Irish Fiscal Advisory Council

Working paper series

The Role of Elasticities in Forecasting Irish Government Revenue

Niall Conroy

Working Paper No. 14

February 2021

Suggested reference:

Conroy, N. (2021). "The Role of Elasticities in Forecasting Irish Government Revenue ". Irish Fiscal Advisory Council Working Paper Series No. 14. Dublin.

The Role of Elasticities in Forecasting Irish Government Revenue

Niall Conroy [1]

February, 2021

Abstract

This paper examines the role of tax elasticities in forecasting Irish government revenue. Tax elasticities give the expected change in revenue for a change in the tax base. We examine how important the choice of elasticity is for forecasting income tax, Value Added Tax (VAT) and Pay Related Social Insurance (PRSI) up to four-years ahead. We find that using policy-adjusted elasticities produces improved forecasts, particularly at longer forecast horizons. This improvement is statistically significant. We also find that using building and construction activity in addition to personal consumption leads to much better forecasts of VAT receipts. These two changes in methodology could lead to improved fiscal forecasts by the Department of Finance. This paper demonstrates a methodology with potential for widespread application in many other countries.

Keywords: Fiscal policy, Elasticity, Discretionary tax measures.

JEL No. E32, E62, H24, H30

This report can be downloaded at www.FiscalCouncil.ie

[1] The author is an Economist at the Irish Fiscal Advisory Council. Email: niall.conroy@fiscalcouncil.ie. The opinions expressed and arguments employed in this paper do not necessarily reflect the official views of the Irish Fiscal Advisory Council. The author would like to acknowledge the kind assistance from members of the Council and Secretariat of the Irish Fiscal Advisory Council. The author would like to acknowledge the useful comments received from Diarmaid Smyth (Department of Finance).

1. Introduction

Forecasting government revenue is key to good budgeting and setting appropriate fiscal policy. Tax elasticities measure the response of revenue to a 1 per cent change in the tax base (also referred to as the macroeconomic driver) if tax policy (for example tax rates and credits) is held fixed. If the elasticity of income tax revenue, for example, is above 1, then you would expect a 1 per cent increase in the tax base to yield an increase in income tax receipts of more than 1 per cent. Changes in tax revenue are determined by two key factors, changes in the tax base and changes in policy. Tax elasticities are key to determining the impact changes in the tax base have on government revenue. As a result, tax elasticities are key to forecasting government revenue.

The tax base is influenced by macroeconomic developments and the tax policy set by government. For example, the amount of income earned in the economy will influence the amount of income tax paid. Policy decisions around what levels of income are eligible for taxation will also impact on the tax base. Meanwhile policymakers' choices on tax rates, bands, and credits all will influence the effective tax rate. There is also feedback between government revenue collected and the macroeconomy, as a rise in tax rates would act as a drag on economic activity and vice versa.

Conroy (2020) showed how more accurate estimates of revenue elasticities can be found using a policy-adjusted revenue series. Previous estimates of revenue elasticities (in Ireland) did not account for the impact of policy changes. This may have biased estimates as policy changes were linked to the economic cycle.

This paper examines if varying the elasticity used has a significant impact on forecast accuracy for government revenue in Ireland. More specifically, would using these recently estimated policy-adjusted elasticities significantly improve forecast accuracy? To the authors knowledge, this is the first paper which examines how varying the elasticity used impacts on the accuracy of forecasts of government revenue. As a result, no other paper has tested if using policy-adjusted elasticities yields superior forecasting performance, relative to unadjusted elasticities.

This analysis is performed for three revenue headings: Income tax, Value-Added Tax (VAT) and Pay Related Social Insurance (PRSI). These three headings accounted for almost two thirds of central government revenue in 2019. We use historical data to see which elasticities would have worked best in forecasting future revenue. In all cases, revenue forecasts take account of the impact of tax policy changes. Using revenue outturns from the previous year (denoted year T-1), we compile forecasts for the current year (denoted year T) and out as far as four-years ahead (denoted year T+4). Most of the previous Irish literature has focused only on shorter forecast horizons (years T and T+1 typically).

We find that the elasticity used is important for forecast accuracy. More specifically, we find that elasticities estimated using policy-adjusted revenue yield superior forecasting performance. These improvements are statistically significant, with reduced bias and smaller absolute errors on average. These improvements are larger at longer forecast horizons, as errors cumulate. These findings are most acute for income tax. The divergence in forecasting performance reflects how policy-adjusted and unadjusted elasticities differ substantially. These elasticities differ because of large and regular income tax policy changes during the sample period examined.

This paper shows that the use of policy-adjusted elasticities could significantly improve forecast accuracy in Ireland. The implications could be wider still, as policy-adjusted elasticities could be estimated and then used for fiscal forecasting in many other countries.

Revenue forecasts are also sensitive to the tax base used. For VAT we find that the choice of tax base also has a significant impact on the accuracy of forecasts. Using building and construction activity in addition to personal consumption leads to significantly more accurate forecasts of VAT receipts. This is mainly due to the sharp fall in VAT receipts in 2009-12, which coincided with a collapse in construction activity (new house purchases are subject to 13.5% VAT). As the fall in personal consumption in this period was more modest, using construction activity as a compliment improves forecasting performance significantly.

Previous Irish studies have focused on in-year and one-year-ahead forecasts. As a result, this is the first Irish study which compares the scale of forecast errors four-years-ahead (T+4) compared to in-year (T) forecasts. We find that typically the longer the forecast horizon, the larger forecast errors are on average. In many cases, four-year-ahead forecast errors are twice as large as those for in-year forecasts.

2. Relevant Literature

There has been some recent work on forecasting Irish government revenue that this paper builds on. Hannon *et al.* (2015) examine forecast errors from tax forecasts of the Department of Finance. They find that while these forecast errors are high by international standards, there is no evidence of bias in the forecasts. They also find that the judgement applied does not introduce bias to the forecasts. One relevant finding from the paper is that, while errors in forecasting the macroeconomic driver and the starting point (the outturn of the previous year) are significant, other factors also contribute. The methodology employed and the elasticity used were also responsible for some of the forecast errors.

Fioramanti *et al.* (2016) examine the fiscal forecast errors of the European Commission. They found that forecasts for Ireland had the largest average absolute forecast errors both when forecasting the current year and one-year-ahead.

The latest Irish tax forecasting methodology review group (TFMRG) report covered the period 2007 – 2018 (TFMRG, 2019). It examined forecasts made for the current year and one-year-ahead at Budget time (which has been held in October since 2013, having previously occurred in December).

A contribution this paper makes to the Irish literature is to extend the horizon of forecasts analysed out to four-years ahead. This allows us to analyse how the size of average errors changes as the forecast horizon expands. We also use formal tests (Diebold-Mariano, 2002) of whether differences in forecasting performance are statistically significant.

The TFMRG (2019) report found that the largest forecast errors were for the years coinciding with the financial crisis (2008 and 2009). Since then, tax forecasts have tended to underestimate revenue, particularly for corporation tax.

The report recommends that VAT forecasts should be supplemented with a housing-specific component, drawing on expected trends in the housing market. A similar recommendation was made in the previous edition of the report (TFMRG, 2008), which was adopted briefly before being discontinued after *Budget 2011*.

Interestingly, this paper finds evidence that using construction activity alongside a measure of personal consumption significantly improves forecasts of VAT (see Section 4.2).

Another relevant part of the literature is studies of forecasting performance of government revenue. Buettner and Kauder (2010) review the practice and performance of forecasting government revenue in OECD countries. They find that uncertainty in forecasting macroeconomic drivers and the timing of forecasts are key factors for forecast accuracy. They also find that independence of forecasting from government has a strong positive effect on the accuracy of forecasts.

Frankel and Schreger (2013) find some evidence of optimism bias in fiscal forecasts of Euro area countries. Merola and Pérez (2013) find that European Commission and OECD fiscal forecast errors are correlated with the electoral cycle of the EU countries examined.[2]

The major contribution of the present paper is examining how sensitive forecasts of government revenue are to the elasticity applied. More specifically, we examine if elasticities based on policy-adjusted data

[2] Fioramanti et al (2016), Frankel and Schreger (2013) and Merola and Pérez (2013) all focus on the general government balance, rather than forecasts of revenue and expenditure separately.

produce superior forecasts. We also test if these improvements are statistically significant.

To the authors knowledge, this is the first paper which examines how varying the elasticity used impacts on the accuracy of forecasts of government revenue. As a result, no other paper has tested if using policy-adjusted elasticities yields superior forecasting performance, relative to unadjusted elasticities.

We find that the choice of elasticity used is an important factor in forecasting government revenue. More specifically, we find that using policy-adjusted elasticities does indeed lead to improved forecasting performance. These improvements are also found to be statistically significant. While the results of this paper relate to Ireland, the findings may be applicable to fiscal forecasting in a wider range of countries.

This paper follows recent work on empirically estimating Irish government revenue elasticities. Conroy (2020) compiled a new dataset of the impact of tax policy changes on various government revenue headings.[3] This dataset is based on budget day estimates of the impact of these policy changes. As a result, policy-adjusted revenue could be used when estimating elasticities. Factoring out tax policy changes allows empirical estimates to more accurately capture the elasticity between changes in the tax base and tax revenue.

Previous estimates of revenue elasticities did not account for the impact of policy changes. This may have biased estimates, as for some revenue headings policy changes have followed the economic cycle. Previous papers such as Acheson *et al.* (2017), Acheson *et al.* (2018) and Deli *et al.* (2017) did not adjust for tax policy changes.

[3] The full dataset is available at https://www.fiscalcouncil.ie/estimating-irelands-tax-elasticities-a-policy-adjusted-approach/

The Conroy (2020) analysis focused on three revenue headings, which are also examined in this paper, namely income tax, VAT and PRSI. As is the case in this paper, income tax referred to here is the aggregate of Pay As You Earn (PAYE) income tax, the Universal Social Charge (USC) and other income tax.

Conroy (2020) found that using policy-adjusted revenue had a significant impact on the elasticity estimated for income tax. A much larger elasticity (1.4, significantly above one) was found using policy-adjusted revenue. This contrasts with estimates using unadjusted revenue (0.8, significantly below one). Income tax policy changes were large, regular and procyclical over the period examined (Figure 3.1), hence there is a significant difference in the elasticity estimated using policy-adjusted or unadjusted data.

An elasticity of 1.4 implies that a 1 per cent increase in income would yield a 1.4 per cent increase in income tax receipts, with an unchanged tax system. This reflects the progressive nature of the Irish income tax system, as marginal tax rates exceed the average rate.

For VAT and PRSI, using policy-adjusted revenue had less of an impact on the elasticities estimated. This is to be expected, as policy changes for those two revenue headings have been much smaller and less frequent (see Figure 3.2 relative to Figure 3.1).

For VAT, both personal consumption and investment in the building and construction sector were used to represent the tax base. Both of these variables were found to be significant predictors of VAT receipts, both in the long run and the short run.

3. Data and Methodology

This section details the data and methodology used to build and assess the tax elasticities we consider.

3.1 Data

We assess three headings of government revenue in this paper: income tax, VAT and PRSI. Revenue data are obtained from the Department of Finance databank and are on an Exchequer basis. For income tax, the figures used include the universal social charge (from 2011) and the income levy (prior to 2011).[4] While henceforth the phrase "income tax" is used, one should remember that this broader definition is being referred to here.[5] Income tax, VAT and PRSI accounted for almost two thirds of central government revenue in 2019.

This paper focuses on how changing the elasticity used impacts on forecasting performance. Using historical data, we can see what forecasts would have been produced when applying different elasticities and evaluate how close they would have been to the eventual outturn.[6] Specifically, we focus on the impact of using different elasticities to forecast each of the revenue headings.

The methodology adopted in this paper to forecast government revenue employs a new dataset (Conroy, 2020) of the impact of tax policy changes. This dataset is made up of official budget day estimates drawing on information from Revenue and the Department of Finance as to what

---

[4] The health levy was abolished and merged into the Universal Social Charge in 2011. The health levy had previously not been included in the category "income tax" receipts in previous years, as it was recorded as a departmental receipt (Department of Health). In 2010, the health levy raised €2.018 billion. To account for this, we add €2.018 billion to the discretionary income tax/USC policy changes listed in the Budget documentation for 2011.

[5] The Universal Social Charge is structured somewhat differently to PAYE income tax. It applies to a wider base of income and has no associated tax credits. USC also typically provides fewer reliefs.

[6] As explained later, this is different to fitting the data, as we only use the revenue outturn for year T-1 and then use this as a base for forecast revenue for the years T to T+4.

the cost/yield of tax policy changes would likely be. These ex-ante estimates of the impact of policy changes will contain errors, as they are not adjusted ex-post. Few, if any, assessments of the actual impact of policy changes on revenues are completed ex-post. This means that there is a significant information gap in relation to the impacts of policy changes on tax revenues.

In the absence of more comprehensive assessments of individual policy changes, these ex-ante estimates provide the best route to correcting government revenue for policy changes made. Initial year, full year and one-off impacts of tax policy changes are recorded. This is done for several revenue headings, including the three examined in this paper.

Figure 3.1: Estimated Impacts of Policy Changes on Income Tax Revenues

€ billions, 1987-2019

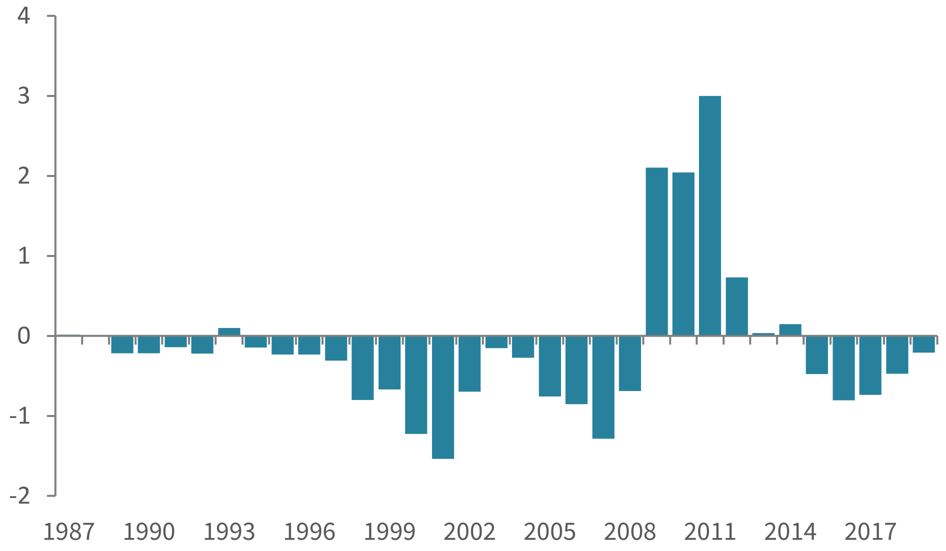

Sources: Department of Finance; and author's own calculations.
Note: Positive values indicate that policy changes were expected to raise revenue overall, or an effective tax rise. This is measured against a no-policy change baseline, which does not include indexation of tax bands or credits for inflation or wage growth. As a result, widening income tax bands and increasing tax credits in line with indexation would be recorded as a revenue-reducing measure here.

Estimates of the impact of tax policy changes have been included in budget day documentation since 1987, so that is when the analysis starts.[7,8] Figure 3.1 shows how significant policy changes have been for income tax over the past thirty years.[9]

In the period of strong economic and income growth preceding the last crisis, there were significant policy changes which reduced the amount of income tax paid. Had these policy changes not been made, revenue would have grown even more rapidly during this period of growth. From 2009 to 2012, significant income tax policy changes were made to raise additional revenue and to reduce a structural government deficit. These changes mitigated the fall in income tax collected in 2009/10 somewhat and aided the increase in receipts in 2011/12. Given how procyclical income tax policy changes were over this period, it is no surprise that estimates of the elasticity of income tax differ greatly if policy-adjusted revenue estimates are used.

The ex-ante estimates of the impact that policy changes have on revenues that we use (taken from annual budget documentation) are based on an assumed “no policy change” baseline. The no policy change baseline used by the Department assumes no automatic indexation of tax bands or credits. This means that any widening of tax bands or increase in credits would be recorded as a revenue-reducing measure. In a growing economy, keeping tax bands and credits fixed will result in more tax being paid at higher rates, resulting in higher revenue.

Figure 3.2 shows that policy changes for VAT and PRSI have been more modest by comparison.[10] As noted in Section 2, this means there is less of

---

[7] One exception is the reduced (9 per cent) rate of VAT (mainly applicable to tourism related activities) introduced midway through 2011, which was not listed in *Budget 2011* documentation. The estimate of the cost of this policy change (€120 million in 2011, €350 million in a full year) is taken from the jobs initiative documentation (Department of Finance, 2011).
[8] 2019 is the final year considered for forecasts.
[9] Section 3.2 describes how policy changes are calculated, using estimates of the initial year and full year impacts.
[10] While Figures 3.1 and 3.2 use absolute (€ billion) amounts, showing policy changes as a percentage of revenue yields a similar pattern, with income tax policy changes being much bigger than those for VAT or PRSI.

a difference between the policy-adjusted and unadjusted estimates of the respective elasticities. With this in mind, one might expect that using policy-adjusted elasticities would have less of an impact on forecasting performance for VAT and PRSI (compared to income tax).

Figure 3.2: Estimated Impacts of Policy Changes on VAT and PRSI Receipts

€ billions, 1987-2019

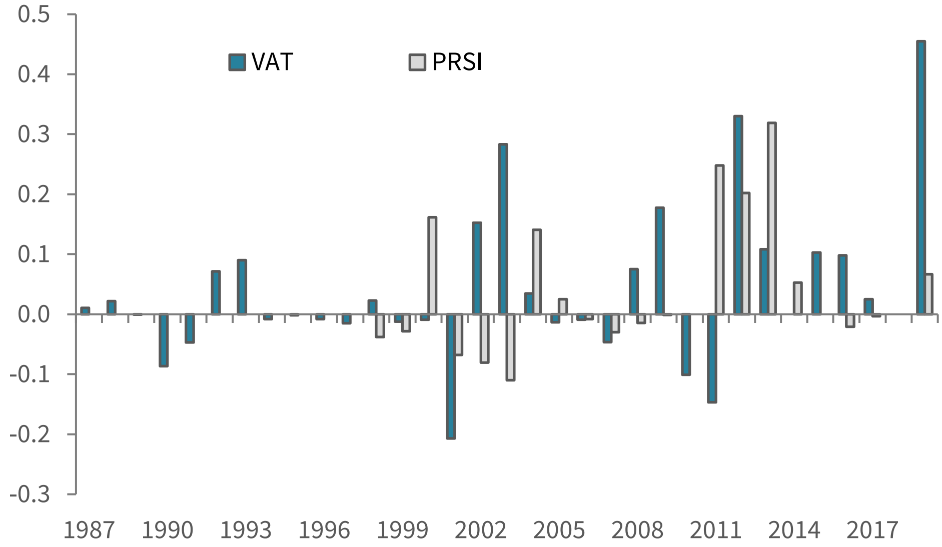

Sources: Department of Finance; and author's own calculations.
Note: Positive values indicate that policy changes are raising revenue overall, or an effective tax rise. This is being measured against a no-policy change baseline. For PRSI, this does not include indexation of bands for inflation or wage growth. This means that widening PRSI bands in line with indexation would be recorded as a revenue-reducing measure here.

Distortions caused by the activities of multinationals mean GDP and GNP are no longer reliable indicators of economic activity in Ireland. Alternative metrics which strip out the impact of foreign-owned multinational enterprises on the economy are more suitable for Ireland. With this in mind, Domestic GVA and modified GNI (GNI*) may be more suitable macroeconomic drivers of income tax and PRSI.

Domestic GVA is a measure that captures the gross value added of sectors that are not dominated by foreign-owned multinational enterprises.[11]

[11] This official measure of economic activity is produced by the Central Statistics Office. The non-domestic sector is defined as the sum of sectors where the turnover of foreign-owned multinational enterprises exceeds 85 per cent of the sector total on average.

GNI* describes Gross National Income excluding factor income of redomiciled companies, depreciation on R&D service imports and trade in intellectual property, and depreciation on aircraft leasing.

As the government revenue figures are in nominal terms, the macroeconomic drivers are also taken in nominal form. Non-agricultural income is used in the baseline income tax and PRSI forecast exercises, with Domestic GVA and GNI* (also in nominal terms) used as robustness checks (which are shown in Appendix A). The non-agricultural income variable comes from Table 1 from the annual National Income and Expenditure accounts. This combines non-agricultural wages and salaries with non-agricultural self-employed earnings.

Figure 3.3: Measures of Income and Output in Ireland

1985-2019, log levels, € millions

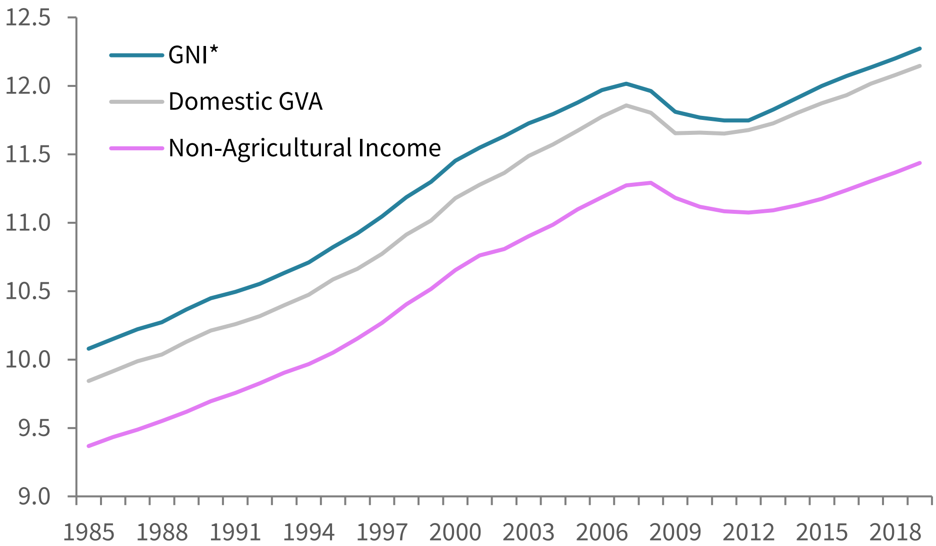

Sources: CSO; and author's own calculations.

Figure 3.3 shows the log of Domestic GVA, GNI* and non-agricultural income over the period 1985 to 2019.[12] The three metrics all show a similar profile. However, we find that using non-agricultural income

[12] The growth rate (given by the difference in logs) of the macroeconomic drivers is used to forecast government revenue, further details in Section 3.2.

yields significantly better forecasts for both income tax and PRSI, compared to using Domestic GVA or GNI* as the macroeconomic driver (see Appendix A).[13]

For VAT, nominal personal consumption and nominal investment in the building and construction (B&C) sector are used. B&C investment is included as the housing sector yields considerable VAT receipts (see Addison-Smyth and McQuinn, 2016). Each of these macroeconomic drivers are taken from the quarterly and annual National Accounts published by the Central Statistics Office.

When compiling forecasts of VAT receipts, the Department of Finance currently uses personal consumption as the sole macroeconomic driver. The 2008 Tax Forecasting Methodology Review Group (TFMRG, 2008) report had suggested using construction activity in addition to consumption. This recommendation was initially adopted but was discontinued after *Budget 2011*. The most recent TFMRG (2019) report also recommended that VAT forecasts be supplemented with a housing element. This could be an important driver of VAT receipts in the coming years, particularly if housing activity continues to grow.

3.2 Methodology

A variety of approaches can be taken to forecast the various government revenue headings. The contribution of this paper is to see how forecasting performance varies depending on the elasticity used. With this in mind, we keep the forecasting methodology constant, only varying the elasticity used. In the case of VAT, we use two different tax bases,

[13] In practical terms, the accuracy of revenue forecasts would also depend on the accuracy of forecasts of the macroeconomic driver.

given that the choice of tax base has such a large impact on forecasting performance.[14]

The methodology used here mirrors that applied by the Department of Finance when compiling revenue forecasts.[15] To forecast revenue for the current year, the outturn from the previous year (excluding any one-off factors) is used.[16] This is combined with expected growth of the macroeconomic driver this year (multiplied by the elasticity), and the assumed impact of tax policy changes. One-off factors are also incorporated into the forecasts. When the Department are compiling forecasts, judgement is often applied. In the forecasting exercise conducted in this paper there is no judgement applied, hence it is excluded from equation (1) below.[17]

$$Revenue_{t,forecast} = (Revenue_{t-1} - one\ off_{t-1}) * (1 + MDgrowth_t * elasticity) + Policy_t + one\ off_t \quad (1)$$

$$Policy_t = Policy\ initial_t + (Policy\ full\ year_{t-1} - Policy\ initial_{t-1}) \quad (2)$$

Where $MDgrowth_t$ represents growth of the macroeconomic driver in year T (given by the logged difference). $Policy\ initial_t$ describes the impact tax policy changes are expected to have in euro terms in the year they are introduced. $Policy\ full\ year_{t-1}$ describes the full year impact policy changes introduced in year T-1 are expected to have (which may differ from the impact they had in the year they were introduced).

[14] For income tax and PRSI, Appendix A shows that forecast accuracy significantly deteriorates if different tax bases are used instead of non-agricultural income.
[15] The Department of Employment Affairs and Social Protection compile the forecasts of PRSI revenue using a similar methodology.
[16] In-year forecasts made in the spring (as part of the SPU publication) follow this methodology. In-year forecasts made for the budget are typically updates of the in-year forecast made in the spring, accounting for the performance of each tax heading and any policy changes which have occurred in the intervening period. Forecasts for one (or more)-year-ahead in both budget and SPU publications follow the methodology used in this paper.
[17] When judgement is applied to official forecasts, it is typically focused on in-year and one-year-ahead forecasts.

$One\ off_t$ describes any one-off factors which impact on revenue in year T.

Tax policy changes can have an impact on revenues, both through the impact of the changes in the initial year, but also through the carryover effect from changes in the previous year (equation 2).

When assessing forecasting performance, normally one would use the vintages of data that were available at the time when forecasts were being compiled.[18] The main focus of this paper, however, is simply to assess how varying the elasticity impacts on the forecasting of government revenue.[19] With this in mind we use the current estimate of the growth of the macroeconomic driver. Similarly, we use the latest estimates of the impact of policy changes, even if these may not have been predicted years in advance. As a result, we can focus solely on the impact different elasticities have on forecasting performance.[20]

Given the extent of revisions to Irish macroeconomic data (Casey and Smyth, 2016) and the volatile nature of these data (Conroy, 2015), large forecast errors for Irish macroeconomic variables are likely. Macroeconomic forecast errors are also likely to be larger over longer forecast horizons, as errors cumulate. Given we are assuming perfect foresight of the macroeconomic drivers, the forecast errors for government revenue in this paper are lower than would be the case if real time forecasts of the macroeconomic drivers were used. This would be more pronounced over longer forecast horizons.

[18] All else being equal, one would expect that using outturns of the macroeconomic drivers as inputs should lead to more accurate forecasts. This is because forecasts of macroeconomic drivers would inevitably include errors, which would (on average) lead to larger errors in forecasting government revenue.
[19] However, it turns out that in the case of VAT the choice of tax base has a much bigger impact than varying the elasticity used.
[20] Diebold-Mariano tests are used to examine if differences in forecast accuracy are statistically significant.

The Department of Finance forecasts various Exchequer tax headings for the current year (T) and the year-ahead (T+1) at Budget time.[21] As a result, the TFMRG (2019) report was only assessing forecasts over this short forecast horizon.

However, longer-term fiscal forecasts are key to support a medium-term orientation of fiscal policy. With that in mind, this paper examines forecasting performance over the horizon of the current year (T) out to four-years ahead (T+4).

For this exercise, the outturns for the previous year (T-1) are the basis for the forecasts over all time horizons (T to T+4). For example, forecasts for the current year (T) are then used as a base for forecasting one-year ahead (T+1). The forecasts for the year T+1 can then be used as a base to forecast T+2 and so on.[22] More generally, when forecasting n years ahead (to year T+n), the forecast for the previous year is used as the base to grow from (as in, year T+n-1). A generic formula for forecasting n years ahead is given in equation (3) below (where n ranges from zero to four).

$$Revenue_{t+n,forecast} = (Revenue_{t+n-1,forecast} - one\ off_{t+n-1}) * (1 + MDgrowth_{t+n} * elasticity) + Policy_{t+n} + one\ off_{t+n} \qquad (3)$$

By using the outturn for the previous year (T+n-1) to forecast revenue for the year T+n, there is a danger of assuming that revenue level in year T+n-1 is sustainable and not being heavily impacted by temporary factors. By subtracting any one-off or temporary factors impacting year T+n-1 which are known, we are mitigating this risk.

[21] General government revenue forecasts are compiled for a longer forecast horizon (typically out to four- or five-years ahead) at budget time. Fiscal forecasts in *Budget 2021* were unusual in that they only covered 2020 and 2021.

[22] While policy changes for future years would not be known when actually making forecasts, we assume they are known in this exercise so that we are isolating the effect of changing the elasticity and/or the tax base used.

We are compiling forecasts for the years 1988 – 2019. So, there are 32 forecasts (and hence errors) for each methodology for in-year forecasts (year T). For the four-year-ahead forecasts (T+4), we have 28 such forecasts and errors.

Table 3.1: Elasticities Used to Forecast Income Tax (including USC)

Elasticity	Origin/Rationale
1.4	Long-Run elasticity estimated using policy-adjusted revenue in Conroy (2020).[23]
0.8	Long-Run elasticity estimated using unadjusted revenue in Conroy (2020).
1.89	Weighted average of the elasticities used by the Department of Finance to forecast Pay As You Earn (PAYE) income tax receipts (2.1) and USC receipts (1.2).

Sources: Various.
Note: Income tax here refers to all income tax (PAYE, USC and other).

We use a variety of elasticities to see what works best for forecasting income tax receipts (again remembering that what we are trying to forecast is total income tax receipts). Table 3.1 explains where each of the elasticities originates from. The Department of Finance forecasts PAYE income tax and USC receipts separately.[24] As a result, the Department of Finance use separate elasticities for PAYE income tax (2.1) and USC (1.2).[25] Table 3.1 shows the elasticity (1.89) that results from taking a weighted average of these two elasticities.[26]

[23] Both long-run and short-run elasticities were estimated in Conroy (2020). As the focus of this paper is on the longer forecast horizon, only the long-run elasticities are used here.
[24] In both cases non-agricultural income is used as the macroeconomic driver.
[25] The difference in elasticities found for USC and PAYE Income Tax using an analytical approach is mainly driven by tax credits, which occur in income tax but not the USC. Tax credits result in less tax being paid by those at the bottom of the income distribution, and also result in a higher marginal tax rate at the income level at which they are exhausted.
[26] Weights are determined by the relative sizes of PAYE income tax receipts and USC receipts over the period 2011-2019.

Historical data of income tax and USC revenue separately (on an Exchequer basis) were not available to the author; hence the aggregate of the two is what is forecast and compared to outturns here. There is data from the office of the Revenue Commissioners available on net receipts of USC.[27] However, this is not on an Exchequer basis and hence is not directly comparable to aggregate income tax and USC used here.[28]

As a robustness check, we forecast PAYE and USC separately using the elasticities commonly used by the Department of Finance. Due to the short period the USC has been in place, we have a limited number of observations for this exercise (Appendix B).

Table 3.2: Elasticities Used for Forecasting VAT

Elasticity	Origin/Rationale
0.80 Consumption, 0.20 Building & Construction	Elasticities estimated using policy-adjusted revenue in Conroy (2020), using consumption and B&C as the tax base.
0.88 Consumption, 0.18 Building & Construction	Elasticities estimated using unadjusted revenue in Conroy (2020), using consumption and B&C as the tax base.
1.1 Consumption	Elasticity estimated in Conroy (2020), using only consumption as the tax base.[29]
1.0 Consumption	Elasticity currently used by the Department of Finance to forecast VAT receipts.

Sources: Various.

[27] See https://www.revenue.ie/en/corporate/information-about-revenue/statistics/receipts/receipts-taxhead.aspx
[28] However, we do use this data to arrive at a weighted average of the PAYE (2.1) and USC (1.2) elasticities used by the Department of Finance.
[29] Using policy-adjusted revenue does not make a significant difference in this instance. Both policy-adjusted revenue and unadjusted revenue result in an elasticity of 1.1.

For VAT, there is more variety in the approaches that can be taken. Two different tax bases are applied to forecast VAT in this paper. First, personal consumption is used as the tax base for VAT receipts. Second, personal consumption and investment in the building and construction sector are used as the tax base. Conroy (2020) empirically estimates elasticities for both choices of tax base. Table 3.2 lists the four different combinations of elasticities used to forecast VAT in this paper, and the rationale underlying them.

Table 3.3 shows the three different elasticities used to forecast PRSI in this paper. Conroy (2020) found that the empirically estimated elasticity of PRSI is not sensitive to whether policy-adjusted revenue or unadjusted revenue is used. In each case, estimates close to one were obtained. The Department of Employment Affairs and Social Protection use an elasticity of one to forecast PRSI receipts.

Table 3.3: Elasticities Used for Forecasting PRSI

Elasticity	Origin/Rationale
1.0	Elasticity estimated using policy-adjusted revenue in Conroy (2020). The Department of Employment Affairs and Social Protection also use an elasticity of one to forecast PRSI.
1.02	Elasticity estimated using unadjusted revenue in Conroy (2020).
1.5	Price *et al.* (2014) estimate for Ireland, using an analytical approach.[30]

Sources: Various.
Note: Policy changes have been relatively limited for PRSI over this period, hence the policy-adjusted elasticity is very similar to the unadjusted elasticity.

It is worth noting that in all cases, attempts to forecast revenue in this paper adjust for the impact of policy changes on revenue. The estimates of the impact of policy changes used in this paper are the same as those

[30] Wages and salaries are used as the tax base in Price *et al.* (2014).

which were used in Conroy (2020) to estimate policy-adjusted elasticities. As a result, one might expect those elasticities to perform well in this forecasting exercise.

4. Results

In this section, we present the results for each of the three revenue headings examined.

4.1 Income tax

Three different elasticities are tested for their suitability in forecasting income tax receipts. In each case we are using the outturn of the macroeconomic driver (non-agricultural income) and budget day estimates of the impact of income tax policy changes. One might expect larger forecast errors if we were using forecasts of the macroeconomic drivers, rather than outturns. This exercise is seeking to find the most appropriate elasticity to apply. To do this, the only thing which varies in each case is the elasticity used. As outlined in Section 3, these forecasts take account of tax policy changes.

Multiple macroeconomic drivers were experimented with, particularly new measures of activity such as Domestic GVA and modified GNI (GNI*). Table A.2 (in Appendix A) shows that forecasts using non-agricultural income are significantly more accurate than those using Domestic GVA or GNI*. As a result, non-agricultural income is used throughout this section.

To test if the errors from each of these elasticities are centred on zero, we examine the time series of the errors produced by each elasticity at each time horizon. So for each elasticity, there are five time series to be examined (one for each year of the forecast horizon).

If the error series is normally distributed, then simple T-tests can be used to determine if forecasts are centred on outturns. If we reject the null hypothesis (that errors are zero on average), then there is evidence that forecasts using that elasticity at that forecast horizon are biased, and not centred on the outturns.

If the forecast error series are not normally distributed, then a nonparametric test (the Wilcoxon signed-rank test) of whether the median is zero is used (detailed result are shown in Appendix C).[31]

We use Diebold-Mariano (2002) tests to see if the differences in forecasting performance of the different elasticities are statistically significant. Diebold-Mariano tests examine whether two competing forecasts have equal predictive accuracy.[32] These tests are performed at each forecast horizon, and also for the pooled sample of forecasts at all horizons (which provides a larger sample size).

The elasticity of 1.4 comes from empirical work by Conroy (2020), estimated using policy-adjusted revenue. When this elasticity is used, both T-tests and Wilcoxon signed-rank tests suggest that there is no significant evidence of bias in these forecasts at any point in the forecast horizon (Appendix C and Figure 4.1).

By contrast, using an elasticity estimated with unadjusted income tax data (0.8) produces forecasts which are biased at all forecast horizons. T-tests and Wilcoxon signed-rank tests both show that forecasts are not centred on outturns (see Table C.1 in Appendix C for full results).[33] A higher elasticity would result in forecasts closer to the outturns on average.

[31] See Wilcoxon (1945) for details.
[32] The test can be summarised as an asymptotic z-test of the hypothesis that the mean of the forecasting loss differential is zero
[33] Forecasts underestimate outturns on average.

Figure 4.1: Average Income Tax Forecast Errors

Average forecast error (per cent of income tax revenue)

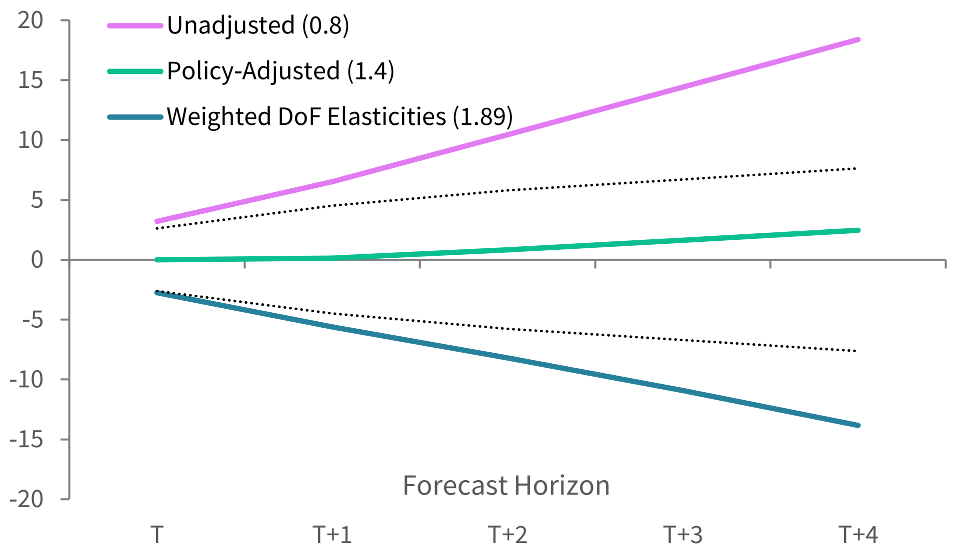

Sources: Author's own calculations.
Note: Values below zero indicate forecasts exceed the outturns on average and vice versa. Errors are calculated using the full sample of forecasts (1988 – 2019). For illustrative purposes, the dashed lines show plus/minus two standard errors, centred on zero. To construct these bands, the largest standard errors are used here (unadjusted elasticity of 0.8). This is purely for illustrative purposes. When formally assessing if average errors are significantly different from zero, T-tests or Wilcoxon signed-rank tests are used.

Next, we examine the elasticities which are used by the Department of Finance. An elasticity of 1.89 is a weighted average of the elasticities used by the Department to forecast PAYE income tax and USC receipts.

Both T-tests and Wilcoxon signed-rank tests indicate that forecasts using an elasticity of 1.89 are not centred on outturns. This suggests that a lower elasticity would result in forecasts closer to the outturns on average.

Looking at the Diebold-Mariano tests, we can see that the policy-adjusted elasticity (1.4) produces significantly better forecasts than the unadjusted elasticity (0.8) at all forecast horizons (Table 4.1). Similarly, the weighted DoF elasticity (1.89) significantly outperform the unadjusted elasticity (0.8) at all forecast horizons.

Table 4.1: Diebold-Mariano Tests, Income Tax

Relative RMSE

Elasticity	T	T+1	T+2	T+3	T+4	Pooled
Policy-Adjusted (1.4) vs Unadjusted (0.8)	0.65**	0.58**	0.53**	0.48**	0.43**	0.49**
Policy-Adjusted (1.4) vs Weighted DoF elasticities (1.89)	1.03	1.01	0.92	0.82	0.74*	0.86*
Weighted DoF elasticities (1.89) vs Unadjusted (0.8)	0.62*	0.58**	0.58**	0.58**	0.58**	0.57**

Sources: Author's own calculations.
Note: Values below one suggest that the forecast named first is superior. Values with ** or * indicates forecasts are significantly different at a 1% or 10% significance level. Pooled results use forecasts over all horizons (N=150).

Comparing the policy-adjusted elasticity to the weighted DoF elasticity, we can see some evidence for the policy-adjusted elasticity performing better. These improvements are statistically significant when pooling forecast errors over all horizons or when looking at four-year-ahead forecasts.

In Appendix B, we use data from the revenue commissioners to examine if forecasts can be improved by forecasting USC and other income tax receipts separately and then using the sum of these two forecasts. The data is not directly comparable to that used here as it is not on an exchequer basis. We find that forecasting USC and other income tax receipts separately (using elasticities of 1.2 and 2.1) and then summing the two yields inferior forecasts to using a weighted average (1.89) of these elasticities.[34] However, as the USC has only been recently

[34] Elasticities of 1.2 and 2.1 are the elasticities used by the Department of Finance to forecast USC and PAYE receipts respectively.

introduced, these results are based on very small samples and hence are not very powerful.

Figure 4.2 shows the forecasting performance of the three elasticities, using the average absolute percentage error to gauge the size of typical forecast errors. As one would expect, the average size of errors increases as the forecast horizon expands. For each of the three elasticities used, the average absolute errors for four-year-ahead (T+4) forecasts are more than double those for in-year forecasts (T). This is an interesting and novel finding, as most of the previous Irish literature on tax forecasting has focused on shorter forecast horizons (up to one-year-ahead typically).

Figure 4.2: Size of Income Tax Forecast Errors

Average absolute forecast error (per cent of income tax revenue)

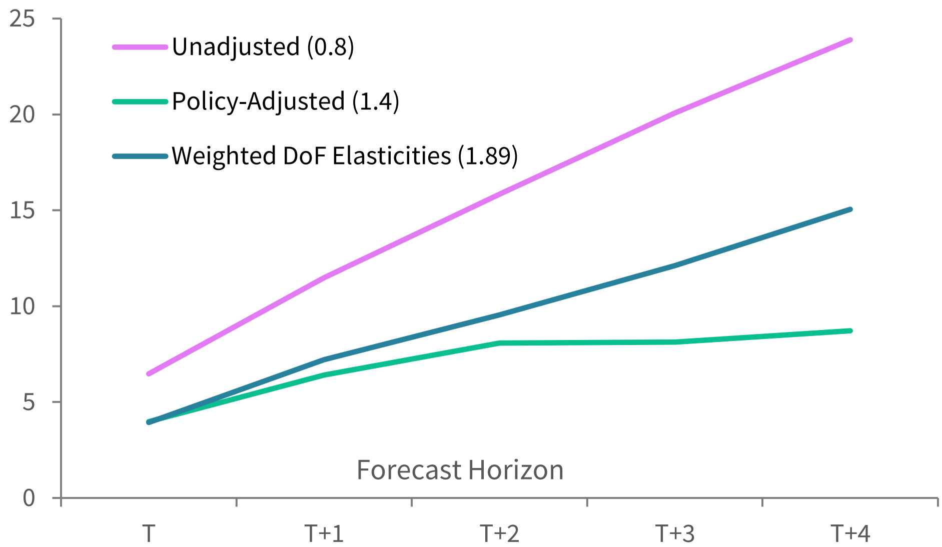

Sources: Author's own calculations. Errors are calculated using the full sample of forecasts (1988 – 2019).

The elasticity estimated using policy-adjusted data (1.4) performs best here, yielding the smallest average absolute errors.[35] This is most noticeable over longer forecast horizons, as errors cumulate. However, it should be kept in mind that these average absolute forecast errors are

[35] The weighted average of elasticities used by the Department of Finance (1.89) does produce slightly lower average absolute errors for in-year forecasts.

still quite large, ranging from 4 per cent (in-year forecasts) to 9 per cent (four-year-ahead forecasts).

The weighted average of elasticities used by the Department of Finance (1.89) works well for in-year forecasts. Over longer time horizons, the size of these forecast errors grows more rapidly than for the policy-adjusted elasticity (Figure 4.2). This is consistent with Table 4.1, which suggested there was little difference in forecasting performance over shorter forecast horizons.

Forecasting with the elasticity estimated using unadjusted data (0.8) gives the largest average absolute errors at all forecast horizons. This again highlights the importance of forecasting using elasticities which were estimated using policy-adjusted data.

To assess the uncertainty around projections from these models, Table 4.2 shows the average absolute errors for each of these different elasticities over different forecast horizons in cash terms (scaled by 2019 receipts). As nominal income tax revenue has been trending upwards over time, if one just took the average absolute error in millions of euros, this would give greater weight to more recent observations. To mitigate this, we take the average absolute error in percentage terms for the whole sample, and then multiply this by the 2019 outturn for income tax. This gives a sense of the typical size of absolute errors in 2019 cash terms. For example, using an elasticity of 1.4 results in an average absolute error of 4.0 per cent when forecasting the current year (T). Multiplying a 4.0 per cent error by the 2019 outturn (€22.9 billion) gives an average absolute error of €913 million.

Table 4.2: Average Absolute Forecast Errors, Income Tax

€ million, scaled by 2019 receipts

Elasticity	T	T+1	T+2	T+3	T+4
Policy-Adjusted (1.4)	913	1,471	1,854	1,864	1,999
Weighted DoF elasticities (1.89)	903	1,655	2,189	2,780	3,452
Unadjusted (0.8)	1,485	2,635	3,633	4,605	5,479

Sources: Author's own calculations.
Note: Values correspond to the average absolute percentage error for that elasticity at that forecast horizon multiplied by the 2019 outturn of income tax receipts (€22.9 billion). Errors are calculated using the full sample of forecasts (1988 – 2019).

We can see that the size of errors varies significantly across the elasticities (Table 4.2). These differences are quite substantial in cash terms and could have significant implications for budgetary planning. For each of the elasticities, the average error size increases as the forecast horizon extends. In line with this, the difference between the performance of the various elasticities also increases as the forecast horizon expands. This illustrates how the choice of elasticity is more important the longer the forecast horizon is, as errors cumulate.

The difference in average error size using the policy-adjusted elasticity (1.4) compared to unadjusted elasticity (0.8) is quite stark here. For a four-year-ahead forecast, the difference in average error size is more than €3 billion. This would be a substantial amount of income tax revenue and would have significant implications for fiscal planning. This shows how these improvements in forecasting are economically significant as well as statistically significant.

A final aspect of forecasting performance considered is the maximum absolute error (as a percentage of income tax) recorded for any year in the period considered (1987 - 2019). The cost of forecast errors may be highly non-linear to the forecaster, with large errors possibly being much more costly from a planning/budgeting point of view. As a result, one may prefer forecasting methods that avoid very large errors. Figure 4.3 shows the largest absolute percentage forecast error made using the various elasticities at each point of the forecast horizon.[36] On this metric, we can see that the policy-adjusted elasticity (1.4) performs best at three of the five forecast horizons examined. For one- and two-year-ahead forecasts, the weighted average elasticity (1.89) performs best. At each forecast horizon, the elasticity estimated using unadjusted data (0.8) performs the worst.

Figure 4.3: Maximum Absolute Error for Income Tax

Largest absolute error over the entire sample period (per cent of income tax revenue)

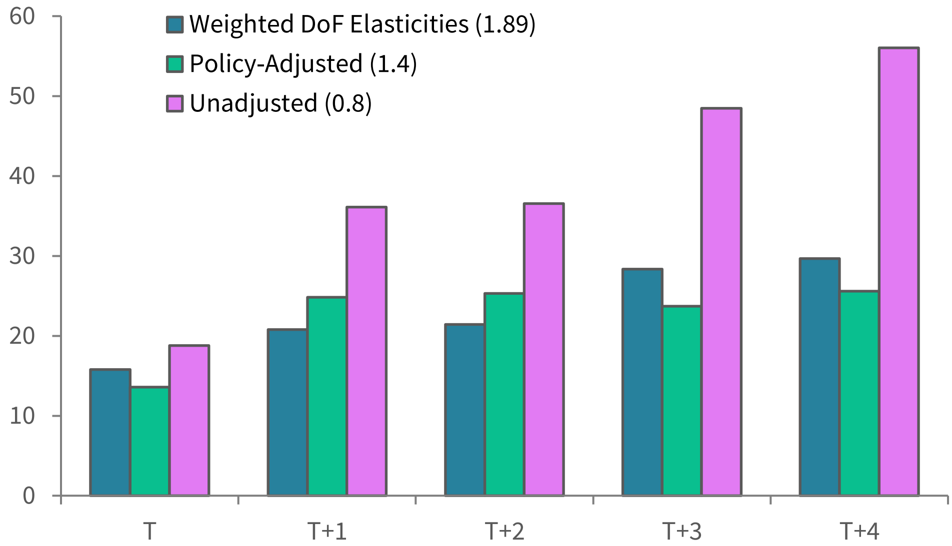

Sources: Author's own calculations.
Note: The sample period considered here is 1988-2019.

In summary, the choice of elasticity used has a significant impact on the income tax forecasts produced. Overall, the policy-adjusted elasticity performs best in forecasting income tax receipts. It produces forecasts which are unbiased, with the smallest errors on average. This is a substantial improvement on using the unadjusted elasticity, highlighting

[36] Examining these instances in detail, a variety of different years give the largest errors for the different elasticities used. In addition, the largest errors come from both underestimation and overestimation.

the importance of adjusting for policy changes. These improvements are largest when forecasting several years ahead, as errors cumulate. They are also found to be statistically significant.

This paper shows that the use of policy-adjusted elasticities could improve income tax forecasts in Ireland. However, the implications could be wider still, as policy-adjusted elasticities could be estimated and then used for fiscal forecasting in many other countries. Policy-adjusted elasticities have been estimated in other countries. However, to the authors knowledge, there has been no previous study of the improvements in forecast accuracy from using policy-adjusted elasticities.

4.2 VAT

Two different approaches to forecasting VAT are examined here. The first approach uses only personal consumption as the tax base. The second approach uses both personal consumption and investment in the Building and Construction (B&C) sector as the tax base for VAT. The rationale for including the B&C sector is that Addison-Smyth and McQuinn (2016) found that the housing sector can yield significant VAT receipts. VAT is charged on new housing purchases (at a rate of 13.5 per cent). In addition, building materials and other inputs are also subject to VAT. Conroy (2020) found that investment in the B&C sector has been a significant predictor of VAT receipts.

For each of these two approaches we use two different sets of empirical estimates of elasticities. The rationale for each of these elasticities is given in Table 3.2. Budget day estimates of the impact of VAT policy changes are incorporated into these VAT forecasts (Section 3.2).

Figure 4.4: VAT Forecasts using only Consumption are Centred on Outturns

Average forecast error (per cent of VAT revenue)

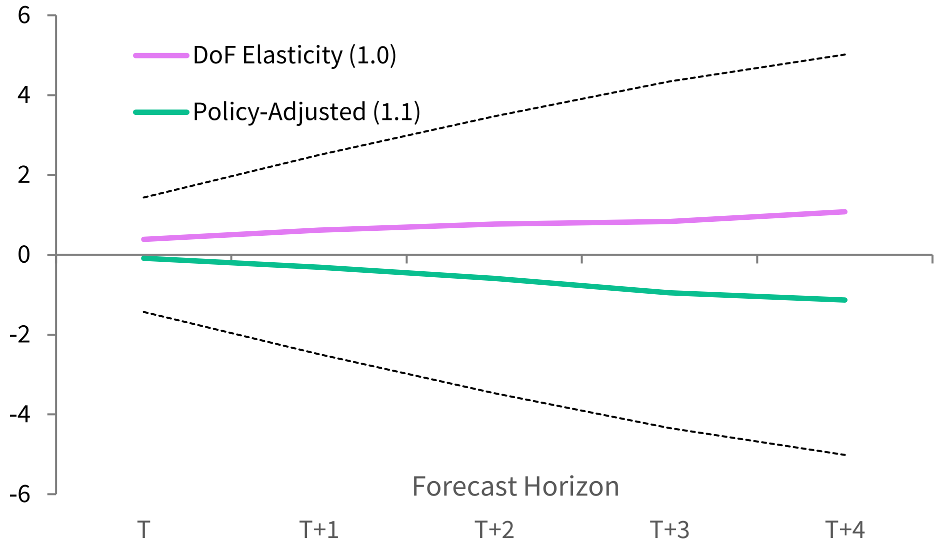

Sources: Author’s own calculations.
Note: Values below zero indicate forecasts exceed the outturns on average and vice versa. Average errors are calculated using the full sample of forecasts (1988 – 2019). For illustrative purposes, dashed lines show plus/minus two standard errors, centred on zero. These are based on errors using the policy-adjusted elasticity (1.1) but are very similar to those using an elasticity of 1.0. When formally assessing if average errors are significantly different from zero, T-tests or Wilcoxon signed-rank tests are used.

Figure 4.4 shows the average errors from forecasts using only personal consumption as the tax base. Both elasticities yield forecasts with errors close to zero on average at all forecast horizons.[37] Table C.2 in Appendix C shows the full results using T-tests and Wilcoxon signed-rank tests.

Forecasts using B&C investment as well as consumption for the tax base (Figure 4.5) are also centred on outturns. This appears to be the case for the full forecast horizon and does not depend on whether or not policy-adjusted or unadjusted elasticities are used.[38]

[37] While it is noticeable that changing from an elasticity of 1.0 to 1.1 (when using consumption alone) results in the sign of the average error changing, this merely reflects that in both cases the average error is very close to zero at all forecast horizons.
[38] Again, Table C.2 shows results of T-tests and Wilcoxon signed-rank tests.

Figure 4.5: VAT Forecasts using Consumption and B&C Investment are Centred on Outturns

Average forecast error (per cent of VAT revenue)

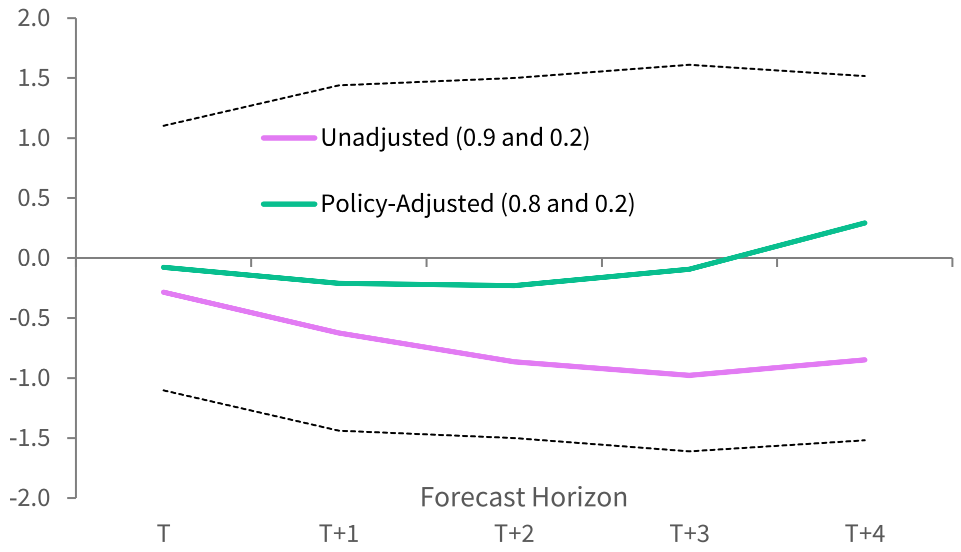

Sources: Author's own calculations.

Note: Values below zero indicate forecasts exceed the outturns on average and vice versa. Errors are calculated using the full sample of forecasts (1988 – 2019). Dashed lines show plus/minus two standard errors, centred on zero. These are based on errors using the policy-adjusted elasticities (0.8 and 0.2) but are very similar to those using the unadjusted elasticities. When formally assessing if average errors are significantly different from zero, T-tests and Wilcoxon signed-rank tests are used.

Looking at the average absolute percentage errors (Figure 4.6), we can see two different trends emerge. The estimates which use B&C investment along with personal consumption as the tax base show much smaller errors, particularly over longer forecast horizons. It is somewhat surprising that the average size of these errors does not significantly increase as the forecast horizon expands. The policy-adjusted elasticities (0.8 and 0.2) and the unadjusted elasticities (0.9 and 0.2) give very similar forecasts and hence very similar forecast errors. This is because the elasticities themselves are quite similar, due to VAT policy changes having been relatively modest over the sample period.

By contrast, when only personal consumption expenditure (PCE) is used, the size of forecast errors increase significantly as the forecast horizon expands. Using the policy-adjusted elasticity (1.1) yields smaller average errors compared to the elasticity currently used by the Department of Finance for forecasting (1.0).

Figure 4.6: Size of VAT Forecast Errors

Average absolute forecast error (per cent of VAT revenue)

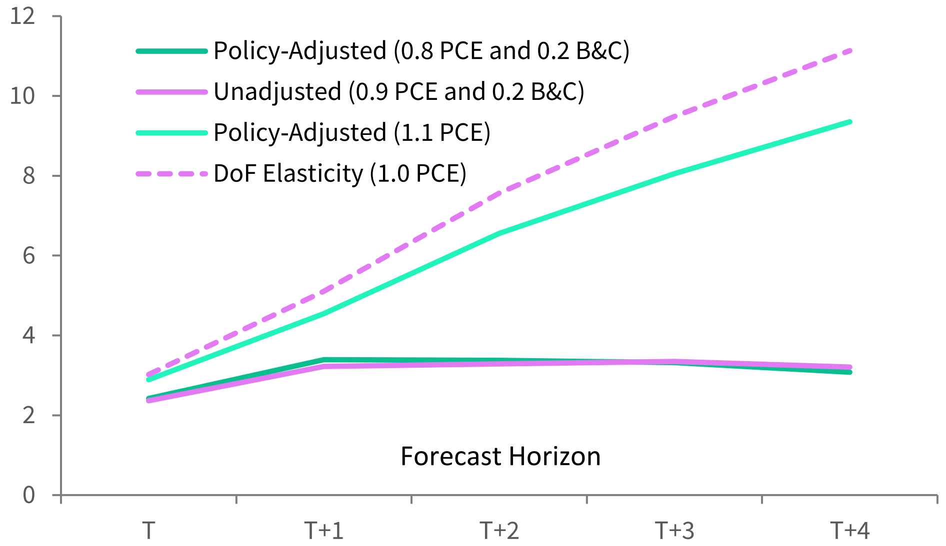

Sources: Author's own calculations.
Note: PCE stands for Personal Consumption Expenditure. B&C stands for investment in the building and construction sector. Errors are calculated using the full sample of forecasts (1988 – 2019).

While both broad approaches (just using consumption or using consumption along with B&C investment) yield forecasts centred on outturns, there is a large disparity in the size of average absolute errors. Figure 4.6 shows that the combined B&C and consumption approach results in smaller forecast errors on average. This is particularly acute when forecasting further ahead. If one uses this approach, the average absolute error (in percentage terms) is smaller than for Income tax or PRSI at all points in the forecast horizon.

Diebold-Mariano tests can show if differences in forecasting performance are statistically significant or not. When using PCE and B&C as macroeconomic drivers, whether or not policy-adjusted elasticities are used does not have a statistically significant impact on forecasting performance (Table 4.3).

When consumption is the only macro driver, using a policy-adjusted elasticity does significantly improve forecasting performance. This is the case at all forecast horizons apart from in-year forecasts.

Table 4.3: Diebold-Mariano Tests, VAT

Relative RMSE

Elasticity	T	T+1	T+2	T+3	T+4	Pooled
Policy-Adjusted (0.80 PCE, 0.20 B&C) vs Unadjusted (0.88 PCE, 0.18 B&C)	1.00	1.00	1.01	1.02	1.05	1.02
Policy-Adjusted (1.1 PCE) vs DoF Elasticity (1.0 PCE)	0.96	0.94**	0.93**	0.93**	0.93**	0.93**
Policy-Adjusted (0.80 PCE, 0.20 B&C) vs Policy-Adjusted (1.1 PCE)	0.67	0.52	0.42*	0.40*	0.33**	0.41**
Unadjusted (0.88 PCE, 0.18 B&C) vs DoF Elasticity (1.0 PCE)	0.64*	0.49*	0.39**	0.36**	0.29**	0.38**

Sources: Author's own calculations.
Note: Values below one suggests that the forecast named first is superior. Values with ** or * indicates forecasts are significantly different at a 1% or 10% significance level. Pooled results use forecasts over all horizons (N=150).

Whether applying policy-adjusted elasticities or not, using PCE and B&C yields significantly better forecasting performance than using PCE alone as the tax base. This finding appears to be strongest at longer forecast horizons. This is consistent with the average absolute forecast errors shown in Figure 4.6.

Interestingly, the last two reports of the Tax Forecasting Methodology Review Group (TFMRG, 2008 and 2019) both suggested that using personal consumption to forecast VAT could be complemented by also using expected trends in the housing market.

Table 4.4 shows the average absolute errors for each of these approaches in cash terms (scaled by 2019 receipts). As we did in Table 4.2, we take the average absolute error in percentage terms for that methodology at that forecast horizon, and then multiply this by the 2019 outturn. For example, when using only consumption and an elasticity of 1.1, the average absolute percentage error when forecasting four-years-ahead (T+4) is 9.4 per cent. Multiplying 9.4 per cent by the 2019 outturn (€15.1 billion) gives a value of €1,414 million.

Table 4.4: Average Absolute Forecast Errors, VAT

€ million, scaled by 2019 receipts

Elasticity and Macro Driver	T	T+1	T+2	T+3	T+4
Policy-Adjusted (0.80 PCE, 0.20 B&C)	366	513	511	502	466
Unadjusted (0.88 PCE, 0.18 B&C)	357	487	497	506	485
Policy-Adjusted (1.1 PCE)	437	688	991	1,218	1,414
DoF Elasticity (1.0 PCE)	457	772	1,144	1,434	1,683

Sources: Author's own calculations.
Note: Values correspond to the average absolute percentage error at that forecast horizon (for that elasticity) multiplied by the 2019 outturn of VAT receipts (€15.1 billion). Errors are calculated using the full sample of forecasts (1988 – 2019).

From Table 4.4 we can see that the average size of errors can vary substantially due to the approach taken. Forecasts using both consumption and B&C activity as the tax base have much smaller average errors than those using just personal consumption. Table 4.4 shows little difference in the forecasting performance due to using the policy-adjusted or unadjusted elasticities when consumption and B&C activity represent the tax base. This is consistent with the earlier Diebold-Mariano tests.

In a couple of cases, the average absolute errors for three-years-ahead (T+3) and four-years-ahead (T+4) forecasts are smaller than those for one-year-ahead (T+1) and two-years-ahead (T+2). While this is quite surprising (and in contrast to the other revenue headings), the magnitude of this difference is small.

Looking at the forecasts using only consumption as the tax base, there are modest gains from using the policy-adjusted elasticity. However, the gains from using a policy-adjusted elasticity here are smaller than the gains from choosing using the correct tax base. Again, this echoes the formal results given by the Diebold-Mariano tests.

The difference in the maximum absolute forecast errors from the two approaches is also quite stark (Figure 4.7). The forecasts using both personal consumption and B&C as the tax base yield much smaller maximum errors at all forecast horizons. This is also contributing to the lower averages seen in Figure 4.6 and Table 4.4. Looking over the various tests used here, it appears that an approach that uses B&C investment as well as personal consumption as the tax base yields superior forecasting performance. The differences in forecasting performance are most apparent over longer forecast horizons.

Figure 4.7: Maximum Absolute VAT Errors

Largest absolute error over the entire sample period (per cent of VAT revenue)

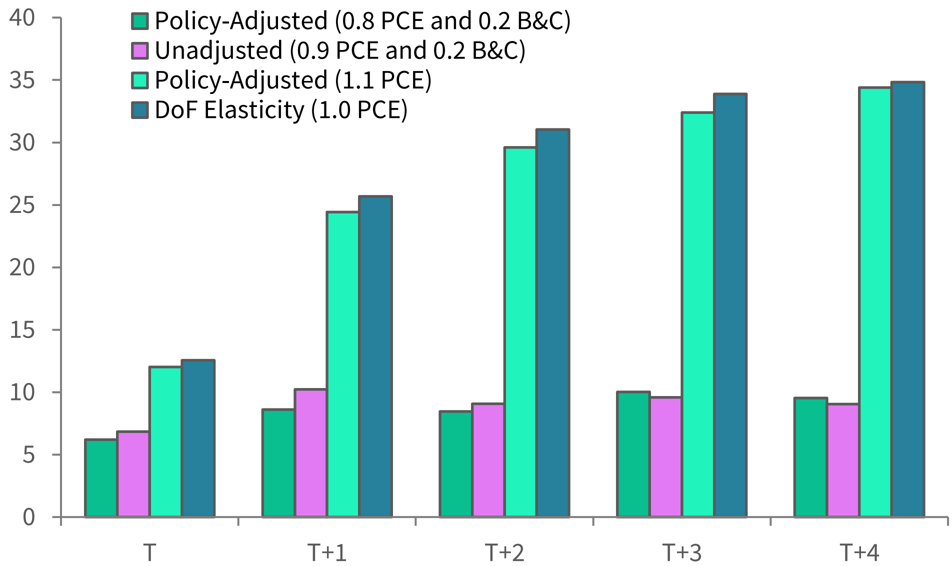

Sources: Author's own calculations.
Note: The sample period considered here is 1987-2019.

Given the stark differences in forecasting performance between the two approaches, it is worth examining what may be driving these differences. Figure 4.8 shows the absolute percentage errors (for four-year-ahead forecasts) in each year using the two approaches. From Figure 4.8, the disparity in forecasting performance is most evident in the period 2008-2012. VAT receipts fell by almost €3 billion (20%) in 2009. Much of this was due to the sudden reduction in construction activity. When using personal consumption and B&C to forecast VAT, this fall in receipts is predicted reasonably accurately (using the outturn data for consumption and B&C). By contrast, when only personal consumption is used to represent the tax base, this fall in VAT is vastly underestimated.[39]

Figure 4.8: VAT Absolute Errors by Year

Absolute errors for the four-year-ahead forecasts (T+4, per cent of VAT revenue)

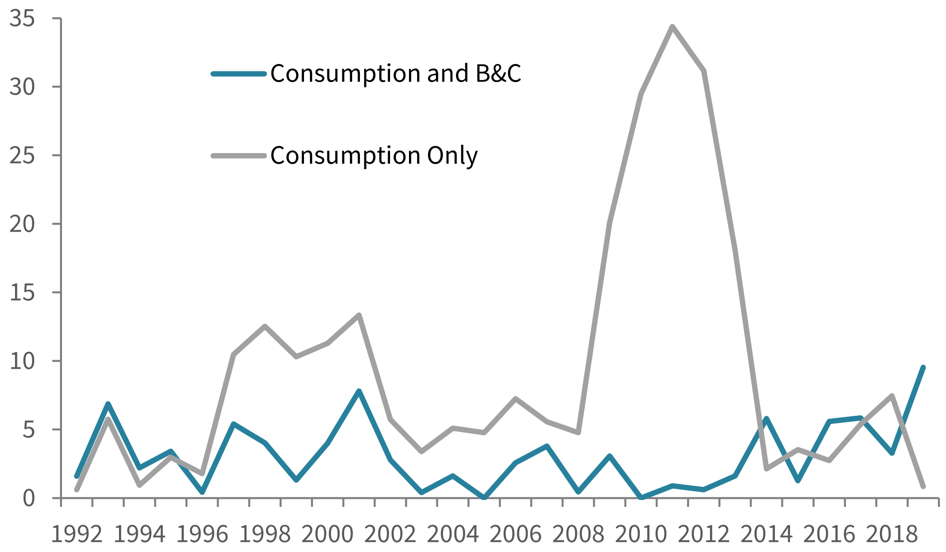

Sources: Author's own calculations.
Note: For simplicity just one elasticity is shown for each method here. In both cases, the policy-adjusted elasticity is used (1.1 for consumption only, 0.8 and 0.2 for the consumption and B&C approach).

[39] This is also reflected in the maximum errors shown in Figure 4.7. When only consumption is used as a macroeconomic driver, the largest errors from that approach occur in the years 2008-11. In each case VAT receipts are forecast to be much higher than the eventual outturn. When consumption and B&C activity are used, the maximum errors are smaller and occur in different parts of the sample, with some in the late 80s/early 90s.

To forecast VAT accurately, it appears the choice of tax base is more important than the elasticity used. Forecast accuracy is improved by using policy-adjusted elasticities, but these improvements are minor relative to using an appropriate tax base. This reflects that VAT policy changes have been relatively minor over the sample period considered.

4.3 PRSI

Three different elasticities for PRSI are tested for their forecasting performance. In Conroy (2020), elasticities estimated using policy-adjusted and unadjusted data were almost identical (close to one). Price *et al.* (2014) estimated an elasticity of 1.5, using an analytical approach.

For all three cases we are using the outturn of the macroeconomic driver (non-agricultural income) and budget day estimates of the impact of PRSI policy changes. The only thing which varies in each case is the elasticity applied.

Figure 4.9: Average PRSI Forecast Errors

Average forecast error (per cent of PRSI revenue)

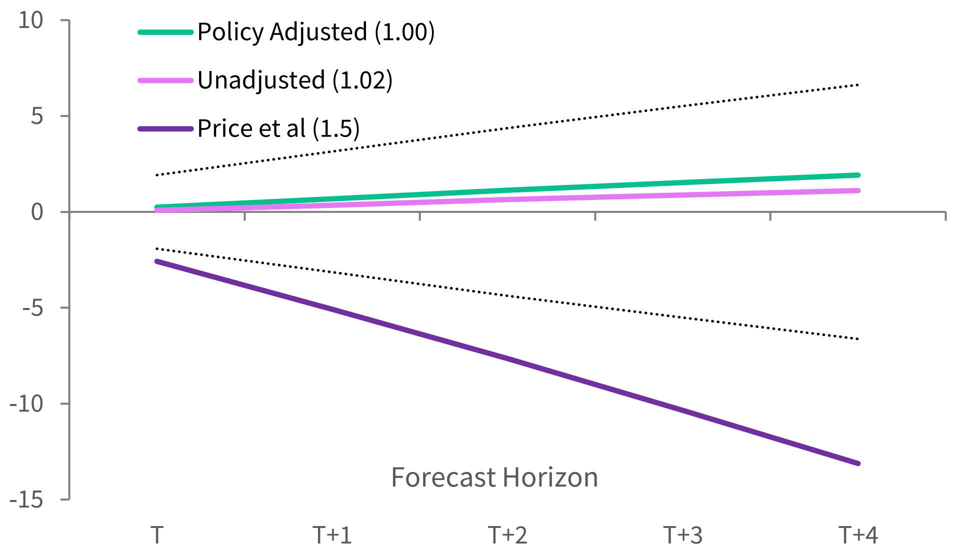

Sources: Author's own calculations.
Note: Values below zero indicate forecasts exceed the outturns on average and vice versa. Errors are calculated using the full sample of forecasts (1988 – 2019). For illustrative purposes, dashed lines show plus/minus two standard errors, centred on zero. These are based on an elasticity of 1.5, which gives the widest error bands of the three elasticities examined. When formally assessing if average errors are significantly different from zero, T-tests and Wilcoxon signed-rank tests are used.

Both the policy-adjusted and unadjusted elasticities estimated in Conroy (2020) produce unbiased forecasts at all forecast horizons. Their average forecast errors are not significantly different from zero. For the larger elasticity from Price *et al.* (2014), average forecast errors are significantly different to zero at all forecast horizons. These results are found using both T-tests and Wilcoxon signed-rank tests (Table C.3 in Appendix C). As shown in Figure 4.9, there is a tendency to overestimate PRSI receipts when using an elasticity of 1.5, particularly for the later forecast years.[40]

Table 4.5: Diebold-Mariano Tests, PRSI

Relative RMSE

Elasticity	T	T+1	T+2	T+3	T+4	Pooled
Policy adjusted (1.00) vs unadjusted (1.02)	1.00	1.00	0.99	1.00	1.02	1.00
Policy adjusted (1.00) vs Price et al (1.5)	0.78*	0.65*	0.55*	0.50**	0.49**	0.54**
Unadjusted (1.00) vs Price et al (1.5)	0.78*	0.65*	0.56*	0.50**	0.48**	0.54**

Sources: Author's own calculations.
Note: Values below one suggest that the first forecast is superior. Values with ** or * indicates forecasts are significantly different at a 1% or 10% significance level. Pooled results use forecasts over all horizons (N=150).

Diebold -Mariano tests show that there is no significant difference in forecasting performance between the policy-adjusted and unadjusted elasticity (Table 4.5). There are significant improvements in performance from using the policy-adjusted or unadjusted elasticity rather than the larger elasticity (1.5). This is the case across the forecast horizon.

From Figure 4.10, we can see that the size of average absolute errors increases as the forecast horizon lengthens for each of the elasticities. The policy-adjusted and unadjusted elasticities produce almost identical forecasts (and hence forecast errors). At each point in the forecast horizon, these errors are smaller on average than those from an elasticity of 1.5.

[40] Price et al (2014) estimates elasticities of social security contributions (and a number of other revenue headings) with respect to income for 28 EU countries. An estimate of 1.5 is found for social security contributions in Ireland.

Figure 4.10: Size of PRSI Forecast Errors

Average absolute forecast error (per cent of PRSI revenue)

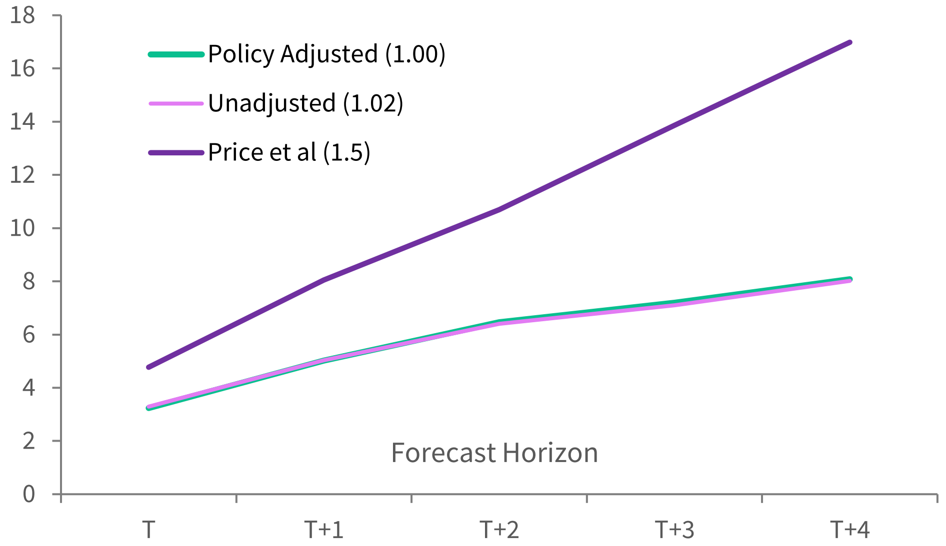

Sources: Author's own calculations.
Note: Errors are calculated using the full sample of forecasts (1988 – 2019).

Looking at the maximum absolute errors, a similar picture emerges, with the policy-adjusted and unadjusted elasticities performing similarly (Figure 4.11). Beyond in-year forecasts, both of these elasticities outperform an elasticity of 1.5.

Figure 4.11: PRSI Maximum Absolute Errors

Largest absolute error over the entire sample period (per cent of PRSI revenue)

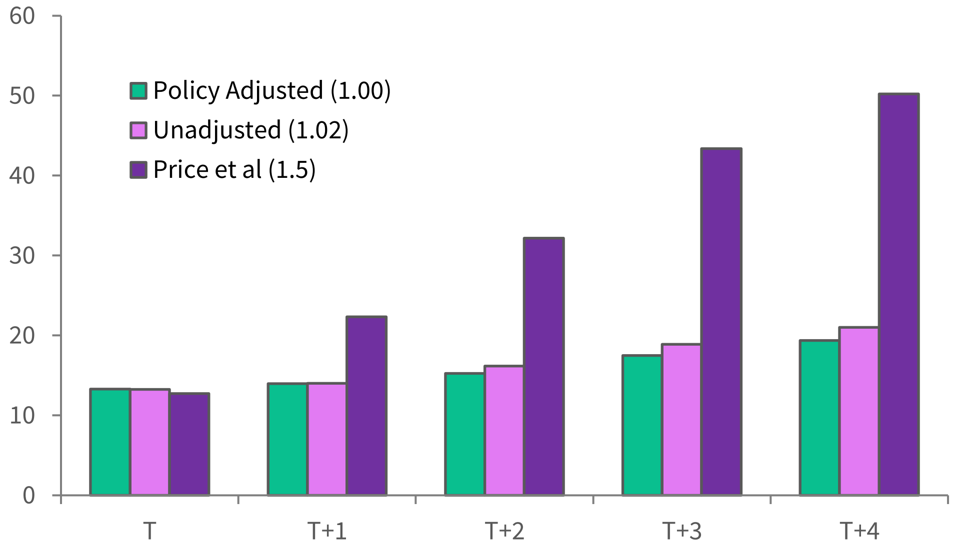

Sources: Author's own calculations.
Note: The sample period considered here is 1988-2019.

Overall, we can see that an elasticity of one (or close to one) performs well in forecasting PRSI receipts. Due to limited policy changes in recent years, the policy-adjusted elasticity is similar to the unadjusted elasticity. As a result, using an elasticity based on policy-adjusted data produces PRSI forecasts which are very similar to that using an elasticity based on unadjusted data. The Department of Employment Affairs and Social Protection uses an elasticity of one when forecasting PRSI receipts, which seems appropriate.

5. Conclusions

Forecasting government revenue is key to good budgeting and setting appropriate fiscal policy. We make several contributes to the literature on forecasting Irish income tax, VAT and PRSI. First, we assess forecasting performance over a longer horizon than has been examined previously, out to four years ahead. Secondly, we show how forecasting performance depends on the elasticity applied. Thirdly, we show that elasticities estimated using policy-adjusted data produce superior forecasts. Finally, we use formal tests to show that these improvements are statistically significant.

We use historical data to see which elasticities would have worked best in forecasting revenue. A new dataset is also used to take account of tax policy changes when making these forecasts. Using revenue outturns for the previous year (T-1), we compile forecasts for the current year (T) and out as far as four-years-ahead (T+4). Previous work on revenue forecasts in Ireland has focused on in-year (T) and one-year-ahead forecasts (T+1). So, this paper makes a key contribution in extending the forecast horizon examined.

Income tax forecasts see the biggest gains from using policy-adjusted elasticities. This reflects significant income tax policy changes which occurred over the sample period. These improvements in forecasts are most evident over longer forecast horizons where errors cumulate. We use Diebold-Mariano tests to formally test that these improvements in forecasting are statistically significant. We find that these forecasting improvements are statistically significant.

For VAT, we find that choosing the right tax base has a bigger impact on forecast accuracy than using a policy-adjusted or unadjusted elasticity. This reflects the relatively minor VAT policy changes which occurred over the sample period. Using investment in the Building and Construction

(B&C) sector in addition to personal consumption to forecast VAT yields a statistically significantly improvement in forecast performance. This improved performance is most noticeable over longer forecast horizons. This is most clearly seen when VAT receipts sharply declined in the last recession, which was partially driven by reduced construction activity.

As a result, it would seem appropriate for forecasts of VAT to reflect not just anticipated personal consumption, but also activity in the B&C sector. This has echoes of recommendations made in the recent report of the Tax Forecasting Methodology Review Group, as well as the previous report published more than ten years ago.

For PRSI, we find that an elasticity of one produces the best forecasts. This is consistent with the current methodology used by the Department of Employment Affairs and Social Protection to forecast PRSI receipts. PRSI policy changes have been relatively modest in the sample period, so policy-adjusted elasticities and unadjusted elasticities are almost equal.

This paper shows that the use of policy-adjusted elasticities could improve forecast accuracy in Ireland. However, the implications could be wider still, as policy-adjusted elasticities could be estimated and then used for fiscal forecasting in many other countries.

6. References

Acheson, J., Y. Deli, D. Lambert and E. Morgenroth (2017). “Income Tax Revenue Elasticities in Ireland: An Analytical Approach”, ESRI Research Series, No. 59. Available at: https://www.esri.ie/system/files?file=media/file-uploads/2017-03/RS59.pdf

Acheson, J., Y. Deli, E. Morgenroth, D. Lambert and M. Murphy (2018). “VAT Revenue Elasticities in Ireland: An Analytical Approach”, ESRI Working paper No. 596. Available at: https://www.esri.ie/system/files/media/file-uploads/2018-10/WP596.pdf

Addison-Smyth, D. and K. McQuinn (2016). “Assessing the sustainable nature of housing-related taxation receipts: the case of Ireland”, Journal of European Real Estate Research, Vol. 9 Issue: 2, pp.193-214. Available at: https://www.emeraldinsight.com/doi/full/10.1108/JERER-01-2016-0004

Buettner, T. and B. Kauder (2010). “Revenue Forecasting Practices: Differences across Countries and Consequences for Forecasting Performance”. Fiscal Studies, Vol. 31, No. 3, September 2010, pp. 313-340. Available at: https://www.jstor.org/stable/24440135?seq=1#page_scan_tab_contents

Casey, E., and D. Smyth (2016). “Revisions to Macroeconomic Data: Ireland and the OECD”. The Economic and Social Review, Vol. 47 (Spring), pp. 33–68. Available at: https://www.esr.ie/article/view/505

Conroy, N. (2015). “Irish Quarterly Macroeconomic Data: A Volatility Analysis.” Special article in Quarterly Economic Commentary, Summer 2015, pp. 44-59. Available at: https://www.esri.ie/pubs/RN20150201.pdf

Conroy, N., (2020). “Estimating Ireland’s Tax Elasticities: A Policy-Adjusted Approach”. The Economic and Social Review, Vol. 51, No.2, Summer 2020, pp. 241–274. Available at: https://www.esr.ie/article/view/1443

Cramér, H., (1928), “On the composition of elementary errors”. Skandinavisk Aktuarietidskrift, Vol. 1, pp. 13–74, Available at: https://www.tandfonline.com/doi/abs/10.1080/03461238.1928.10416862

Deli Y.D., D. Lambert, M. Lawless, K. McQuinn and E. Morgenroth (2017). “How Sensitive Is Irish Income Tax Revenue to Underlying Economic Activity?” The Economic and Social Review, Vol. 48, No. 3, Autumn, 2017, pp. 317-336. Available at: https://www.esr.ie/article/view/772

Department of Finance (2011). “Jobs initiative”. May 2011. Available at: https://merrionstreet.ie/en/wp-content/uploads/2011/05/Jobs_Initiative_Booklet_10_May_2011.pdf

Diebold, F.X. and Mariano, R.S., (2002). “Comparing predictive accuracy”. Journal of Business & economic statistics, Vol. 20, No. 1, pp.134-144. Available at: https://www.tandfonline.com/doi/abs/10.1198/073500102753410444

Fioramanti M., L. González-Cabanillas, B. Roelstraete and S.A. Ferrandis-Vallterra (2016) "European Commission's Forecasts Accuracy Revisited: Statistical Properties and Possible Causes of Forecast Errors", European Economy - Discussion Paper N.027. Available at: https://ec.europa.eu/info/sites/info/files/dp027_en.pdf

Frankel, J., and J. Schreger (2013). “Over-Optimistic Official Forecasts and Fiscal Rules in the Eurozone.” Review of World Economics 149(2): 247–72. Available at: https://link.springer.com/article/10.1007/s10290-013-0150-9

Hannon, H., E. Leahy and R. O’Sullivan (2015). “An Analysis of Tax Forecasting Errors in Ireland”. Irish Fiscal Advisory Council Working Paper No. 3, September 2015. Dublin: Irish Fiscal Advisory Council. Available at: https://www.fiscalcouncil.ie/wp-content/uploads/2012/01/An-Analysis-of-Tax-Forecasting-Errors-in-Ireland-final.pdf

Merola, R and J. Pérez (2013). "Fiscal forecast errors: Governments versus independent agencies?" European Journal of Political Economy, Elsevier, vol. 32(C), pages 285-299. Available at: https://www.sciencedirect.com/science/article/pii/S0176268013000682

Price, R., T. Dang and Y. Guillemette (2014). "New Tax and Expenditure Elasticity Estimates for EU Budget Surveillance", OECD Economics Department Working Papers, No. 1174, OECD Publishing, December 2014. Available at: http://www.oecd.org/officialdocuments/publicdisplaydocumentpdf/?cote=ECO/WKP(2014)70&docLanguage=En

Tax Forecasting Methodology Review Group, 2008. "Report of the Tax Forecasting Methodology Review Group". Available at: http://s3.amazonaws.com/zanran_storage/www.finance.gov.ie/ContentPages/8987749.pdf

Tax Forecasting Methodology Review Group, 2019. "Tax Forecasting Methodological Review 2019". Available at: https://www.gov.ie/en/publication/76468a-tax-forecasting-methodological-review-2019/

Wilcoxon, F., (1945). "Individual comparisons by ranking methods". Biometrics Bulletin, vol. 1, No.6, pp 80–83. Available at: https://www.jstor.org/stable/3001968?seq=1

Appendix A: Results using Alternative Macroeconomic Drivers

This appendix shows the results of forecasting income tax and PRSI using different macroeconomic drivers. From Table A.1 we can see that non-agricultural income performs best in forecasting income tax receipts.

Table A.1: Income Tax Forecast Errors using Different Macroeconomic Drivers

Forecast errors are as a percentage of income tax

Macro Driver	Elasticity[1]	Average Absolute Error[2]
Non-Agri income	1.5	6.6
Domestic GVA	1.3	8.8
GNI*	1.3	8.7

Sources: CSO, Department of Finance and author's calculations.
Note: 1 This is the elasticity which is calibrated to minimise the average absolute error over the whole forecast horizon for that macroeconomic driver. 2 Equal weight is given to the average absolute errors (as a percentage of income tax revenue) for each of the five points on the forecast horizon (years T to T+4).

Using non-agricultural income produces smaller forecast errors at all points of the forecast horizon, from in year forecasts (T) to four-year-ahead forecasts (T+4). Using Diebold-Mariano tests, we can see that these differences in forecasting performance are statistically significant. (Table A.2) There is no significant difference between using Domestic GVA or GNI* for forecasting.

Table A.2: Diebold-Mariano Tests of Macroeconomic Drivers, Income Tax

Relative RMSE

Macroeconomic driver	Relative RMSE
Non-Agri income (1.5) vs GNI* (1.3)	0.69**
Non-Agri income (1.5) vs Domestic GVA (1.3)	0.67**
Domestic GVA (1.3) vs GNI* (1.3)	0.98

Sources: Author's own calculations.
Note: Values below one suggest that the first forecast is superior. Values with ** or * indicates forecasts are significantly different at a 1% or 10% significance level.

For PRSI, we also find that non-agricultural income produces forecasts with the smallest average absolute errors (Table A.3). Using Diebold-Mariano tests, we can see that improvement in forecasting performance from using non-agricultural income as the macroeconomic driver is statistically significant (Table A.4).

Table A.3: PRSI Forecast Errors using Different Macroeconomic Drivers

Forecast errors are as a percentage of PRSI

Macro Driver	Elasticity[1]	Average Absolute Error[2]
Non-Agri income	1.0	6.0
Domestic GVA	0.9	7.4
GNI*	0.8	8.7

Sources: CSO, Department of Finance and author's calculations.
Note: 1 This is the elasticity which is calibrated to minimise the average absolute error over the whole forecast horizon for that macroeconomic driver. 2 Equal weight is given to the average absolute errors (as a percentage of PRSI revenue) for each of the five points on the forecast horizon (years T to T+4).

Table A.4: Diebold-Mariano Tests of Macroeconomic Drivers, PRSI

Relative RMSE

Macroeconomic driver	Relative RMSE
Non-Agri income (1.0) vs GNI* (0.8)	0.60**
Non-Agri income (1.0) vs Domestic GVA (0.9)	0.69**
Domestic GVA (0.9) vs GNI* (0.8)	0.87**

Sources: Author's own calculations.
Note: Values below one suggest that the first forecast is superior. Values with ** or * indicates forecasts are significantly different at a 1% or 10% significance level.

Appendix B: Forecasting USC and Income Tax Separately

This appendix shows the results of attempts to forecast USC and other income tax in the short period that the USC has been in place. For the years 2011-2019, using data from the office of the Revenue Commissioners, we can separate out net receipts of USC from the rest of income tax. While the income tax and USC data used here are not directly comparable to that used in the main results section (it is not on an exchequer basis), it allows for some comparison of the different forecast methodologies. Two approaches are taken to forecast income tax for the period 2012-2019.

First, we use the elasticities used by the Department of Finance to forecast USC receipts and other income tax receipts separately (1.2 and 2.1 respectively). We then aggregate these two forecasts and compare them to the outturns. Second, we use a weighted average of these two elasticities (1.89) to forecast the aggregate of USC and other income tax.

Figure B.1 shows the average absolute percentage errors for both of these approaches. At each forecast horizon, the weighted elasticity (1.89) performs better, with smaller errors on average. It is perhaps surprising that the weighted elasticity outperforms the two separate elasticities. However, all findings here are tempered by there being so few observations, due to the USC being only recently introduced.[41]

[41] Due to the very small sample sizes, Diebold-Mariano tests are not performed.

Figure B.1: Income Tax Forecast Errors

Average absolute forecast error (per cent of revenue)

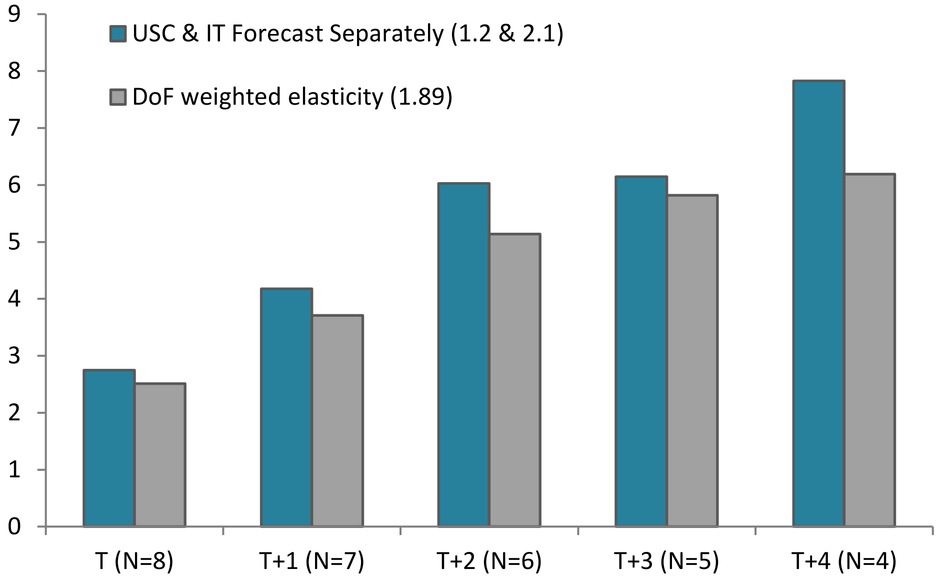

Sources: Author's own calculations and Office of the Revenue Commissioners.
Note: The dark blue bars show the average absolute forecast errors when forecasting USC and other income tax receipts separately, and the combining these for a forecast of aggregate income tax receipts.

Appendix C: Assessing Bias in forecasts

This appendix shows the results of various attempts to examine average forecast errors for the different revenue headings examined. We want to establish if average errors are significantly different from zero. If the errors are normally distributed, a simple T-test that the average error equals zero will tell us if the forecasts are biased. Due to our small sample sizes, our tests have relatively low power. Thus, a failure to reject a null hypothesis cannot be considered conclusive evidence that the null is true. Instead, this is an indication that within our sample, we cannot find enough evidence to support rejecting the null hypothesis.

We test for normality using the Cramér–von Mises criteria (Cramér, 1928). If we find evidence to reject the null (that errors are distributed normally), then we employ a nonparametric test (the Wilcoxon signed-rank test) of whether the median is zero. The null hypothesis for this test is that errors are centred on zero.

Table C.1: Assessing Average Income Tax Forecast Errors

P-values

Elasticity	Forecast Horizon	T-test	Normality test	Wilcoxon signed-rank test
1.4	T	0.99	0.00**	0.33
1.4	T+1	0.94	0.06	0.56
1.4	T+2	0.65	0.37	0.41
1.4	T+3	0.41	0.85	0.50
1.4	T+4	0.24	0.62	0.31
1.89	T	0.00**	0.03*	0.00**
1.89	T+1	0.00**	0.25	0.00**
1.89	T+2	0.00**	0.53	0.00**
1.89	T+3	0.00**	0.69	0.00**
1.89	T+4	0.00**	0.41	0.00**
0.83	T	0.02*	0.00**	0.00**
0.83	T+1	0.01*	0.00**	0.00**
0.83	T+2	0.00**	0.00**	0.00**
0.83	T+3	0.00**	0.01*	0.00**
0.83	T+4	0.00**	0.04*	0.00**

Sources: CSO, Department of Finance and author's calculations.
Note: * and ** indicate statistical significance at a 5% and 1% level respectively. Rejecting the null in the normality test suggest that the forecast errors are not normally distributed and hence the Wilcoxon signed-rank test may be a better way to assess if errors are centred on zero. For both the T-test and the Wilcoxon signed-rank test, the null is that errors are centred on zero.

For income tax, we find that forecast errors from elasticities of 1.4 and 1.89 are broadly normal, hence T-tests may be appropriate. In any event, the T-tests and Wilcoxon signed-rank test give similar results. Forecast errors using the policy adjusted elasticity (1.4) are centred on zero. This is the case at all forecast horizons. For the weighted elasticity (1.89), we find that errors are not centred on zero at all forecast horizons.

For the unadjusted elasticity (0.83), it appears that the forecast errors are not normally distributed. As a result, it is best to focus on the Wilcoxon signed-rank test. Looking at this, the median forecast errors are significantly different to zero at all forecast horizons.

Table C.2: Assessing VAT Forecast Errors

P-values

Elasticity	Forecast horizon	T-test	Normality test	Wilcoxon signed-rank test
0.8 & 0.2	T	0.89	0.68	0.99
0.8 & 0.2	T+1	0.77	0.19	0.85
0.8 & 0.2	T+2	0.76	0.74	0.81
0.8 & 0.2	T+3	0.91	0.63	0.97
0.8 &0.2	T+4	0.70	0.38	0.45
0.9 & 0.2	T	0.60	0.57	0.74
0.9 & 0.2	T+1	0.38	0.44	0.52
0.9 & 0.2	T+2	0.25	0.61	0.39
0.9 & 0.2	T+3	0.22	0.54	0.32
0.9 & 0.2	T+4	0.26	0.99	0.28
1.0	T	0.62	0.02*	0.25
1.0	T+1	0.65	0.00**	0.10
1.0	T+2	0.69	0.00**	0.06
1.0	T+3	0.73	0.00**	0.12
1.0	T+4	0.70	0.00**	0.15
1.1	T	0.90	0.05	0.65
1.1	T+1	0.81	0.00**	0.44
1.1	T+2	0.74	0.00**	0.37
1.1	T+3	0.66	0.00**	0.39
1.1	T+4	0.65	0.00**	0.39

Sources: CSO, Department of Finance and author's calculations.
Note: * and ** indicate statistical significance at a 5% and 1% level respectively. Rejecting the null in the normality test suggest that the forecast errors are not normally distributed and hence the Wilcoxon signed-rank test may be a better way to assess if errors are centred on zero. For both the T-test and the Wilcoxon signed-rank test, the null is that errors are centred on zero.

Looking at Table C.2, the first two sections focus on results using consumption and B&C investment. In both cases it appears that the forecast errors are normally distributed. In any event, the T-tests and Wilcoxon signed-rank tests both suggest that errors are centred on zero when this approach is used.

The bottom half of Table C.2 shows the results when using only personal consumption to forecast VAT receipts. There is evidence that these forecast errors are not normally distributed. When looking at the Wilcoxon signed-rank test, it would appear that these forecast errors are centred on zero. In summary, all four sets of forecast errors appear to be unbiased.

Table C.3: Assessing Average PRSI Forecast Errors

P-values

Elasticity	Forecast horizon	T-test	Normality test	Wilcoxon signed-rank test
1.0	T	0.76	0.22	0.67
1.0	T+1	0.55	0.88	0.38
1.0	T+2	0.43	0.48	0.37
1.0	T+3	0.36	0.62	0.25
1.0	T+4	0.30	0.22	0.25
1.02	T	0.91	0.30	0.83
1.02	T+1	0.75	0.82	0.59
1.02	T+2	0.64	0.45	0.48
1.02	T+3	0.59	0.49	0.40
1.02	T+4	0.54	0.13	0.36
1.5	T	0.01*	0.77	0.01*
1.5	T+1	0.00**	0.71	0.01*
1.5	T+2	0.00**	0.53	0.00**
1.5	T+3	0.00**	0.60	0.00**
1.5	T+4	0.00**	0.16	0.00**

Sources: CSO, Department of Finance and author's calculations.
Note: * and ** indicate statistical significance at a 5% and 1% level respectively. Rejecting the null in the normality test suggest that the forecast errors are not normally distributed and hence the Wilcoxon signed-rank test may be a better way to assess if errors are centred on zero. For both the T-test and the Wilcoxon signed-rank test, the null is that errors are centred on zero.

Looking at Table C.3, the policy-adjusted elasticity and unadjusted elasticity produce very similar results. In both cases, forecast errors appear to be normally distributed. In any event, both the T-tests and Wilcoxon signed-rank tests indicate that both sets of forecasts are centred on the outturns. This is consistent across the whole forecast horizon. For the Price et al (2015) elasticity (1.5), we also find no significant evidence that the errors are not normally distributed. Both the T-tests and the Wilcoxon signed-rank tests indicate that errors are not centred on zero. This appears to be the case at all points in the forecast horizon.