Improved productivity measurement in New Zealand's Longitudinal Business Database

Richard Fabling and David C. Maré

Motu Working Paper 19-03

Motu Economic and Public Policy Research

April 2019

Document information

Author contact details

Richard Fabling
Independent Researcher
richard.fabling@xtra.co.nz

David C. Maré
Motu Economic and Public Policy Research
dave.mare@motu.org.nz

Acknowledgements

We gratefully acknowledge funding from Productivity Hub agencies – Ministry of Business, Innovation & Employment; Productivity Commission; The Treasury; and Stats NZ. Thanks to Lynda Sanderson (MBIE), Sandra Watson (IRD) and Sri Farley (IRD) for comments on the paper.

Disclaimer

The results in this paper are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI), managed by Stats NZ. The opinions, findings, recommendations, and conclusions expressed in this paper are those of the author(s), not Stats NZ, other Productivity Hub agencies, or Motu. Access to the anonymised data used in this study was provided by Stats NZ under the security and confidentiality provisions of the Statistics Act 1975. Only people authorised by the Statistics Act 1975 are allowed to see data about a particular person, household, business, or organisation, and the results in this paper have been confidentialised to protect these groups from identification and to keep their data safe. Careful consideration has been given to the privacy, security, and confidentiality issues associated with using administrative and survey data in the IDI. Further detail can be found in the Privacy impact assessment for the Integrated Data Infrastructure available from www.stats.govt.nz. The results are based in part on tax data supplied by Inland Revenue to Stats NZ under the Tax Administration Act 1994. This tax data must be used only for statistical purposes, and no individual information may be published or disclosed in any other form, or provided to Inland Revenue for administrative or regulatory purposes. Any person who has had access to the unit record data has certified that they have been shown, have read, and have understood section 81 of the Tax Administration Act 1994, which relates to secrecy. Any discussion of data limitations or weaknesses is in the context of using the IDI for statistical purposes, and is not related to the data's ability to support Inland Revenue's core operational requirements.

Motu Economic and Public Policy Research

PO Box 24390	info@motu.org.nz	+64 4 9394250
Wellington	www.motu.org.nz	
New Zealand		

 ISSN 1176-2667 (Print), ISSN 1177-9047 (Online).

Abstract

Accounts information that businesses supply to Inland Revenue for tax purposes provide over 96% of the observations in the productivity dataset in the Longitudinal Business Database. In 2013, material changes in the data collected halted the annual updating of the productivity dataset. This paper describes a method for accounting for these raw data discontinuities, and revisits the prior productivity dataset methodology, implementing wholesale changes that improve the overall quality of the data and the versatility of the productivity dataset.

JEL codes

D20; D24

Keywords

Stats NZ Longitudinal Business Database; production function; multifactor productivity; administrative data

Summary haiku

Tax data changes

Experience tree bears fruit

Measurement improves

1 Motivation

Major changes in the underlying business tax collection from the 2013 (March) tax year have prevented the simple updating of the micro productivity dataset (Fabling 2016).[1] Since tax data account for over 96 percent of observations and almost two thirds of (unweighted) total labour input, it is important that the changes in data collection are adequately accounted for. In this paper, we outline modifications to the productivity dataset methodology that enable old and new IR10 tax form data to be used together, extending the consistent data coverage to 16 consecutive tax years (2001-2016), with approximately 200,000 firm-level observations per year.

Subject to stable and timely input data updates in the future, these methodology changes also re-enable annual updating of the productivity database – in line with the remainder of the Longitudinal Business Database (LBD) – adding value to other data sources, particularly the Business Operations Survey which was streamlined in 2009 on the assumption that financial data would be sourced primarily from linked administrative tax data in the future (Fabling and Sanderson 2016).

In addition to accounting for tax data changes, we take the opportunity to improve the existing method (described in Fabling and Maré 2015b) based on our experience with using the data, and to enable a broader set of empirical research questions to be addressed in a way that is internally consistent with the assumptions made in the productivity dataset production. The primary innovation is to replace the Stats NZ tax data cleaning processes with new processes that are more consistent with good microeconomic research practices. These new processes apply quality criteria based on relative, rather than absolute, error size and minimise the risk associated with making arbitrary choices as to where errors have been made in internally inconsistent tax returns. The new processing also identifies and repairs common fixable filing errors that previously resulted in observations being dropped from the dataset. In conjunction with these changes, we introduce a new suite of data flags that identify which processing steps have been applied to each observation allowing researchers to investigate how important these transformations are in their particular analysis sample.

In terms of expanding the usefulness of the micro productivity dataset,

[1] Fabling (2016) completed an (unpublished) initial investigation of the tax data continuity issues, which is available from the corresponding author on request. Our paper and associated methodology changes draw heavily on that initial issues assessment, which was funded by Stats NZ as part of their investment in the Longitudinal Business Database.

we: add population weights based on firm size (total labour input), entrant/exiter status and industry; test for the internal consistency of tax-based wage and salary data, improving the internal consistency of value-added and profit measures; and add additional capital stock variables, enabling alternative measures of K in production functions and a deeper understanding of the composition of capital including intangible assets.

Overall, these methodological changes result in greater exclusion of low quality observations, but a net increase in the annual number of observations in all but the first three years of data. This dual outcome of better and more data is achieved by removing previous data exclusions that had no direct implications for the quality of productivity components (eg, balance sheet assets equalling liabilities plus equity), and the effect of changing the data cleaning process. Dispersion in multifactor productivity (MFP) across firms is reduced, partly reflecting the removal of previously high dispersion observations. This outcome may relate to the previous cleaning processes being more forgiving of small businesses, which also tend to have greater variability in productivity (Fabling and Sanderson 2014).

To avoid repetition and lack of clarity around what has changed, this paper focusses on the new methodological developments and – where applicable – the processes that they replace, rather than providing a self-contained description of the entire productivity dataset method. The paper should be read in conjunction with the original methodology summarised in Fabling (2011) (permanent firm identifiers), Fabling and Maré (2015a) (measurement of labour, including working proprietors) and Fabling and Maré (2015b) (productivity components, population of interest, and estimation issues). To understand the broader data environment the productivity dataset sits within, particularly the opportunities for linking to other firm-level data, see also Fabling and Sanderson (2016) and references within that paper.[2]

Methodological changes are split into three sections – general methodological improvements based on experience using the data (sections 2 & 4), and methods specifically addressing the change in the administrative tax form (section 3). Section 5 summarises the impact of these changes and provides examples of how new variables in the dataset might be used, while section 6 wraps up and briefly outlines outstanding issues.

[2] Fabling and Sanderson (2016) also discuss Stats NZ access rules for business tax data-based research, which have recently been relaxed to improve access for university academics, subject to meeting standard microdata access requirements.

2 General methodological improvements

This section covers changes to the productivity methodology based on experience with using the data, particularly in relation to the identification of low quality data and a desire to make the productivity dataset more capable of addressing other research questions (eg, in relation to profitability). Two changes come from identified errors in the prior method: failure to remove gains and losses on sales of fixed assets; and incorrect adjustment of working proprietor counts on entry and exit.

Most of the changes discussed in this section were inspired by re-examining the original method in light of the change in tax form and associated guides (Fabling 2016). For example, investigation of how to account for the absence of a GST-inclusive indicator on the new tax form led to an understanding that misreporting of GST status was occurring under the old IR10 tax form. Reading of the associated tax guides and their changes, led to an understanding that intermediate consumption as measured might include direct labour costs for some firms. Attempting to modify the Stats NZ cleaning rules to account for the new data source led to a complete first principles revision of the edit rules from a microeconomic perspective.

Sequentially, the productivity code now does the following steps, which are labelled to reflect whether the processing has changed ("OLD" for old (original) method; "MOD" for steps that existed in the original method, but that have since been modified; and "NEW" for steps that did not exist previously):

OLD Identify population ("measured sector," private-for-profit, $L > 0$)

NEW Edit IR10s to correct obvious filing errors

MOD Apply quality threshold to determine usable IR10s

NEW Exclude non-itemised IR10s (totals, but not components, reported)

OLD Determine data source(s) – AES+IR10, AES-only or IR10-only

NEW Remove gains/losses on sale of fixed assets from Y/M (IR10)

NEW Remove direct labour costs from M (IR10)

NEW Drop IR10 observations where W&S is inconsistent with PAYE (EMS)

MOD Remove GST from GST-inclusive returns (IR10)

NEW Harmonise IR10 (form changes from tax to accounting variables)

MOD Calculate active years from L, GST and AES/IR10

OLD Use $t-1$ active status & AES/IR10 to construct lagged K

MOD Use $t+1$ active status to adjust WP labour input (final year only now)

NEW Use industry aggregate K shares to create capital goods deflator

MOD Make industry-year adjustment for components not in IR10

MOD Model rental, leasing & rates expenses for AES (move from M to K)

MOD Remove firms with implausibly large changes in Y, M or K

NEW Calculate population weights, stratifying on industry and firm size

Each of these modified and new processes is discussed sequentially in the following subsections. In addition to these methodology changes, the output dataset now includes a number of additional variables, which are itemised in the table codebooks at the end of the paper (appendix B).

2.1 IR10 edits and quality restrictions

After identifying the productivity population – based on industry (Stats NZ "measured sector"), private-for-profit status and positive labour input (L) – the first step in the productivity code is to identify what financial data are available to construct productivity components (gross output Y, intermediate consumption M, and capital K). For AES-based data, this part of the process remains largely unchanged,[3] and the focus of this subsection is to motivate and explain the complete reworking of the IR10 processing. Representative copies of the IR10 old and new form are included in appendix A for reference.

[3]The primary change to AES is to now drop a small number of individual returns where these cause firms to fail the implausibly large change in components rule ($|\Delta \ln(\text{component})| > 4$). Dropping individual AES returns at this stage avoids losing the entire firm from the productivity dataset, allows available IR10 data to be used instead, and is done only in cases where there is a single obvious outlier AES return. Other minor changes affecting the AES data are: changing the prioritisation of AES and IR10, so that AES are always prioritised (previously rare multi-ENT cases could prioritise IR10s when more IR10s were available than AES returns); dropping AES observations where closing asset book value components are negative (previously negatives were set to zero and the AES return was retained); rounding data to be consistent with the form (AES to nearest $1,000; IR10 to nearest $1 – probably inconsistent with form due to Stats NZ manipulations, and previously left unaltered).

Under the original methodology, IR10s were cleaned by Stats NZ using rules applied in the production of national accounts. Those rules checked for internal inconsistencies within tax returns – eg, totals not matching summed components – and then adjusted components to make returns internally consistent, provided the identified inconsistencies were within certain fixed tolerance limits. If all manipulations required to create an internally consistent return were within tolerance limits then the return was assigned a "pass," in which case the return – along with its adjusted components – was assumed to be adequate for productivity estimation. Returns that did not pass these tests held either a designation of "fail" (where tolerance limits were exceeded) or "zero-punched" where either or both return pages had no completed fields.[4]

The work to adapt the IR10 cleaning methodology to account for the new IR10 tax form led to a first-principles examination of the appropriateness of the existing IR10 data cleaning for microeconomic analysis, finding the method wanting on a number of dimensions. Of particular concern were the (Fabling 2016):

- Use of absolute dollar value thresholds for passing returns, excluding large firms with small *relative* errors and including small firms with large *relative* errors
- Choice to lump adjustments into "other" categories, arbitrarily affecting which productivity component gets adjusted
- Inclusion of "other information" fields in the zero-punched rule, passing returns with zero balance sheet components
- Inclusion of irrelevant (to productivity measurement) consistency tests, excluding usable data
- Failure to account for common respondent errors before cleaning (eg, non-reporting of totals), excluding usable data
- Lack of tests for partial (non-itemised) returns, overestimating M
- Failure to flag edits on passed returns, restricting users' ability to assess data quality
- Sequential editing and testing of returns, limiting the usefulness of "reasons for failure" data flags and resulting in unexpected outcomes

[4] In the data received from Inland Revenue, item non-response and zero responses are both denoted by the value zero, hence the terminology "zero-punched" meaning all fields on a page are observed as zero.

(eg, zero-punched returns being incorrectly misclassified as completed)

- Not allowing for rounding error, leading to incorrect adjustment of returns (eg, double-counting gains/losses on sales of fixed assets)

Since the original intention of these edit processes was to feed supplemental data into National Accounts for small business, the identified issues may not have a significant impact on derived aggregates. This is because AES postal returns are used for large firms, and because imputation and weighting account for missing data. However, for microeconomic research purposes, official-statistics-based processes fall well short of best practice, having been adopted largely as an expedient technology when the LBD was first developed (Fabling 2016).[5]

Because of the identified issues, we replace the entire IR10 processing system with a fit-for-purpose alternative. We also apply a principle of flagging all significant manipulations of the data in the final productivity dataset and note the related flag in the main text and in the table codebooks.[6] First, we repair what we believe are three common filing errors – negatives in fields that cannot be negative, summed income components being wrongly reported in "other income," and inconsistency in the gross profit calculation.

For illegitimate negative fields, we simply take absolute values, relying on subsequent quality checks to determine whether this is appropriate. This approach was also taken in the old processing system, and is forced on us for consistency because IR imposes this on post-2012 (new form) data before it is supplied to Stats NZ. That is, post-2012 there are no *observed* illegitimate negative values in the raw LBD IR10 data, though they may continue to exist in the filed data.[7]

For "other income," we check that this hasn't been incorrectly reported as a summation of other components (total income, gross profit or income sub-components following the gross profit calculation). In cases where this appears to be the case we simply set "other income" to zero (edit indicated

[5] At that time, it was also believed that the LBD might serve a dual purpose as both a research dataset and a dataset that Stats NZ might use for official statistical purposes. In that latter setting, there is some benefit in retaining a consistent data treatment with National Accounts.

[6] All tests of internal consistency allow for rounding error up to $5 so that edit flags do not inaccurately overestimate the level of inconsistency in the data.

[7] This manipulation of the data is not flagged, because it happens post-2012 but cannot be observed in that period. Since we only retain observations where the sign of the variable appears to have been incorrectly reported, observing where these manipulations have occurred may be unnecessary for understanding final data quality.

by *flag_i10_othinc_edited*).[8]

For gross profit, it is allowed (in the pre-2013 IR10 guide instructions) for this to have been derived outside of the IR10 from an accounting cost-of-goods sold calculation. In such a situation, derived gross profit (ie, from summing and subtracting components in the IR10) and reported gross profit (ie, as recorded in box 6) may legitimately differ, since reported purchases may differ from cost-of-goods sold (COGS) as reported in financial accounts. In particular, any difference between derived and reported gross profit may reflect the presence of direct labour costs in the COGS calculation, and not reflected in the purchases variable. Thus, the derivation method for gross profit affects the consistency of the definition of the purchases variable between firms and, potentially, over time for the same firm.

If the purchases variable includes additional COGS-related components, those expenses will not be reported in the expense side of the profit-and-loss statement, potentially affecting the allocation of expenses to M. In particular, as discussed later in the paper, the inclusion of direct labour costs – a component of COGS – in the purchases variable presents a serious measurement issue for intermediate consumption, since wages and salaries (W&S) should not be included in M. We address the contamination of purchases data with W&S by referencing independently reported wage and salary (W&S) costs from the Employer Monthly Schedule (EMS). But, in order to make this W&S adjustment consistently across firms, we must first harmonise the purchases variable so that it reflects a COGS calculation for all firms. We do this by assuming that any inconsistency in the gross profit calculation is due to the external derivation of gross profit from a COGS calculation – that is, we set purchases to the value necessary to make the gross profit calculation internally consistent (edit indicated by *flag_i10_purch_edited*), conditional on gross profit being non-zero and the resulting purchases being non-negative.

Figure 1 shows the proportion of IR10-based observations in the final productivity dataset that have been edited. In addition to the clear downward trend in the necessity for editing, there is a further downward jump in the editing rate going from 2012 to 2013. This latter drop is likely due to change in the IR10 guide that required firms to report purchases as in a COGS calculation, and at the same time made it clearer to respondents that COGS-based "purchases" should include direct labour costs. Both these requirements appear to have reduced the level of internal inconsistency in the

[8] This issue does not appear to be common for "other expenses," perhaps because the expense itemisation in the IR10 does not require any intermediate summation, such as that generated by gross profit calculation.

gross profit calculation. Overall, the edit rate drops from around 13 percent in 2001 to less than 2 percent in 2016. Under the old processing system, these returns would have been dropped if the edit exceeded the tolerance threshold for inconsistency in that section of the form. Consequently, the new pre-quality editing results in more data being retained for productivity estimation.

Screening of IR10s to remove low quality data relies on two steps – the identification of incomplete returns, and an assessment of whether returns are internally consistent within a certain *relative* tolerance limit. In the old methodology, incomplete returns were judged based solely on whether the return was zero-punched on either or both pages. We extend this approach in two important ways – firstly, we also require that asset data is supplied on the back page to avoid using returns where the only back page information provided is non-balance sheet "other information."[9] This latter requirement is important for the new (post-2013) IR10 form since the tax depreciation variable moves to the back page other information section increasing the probability that responses with no balance sheet might report non-zero "other information," invalidating the intent of the zero-punched back page test.

Table 1 shows the prevalence of zero-punched data (including missing asset data) as a proportion of the IR10s that might be usable for productivity measurement – around 10 percent of forms are lost in each year and this is predominantly because of the absence of a balance sheet, which may reflect the number of sole proprietors in the dataset for whom the business balance sheet is not separate from the individual assets, or it may reflect an expectation amongst respondents that the profit and loss statement is a sufficient data supply for IR to assess tax liabilities.

The other form of non-completion that we identify could be dubbed "half-hearted compliance," where sufficient data are supplied to derive taxable profit but income and/or expenditure is reported only as a total or is reported entirely under the relevant "other" category minimising the amount of box-filling required. We assume that it is unlikely that firms have only other income and other expenses, implying that non-itemisation inhibits our ability to convincingly separate income and expenses into appropriate productivity components. For that reason, we discard these returns, which affects between 1.4 and 6.3 percent of observations (table 1, column 3). The rate of non-itemisation appears to fall following the form change, which may partly reflect the switch to accounting, rather than tax, variables making full

[9]Formally, we require that some component of assets is positive, or that positive liabilities equals negative equity so that total assets is zero if balance sheet equality holds.

compliance easier for filers.[10] Under the old processing system, unitemised returns were likely to pass quality tests if other categories and totals were reported, and to fail tests if only totals were reported and those exceeded relevant thresholds for inconsistency.

Table 2 shows how *raw* productivity components are calculated from IR10s for both the old and new form. These components are raw in the sense that they will be subsequently adjusted in a number of ways before becoming usable for productivity analysis. In particular, at this stage, M_{raw} includes rental, leasing & rates expenses (RLR) which will ultimately form part of K. RLR appears in intermediate consumption at this stage for two reasons – firstly, it is easier for quality checks to be run independently for profit and loss and balance sheet; and, secondly, not all firms report RLR expenses separately so that some of the transfer of this component to K will come from a modelling step later in the process. Put another way, some IR10-based firms (and all AES-based firms) cannot have this component shifted to K from the outset and it is, therefore, more practical to retain a consistent measure of M that includes RLR up to the point where RLR can be removed from all observations of M. At this point, and for similar reasons, K_{raw} also exclude the (tax) depreciation component (since this may need to be modelled from accounting depreciation under the new form).

All raw productivity variables are derived by summing the indicated components in the first two columns of table 2 but, since the dataset includes totals and intermediate calculations, simple summation is not the only way to derive the raw productivity components. Indeed, under the old processing system, corrections to IR10 forms involved adjusting other income/expense categories to match totals, which is mathematically equivalent to deriving productivity components by taking stated totals and deducting excluded subcomponents (since other income/expenses are each included in their respective productivity component). The final column of table 2 lists these alternative approaches and, in an internally consistent IR10, the approaches yield identical estimates of the raw productivity component.[11]

A key innovation in the quality assessment method is to test whether the data allow us to be indifferent to how each productivity component is

[10] In the first four years of data, where unitemised rates are the highest, the issue is largely associated with reporting of expenses (84% of cases), which are perhaps more likely to be harder to itemise if accounting expense categories do not align directly to the categories required by the IR10.

[11] The new IR10 form does not require the separate reporting of total fixed assets, reducing the scope for testing the internal consistency of the fixed asset schedule after 2012. We maintain a consistent set of quality tests over time.

constructed, since we don't know which (if any) approach is superior when these aggregations yield different results. We define a sufficient level of indifference as the maximum raw in component calculations being at most one percent of the average of the minimum and maximum value, ie:

$$\text{diff}_{max} = \max\{2 \times \frac{\max(x_i) - \min(x_i)}{\min(x_i) + \max(x_i)}, x \in Y_{raw}, M_{raw}, K_{raw}\}, \tag{1}$$

where i indexes over the methods for calculating a productivity component.

Table 1 shows the impact of the one percent relative difference cut-off on data loss. Returns that fall outside this cut-off are labelled inconsistent, while those that are internally inconsistent – beyond rounding error – but are sufficiently close to exact are labelled inexact (indicated by *flag_i10_fp_inexact* and/or *flag_i10_bp_inexact*). As the table shows, this suite of tests results in most internally inconsistent returns being discarded. On average, 5% of observations are dropped because of inconsistency, with a further 1.1% retained as sufficiently accurate. Again, there is a very clear increase in data quality when the new IR10 form is introduced in 2013. Overall, the new data cleaning processes retain 82% of IR10 forms (final two column of table), though that average rate is over 5pp higher post-2012, mainly due to increases in the internal consistency of the data.

Figure 2 summarises how this relative quality cut-off affects the quality of retained observations, compared to the prior methodology. We use diff_{max} as the relevant quality metric and restrict attention to correctly itemised returns, so that we don't take account of the quality improvements arising from correctly specifying the zero-punched rule, excluding non-itemised returns and the initial editing of the data for common filing errors. Unlike in table 1, we also restrict the analysis to $L > 0$ firms to aid comparison to the final data coverage of the prior methodology. The solid and dashed black lines in figure 2 plot the cumulative distribution of diff_{max} for raw observations from old and new IR10 forms respectively.

As table 1 demonstrated, data quality is significantly worse prior to 2013, with the cumulative distribution of raw data quality for old form observations (solid black line) significantly below the corresponding cumulative distribution of new form observations (dashed line). As a consequence of these quality differences, the cut-off at 1% of diff_{max} (represented by black dots) removes significantly more old form observations than new form observations.

For comparison with the old methodology, the solid grey line shows the cumulative quality distribution restricting the old form observations to the subset ultimately included in the prior productivity dataset. Under the old methodology, the Stats NZ quality checks and other data dropping steps, remove a substantial proportion of low quality IR10 forms (ie, the grey line sits above the solid black line). However, since the cut-off does not apply, the old methodology retains a sizeable sub-population where the choice of aggregation method has an important effect on the (relative) value of the productivity components. Overall, 3.7% of retained old productivity dataset observations exceed the new (1% diff_{max}) quality threshold. One percent of observations in the old productivity dataset have a $\text{diff}_{max} > 0.68$ implying that, for at least one component, one method of derivation yields a component value at least double what would be calculated by an alternative derivation method for that component.

Figure 3 demonstrates the permissiveness of the old methodology by plotting the proportion of raw IR10 observations at a given quality level that are included in the old productivity dataset. As hoped, this line is downward sloping indicating that the old methodology is less likely to include low quality – relative to high quality – observations. However, a significant proportion of returns at lower quality levels are still included in the old productivity dataset. For example, at $\text{diff}_{max} = 0.68$ the old method includes more than half of the available observations in the final dataset. Figure 3 implies, therefore, that a key reason why the old processing rules don't yield worse average quality than they do is due to the relative scarcity of poor quality raw data, rather than that the old quality checks do a particularly good job of screening out low quality observations. Consequently, data quality in the old productivity dataset varies over time with the underlying variation in raw data quality. In contrast, the new productivity dataset varies in data coverage rather than data quality (table 1).

The new quality method has at least two potential shortcomings. Firstly, because we cannot tell the difference between reported zeros and missing fields in the IR10 data, we do not apply the related consistency checks when necessary totals are zero/missing (indicated by *flag_i10_total_missing*). In such cases data are not exposed to the full suite of tests. If remaining tests are passed, the absence of totals does not affect derived raw productivity components because of the choice to sum from subcomponents. In practice this is equivalent to assuming that these totals are missing and actually equal to the summed subcomponents.

Secondly, the new testing approach restricts the consistency checks to elements that directly affect productivity component measurement. Unlike the prior methodology, it does not require the firm balance sheet to balance, that is, for assets to equal liabilities plus equity. Inconsistency that doesn't affect productivity measurement could still indicate poorer quality due to accounting systems, understanding or care on the part of the respondent. However, even if true, this inferior quality hasn't manifested directly in productivity data inconsistencies, which we treat as the relevant metric in the analysis. Furthermore, some firms appear to have problems with balance sheet identities that may be unrelated to productivity component quality – for example, sole proprietors who may be able to identify specific assets used in the production process, but who may not have a formal business balance sheet separate from personal assets. As we will show later, added observations do not increase multifactor productivity dispersion, suggesting that the gains in sample size from relaxing this test does not come at the cost of data quality.

2.2 Gains and losses on sales of fixed assets

In the previous version of the productivity code we assumed that gains and losses on the sale of fixed assets (gains/losses) were not reported in other income and other expenses, respectively. Under that assumption, no adjustment was made to those components before they were included in Y and M respectively. Inspection of the Stats NZ processing showed that, in instance where other income/expenses were not reported by respondents at levels that accommodated gains/losses (as we'd assumed), those gains/losses were added to other income/expenses. As a result the productivity code was, in fact, counting gains and losses on the sale of fixed assets as part of gross output and/or intermediate consumption. In principle, these components should be excluded since they do not reflect part of the production process of the firm. Making matters worse, the sequential editing in the Stats NZ processes coupled with rounding error, sometimes meant that adjustments resulted in double-counting of gains/losses.[12]

While removing the Stats NZ processing step fixes some of this issue, the tax guides and data suggest that gains and losses are being reported in other income/expenses by at least some firms, and that we should attempt to deduct it. To remove these non-productivity components, we deduct reported gains/losses from other income/expenses where this doesn't result in negative "other" categories. We make one exception to this adjustment – in the case where the reported loss is equal to reported depreciation, we assume that

[12]Presumably, AES/National Accounts processes then deduct this back out from Y and M, at least in the instances where double-counting hasn't happened.

the respondent is confused, and make no adjustment (since we don't believe other expenses is contaminated with asset sales).

The adjustment process is complicated by the introduction of the new IR10 form, since gains/losses should then – according to the tax guide – be reported in the tax adjustment box.[13] Therefore, post-2012, the adjustment for the net difference between gains and losses occurs to the tax adjustment variable if that is non-zero, and to other income/expenses if the tax adjustment variable is zero (and other income/expenses category can accommodate the adjustment). This expanded process necessitates three data flags for each of gains and losses, capturing when reported gains/losses have been removed from the tax adjustment variable (*flag_i10_gain_adj_taxadj*, *flag_i10_loss_adj_taxadj*) or other income/expenses (*flag_i10_gain_adj_othinc*, *flag_i10_loss_adj_othexp*), or haven't been removed because other income/expenses cannot accommodate the adjustment (*flag_gain_unadj*, *flag_loss_unadj*).

Table 3 shows the proportion of IR10 observations that fall into each of these categories, as well as the mean adjustment to the relevant productivity component, conditional on adjustment. The star notation added to component variables in this table – and subsequent tables – denote that the adjustment is measured as a proportion of the component value as at this processing step (and prior to adjustment). The first thing to observe is that losses are far more prevalent than gains.[14]

Under the old IR10 form, the majority of gains on sales of fixed assets are not adjusted (because reported gains exceed other income), while almost all losses are adjusted. With the introduction of the new form, the overall rate of adjustment remains similar for losses, though the prioritisation of the tax adjustment over other expense means that a large proportion of the adjustment shift to the tax adjustment variable. Because there is subsequently an overall adjustment to M based on the tax adjustment term, the overall impact of changing which component is adjusted may be minimal, depending on the sign and magnitude of the (net) tax adjustment term.

[13] Gains/losses could also be reported in exceptional items if assets were sold as part of the selling off of part (or all) of the business as a going concern. For consistency of treatment between the old and new form data, we ignore this possibility, since going concern sales should also result in non-reporting of gains/losses in the profit and loss statement under the old IR10 form. That is, we consistently assume that sales of going concerns is an ignorable event in the context of asset sales.

[14] Among other things, this may reflect systematic differences between the rate at which tax depreciation is allowed on assets, and the corresponding market value of assets at that vintage.

For gains, the ability to remove this component from the tax adjustment term increases the overall rate of adjustment, meaning that the majority of gains are now adjusted for, and the primary mechanism for adjustment is via the tax adjustment term. Overall, the effect on other income results in fewer adjustments to Y, since any subsequent tax adjustment correction is only made to M (see section 3).

Whether the overall level of adjustment made for gains and losses of the sale of fixed assets is optimal is hard to judge. Conditional on adjustment, the change to Y^* and M^* is certainly material (an average 3.7% and 4.1% respectively), so the question is an important one when we consider that a significant proportion of firms have M adjusted for reported losses. For gains, it is unclear whether the change in adjustment rates is due to the permissiveness of an extra adjustment variable, or increased compliance in reporting, particularly since the tax adjustment variable is a net term making it's value uninformative as to whether the gain has been reported.

For losses, it seems possible that in some instances the other expense category permits the adjustment only because respondents may lump expenses into this category,[15] rather than because the loss was reported in other expenses. However, respondents also have a clear incentive to net these losses off income, since losses as reported on the back page of the IR10 are supposed to be *tax-deductible*. That logic implies that it is appropriate to adjust M for losses and that we should expect the rate of non-adjustment on the loss side to be very close to zero, which is what we observe.

2.3 Direct labour costs and consistency of labour cost measures

Through clarification of the preferred COGS calculation and an expressed desire for the gross profit calculation to be internally consistent, the new IR10 guide makes clear that direct labour costs (among other things) should be included in the purchases variable (Fabling 2016). As figure 1 shows, the change in form and instructions appear to have reduced the rate of inconsistency in the gross profit calculation, and we have taken effort to adjust the purchases variable to ensure that consistency across all IR10 responses (both old and new form). While this process harmonizes the purchases variable over time, it also presents a fundamental challenge to accurately measuring inter-

[15] Bearing in mind that we dropped non-itemised returns, which reflect the most extreme example of this filing behaviour.

mediate consumption, since the inclusion of any labour costs in M results in a (cost-weighted) misallocation and double-counting of L in the production function. This issue was previously unrecognised, meaning that the misallocation is present for the subset of firms reporting direct labour costs in the purchases variable, and has only been exacerbated by the introduction of the new tax form and guide.

To some extent, the labour cost contamination risk exists for two further expense categories that contribute to M. For R&D expenses, it becomes an issue if firms allocate labour costs (which is the predominant cost in aggregate R&D) to the R&D expenses category, and for other expenses, it becomes an issue for employing firms if W&S are not separately itemised. The previous version of the productivity code did nothing to remove this contamination from M.

For employing firms, we correct these issues by using an alternative measure of the total wage bill – total gross earnings from the Employer Monthly Schedule (EMS), labelled W_{EMS}. For firms with working proprietors (WPs), we have no equivalent independent data with which to estimate the extent to which related party remuneration has been included in expense categories other than related party remuneration. Furthermore, because we cannot be completely confident that all WPs have been identified and removed from the EMS data by prior processing (Fabling and Maré 2015a), and because pre-2013 IR10 forms do not distinguish related party remuneration separately from W&S, we combine IR10 W&S and related party remuneration together for comparison with EMS gross earnings (labelling this combined cost W_{IR10}).[16]

Because W_{EMS} should have most working proprietor labour income removed, either through processing or because WP remuneration usually occurs outside of the PAYE system, and because the W_{IR10} should reflect a total remuneration measure, we expect $W_{EMS} \leq W_{IR10}$.[17] Thus, where we observe $W_{EMS} > W_{IR10}$, we attribute this discrepancy to the reporting of

[16] The related party remuneration variable was added to the IR10 in 2013, but was renamed associated persons' remuneration in 2016. We treat these two new form variables as equivalent. By combining related party REM and W&S data and by treating the old and new form as equivalent, we are assuming that working proprietor labour income (if any) was reported as W&S prior to 2013. In addition, because the new form W&S variable is total remuneration, we add fringe benefit tax to total labour costs prior to 2013, for consistency.

[17] W_{IR10} as reported on the IR10 form is included in the final productivity dataset – combined with it's AES analogue – as *WS_RelPartyRem_nom*, which enables comparison with W_{EMS} (stored in the labour table as *total_gross_earn*).

labour costs elsewhere, reducing relevant expense categories until the two earnings measures are reconciled. We adjust expense categories in a prioritised order based on the likelihood that labour costs have been reported there: firstly purchases, then R&D, and then finally other expenses, which is only adjusted if $W_{IR10} = 0$ (ie, if it looks like W&S has not been itemised). Each of these adjustments is identified by a flag in the final dataset, respectively, *flag_i10_WS_adj_purch*, *flag_i10_WS_adj_rd*, and *flag_i10_WS_adj_othexp*.

Despite the fact that both EMS and IR10 data are tax-based, the sources may be inconsistent because of timing (eg, incorrect attribution to financial year) and definitional differences. In general, we expect these definitional differences to results in under-adjustment of non-W&S expense categories because of the inclusion of WP earnings and additional total REM components (eg, fringe benefit tax). However, in some cases W_{EMS} exceeds the combined total of all potentially W&S including expense categories. To enable users to be confident that the labour and productivity data are consistent for analyses that examine workers and firms simultaneously (including the ability to use W&S data to create a plausible profit metric), we discard observations where W_{EMS} exceeds the total of IR10 expense categories where W&S might be reported by more than 5%.[18]

Table 4 shows the proportion of employing firm IR10 data that is dropped or adjusted. Overall, we drop 3.4% of observations for consistency reasons, and adjust a further 18% of observations (on average). The rate of adjustment is somewhat lower from 2014 onwards though, when adjustments are made, they represent a significantly larger average proportion of M^* under the new form (10%) compared to the old form (7%). The final column of table 4 shows the adjustment as a proportion of W_{EMS}, which also increases substantially post-2012. Together, these increases, coupled with the increased internal consistency of the COGS calculation in the IR10, are consistent with the IR10 guide changes having the desired effect of increasing the reporting of direct labour costs in purchases and, therefore, an increasing need for the productivity code to address the contamination of M with labour costs.

[18] That is, we drop observations where $W_{EMS} > 1.05 \times (W_{IR10} +$ purchases + R&D), or $W_{EMS} > 1.05 \times$ (purchases + R&D + other expenses) when $W_{IR10} = 0$. AES data do not require a W&S adjustment, but there are similar inconsistencies between W_{EMS} and W_{AES}. We do not drop AES observations, since the inconsistency does not affect the adjustment of components, though it will impact on the plausibility of combining EMS gross earnings values with AES financial data in, eg, a profit calculation.

2.4 GST reporting and adjustment method

Under the old IR10 guide, respondents were encouraged to report their accounts on a GST-exclusive basis, though it was acknowledged that (mainly small) businesses may not have such accounts and that GST-inclusive reporting was acceptable in such situations. To identify which type of filing was provided, a GST-exclusive indicator question appears on the back page of the form (see appendix A), and the old Stats NZ processing took this response into account, making a simple deduction from all income and expense categories where GST is expected to be reported. Under the new tax form guide, GST-inclusive reporting is discouraged and the indicator question on the form has been removed, though some firms may continue to report on a GST-inclusive basis. While trying to establish a method for identifying GST-inclusive reporting under the new form, three issues were identified with the existing adjustment method (Fabling 2016):

- Ignoring the net GST term that makes taxable profit correct, overestimating Y or M depending on whether the firm made a net payment or refund of GST (ie, introducing a bias correlated with profit)
- Assuming that the GST rate is constant across firms, ignoring exporting (exports are zero-rated, so Y is systematically biased downwards for adjusted exporters), and balance date effects when the GST rate changes (systematically affecting industries, eg, agriculture, with non-March balance dates)[19]
- Assuming that respondents accurately report GST-inclusiveness, and so removing non-existent GST components (underestimating Y & M)

Figure 4 illustrates the last of these issues by plotting the (log) ratio of IR10 "GST-equivalent" sales (ie, sales+rent+other income) to GST-exclusive sales from GST returns, for those firms that reported their IR10 return was GST-inclusive.[20] Ignoring consistency issues between the two measures – which might arise from, eg, the apportionment of GST returns to financial year and differences in the scope of the two sales measures – an IR10 return that includes GST should show a ratio equal to the current GST rate,[21] while

[19] Stats NZ processing assumes the GST rate in the transition year is the average of the before (12.5%) and after (15%) rate in the transition year, which reflects the fact that there are six months at each rate in the 2011 *March* balance date year (since the GST rate changed on 1st October 2010).

[20] These results exclude firms with zero-rated GST to rule out the possibility that the spike at zero is due to export-only firms, rather than incorrect filing.

[21] Technically, the density should be clustered around $\ln(1 + r_{\mathrm{GST}})$, which has a value of

a return that does not include GST should show a (log) ratio of zero. Since we have included only the GST-inclusive reporters, there should be a single spike in the data at the relevant GST rate.

Ignoring returns where the two sales measures are very different from each other – that is, the spikes of pooled observations at either end of the distribution – the data suggest that respondents are largely accurate in their reporting of GST-inclusiveness in 2001 & 2002 (top left panel of figure 4).[22] However, there is a small spike at zero suggesting that the IR10 return is actually GST-exclusive. Over time two things happen – firstly the total number of GST-inclusive reporting firms declines and, secondly, the proportion of inaccurate declarations increases. By 2009-2010, it is more likely that reporting is wrong than right (bottom left panel of figure 4).

Unfortunately, because the distribution of sales ratios for correct and incorrect GST-status filers appear to overlap substantially in those later years, we cannot rely of the GST ratio to delineate correct and incorrect reporting. Instead, we rely on the consistency of firm reporting and the timing of that reporting – ie, is it in a period when they are statistically more likely to be accurate – to determine at a firm-level whether we trust reported GST-inclusiveness.[23] Specifically, to treat *all* GST-inclusive reported returns of a firm as GST-inclusive we require either that they always report that their returns are GST-inclusive or that

$$\sum_{t=1999}^{2012} \frac{2}{11}[2012 - t] \times \delta_t(\text{inclusive}) \geq 3.9,$$

which gives a "consistency" weight of two to a GST-inclusive observation in 2001 and a weight of zero to an observation in 2012.[24] The cut-off of

0.118 when the GST rate is 12.5%, and 0.140 when the GST rate is 15%.

[22] This test is too blunt to assess the accuracy of GST-exclusive reporting because natural variation in the ratio for accurate reporters (centred on zero) swamps any spike we might observe from a small number of inaccurate reporters. Put another way, the method tells us that the proportion of inaccurate GST-exclusive filers is not large enough to identify it from the GST sales ratio. Fabling (2016) uses the natural variation in the GST rate in 2011 to show that the misreporting of GST status happens for a small proportion of GST-exclusive reporting firms. We treat this as ignorable given the limitations in detecting affected returns.

[23] Fabling (2016) shows that consistently reported inclusive filing over time is more likely to be accurate, using the sales ratio method to judge accuracy, which is consistent with the behaviour being driven by the non-availability of GST-exclusive financial accounts.

[24] We use all non-zero punched IR10s to calculate the consistency of filing, which includes returns already dropped for quality reasons, and returns from non-employing years and 1999 (where employment data is unavailable).

3.9 is chosen so that a firm that files as GST-inclusive in both 2001 and 2002 achieves the necessary weight, having reported GST-inclusiveness in two time periods where it was likely that this reporting was correct (top left panel of figure 4). While all GST-inclusive reported returns for firms that achieve the cut-off weight are treated as GST-inclusive, returns reported as GST-exclusive will *always* be treated as GST-exclusive. Raw and final GST-inclusive status are identified by *flag_i10_gst_incl_raw* and *flag_i10_gst_incl* respectively in the final dataset.

Table 5 shows the proportion of raw GST-inclusive reporting firms, and the proportion that we change from inclusive to exclusive. As the weighting rule is designed to achieve, the proportion of incorrect filing increases over time, roughly mimicking the patterns in figure 4. To verify that the actual allocation of firms to GST-inclusive/exclusive is generally an improvement over the stated GST treatment, figure 5 shows the sales ratio based on this final allocation (top two panels respectively). While there is still some misallocation present, returns assigned to GST-exclusive generally appear to be so (concentrated at a ratio of zero), while returns that retain the GST-inclusive designation generally appear to cluster at the GST rate.

Where GST adjustment is necessary, Y and M are both decreased using a GST rate derived directly from year-specific GST filing, separately for income and expenses. Making the adjustment based on firm-level GST filing solves the problem of overestimating GST rates on sales for exporters, and the lack of firm-specific GST rates in the 2011 tax year caused by the change in GST rate.

The final adjustment made to the GST methodology accounts for the reporting of net GST in the profit and loss statement. Regardless of actual GST treatment in the IR10, taxable profit must be GST-exclusive. To achieve this, a firm reporting GST-inclusive accounts must report the net GST payment (refund) as other expenses (income). We make this additional adjustment, dropping observations where the other income or expense category is insufficiently large to accommodate the adjustment, which may occur because we've made a prior over-adjustment (eg, removed gains/losses of sales of fixed assets when they weren't reported), the return is still incorrectly classified (ie, it is actually GST-exclusive), or because the respondent has not reported the necessary adjustment or has reported it in an incorrect income/expense category.

The bottom panel of figure 5 shows the distribution of the sales ratio for dropped observations. Overall, we drop 16% of returns that we have treated as GST-inclusive, and this proportion (grey line) is higher at sales ratios where the classification appears to still be incorrect, suggesting that the dropping procedure provides a second line of defence against including overadjusted returns in the final dataset.

The new IR10 form doesn't have an indicator of GST-inclusive reporting on which to base an adjustment. Since we lack an appropriate technology for identifying these returns, we assume all returns are GST-exclusive under the new form. Using the sales ratio approach, we estimate that approximately half of previously GST-inclusive filers switch to GST-exclusive when the new guide and form comes into effect. On that basis, a ballpark figure for the rate of GST-inclusive filing in 2013 is around 1.4% of returns as at this processing step (assuming the dropping procedure is identifying incorrectly classified returns). Users may need to exercise caution when examining productivity growth rates for (treated as) GST-inclusive filers that span the form change, as some of these will exhibit growth in Y and M due to the discontinuation of the GST adjustment process.

3 Accounting for IR10 form changes

Aside from expense category changes, the main change to the profit and loss component of the IR10 is the switch from reporting tax variables to accounting variables – see appendix A for Fabling's (2016) full summary of the changes. Tax and accounting versions of variables may differ because of, eg, tax rules around what constitutes a deductable expense. Since taxable profit is still the ultimate target variable for reporting, the aggregate difference between tax and accounting variables is reflected in a "tax adjustment" variable which can then be deducted from net profit/loss to yield taxable profit.

An important example of this change is depreciation expenses, since Inland Revenue sets out specific depreciation rates for asset classes that may differ from those used by an accountant, and we rely on IR's specification of economic depreciation rates to enable us to directly include depreciation as a major component of K. Fortunately, both tax and accounting depreciation expenses are collected on the new IR10 form, enabling a specific methodology to harmonise depreciation expenses over time and to test for the consistency of the two measures. Following the specific treatment for depreciation, we make a general adjustment to M using the tax adjustment variable. The methodology in this section has no prior comparator, since the change in IR10 form halted the annual updating of the productivity dataset. In addition, Stats NZ processes treat the old and new form data as equivalent.

3.1 Tax depreciation

In the new IR10 form, because profit and loss variables are now accounting-based, tax depreciation moves to the "other information" section of the back page of the form (see appendix A), with the front page variable now labelled "accounting depreciation and amortisation." Tax and accounting depreciation may differ in terms of applicable rates as well as scope (eg, the amortisation costs of some intangible assets are not tax-deductible).

While both tax and accounting depreciation should be reported on the IR10 form, the movement of the tax depreciation variable to the back page of the return has resulting in a substantial increase in the proportion of observed zeros in the tax depreciation variable, reflecting increased non-response (Fabling 2016). Because of this we cannot use just the tax depreciation (δ_{TAX}) variable in isolation and must rely on accounting depreciation (δ_{ACC}) where tax data are missing. The source for tax depreciation is determined by the following prioritisation:

1. δ_{TAX}, where tax is reported (ie, $\delta_{\text{TAX}} > 0$)
2. δ_{TAX}, where $\delta_{\text{TAX}} = \delta_{\text{ACC}} = 0$
3. δ_{ACC}, where there is another year where both depreciation measures are non-zero and within 1% of each other
4. δ_{ACC}, where there is a year where $\delta_{\text{ACC}} > 0$, $\delta_{\text{TAX}} = 0$ and the tax adjustment is zero (implying $\delta_{\text{ACC}} = \delta_{\text{TAX}}$ and δ_{TAX} is unreported)
5. $\text{TFA} \times \bar{d}/(1-\bar{d})$, where the closing book value of total fixed assets is non-zero ($\text{TFA} > 0$) and $\bar{d}$ is the firm-level mean of $\delta_{\text{TAX}}/(\delta_{\text{TAX}} + \text{TFA})$ for years where $\delta_{\text{TAX}} > 0$ & $\text{TFA} > 0$
6. δ_{ACC}

Table 6 show the proportion of tax depreciation observations derived from each of the methods, where the two tests of equivalence (step 3 and 4) are combined into a single category ("$\delta_{\text{ACC}} > 0$ & $\delta_{\text{ACC}} = \delta_{\text{TAX}}$ in another year"). If we accept those equivalence assumptions are correctly repairing non-response, imputation is extremely limited ($\sim$0.2% of observations, indicated in the dataset by *flag_i10_dep_imp_TFA* and *flag_i10_dep_imp_acc*).

Figure 6 compares the distribution of the approximate depreciation rate (d, as defined above) in the four years prior to the new form (using reported tax depreciation), with the four years of new form data (using derived tax depreciation). The two distributions look very similar, suggesting that, not only are the rules for deriving tax depreciation post-2012 adequate, but that fixed assets values themselves have not been radically altered by the shift to accounting values. This is a natural consequence of the majority of firms having accounting depreciation rates consistent with tax depreciation rates – an approach which minimises compliance costs associated with tax filing.

3.2 Other differences between tax and accounting variables

To improve the consistency of tax and accounting variables across form vintages, we adjust M using the tax adjustment variable, taking into account components of that adjustment that have already been accounted for – ie, any difference between tax and accounting depreciation and any reported gains or losses from the sale of fixed assets. In making this adjustment, we assume that the tax adjustment variable applies only to intermediate consumption-related components of expenses, rather than non-M expenses or any component of income. Because these assumptions are quite blunt, and may become less tenable as the relative size of the adjustment increases, we restrict the scope of the tax adjustment to be no more than 5% of M^* (averaged over it's before and after adjustment value), dropping observations when that threshold is exceeded.

Table 7 shows the proportion of firms that are dropped or that have M adjusted. Overall, we drop 5.2% of IR10 observations by excluding relatively large tax adjustments. A further 16.9% of observations are adjusted, with a significant bias towards the tax adjustment variable being positive and, therefore, accounting expenses tending to be larger than tax expenses (after depreciation and gain/loss effects). Consequently, most adjustments to M are negative, averaging around 1% of M^*. Flags in the final dataset identify observations adjusted at this step, including whether the tax adjustment term is positive or negative (*flag_i10_taxadj_pos* and *flag_i10_taxadj_neg*).

Table 8 shows the proportion of IR10-only firms in the final productivity dataset that are affected by each adjustment step, additionally breaking each of these adjustments down by the share of associated L (right panel), and by old versus new form (bottom of table). The table also reports the rate of imputed/modelled rental, leasing and rates (RLR), which is discussed in the following section. Comparison of the left- and right-hand panels of table 8 show that, aside from the GST and RLR adjustment, these data adjustments are generally associated with larger firms consistent with such firms having more complex balance sheets. A major improvement to the productivity dataset is the inclusion of the edit flags, enabling summary statistics like table 8 to be easily produced, and for subsets of the data to be examined for quality issues.

4 Other methodology improvements

This section explains additional methodological improvements that either tweak previously existing systems, or add new abilities to the data collection. The main effects are to: create a productivity population dataset, enabling the construction of population weights in the productivity dataset; harmonise the adjustment of WP labour input to match the new processes run on the IDI for WPs in the labour dataset; and to allow users to construct alternative measures of K from within the productivity data.

To create a productivity population dataset, we extend the derivation of firm entry and exit to all firms in the productivity population and combine this data with the permanent industry table on the LBD to identify the subset of private-for-profit firms in productivity industries. As in the previous instance of the productivity code, firm entry and exit is a year-to-year concept relying on activity indicators in adjacent years for firms who have employment in the current year. Adjacent activity is defined as the presence of L, Y or M in the year. Since some firms do not appear in the productivity data in all years they are active, we also use GST sales and purchases to identify the likely presence of Y and M. We expect GST to be full coverage due to the mandatory nature of GST filing. The processing at this step changes from using Stats NZ's Business Activity Indicator series, which is no longer supported by Stats NZ, to a direct method of allocating raw GST returns to financial years implemented within the productivity code suite. The resulting apportioned GST table is included in the productivity suite at both the monthly and annual level (see appendix B).

The productivity population dataset also includes labour variables, since the WP count may be modified to account for firm exit. In previous instances of the productivity code, all WP counts were halved in years adjacent to firm entry/exit on the assumption that WPs were unlikely to work full years in such firms. The latest instance of the labour dataset methodology pushes this logic down to the individual level – ie, adjusting for entry/exit of WPs from firms – which is likely to produce more accurate WP counts as well as solving issues of over-adjustment related to prior method.[25,26] Because WP identification in the IDI relies on annual tax filings for individuals, identification of potential WP exit in the last year of productivity data is sometimes not possible (affected WPs counts are identified by *WP_unknown_trans* in the LBD labour table).[27] For these observations, we use firm exit to determine whether an adjustment should be made, indicating this transformation by renaming the variable as *wp_adj* in productivity tables.

Additionally, we now retain productivity observations with zero component dollar values (previously dropped), so that weights in the dataset simply account for missing data (not zeros). Retained observations with zeros in one or more components may also be useful for, eg, calculating: industry-level aggregates; labour productivity; value-added production functions; and profit metrics.

We now include nominal values for variables so that users can choose to apply different deflators and/or manipulate the data prior to deflating (eg, create a nominal profit measure). In addition, as part of the improvements to the scope of the capital stock data, we now create an industry-specific deflator for K, where we previously used Stats NZ's all-goods deflator in each industry. Recognising that industries can have very different asset composition, we construct each deflator by aggregating the (real) value across all years in each asset class for each industry, using a representative asset deflator.[28] The industry deflator for K is then the sum of the representative asset deflators weighted by that asset classes share of total industry fixed assets.[29]

[25] Specifically, under the old methodology, a WP who switches from one firm to another during a year has 0.5FTE allocated to each firm, since they are a WP in two firms during the year. If, say, the first firm also closed that year, then the half FTE at that firm would be halved again by the productivity code.

[26] The change in approach also fixes an error in the old code, which resulted in incorrect WP adjustment for some firms. All comparisons in this paper between old and new productivity data use fixed WP counts for observations where incorrect adjustments were previously made.

[27] This identification issue also exists in the 2000 financial year (for potential WP entrants), but that year is not used in the productivity dataset.

[28] The "representative" asset class deflators used are: transport equipment (for vehicles); plant, machinery & equipment (for plant & machinery plus other fixed assets); furniture (for furniture & fittings); and non-residential buildings (for land & buildings).

[29] This approach has an underlying assumption that the depreciation and rental, lease &

These deflators – together with the official Stats NZ deflators for inputs and outputs – are stored in a new table on the LBD (see appendix B).

The inclusion of nominal values for K and the components of K – including average book values of each fixed asset type – allow the end user to construct alternative deflators, or alternative capital stock measures as illustrated in the following section.[30] Finally, to further aid that flexibility, we add reported nominal intangible assets to the productivity table, taking care to do this in a manner consistent with the fixed asset data.

The method for deriving the rental, leasing and rates (RLR) component of K for AES-based returns (and for IR10 returns where it is missing/zero) is altered, reflecting the discovery that the constructed "AES other expenses" variable does not satisfy the necessary internal consistency properties for it to be effective. Specifically, the value of "AES other expenses" sometimes exceeds $M^\star$, meaning it isn't a subcomponent of $M^\star$ comparable to the IR10 equivalent, which itself is now always internally consistent due to quality improvements. The new method we adopt is simpler and better – we start by applying AES-based industry aggregate adjustments to IR10s (previously done after RLR imputation) and then use the now-consistent (across AES & IR10) measure of $M^\star$ as the denominator in the RLR modelling exercise.[31]

The final step we complete before deriving population weights is to drop firms with extreme changes in any of the (non-L) productivity components. Because we now include observations with productivity components equal to zero, we must modify the rules in this step, which previously depended on calculating one-year log-changes.[32] In changing the method, we also recognise that by focussing on year-on-year changes to assess potential data issues, we failed to adequately test the quality of data for intermittently reporting firms. Those firms may have worse quality data overall, particularly if their intermittent data is due to failing quality checks in some years and not others, or repeated exit/entry.

Instead of a year-on-year test, we now look at consecutive non-zero

rates components of K have the same within-industry distribution across asset classes.

[30] The depreciation component of K can be backed out from the included components, remembering that the cost of capital component is 10% of the ABV of total fixed assets.

[31] Fabling and Maré (2015b) explain the purpose and detail of the AES aggregate adjustments and RLR modelling. The string variable used in the previous method to capture RLR and IR10 book value treatment has been separated out into indicator variables – *flag_i10_nolagCBV*, *flag_RLR_avg*, *flag_RLR_model* – to improve utility and consistency with the other new flags in the productivity dataset.

[32] Recall, we also added an earlier step for AES returns that dropped outliers so that the number of AES observations dropped at this step is lower.

data observations for a firm (ie, ignoring the time between observations), still applying a threshold for dropping firms as a log change greater than four. To recognise that small firms may reasonably achieve very large log changes, we do not drop observations if the absolute change is below a value equal to the distance between the 50th and 25th percentile of the relevant component variable (scaled by the number of years between observations).[33]

The methodologies used to estimate multifactor productivity – unweighted OLS & firm fixed effect gross output Cobb-Douglas, and OLS gross output translog – remain unchanged, though we add a table to the collection reporting estimated coefficients for reference. In particular, these estimates do not use population weights as that would require additional assumptions about the correct weighting for the fixed effects estimates (see below for a particular choice).

5 Outcomes and applications

5.1 Productivity data coverage

Final coverage rates for the productivity dataset are reported in table 9 as a proportion of firms and as a proportion of total employment. These proportions are relative to the full population of employing enterprises, meaning that there are three overarching reasons for incomplete coverage. Firstly, since we include all enterprises, an average of 17% (27.5%) of observations (employment) is not in the private-for-profit measured sector, with the bulk of employment not covered being in the public sector. Secondly, in-scope firms may have no AES/IR10 data,[34] making it impossible to include them in the dataset without full imputation of productivity components. Missing raw data is more of an issue for smaller firms, because AES has higher coverage for large firms so that, on average, 19.8% of observations are dropped because of no data, but only 8.6% of total employment is dropped. Finally, the processing steps we implement to clean the data and remove low quality observations result in a further 9.0% of observations being dropped. Because

[33]In the prior methodology, these thresholds were fixed dollar values and the same for all three components. The prior method also tested for large changes in L, which we dispense with because the 25th-50th threshold approach is not satisfactory for L and because labour inputs are less susceptible to the issues that cause implausible component changes.

[34]The table note defines what we mean by available data. In particular, zero-punched IR10s are not counted as available data.

most of the dropping of data is related to IR10 processing, the proportion of employment dropped is lower, averaging 5.9% of the total.

The coverage of the dataset improves over time due to both an increase in the availability of raw data (except in 2016, which may be due to late filing), and an increase in raw data quality (ie, a declining drop rate). As a proportion of the in-scope population, the productivity dataset now covers, on average, 65% of firms and 80% of total employment. Table 10 breaks productivity population coverage down by employer status – ie, employer or WP-only firm. The overall average coverage rate is around 11pp higher for employing firms than WP-only firms. This gap is driven by differences in the availability of raw data, probably due to the additional (non-IR10) tax filing options available to sole-proprietors and partnerships.[35]

5.2 Comparison between old and new productivity data

We compare the new methodology productivity dataset to the old methodology dataset in two ways. Firstly, we compare the standard deviation of MFP in the two datasets – on average and then decomposed by whether observations are dropped, gained or retained when the method changes. MFP is estimated from an unweighted gross output Cobb Douglas production function with firm fixed effects, where MFP is the combination of the estimated fixed effect (δ_i) and residual (ϵ_{it}).[36] Secondly, we examine the autocorrelation of MFP and productivity components over time.

Table 11 reports the standard deviation of MFP by year, noting that the old productivity dataset extends only to 2012 because of the IR10 form change. In all years the standard deviation of MFP is lower in the new dataset, compared to the old dataset, driven by declines in the variability in the estimated firm fixed effect. Since the bulk of firm-year observations remain the same each year, we interpret the decline in the standard deviation as an improvement in the signal-to-noise ratio in the MFP data. Table 12 decomposes this change into components due to gains and losses of firms from the sample, and to changes in estimated MFP for retained firms (between the two data vintages). Comparison of the first two columns of table 12 indicates

[35] In some instances – eg, for farm or rental income declarations – the original (IR3F/R) tax return may be transcribed to an IR10 by Inland Revenue. However, these transcribed IR10s don't have balance sheets because that data is not collected on the IR3F/R form.

[36] In the new productivity dataset, this MFP estimate is labelled *mfp_go_fe* with corresponding fixed effect *go_fe*. The old productivity dataset has the same variables, though we use re-estimated MFP after correcting for the error identified in the WP count.

that the contribution of common sample firms to the observed reduction in MFP dispersion is minimal (ie, values in the two columns are similar).

The next two columns show the number of gained and lost observations, where two features are relevant to the decomposition. First, the number of gained and lost observations is significant relative to the respective datasets they belong too – gained observations on average represent 13.8% of the new productivity data, while 11.5% of observations are lost from the old dataset. Second, except in the first three years where the new quality checks are more often binding (see table 9), the new methodology yields a net gain in observations totalling over 60,000 over the full twelve year period. Because the aggregate change in composition is material, quality differences between dropped and gained observations affect the overall standard deviation of MFP. The last two panels of data in table 12 align with the view that the data process changes have improved the average quality of the data by weeding out low quality IR10 observations. Specifically, the standard deviation of MFP for lost firms is much higher than the standard deviation of both retained and gained firms, with the higher variation due to both greater variance in the fixed effect and the residual component. Conversely, gained observations are more similar to retained firms indicating that the increasing data coverage has not come at much of a quality cost.

Figure 7 & 8 show the autocorrelation of MFP and individual productivity components, respectively, with solid lines for the new productivity data (to 2016) and dashed lines for the old productivity data (to 2012). Comparing old and new data, two key differences stand out – firstly, the new data show a lower correlation in MFP over time, which is not reflected in lower correlations over time in the individual component series. This likely relates to the high dispersion in the fixed effects of the dropped observations, since larger fixed effects (relative to residuals) will tend to increase the observed autocorrelation in MFP. Secondly, the autocorrelation of L is significantly higher in the new data, which may reflect composition changes also (dropping more small firms with high variation in L) as well as the change in method to an individual-based method of adjusting WP labour input.

Focussing on the new data, the autocorrelation in MFP drops significantly in 2013 before recovering (figure 7), suggesting that the harmonisation process between old and new form has not been completely successful. Intermediate consumption and capital are subject to harmonisation, with the former bearing most of the adjustments and appearing to be mainly responsible for the observed loss of continuity (figure 8, top right panel).

5.3 Population weights and aggregates

The primary goal we set for the population weights in the productivity dataset is that they aggregate to productivity industry-level annual firm counts in the productivity population. That goal can be achieved simply by stratifying the data on industry and year and applying constant within-cell weights equal to the inverse of the coverage rate within the cell.

As a secondary goal, we target replicating aggregate L using the population weights. We choose L as the target because we have complete population data for L (ie, we can verify how successful we've been), and because accurately estimating one productivity component is likely to achieve approximately correct aggregates in other components. Entering and exiting firms are known to display different productivity dynamics from incumbent firms, and we also apply special data treatments to entrants and exiters. For those reasons, we also place entrants and exiting firms in distinct cells. Similarly, concerns around estimation of the labour input and other components for WP-only firms (Fabling and Sanderson 2014) suggest it is prudent to put WP-only firms in a separate size category, leaving open the option of excluding them from an analysis without a need to recalculate weights.

To satisfy the goal of matching aggregate L, we start by examining the coverage rate of the productivity data by firm size. Figures 9 & 10 do this for employer and WP-only firms respectively. Figure 9 shows the distribution of employers in the population and in the productivity dataset by data source (top and middle panels respectively), and the coverage rate by firm size (bottom panel, solid line) together with the proportion of productivity data from AES (hollow line), and the firm entry and exit rate (dashed and dotted lines respectively). Figure 10 compresses most of this information into a single graph, since WP labour input is lumpy (by construction).

The inclusion of AES data ensures that the coverage rate is almost 90% for employers with $L \geq 100$. Coverage is lowest for micro-enterprises – both for employing firms and WP-only firms, which have coverage rates between 50% and 70% for firms where $L \leq 2$. For WP-only firms, the higher coverage rate for $L = 2$ firms reflects our expectation that partnerships may be more likely to file IR10s than sole-proprietors, with the relatively low coverage rate for the latter group negatively impacting the overall coverage of WP-only firms (as shown in table 10).

The skewedness of the employer firm size distribution, coupled with the increasing coverage rate by firm size, mean that weighted aggregate L will overestimate true L if we use constant weights for firms in the same industry-year. To avoid this, we stratify employers into seven firm size groupings as shown in table 13. Since entry and exit are largely associated with employing firms of $L \leq 2$ (figure 9, bottom panel) and for WP-only firms, we stratify on entry/exit dynamics only for those two firm size groups. A small group of firms are "one-year" firms – ie, both entrants and exiters – and we group these with exiters because these two groups have similarly low productivity data coverage.

The final column of table 13 shows the average coverage rate within firm size cells, reflecting the coverage rates reported in figures 9 & 10. The inverse of industry-firm-size-year-specific analogues of these coverage rates are the productivity population weight included in the productivity dataset as *pop_weight*. From the final column of table 13 it is clear that alternative tax filing options are not the only reason for lower coverage rates for micro enterprises. For both $L \leq 2$ employers and WP-only firms, exiters are a substantial proportion of the population (over 8%), and have coverage rates below 30%. Undercoverage for these firms may reflect lower data quality in the exit year; preference for non-IR10 tax filing; non-compliance with filing requirement; timing differences that incorrectly identify the exit year (eg, payment in arrears to employees); and/or a failure to identify enterprise number continuity for micro enterprises.[37]

Figures 11–13 show three simple applications of the data weights. Figure 11 provides a test of the ability of the weighted data to replicate aggregate industry employment. Pooling all years, we plot the unweighted productivity data aggregate L (x-axis) against the weighted productivity data aggregate L (y-axis), each as a proportion of true industry aggregate L (bubble area scaled to true industry size). With the exception of one industry – the uppermost observation (telecommunications, JJ12) – weighting the productivity data moves the estimated aggregate closer to its true value. In most cases the resulting weighted L is within two percent of the true value (represented by the horizontal dashed grey line), with a bias towards overestimation.[38] Forestry & logging (AA21) is the main exception to this general case, where weighting moves the estimate from 58% to only 85%, which is still a substantial improvement.

[37] In relation to the last of these issues, the PENT technology does not repair enterprise number continuity for micro enterprises because it is based on Stats NZ repairs to plant identifiers for plants with at least three employees (Fabling 2011).

[38] The dotted gray line represents the 45° equivalence line where unweighted and weighted L are equal.

Figure 12 shows the effect of using the weights to aggregate other productivity components, producing an index of labour productivity (LP) that can be compared to Stats NZ official statistics. In all three panels, solid black lines are official series, solid and dashed grey lines reflect alternative ways of combining productivity-population weighted value-added with L (true population L or weighted L), and grey dotted lines show unweighted productivity dataset aggregates (all series indexed to 100 in 2001).

The unweighted labour productivity calculation deviates substantially from the other series, showing much slower LP growth over the full 16-year period. This effect is caused by an increase over time in the coverage rate of relatively low productivity firms, resulting in an overestimate of aggregate growth in labour relative to output (comparing trajectories of the dotted lines in the middle and bottom panels, relative to the solid black line). The weighted productivity dataset accounts for changing coverage rates of small (lower productivity) firms, doing a better job of matching aggregate LP growth over the full 16-year period (top panel).

Looking at the detailed time series properties, the weighted and official LP series track each other reasonably well until 2008 and then deviate from each other markedly in 2009, with the official series showing a far milder downturn in labour productivity than the weighted productivity data. Looking at the middle and bottom panels of figure 12 these differences do not arise from the labour series (due, say, to the FTE measure not tracking hours adjustment well), but rather are due to differences in aggregate value-added growth post-GFC.

While this divergence deserves further investigation, it may be impossible to reconcile using the LBD. Re-estimating value-added using raw AES data in the LBD – including imputed- and tax-based data with AES dataset weights – yields a very similar series to the productivity population-weighted productivity data. Since the weighted AES data forms the starting point for National Accounts estimates, the consistency of the AES- and productivity-weighted data seems to imply that the deviation of the official and weighted productivity data value-added series is due to macro adjustments subsequently applied to the AES data, which the LBD microdata cannot shed light on. Despite this, population weights on the productivity dataset help narrow down the plausible set of reasons why micro and macro methods reach different conclusions about aggregate productivity growth.

Figure 13 shows how commonly-used measures of productivity dispersion are affected by weighting the data. The top two panels of figure 13 show the 10th, 25th, 75th & 90th percentile of the MFP distribution normalised so that the 50th percentile is zero in each year to aid visualisation of changes in productivity dispersion. The left panel is unweighted and the right panel uses productivity population weights. The bottom two panels use the same data, comparing the interquartile range (ie, the 75th-25th percentile value) and the 90th less 10th percentile. For both metrics, the weighted measures of dispersion (solid lines) is higher than the unweighted metrics (dashed lines) because small firm – particularly WP-only firm – productivity dispersion is higher than large firm dispersion (see, eg, Fabling and Sanderson 2014), and weights are higher for small firms. Both measures show similar trends though with productivity dispersion stable or slightly increasing until 2010, before steadily declining.

5.4 Alternative measures of capital

While theoretically well-founded, the measure of K in the dataset is non-standard, particularly because we include the RLR contribution from non-owned assets. Additionally, because RLR is not directly observed in the AES return and is missing (appears as zero) in some IR10 response, a population-weighted average of 13% of RLR observations are modelled (table 14, final column) using the methodology outlined in Fabling and Maré (2015b). A further 15% is imputed using firm-level data from other years. In aggregate, this means that over 11% of K is either imputed or modelled, with RLR in total constituting around 30% of aggregate K (table 15). A further 29% of K is aggregate depreciation costs, with the (more traditional) cost-of-capital component making up the remaining 41% of K.

Because users may want to take a standard approach to the capital stock, we now include the raw data necessary to enable that kind of estimation, as well as including the reported stock of intangible assets for augmented production function estimation. Figure 14 demonstrates the increasing relevance of intangibles in the total capital stock, by reporting the (weighted) aggregate real average book value of intangibles (weighted and deflated with the capital deflator), intangibles share of productive capital (solid line, RHS axis), and the proportion of firms with intangible assets (dashed line). The value of intangibles grows rapidly over the 2004-2007 period, peaking at over \$30 billion (2016 dollars) and 15% of the total stock of productive capital.

Table 16 reports a set of alternative specifications for K in a gross output Cobb-Douglas production function (estimated with firm fixed effects and year dummies). Unlike the estimation approach used to generate MFP in the productivity dataset, we pool all industries (for illustrative purposes).

Aside from that, columns (1) and (2) are the standard specification of K – ie, including all three components of capital services – on an unweighted and population-weighted basis respectively. For weighted estimates, we average the firm weight across all years to satisfy the condition that firms have constant weights for fixed effect estimation. Column (1) is the benchmark for comparison between measures of MFP (fixed effect plus residual), and the second-from-bottom panel of table 16 reports (population-weighted) correlations between this MFP and each alternative specification, alongside summary dispersion statistics.[39] Both the correlations and the estimated coefficients in column (2) suggest that the decision whether to estimate weighted or unweighted has little impact on MFP, at least when industries are pooled.

Columns (3)-(5) of table 16 illustrate alternative specifications of the capital component, all unweighted and estimated on the same sample of firms.[40] Columns (3) replaces K with average total fixed assets. To be consistent with this approach, we add RLR costs back into M, which unwinds any imputation or modelling that may have been done to separately identify the RLR component. Depreciation costs are excluded in this "standard" approach. While returns to scale remain similar comparing this specification to that in column (1), the capital coefficient now seems implausibly low.

Columns (4) and (5) revert to the preferred measure of K and add intangibles to the production function, either as a separate production factor or together with K. In the latter case, to be consistent with the methodology for deriving K, we add $K_{\text{Intang}} = [0.1\times$ Intangibles$]$ to K representing the cost-of-capital component.[41] Column (4) indicates a positive and significant correlation between gross output and intangible usage, after controlling for other inputs, and we expect this relationship to vary in strength across industries. The choice of whether to include intangibles as a separate input or together with K will depend on the researchers' preferences and research question, though both methods yield similar MFP estimates to the base case of excluding intangibles.

[39] All summary statistics are calculated after industry-year demeaning the data for consistency with other MFP-based statistics reported in the paper.

[40] To retain a constant sample, columns (3) and (4) include indicator variables for when TFA or intangibles are zero, respectively, and the logs of those variables are set to zero in those cases (rather than missing). An alternative specification of column (3) where we exclude firms with zero TFA (4% of obs) yields very similar results to those reported. Estimating the same variant for column (4) is not recommended, since selection effects become important (only around a quarter of the sample having intangibles, figure 14)

[41] Amortization of some intangibles may already be included in the depreciation component of K, and the RLR component may already include licensing fees and royalties paid for non-owned intellectual property.

5.5 Profit and mark-up

Excluding IR10-based observations where EMS gross earnings and IR10 total remuneration disagree increases the legitimacy of deriving profit from the productivity dataset. To illustrate this application, we construct a measure of mark-up as

$$\mu = \frac{Y}{M + \hat{W} + K} - 1, \tag{2}$$

where $\hat{W}$ is total gross earnings from the EMS, including imputed earnings for working proprietors. Note, it important to use the raw EMS in this calculation, because the IR10 total REM variable on the productivity dataset is the raw variable – ie, unadjusted for reallocation of labour costs from other expense components – and may inconsistently include WP remuneration (with or without a return on capital component).

Imputing WP earnings improves consistency across firms and over time, but comes at the cost of having to make modelling assumptions. For employers, imputed WP earnings are estimated from N(WP)×[total employee earnings/total employee FTE]. For WP-only firms, imputed earnings are the summed industry-year analogue of the firm-level calculation. Table 17 reports unweighted medians of μ by year and employer status. The first column shows firms where FTE>0, and the second column shows WP-only firms applying the industry-year imputation for $\hat{W}$. For employing firms, the median mark-up seems plausible, as does its decline and recovery following the Global Financial Crisis (GFC). The imputation of WP-only earnings, though, seems completely unsatisfactory and is consistent with the possibility that the adjusted WP count overestimates true WP labour input.[42] The final column of table 17 shows the median mark-up for WP-only firms if we take an alternative – extreme – assumption that $\hat{W} = 0$ for WPs. These numbers are also implausible, but with the opposite sign, suggesting there may be a sensible middle ground that could be uncovered by a better imputation methodology.

Restricting attention to firms with employees, figure 15 plots the distribution of mark-ups using population weights (grey area) and weighted by population-weights and real costs (solid line) for all years pooled.[43] This

[42] It is also consistent with the literature on self-employment suggesting labour earnings are lower for WPs compared to employees with similar characteristics.

[43] Real costs are the denominator in the mark-up equation, ie, $M + \hat{W} + K$, where we deflate $\hat{W}$ using the input PPI (M deflator).

distribution looks plausible, particularly the fact that the real-cost weighted version is more compressed towards zero, reflecting the fact that firms with unusually low mark-ups tend to be relatively small (in cost terms).

6 Conclusions

The change in 2013 to the IR10 tax form has sparked a radical reworking of the productivity dataset methodology, based on rethinking whether the prior system based on existing Stats NZ technologies was fit-for-purpose. The resulting changes in approach have yielded a more complete dataset with higher quality, achieved by removing data exclusions that had no direct implications for quality, and by changing the data cleaning process to reflect best microeconomic research practices. At the same time, the usefulness of the dataset has been improved by adding population weights, more data on the composition of capital (and better deflators for K), detailed flags for tracking edits to the raw data, and from imposing greater consistency between the productivity and labour datasets.

Some issues remain unresolved, notably the reconciliation of productivity aggregates to official statistics post-GFC, the reasons why EMS and IR10/AES earnings data are often in disagreement, and the incomplete harmonisation of M between new and old IR10 form observations (evidenced by the dip and recovery in autocorrelation between 2012 and 2014, figure 7). Increased use and feedback may suggest further issues and improvements as subsets of the data are explored in more detail (eg, industry-specific issues). Subject to ongoing stability in the raw data feeds, though, the current suite of improvements establishes a methodology that should provide researchers with years of future data to aid their understanding of New Zealand firm performance.

References

Fabling, R. (2011). Keeping it together: Tracking firms in New Zealand's Longitudinal Business Database. Working Paper 11-01, Motu Economic and Public Policy Research.

Fabling, R. (2016). Changes to Inland Revenue's financial statements summary: Issues for the Longitudinal Business Database and its users. Mimeo, Statistics New Zealand.

Fabling, R. and D. C. Maré (2015a). Addressing the absence of hours information in linked employer-employee data. Working Paper 15-17, Motu Economic and Public Policy Research.

Fabling, R. and D. C. Maré (2015b). Production function estimation using New Zealand's Longitudinal Business Database. Working Paper 15-15, Motu Economic and Public Policy Research.

Fabling, R. and L. Sanderson (2014). Productivity distributions in New Zealand: The dangers of international comparisons. Working Paper 14-16, Motu Economic and Public Policy Research.

Fabling, R. and L. Sanderson (2016). A rough guide to New Zealand's Longitudinal Business Database (2nd edition). Working Paper 16-03, Motu Economic and Public Policy Research.

Maré, D. C., D. R. Hyslop, and R. Fabling (2017). Firm productivity growth and skill. *New Zealand Economic Papers 51*(3), 302–326.

Tables

Table 1: IR10 loss due to quality restrictions

		Proportion				
		Zero-			Passed	
	N(firms)	punched	Unitemised	Inconsistent	Inexact	Exact
2000	243,540	0.108	0.035	0.083	0.014	0.759
2001	270,987	0.120	0.063	0.075	0.013	0.729
2002	271,812	0.116	0.050	0.074	0.014	0.746
2003	276,213	0.113	0.044	0.065	0.012	0.767
2004	276,351	0.108	0.029	0.058	0.012	0.793
2005	283,836	0.107	0.029	0.056	0.011	0.796
2006	284,895	0.108	0.017	0.053	0.011	0.811
2007	287,490	0.106	0.018	0.050	0.012	0.814
2008	295,392	0.102	0.018	0.055	0.012	0.813
2009	292,329	0.100	0.020	0.056	0.011	0.813
2010	290,295	0.101	0.020	0.054	0.012	0.813
2011	291,681	0.096	0.020	0.054	0.011	0.819
2012	292,503	0.100	0.021	0.054	0.011	0.814
2013	294,531	0.097	0.017	0.018	0.006	0.862
2014	296,766	0.095	0.016	0.015	0.006	0.868
2015	300,099	0.094	0.015	0.032	0.009	0.850
2016	277,605	0.097	0.014	0.012	0.004	0.874
Total	**4,826,325**	**0.104**	**0.026**	**0.050**	**0.011**	**0.810**

Restricted to firms in the productivity population (private-for-profit, productivity industry, $L > 0$ in at least one year between 2001 and 2016). Not restricted to $L > 0$ in current year because prior period total fixed assets is used to derive average K (and may have $L = 0$). Zero-punched returns are where all responses are zero/missing and include returns where the only non-zero items on the back page are "other information" responses. Unitemised returns are where only other or total categories are reported in income and/or expenses. Inconsistent returns have a maximum difference greater than 0.01 in at least one raw productivity components (gross output, intermediate consumption, or total fixed assets). Passed, but inexact returns have a maximum difference in $(0, 0.01]$. Passed exact returns are internally consistent to within rounding error (assumed to be at most \$5).

Table 2: Raw productivity components derived from IR10s, and alternative definitions used in quality checks

Component	Old form (pre-2013)	New form (post-2012)	Alternative derivations (quality checks)
Y_{raw}	sales; Δstock; rental & lease inc; other inc	sales; Δstock; rental, lease and licence inc; other inc	total inc + purchases − non-Y inc; gross profit (GP) + purchases + rent + other; rent + other inc derived from [total inc − GP − non-Y inc]
M_{raw} (+RLR)	purchases; entertainment; insurance; legal; rates; rental & lease payments; repairs & maintenance; R&D; subcontractors; travel & accom.; vehicle; other exp	purchases; insurance; professional & consulting fees; rates; rental, lease & licence payments; repairs & maintenance; R&D; contractors & subcontractors; other exp	non-purchases M exp derived from [total exp − non-M exp]; purchases derived from gross profit; both of above
K_{raw}(TFA)	vehicles; plant & machinery; furniture & fittings; land & buildings; other fixed assets	vehicles; plant & machinery; furniture & fittings; land; buildings; other fixed assets	Total assets − non-K assets

Δstock is closing stock less opening stock. Rental, leasing & rates is included in M_{raw} for consistency tests, because these are expense components and may be missing (zero), but is ultimately moved to be a component of K, along with depreciation costs, later in the data processing. The consistency of both RLR and depreciation costs (tax in old form and accounting in new form) are, therefore, implicitly tested in these quality checks. In the previous instance of the productivity method, fringe benefit tax (FBT) was incorrectly included in M. It is now included in labour compensation – which is a manual addition under the old form, and is directly included in salaries & wages by respondents under a broader (total REM) definition of that variable in the new form. The new form does not have a box for total fixed assets, reducing the number of consistent (over time) alternative derivations of K_{raw}. Under the old productivity method, M_{raw} was derived from total expenses less non-M expenses. For consistency across components, we switch to summation of included components in this version of the method. The quality checks ensure that the impact of this change is minimal.

Table 3: Adjustment for gains & losses on sales of fixed assets

		Proportion with gain				Proportion with loss			
	N(firms)	Not adjusted	Oth inc	Adjusted Tax adj	$\overline{\text{adj}/Y}^{*}$	Not adjusted	Oth exp	Adjusted Tax adj	$\overline{\text{adj}/M}^{*}$
2001	193,950	0.032	0.017		0.036	0.004	0.188		0.046
2002	198,414	0.036	0.018		0.038	0.004	0.189		0.045
2003	204,738	0.036	0.017		0.042	0.005	0.187		0.045
2004	210,183	0.037	0.016		0.042	0.004	0.198		0.045
2005	215,016	0.033	0.013		0.039	0.004	0.196		0.044
2006	218,463	0.032	0.012		0.035	0.004	0.192		0.044
2007	221,169	0.030	0.012		0.036	0.004	0.184		0.043
2008	226,014	0.029	0.013		0.033	0.004	0.181		0.041
2009	222,654	0.023	0.012		0.034	0.003	0.171		0.040
2010	219,972	0.020	0.012		0.036	0.003	0.159		0.036
2011	221,151	0.018	0.012		0.031	0.003	0.160		0.035
2012	220,776	0.019	0.014		0.034	0.003	0.161		0.032
2013	232,362	0.017	0.008	0.012	0.039	0.002	0.110	0.050	0.038
2014	235,914	0.015	0.007	0.013	0.037	0.002	0.109	0.056	0.037
2015	234,354	0.013	0.007	0.013	0.037	0.002	0.105	0.061	0.039
2016	220,869	0.014	0.007	0.014	0.041	0.002	0.107	0.065	0.041
Total	**3,495,999**	**0.025**	**0.012**	**0.013**	**0.037**	**0.003**	**0.161**	**0.058**	**0.041**

Restricted to firms with usable (at this processing step) IR10s and $L > 0$ in the current and/or and adjacent year. The average adjustment, conditional on adjustment, is reported as a proportion of the affected component (Y or M) as at this processing step (denoted by a star) and prior to the adjustment.

Table 4: Removal of direct labour costs from purchases

	N(firms)	Proportion Dropped	Adjusted	$\overline{\text{adj}/M}^{*}$	$\overline{\text{adj}/W}_{\text{EMS}}$
2001	89,778	0.034	0.201	0.067	0.340
2002	91,605	0.036	0.203	0.067	0.326
2003	96,531	0.037	0.195	0.069	0.328
2004	100,263	0.036	0.193	0.069	0.318
2005	103,395	0.036	0.198	0.070	0.315
2006	105,804	0.036	0.198	0.071	0.312
2007	107,700	0.036	0.191	0.070	0.309
2008	110,289	0.039	0.197	0.070	0.300
2009	107,268	0.038	0.196	0.070	0.301
2010	103,779	0.033	0.173	0.072	0.303
2011	103,752	0.032	0.168	0.071	0.295
2012	104,052	0.032	0.167	0.072	0.298
2013	110,466	0.031	0.173	0.094	0.371
2014	113,133	0.028	0.154	0.099	0.383
2015	114,498	0.029	0.150	0.103	0.382
2016	109,851	0.028	0.153	0.108	0.388
Total	**1,672,164**	**0.034**	**0.181**	**0.077**	**0.328**

Restricted to firms with usable (at this processing step) IR10s and FTE employment greater than zero in the current year. IR10s are dropped as inconsistent if total gross earnings from the Employer Monthly Schedule (W_{EMS}) are more than 105% of an upper bound W&S derived from IR10 components. The average adjustment to M, conditional on adjustment, is reported as a proportion of M as at this processing step (denoted by a star) and prior to the adjustment, and as a proportion of W_{EMS}.

Table 5: Final GST status of IR10 after consistency rules applied

	N(firms)	Reported as incl.	Proportion Changed incl.→excl.	Treated as incl.
2001	186,261	0.108	0.015	0.095
2002	187,302	0.101	0.017	0.084
2003	192,849	0.092	0.017	0.075
2004	197,307	0.087	0.021	0.067
2005	200,130	0.079	0.021	0.058
2006	204,138	0.076	0.026	0.051
2007	206,304	0.075	0.029	0.046
2008	210,453	0.072	0.031	0.041
2009	207,171	0.068	0.031	0.037
2010	204,297	0.067	0.033	0.034
2011	206,553	0.068	0.034	0.034
2012	205,905	0.066	0.032	0.034
Total	**2,408,670**	**0.079**	**0.026**	**0.054**

Restricted to firms with usable (at this processing step) IR10s and $L > 0$ in the current year. The proportion "changed" doesn't quite equal the difference between the "reported as" and "treated as" inclusive proportions because a small number of missing responses are imputed as inclusive, which are included in the "treated as" proportion.

Table 6: Source for tax depreciation under new IR10 form, 2013-2016 pooled

Data source	N(obs)	Proportion
Reported tax depreciation		
$\delta_{\text{TAX}} > 0$	662,436	0.766
$\delta_{\text{TAX}} = \delta_{\text{ACC}} = 0$	61,809	0.071
Reported accounting depreciation ($\delta_{\text{TAX}} = 0$)		
$\delta_{\text{ACC}} > 0$ & $\delta_{\text{ACC}} = \delta_{\text{TAX}}$ in another year	139,128	0.161
$\delta_{\text{ACC}} > 0$ & $\delta_{\text{ACC}} \neq \delta_{\text{TAX}}$ in all years	546	0.001
Imputed from CBV of total fixed assets	660	0.001
Total	**864,579**	**1.000**

Restricted to firms with usable (at this processing step) IR10s and $L > 0$ in the current year (2013-2016). δ_{ACC} and δ_{TAX} denote accounting and tax depreciation expenditure respectively. $\delta_{\text{ACC}} = \delta_{\text{TAX}}$ is defined as either: both tax and accounting depreciation non-zero and different by $\leq 1\%$ (using the average of the two measures as the denominator); or $\delta_{\text{ACC}} > 0$, $\delta_{\text{TAX}} = 0$ and the tax adjustment variable is zero (implying that tax depreciation is missing and equal to the accounting value). Closing book value (CBV) imputation uses the firm-level average of $\delta_{\text{TAX}}/(\delta_{\text{TAX}} + \text{TFA})$ from years where $\delta_{\text{TAX}} > 0$ and $\text{TFA} > 0$ (including years prior to 2013).

Table 7: Effect of tax adjustment on data coverage and M

	N(firms)	Dropped	Proportion		$\overline{\text{adj}/M}^{*}$	
			$\text{Tax}_{\text{adj}} > 0$	$\text{Tax}_{\text{adj}} > 0$	$\text{Tax}_{\text{adj}} > 0$	$\text{Tax}_{\text{adj}} > 0$
2013	216,060	0.051	0.115	0.026	-0.010	0.011
2014	219,990	0.049	0.131	0.031	-0.010	0.011
2015	218,688	0.053	0.144	0.036	-0.010	0.011
2016	209,841	0.056	0.162	0.028	-0.010	0.012
Total	**864,579**	**0.052**	**0.138**	**0.031**	**-0.010**	**0.011**

Restricted to firms with usable (at this processing step) IR10s and $L > 0$ in the current year. Dropped observations have a tax adjustment greater than 5% of intermediate consumption as at this processing step (M^*), as a percentage of the centred average of M^* before and after adjustment. Reported mean changes in M^* are calculated using the same centred average denominator. A positive (negative) tax adjustment is assumed to result from over-reporting (under-reporting) of accounting expenses relative to tax deduction rules, implying the need to deduct (add) the tax adjustment from (to) M^* to get a better tax-based measure of intermediate consumption.

Table 8: Proportion of final IR10-based productivity observations affected by edit processes

		Proportion adjusted					Proportion adjusted (L-weighted)				
	N(firms)	Gain/loss	W&S	GST	Tax adj	RLR	Gain/loss	W&S	GST	Tax adj	RLR
2001	173,403	0.195	0.094	0.081		0.239	0.283	0.168	0.057		0.132
2002	174,357	0.198	0.097	0.070		0.241	0.285	0.171	0.048		0.130
2003	179,736	0.197	0.095	0.062		0.242	0.290	0.173	0.041		0.128
2004	184,050	0.207	0.096	0.054		0.245	0.299	0.179	0.036		0.128
2005	187,035	0.204	0.100	0.047		0.248	0.301	0.184	0.031		0.127
2006	190,449	0.199	0.101	0.039		0.256	0.297	0.187	0.025		0.133
2007	193,476	0.191	0.098	0.036		0.258	0.291	0.185	0.022		0.130
2008	197,790	0.189	0.101	0.031		0.265	0.288	0.191	0.019		0.132
2009	194,733	0.176	0.099	0.028		0.269	0.265	0.186	0.017		0.141
2010	193,377	0.165	0.086	0.027		0.270	0.255	0.161	0.016		0.142
2011	195,618	0.167	0.083	0.026		0.275	0.256	0.149	0.015		0.145
2012	195,201	0.168	0.082	0.025		0.279	0.257	0.153	0.014		0.144
2013	197,085	0.160	0.085		0.143	0.269	0.258	0.160		0.280	0.140
2014	201,213	0.164	0.076		0.163	0.268	0.267	0.142		0.313	0.139
2015	199,158	0.165	0.074		0.182	0.264	0.272	0.143		0.344	0.135
2016	190,401	0.167	0.075		0.193	0.263	0.276	0.145		0.360	0.134
Total (01-12)	**2,259,225**	**0.188**	**0.094**	**0.043**		**0.258**	**0.280**	**0.174**	**0.028**		**0.134**
Total (13-16)	**787,857**	**0.164**	**0.077**		**0.170**	**0.266**	**0.268**	**0.147**		**0.325**	**0.137**

Restricted to non-AES final productivity dataset observations. Adjustments: remove gains/losses from fixed asset sales from Y/M (gain/loss); remove direct labour costs from purchases (W&S); remove GST from Y and M (GST); add/remove the difference between tax and accounting values to/from M (tax adj); and reallocate imputed unitemised rental, leasing & rates expenses from M to K (RLR). GST-inclusive reporting is estimated to be minor post-2012 and is ignored. Tax adjustment is necessary only from 2013 under the new IR10 form.

Table 9: Productivity dataset coverage of all firm population, by year

	All firms ($L > 0$)	Productivity Population	Productivity Dataset	Proportion of all firms/employment Included	Out of scope	No AES or IR10	Dropped
			N(firms)				
2001	361,062	296,532	181,143	0.502	0.179	0.187	0.132
2002	358,263	295,029	181,803	0.507	0.177	0.193	0.123
2003	361,887	298,863	187,146	0.517	0.174	0.198	0.111
2004	366,588	303,429	191,658	0.523	0.172	0.211	0.094
2005	369,183	306,174	194,658	0.527	0.171	0.210	0.092
2006	373,515	310,563	198,945	0.533	0.169	0.218	0.081
2007	375,822	312,711	201,108	0.535	0.168	0.218	0.079
2008	377,415	314,004	205,356	0.544	0.168	0.205	0.083
2009	372,327	309,690	202,725	0.544	0.168	0.206	0.082
2010	364,461	302,679	200,169	0.549	0.170	0.203	0.078
2011	364,110	302,520	202,425	0.556	0.169	0.197	0.078
2012	361,764	300,405	201,720	0.558	0.170	0.195	0.078
2013	359,808	299,472	203,586	0.566	0.168	0.183	0.083
2014	359,652	299,595	207,777	0.578	0.167	0.175	0.080
2015	357,831	298,008	205,878	0.575	0.167	0.163	0.094
2016	354,969	295,569	197,901	0.558	0.167	0.198	0.078
Total	**5,838,657**	**4,845,243**	**3,163,998**	**0.542**	**0.170**	**0.198**	**0.090**
			Total employment (L)				
2001	1,593,500	1,181,600	896,200	0.562	0.258	0.083	0.096
2002	1,620,200	1,198,200	905,700	0.559	0.260	0.092	0.089
2003	1,658,600	1,225,100	948,600	0.572	0.261	0.089	0.077
2004	1,706,200	1,259,300	990,100	0.580	0.262	0.093	0.064
2005	1,749,000	1,292,000	1,020,000	0.583	0.261	0.094	0.062
2006	1,791,400	1,321,800	1,054,300	0.589	0.262	0.096	0.054
2007	1,814,600	1,335,800	1,066,100	0.588	0.264	0.098	0.051
2008	1,847,700	1,355,900	1,083,900	0.587	0.266	0.092	0.055
2009	1,844,900	1,337,700	1,083,600	0.587	0.275	0.086	0.052
2010	1,800,700	1,284,800	1,038,800	0.577	0.286	0.086	0.050
2011	1,809,200	1,286,300	1,039,100	0.574	0.289	0.086	0.050
2012	1,823,600	1,297,000	1,060,500	0.582	0.289	0.081	0.048
2013	1,842,500	1,310,300	1,066,600	0.579	0.289	0.080	0.052
2014	1,878,200	1,338,200	1,101,400	0.586	0.288	0.075	0.051
2015	1,916,200	1,367,400	1,125,900	0.588	0.286	0.068	0.058
2016	1,964,900	1,394,100	1,129,800	0.575	0.290	0.085	0.049
Total	**28,661,400**	**20,785,500**	**16,610,600**	**0.580**	**0.275**	**0.086**	**0.059**

The productivity population is all $L > 0$ private-for-profit firms in productivity industries (Stats NZ "measured sector"). Out of scope firms are, therefore, businesses in the "all firm" population that are not in the productivity population, including the public sector. A firm is deemed to have financial data if they have a non-imputed AES postal return and/or an IR10 where neither the front nor back page are zero-punched (ignoring the "other information" rule). Unlike the productivity dataset construction, this test counts financial filing from all enterprises within a permanent enterprise (PENT), rather than just the current enterprise in the PENT chain.

Table 10: Coverage of productivity population, by employer status and year

	Productivity Population	Productivity Dataset	Proportion of productivity pop. Included	Proportion of productivity pop. No AES or IR10	Proportion of productivity pop. Dropped
			N(employers)		
2001	127,764	86,274	0.675	0.152	0.173
2002	129,657	87,678	0.676	0.160	0.163
2003	133,275	92,280	0.692	0.165	0.142
2004	137,394	95,967	0.698	0.182	0.119
2005	141,414	99,030	0.700	0.184	0.116
2006	144,450	101,409	0.702	0.196	0.101
2007	146,082	103,119	0.706	0.197	0.098
2008	147,714	105,327	0.713	0.181	0.106
2009	144,402	103,017	0.713	0.182	0.104
2010	138,546	100,188	0.723	0.178	0.099
2011	137,409	100,368	0.730	0.171	0.098
2012	137,037	100,566	0.734	0.168	0.098
2013	137,646	101,667	0.739	0.160	0.102
2014	139,311	104,646	0.751	0.152	0.097
2015	142,206	105,357	0.741	0.141	0.118
2016	144,582	101,415	0.701	0.206	0.092
Total	**2,228,889**	**1,588,308**	**0.713**	**0.174**	**0.114**
			N(WP-only firms)		
2001	168,768	94,869	0.562	0.286	0.152
2002	165,372	94,125	0.569	0.293	0.138
2003	165,588	94,866	0.573	0.299	0.128
2004	166,035	95,691	0.576	0.314	0.109
2005	164,760	95,628	0.580	0.313	0.106
2006	166,113	97,536	0.587	0.320	0.093
2007	166,629	97,989	0.588	0.320	0.092
2008	166,290	100,029	0.602	0.304	0.094
2009	165,288	99,708	0.603	0.304	0.093
2010	164,133	99,981	0.609	0.301	0.090
2011	165,111	102,057	0.618	0.292	0.090
2012	163,368	101,154	0.619	0.291	0.090
2013	161,826	101,919	0.630	0.272	0.098
2014	160,284	103,131	0.643	0.260	0.096
2015	155,802	100,521	0.645	0.247	0.108
2016	150,987	96,486	0.639	0.267	0.094
Total	**2,616,354**	**1,575,690**	**0.602**	**0.293**	**0.105**

See table 9 for notes. In contrast to table 9, proportions are relative to the productivity population, rather than the all firm population.

Table 11: Standard deviation of multifactor productivity – new vs old dataset

	New productivity data			Old productivity data		
	σ(MFP)	$\sigma(\delta_i)$	$\sigma(\epsilon_{it})$	σ(MFP)	$\sigma(\delta_i)$	$\sigma(\epsilon_{it})$
2001	0.716	0.661	0.404	0.766	0.714	0.413
2002	0.714	0.636	0.393	0.759	0.681	0.400
2003	0.717	0.621	0.391	0.765	0.663	0.399
2004	0.709	0.608	0.389	0.755	0.648	0.393
2005	0.696	0.594	0.383	0.747	0.640	0.386
2006	0.696	0.583	0.387	0.746	0.626	0.391
2007	0.694	0.576	0.382	0.745	0.621	0.388
2008	0.700	0.575	0.381	0.748	0.617	0.385
2009	0.700	0.568	0.380	0.746	0.614	0.386
2010	0.698	0.560	0.380	0.748	0.609	0.390
2011	0.695	0.560	0.369	0.740	0.611	0.380
2012	0.684	0.555	0.367	0.728	0.619	0.385
2013	0.688	0.556	0.368			
2014	0.674	0.552	0.358			
2015	0.672	0.556	0.364			
2016	0.678	0.571	0.381			

MFP= $\delta_i + \epsilon_{it}$ where ϵ_{it} is the residual from an industry-specific gross output Cobb-Douglas production function with firm fixed effects (δ_i). MFP for the old productivity dataset is estimated using corrected (old method) WP counts.

Table 12: Standard deviation of multifactor productivity, by continuity of appearance in dataset

	Retained firms		N(firms)		Gained firms			Lost firms		
	$\sigma(\text{MFP}_{new})$	$\sigma(\text{MFP}_{old})$	Gained	Lost	$\sigma(\text{MFP}_{new})$	$\sigma(\delta_i)$	$\sigma(\epsilon_{it})$	$\sigma(\text{MFP}_{old})$	$\sigma(\delta_i)$	$\sigma(\epsilon_{it})$
2001	0.714	0.720	22,104	32,328	0.731	0.626	0.477	0.946	0.856	0.515
2002	0.706	0.712	22,875	29,040	0.771	0.627	0.489	0.965	0.830	0.523
2003	0.713	0.717	22,965	26,847	0.751	0.594	0.473	0.994	0.814	0.543
2004	0.705	0.709	24,267	22,965	0.740	0.571	0.482	1.014	0.812	0.561
2005	0.690	0.699	24,921	22,119	0.741	0.565	0.481	1.028	0.823	0.563
2006	0.692	0.700	25,410	19,335	0.731	0.567	0.472	1.057	0.830	0.578
2007	0.686	0.695	26,139	19,005	0.748	0.558	0.474	1.087	0.836	0.599
2008	0.693	0.701	26,904	19,518	0.752	0.560	0.467	1.076	0.818	0.586
2009	0.695	0.702	27,357	18,891	0.733	0.540	0.456	1.059	0.807	0.593
2010	0.692	0.699	29,727	17,928	0.731	0.543	0.455	1.093	0.824	0.609
2011	0.688	0.695	33,123	17,517	0.733	0.548	0.435	1.069	0.812	0.594
2012	0.680	0.688	37,989	17,142	0.698	0.537	0.409	1.024	0.828	0.572

See table 11 for notes. "Retained" observations appear in both the old and new productivity datasets, but with potentially different estimated MFP. "Gained" ("lost") observations appear in the new (old) productivity dataset but not the old (new) dataset. A small percentage of lost observations (3.1%) are lost, not because of productivity data changes, but because a firm exits the productivity population due to revised industry or private-for-profit status, loss of L, or changes in the permanent enterprise identifier.

Table 13: Coverage of productivity population, by size strata

	N(observations) Population	Dataset	Coverage rate
WP-only			
Entry	225,834	119,562	0.529
Continuer	2,169,579	1,392,111	0.642
Exit/1yr	220,947	64,026	0.290
Subtotal WP-only	**2,616,360**	**1,575,699**	**0.602**
Employer			
$L \in (0, 2]$ & entry	107,532	68,946	0.641
$L \in (0, 2]$ & continuer	675,006	461,634	0.684
$L \in (0, 2]$ & exit/1yr	75,267	19,323	0.257
$L \in (2, 5]$	809,142	603,924	0.746
$L \in (5, 10]$	307,335	236,250	0.769
$L \in (10, 20]$	144,090	110,490	0.767
$L \in (20, 50]$	71,457	54,732	0.766
$L \in (50, 100]$	21,324	17,151	0.804
$L \in (100, \text{inf})$	17,736	15,852	0.894
Subtotal employer	**2,228,889**	**1,588,302**	**0.713**
Total	**4,845,249**	**3,164,001**	**0.653**

"Entry" ("exit") firms are not active in the immediately prior (following) year. Activity is assessed using L, GST and IR10/AES data. Because the first two of these dataset are essentially available for the full population of firms, assessing activity in adjacent years is not reliant on the presence of productivity data in those years. "Continuers" are neither entrants nor exiters. "1yr" firms are a relatively small group that are both entrants and exiters. These firms are grouped with exiters due to the similarity in data coverage rate between the two groups.

Table 14: Imputation rate for rental, leasing and rates

	Proportion of firms		
	Unimputed	Mean-imputed	Modelled
2001	0.733	0.110	0.156
2002	0.733	0.126	0.141
2003	0.733	0.136	0.131
2004	0.729	0.142	0.129
2005	0.726	0.146	0.129
2006	0.720	0.151	0.129
2007	0.718	0.155	0.128
2008	0.710	0.160	0.129
2009	0.706	0.165	0.129
2010	0.706	0.166	0.128
2011	0.701	0.170	0.129
2012	0.697	0.173	0.131
2013	0.709	0.160	0.131
2014	0.711	0.153	0.136
2015	0.715	0.144	0.141
2016	0.713	0.133	0.153
Total	**0.716**	**0.150**	**0.134**

Productivity population-weighted and restricted to firms with $K > 0$. Mean-imputed rental, leasing and rates (RLR) uses the firm-level average of $\mathrm{RLR}_{\mathrm{IR10}}/M_{\mathrm{IR10}}$ from years where $\mathrm{RLR}_{\mathrm{IR10}}$ is reported (ie, non-zero). Firms with an AES return and an IR10 return with reported RLR are counted as unimputed ($\mathrm{RLR} = \mathrm{RLR}_{\mathrm{IR10}} \times M_{\mathrm{AES}}/M_{\mathrm{IR10}}$).

Table 15: Contribution of capital components to aggregate K

	Proportion of aggregate K					
	Cost of capital	Depreciation	Rental, leasing & rates (RLR)	RLR Unimputed	RLR Mean-imputed	RLR Modelled
2001	0.395	0.320	0.284	0.164	0.030	0.090
2002	0.397	0.321	0.281	0.169	0.031	0.081
2003	0.393	0.330	0.276	0.175	0.036	0.065
2004	0.386	0.325	0.289	0.179	0.037	0.073
2005	0.387	0.317	0.296	0.184	0.037	0.076
2006	0.392	0.309	0.299	0.185	0.036	0.077
2007	0.406	0.300	0.294	0.188	0.036	0.070
2008	0.401	0.300	0.300	0.191	0.036	0.073
2009	0.398	0.292	0.310	0.190	0.039	0.081
2010	0.409	0.287	0.305	0.188	0.039	0.078
2011	0.415	0.282	0.303	0.187	0.039	0.076
2012	0.425	0.271	0.304	0.190	0.039	0.074
2013	0.422	0.274	0.304	0.190	0.040	0.074
2014	0.418	0.273	0.310	0.192	0.041	0.077
2015	0.416	0.271	0.312	0.196	0.036	0.081
2016	0.418	0.274	0.308	0.185	0.040	0.083
Total	**0.407**	**0.293**	**0.300**	**0.186**	**0.037**	**0.077**

Productivity population-weighted. Mean-imputed rental, leasing and rates (RLR) uses the firm-level average of $\mathrm{RLR}_{\mathrm{IR10}}/M_{\mathrm{IR10}}$ from years where $\mathrm{RLR}_{\mathrm{IR10}}$ is reported (ie, non-zero). Firms with an AES return and an IR10 return with reported RLR are counted as unimputed ($\mathrm{RLR} = \mathrm{RLR}_{\mathrm{IR10}} \times M_{\mathrm{AES}}/M_{\mathrm{IR10}}$).

Table 16: Alternative Cobb-Douglas production function specifications for K (firm fixed effects & pooled industries)

	(1)	(2)	(3)	(4)	(5)
$\ln(M)$	0.597***	0.599***		0.596***	0.597***
	[0.00131]	[0.00101]		[0.00131]	[0.00131]
$\ln(M + \text{RLR})$			0.690***		
			[0.00135]		
$\ln(L)$	0.237***	0.238***	0.237***	0.236***	0.237***
	[0.00118]	[0.000970]	[0.00117]	[0.00118]	[0.00118]
$\ln(K)$	0.141***	0.142***		0.140***	
	[0.00108]	[0.000891]		[0.00108]	
ln(TFA)			0.0499***		
			[0.000647]		
ln(Intangibles)				0.00980***	
				[0.000566]	
$\ln(K + K_{\text{Intang}})$					0.143***
					[0.00110]
N(observations)	3,112,464	3,112,464	3,112,464	3,112,464	3,112,464
Returns to scale (RTS)	0.974	0.979	0.977	0.982	0.977
σ(MFP)	0.715	0.714	0.702	0.714	0.714
75^{th}-25^{th}(MFP)	0.560	0.558	0.545	0.559	0.559
90^{th}-10^{th}(MFP)	1.339	1.335	1.304	1.338	1.336
Pearson's correlation with (1)		1.0000	0.9868	0.9998	0.9995
Spearman rank correlation with (1)		0.9999	0.9718	0.9993	0.9986
Population-weighted	N	Y	N	N	N
δ("K" $= 0$)	N	N	Y	Y	N

Estimation is OLS with firm fixed effects and year dummies (robust standard errors in brackets; *** significantly different from zero at 1% level). Returns to scale are always different from one ($p = 0.000$). δ("K" $= 0$)=Y indicates a specification that includes an (unreported) indicator variable for firms with zero "K" (the base population is $K > 0$). Real average book values are used for total fixed assets (TFA) & intangibles (intangibles deflated using the capital deflator for internal consistency, and the lack of a better alternative). To be consistent with the cost-of-capital calculation $K_{\text{Intang}} = 0.1\times$ Intangibles. MFP-related summary statistics (third panel) are population-weighted. Prior to calculating summary statistics, all MFP estimates are industry-year demeaned (consistent with the estimates in the productivity dataset).

Table 17: Median mark-up by year and employer status

	Employing firms	WP-only firms $\hat{W} > 0$	WP-only firms $\hat{W} = 0$
2001	0.043	-0.343	0.299
2002	0.053	-0.334	0.334
2003	0.044	-0.338	0.337
2004	0.042	-0.332	0.360
2005	0.042	-0.328	0.377
2006	0.032	-0.334	0.395
2007	0.029	-0.338	0.392
2008	0.033	-0.336	0.419
2009	0.015	-0.356	0.404
2010	0.014	-0.376	0.391
2011	0.026	-0.357	0.428
2012	0.033	-0.342	0.465
2013	0.031	-0.344	0.488
2014	0.049	-0.323	0.535
2015	0.051	-0.319	0.562
2016	0.049	-0.322	0.589

Unweighted medians where mark-up is defined as $(Y/[M+\hat{W}+K]-1)$ and $\hat{W}$ is total gross earnings, including imputed earnings for working proprietors (WPs). For FTE>0, imputed WP earnings are N(WP)×[total employee earnings/total employee FTE]. For WP-only firms (FTE=0), imputed earnings are the summed industry-year analogue (median related mark-up reported in the $\hat{W} > 0$ column). The final column reports an alternative (extreme) assumption of zero labour compensation for WPs in WP-only firms ($\hat{W} = 0$).

Figures

Figure 1: Proportion of IR10-based productivity observations with raw edits

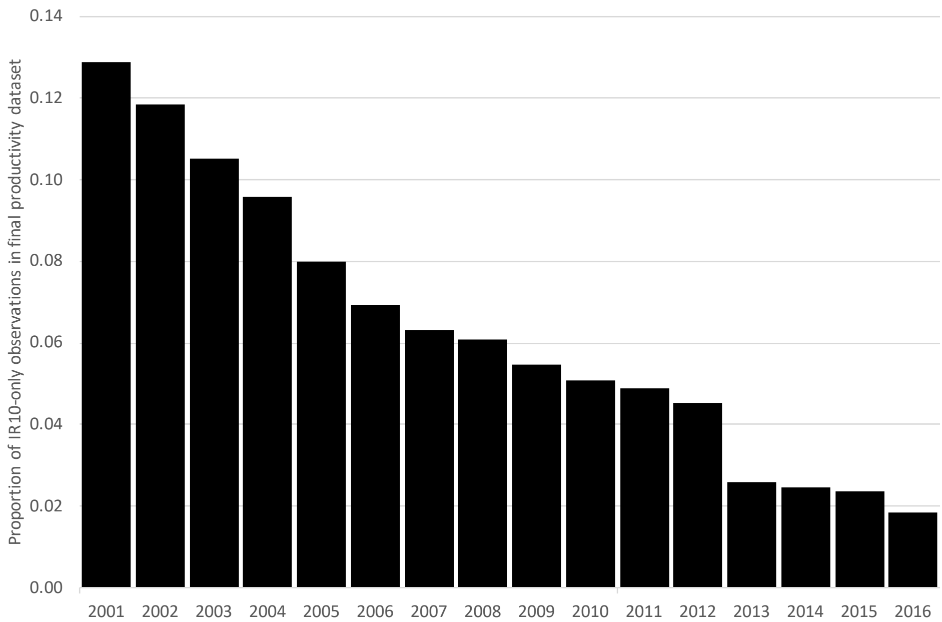

Raw edits adjust: purchases to be consistent with gross profit (*flag_i10_purch_edited*) and/or other income to fix misreporting of totals in wrong boxes (*flag_i10_othinc_edited*). Denominator restricted to IR10-only productivity observations – ie, where productivity components are based on IR10 responses, not joint AES-IR10 returns where AES is the primary data source.

Figure 2: Cumulative distribution of IR10 quality by vintage & prior usage

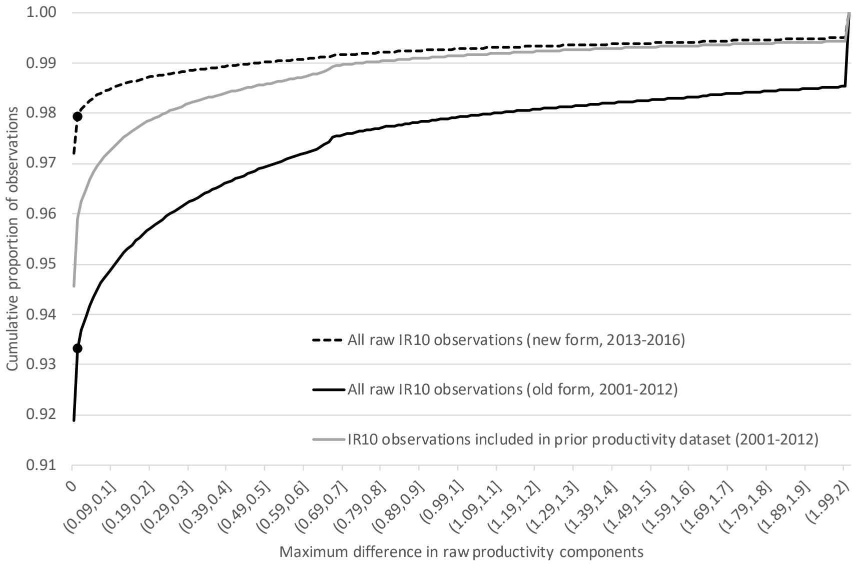

Difference measured as $2 \times (x_{max} - x_{min})/(x_{min} + x_{max})$, where x is the initial (raw) estimate of gross output, intermediate consumption or total fixed assets. Excludes IR10s that are discarded because they are zero-punched or unitemised. Restricted to $L > 0$ observations to aid comparison to the prior final productivity dataset (2001-2012). Solid dots indicate the 1% quality cut-off in new productivity dataset methodology.

Figure 3: Proportion of IR10s used in old productivity dataset by quality

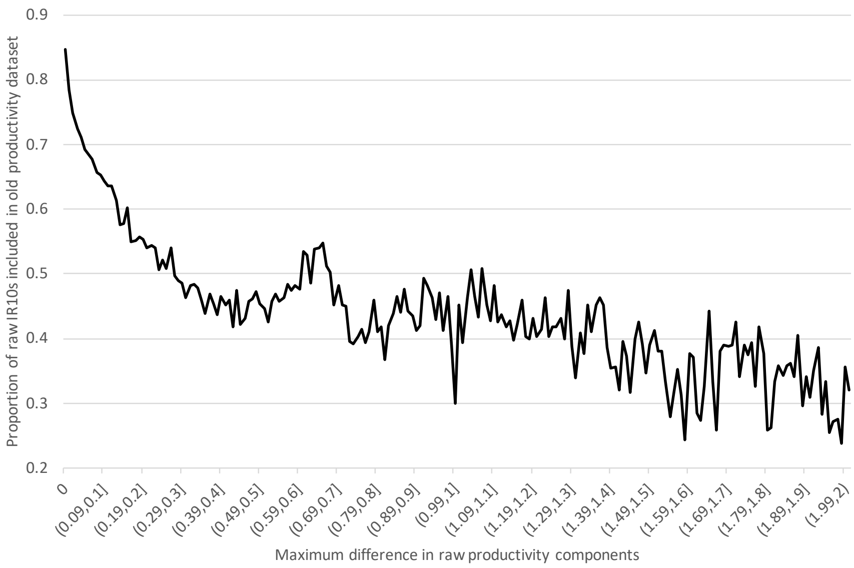

See figure 2 for notes.

Figure 4: Ratio of IR10 sales to GST sales, for GST-inclusive reporting firms

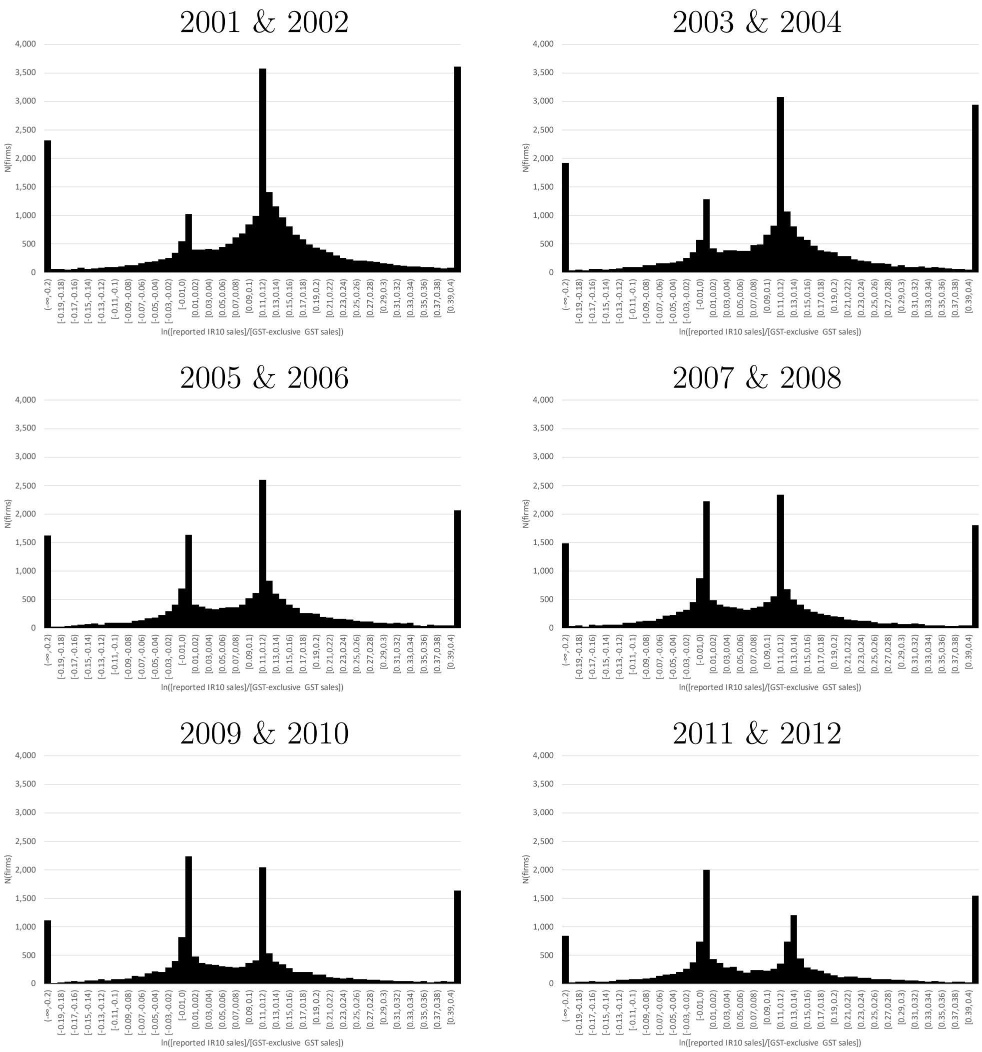

Restricted to firms with $L > 0$ and a usable (at this processing step) IR10 which reported that the return was GST-inclusive (GST field not present on new form). Pairs of years are pooled. IR10 "GST-equivalent" sales are defined as (sales+rent+otherinc) and, if reported consistently, will be GST-inclusive. Annual GST sales are derived from one/two/six-monthly mandatory GST returns (frequency related to firm size and generally aligned to firm balance date) and are GST-exclusive. To calculate the log ratio of the two measures, we require both IR10 and GST sales to be positive. We further restrict the analysis to firms with no zero-rated GST sales (which excludes, eg, exporters), so that all sales attract GST, eliminating zero-rated GST as an explanation for equality of the two sales measures.

Figure 5: Sales ratio for GST-inclusive reporting firms, by final treatment

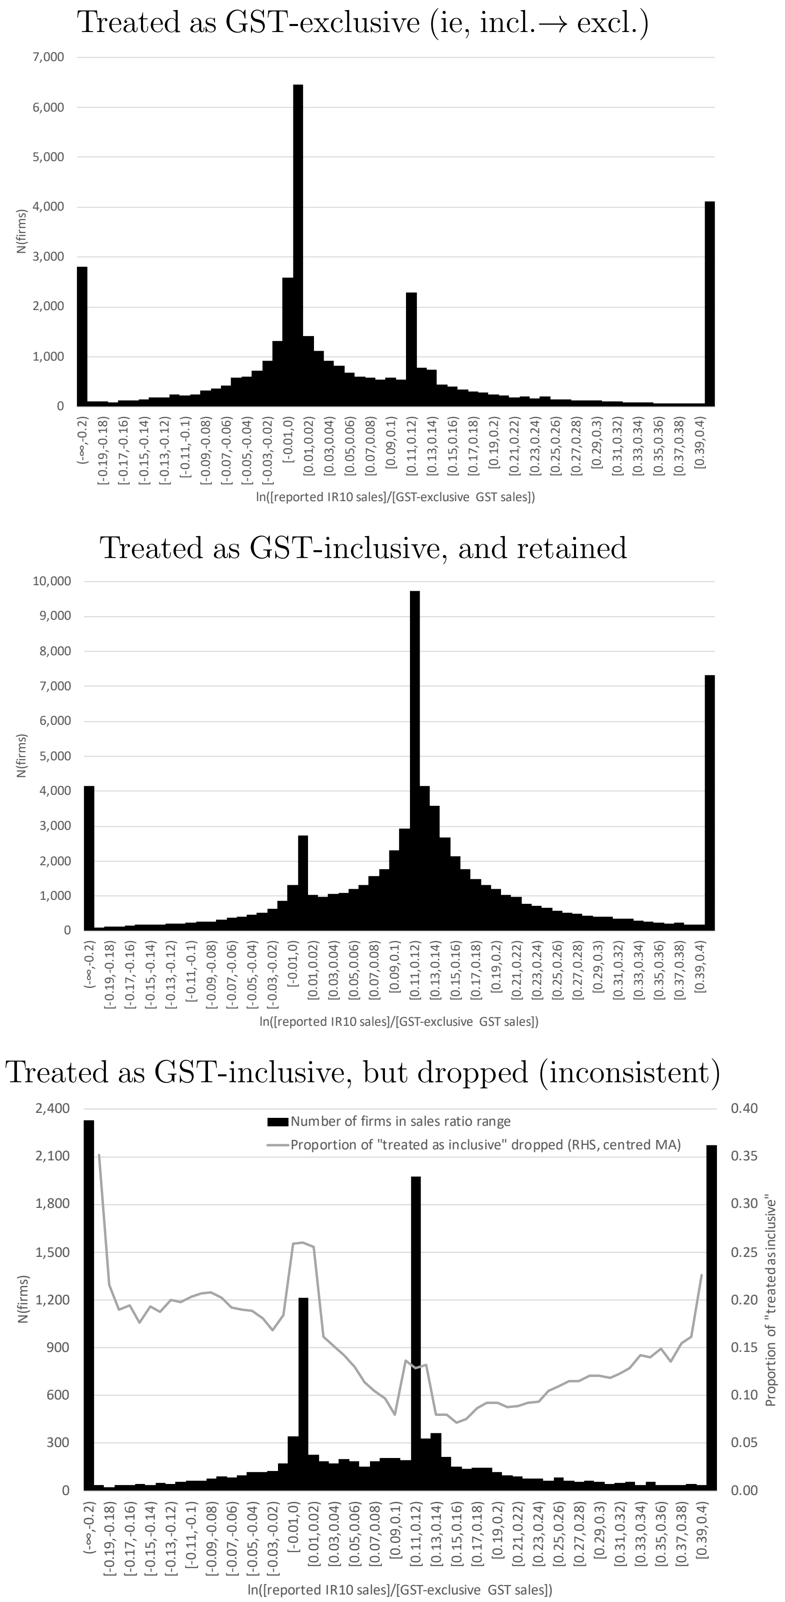

See figure 4 for notes. Bottom panel is IR10s treated as GST-inclusive, but dropped because other expenses (income) is smaller than the net GST payment (receipt) adjustment implying a potential misclassification. The proportion of "treated as inclusive" dropped is shown as a three-category centered moving average (MA).

Figure 6: Approximate distribution of tax depreciation rate, by IR10 form

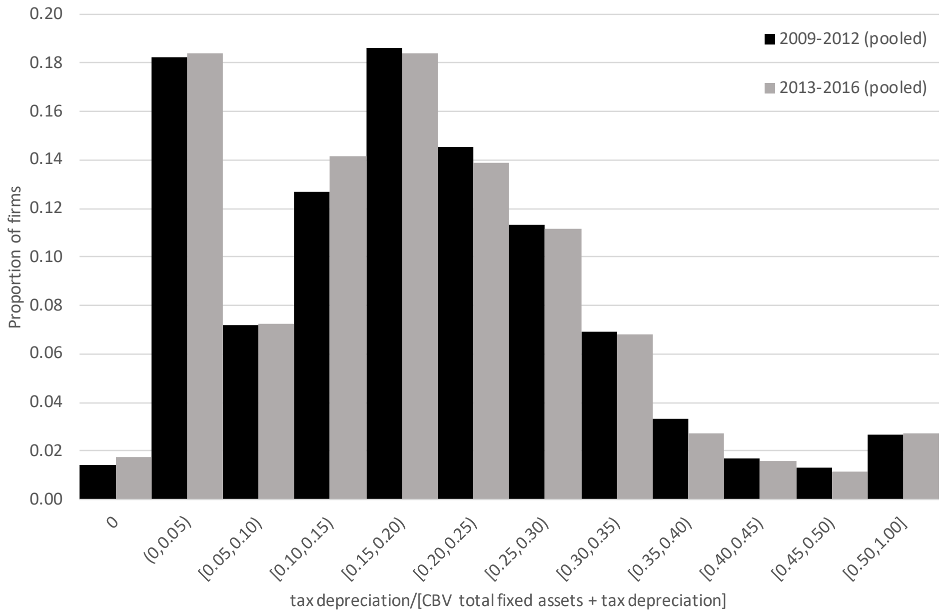

Restricted to firms with a usable (at this processing step) IR10, $L > 0$ and non-zero total fixed assets. The ratio of depreciation to depreciation plus closing book value (CBV) of total fixed asset (TFA) approximates the depreciation rate on the opening book value of TFA (ignoring additions, disposals and revaluations of fixed assets). Four-year periods are pooled by form type. In the later time period (2013-2016), tax depreciation is derived from multiple sources as summarised in table 6 and associated table note.

Figure 7: Autocorrelation in MFP – new vs old data

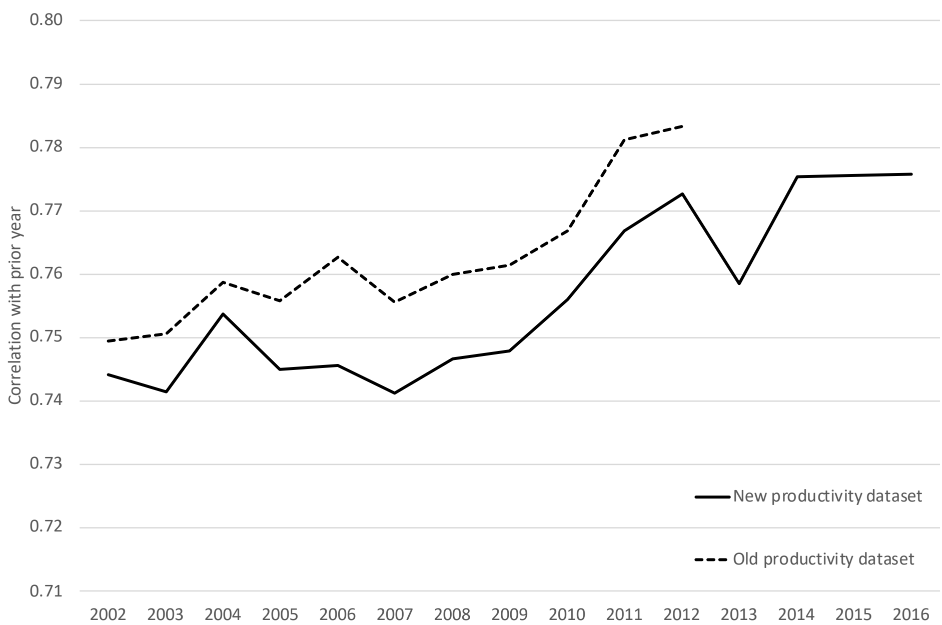

Multifactor productivity (MFP) is the estimated residual plus firm fixed effect from a productivity industry-specific gross output Cobb-Douglas production function with firm fixed effects. "New" uses latest dataset, while "old" uses the previous productivity dataset with corrected (old method) working proprietor count.

Figure 8: Autocorrelation in productivity components – new vs old data

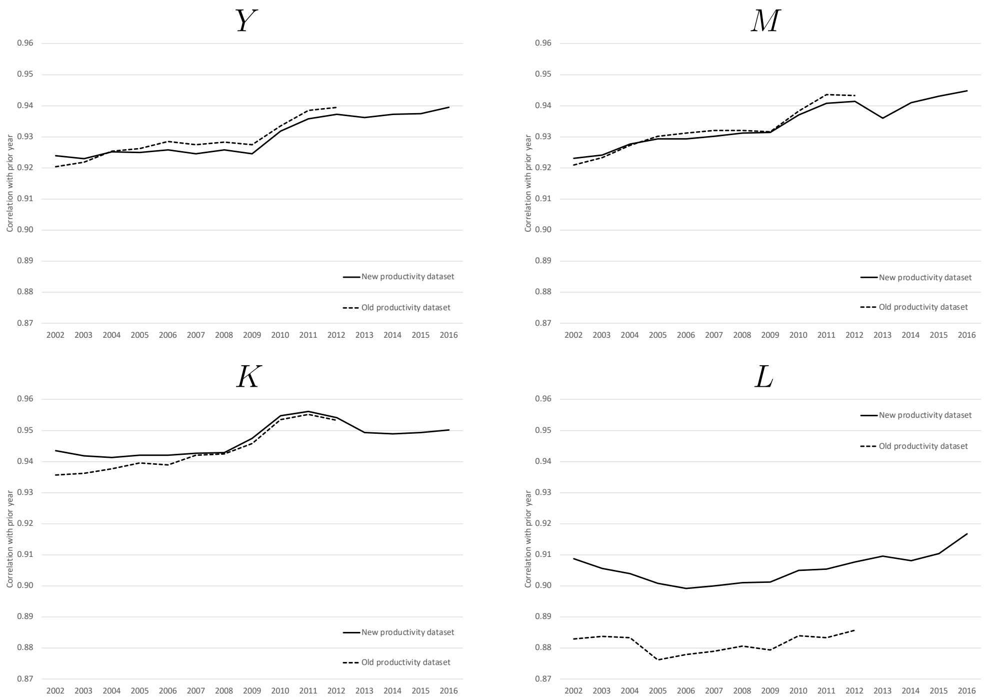

See figure 7 for notes.

Figure 9: Productivity dataset coverage for employers, by firm size

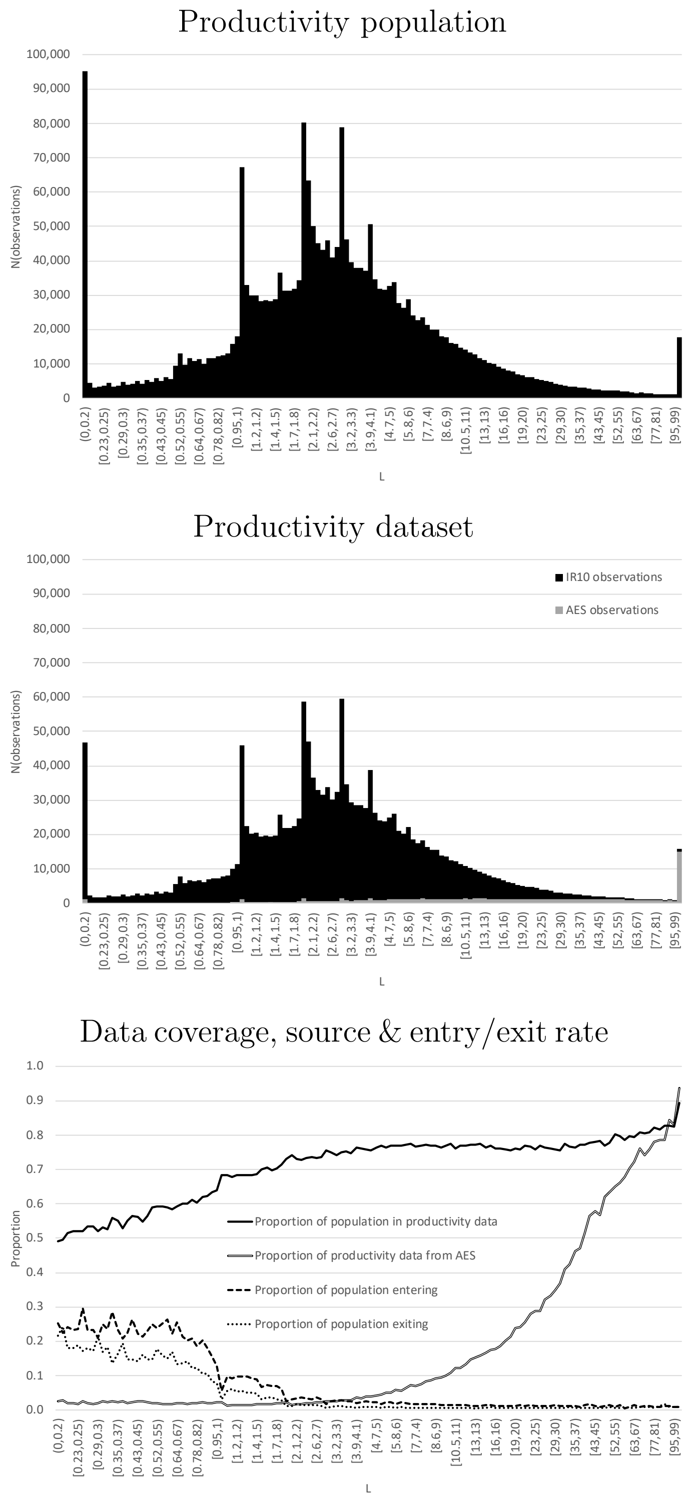

Restricted to firms with FTE>0. L groupings are equally log-spaced. "Entering" ("exiting") firms are not active in the immediately prior (following) year. Activity is assessed using L, GST and IR10/AES data. Because the first two of these dataset are essentially available for the full population of firms, assessing activity in adjacent years is not reliant on the presence of productivity data in those years.

Figure 10: Productivity dataset coverage for WP-only firms, by firm size

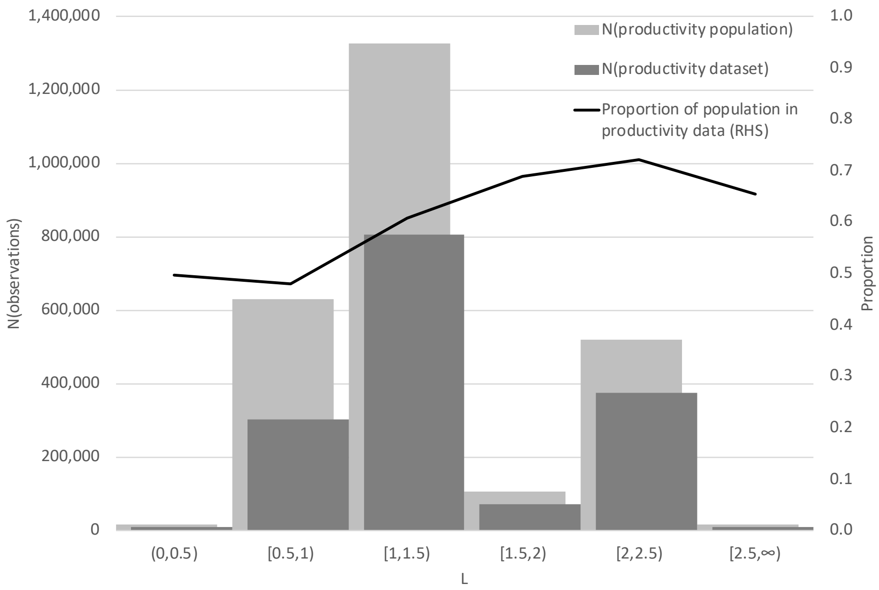

Restricted to firms with FTE= 0.

Figure 11: Effect of weighting on estimated industry L, all years pooled

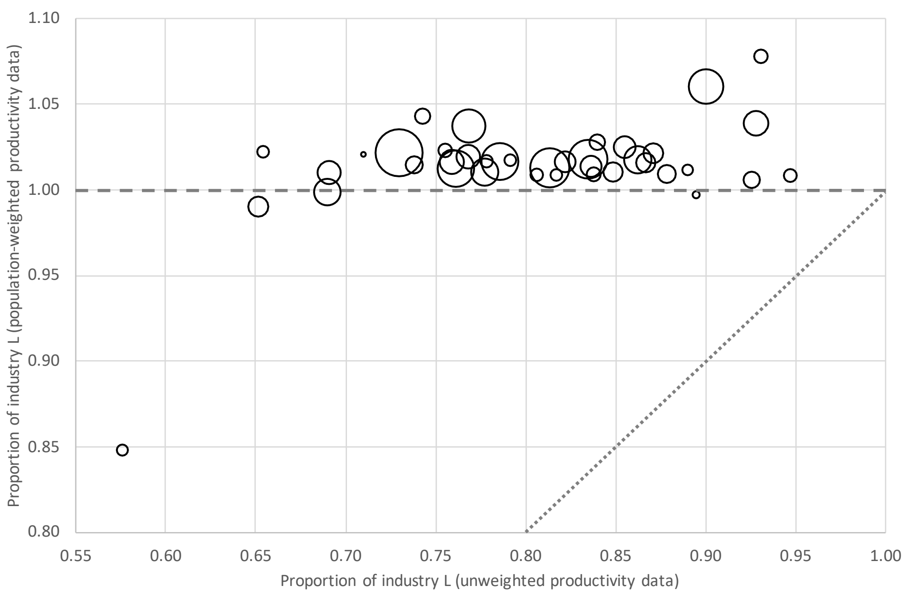

Restricted to productivity industries. Productivity population-weighting exactly replicates firm counts within industry-year-firm-size cell, where the latter is described in table 13. True industry L (denominator) comes from the productivity population dataset, and bubble area is scaled to true industry size. Dashed and dotted lines show, respectively, exact replication of true industry L (proportion of one) and the equivalence line (unweighted=weighted).

Figure 12: Labour productivity – official vs productivity dataset

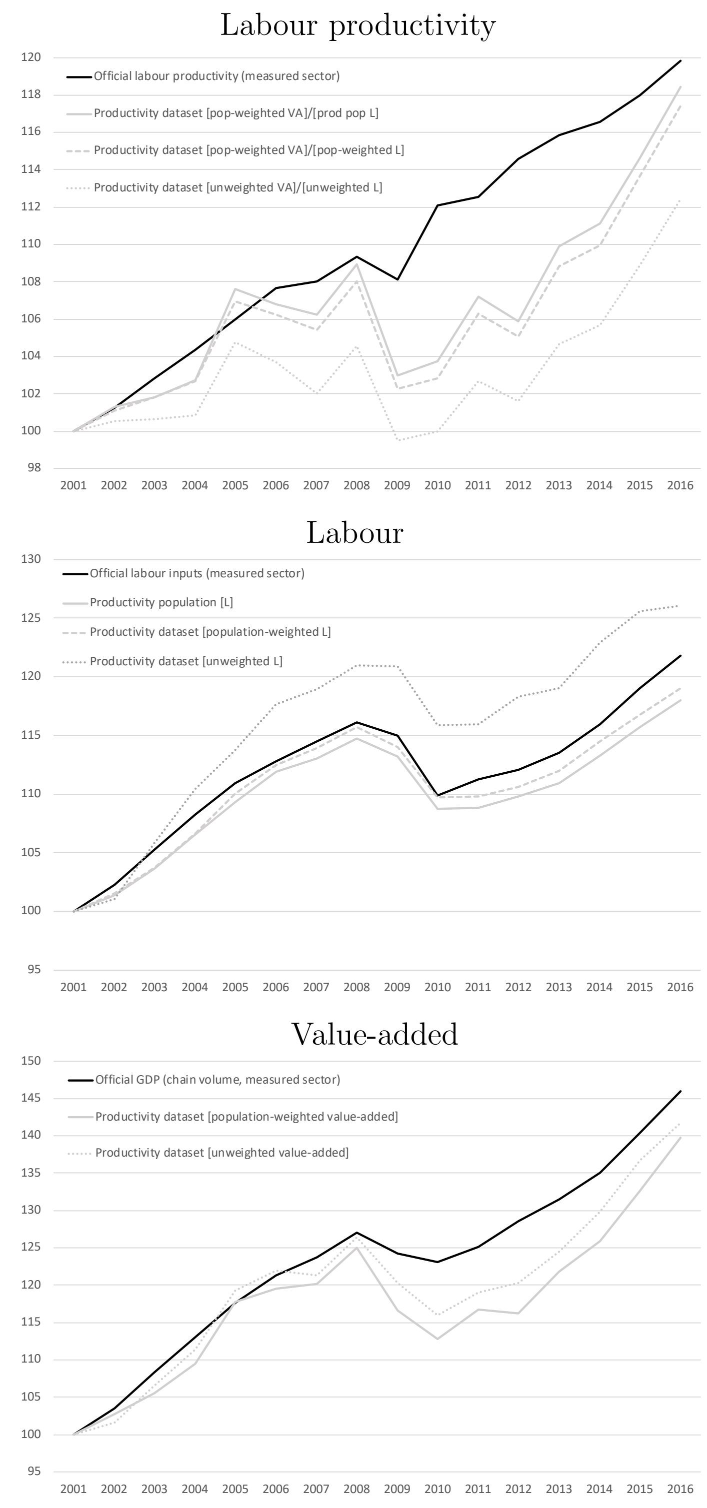

The two weighted estimates of aggregate (measured sector) labour productivity (top panel) differ in the construction of the labour series. One uses actual aggregate productivity population L ("prod pop L"), while the other uses the population-weighted productivity dataset aggregate L ("pop-weighted L"). The former is exact (for L), while the latter is internally consistent with estimated output since, in both cases, value-added is estimated as the population-weighted productivity dataset aggregate. Both labour series are shown in the middle panel of the figure.

Figure 13: MFP dispersion – unweighted vs weighted

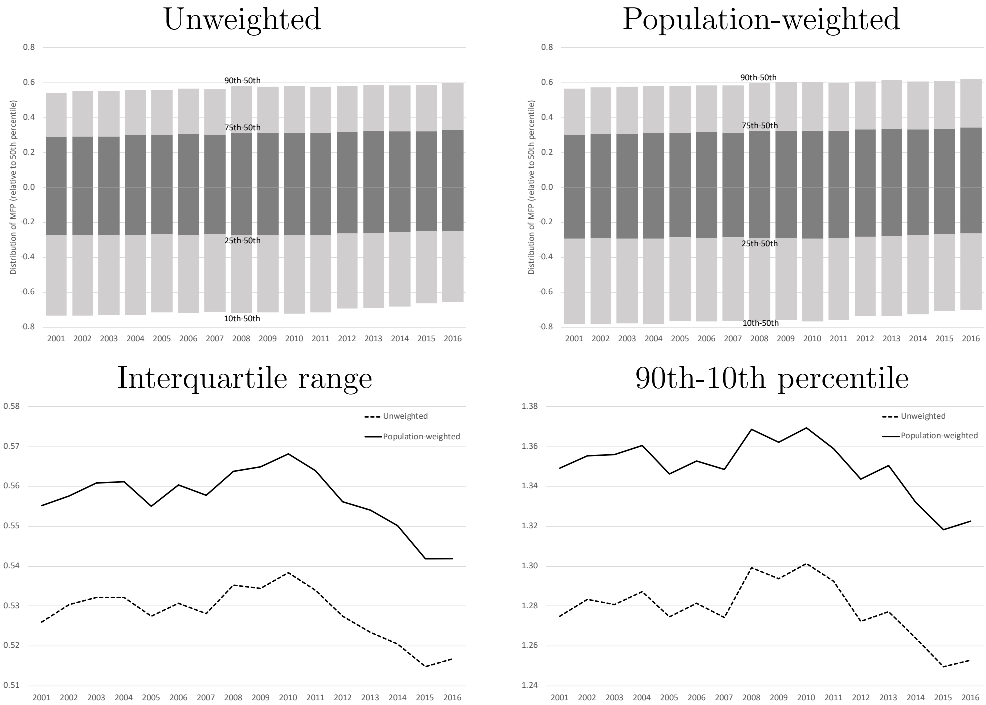

MFP= $\delta_i + \epsilon_{it}$ where ϵ_{it} is the estimated residual from an industry-specific gross output Cobb-Douglas production function with firm fixed effects (δ_i). Weighted estimates use productivity population weights. Productivity distribution in upper two panels is reported relative to median in each year (ie, as percentile less median).

Figure 14: Aggregate intangibles in productivity population industries

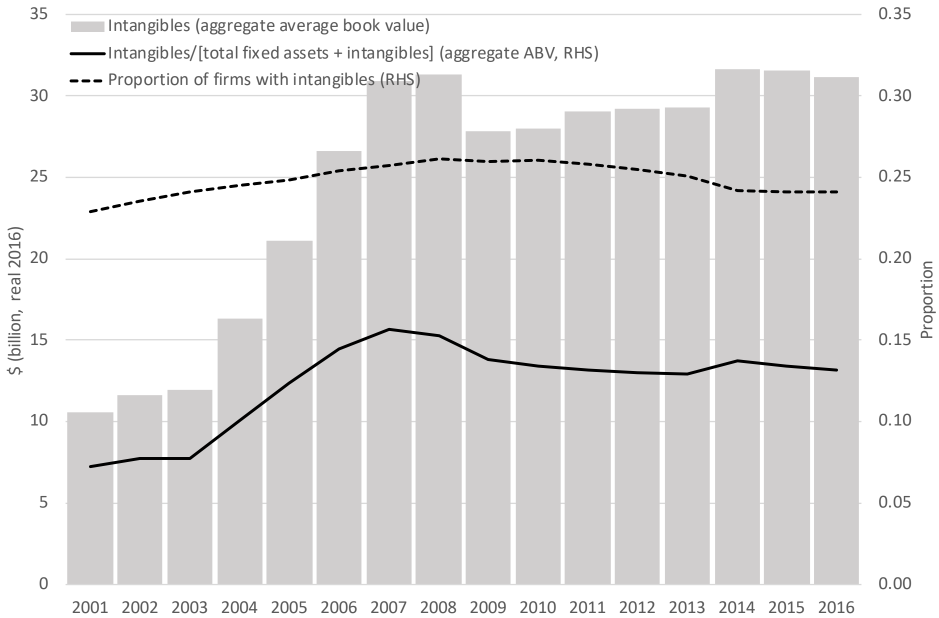

Productivity population-weighted. Intangibles deflated with industry-specific fixed capital deflator (ie, consistent with total fixed asset deflation). The proportion of firms with intangibles is as a proportion of firms with (TFA+intangibles) > 0. Fixed asset and intangibles are each measured using average book values (ABVs), where opening book values generally come from prior year closing book values.

Figure 15: Mark-up distribution – employing firms only

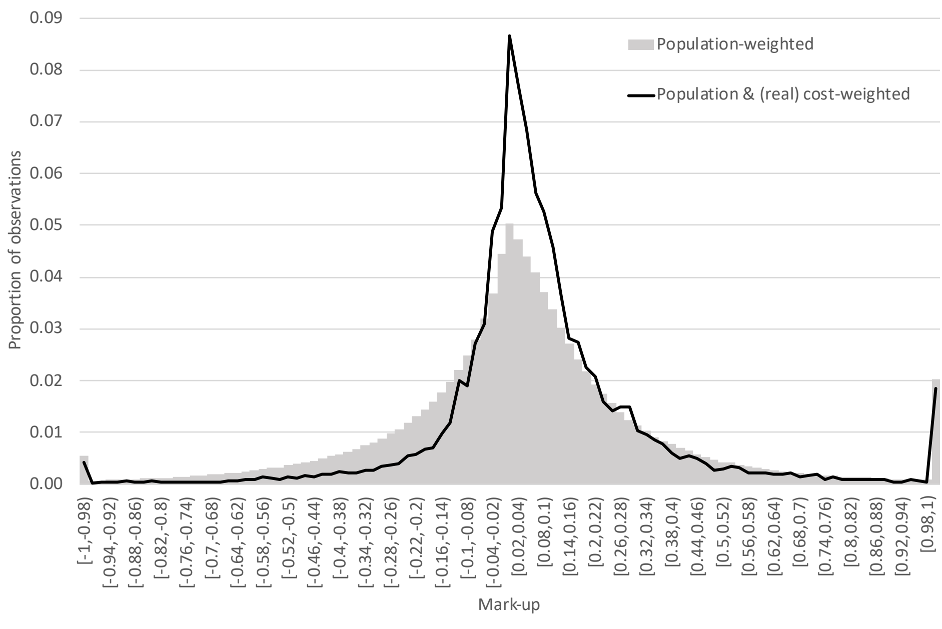

Productivity population-weighted and restricted to FTE>0. Mark-up is defined as $(Y/[M + \hat{W} + K] - 1)$ where $\hat{W}$ is total gross earnings, including imputed earnings for working proprietors (WPs). Imputed WP earnings are N(WP)×[total employee earnings/total employee FTE]. All years (2001-2016) pooled, and PPI(input) used to deflate $\hat{W}$ when additionally applying real cost weights.

Appendix A – IR10 form (2012)

Inland Revenue
Te Tari Taake

Accounts information

IR 10
November 2008

Year ended 31 March

- You only need to complete this form if you are in business. Please refer to the notes on page 3 – "About the IR 10".
- Please complete both pages of this form. Copy each amount from your financial statements.

Your full name

Your IRD number (8 digit numbers start in the second box. 1 2 3 4 5 6 7 8)

Multiple activity indicator	Read Note 1 on page 3.	1	Yes No
Gross income from	**Sales and/or services**	2	$, , .0 0
Cost of goods sold	Opening stock (include work in progress)	3	$, , .0 0
	Purchases	4	$, , .0 0
	Closing stock (include work in progress)	5	$, , .0 0
Gross profit	**(if a loss put a minus sign in the last box)**	6	$, , .0 0
Other gross income	Interest received	7	$, , .0 0
	Dividends	8	$, , .0 0
	Rental and lease payments	9	$, , .0 0
	Other income	10	$, , .0 0
Total income	**(add up all income entered in Boxes 6 to 10 —if a loss put a minus sign in the last box)**	11	$, , .0 0
Expenses (do not include amounts treated as non-deductible for tax purposes)	Bad debts (written off)	12	$, , .0 0
	Depreciation	13	$, , .0 0
	Entertainment	14	$, , .0 0
	Fringe benefit tax	15	$, , .0 0
	Insurance (exclude ACC levies)	16	$, , .0 0
	Interest expenses	17	$, , .0 0
	Legal expenses	18	$, , .0 0
	Rates	19	$, , .0 0
	Rental and lease payments	20	$, , .0 0
	Repairs and maintenance	21	$, , .0 0
	Research and development	22	$, , .0 0
	Salaries and wages	23	$, , .0 0
	Subcontractor payments	24	$, , .0 0
	Travel and accommodation	25	$, , .0 0
	Vehicle (excluding depreciation)	26	$, , .0 0
	Other expenses	27	$, , .0 0
Total expenses	**(add up all expenses entered in Boxes 12 to 27)**	28	$, , .0 0
	Total current year taxable profit (if a loss put a minus sign in the last box)	29	$, , .0 0

Balance sheet items

Current assets (as at balance date)	Accounts receivable (debtors)	30	$, , .0 0
	Bank accounts	31	$, , .0 0
	Other current assets	32	$, , .0 0
	Total current assets	33	$, , .0 0
Fixed assets (closing tax value)	Vehicles	34	$, , .0 0
	Plant and machinery	35	$, , .0 0
	Furniture and fittings	36	$, , .0 0
	Land and buildings	37	$, , .0 0
	Other fixed assets	38	$, , .0 0
	Total fixed assets	39	$, , .0 0
Other assets (as at balance date)	Intangibles	40	$, , .0 0
	Preference shares	41	$, , .0 0
	Shares and debentures	42	$, , .0 0
	Term deposits	43	$, , .0 0
	Other assets	44	$, , .0 0
Total assets		45	$, , .0 0
Current liabilities (as at balance date)	Accounts payable (creditors)	46	$, , .0 0
	Bank accounts	47	$, , .0 0
	Other current liabilities	48	$, , .0 0
	Total current liabilities	49	$, , .0 0
	Term liabilities	50	$, , .0 0
Total liabilities		51	$, , .0 0
Proprietor or shareholder funds	Drawings	52	$, , .0 0
	Current accounts closing balance **(if a debit put a minus sign in the last box)**	53	$, , .0 0
Total proprietor or shareholder funds	**(if a debit put a minus sign in the last box)**	54	$, , .0 0
Other information	Deductible loss on disposal of fixed assets	55	$, , .0 0
	Capital gain on disposal of fixed assets	56	$, , .0 0
	Dividends paid	57	$, , .0 0
	Are your accounts GST-exclusive?	58	Yes No
	Are your accounts for a period of 12 months?	59	Yes No

Cut off this page and staple it to page 3 of your return. You do not need to send in your financial records as well.

2

IR10 form (2016)

Inland Revenue
Te Tari Taake

Financial statements summary

IR10
March 2016

Year ended 31 March

- You only need to complete this form if you are in business.
- Please complete both pages of this form. Copy each amount from your financial statements.

Your full name			
Your IRD number	(8 digit numbers start in the second box 1 2 3 4 5 6 7 8)		
Multiple activity indicator		1	Yes / No

Profit and loss statement

Gross income from	Sales and/or services	2	$ .00
Cost of goods sold	Opening stock (include work in progress)	3	$ .00
	Purchases	4	$ .00
	Closing stock (include work in progress)	5	$ .00
Gross profit	(if a loss, put a minus sign in the last box)	6	$ .00
Other gross income	Interest received	7	$ .00
	Dividends received	8	$ .00
	Rental, lease and licence income	9	$ .00
	Other income	10	$ .00
Total income	Add up all income entered in **Boxes 6 to 10** (if a loss, put a minus sign in the last box)	11	$ .00
Expenses (as per financial statements)	Bad debts	12	$ .00
	Accounting depreciation and amortisation	13	$ .00
	Insurance (exclude ACC levies)	14	$ .00
	Interest expense	15	$ .00
	Professional and consulting fees	16	$ .00
	Rates	17	$ .00
	Rental, lease and licence payments	18	$ .00
	Repairs and maintenance	19	$ .00
	Research and development	20	$ .00
	Associated persons' remuneration	21	$ .00
	Salaries and wages paid to employees	22	$ .00
	Contractor and sub-contractor payments	23	$ .00
	Other expenses	24	$ .00
Total expenses	Add up all expenses entered in **Boxes 12 to 24**	25	$ .00
Exceptional items	(if there is a negative amount, put a minus sign in the last box)	26	$ .00
Net profit/loss before tax	**Box 11** less **Box 25**, add **Box 26** (if positive) or deduct **Box 26** (if negative)	27	$ .00
Tax adjustments	(if there is a negative amount, put a minus sign in the last box)	28	$ .00
Current year taxable profit/loss	(if a loss, put a minus sign in the last box)	29	$ .00

Balance sheet items

Current assets (as at balance date)	Accounts receivable (debtors)	30	$ .00
	Cash and deposits	31	$ .00
	Other current assets	32	$ .00
Fixed assets (closing accounting value)	Vehicles	33	$ .00
	Plant and machinery	34	$ .00
	Furniture and fittings	35	$ .00
	Land	36	$ .00
	Buildings	37	$ .00
	Other fixed assets	38	$ .00
Other non-current assets (as at balance date)	Intangibles	39	$ .00
	Shares/ownership interests	40	$ .00
	Term deposits	41	$ .00
	Other non-current assets	42	$ .00
Total assets	Add up all assets entered in **Boxes 30 to 42**	43	$ .00
Current liabilities (as at balance date)	Provisions	44	$ .00
	Accounts payable (creditors)	45	$ .00
	Current loans	46	$ .00
	Other current liabilities	47	$ .00
Total current liabilities	Add up all liabilities entered in **Boxes 44 to 47**	48	$ .00
Non-current liabilities (as at balance date)		49	$ .00
Total liabilities	Add **Box 48** to **Box 49**. Print your answer in **Box 50**	50	$ .00
Owners' equity	(if in debit, put a minus sign in the last box)	51	$ .00

Other information

Tax depreciation	52	$ .00
Untaxed realised gains/receipts	53	$ .00
Additions to fixed assets	54	$ .00
Disposals of fixed assets	55	$ .00
Dividends paid	56	$ .00
Drawings	57	$ .00
Current account year-end balances (if in debit, put a minus sign in the last box)	58	$ .00
Tax-deductible loss on disposal of fixed assets	59	$ .00

Cut off this page and staple it to page 3 of your return. You do not need to send in your financial statements as well.

2

Changes to IR10 form & guide (Fabling 2016)

Variable(s)	Summary of change(s) made
Expense categories; fixed asset components	Change from tax-deductable (tax) to financial accounts (accounting) reporting basis. Depreciation expense is now collected on both bases, with the accounting variable including amortisation (of intangible assets)
Expenses (entertainment; FBT; legal; travel & accom.; vehicle)	No longer collected
Expenses (professional & consulting fees; related party remuneration)	New variables added
Rental & lease income/expenses	Changed to include licensing income/expenditure (includes franchise fees & royalties)
"Other" categories (income; expenses; assets; current liabilities; non-current assets)	Scope changed by changes to separately itemised components. Other income now excludes exceptional items (now reported separately)
Exceptional items	New variable added
Salaries & wages	Definition changed to include ACC & super (previously in other expenses). Working proprietor remuneration removed (now in related party remuneration)
Purchases	Instruction in guide makes clear that direct labour costs should be included, as in an accounting COGS calculation
Dividends received	Inter-group dividends now included
Net profit/loss before tax; tax adjustment	New variables added. The latter reflects the overall difference between accounting profit (before tax) and taxable profit
Other current assets/liabilities; total assets/liabilities	Now includes owners' current account
Total current assets; total fixed assets	Totals no longer collected
Land & buildings	Now two separate variables
Other assets (preference shares; shares & debentures)	No longer collected. Replaced with shares/ownership interests, which is broader than the two dropped categories (eg, includes interests in partnerships, JVs & trusts)
Current liability (provisions)	New variable added
Current liability (bank accounts)	No longer collected. Replaced with current loans, which is broader
Current loans; untaxed realised gains/receipts; additions to/disposals of fixed assets	New variables added. Additions and disposals are two separate variables, but not asset specific
Capital gain on disposal of fixed assets	No longer collected. Are a subset of untaxed realised gains
Dividends paid	Change in coverage (now excludes proposed, but not paid, and includes non-cash dividends)
Indicators (are accounts GST-exclusive?; for a period of 12 months?)	No longer collected

Appendix B – Codebooks for tables stored on IBULDD_Research_Datalab

List of tables

Permanent enterprise (PENT) characteristics

PENT relationship to enterprises – pent_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	enterprise_nbr	char(10)	N	Associated enterprise numbers on the BR in format "EN" followed by 8-digit number
	start_month	int	N	First month that enterprise_nbr is current enterprise in pent (in format YYYYMM)
	end_month	int	N	Last month that enterprise_nbr is current enterprise in pent (in format YYYYMM)

See Fabling (2011) for permanent enterprise number (PENT) methodology.

Permanent balance date, industry & sector – pent_bal_ind_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
	balance_month_nbr	tinyint	N	Permanent balance month (final month in tax/financial year, usually March= 3)
	bal_method	varchar(7)	N	Method† used to determine permanent balance month
	anz06_4d	char(5)	Y	Permanent 4-digit ANZSIC'06 industry (NULL if ANZSIC'96 & ANZSIC'06 always NULL on BR)
	anz06_method	varchar(7)	Y	Method† used to determine permanent industry (NULL if ANZSIC'06 NULL)
	imp_link_strength	float	Y	Proportion of single industry firms that have modal mapping from ANZSIC'96 to ANZSIC'06 (NULL if anz06_method not "imputed")
	nzsioc_lvl3	char(4)	Y	NZSIOC (level 3), derived from ANZSIC'06 (NULL if ANZSIC'06 NULL)
	pf_ind	varchar(4)	Y	Production function industry, derived from NZSIOC (NULL if ANZSIC'06 NULL, or firm in non-productivity industry)
	always_private_for_profit	tinyint	Y	Indicator= 1 if firm is always private-for-profit (= 0 otherwise, NULL if institutional sector or business type ever NULL on BR)

† Where raw data exist, the **method** – ordered from first to last tiebreaker – for choosing permanent balance date/industry is: "one" – only one available; "emp"/"fte" – predominant based on employment share (headcount for balance date & FTE for industry); "emp mth"/"fte mth" – predominant based on number of employing months amongst previously tied; "act mth" – predominant based on number of active (on BR) months amongst previously tied; "last" – most recently observed amongst previously tied. Missing balance date is "imputed" to the mode (March). Missing ANZSIC'06 is "imputed" from ANZSIC'96 (using the modal mapping for single industry firms) when ANZSIC'96 is available. Industry imputation largely affects firms that ceased activity (as measured on the BR) prior to the implementation of ANZSIC'06 and, therefore, were not dual-coded by Stats NZ (Fabling and Sanderson 2016).

Apportioned GST return

Monthly apportioned GST – pent_month_GST_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	dim_month_key	int	N	Calender month (in format YYYYMM)
	dim_year_key	int	N	Tax/financial year containing dim_month_key (aligned to closest March year, in format YYYY03)
	sales	float	N	Total GST sales (GST-inclusive, apportioned to month)
	zero	float	N	Total zero-rated GST sales (apportioned to month)
	purch	float	N	Total GST purchases (GST-inclusive, apportioned to month)
	gst_on_sales_ex_adj	float	N	Total GST on GST sales (excluding adjustments, apportioned to month)
	gst_on_purch_ex_adj	float	N	Total GST on GST purchases (excluding adjustments, apportioned to month)
	min_gst_freq_impute	varchar(7)	Y	Minimum of method† used to impute GST frequency (NULL if GST frequency is unimputed)
	max_gst_freq_impute	varchar(7)	Y	Maximum of method† used to impute GST frequency (NULL if GST frequency is unimputed)

† **Method** – ordered from first to last priority rule – for determining missing GST frequency is: "one" – only one used; "one adj" – only one adjacent frequency; "t adj" – multiple adjacent but timing between prior and current filing matches one of those frequencies; "min adj" – minimum periodicity of adjacent frequencies; "t noadj" – frequency never supplied, but timing of filing implies frequency; "assumeM" – assumed to be monthly. The minimum and maximum method do not hold special significance – both are included to alert users to the potential that multiple methods have been used (ie, when minimum and maximum differ).

Annual apportioned GST – pent_year_GST_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	dim_year_key	int	N	Tax/financial year (aligned to closest March year, in format YYYY03)
	sales	float	N	Total GST sales (GST-inclusive, apportioned to month)
	zero	float	N	Total zero-rated GST sales (apportioned to month)
	purch	float	N	Total GST purchases (GST-inclusive, apportioned to month)
	gst_on_sales_ex_adj	float	N	Total GST on GST sales (excluding adjustments, apportioned to month)
	gst_on_purch_ex_adj	float	N	Total GST on GST purchases (excluding adjustments, apportioned to month)
	min_gst_freq_impute	varchar(7)	Y	Minimum of method† used to impute GST frequency (NULL if GST frequency is unimputed)
	max_gst_freq_impute	varchar(7)	Y	Maximum of method† used to impute GST frequency (NULL if GST frequency is unimputed)

† **Method** – ordered from first to last priority rule – for determining missing GST frequency is: "one" – only one used; "one adj" – only one adjacent frequency; "t adj" – multiple adjacent but timing between prior and current filing matches one of those frequencies; "min adj" – minimum periodicity of adjacent frequencies; "t noadj" – frequency never supplied, but timing of filing implies frequency; "assumeM" – assumed to be monthly. The minimum and maximum method do not hold special significance – both are included to alert users to the potential that multiple methods have been used (ie, when minimum and maximum differ).

Reference tables

Productivity component deflators – ppi_prod_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	nzsioc_lvl3	char(4)	N	Production function industry, derived from NZSIOC
⋆	dim_year_key	int	N	March balance month tax/financial year (in format YYYY03)
	ppii	float	N	Producer Price Index (Inputs) used to deflate M (average of quarterly official industry series, normalised to 100 in final productivity year)
	ppio	float	N	Producer Price Index (Outputs) used to deflate Y (average of quarterly official industry series, normalised to 100 in final productivity year)
	ppik	float	N	Producer Price Index (Capital) used to deflate K (average of quarterly official asset series weighted by industry asset shares, normalised to 100 in final productivity year)

Labour tables

Plant-level monthly employment – pent_pbn_month_L_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	pbn_nbr	char(10)	N	Permanent business number (plant id) in format "PB" or "PX" followed by 8-digit number
⋆	dim_month_key	int	N	Calender month (in format YYYYMM)
	dim_year_key	int	N	Tax/financial year containing dim_month_key (aligned to closest March year, in format YYYY03)
	fte	float	N	Total FTE employment in month
	employee_count	int	N	Headcount of employees during month

See Fabling and Maré (2015a) for FTE methodology. A small number of workers appear to be employed in multiple plants in the same firm in the same month. They are allocated to the minimum PBN number plant in that month to avoid double-counting in the headcount (if aggregated to the PENT level).

Firm-level annual employment – pent_year_L_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	dim_year_key	int	N	Tax/financial year (aligned to closest March year, in format YYYY03)
	fte	float	N	Average monthly FTE employment in dim_year_key
	wp	float	N	Total WP labour input in dim_year_key, *unadjusted* for WP_unknown_trans
	WP_unknown_trans	float	N	Total WP labour input potentially overestimated due to unobservable WP exit transitions (final year only)
	rme_no_WP	float	N	Average monthly employee headcount in dim_year_key (rolling mean employment)
	total_gross_earn	decimal(13,2)	N	Total EMS gross earnings in dim_year_key
	ffe	float	Y	Firm fixed effect (FE) from two-way wage FE model (NULL for firms never employing worker with age & sex in estimation period)
	fe_group	int	Y	Group id for firms connected by worker transitions (NULL for firms never employing worker with age & sex in estimation period)
	fte_with_wfe	float	N	Average monthly FTE employment of workers with FE (FE not estimated for workers with NULL age or sex)
	avg_wfe	float	Y	Average (FTE-weighted) worker FE from two-way wage FE model (NULL for firms with fte_with_wfe= 0)
	avg_xb	float	Y	Average (FTE-weighted) worker observables component (sex-specific age profile) from two-way wage FE model (NULL for firms with fte_with_wfe= 0)

See Fabling and Maré (2015a) for FTE/WP methodology, and Maré et al. (2017) for two-way wage fixed effects methodology.

Productivity tables

Productivity population employment – pent_prod_pop_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	dim_year_key	int	N	Tax/financial year (aligned to closest March year, in format YYYY03)
	active_prior_year	tinyint	N	Indicator= 1 if firm active in prior tax/financial year (0 otherwise)
	active_next_year	tinyint	N	Indicator= 1 if firm active in next tax/financial year (0 otherwise)
	pf_ind	varchar(4)	N	Production function industry, derived from NZSIOC
	fte	float	N	Average monthly FTE employment in dim_year_key
	wp_adj	float	N	Total WP labour input in dim_year_key, (*adjusted* for final year exit transitions)

See Fabling and Maré (2015a) for FTE/WP methodology. Final year WP count adjusted for firm-level exit where individual-level WP exit is unobservable.

Productivity dataset – pent_prod_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pent	char(10)	N	Permanent enterprise number (firm id) in format "EN" followed by 8-digit number
⋆	dim_year_key	int	N	Tax/financial year (aligned to closest March year, in format YYYY03)
	active_prior_year	tinyint	N	Indicator= 1 if firm active in prior tax/financial year (0 otherwise)
	active_next_year	tinyint	N	Indicator= 1 if firm active in next tax/financial year (0 otherwise)
	pf_ind	varchar(4)	N	Production function industry, derived from NZSIOC
	size_stratum	varchar(17)	N	Size stratum† for population-weighting
	pop_weight	float	N	Productivity population weight (constant within dim_year_key, pf_ind, size_stratum)
	fte	float	N	Average monthly FTE employment in dim_year_key
	wp_adj	float	N	Total WP labour input in dim_year_key, (*adjusted* for final year exit transitions)
	go_nom	bigint	N	Nominal gross output (Y)
	M_nom	bigint	N	Nominal intermediate consumption (M)
	K_nom	bigint	N	Nominal capital services (K)
	lngo_real	float	Y	Real $\ln Y$ (NULL if $Y = 0$)
	lnM_real	float	Y	Real $\ln M$ (NULL if $M = 0$)
	lnK_real	float	Y	Real $\ln K$ (NULL if $K = 0$)
	lnL	float	Y	$\ln L = \ln($ fte + wp$)$
	mfp_go_cd	float	Y	Estimated multifactor productivity (MFP) = residual from industry-specific gross output Cobb-Douglas production function (OLS, NULL if Y, M, K zero)
	mfp_go_tl	float	Y	Estimated multifactor productivity (MFP) = residual from industry-specific gross output Translog production function (OLS, NULL if Y, M, K zero)
	mfp_go_fe	float	Y	Estimated multifactor productivity (MFP) = (residual+fixed effect) from industry-specific gross output Cobb-Douglas production function (firm fixed effects, NULL if Y, M, K zero)
	go_fe	float	Y	Estimated firm fixed effect from industry-specific gross output Cobb-Douglas production function (firm fixed effects, NULL if Y, M, K zero)

Continued on next page. . .

Continued from previous page...

Key	Variable	Format	NULL	Description
	RLR_nom	bigint	N	Nominal rental, leasing & rates expense
	abv_intang_nom	bigint	N	Nominal average of opening and closing book value (ABV) of intangible assets
	abv_vehicle_nom	bigint	N	Nominal ABV of vehicle fixed assets
	abv_PM_other_nom	bigint	N	Nominal ABV of plant, machinery & other fixed assets
	abv_furnfit_nom	bigint	N	Nominal ABV of furniture & fittings fixed assets
	abv_landbuild_nom	bigint	N	Nominal ABV of land & buildings fixed assets
	WS_RelPartyRem_nom	bigint	N	Nominal labour compensation from AES/IR10 (total remuneration for employees + related party remuneration)
	flag_in_aes	tinyint	N	Indicator= 1 if usable AES return (0 otherwise)
	flag_in_i10	tinyint	N	Indicator= 1 if usable IR10 return (0 otherwise)
	flag_i10_purch_edited	tinyint	Y	Indicator= 1 if IR10 purchases edited to reflect gross profit (0 otherwise, NULL if no usable IR10)
	flag_i10_othinc_edited	tinyint	Y	Indicator= 1 if IR10 other income set to zero because it incorrectly reported an income (sub-)total
	flag_i10_fp_inexact	tinyint	Y	Indicator= 1 if IR10 front page not internally consistent, but within 1% accuracy
	flag_i10_bp_inexact	tinyint	Y	Indicator= 1 if IR10 back page not internally consistent, but within 1% accuracy
	flag_i10_total_missing	tinyint	Y	Indicator= 1 if any IR10 totals used in testing are zero (assumed missing)
	flag_i10_gain_unadj	tinyint	Y	Indicator= 1 if IR10 gain on sale of fixed asset, but no adjustment
	flag_i10_gain_adj_othinc	tinyint	Y	Indicator= 1 if IR10 gain on sale of fixed asset, and removed from other income
	flag_i10_gain_adj_taxadj	tinyint	Y	Indicator= 1 if IR10 gain on sale of fixed asset, and removed from tax adjustment
	flag_i10_loss_unadj	tinyint	Y	Indicator= 1 if IR10 loss on sale of fixed asset, but no adjustment
	flag_i10_loss_adj_othexp	tinyint	Y	Indicator= 1 if IR10 loss on sale of fixed asset, and removed from other expenses
	flag_i10_loss_adj_taxadj	tinyint	Y	Indicator= 1 if IR10 loss on sale of fixed asset, and removed from tax adjustment

Continued on next page...

Continued from previous page. . .

Key	Variable	Format	NULL	Description
	flag_i10_WS_adj_purch	tinyint	Y	Indicator= 1 if IR10 labour costs removed from purchases
	flag_i10_WS_adj_rd	tinyint	Y	Indicator= 1 if IR10 labour costs removed from R&D expense
	flag_i10_WS_adj_othexp	tinyint	Y	Indicator= 1 if IR10 labour costs removed from other expenses
	flag_i10_gst_incl_raw	tinyint	Y	Indicator= 1 if IR10 reported as GST-inclusive
	flag_i10_gst_incl	tinyint	Y	Indicator= 1 if IR10 treated as GST-inclusive (ie, GST removed)
	flag_i10_dep_imp_TFA	tinyint	Y	Indicator= 1 if IR10 tax depreciation imputed from firm-level average tax depreciation rate in other years
	flag_i10_dep_imp_acc	tinyint	Y	Indicator= 1 if IR10 tax depreciation imputed from accounting depreciation (where no evidence that these are equivalent)
	flag_i10_taxadj_neg	tinyint	Y	Indicator= 1 if IR10 tax adjustment is negative, so M increased
	flag_i10_taxadj_pos	tinyint	Y	Indicator= 1 if IR10 tax adjustment is positive, so M decreased
	flag_i10_nolagCBV	tinyint	Y	Indicator= 1 if firm is non-entrant and lagged IR10 missing, so Average Book Value (ABV) of fixed assets set to Closing BV (ie, Opening=Closing BV)
	flag_RLR_avg	tinyint	N	Indicator= 1 if rental, leasing and rates expense imputed from firm-level average in other years
	flag_RLR_model	tinyint	N	Indicator= 1 if rental, leasing and rates expense modelled

See Fabling and Maré (2015a) for FTE/WP methodology. Final year WP count adjusted for firm-level exit where individual-level WP exit is unobservable. † Size stratum for population-weighting defined in table 13.

Production function coefficients – pent_prod_est_betas_IDI_20171020

Key	Variable	Format	NULL	Description
⋆	pf_ind	varchar(4)	N	Production function industry, derived from NZSIOC
	lnM_go_cd	float	N	Estimated coefficient on M in gross output Cobb-Douglas (GO CD) production function (OLS)
	lnL_go_cd	float	N	As above (L)
	lnK_go_cd	float	N	As above (K)
	t200203_go_cd	float	N	Estimated coefficient on 2002 tax/financial year indicator in GO CD production function (OLS, 2001 is base year)
	t200303_go_cd	float	N	As above (2003 year)
	⋮	⋮	⋮	⋮
	t201603_go_cd	float	N	As above (2016 year)
	lnM_go_cd_fe	float	N	Estimated coefficient on M in gross output Cobb-Douglas (GO CD) production function with firm fixed effects (FFE)
	lnL_go_cd_fe	float	N	As above (L)
	lnK_go_cd_fe	float	N	As above (K)
	t200203_go_cd_fe	float	N	Estimated coefficient on 2002 tax/financial year indicator in GO CD production function (FFE, 2001 is base year)
	t200303_go_cd_fe	float	N	As above (2003 year)
	⋮	⋮	⋮	⋮
	t201603_go_cd_fe	float	N	As above (2016 year)
	lnM_go_tl	float	N	Estimated coefficient on M in gross output Translog (GO TL) production function (OLS)
	lnL_go_tl	float	N	As above (L)
	lnK_go_tl	float	N	As above (K)
	lnM_x_lnM_go_tl	float	N	As above (M^2)
	lnM_x_lnL_go_tl	float	N	As above ($M \times L$)
	lnM_x_lnK_go_tl	float	N	As above ($M \times K$)
	lnL_x_lnL_go_tl	float	N	As above (L^2)
	lnL_x_lnK_go_tl	float	N	As above ($L \times K$)
	lnK_x_lnK_go_tl	float	N	As above (K^2)
	t200203_go_tl	float	N	Estimated coefficient on 2002 tax/financial year indicator in GO TL production function (OLS, 2001 is base year)
	t200303_go_tl	float	N	As above (2003 year)
	⋮	⋮	⋮	⋮
	t201603_go_tl	float	N	As above (2016 year)

All coefficients are from unweighted regressions, estimated separately for each industry (pf_ind).

Recent Motu Working Papers

All papers in the Motu Working Paper Series are available on our website www.motu.nz, or by contacting us on info@motu.org.nz or +64 4 939 4250.

19-04 Kerr, Suzi, Steffen Lippert and Edmund Lou. 2019. "Financial Transfers and Climate Cooperation."

19-03 Fabling, Richard and David C Maré. 2019. "Improved productivity measurement in New Zealand's Longitudinal Business Database." (forthcoming)

19-02 Sin, Isabelle and Judd Ormsby. 2019. "The settlement experience of Pacific migrants in New Zealand: Insights from LISNZ and the IDI"

19-01 Benjamin Davies and David C Maré. 2019. "Relatedness, Complexity and Local Growth."

18-16 Hendy, Jo, Anne-Gaelle Ausseil, Isaac Bain, Élodie Blanc, David Fleming, Joel Gibbs, Alistair Hall, Alexander Herzig, Patrick Kavanagh, Suzi Kerr, Catherine Leining, Laëtitia Leroy, Edmund Lou, Juan Monge, Andy Reisinger, Jim Risk, Tarek Soliman, Adolf Stroombergen, Levente Timar, Tony van der Weerdan, Dominic White and Christian Zammit. 2018. "Land-use modelling in New Zealand: current practice and future needs."

18-15 White, Dominic, Niven Winchester, Martin Atkins, John Ballingall, Simon Coates, Ferran de Miguel Mercader, Suzie Greenhalgh, Andrew Kerr, Suzi Kerr, Jonathan Leaver, Catherine Leining, Juan Monge, James Neale, Andrew Philpott, Vincent Smart, Adolf Stroombergen, and Kiti Suomalainen. 2018. "Energy- and multi-sector modelling of climate change mitigation in New Zealand: current practice and future needs."

18-14 Preston, Kate, David C Maré, Arthur Grimes and Stuart Donovan. 2018. "Amenities and the attractiveness of New Zealand cities."

18-13 Alimi, Omoniyi, David C Maré and Jacques Poot. 2018. "Who partners up? Educational assortative matching and the distribution of income in New Zealand."

18-12 Fabling, Richard. 2018. "Entrepreneurial beginnings: Transitions to self-employment and the creation of jobs."

18-11 Fleming, David A and Kate Preston. 2018. "International agricultural mitigation research and the impacts and value of two SLMACC research projects." (also a Ministry for Primary Industries publication)

18-10 Hyslop, Dean and David Rea. 2018. "Do housing allowances increase rents? Evidence from a discrete policy change."

18-09 Fleming, David A., Ilan Noy, Jacob Pástor-Paz and Sally Owen. 2018. "Public insurance and climate change (part one): Past trends in weather-related insurance in New Zealand."

18-08 Sin, Isabelle, Kabir Dasgupta and Gail Pacheco. 2018. "Parenthood and labour market outcomes." (also a Ministry for Women Report)

18-07 Grimes, Arthur and Dennis Wesselbaum. 2018. "Moving towards happiness."

18-06 Qasim, Mubashir and Arthur Grimes. 2018. "Sustainable economic policy and well-being: The relationship between adjusted net savings and subjective well-being."

18-05 Clay, K Chad, Ryan Bakker, Anne-Marie Brook, Daniel W Hill Jr and Amanda Murdie. 2018. "HRMI Civil and Political Rights Metrics: 2018 Technical Note."

18-04 Apatov, Eyal, Nathan Chappell and Arthur Grimes. 2018. "Is internet on the right track? The digital divide, path dependence, and the rollout of New Zealand's ultra-fast broadband." (forthcoming)

18-03 Sin, Isabelle, Eyal Apatov and David C Maré. 2018. "How did removing student allowances for postgraduate study affect students' choices?"

18-02 Jaffe, Adam B and Nathan Chappell. 2018. "Worker flows, entry, and productivity in New Zealand's construction industry."

18-01 Harris, Richard and Trinh Le. 2018. "Absorptive capacity in New Zealand firms: Measurement and importance."

Motu