effects, and  $u_{it}$  and  $\omega_{it}$  are random errors, assumed i.i.d. The optimal lag lengths were chosen using Akaike’s Information Criterion (AIC). Specifically, for each of the regressions in Table 4 below, we considered four possibilities—with either one or two lags of each of *R* and *RV*—and report the specification that minimizes the AIC. Regressions 1 and 2 have no additional controls other than country dummies, while 3 and 4 introduce in addition the standard controls for a tax effort equation, as discussed above. Note that when the AIC specifies that two lags of the non-dependent variable should be included, (as in regressions 2–4), we test for the joint significance of these lags using an F-test, as reported in the table.

Without controls, there appears to be two-way Granger causality (at the standard five percent significance level) between *R* and *RV*: that is, lagged values of *RV* help determine *R* and vice-versa (for the latter case, the F-test shows the coefficients on lagged revenue to be jointly significant, though not individually so). When country controls are introduced, however, causality runs only one way, from total revenue to VAT revenue. At ten percent, however, two-way causality cannot be rejected.

In a purely statistical sense, there is, thus, no strong evidence that the VAT has in itself caused the growth of government.

*Has Increased VAT Revenue Been Offset by Reduced Revenue from Other Taxes?*

The second approach is more structural, exploiting the result established above:

<sup>23</sup> The qualification is that if agents choose *X* in anticipation of future values of *Y*, with the expectation of the latter based on its own past values, then lagged values of *Y* will appear to Granger-cause *X*, even though true causation runs the other way. For an example of this kind, see Hamilton (1994, p. 306).

**TABLE 4**  
GRANGER–CAUSALITY TESTS<sup>1</sup>

Dependent variable	1	2	3	4
	R	RV	R	RV
R-1	0.916 (16.47)**	0.007 (0.29)	0.889 (14.31)**	-0.004 (0.14)
R-2	-0.016 (0.30)	0.025 (1.16)	-0.031 (0.55)	0.035 (1.50)
RV-1	0.07 (2.66)**	0.865 (26.12)**	-0.01 (0.15)	0.935 (21.31)**
RV-2			0.078 (1.25)	-0.074 (2.42)*
Ln(YPC)			-0.487 (1.36)	-0.323 (1.99)*
POP			0.007 (0.89)	0 (0.15)
OPEN			-0.729 (1.91)	-0.204 (0.81)
AGR			-0.104 (3.92)**	-0.041 (2.59)**
Observations	971	969	848	847
R <sup>2</sup>	0.98	0.95	0.98	0.95
F-test for Granger causality	n.a.	F(2, 936) = 3.76 Prob > F = 0.0236	F(2, 810) = 2.36 Prob > F = 0.0951	F(2, 809) = 3.21 Prob > F = 0.0409

Note:

<sup>1</sup>Robust t-statistics in parentheses; country dummies included in all regressions; and \*\* indicates significance at 1 percent, \* at 5 percent.

If the greater efficiency of the VAT itself explains the growth of government, then any increase in total revenue should be less than that from the VAT itself, with that greater efficiency reflected in part in reduced reliance on other forms of tax.

Is there any sign that there has, indeed, been such offsetting in OECD countries? To explore this, we estimate a variety of specifications of the general form:

$$[6] \quad R_{it} = \delta R_{i,t-1} + \gamma_V RV_{it} + \sigma V_{it} + \gamma' Z_{it} + \mu_i + \zeta_t + \xi_{it},$$

with *RV*, as before, denoting revenue from the value added tax (as a share of GDP) *Z*; a vector of additional variables  $\delta, \gamma, \sigma$  and  $\gamma$  parameters to be estimated; and the last three terms again being country- and time-effects and an idiosyncratic error. The dataset used for this exercise is the

same unbalanced panel of all current OECD members as used above. Note that we include all observations in this estimation, including those in which no VAT was present, and include the simple VAT dummy *V* to allow for an effect on other revenue from the presence of the VAT that is independent of the revenue it raises—a common claim, for example, is that implementation of a VAT also provides information useful for the enforcement of the personal income tax. This device also provides a simple way of allowing for some non-linearity in the relationship between total and VAT revenues.

Interest centers on the extent to which revenue raised by the VAT is offset (or, conversely, matched) by reductions (increases) in revenue from other taxes. Once a VAT has been adopted, this is given by  $\gamma_v$  in the short run and by

$$[7] \quad \phi \equiv \frac{\gamma_V}{1-\delta}$$

in the long run. Thus  $\phi = 1$  corresponds to zero offsetting of increased VAT revenues, at the (intensive) margin, with  $\phi < 1$  corresponding to some marginal offset and  $\phi > 1$  to increases in revenue from the VAT being accompanied by increased revenue from other taxes too. Assessing the full

revenue effect of the VAT also requires taking account, however, of any discrete effect  $\sigma$  from its presence, a point to which we shall return.

Results are reported in Table 5.<sup>24</sup> The first column reports OLS estimates of [6], with the lagged dependent variable suppressed. This suggests that a one-point increase in the revenue raised by an existing VAT, rela-

TABLE 5  
RELATING TOTAL REVENUE TO VAT REVENUE<sup>1</sup>

	1 <sup>2</sup>	2	3	4
<i>R-1</i>		0.812** (0.017)	0.816** (0.017)	0.836** (0.181)
<i>RV</i>	0.598** (0.078)	0.172** (0.038)	0.172** (0.039)	0.137** (0.041)
<i>V</i>	-2.414** (0.544)	-0.835** (0.258)	-0.795** (0.258)	
<i>ln(YPC)</i>	-5.990** (1.041)	-0.980* (0.384)	-0.817* (0.0368)	-1.390** (0.444)
<i>OPEN</i>	-2.726** (0.951)	-0.980* (0.401)	-0.894* (0.408)	-0.687 (0.431)
<i>AGR</i>	-0.432** (0.057)	-0.095** (0.030)	-0.103** (0.027)	-0.163** (0.037)
<i>Ln(POP)</i>	1.225 (0.881)	0.346 (0.458)		
<i>DEPOLD</i>	71.48* (12.619)	17.719** (5.617)	19.785** (4.964)	13.266* (5.795)
<i>DEPYOUNG</i>	-7.577 (7.175)	-3.042 (3.736)		
$\phi$	0.598** (0.078)	0.913** (0.191)	0.935** (0.196)	0.835** (0.245)
Observations	864	825	825	630
R <sup>2</sup>	0.944			
Serial correlation	F(1,29)=1,159.77 p=0.000			
Sargan <sup>3</sup>		1.000	1.000	1.000
$m_1$ <sup>3,4</sup>		0.000	0.000	0.000
$m_2$ <sup>3,4</sup>		0.239	0.233	0.962

Notes:

<sup>1</sup>Both in percent of GDP; robust z-statistics in parentheses; \*\* indicates significance at 1 percent, \* at 5 percent.

<sup>2</sup>Country and time dummies included (the former in all regressions) but not reported.

<sup>3</sup>p-values.

<sup>4</sup>The  $m_1$  and  $m_2$  statistics test for first- and second-order serial correlation in the equation estimated in first differences, with the former present and the latter absent if the equation is well-specified.

<sup>24</sup> Again, we shall not discuss in any detail the point estimates on the other variables. That the coefficient on openness is more firmly negative than in Table 3 reinforces the surprising result there, but is consistent with the common presumption—some evidence for which is given in Ebrill et al. (2001)—that the ease of bringing imports into the VAT makes it an especially effective tax in more open economies.

tive to GDP, is generally offset by a reduction in revenue from other taxes of about 0.4 points, so that while revenue increases, it does so by only 0.6 points of GDP. The presence of a VAT in itself, however, has a significantly negative impact on total revenue, suggestive of non-linearity in the relationship. But the F-test also indicates significant first-order serial correlation, pointing again to more complex dynamics. A further concern is the potential for endogeneity bias arising from common shocks to VAT and total revenue.

The second column, therefore, introduces the lagged dependent variable, using the Arrelano–Bond (1991) GMM estimator so as to deal with the potential bias from the lagged dependent variable, and addressing the endogeneity issue by treating VAT revenue as predetermined. Now the degree to which increased revenue from an existing VAT is offset appears to be much greater in the short term, and much smaller in the long term. The point estimate of  $\phi$  is significantly different from zero—so that marginal increases in VAT revenue are indeed associated with increases in total tax revenue—but is also less than one. This suggests that increases in VAT revenue have not simply occurred in tandem with increases in revenue from other taxes but rather, at the margin, have been used to reduce reliance on these alternatives. Note too that the discrete impact of the VAT dummy remains significantly negative.

Column 3 reports the results of eliminating the variables in column 2 that proved insignificant at ten percent, the results being broadly unchanged. Finally, column 4 reports the same specification estimated only on observations for which a VAT was in place: as one would expect given the negative discrete effect of the VAT, the

marginal replacement that emerges now appears somewhat lower.

What does this imply for the strong money machine hypothesis, the key prediction of which, recall, is that revenue from the VAT will be in part offset by reduced revenue from other sources?

At the margin—that is, considering increased revenue from an existing VAT—this does, indeed, appear to be the case, though the degree of offset is fairly small: the point estimate of  $\phi$  is fairly robustly less than unity. The hypothesis that it equals unity cannot be rejected, but the estimates are far from being so much in excess of unity that a demand-led explanation of marginal increases in VAT revenue—that this has been just one way in which a stronger taste for government has been met—appears clearly the less plausible of the two. The discrete negative revenue impact of the presence of a VAT, suggestive of an underlying non-linearity, complicates but does not overturn this interpretation. To see this, note first that [3] implies the long-run impact on total revenue of introducing a VAT that raises revenue  $RV$  to be

$$[8] \quad \Delta R = \frac{\gamma_V RV + \sigma}{1 - \delta}.$$

Using the point estimates in column 3,  $\Delta R$  will, thus, be positive so long as revenue from the VAT exceeds about 4.6 percent of GDP—which, as can be seen from Table 5, it indeed does in almost all OECD countries. The results are, thus, broadly consistent with the VAT at least having been a net addition to revenue. But—and consistent with our earlier results on the weak money machine, reported in Table 3—this addition may in many cases be quite small, since the degree to which increased VAT revenue has been offset by reductions in other taxes has tended to be quite large.<sup>25</sup> Consider, for

<sup>25</sup> While this runs counter to Kenny and Winer's (2006) empirical support for the scale effect—higher total revenue being found there to be associated with higher revenue, relative to GDP, from all main tax categories—their work focuses on a much wider range of countries and does not distinguish between the VAT and other taxes on goods and services.



example, the “average” OECD country, which (from Table 1) collects about 7.2 percent of its GDP in VAT revenue. Using again in [6] the point estimates of column 3, the associated long-run increase in total revenue is about 2.4 percent of GDP: around two-thirds of the revenue raised by the VAT is, thus, offset by reduced revenue from other taxes.

### POLITICS AND MONEY MACHINES

Suppose then that, as the results above suggest is, indeed, the case, the VAT is, indeed, a money machine, in the sense of being an especially effective form of taxation. How persuasive are the political economy arguments that it would, as a consequence, be a good idea to prevent a government from adopting one?

The clearest statement of the view that it may be wise to preclude the use of efficient tax instruments is that of Brennan and Buchanan (1977). The essence of their argument is that the citizenry may benefit by imposing restrictions, at the constitutional phase, on the set of tax instruments available to a revenue-maximizing Leviathan who will be essentially unconstrained in the post-constitutional phase.<sup>26</sup> In this way, they can beneficially limit the resources that the Leviathan will be able to extract from them.

This line of argument has proved extremely influential. It clearly reflects, however, a quite restrictive view of both

the objectives of policy makers and the constraints, notably electoral, under which they operate. The more recent political economy literature, largely spawned by this work, suggests a series of insights as to how these further considerations may affect the case for entrusting policy makers with effective tax instruments.<sup>27</sup>

Consider first the implications of simply relaxing the view that policy makers attach no explicit value to the welfare of the citizenry. Suppose, for example, that policy makers seek to maximize some function  $W(C,U)$  defined not only over the tax revenue  $C$  that they can divert to their own private benefit (which is the sole concern of the classic Leviathan), but also, at least to some degree—perhaps only in order to deter their own overthrow—about the welfare of the citizenry,  $U$ . To see the implications, note that in the framework above (now, for simplicity, assuming a single tax instrument), private utility would then be

$$[9] \quad U = V(R - C) - \left(\frac{1}{2}\right)\theta R^2.$$

Consider now how an increase in the efficiency of the tax system—a reduction in  $\theta$ —affects the policy maker’s possibility frontier in  $(C,U)$ -space. Clearly, it shifts unambiguously outwards:<sup>28</sup> a more efficient tax instrument enables policy makers to leave the citizenry better off for any given level of resources enjoyed by themselves. So long as  $U$  is normal in the policy makers’ preferences, this income

<sup>26</sup> The formal structure of their argument is simple and, by the standards of the later literature, ad hoc. Knowing that policy makers will divert some fixed proportion of tax revenue to their own use, citizens restrict the tax instruments available to them in such a way that the maximum revenue which can be raised, net of this diversion, will be just such as to finance their desired level of public spending.

<sup>27</sup> One vein of the literature not pursued in detail here focuses on the role of interest groups. Becker and Mulligan (2003), in particular, explore a framework in which the size of government is determined by non-cooperative strategic interaction between taxpayers and the beneficiaries of government spending, both of whom can expend resources to affect levels of revenue and spending. Like us, they consider the effect of a change in the efficiency of tax instruments on the equilibrium size of government, finding that an exogenous increase in the efficiency of the tax system does not necessarily make both groups better off. This has the same flavor as the results here, and in that sense is consistent with our broad conclusions, but the mechanism at work is entirely different.

<sup>28</sup> This and other results asserted in this paragraph and the next are proved in Appendix B.

effect of the tax innovation thus leads to an increase in citizens' well-being.

But there is also a substitution effect at work, and this tends to act in the opposite direction: raising the revenue to finance a marginal increase in  $C$  now has a smaller efficiency cost, and so requires less of a reduction in  $U$ . Thus, as well as shifting outwards, the possibility frontier thus becomes flatter (here visualizing  $C$  as being on the horizontal axis), which in itself inclines policy makers to reduce  $U$ . The extent of this flattening, it can be shown, is greater the more rapidly citizens' marginal valuation of the spending from which they benefit,  $V'$ , declines with the level of spending: intuitively, the increased provision of the public good made possible, for given  $C$ , by the increased efficiency of the tax system then leads to a greater reduction in that marginal valuation and, hence, increases the rent extraction that the policy maker must forego in order to achieve some given increase in private utility.

The impact on citizens' utility of access to a money machine is, thus, in this simple case, ambiguous. Broadly speaking, it is more likely to be positive the greater is the income elasticity of the policy makers' demand for citizens' utility, and the smaller is  $V''$ . It is hard to translate these concepts into hard numbers. What is clear, however, is that even with policy makers who look largely to their own narrow interests, allowing them access to efficient tax instruments may well increase citizens' welfare. Indeed, this is sure to be the case, for example, if  $V'' = 0$ , since then (as the intuition above suggests and is readily verified) the slope of the possibility frontier remains unchanged as  $\theta$  falls, so that the substitution effect vanishes. In this sense, even an only slightly less pessimistic view of policy makers' preferences

can imply a much more optimistic outlook for the consequences of entrusting them with a money machine.

Policy makers may also be faced with a series of constraints under which they operate. A natural response to a fear that government is inclined to tax and spend too much, for instance, is to impose direct limits on the level of spending, so providing some protection whilst also enabling whatever level of revenue is needed to be raised in the most effective way. Several countries, such as Sweden for example, do exactly this.

Elections, of course, are a key device for restraining abusive policy making. What then if policy makers have to face elections?

Consider first the Downs (1957) model of electoral competition, the key feature of which is that successful candidates for office are obliged to implement the policies that they announce before the election. More precisely, suppose that, in the notation above, two identical non-benevolent candidates  $i = A, B$  simultaneously propose policies  $(R_i, C_i)$ , and then an election follows. Candidates are precommitted to implementing their proposed policy if elected, care about holding office in itself—from which they derive some positive non-monetary "ego-rent"  $E$ —and have preferences over policy given by  $W(C, U)$  if they win the election (with a payoff of zero if not elected). Elected politicians' interests are, thus, not fully aligned with those of the electorate. Nevertheless, if voters care only about policies (that is, have no bias, for ideological or other reasons, in favor of one candidate or another), and have identical preferences, then it is easily seen that the only possible equilibrium outcome<sup>29</sup> is one that maximizes voter utility  $U$ . Therefore, in this equilibrium,

<sup>29</sup> To see this, suppose first that both candidates propose  $(\hat{R}, \hat{C})$  that maximizes  $U$ . Each wins the election with a probability of one-half. If one candidate deviates to some other policy, he will certainly lose the election and so his payoff must fall. Thus,  $(\hat{R}, \hat{C})$  is certainly an equilibrium. A similar argument implies that it is unique.

there is no rent-diversion in equilibrium. In the absence of voter biases, electoral competition completely eliminates the discretion that policy makers have to exploit the population—and an increase in the efficiency of available tax instruments undoubtedly benefits the voters.

This conclusion is modified if some or all of the voters are in part motivated by factors other than the policies at immediate issue. To take a very simple example (based on Dixit and Londregan (1996)), suppose now that voters have an identical ideological preference parameter<sup>30</sup> in favor of (say) candidate *A*, which is drawn from a uniform distribution on  $[-B, B]$ . Thus *B* measures the ex ante degree of voter bias: the greater is *B*, the greater is the expected bias of all voters in favor of one candidate or the other.<sup>31</sup> The sequence of events we consider is again the Downsian one: politicians first simultaneously propose policies, to which they are committed; the ideology parameter is then realized; and the vote then takes place. It is easily calculated<sup>32</sup> that in this case there is a symmetric equilibrium in which both candidates propose a policy ( $C^*, R^*$ ) and each is elected with ex ante probability (prior to the realization of the ideology parameter) of one-half.

The key property of the political equilibrium that emerges in this case, for present purposes, is that (under a weak technical condition) rent-extraction  $C^*$  is strictly positive—in sharp contrast to the simple case above—but less than it would be without elections. Intuitively, the presence of an ideological bias provides some cover behind which policy makers can extract surplus for themselves without

doing excessive damage to their electoral prospects; at the same time, the prospective ego-rents provide an incentive not to jeopardize those prospects by paying too little attention to voters' well-being. Thus, rent diversion can be shown to be lower the greater are the ego-rents from office, *E*, and the less biased are voters (the lower is *B*). The stronger is electoral competition (the higher is  $E/B$ ), the more likely it is, other things equal, that an increase in the efficiency of the tax system will translate in equilibrium into a welfare gain for voters. Loosely put, if politicians are self-important rather than venal, or if citizens vote largely on policies rather than personalities, then the case for denying them efficient tax instruments is weakened.

An unattractive feature of the Downsian framework, however, is the assumption that electoral candidates can precommit to pursue particular policies if elected (with the further and unrealistic implication in the present context that the degree of rent extraction *C* must be observable), irrespective of their own preferences. This in turn precludes any role for such realistic behavior as voting based on the past performance of the incumbent. Both of these features can be relaxed. Besley and Smart (2003), in particular, consider a simple framework of this kind in which there are two types of politicians—some pure Leviathans, concerned only with the surplus *C* they can extract from themselves, some wholly benevolent—competing for office in a world with a two-period term limit. Voters do not directly observe politicians' types, and while they can observe the taxes they pay and the public services

<sup>30</sup> This is, of course, unrealistic and means that in equilibrium all voters vote in the same way: almost always, some candidate will get 100 percent of the votes, because the common bias of the voters will generically differ from the difference in voter utility from the two policy proposals. This could be avoided, as in Dixit and Londregan (2002), by also introducing individual-specific randomness to smooth the outcome, but this also complicates the model.

<sup>31</sup> Denoting by  $\beta$  the bias variable distributed on  $[-B, B]$ , the bias in favor of some candidate (either *A* or *B*) is simply the absolute value of  $\beta$ , which, given the uniform distribution of  $\beta$ , has expected value of  $B/2$ .

<sup>32</sup> For the proof of this, and of the claims that follow, see Appendix C.

they enjoy, they cannot observe the cost of providing those services or, hence, the surplus that the incumbent policy maker extracts for themselves.

There are then broadly two types of outcome, depending on the parameters of the model. One possibility is that Leviathan incumbents “go for broke,” extracting as much revenue as they can when in office<sup>33</sup> and accepting that in doing so they will reveal their identity as Leviathans, and consequently not be re-elected (a separating equilibrium). The other possibility is that Leviathan incumbents will restrain the amount of revenue they raise so as to mimic the behavior of a benevolent policy maker faced with an adverse cost shock, and so improve their chances of being re-elected and extracting as much surplus as they can in a final period of office (a pooling equilibrium).

Within this framework, Besley and Smart (2003) directly address the question of interest here: might an increase in the efficiency of the tax system actually reduce voter welfare (evaluated *ex ante* before the type of the first-period incumbent is known)? A key result is that this cannot be the case if the nature of the equilibrium does not switch from pooling to separating or vice versa: in a separating equilibrium, for example, an increase in tax efficiency generates an evident gain for voters if the incumbent policy maker is benevolent—and if they are not, they continue to simply raise and spend on themselves as much revenue as they can. Interestingly, this conclusion of an unambiguous welfare gain from access to a more efficient tax instrument rests on an assumption that  $V'' = 0$  that was seen earlier to be sufficient to ensure a welfare gain in the simple model of unconstrained but partly benevolent policy makers above.

An increase in the efficiency of the tax system may, however, *reduce* voter welfare if it leads to a change in the nature of the equilibrium. Since such an increase makes it more attractive for Leviathan to mimic a benevolent policy maker—the later would now choose a higher level of public good provision, which enables the former to extract more rent by pretending that its cost has proved high—the relevant possibility is a shift from separating to pooling. The additional discipline this exerts on an incumbent Leviathan clearly benefits the voter. Against this, however, the electoral process now becomes less effective at removing Leviathans (since they no longer reveal themselves), and so creating more risk of abuse in the final term of office. This source of loss is greater the higher the likelihood that a candidate with no record of office would prove to be benevolent.<sup>34</sup> For this reason—and counter perhaps to simple intuition—an increase in tax efficiency that shifts the qualitative equilibrium is more likely to reduce voter welfare the *fewer* politicians are potential Leviathans.

## CONCLUSIONS

The empirical analysis of the OECD experience reported here suggests that the answer to the question posed in our title is: “Yes, but....” The VAT does, indeed, appear to have been a “money machine” in both senses of the term defined here. It seems to have been a money machine in the weak sense that countries with a VAT tend to raise more revenue, all else equal, than do those without. And it seems also to have been a money machine in the stronger sense that, although the VAT does not appear to have statistically “caused” an increase

<sup>33</sup> The Besley–Smart model has an exogenous upper limit on the amount of revenue that can be raised.

<sup>34</sup> It is also greater the lower is the voter’s discount rate, since then the present value cost of a future unrestrained Leviathan is greater.

in government size, the revenue that it raises has to some degree been offset by reduced revenues from other taxes—suggesting that its use has been driven largely by the desire to exploit its greater effectiveness rather than by generalized pressures to finance bigger government. The primary “but”—there are others—is that the association between the presence of the VAT and total tax revenue is not simple (but rather depends on country circumstances), is not always statistically significant at five percent (though it usually is, and failures are fairly marginal), and may in any event be small. This relative weakness of the evidence for the weak form of the money machine hypothesis is consistent, however, with the relative strength of that for the strong form. The picture that emerges is that the VAT has proved to be a particularly effective form of taxation, but the impact of this on the overall size of government has been substantially diluted—making evidence for the weak form harder to detect—by a tendency to take these gains in large part in the form of reduced use of less effective tax instruments.

As for politics, one certainly find cases in which access to a more efficient tax instrument, such as the VAT, reduces citizen/voter welfare. But these seem to us to be somewhat strained. This conclusion emerges, in particular, from the discussion above of two models of the political process that capture some essentials of the debate, and have particular resonance in the US context. Both have the feature that politicians cannot precommit to pursue particular policies if they come to office. In one, they are also electorally unconstrained, but attach some weight not only to the surplus that—in Leviathan spirit—they can extract for themselves, but also from citizens’ welfare. In the other, politicians differ in type—some

are pure Leviathans, the others, wholly benevolent—but are subject to re-election to a second and final term. In both models, and subject to one qualification for the latter, more efficient tax instruments lead to an increase in citizen/voter welfare:<sup>35</sup> the increase is not as large as it would be if all politicians were benevolent, but, nevertheless, it is an increase. The qualification is that in the second model, welfare can (but need not) fall if increased tax efficiency changes the nature of the political equilibrium in such a way that past performance in office becomes uninformative as to a politician’s true type (and, hence, likely behavior in a final term of office). And this cost can be large enough to outweigh other potential benefits only if there is, indeed, a high-enough chance that a random candidate would have proved to be “good.” That is, precluding the use of a more tax efficient instrument can reduce welfare only if the number of potential Leviathans is relatively small. In that sense, the more pessimistic is one’s view of politicians, the weaker is the case for such restriction. No doubt there exist other models that will yield different conclusions. But the presumption that emerges, for us at least, is that if the VAT is, indeed, a money machine—as our empirical results suggest to be the case—then that is an argument for, not against, its adoption.

### Acknowledgments

We are grateful to Rosanne Altshuler, Isaias Coelho and Victoria Perry for many helpful suggestions, and to Assi Woldemariam for diligent research assistance. Views and errors are ours alone, and should not be attributed to the International Monetary Fund, its Executive Board, or its management.

---

<sup>35</sup> So long, that is, as their marginal valuation of the public spending from which they benefit does not decline too rapidly.

REFERENCES

Arrelano, Manuel, and Stephen Bond.  
 "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies* 58 No. 2 (April, 1991): 277–97.

Becker, Gary S., and Casey B. Mulligan.  
 "Deadweight Costs and the Size of Government." *Journal of Law and Economics* 46 No. 2 (October, 2003): 293–340.

Besley, Timothy, and Michael Smart.  
 "Fiscal Restraints and Voter Welfare." London School of Economics. Mimeo, 2003.

Brennan, Geoffrey, and James M. Buchanan.  
 "Towards a Tax Constitution for Leviathan." *Journal of Public Economics* 8 No. 3 (December, 1977): 255–73.

Dixit, Avinash, and John Londregan.  
 "The Determinants of Success of Special Interests in Redistributive Politics." *Journal of Politics* 58 No. 4 (November, 1996): 1132–55.

Downs, Anthony.  
*An Economic Theory of Democracy*. New York: Harper and Row, 1957.

Ebrill, Liam, Michael Keen, Jean–Paul Bodin, and Victoria Summers.  
*The Modern VAT*. Washington, D.C.: International Monetary Fund, 2001.

Emran, Shahe M., and Joseph E. Stiglitz.  
 "On Selective Indirect Tax Reform in Developing Countries." *Journal of Public Economics* 89 No. 4 (April, 2005): 599–623.

Hamilton, James D.  
*Time Series Analysis*. Princeton, NJ: Princeton University Press, 1994.

Hicks, John.  
*Value and Capital*. Oxford: Oxford University Press, 1939.

Keen, Michael, and Ben Lockwood.  
 "The Value Added Tax: Its Causes and Consequences." International Monetary Fund. Mimeo, 2006.

Keen, Michael, and Stephen Smith.  
 "What Do We Know about VAT Evaluation?" University College London. Mimeo, 2006.

Kenny, Lawrence W., and Stanley L. Winer.  
 "Tax Systems in the World: An Empirical Investigation into the Importance of Tax Bases, Administration Costs, Scale and Political Regime." *International Tax and Public Finance* 13 (2006): 181–215.

Nellor, David.  
 "The Effect of the Value–Added Tax on the Tax Ratio." Working Paper No. WP/87/47. Washington, D.C. International Monetary Fund, 1987.

President’s Advisory Panel on Federal Tax Reform.  
*Simple, Fair & Pro–Growth: Proposals to Fix America’s Tax System*. Washington, D.C., 2006.

Rodrik, Dani.  
 "Why Do More Open Economies Have Bigger Governments?" *Journal of Political Economy* 106 No. 5 (October, 1998): 997–1032.

Stockfish, J.A.  
 "Value Added Taxes and the Size of Government: Some Evidence." *National Tax Journal* 38 No. 4 (December, 1985): 547–52.

Wooldridge, Jeffrey M.  
*Econometric Analysis of Cross–Section and Panel Data*. Cambridge, MA: MIT Press, 2002.

APPENDIX A

Perturbing [2] and setting  $\lambda = 1$  gives:

$$[A1] \begin{bmatrix} V'' - \theta_A & V'' \\ V'' & V'' - \theta_B \end{bmatrix} \begin{bmatrix} dR_A \\ dR_B \end{bmatrix} = \begin{bmatrix} -V' & R_A \\ -V' & 0 \end{bmatrix} \begin{bmatrix} d\lambda \\ d\theta_A \end{bmatrix},$$

from which

$$[A2] \begin{bmatrix} dR_A \\ dR_B \end{bmatrix} = \frac{1}{\Delta} \begin{bmatrix} -(V'' - \theta_B)V' + V''V' & (V'' - \theta_B)R_A \\ V''V' - (V'' - \theta_A)V' & -V''R_A \end{bmatrix} \begin{bmatrix} d\lambda \\ d\theta_A \end{bmatrix}$$

$$[A3] = \frac{1}{\Delta} \begin{bmatrix} \theta_B V' & (V'' - \theta_B)R_A \\ \theta_A V' & -V''R_A \end{bmatrix} \begin{bmatrix} d\lambda \\ d\theta_A \end{bmatrix},$$

where  $\Delta \equiv (V'' - \theta_A)(V'' - \theta_B) > 0$ . Thus, as claimed in the text,

$$[A4] \quad \frac{dR_i}{d\lambda} = \left( \frac{\theta_i V'}{\Delta} \right) > 0;$$

$$\frac{dR_A}{d\theta_A} = \frac{(V'' - \theta_B)R_A}{\Delta} < 0;$$

$$\frac{dR_B}{d\theta_A} = \frac{-V''R_A}{\Delta} > 0.$$

The last term shows, in particular, that an increase in the efficiency of one instrument optimally reduces the revenue raised by the other. It, nevertheless, increases total revenue, as one would expect, since the last two parts of [A4] imply that:

$$[A5] \quad \frac{dR}{d\theta_A} = \frac{-\theta_B R_A}{\Delta} < 0.$$

**APPENDIX B**

Substituting from [6] into  $W(C,U)$ , the necessary conditions for the policy maker’s choice of  $R$  and  $C$  are

$$[B1] \quad V'(R - C) - \theta R = 0, \text{ and}$$

$$[B2] \quad -W_U(C,U)V' + W_C = 0.$$

To establish the effect of a change in  $\theta$  on the possibility set in  $(C,U)$ -space, note that [B1] defines  $R(C,\theta)$ , with

$$[B3] \quad R_\theta = \frac{R}{V'' - \theta}.$$

Writing the frontier as

$$[B4] \quad u(C,\theta) \equiv U[R(C,\theta) - C, C] - \frac{1}{2}\theta[R(C,\theta)]^2$$

the results claimed in the text follow on noting, using [B1], that

$$[B5] \quad u_\theta = -\frac{1}{2}R^2 < 0,$$

and that the slope of the frontier is given by

$$[B6] \quad u_C = -V',$$

so that, using also [B3],

$$[B7] \quad u_{C\theta} = -V''R_\theta = \frac{-V''R}{V'' - R} < 0.$$

That a reduction in  $\theta$  will lead to an increase in the citizen’s welfare if  $V'' = 0$  is intuitively apparent from [B6], which shows that in this case the slope of the possibility frontier is independent of  $\theta$ . To see the point somewhat more formally, note in this case perturbing [B2], substituting for  $V'$  from [B2] and collecting terms gives:

$$[B8] \quad dU = \frac{(W_{CC} - W_{UC}(W_C / W_U))}{(W_C / W_U)(W_{UU} - W_{CU}(W_U / W_C))} dC.$$

If both  $C$  and  $U$  are normal in policy maker’s preferences, the terms in numerator and denominator on the right of [B8] are both negative (see, for example, Hicks (1939)), so that  $dU$  and  $dC$  have the same sign. Since a reduction in  $\theta$  raises the policy maker’s welfare—recall that the possibility set unambiguously expands—it follows that  $U$ , along with  $C$ , must increase.

**APPENDIX C**

In symmetric equilibrium, both candidates propose the same policies, denoted  $(C^*, R^*)$ . Given this, the payoff to candidate  $A$  (for example) from deviating to some other proposal  $(C', R')$  is calculated as follows. First, let  $\beta$  be the additional utility that any voter receives if candidate  $A$  rather than candidate  $B$  implements a particular policy; that is, their bias in favor of  $A$ . Then, given  $\beta$ , all voters will vote for  $A$  rather than  $B$  iff  $U(C', R') + \beta \geq U(C^*, R^*)$ . So, recalling that  $\beta$  is uniformly distributed on  $[-B, B]$ , the probability that  $A$  wins the election by deviating to  $(C', R')$  is

$$[C1] \quad \Pr[\beta \geq U(C', R') - U(C^*, R')] = \frac{1}{2} + \frac{U(C', R') - U(C^*, R^*)}{2B} \equiv P(C', R').$$

Assuming that a candidate who loses the election has payoff of zero,  $A$ 's overall payoff from deviating to some policy  $(C, R)$  is thus

$$[C2] \quad \Delta(C, R) = P(C, R)(w(C, R) + E),$$

where  $w(C, R) \equiv W[U(C, R), C]$ . At equilibrium, the derivatives of  $\Delta(C, R)$  with respect to the variables  $C$  and  $R$  must be zero. By straightforward calculation from [C1] and [C2], this requires

$$[C3] \quad \Delta_i(C^*, R^*) = Pw_i + \frac{\partial P}{\partial X}(W + E) \\ = \frac{1}{2}w_i + \frac{U_i}{2B}(w + E) = 0, \quad i = R, C$$

where the second equality uses that facts that at equilibrium,

$$[C4] \quad P = \frac{1}{2}, \quad \frac{\partial P}{\partial X} = \frac{U_i}{2B}, \quad i = R, C.$$

First note that since  $w_R = W_U U_R$ , [C3] implies that  $U_R = 0$  at equilibrium. So, let  $R(C)$  be the equilibrium choice of  $R$  for any  $C$ :  $R(C) = \arg \max_R \{U(C, R)\}$ . Then we can rewrite [C3], for  $i = C$ , as

$$[C5] \quad \omega'(C) + \frac{U_C(R(C), C)}{B}(\omega(C) + E) = 0,$$

where  $\omega(C) = w(R(C), C)$ . Note that [C5] determines the equilibrium value of rent-diversion,  $C^*$ .

Now, note from [6] that  $U_C = -V' < 0$ . Then, from [C5] and  $U_C < 0$ , and using  $E, B, \omega > 0$ , we see that  $\omega'(C^*) > 0$ . A policy maker who does not have to face an election, on the other hand, sets  $\omega'(\hat{C}) = 0$ . So, if we assume that  $\omega(C)$  is strictly concave,  $C^*$  must be lower than  $\hat{C}$ , as claimed in the text.

Finally, total differentiation of [C5] with respect to  $E$  and  $B$ , recalling  $U_R = 0$ , gives

$$[C6] \quad \frac{dC^*}{dE} = \frac{-U_C}{\omega''B + \omega'U_C + \omega U_{CC}}, \\ \frac{dC^*}{dB} = \frac{U_C(\omega + E)/B}{\omega''B + \omega'U_C + \omega U_{CC}}.$$

To sign the terms in [C6], recall that  $U_C < 0$ ,  $\omega' > 0$ ,  $\omega'' < 0$ , the last by assumption. Moreover, as  $U_C = V'$ , it follows that  $U_{CC} = V'' < 0$ . This, along with  $E, B, \omega > 0$ , implies that, as claimed,  $dC^*/dE < 0$  and  $dC^*/dB > 0$ .



Copyright of National Tax Journal is the property of National Tax Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.