

The Anticipation Content of Fiscal Policy: UK Tax Measures, 1945–2019

Alessandro Conway

Working Paper AC-WP-2021-01

Version 2.0 (revised September 2026)

Not peer reviewed. Comments welcome.

Abstract

Narrative tax shocks are dated by implementation, and the resulting series are pooled across decades. Using 2,252 UK tax measures announced between 1945 and 2019, I show that the interval between announcement and implementation has changed enough for such a series to mix two economically different objects. The share of the tax impulse announced more than 120 days before taking effect rises from 0.18 in 1945–79 to 0.70 since 2000, and the unanticipated component of the quarterly impulse falls roughly fivefold. Computing the same quantity from the century-long dataset of Cloyne et al. (2025), on their coding and their threshold, reproduces the series and extends it back to 0.02 in 1920–44. Its length is also systematic. Within the same Budget, National Insurance changes are 39 percentage points more likely than excise duties to arrive with long notice, the fiscal-year calendar produces long notice only when the Budget falls late in the year, and tax rises are four times more likely than cuts to take effect after an election. The rise dates to the early 1990s, when the Budget process was reformed in the name of scrutiny and predictability, and its macroeconomic consequences were never assessed; the crises of the period do not appear to have shortened it.

Keywords: narrative tax shocks, fiscal foresight, tax policy making, announcement effects, fiscal multipliers

JEL classification: D84, E62, E65, H20, H30

Suggested citation: Conway, Alessandro (2021). “The Anticipation Content of Fiscal Policy: UK Tax Measures, 1945–2019.” Working Paper AC-WP-2021-01, Version 2.0.

© Alessandro Conway 2026. Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0).

The Anticipation Content of Fiscal Policy: UK Tax Measures, 1945–2019

Alessandro Conway*

First draft December 2021

This version September 2026

Abstract

Narrative tax shocks are dated by implementation, and the resulting series are pooled across decades. Using 2,252 UK tax measures announced between 1945 and 2019, I show that the interval between announcement and implementation has changed enough for such a series to mix two economically different objects. The share of the tax impulse announced more than 120 days before taking effect rises from 0.18 in 1945–79 to 0.70 since 2000, and the unanticipated component of the quarterly impulse falls roughly fivefold. Computing the same quantity from the century-long dataset of Cloyne et al. (2025), on their coding and their threshold, reproduces the series and extends it back to 0.02 in 1920–44. Its length is also systematic. Within the same Budget, National Insurance changes are 39 percentage points more likely than excise duties to arrive with long notice, the fiscal-year calendar produces long notice only when the Budget falls late in the year, and tax rises are four times more likely than cuts to take effect after an election. The rise dates to the early 1990s, when the Budget process was reformed in the name of scrutiny and predictability, and its macroeconomic consequences were never assessed; the crises of the period do not appear to have shortened it.

Keywords: narrative tax shocks, fiscal foresight, tax policy making, announcement effects, fiscal multipliers.

JEL classification: D84, E62, E65, H20, H30.

*Working paper. The first draft dates from December 2021 and was not circulated; this version revises the analysis throughout and cites work published since. I thank James Cloyne for making the underlying narrative dataset public, without which this paper could not exist.

1 Introduction

Every tax measure has two dates, the day it is announced and the day it takes effect. The narrative tradition in fiscal policy has recorded both for forty years, but it has not treated the interval between them as a quantity in its own right. Romer and Romer (2010) identify legislated US tax changes from the contemporaneous record, code their motivation, and date them by when they take effect, and Cloyne (2013) builds the corresponding UK account for 1945 to 2009. That record has since been extended in both directions, back into the interwar period by Cloyne et al. (2023) and forward to 2020 by Hürtgen et al. (2024). A narrative account of British tax policy now spans roughly a century, and Cloyne et al. (2025) have used it at that length.

Throughout that work the implementation date is a dating convention, an input to the construction of a shock series and of no interest in itself. The announcement date is used, where it is used at all, to separate measures that were foreseen from those that were not.

The literature on fiscal foresight puts the same dates to a sharper use, asking whether anticipation undermines identification. Yang (2005) shows that policy foresight changes the dynamic response to a tax change. Leeper et al. (2013) establish that foresight places the econometrician behind the agent’s information set, leaving a non-invertible representation from which the shock cannot be recovered in the usual way. Mertens and Ravn (2011, 2012) construct separate anticipated and unanticipated US tax series in order to estimate their effects apart. Ramey (2011) makes the parallel argument for government spending, that the timing of the news rather than the timing of the outlay identifies the shock. Favero and Giavazzi (2012) and Alesina et al. (2015) move closer to the policy process itself, the latter treating fiscal consolidations as multi-year plans with announced and unexpected components rather than as sequences of independent measures. Of the existing work, theirs comes nearest to the treatment taken here.

Both bodies of work handle anticipation at the level of the individual measure. Foresight is something to be recorded, corrected for, or modelled, and the correction is applied measure by measure within a sample treated as homogeneous. Neither assembles the anticipation of individual measures into a property of the policy regime, or asks how much of it there is in a given decade, whether it falls disproportionately on particular instruments, or whether the quantity has moved over the life of the sample. That it has not moved is assumed throughout, although the assumption is never stated and so never defended.

A second literature, entirely separate from the first, does take the timing of tax policy as its subject, although it treats timing as a matter of procedure. The parliamentary record and the Treasury’s own accounts of how tax policy is made document four reforms of the Budget process: the unification of the Budget and the Autumn Statement announced in 1992, the consultation framework of 2010 and 2011, and the single fiscal event of 2016. These documents are explicit about lead times, and the Treasury now states a target of sixteen months from announcement to effect.

The two literatures do not refer to each other. The institutional record justifies longer lead times by reference to parliamentary scrutiny and the predictability of the tax system. It contains no assessment of what a longer lead time does to the transmission of a fiscal impulse,

which was never its purpose.

This paper builds on both. The macroeconomic literature has established that foresight matters and has built the tools to handle it, one measure at a time and within a sample assumed to be stationary in this respect. The institutional record shows that the process generating foresight was deliberately reformed, twice, without any consideration of the consequence. Neither has measured foresight itself as a time series, asked whether it is systematically distributed, or connected its movement to the reforms that produced it, and this paper attempts each in turn.

On 2,252 UK tax measures across 122 Budget events, the share arriving with at least 120 days' notice runs at roughly 15 per cent from the 1950s through the 1980s and reaches 65 per cent in the 2010s. Weighted by revenue, the share of the tax impulse announced more than 120 days ahead rises from 0.18 in 1945–79 to 0.70 since 2000, and the unanticipated component of the quarterly impulse falls by a factor of about five. The quantity also replicates on data I did not build: computing it from the century-long series published by Cloyne et al. (2025), on their coding and their threshold, gives 0.25, 0.42, and 0.70 across the same eras, and extends the finding back to 0.02 in 1920–44.

The interval is also systematic, and the comparisons that show it hold within Budget events, so that measures announced by the same Chancellor on the same afternoon are compared with one another. National Insurance changes are 38.5 percentage points more likely than excise duties to arrive with long notice, income tax changes 12.8 points more likely, and capital gains changes 11.2 points less likely. Among measures pinned to early April by the fiscal year, 32 per cent arrive with long notice after a spring Budget and 86 per cent after an autumn one, so the same rule applied to the same measure gives three weeks or five months of warning depending only on when the Budget falls. Among measures announced six to twenty-four months before a general election, 12.5 per cent of tax rises take effect after the country has voted, against 3.1 per cent of tax cuts.

Crises do not shorten the interval on any of three definitions of treatment, and each estimate is either null or of the opposite sign.

Section 2 describes the data, what the source can and cannot record, and the tests that establish the trend is not an artefact of joining two sources. Sections 3 to 6 present the four results and Section 7 the crisis null. Section 8 converts the measure-level findings into the anticipation content of the revenue impulse and reports the external replication. Section 9 dates the change and sets it against the institutional record.

2 Data

2.1 Construction

Cloyne (2013) codes UK tax measures from 1945 to 2009 from Financial Statement and Budget Reports, recording an announcement date, an implementation date, a revenue costing, and a motivation classification. I extend the record from Budget 2004 to Autumn 2018 from Treasury scorecards, using the same concepts and the same fields. The two are chained at 1 January 2004, Cloyne's coding before and mine after, which gives 2,252 datable measures across 122

Budget events.

The announcement date recorded here is the date of the scored, dated commitment, taken from the Budget report or the scorecard, and 90 per cent of measures carry Budget day itself. Agents may learn earlier that a measure is coming: since 2010 the tax policy making framework has placed a consultation stage ahead of that date for many measures, so for those the information set moved earlier than the recorded date.

The measured lead is therefore a lower bound on foresight, and the bound is loosest in the period where the measured lead is longest, so the rise documented below understates the rise in true foresight. Every quantity reported is the formal lead time of scored policy, which coincides with foresight for the measures that arrive without prior consultation.

The sample is every datable, non-reversal measure regardless of exogeneity. How long the policy process takes is a descriptive question with no endogeneity to purge. Imposing the identifying restriction that a causal estimate would require would discard 41 per cent of Cloyne’s datable measures and 69 per cent of mine, for no methodological gain. Exogeneity is carried through the data because Section 8 needs it and because the contrast between exogenous and endogenous measures is itself informative.

2.2 The outcome

The primary outcome is an indicator for a gap of at least 120 days between announcement and implementation. The concept comes from Cloyne (2013), who, following Mertens and Ravn (2012), classifies a tax change as anticipated when its implementation lag exceeds 90 days, a convention the literature has used since. The 120-day threshold used here is a variant of it, widened by one month for a reason specific to the British fiscal calendar.

I considered three ways of measuring the interval. A fiscal-year indicator, recording whether implementation falls in a later fiscal year than announcement, is partly artefactual. A Budget delivered in March and implementing on 6 April crosses the boundary on three weeks’ notice, and the frequency of that case drifts from 0.72 of such crossings in the 1980s to 0.11 in the 2010s, so the indicator changes meaning over the period being studied. The raw lag in quarters fails for the opposite reason, with 42 per cent of its mass at zero. A count of days avoids both problems, and it is the measure used here.

Within a day count, the choice between 90 and 120 days depends on the same fiscal calendar. A Budget delivered in late November or December that implements on the following 1 or 6 April sits about 110 days out, so a 90-day rule records it as anticipated when it is simply the routine start of the next fiscal year. Widening the threshold to 120 days removes that case, and the correction is small: 94 measures, 4.2 per cent of the sample, fall between 91 and 119 days, and 32 of them are autumn or winter Budgets pinned to the following April. Table 1 reports the two headline results on all four outcomes, including Cloyne’s own 90-day rule, and the sign and significance are the same throughout.

All inference is unweighted and clustered on the Budget event. Revenue weighting reduces the Kish effective sample from 2,252 to 251, so weighting is used for descriptive magnitudes only, and never for tests.

Table 1: The headline results on four candidate outcomes

Outcome	2010s against the 1940s		Social security within Budget	
	Effect	z	Effect	z
120 days or more (primary)	+0.483	6.84	+0.385	7.48
90 days or more (Cloyne, 2013)	+0.506	7.96	+0.356	6.93
Crosses the fiscal year	+0.634	9.68	+0.543	11.59
Lag in quarters	+2.917	11.66	+1.823	3.90

The first three outcomes are indicators and the fourth is a count of quarters, so only the first three rows are on a common scale. Standard errors are clustered on the Budget event.




Figure 1: Announcement-to-implementation gap. The mass below thirty days is the March Budget to 6 April convention, a fiscal-year crossing on three weeks' notice.

2.3 Scope of the scorecards

A scorecard records commitments, and it contains nothing on whether they were delivered. Across 3,092 rows there is not one case in which an expiry date reflects a later decision rather than a sunset announced at the outset. Every reversal row is announced on or before its parent and implements exactly on the parent’s stop date. A measure that was announced and later dropped without ceremony therefore leaves no trace in the source.

Every claim below concerns scored, dated commitments. Where a sunset is announced it is recorded, and sunsets have become more common: 10.6 per cent of post-2004 measures carry one, against 3.6 per cent before.

2.4 Joining two sources

The trend crosses a join at 2004, with Cloyne’s coding before it and mine after. If the two codings recorded announcements or itemised Budgets differently, part of the rise could be an artefact of the change of coder. Six tests address that possibility, and none supports it.

Both codings place roughly nine tenths of announcements on the Budget day itself, 0.916 for Cloyne and 0.888 for mine, so the announcement date is the same object in each. The years 2004 to 2009 are coded twice. On the eleven Budgets both records contain, Cloyne’s coding gives a long-notice share of 0.460 and mine 0.496, or 0.627 and 0.615 weighted by revenue, with a paired within-event difference of -0.021 and a t statistic of -0.46 . Cloyne splits those same Budgets 2.8 times more finely, at 32.0 measures per event against 11.5, and still returns the same share, so differences in itemisation are not behind the result.

The trend itself holds unweighted, from 0.162 to 0.646, at the level of the Budget event, from 0.195 to 0.618, and weighted by revenue, from 0.285 to 0.740. Holding the 1945–79 instrument mix fixed and letting only the within-instrument rates move changes the 2010s figure by 0.019, so the rise does not come from a change in which taxes governments use. Taking the twenty largest measures of each decade, an equal count that differences in itemisation cannot affect, gives 0.25, 0.10, 0.20, 0.30, and 0.20, then 0.70, 0.80, and 0.80.

The most direct test is that the post-1990 step is $+0.206$ ($z = 5.18$) within Cloyne’s coding alone, and $+0.175$ ($z = 4.43$) with instrument controls; since the step lies entirely on one side of the join, the join cannot have produced it.

3 Advance notice over time

Table 2 reports the share of measures arriving with at least 120 days’ notice, by decade of announcement. The series is flat through the 1980s and rises from the 1990s, and no decade depends on a single Budget, since the largest event accounts for between 11 and 25 per cent of a decade’s measures.

4 Instruments

Comparing measures announced in the same Budget holds constant the Chancellor, the political situation, and the macroeconomic conditions of the day. The comparison still yields only an

Table 2: Share of measures with at least 120 days between announcement and implementation

	1940s	1950s	1960s	1970s	1980s	1990s	2000s	2010s
Share	0.162	0.039	0.129	0.182	0.145	0.363	0.393	0.646
Measures	80	102	178	203	394	534	377	384
Budget events	6	13	17	21	13	17	22	20




Figure 2: The long-notice share in five-year blocks.

association between instrument and notice, since measures within a Budget differ in administrative, legislative, and statutory-calendar characteristics, and Section 5 shows that the fiscal calendar accounts for much of it. Table 3 reports the effects relative to excise duties. The joint test gives $F = 11.7$ on 10 degrees of freedom, and R^2 rises from 0.274 with Budget fixed effects alone to 0.314 once instrument is added.

Table 3: Instrument effects on the long-notice indicator, Budget-event fixed effects

Instrument	Effect	z	p
Social security (NICs)	+0.385	7.48	0.000
Property, recurrent	+0.148	2.03	0.042
Income	+0.128	3.46	0.001
VAT	+0.096	2.27	0.023
Corporate	+0.036	0.83	0.406
Capital gains	−0.112	−3.25	0.001

The reference category is excise and duties. Oil, property transaction, inheritance, and other are estimated and insignificant. Standard errors are clustered on the Budget event.

The specification carries 128 parameters against 118 clusters, so the cluster-robust variance matrix is rank deficient and the standard errors are suspect in principle. A pairs cluster bootstrap over 600 replications, resampling whole Budget events, returns standard deviations of 0.046, 0.076, 0.035, 0.043, and 0.034 for social security, recurrent property, income, VAT, and capital gains, against cluster-robust standard errors of 0.051, 0.073, 0.037, 0.042, and 0.034. A within-transformed specification, demeaning by Budget event and so carrying ten parameters against the same 118 clusters, returns identical point estimates and standard errors of 0.050, 0.071, 0.036, 0.041, and 0.033. The result is not an artefact of the estimator, with one qualification. Recurrent property is marginal on every version of the variance, and 74 of the 600 bootstrap samples contain no recurrent property measure at all, so that row is the least secure of the five.

5 The fiscal-year calendar

Income tax and National Insurance changes are pinned to early April by the fiscal year, which is why they sit at the top of Table 3. The pinning produces long notice only when the Budget falls late in the year, and this can be tested directly on the 1,050 measures implementing between 1 and 7 April. Among those measures the long-notice rate is 0.32 after a spring Budget and 0.86 after an autumn one. The raw difference of 0.54 overstates the effect, because autumn Budgets are concentrated in the later decades when notice was longer for every instrument; conditioning on decade gives +0.365 with $z = 5.14$, clustered on the Budget event. For the same rule and the same measure, the warning a taxpayer receives therefore differs by five months according to where in the year the Budget falls.

The calendar explains variation across measures, but very little of the trend: adding calendar controls raises R^2 from 0.155 to 0.214 and absorbs only 9 per cent of the 2010s decade effect. The implementation date has also migrated within April, the 6 April share falling from 0.47 in the 1970s to 0.06 in the 2010s while the 1 April share rises from 0.10 to 0.53, which reflects a




Figure 3: Instrument effects with clustered 95 per cent confidence intervals.

change in administrative convention rather than in notice.

6 Elections

Among measures announced between six and twenty-four months before a general election, 3.1 per cent of tax cuts and 12.5 per cent of tax rises take effect after the election. Within Budget the difference is $+0.075$ ($z = 2.37$, $p = 0.018$) on 751 measures across 49 events, and on the wider window of nought to twenty-four months it is $+0.060$ ($z = 2.26$, $p = 0.024$).

The pattern is a descriptive association, and for several reasons it cannot yet be read as deliberate scheduling. Election dates were not fixed before 2011, so the vote a measure is compared against was itself partly a government choice, and the causation can run from the fiscal calendar to the election date rather than the other way. The date was in any case not known when most of these measures were announced, so the test conditions on a realised election rather than an expected one. Budget fixed effects also leave within-Budget differences in instrument mix and statutory calendar untouched, and both correlate with sign, since tax rises are concentrated in the fiscal-year-pinned instruments of Section 5, which take effect later by construction.

Restricting the sample to the fixed-term period would deal with the first two objections, but eight years of data are too few for the test, so the pattern stands as an association whose interpretation remains open.

Why instruments differ: the fiscal-year clock

Measures pinned to early April, by Budget season.
The same measure is three weeks from a March Budget and five months from a November one.




Figure 4: April-pinned measures by Budget season.

Tax rises are scheduled past the next election

Share of measures whose implementation date falls after the next general election, for measures announced 6-24 months before that election




Figure 5: Tax rises and cuts taking effect after the next general election.

7 Crises

Taking the years 1974–76, 1992–93, and 2008–09 as crisis windows and conditioning on decade, the estimate is +0.047 on 753 endogenous measures of which 159 are in a window ($p = 0.57$, 95 per cent interval -0.113 to $+0.206$), and +0.070 on 1,208 exogenous measures of which 84 are in a window ($p = 0.65$, interval -0.233 to $+0.372$). Both are positive and neither is significant, and earlier tests on the population of crisis-response measures, and on years containing three or more fiscal events, are either null or of the wrong sign.

The intervals rule out a crisis shortening of more than 11 percentage points on the endogenous sample and 23 on the exogenous one. The smallest true shortening this design would detect four times in five is 23 and 43 points respectively, however, so a moderate crisis effect could sit inside those bounds.

The descriptive record, if anything, points more strongly the other way. Of the 22 exogenous measures announced in 2008–09, 95.5 per cent carried long notice, against 34.4 per cent of the 195 announced elsewhere in that decade. Twenty-two measures are too few for a formal test, but the contrast has the same sign as the estimates. In this sample crises bring more pre-commitment, which is the opposite of what one would expect if short notice were a capacity held in reserve for emergencies.

8 The anticipation content of the fiscal impulse

Sections 3 to 6 count measures, whereas the object other researchers use is the revenue impulse, and the two need not move together, since a Budget can pre-announce many small measures and spring one large one. Table 4 reports the revenue-weighted share of the impulse, dated by implementation, that was announced at least 120 days earlier. Measures implementing after 2019 are dropped, since they are the long-lead measures still pending when the sample ends.

Because these quantities are the ones another researcher would want to reconstruct, the aggregation rules are stated here. Write v_i for the costing of measure i , signed so that a rise is positive, and ℓ_i for the indicator that its announcement preceded implementation by at least 120 days. The anticipation share in Table 4 is $\sum_i |v_i| \ell_i / \sum_i |v_i|$, taken over measures whose implementation date falls in the era, and the revenue-weighted lead is $\sum_i |v_i| d_i / \sum_i |v_i|$ with d_i the gap in days censored below at zero. Both use absolute costings, so a rise and a cut of equal size contribute equally and neither nets against the other.

Costings are the peak-year figure at announcement, expressed as a share of that year's nominal GDP and never revised, so subsequent re-costings do not enter. A measure with several implementation dates is dated by the first. A measure with a pre-announced sunset is counted once, at implementation.

The quarterly decomposition below is a different object and nets within the quarter. Let F_t be the sum of v_i over measures implementing in quarter t with $\ell_i = 1$, and U_t the same sum over those with $\ell_i = 0$, each divided by nominal GDP in that quarter. The paired figures reported by subperiod are $\overline{|F_t|}$ and $\overline{|U_t|}$, the mean absolute net quarterly impulse in each component, and the dispersion figure is the standard deviation of the signed U_t . The two statistics measure different things, so the shares implied by them differ slightly: the first is the fraction of the

gross flow that was pre-announced, the second the size of a typical quarter’s net surprise.

Table 4: Share of the tax impulse announced at least 120 days ahead

Sample	1945–79	1980–99	2000–19
All measures	0.184	0.384	0.699
Household, exogenous	0.202	0.364	0.815
Revenue-weighted lead, days (all measures)	102	147	322

Decomposing the quarterly impulse into a foreseen and an unforeseen component, as a share of nominal GDP, gives 0.041 and 0.077 in 1945–79, 0.049 and 0.083 in the 1980s, 0.051 and 0.030 from 1990 to 2003, and 0.030 and 0.015 from 2004 to 2018. Most of the change is in the surprise component. The unforeseen impulse falls by a factor of about five, and its standard deviation from 0.240 to 0.033 per cent of GDP, while the foreseen impulse falls by much less, so the foreseen share rises. The fall occurs between the 1980s and the period from 1990 to 2003, a window lying entirely within Cloyne’s coding, so it cannot be a property of my own data.

8.1 External replication

The result above is computed on a series I assembled, and it can also be computed on someone else’s. Cloyne et al. (2025) use a narrative UK dataset covering roughly 1918 to 2020. Their replication package publishes the quarterly exogenous shock series in two variants, a baseline containing all narrative shocks and an unanticipated variant restricted to measures implemented within Cloyne’s 90-day window. Writing B_t for the baseline series and U_t for the unanticipated variant, the anticipation share for an era is $1 - \sum_t |U_t| / \sum_t |B_t|$, summing absolute quarterly values over the quarters in the era, so signs, within-quarter netting, and zero denominators do not arise. The share is therefore built by other researchers, on their own coding and threshold, and over a longer sample.

Table 5: Anticipation share of the tax impulse, two independent datasets

Era	Cloyne et al. (2025), exogenous, 90-day rule	This paper, all measures, 120-day rule
1920–44	0.022	not covered
1945–79	0.248	0.184
1980–99	0.420	0.384
2000–20 / 2000–19	0.697	0.699

The two columns are built differently: their series covers exogenous measures only and runs to 2020, while the column from this paper is the all-measures row of Table 4 and stops in 2019. The household-exogenous row of Table 4, which is closer to their sample in construction, gives 0.202, 0.364, and 0.815. The comparison shows that the level and shape of the series hold when the coding, the threshold, and the sample all change.

The two agree in shape and in level, and most closely where the sample is densest. Their decade series also contains the break, running 0.286 in the 1980s and 0.614 in the 1990s, which is where Section 9 dates it independently from measure-level data. Their longer sample also reaches back to 1920–44, when the anticipation share was 0.022, so over a century the share has risen from almost nothing to roughly four fifths.

The stock of announced but not yet implemented tax change rises with it. The median pending stock rises from 0.026 per cent of GDP in 1945–79 to 0.991 per cent in 2004–18, and the mean number of pending measures from 2.0 to 26.4. The median is the better statistic here, because the mean is dominated by the introduction of VAT, announced in March 1971 and in force in April 1973, which on its own left roughly 11 per cent of GDP pending. At the time a lead of that length was exceptional, and by the end of the sample long leads had become routine.

For applied work, the consequence is that a multiplier estimated on a UK narrative series pooled across the post-war period averages two regimes, one in which tax changes mostly arrived as surprises and one in which they mostly did not. The mix between them changes monotonically over the sample.

Whether that pooling biases the estimate, and in which direction, depends on how news enters the estimator and on whether anticipated and unanticipated UK measures generate different macroeconomic responses, which is the question of the companion paper. This paper establishes only that the information content of the series is not stationary, which is a property of the instrument itself. The announcement-dated and implementation-dated series built here are supplied separately, so that the two can be estimated apart, and Leeper et al. (2013) give the reasons for expecting the distinction to matter.

9 Dating the change

A grid search over single break years maximises fit at 1993, with an R^2 of 0.142, and the annual series is sharp around it, running 0.047 in 1992, 0.372 in 1993, 0.434 in 1994, and 0.600 in 1995. The break year is searched rather than assumed, so its nominal fit is not a test statistic and no p -value is attached to it. The five best-fitting years span 1989 to 1995, so the change is best dated to the early 1990s as a whole. Table 6 then tests each named candidate locally, on the eight years either side, taking the date as given.

Table 6: Institutional candidates, local windows of eight years either side

Candidate	Year	Estimate	z	p
Unified autumn Budget, first held Nov. 1993	1993	+0.288	6.96	< 0.0001
Return to a spring Budget	1997	+0.026	0.36	0.716
Tax policy making framework	2010	+0.257	3.20	0.001
Draft-clause consultation	2011	+0.157	1.80	0.072
Single fiscal event	2017	+0.108	1.20	0.231

Two of the candidates hold on the count of measures, 1993 and 2010, but only one of them survives revenue weighting, which matters because the object of Section 8 is the money rather than the count. Repeating the exercise on the annual revenue-weighted long-notice share, the grid search again picks 1993, with a higher R^2 of 0.431 and the same 1989 to 1995 spread, and the local test at 1993 strengthens to +0.429 ($t = 4.02$, $p = 0.001$). The 2010 candidate, by contrast, does not hold, at +0.190 with $p = 0.116$, because the 2010 framework moved the number of measures given long notice without detectably moving the share of the revenue. It

therefore belongs in the account of the trend but not in the account of the impulse, and only the 1993 date is supported on both.

Nor is the 1993 step an artefact of the calendar. Restricting the sample to spring Budgets alone, the long-notice share still runs 0.143 in the 1980s, 0.275 in the 1990s, and 0.657 in the 2010s, with a post-1990 step of +0.275 ($z = 5.56$). Over the last three decades the autumn-minus-spring gap runs +0.21, +0.17, and -0.03 , so the calendar effect disappears while the level continues to rise. The rejection of 1997 points the same way: when the Budget returned to spring that year, the long-notice share stayed where it was.

Both dates correspond to reforms of the Budget process. Norman Lamont announced the first in the Budget Statement of 10 March 1992, saying that “next year’s Budget will be the last spring Budget” and that “from then on the annual Budget will be in December, and it will cover not just taxation but also public expenditure”. He added that the 1994 Finance Bill would be presented in January rather than April, which moved the legislative timetable as well as the announcement. The date was brought forward to November in the Budget Statement of 16 March 1993, and the first unified Budget was delivered by Kenneth Clarke on 30 November 1993.

Budget 2010 set out predictability, stability, and simplicity as objectives and published proposals “to improve the way tax policy is made to support these objectives”. The consultation response of December 2010 recorded that confirming most intended tax changes “at least three months ahead of publication of the Bill supports predictability in the tax system”. The move to a single fiscal event in 2016 was made “to promote certainty and simplicity within the tax system”, and the Treasury’s current statement of the resulting timetable is that “most policies will be announced at least 16 months before they come into effect at the start of the next tax year”.

The stated objectives across 1992, 2010, 2016, and 2017 are parliamentary and external scrutiny, predictability, stability, certainty, simplicity, and time for stakeholders to comment, and none of them is macroeconomic. Each reform was argued for in public and chosen deliberately, yet the anticipation content of fiscal policy appears nowhere in those arguments.

The design cannot establish that the reforms caused the shift, since there is one country and one series, the break date is searched, candidate reforms in adjacent years cannot be separated on annual data, and a strong secular trend runs underneath. It does establish that the change coincides with the 1993 Budget-process reform, that the coincidence holds under revenue weighting and once the calendar effect is removed, and that the 1997 return to a spring Budget did not reverse it. Whatever caused the change, the anticipation content of British fiscal policy is now different from what the pooled samples of the literature assume, and the documents above give no sign that anyone intended it.

10 Conclusion

The interval between announcing a British tax change and its taking effect has roughly quadrupled since the middle of the last century, and the unanticipated component of the tax impulse has fallen by about five to one. Over a century, the anticipation share has risen from close to

zero to roughly four fifths. How long the interval is depends on the instrument chosen, the point in the year at which the Budget falls, and whether an election is approaching, while crises appear to leave it unchanged.

The distinction between anticipated and unanticipated tax measures, and the tools for handling foresight, come from the literature. The contribution here is to treat anticipation as a property of the policy regime, measured as a time series, shown to be systematically distributed, and dated to identifiable reforms. The series needed to do so were already in a public replication archive, where they had served as a robustness check.

For applied work, a narrative tax series pooled across the post-war period is not a series of comparable objects, and the announcement-dated and implementation-dated series built here allow the two to be estimated apart. For policy, the timing of a tax change has economic content, yet at present it is largely a by-product of instrument choice and Budget scheduling, with some political influence at the margin.

Notice is of use only to households able to act on it, which takes savings, flexible income, or paid advice, so a large increase in notice is a large increase in a benefit that is unevenly distributed by construction. Whether the unevenness matters for household behaviour cannot be tested in this dataset, and the simplest version of the claim fails on what can be tested. The gap in notice between regressive and progressive instruments was significant between 1980 and 1999 and is not significant today. The raw difference persists, with 43 per cent of regressive measures arriving on under a month's warning against 33 per cent of progressive ones. The companion paper takes up who receives the notice.

Britain reformed the timing of its tax policy for reasons of parliamentary scrutiny and legal certainty, and in doing so transformed the anticipation content of fiscal policy, a consequence that none of the documents making the case for the reforms mentions.

References

- Alesina, A., Favero, C. and Giavazzi, F. (2015), 'The output effect of fiscal consolidation plans', *Journal of International Economics* **96**, S19–S42.
- Cloyne, J. (2013), 'Discretionary tax changes and the macroeconomy: new narrative evidence from the United Kingdom', *American Economic Review* **103**(4), 1507–1528.
- Cloyne, J., Dimsdale, N. and Postel-Vinay, N. (2023), 'Taxes and growth: new narrative evidence from interwar Britain', *Review of Economic Studies* .
- Cloyne, J., Hürtgen, P. and Dimsdale, N. (2025), 'Are tax cuts contractionary at the zero lower bound? Evidence from a century of data', *Journal of Political Economy* **133**(2), 568–603.
- Favero, C. and Giavazzi, F. (2012), 'Measuring tax multipliers: the narrative method in fiscal VARs', *American Economic Journal: Economic Policy* **4**(2), 69–94.
- Hürtgen, P., Cloyne, J., Dimsdale, N. and Postel-Vinay, N. (2024), Tax changes in the United Kingdom 2009–2020: a new narrative account and dataset, Technical report, SSRN Working Paper 4909051.

- Leeper, E. M., Walker, T. B. and Yang, S.-C. S. (2013), ‘Fiscal foresight and information flows’, *Econometrica* **81**(3), 1115–1145.
- Mertens, K. and Ravn, M. O. (2011), ‘Understanding the aggregate effects of anticipated and unanticipated tax policy shocks’, *Review of Economic Dynamics* **14**(1), 27–54.
- Mertens, K. and Ravn, M. O. (2012), ‘Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks’, *American Economic Journal: Economic Policy* **4**(2), 145–181.
- Ramey, V. A. (2011), ‘Identifying government spending shocks: it’s all in the timing’, *Quarterly Journal of Economics* **126**(1), 1–50.
- Romer, C. D. and Romer, D. H. (2010), ‘The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks’, *American Economic Review* **100**(3), 763–801.
- Yang, S.-C. S. (2005), ‘Quantifying tax effects under policy foresight’, *Journal of Monetary Economics* **52**(8), 1557–1568.