Motu Working Paper 24-06

Still medalling: Productivity gets a bronze (data source)

Richard Fabling

October 2024

Document information

Author contact details

Richard Fabling
Senior Research Associate, Motu Research
richard.fabling@xtra.co.nz

Acknowledgements

The bulk of this work was undertaken while Richard was an employee of the New Zealand Productivity Commission, and would not have been possible without that organisation's strong commitment to productivity research. Richard thanks Amelia Guha for able research assistance, and Dave Maré for feedback on the right-of-use adjustment process (described in appendix A).

Disclaimer

These results are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI) and Longitudinal Business Database (LBD) which are carefully managed by Stats NZ. For more information about the IDI please visit https://www.stats.govt.nz/integrated-data/.
The results are based in part on tax data supplied by Inland Revenue to Stats NZ under the Tax Administration Act 1994 for statistical purposes. 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. The results presented in this study are the work of the author, not Stats NZ or individual data suppliers.

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

Productivity data is missing from the Longitudinal Business Database (LBD) for over a third of firm-year observations in “measured sector” industries, equating to a fifth of total labour in those industries. We develop a method to fill these data gaps using an additional (third) data source – firm-level annualised goods and services tax (GST) returns. Coupled with additional modelling using full-coverage employment information, the resulting “complete” productivity dataset provides additional avenues for researchers to test the robustness of their results to the inclusion of firm types previously underrepresented in the productivity data – particularly new and owner-operated firms.

JEL codes

D20; D24

Keywords

Longitudinal Business Database; administrative data; productivity

Summary haiku

Missing-at-random
would be a wonderful thing,
but it just ain’t so

1 Motivation

Productivity growth is an important determinant of output growth in the economy and, therefore, aggregate improvements in income and well-being. The proliferation of high quality worker- and firm-level data, places empirical microeconomics at the forefront of understanding the dynamics of productivity growth and identifying obstacles to higher productivity (Syverson 2011; De Loecker and Syverson 2021).

In New Zealand, such work is enabled by the Longitudinal Business Database (LBD) and the Fabling-Maré firm-level labour and productivity datasets (Fabling 2011; Fabling and Maré 2015a, 2015b, 2019). Over the last decade, these data have been combined with other LBD components to study a range of important productivity-related topics,[1] including:

- Exporting/foreign-ownership (Fabling and Sanderson 2013, 2014a);
- Rent-sharing/wages (Criscuolo et al. 2020; Allan and Maré 2021, 2022; Sin et al. 2022; Fabling and Maré 2024);
- Market concentration/spillovers (Maré and Graham 2013; Conway et al. 2015; Zheng 2016; Maré 2016; Maré and Fabling 2019);
- Business practices/intangibles (Fabling and Grimes 2014, 2021; Chappell and Jaffe 2018; Fabling 2021a);
- Productivity dispersion (Fabling and Sanderson 2014b; Meehan 2020);
- Worker skills (Maré and Fabling 2013; Maré, Hyslop, and Fabling 2017; Maré et al. 2017; Fabling et al. 2022; Kirker and Sanderson 2022); and
- Industry studies (Apatov et al. 2015; Jaffe et al. 2016; Chappell et al. 2018).

These analyses face a common issue – incomplete productivity data – arising from two main sources: the absence of survey responses (due to non-response or non-sampling) and tax returns (due to alternative filing methods); and the removal of low-quality (partial, inconsistent or implausible) observations.[2] As these mechanisms might imply, productivity data aren't *missing-at-random*, raising the spectre of bias in any statistics produced from them.

Table 1 illustrates two potential sources of bias – differences in firm-year coverage by industry and by firm size – of Annual Enterprise Survey (AES) and IR10 tax returns, which we label "tier 1" productivity data. For firm size, the table reports productivity coverage rates separately for three non-overlapping groups: working-proprietor-only (WP-only) firms; micro employers; and larger employers.[3] At the population level (bottom row), WP-only firms have a 60.8% cover-

[1] Fabling and Sanderson (2016) describe the LBD components and provide tips for new users.

[2] Misidentification of the population may also be important, though the extent of this issue is difficult to quantify. In particular, since labour input is hard to accurately measure in working-proprietor-only firms, some firms may be classified as active when they are not. We ignore this issue, noting that identifying the population better might result in a lower *or* higher rate of missingness.

[3] The majority of WP-only firms have one or two active working proprietors.

age rate, which is nine percentage points (pp) lower than the rate for micro employers and 17pp lower than the rate for larger employers. The overall productivity coverage rate is 66%, reflecting the relatively high proportion of micro enterprises in the New Zealand private sector (compared to other OECD countries). Failure to account for variation in coverage by firm size may, therefore, lead to statistics that are less representative of micro enterprises and more representative of large firms than they should be given the population distribution.

Focussing on the three largest (by observations) industry divisions – agriculture, construction and professional services – the coverage rate for the latter is 59.8%, which is 6pp below the overall mean rate, while agriculture sits at the overall rate and construction is 2pp above that rate. These gaps are partially attributable to differences in the firm size distribution across industries. However, differences in the WP-only firm coverage rate across industries suggest an independent role for industry in coverage, which may reflect variation in tax filing norms, or unequal effort in statistical collections across industries.[4]

Approaches of varying complexity have developed to address missing data bias. At the minimal effort end of the scale, the raw sample is used acknowledging that resulting statistics represent a non-random sample, not the population. This approach is common when productivity data is linked to another source within the LBD. Fabling and Grimes (2014) is a representative example of this approach, where Business Operations Survey (BOS) responses are linked to productivity data. BOS is a sample survey with full coverage of the largest firms in the economy, no sampling of firms below six employees, and random sampling of other firms. The BOS small business exclusion increases the match rate of survey and productivity data, while imposing that the analysis is only of larger firms, not micro-enterprises.

Maré et al. (2017) take an alternative approach of weighting statistics by (predicted) output. This approach reflects their desire to capture the relative contribution of firm types to aggregates. "Importance-weighting" in this manner reduces the potential impact of coverage on statistical bias, since relatively high weights are assigned to firms that are relatively more likely to have productivity data. Maré et al. (2017) also explicitly track firms that join and leave the estimation sample due to changes in firm-level filing patterns over time.

A third approach is cell-based weighting using the inverse of the productivity coverage rate (ie, assuming that, within cell, unobserved firm data is similar to observed firm data). Fabling and Maré (2019) show how firm size by industry cell weights can be used to generate macroeconomic aggregates from the productivity data that approximately mimic growth rates in comparable official statistics. While their cells include separate groupings for entering and exiting firms, the weights are cross-sectional and not explicitly designed for longitudinal analysis.

This paper introduces a new approach that yields a full-coverage productivity dataset by adding a third ("bronze") input data source. The method hinges on the presence of Goods and Services Tax (GST) data in the LBD, which is substantially closer to full coverage than the productivity data because the New Zealand GST system is broad-based with mandatory filing above a (time-

[4]For example, agriculture is excluded from the AES sample.

varying) threshold. Table 2 illustrates this coverage for productivity population observations where tier 1 data are missing – ie, for the gaps we fill using GST data.[5] The overall coverage rate is 83.4%, with similarly high coverage in all industries except finance and insurance, which is lower because financial services do not attract GST. GST coverage rates are lower for WP-only firms, but still average 77.9%.

Figure 1 plots the coverage rates reported in tables 1 and 2, with bubble size reflecting the gap left by tier 1 productivity data. The five largest gaps (shaded bubbles) relate to WP-only or micro employers in the largest industries where, even in the worst case scenario of WP-only construction firms, GST data is present for 77.5% of the missing tier 1 productivity observations. Figure 1 also demonstrates an unweighted correlation between coverage rates (dashed line), which may reflect common reasons why firms don't have linked tax records (eg, activity below the mandatory filing threshold, inactive firms, incorrect tax id links to the Business Register).

While sales and purchases are not equivalent to gross output and intermediate consumption, there is a pre-LBD history of using net (sales less purchases) GST as a proxy for value-added (Law and McLellan 2005; Maré and Timmins 2006; Grimes et al. 2012). Our main innovation, therefore, is to apply a hierarchy of industry-specific regression specifications to integrate the GST data into the tier 1 productivity data. We have no additional administrative source for capital services, but we expect this productivity component to be the most stable and, therefore, for firm fixed effects to capture a significant proportion of the variation across firms. With that expectation, we also treat capital services as a function of GST components. The remaining gap in productivity coverage is small, and we close it using models based on labour input. "Completing" the productivity dataset in this way, enables simple testing for statistical bias from excluded observations and, subject to a case-by-case assessment of suitability, alternative population statistics.

Sections 2 provides a brief description of the productivity and GST data, while section 3 describes the method for integrating the latter into the former. Section 4 summarises and touches on possible extensions.

2 Data

This section briefly summarise the Fabling-Maré productivity tables and the production function variables available. The construction of these data is covered in detail elsewhere, and LBD users are encouraged to read those earlier papers (Fabling 2011; Fabling and Maré 2015a, 2015b, 2019) prior to using the productivity dataset.[6]

[5]GST coverage is defined as non-zero sales and/or purchases within the year. Filed returns where sales and purchases are both zero are excluded.

[6]Results in this paper are based on the 202310 instance of the tables. Innovations that have been added to the production process since 2019, including adjustments for recent right-of-use accounting changes for large firms, are summarised in appendix A.

The productivity tables currently include twenty-two years of data covering 2001-2022.[7] The productivity population is defined by industry and employment for private-for-profit firms. Following Stats NZ official statistics, we exclude industries where prices are hard to measure due to the dominance of non-market providers (government administration, health and education sectors). Since we define industry and private-for-profit status as permanent business characteristics, entry and exit from the population is due entirely to transitions in and out of having positive labour input (L).

L is calculated from full-coverage monthly payroll (Employer Monthly Schedule) tax returns for employees (adjusted for part-time work, multiple job holding and minimum wages) combined with working proprietor (WP) counts derived from annual tax returns specific to business type (sole trader, partnership and company).[8]

Other tier 1 productivity components – gross output (Y), intermediate consumption (M), capital services (K) – are derived from combining Annual Enterprise Survey (AES) and administrative tax returns (IR10s), and are subject to missingness. According to Fabling and Maré (2019), the absence of both AES and IR10 data accounts for almost 70% of missing observations, due to AES being a sample survey with higher sampling rates for large firms, and IR10s being only one of several methods for firms to file mandatory tax returns. The other 30% of missingness is due to observations being dropped due to tax returns being incomplete or internally inconsistent. Fabling and Maré (2019) find that incomplete returns is the primary issue, with internal inconsistency becoming less of an issue over time, coinciding with changes in the IR10 form and the introduction of electronic filing.

AES and IR10 data are harmonised via a series of adjustments at the firm and productivity industry levels.[9] Through harmonization, gross output and intermediate consumption adhere to national accounting definitions, except that we choose to exclude rental, leasing and rates costs (RLR) from M. We then calculate K as the combination of RLR, depreciation, and a cost of borrowing based on average book values of the capital stock. This aggregation is non-standard, but allows for the equal treatment of owned and rented capital.

GST returns provide sales (S) and purchases (P) and, like labour input, GST is approximately full coverage since the tax has few exclusions (the main one being financial services), filing is mandatory following GST registration, and registration is used to populate Stats NZ's Business Register, meaning that firms in the productivity population are likely to be GST-registered.[10] GST returns are apportioned to the appropriate financial year (see Appendix A) and then deflated

[7] Productivity years use financial year ends (31 March for most firms). For example, the 2020 productivity year is mostly in the 2019 calendar year (ie, 2019Q2 to 2020Q1).

[8] The three WP tax returns (IR3, IR4, IR7) report taxable profit. Any benefit that might be gained from incorporating this profit measure into the estimation of productivity components is reduced by: partnership and sole proprietor returns being individual-level data (ie, in the IDI, not the LBD); net profit being less informative than gross profit components; and, a significant proportion of firms having negative taxable profit, complicating their inclusion in a log model.

[9] There are 39 productivity industries, grouped based on the availability of official price indices and having sufficient sample size to support the adjustment process.

[10] Employing for the first time is also a trigger for inclusion on the Business Register.

following the same approach as productivity components.

We expect $S \neq Y$ and $P \neq M$ for both conceptual and data processing reasons, primarily:

- Y includes changes in stocks of finished goods, whereas S only counts goods sold
- P includes expenses that are excluded from M
- S and P may be more volatile due to one-off items that would be itemised separately (and excluded from productivity components) in a complete set of financial accounts
- Y and M are cleaned as part of the productivity dataset processing, including industry adjustment to improve the consistency of AES- and IR10-based observations

Table 3 tests these expectations, reporting means and standard deviations for GST and tier 1 productivity variables, plus the capital-labour ratio and labour productivity derived using productivity components ($v_{ym} - l$) and GST sales and purchases ($v_{sp} - l$). All variables are in natural logs (denoted by the use of lower case) and the table includes levels as well as year-on-year growth rates (log changes). These statistics are estimated on the full sample, meaning that coverage differs across variables, as well as across levels and changes for the same variable.

Bearing that caveat in mind, the means of gross output and GST sales are within one log point of each other, while mean GST purchases is 24 log points higher than intermediate consumption. As a consequence, mean labour productivity is higher when calculated with tier 1 productivity data, that when calculated using GST data. Figure 2 reports the ratio of aggregate S (P) and tier 1 Y (M) for the common sample. This comparison suggests that sales and purchases capture substantially more than gross output and intermediate consumption do (ie, the ratios are well above one). Consistent with table 3 means, this ratio is higher for P/M than it is for S/Y. It follows that the within-source ratio of aggregate tier 1 Y to M is substantially higher than the S to P ratio for GST. Reassuringly, though, pairs of within- and across-source ratios have similar time variation. It is these systematic levels differences – driven by differences in what each component captures – that lead us towards a regression-based harmonisation approach, and it is the signal present in the time variation of S and P that makes this effort worthwhile.

Another important difference between GST and tier 1 productivity components is the higher standard deviation in both levels and growth rates that are apparent in the GST data, which likely stem from the inclusion of infrequent (non-productivity) income and expense items in the GST returns, the greater coverage of micro-enterprise in the GST data, and the attribution of "lumpy" activity across time periods. Our goal is to maintain a high quality standard in the productivity data, and these explanations suggest measurement error is an issue for GST data. To account for the extremes of this issue, we clean the GST data using the same methodology used for tier 1 productivity components (Fabling and Maré 2019), removing one-off observations that imply implausibly large log changes in either sales or purchases.[11]

The two rightmost columns of table 3 report summary statistics for growth rates showing that

[11] Statistics in table 3 reflect GST components prior to data cleaning, whereas tier 1 productivity components have already been subjected to cleaning.

mean growth rate in p also deviates substantially from that for m – with the former negative (-4.1%) and the latter positive (1.0%). Our subsequent testing focuses on the consistency of changes in productivity components, both within the non-tier 1 data, but also when a firm transitions into and out of tier 1.

The correlations in table 4 strongly suggest that sales and purchases are closely related to output and intermediate consumption at the firm level, both in levels and in growth rates. Focussing on year-on-year changes, Δy has a correlation of 0.67 with Δs and 0.48 with Δp, and Δm has a correlation of 0.47 with Δs and 0.55 with Δp. These correlations suggest that s and p (which have a correlation of 0.54 in growth rates) might be used together to derive productivity component.

A further reason to include both s and p in our empirical estimates is so that models can accommodate instances where one of the GST components is zero. Table 5 and figures 3 and 4 summarise the properties of y and m when GST components are zero. Focussing first on the observation count column of table 5, instances of reported zero GST are less frequent than non-reporting of GST. However, in both cases, firms are less likely to have tier 1 productivity data (see the following column), than firms with $S > 0$ and $P > 0$, particularly when $P = 0$. Conditional on observing productivity data, zero S or zero P are not very good indicators of either Y or M also being zero (the two rightmost columns). At best, observed output is zero in 52% of cases where we observe tier 1 productivity and GST sales are zero, but GST purchases isn't.[12] The top row of table 5 shows that not having a GST return is very rarely associated with zero output and intermediate consumption, which is plausible when we consider that firms in the population have non-zero labour input.

Outside of the finance sector, the most likely explanation for non-reporting of GST components is that a firm is currently below the threshold for mandatory filing. Figures 3 and 4 show the distribution of tier 1 y and m respectively when the corresponding GST component is zero, non-zero or not present. The distribution of no-GST observations (dotted lines) confirm the view that low sales is a plausible explanation of non-filing, with most of the observed density of output below the (current) mandatory filing threshold of $s = 11$ (ie, turnover of $60,000, indicated by the vertical line). The distribution being closely aligned to the threshold, rather than distributed more broadly below the threshold is likely due to the fact that firms are not usually added to the Business Register until they have registered for GST.

Consistent with the importance of the GST threshold, when sales are non-zero (solid line), y is usually above the filing threshold. The $S = 0$ and $P = 0$ (dashed) lines in figures 3 and 4 reinforce the view that reported GST zeros don't translate consistently to tier 1 productivity component values. For both figures, the distribution of productivity components is largely below the mandatory filing threshold, with significant density to the left of the no-GST group. In the case of $S = 0$, 60% of observations have $P > 0$ (table 5), whereas $S > 0$ is only true for 47% of observations when $P = 0$, which may reflect a bias towards continuing to collect

[12] In the empirical modelling and the complete productivity dataset, we assign tier 1 zeroes a log value of zero ($1) so that these observations are not treated as missing.

and file GST when below the threshold if that would generate a GST refund (ie, $P > S$). This hypothesis is supported by the left-shifting of the zero groups, relative to the other groups.

Table 6 summarises tier 1 missingness patterns over three consecutive years. Since not all firms are present in all three years, these patterns are distinguished into distinct firm continuity groups at time t. The top panel looks at “ongoing” firms (ie, that are also in the population at $t-1$ and $t+1$), with the next three panels looking at entering firms that continue, exiting firms that were formerly continuing, and “one-year” transitory firms. Each panel divides its group by the observed pattern of tier 1 data at $t-1$ and $t+1$ (leftmost two columns) and then shows group size, the probability of tier 1 data being present at t, and the contribution of the group to overall tier 1 productivity missingness.

Ongoing firms are the dominant firm type and demonstrate persistence in tier 1 data supply, which is likely driven by persistence in tax filing methods and in the quality of those filings. Firms that have tier 1 productivity at both $t-1$ and $t+1$ have a 93% probability of also having it at time t. As a consequence, one-year gaps in ongoing firm filing only account for 9% of overall missingness. The flipside of this persistence is that ongoing firms that don’t have tier 1 productivity over the three-year period account for 43% of missingness. For continuing firms, switchers into or out of the tier 1 productivity dataset make up a combined 16% of missing observations.

Entering and exiting firms display weaker persistence, perhaps reflecting changes to firm operations during transitions and/or the potential misidentification of entry/exit years. Entrants and exiters account for 28% of missingness, primarily because of persistent non-filing (19pp) and the absence of productivity data in the (expected) exit year (6pp). Finally, one-year firms have a 39% probability of having tier 1 data in that year, which is similar to the overall rate for exiting firms (44%). Their small group size, however, means that they only account for 4% of overall missingness.

The persistence of presence in the tier 1 productivity data is summarised in the bottom three rows of table 6 – nearly two thirds of missing observations are associated with firms that don’t have observations in an adjacent year (either because of non-filing, input data quality concerns, or because the firm was inactive).

3 Integrating GST data into the productivity dataset

Converting GST sales and purchases into output- and intermediate consumption-like components relies on the dominant group of firms that have both GST and tier 1 data. For firms that transition between having and not having tier 1 data, firm fixed effects improve the integration of the GST data. Specifically, for firm i at time t, we estimate the following regressions separately for each production function industry (separable into $j(i)$ detailed industries).

Where firms have at least one non-zero GST component, we estimate:

$$z_{it} = \sum_{\tau} \delta(t=\tau)\left[\alpha_\tau + \beta^s_\tau s_{it} + \beta^p_\tau p_{it}\right] + \delta(P_{it}=0)\left[\lambda^s + \gamma^s s_{it}\right] + \delta(S_{it}=0)\left[\lambda^p + \gamma^p p_{it}\right] + \eta_{it} \tag{1}$$

$$\eta_{it} = \begin{cases} \delta_i + \epsilon_{it} & \text{(2a)} \\ \delta_{j(i)} + \epsilon_{it} & \text{(2b)} \end{cases}$$

where $z_{it} \in (y_{it}, m_{it}, k_{it})$; δ_i is a set of firm fixed effects; $\delta_{j(i)}$ is four-digit ANZSIC industry controls; and ϵ_{it} is an independent and identically distributed (iid) error term. Firm fixed effects is the preferred model, but only applies to firms where tier 1 $y/m/k$ is observed contemporaneously with s/p. The detailed industry control model, therefore, applies to firms that never have tier 1 productivity observations at the same time as they have non-zero GST. The summation term in equation (1) allows the relationship between tier 1 and GST components to vary by year (within productivity industry). Variation in that relationship may result from real world changes over the business cycle in the proportion of sales/purchases that should be included in output/intermediate consumption. Industry-year level adjustments made to improve the consistency of AES and IR10 productivity components may also induce time variation in the relationship with GST components. The remaining terms allow the coefficients on sales/purchases to differ in the special case where the other GST component is zero.[13]

Appendix B (table 10) reports estimated coefficients from a pooled industry version of equation (1) where we impose time-invariant coefficients on s and p.[14] The purpose of the appendix is to provide evidence of proof-of-concept rather than coefficients that relate directly to the integration of GST-based observations. Four points are noteworthy – first, adjusted R^2s are high, reflecting the strong relationship between tier 1 productivity and GST components; second, the inclusion of s in regressions for m seems more valuable than the inclusion of p in regressions for y, probably reflecting the greater opportunity for GST to capture non-productivity components in purchases compared to sales; third, that the absence of the primary GST component is compensated for by a higher coefficient on the other GST component for y and m; and, finally, the model for k has weaker explanatory power than is the case for other components, but is bolstered by the inclusion of industry controls and firm fixed effects.[15]

Where we don't have at least one non-zero GST component, we estimate:

$$z_{it} = \sum_{\tau} \delta(t=\tau)\left[\alpha_\tau + \beta_\tau l_{it}\right] + \delta(L_{t-1}=0)\left[\alpha_{\text{entry}} + \beta_{\text{entry}} l_{it}\right] + \delta(L_{t+1}=0)\left[\alpha_{\text{exit}} + \beta_{\text{exit}} l_{it}\right] + \delta(\text{WP}>0 \wedge \text{FTE}>0)\left[\alpha_{\text{mix}} + \beta_{\text{mix}} l_{it}\right] + \delta(\text{WP}=0)\left[\alpha_{\text{emp}} + \beta_{\text{emp}} l_{it}\right] + \eta_{it} \tag{3}$$

[13] The zero GST component in these cases is set to log zero ($1).

[14] Figure 12 reports estimated time-varying coefficients from firm fixed effect pooled industry models for equation (1) and (3).

[15] The table note reports adjusted R^2 values for models with only year controls. Industry controls improve the R^2 by 15pp for k, 6pp for m and 3pp for y.

where $z_{it} \in (y_{it}, m_{it}, k_{it})$, η_{it} is defined as in equation (2), and the summation provides year-specific intercepts and main coefficients on l_{it}. The following terms in equation (3) allow separate intercepts and coefficients on labour input for firms: entering the population; exiting the population; with mixed labour (employees and WPs); and employee-only firms (WP-only and ongoing firms are the reference groups).

Since all firms have l, these equations are estimated over all firms with tier 1 productivity data, and can be used to estimate all three productivity components for any firm in the productivity population. The interaction terms in equation (3) allow for the relationship between z and l to differ in transition years and when working proprietor counts are non-zero, both of which cases may result in systematic mismeasurement of l, and/or "non-representative" (for the firm) output or other productivity components.

Appendix B (table 11) presents industry control and fixed effects model estimates (with all industries pooled and time-invariant coefficients). As expected, for y and m, these models do not perform as well at those based on s and p (table 10) with adjusted R^2s around 30pp lower across comparable specifications. For k, the drop in R^2 is lower (8-10pp). With industry controls, therefore, the model for k has the greatest explanatory power – though only slightly over m – likely reflecting the fact that both m and l are more flexible inputs that adjust to shocks more quickly than k. One final feature of the flexibility introduced by interacting l with firm type and state is that we estimate substantially different coefficients on l across groups.

As with the GST model, we prioritise firm fixed effects over industry controls since the former introduce a firm-level adjustment to the relationship between l and other productivity components. Using that quality prioritisation, table 7 shows the relative contribution of each data source and integration method to the complete productivity dataset, with the rows ordered from highest to lowest priority based on data quality. The top three rows, therefore, capture the pre-existing tier 1 productivity data following the prioritisation of Fabling and Maré (2015b), who prefer duel AES-IR10 observations where both are available. As we saw in table 1, tier 1 data covers 65.9% of observations, and this coverage is largely due to IR10-only observations. Since we now have complete productivity data, we can also estimate the contribution of tier 1 data to each of the productivity components, which confirms the importance of AES in aggregates. While IR10-only firms have 41% of employment, they account for 28-34% of other components. On the other hand, AES-only firms, while only 0.8% of observations, account for 43-47% of aggregates other than labour.

We define three additional tiers of data quality based on data source/model permutation for y, m and k. Tier 2 productivity consists of GST-based observations where firm fixed effects are estimated – ie, where there is an overlap between tier 1 and GST at the firm-year level. Tier 2 is the highest quality new addition to the productivity dataset, and also the largest addition (based on all metrics in table 7). Around 19% of observations are tier 2, capturing 12% of labour input and 8-10% of other productivity components. Tier 3 also depends on GST data in conjunction with detailed industry controls, and constitutes 9% of observations, 6% of L and 3-5% of other components. Tier 4 depends on labour input and only accounts for 6% of observations, 2.7% of L and less than 2% of other components. Even in tier 4, the fixed effect-based approach is feasible in the majority of cases, particularly for relatively large firms in the tier. Table 8 reports tier shares for L and a composite ($Y + M + K$) of other productivity components by industry division, which vary systematically with GST coverage (as reported in table 2).

Together, figures 5–7 show how the composition of productivity data has changed over time. Figure 5 uses full coverage data – population count, labour and GST components – to demonstrate how tier 1 productivity coverage has increased over time. Specifically, population coverage has risen by 10pp over the last two decades, while coverage estimated using labour and GST components has increased by 6-7pp, with the increase driven by a mix of higher IR10 filing rates and improved average quality in filed IR10s (Fabling and Maré 2019).

Within the non-tier 1 subset of the data, figure 6 shows the proportion of aggregates that are derived from tier 2. Between 2004 and 2016, these ratios are reasonably stable for Y, M and L, and slowly increasing for K. From 2016 onwards, the tier 2 share of non-AES/IR10 productivity has declined by 6-8pp. As figure 7 shows, the relative decline in tier 2's contribution over this period is due to that tier being the one that is shrinking as the tier 1 productivity share increases. Tier 3 and 4, meanwhile, are a stubbornly stable share of each aggregate. Setting aside churn in the population, it makes sense that the main source of additional tier 1 observations is firms that have a past history of filing tier 1 data and are likely to be above the mandatory GST filing threshold (ie, tier 2 observations).[16] Prior to 2016, declines in the aggregate share of tier 3 observations exceeds the decline in tier 2 observations.

To test data quality, we primarily rely on productivity component changes for firms moving between productivity tiers – specifically in and out of tier 1.[17] Movement between tiers is likely to be a function of firm performance – eg, crossing the mandatory GST threshold – and, therefore, we expect firm growth rates to differ between firms remaining in any tier, and those that drop or rise. Figure 9 confirms this suspicion by plotting the mean growth rate in l by transition type and over time. The solid line is the mean for firms that are in tier 1 in both t and $t-1$, the dashed line is for firms transitioning out of tier 1, the dash-dotted line is for transitions into tier 1, and the dotted line is for firms remaining outside tier 1 in both years. Firms dropping out of tier one have an average labour growth rate of -15%, while firms joining tier one have an average growth rate of 6%. To compensate for this, we focus on (log) growth rates relative to changes in l in the hope that this partially accounts for actual performance differences that might affect comparisons. We start with time series statistics averaged across firms, before looking at the firm-level distribution of growth rates and within-firm variability in levels by firm size.

Figure 9 reports the mean change in the output, intermediate consumption and capital ratio to labour by t. Across all three ratios, there is a clear difference in mean growth rates for firms

[16] These firms are more likely to be in tier 2 over tier 3 because the presence of tier 1 data enables firm fixed effects estimation (unless GST is zero/missing in tier 1 years).

[17] This is the main transition boundary because tier 1 is relatively large and because, by construction, firms cannot transition between tier 2 and tier 3 (the next two largest tiers).

transitioning into versus out of tier 1. Firms moving to tier 2 or lower have higher mean growth rates than firms moving into tier 1 from a lower tier. In contrast, firms that stay within tier 1 or within a lower tier have similar mean growth rates over time. As already noted, we cannot reject the possibility that mean growth rate differences are due to real world outcomes that are correlated with changing filing patterns. We take comfort from the fact that all subsets of the data reflect similar time series patterns, which is most evident for the output-labour ratio which dropped markedly during the Global Financial Crisis and, to a lesser extent, during the Covid-19 pandemic. Further, year-on-year changes are dominated by firms remaining within or outside tier 1 – averaging 83% of observations – reflecting the persistence of tier 1 filing (table 6).

Figure 10 uses the same firm groups to show the firm-level distribution of these annual changes. A key feature of these figures is that each of the ratios is more stable for stayers, than it is for tier changers – and is more stable for tier 2$+$ stayers than it is for tier 1 stayers (ie, the density around zero is lower for tier 1 than higher tier stayers). To complement this finding, figure 11 looks at the within-firm standard deviation of the three ratios by tier 1 versus lower tiers. The focus here is on whether the within-firm variation in components varies systematically by firm size (mean l), extending the analysis of Fabling and Sanderson (2014b) who showed that tier 1 micro enterprises display greater productivity dispersion (across firms) than larger firms. For all three ratios, we see greater within-firm variation in micro firms, with similar patterns for tier 1 and lower tier firms. In the case of $k-l$ and (to a lesser extent) $y-l$, within-firm variation is lower in tier 2$+$ micro firms compared to tier 1 micro firms. This is, perhaps, unsurprising given that tier 4 (l-based) modelling is relatively more common for micro enterprises.

4 Conclusions

The “complete” productivity table puts another tool in the researcher toolbox. Table 9 provides metadata for this table, which is freely available to LBD users.[18] These data provide improved scope for aggregate statistics, particularly for longitudinal statistics where the current cross-sectional weights may be inadequate.

Only 6% of firm-year productivity observations, corresponding to 3% of aggregate labour input, are modelled from labour inputs (table 7), of which the majority are calibrated against at least one year of firm-level tier 1 productivity filing (ie, firm fixed effects apply). The remaining data rely on AES, IR10 or GST returns, where the latter is used for 28% of observations, and 18% of labour input. These numbers suggest that the complete productivity table is suitable for applications focussed on employees, particularly where incompleteness would complicate the analysis (eg, where job changers are an important subset as in, eg, Fabling and Maré 2024).

Our analysis suggests several words of caution. Firstly, firms moving productivity tier have sys-

[18] For researchers wishing to understand the tier 1 data processing, or use non-productivity variables (eg, capital stock composition), the original AES/IR10 productivity table is unchanged.

tematically different mean growth rates. At least some of these differences are likely to be real and correlated with filing change (see figure 9). They could also reflect inadequacies in the GST harmonisation process, which could be driven by the GST data, or by differences between the firms used to estimate the models and the firms with productivity gaps. While hard to pin down, there is likely to be some truth to these additional explanations. Secondly, the addition of GST-based observations doesn't fix pre-existing micro enterprise measurement issues, which are likely exacerbated by relatively high rates of WP-only firms, and of firm churn (Fabling and Sanderson 2014b). While greater within-firm variation in productivity components might be expected for micro businesses, it may be prudent to exclude entry/exit years from analysis, particularly for WP-only firms. Finally, the construction of tier $2+$ k suffers from the lack of a corresponding administrative data source. Firm-level stability in the capital stock – at least as a proportion of labour – makes k easier to predict, but users should keep in mind that sales and purchases have been used to estimate k.

With those caveats, the complete dataset presents an opportunity for researchers to test the robustness of their results to the inclusion/exclusion of tier $2+$ observations with a harmonised set of productivity components, and to explore the potential for bias that non-random missingness introduces.

The success of the regression-based approach suggests it may be feasible to use GST/EMS data to predict aggregate productivity components from the bottom up.[19] Because both EMS and GST are high-frequency and timely data, such an approach would amount to "nowcasting," which may be valuable given the substantial delays between real world events and official statistics.[20]

[19] A non-exhaustive list of issues that would need to be addressed includes: the use of t-specific coefficients; the availability of deflators; and quarterly vs annual frequency. In addition, Fabling and Maré (2019) conclude that post-AES adjustments by Stats NZ affect the comparability of micro aggregates and official statistics in non-trivial ways.

[20] Even in the absence of full forecasts, micro data has been used to improve productivity growth forecasts (see, eg, Bartelsman and Wolf 2014).

References

Allan, C. and D. C. Maré (2021). Do workers share in firm success? Pass-through estimates for New Zealand. Discussion Paper 14764, Institute of Labor Economics (IZA).

Allan, C. and D. C. Maré (2022). Who benefits from firm success? Heterogenous rent sharing in New Zealand. Discussion Paper 15264, Institute of Labor Economics (IZA).

Apatov, E., R. Fabling, A. B. Jaffe, M. Morris, and M. Thirkettle (2015). Agricultural productivity in New Zealand: First estimates from the Longitudinal Business Database. Working Paper 15-13, Motu Economic and Public Policy Research.

Bartelsman, E. J. and Z. Wolf (2014). Forecasting aggregate productivity using information from firm-level data. *The Review of Economics and Statistics 96*(4), 745–755.

Chappell, N. and A. Jaffe (2018). Intangible investment and firm performance. *Review of Industrial Organization 52*(4), 509–559.

Chappell, N., A. B. Jaffe, and T. Le (2018). Worker flows, entry and productivity in the New Zealand construction industry. Working Paper 24376, National Bureau of Economic Research.

Conway, P., L. Meehan, and G. Zheng (2015). Do New Zealand firms catch up to the domestic productivity frontier? Working Paper 2015/3, New Zealand Productivity Commission.

Criscuolo, C., A. Hijzen, C. Schwellnus, E. Barth, W.-H. Chen, R. Fabling, P. Fialho, K. Grabska, R. Kambayashi, T. Leidecker, O. Nordström Skans, C. Riom, D. Roth, B. Stadler, R. Upward, and W. Zwysen (2020). Workforce composition, productivity and pay: The role of firms in wage inequality. Discussion Paper 13212, Institute of Labor Economics (IZA).

De Loecker, J. and C. Syverson (2021). An industrial organization perspective on productivity. In K. Ho, A. Hortaçsu, and A. Lizzeri (Eds.), *Volume 4*, Number 1 in Handbook of Industrial Organization, Chapter 3, pp. 141–223. Elsevier.

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. (2021a). Living on the edge: An anatomy of New Zealand's most productive firms. Working Paper 21-01, Motu Economic and Public Policy Research.

Fabling, R. (2021b). Of interest? Estimating the average interest rate on debt across firms and over time. Working Papers 21-05, Motu Economic and Public Policy Research.

Fabling, R. and A. Grimes (2014). The "suite" smell of success: Personnel practices and firm performance. *Industrial and Labor Relations Review 67*(4), 1095–1126.

Fabling, R. and A. Grimes (2021). Picking up speed: Does ultrafast broadband increase firm productivity? *Information Economics and Policy 57*(C), 1–13.

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 D. C. Maré (2019). Improved productivity measurement in New Zealand's Longitudinal Business Database. Working Paper 19-03, Motu Economic and Public Policy Research.

Fabling, R. and D. C. Maré (2024). Pay in Māori-led firms. Working Paper 24-05, Motu Economic and Public Policy Research.

Fabling, R., D. C. Maré, and P. Stevens (2022). Migration and firm-level productivity. Discussion Paper 15482, Institute of Labor Economics (IZA).

Fabling, R. and L. Sanderson (2013). Exporting and firm performance: Market entry, investment and expansion. *Journal of International Economics 89*(2), 422–431.

Fabling, R. and L. Sanderson (2014a). Foreign acquisition and the performance of New Zealand firms. *New Zealand Economic Papers 48*(1), 1–20.

Fabling, R. and L. Sanderson (2014b). 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.

Grimes, A., C. Ren, and P. Stevens (2012). The need for speed: Impacts of internet connectivity on firm productivity. *Journal of Productivity Analysis 37*(2), 187–201.

Jaffe, A., T. Le, and N. Chappell (2016). Productivity distribution and drivers of productivity growth in the construction sector. Working Paper 16-08, Motu Economic and Public Policy Research.

Kirker, M. and L. Sanderson (2022). Firm productivity growth and the knowledge of new workers. Working Paper 22/01, New Zealand Treasury.

Law, D. and N. McLellan (2005). The contributions from firm entry, exit and continuation to labour productivity growth in New Zealand. Working Paper 05/01, New Zealand Treasury.

Maré, D. C. (2016). Urban productivity estimation with heterogeneous prices and labour.

Working Paper 16-21, Motu Economic and Public Policy Research.

Maré, D. C. and R. Fabling (2013). Productivity and local workforce composition. In R. Crescenzi and M. Percoco (Eds.), *Geography, institutions and regional economic performance*, Advances in spatial science, pp. 59–76. Berlin and New York: Springer.

Maré, D. C. and R. Fabling (2019). Competition and productivity: Do commonly used metrics suggest a relationship? Working Paper 19-16, Motu Economic and Public Policy Research.

Maré, D. C. and D. J. Graham (2013). Agglomeration elasticities and firm heterogeneity. *Journal of Urban Economics 75*(C), 44–56.

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

Maré, D. C., T. Le, R. Fabling, and N. Chappell (2017). Productivity and the allocation of skills. Working Paper 17-04, Motu Economic and Public Policy Research.

Maré, D. C. and J. Timmins (2006). Geographic concentration and firm productivity. Working Paper 06-08, Motu Economic and Public Policy Research.

Meehan, L. (2020). Productivity in New Zealand: The role of resource allocation among firms. *New Zealand Economic Papers 54*(1), 39–66.

Sin, I., S. Stillman, and R. Fabling (2022). What drives the gender wage gap? Examining the roles of sorting, productivity differences, bargaining, and discrimination. *The Review of Economics and Statistics 104*(4), 636–651.

Syverson, C. (2011). What determines productivity? *Journal of Economic Literature 49*(2), 326–365.

Zheng, G. (2016). Geographic proximity and productivity convergence across New Zealand firms. Working Paper 2016/04, New Zealand Productivity Commission.

Tables

Table 1: Coverage of tier 1 productivity data by industry and firm size

			p(observe tier 1 prod)			
				Employer		All
	Industry	N(obs)	WP-only	L<=5	L>5	firms
A	Agriculture, Forestry & Fishing	1,380,219	0.633	0.693	0.725	0.658
B	Mining	8,439	0.623	0.680	0.807	0.694
C	Manufacturing	468,432	0.653	0.731	0.797	0.718
D	Electricity, Gas, Water & Waste Services	17,898	0.719	0.707	0.835	0.738
E	Construction	1,143,921	0.638	0.720	0.803	0.681
F	Wholesale Trade	302,052	0.652	0.720	0.781	0.710
G	Retail Trade	550,359	0.620	0.742	0.822	0.712
H	Accommodation & Food Services	398,697	0.621	0.691	0.782	0.694
I	Transport, Postal & Warehousing	332,070	0.581	0.652	0.755	0.618
J	Information Media & Telecommunications	81,804	0.560	0.644	0.718	0.598
K	Financial & Insurance Services	123,417	0.571	0.584	0.698	0.588
L	Rental, Hiring & Real Estate Services	73,107	0.569	0.684	0.749	0.627
M	Professional, Scientific & Technical Services	1,111,188	0.561	0.666	0.726	0.598
N	Administrative & Support Services	331,464	0.578	0.650	0.694	0.611
R	Arts & Recreation Services	138,477	0.510	0.653	0.744	0.557
S	Other Services	377,979	0.639	0.734	0.805	0.700
	Total	**6,839,523**	**0.608**	**0.700**	**0.775**	**0.659**

Population defined as private-for-profit firms in "market sector" industries and non-zero labour input (L). Coverage rates based on productivity population firm-years (2001-2022). Results reported by ANZSIC'06 division excluding divisions that have no productivity industries. Observed productivity components are derived from AES and/or IR10 (described as "tier 1" in this paper).

Table 2: Coverage of GST data where tier 1 productivity data missing

			p(GST>0 | no tier 1 prod)			
				Employer		All
	Industry	N(obs)	WP-only	L<=5	L>5	firms
A	Agriculture, Forestry & Fishing	471,786	0.831	0.950	0.962	0.873
B	Mining	2,586	0.827	0.909	0.779	0.849
C	Manufacturing	132,012	0.777	0.941	0.930	0.859
D	Electricity, Gas, Water & Waste Services	4,683	0.818	0.931	0.901	0.878
E	Construction	365,301	0.775	0.927	0.956	0.828
F	Wholesale Trade	87,663	0.817	0.932	0.951	0.885
G	Retail Trade	158,418	0.744	0.941	0.969	0.850
H	Accommodation & Food Services	122,121	0.723	0.934	0.948	0.874
I	Transport, Postal & Warehousing	126,774	0.828	0.936	0.928	0.860
J	Information Media & Telecommunications	32,898	0.761	0.898	0.915	0.803
K	Financial & Insurance Services	50,790	0.574	0.627	0.745	0.606
L	Rental, Hiring & Real Estate Services	27,240	0.784	0.929	0.907	0.835
M	Professional, Scientific & Technical Services	447,159	0.802	0.899	0.936	0.827
N	Administrative & Support Services	129,018	0.721	0.879	0.929	0.780
R	Arts & Recreation Services	61,278	0.681	0.895	0.909	0.728
S	Other Services	113,394	0.645	0.927	0.962	0.778
	Total	**2,333,121**	**0.779**	**0.920**	**0.941**	**0.834**

GST coverage defined as non-zero GST sales and/or purchases in the year. Firm-year population as in table 1 excluding observations where tier 1 (AES/IR10) productivity components are present.

Table 3: Mean and standard deviation of GST and tier 1 productivity data

		Level		Growth rate	
	(log) variable	Mean	St.dev.	Mean	St.dev.
l	Labour	0.425	1.092	-0.001	0.536
y	Gross output	12.246	1.685	0.015	0.606
m	Intermediate consumption	11.266	1.787	0.010	0.626
k	Capital services	10.217	1.704	0.024	0.552
s	GST sales	12.237	1.807	0.005	0.800
p	GST purchases	11.505	1.923	-0.041	0.761
$k-l$	Capital-labour ratio	9.657	1.408	-0.004	0.594
$v_{ym}-l$	Labour productivity (prod)	11.149	1.134	-0.003	0.686
$v_{sp}-l$	Labour productivity (GST)	11.040	1.301	0.027	0.916

Value-added is output less intermediate consumption (v_{ym}), or sales less purchases (v_{sp}). Logging excludes non-positive observations. Growth rate is log first difference. All statistics estimated on full sample, ie, coverage varies across variables (AES/IR10 for productivity and GST for sales/purchases), and across levels and growth rates for the same variable.

Table 4: GST and tier 1 productivity correlations

	Level					
	l	y	m	k	s	p
l	1.000					
y	0.678	1.000				
m	0.633	0.881	1.000			
k	0.568	0.649	0.687	1.000		
s	0.653	0.931	0.837	0.636	1.000	
p	0.621	0.821	0.890	0.721	0.848	1.000
	Growth rate					
	Δl	Δy	Δm	Δk	Δs	Δp
Δl	1.000					
Δy	0.383	1.000				
Δm	0.299	0.636	1.000			
Δk	0.319	0.385	0.386	1.000		
Δs	0.373	0.670	0.474	0.334	1.000	
Δp	0.331	0.479	0.552	0.360	0.544	1.000

Pairwise correlation (ie, using all observations for each pair) reported. Casewise correlation (ie, a common sample for all pairs) produces similar results. Growth rate (Δ) is log first difference.

Table 5: Pattern of zeros in GST and tier 1 productivity

$S > 0$	$P > 0$	N(obs)	with prod	Proportion of prod where $Y = 0$	$M = 0$
No GST		511,038	0.333	0.030	0.020
N	N	65,244	0.283	0.255	0.151
N	Y	101,373	0.455	0.520	0.015
Y	N	58,830	0.243	0.030	0.173
Y	Y	6,103,038	0.698	0.005	0.001
Total		**6,839,523**	**0.659**	**0.012**	**0.003**

Final two columns are proportions conditional on being present in the tier 1 productivity data.

Table 6: Patterns in tier 1 data availability

In prod $t-1$	$t+1$	N(obs)	p(prod) t	Share of missing
Ongoing firms				
N	N	909,639	0.102	0.431
N	Y	411,480	0.639	0.078
Y	N	463,005	0.657	0.084
Y	Y	2,579,139	0.931	0.094
Entering firms				
.	N	217,206	0.245	0.086
.	Y	349,047	0.838	0.030
Exiting firms				
N	.	221,394	0.153	0.099
Y	.	315,102	0.644	0.059
One-year firms				
.	.	122,700	0.392	0.039
Totals				
Two adj. prod		2,579,139	0.931	0.094
One adj. prod		1,538,634	0.691	0.251
No adj. prod		1,470,939	0.722	0.655

Ongoing firms are in the productivity population for three consecutive years ($t-1$, t, $t+1$). Entering, exiting, and one-year firms are only present in (t, $t+1$), ($t-1$, t), and (t) respectively. Analysis restricted to $t \in [2003, 2020]$.

Table 7: Productivity aggregate shares by source

Tier	Source y, m, k	Total N(obs)	Total L	Proportion of total N(obs)	L	Y	M	K
1	AES + IR10	101,475	3,167,600	0.015	0.106	0.109	0.113	0.074
1	AES	55,464	8,281,500	0.008	0.276	0.431	0.475	0.437
1	IR10	4,349,466	12,335,500	0.636	0.411	0.295	0.276	0.342
2	$s, p +$ ffe	1,270,224	3,611,700	0.186	0.120	0.100	0.079	0.103
3	$s, p +$ ind	641,049	1,776,300	0.094	0.059	0.048	0.040	0.028
4	$l +$ ffe	236,949	424,100	0.035	0.014	0.011	0.012	0.011
4	$l +$ ind	184,887	380,800	0.027	0.013	0.006	0.005	0.005

"+ FFE"/"+ industry" indicates models that include firm fixed effects/detailed industry dummies respectively, based on equations (1)–(3). Row order in the table reflects our prioritisation of data source and model. Tier 1 is data derived from AES/IR10; tier 2 relies on GST with firm fixed effects; tier 3 relies on GST with detailed industry controls; and tier 4 relies on labour input.

Table 8: Productivity aggregate shares by industry and tier

	Industry	Tier 2 share L	Tier 2 share YMK	Tier 3 share L	Tier 3 share YMK	Tier 4 share L	Tier 4 share YMK
A	Agriculture, Forestry & Fishing	0.203	0.196	0.082	0.058	0.034	0.022
B	Mining	0.041	0.068	0.033	0.010	0.074	0.028
C	Manufacturing	0.092	0.049	0.046	0.032	0.021	0.013
D	Electricity, Gas, Water & Waste Services	0.030	0.005	0.015	0.003	0.008	0.002
E	Construction	0.132	0.111	0.052	0.045	0.026	0.017
F	Wholesale Trade	0.118	0.119	0.064	0.057	0.017	0.015
G	Retail Trade	0.087	0.089	0.041	0.032	0.016	0.011
H	Accommodation & Food Services	0.135	0.144	0.055	0.051	0.024	0.019
I	Transport, Postal & Warehousing	0.099	0.069	0.043	0.050	0.025	0.012
J	Information Media & Telecommunications	0.074	0.042	0.036	0.037	0.019	0.007
K	Financial & Insurance Services	0.049	0.045	0.058	0.051	0.057	0.053
L	Rental, Hiring & Real Estate Services	0.130	0.096	0.068	0.042	0.032	0.024
M	Professional, Scientific & Technical Services	0.155	0.170	0.084	0.070	0.036	0.026
N	Administrative & Support Services	0.111	0.120	0.105	0.107	0.034	0.023
R	Arts & Recreation Services	0.129	0.163	0.076	0.055	0.052	0.021
S	Other Services	0.148	0.135	0.043	0.044	0.028	0.017
	Total	**0.120**	**0.093**	**0.059**	**0.044**	**0.027**	**0.017**

"YMK" column is the share of the sum of Y, M, and K. Tier 2 relies on GST with firm fixed effects; tier 3 relies on GST with detailed industry controls; and tier 4 relies on labour input. Tier 1 (derived from AES/IR10) share not reported (since shares sum to one).

Table 9: [ibuldd_research_datalab].[STATSNZ\dl_RFabling].[pent_prod_pop_IDI_202310]

	Variable	Type		Description
KEY	pent	char(10)	NOT NULL	Firm id (permanent enterprise number)
KEY	dim_year_key	int	NOT NULL	March year (YYYY03)
	entry_Lpop	tinyint	NOT NULL	Indicator for $L_{t-1} = 0$ (ie, pop entry)
	exit_Lpop	tinyint	NOT NULL	Indicator for $L_{t+1} = 0$ (ie, pop exit)
	active_prior_year	tinyint	NOT NULL	Indicator for any activity at $t - 1$ (L,GST,prod)
	active_next_year	tinyint	NOT NULL	Indicator for any activity at $t + 1$ (L,GST,prod)
	pf_ind	varchar(4)	NOT NULL	Production function industry
	fte	float	NOT NULL	FTE employment
	wp_adj	float	NOT NULL	WP labour adjusted for firm entry/exit
	size_stratum	varchar(17)	NOT NULL	Stratification by industry and firm size
NEW	lnY	float	NOT NULL	Gross output (real, logged)
NEW	lnM	float	NOT NULL	Intermediate consumption (real, logged)
NEW	lnK	float	NOT NULL	Capital services (real, logged)
NEW	lnL	float	NOT NULL	Labour input (logged)
NEW	prod_tier	tinyint	NOT NULL	Productivity tier (values 1-4)
NEW	source_ymk	varchar(8)	NOT NULL	Source for $Y/M/K$

This table is available to all researchers who have Longitudinal Business Database access. Productivity tier (prod_tier) and source (source_ymk) correspond to the ranking in columns one and two of table 7 respectively.

Figures

Figure 1: Coverage of tier 1 productivity and GST data by industry and firm size

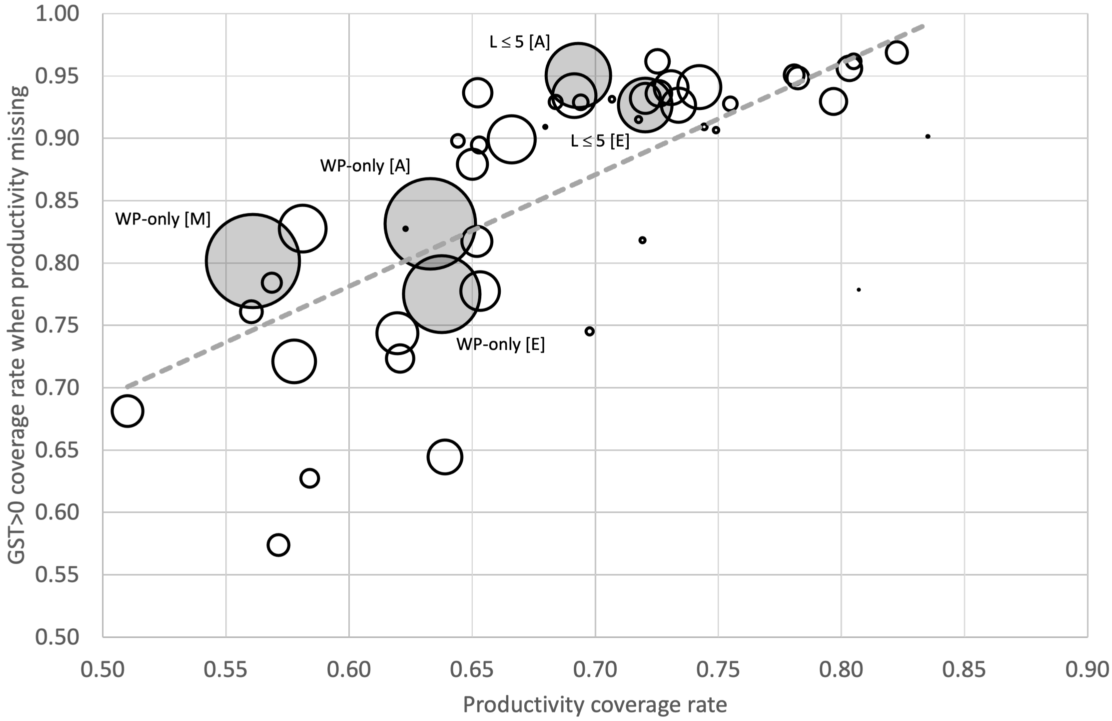

Industry/firm size groups as in table 1. Bubble size reflects the total number of observations in the group that are not tier 1 (ie, observed using AES/IR10). Shaded bubbles highlight the five largest gaps (labelled). Dashed line is the unweighted linear relationship between coverage rates.

Figure 2: Ratio of GST and tier 1 productivity aggregates, where both observed

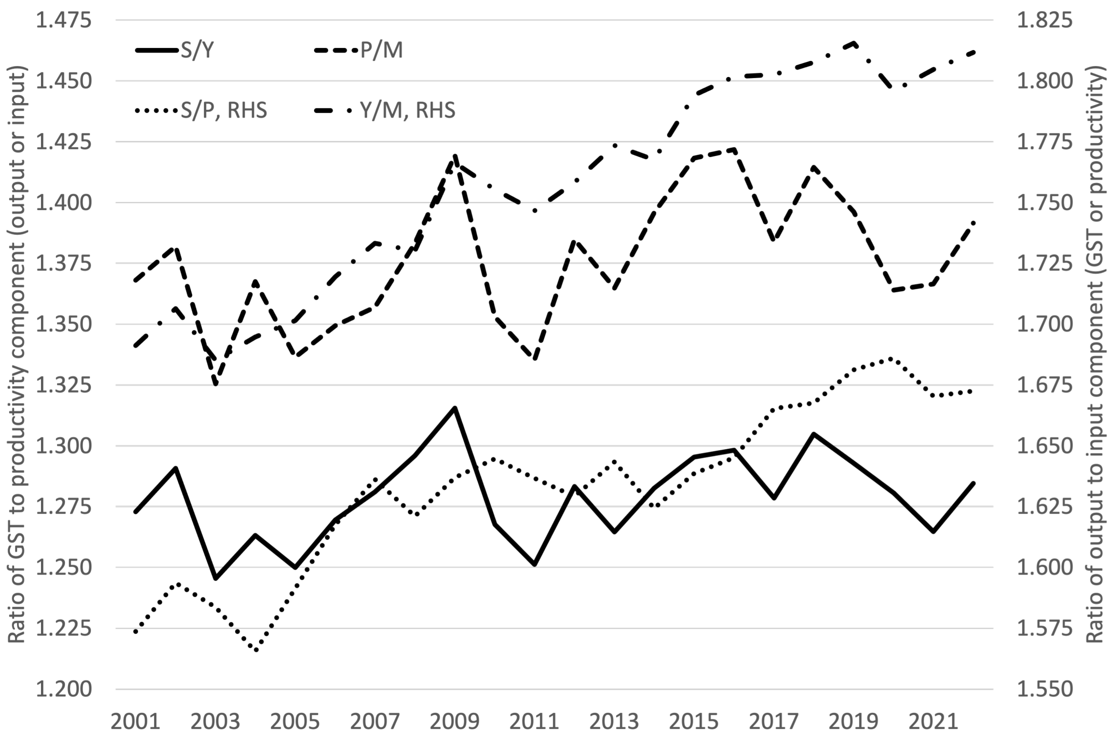

Ratios based on tier 1 productivity observations with non-zero GST sales or purchases.

Figure 3: Distribution of y by GST sales status

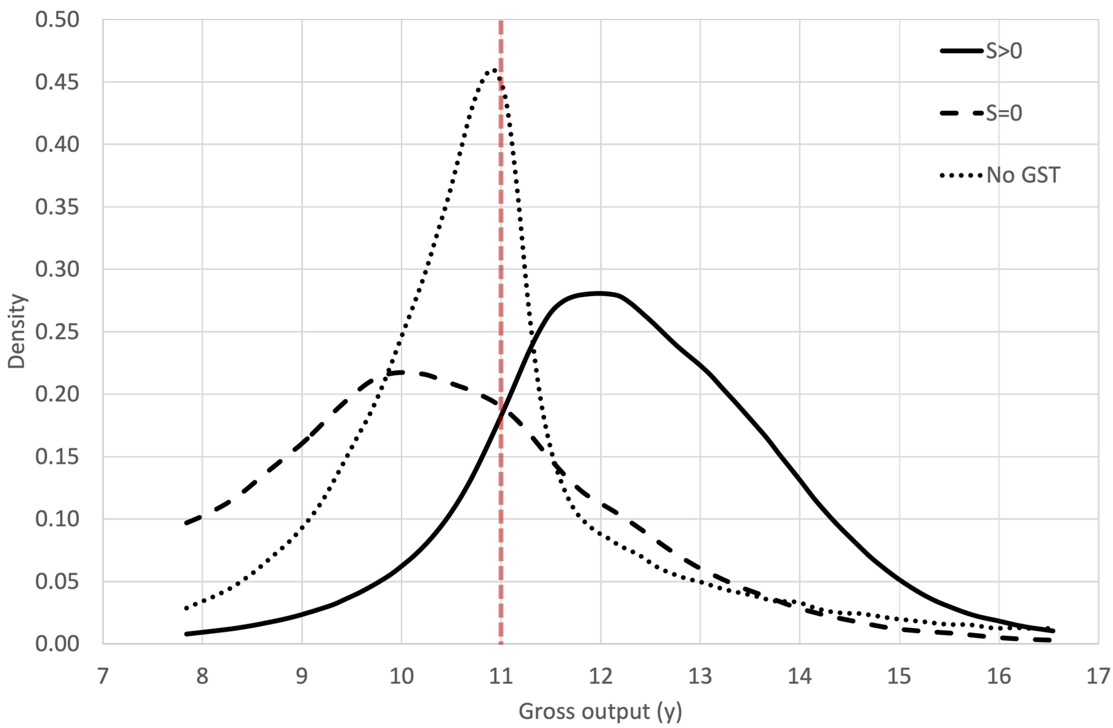

Kernel density (Stata default parameters) truncated for confidentiality and clarity. Vertical dashed line is current mandatory GST threshold ($60,000).

Figure 4: Distribution of m by GST purchases status

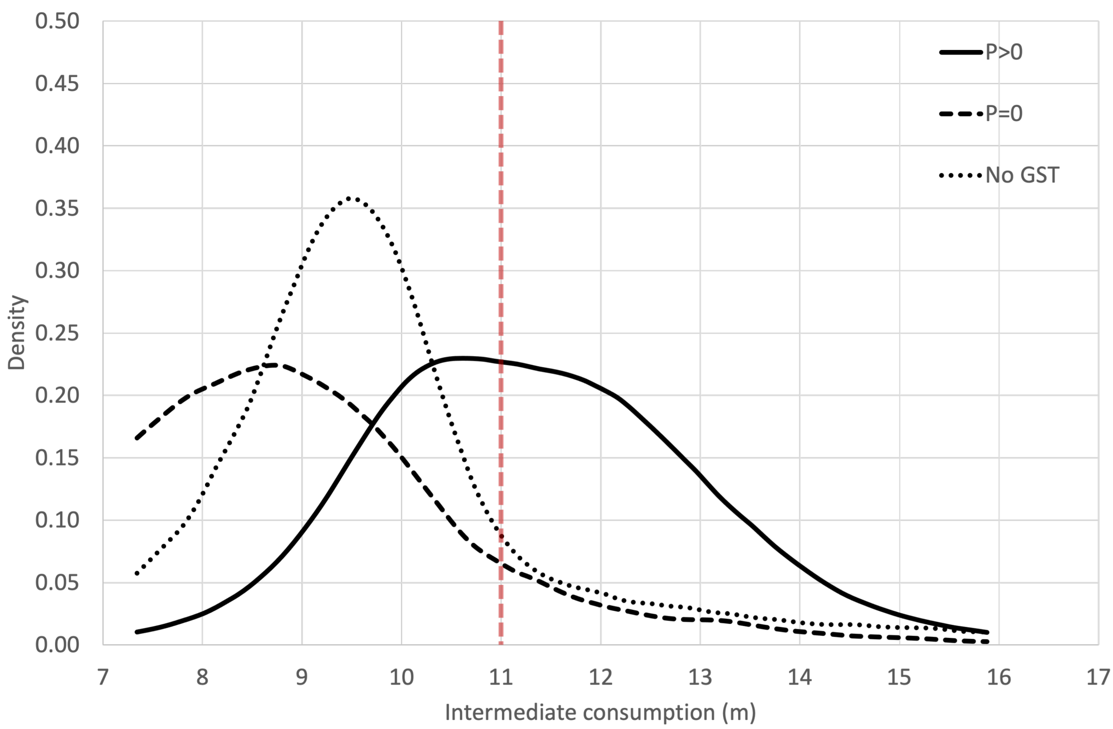

See figure 3 notes.

Figure 5: Tier 1 productivity coverage estimated using aggregate L, S and P

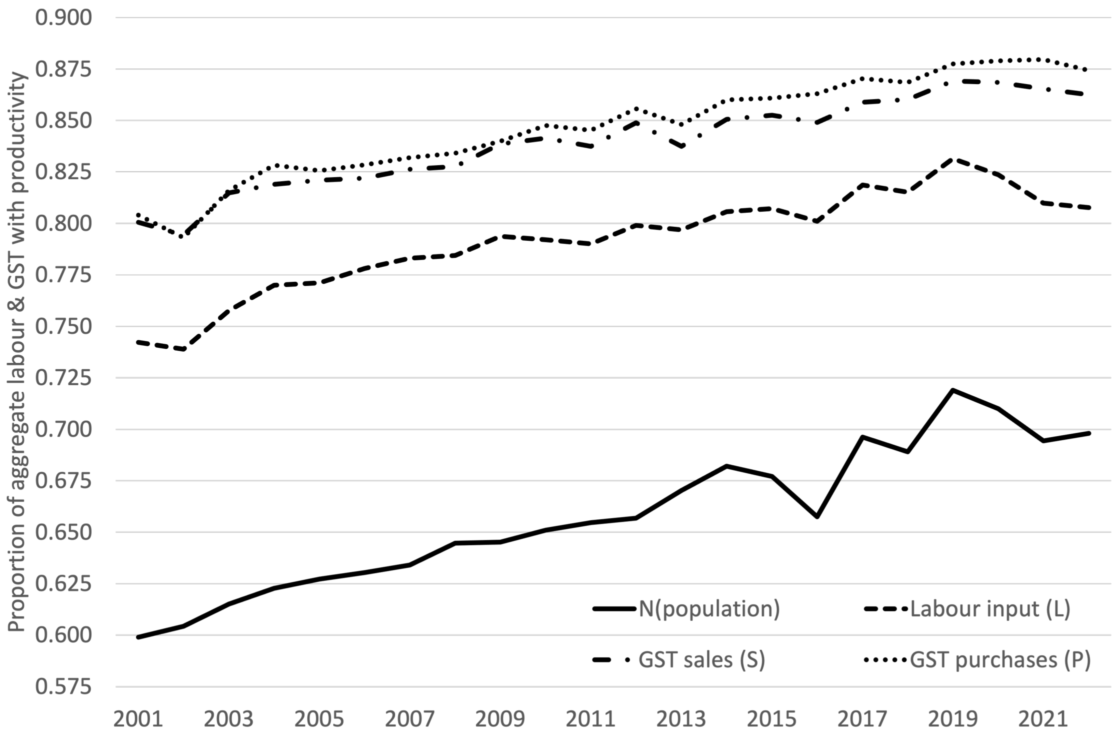

L, S, P and the population size are treated as full coverage due to mandatory tax filing thresholds coupled with the restriction to the productivity population.

Figure 6: Tier 2 share of tier 2+ aggregate productivity components

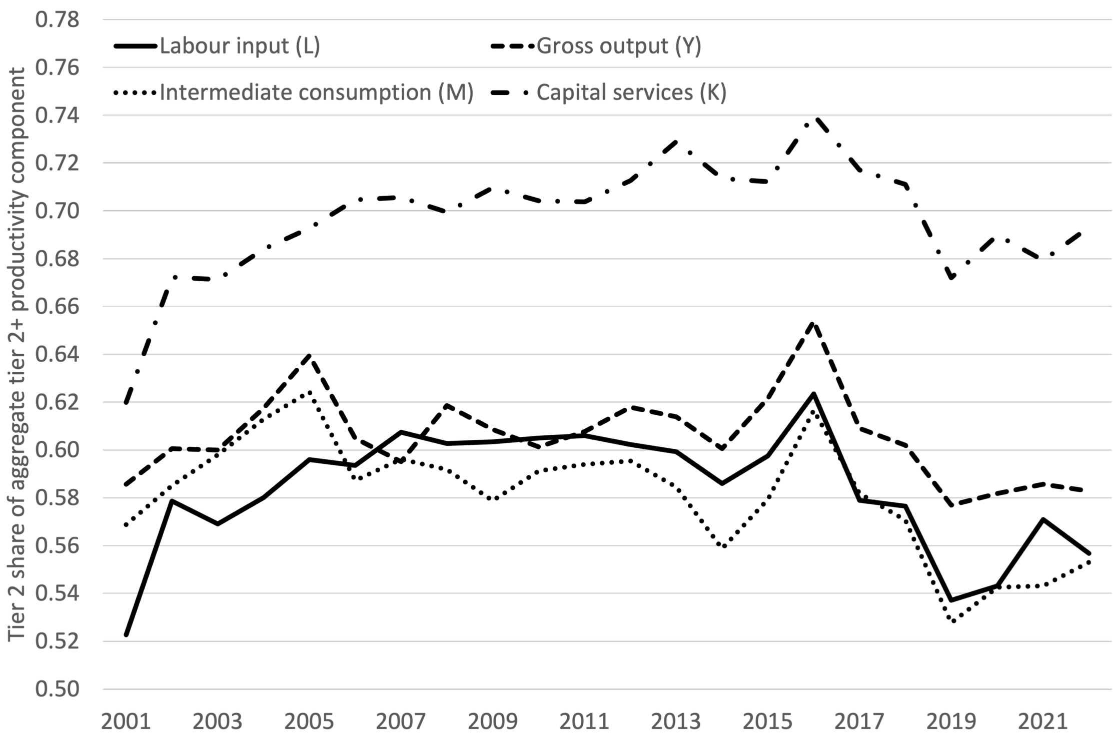

Figure 7: Share of aggregate productivity components for tiers 2+ by t

Gross output (Y)

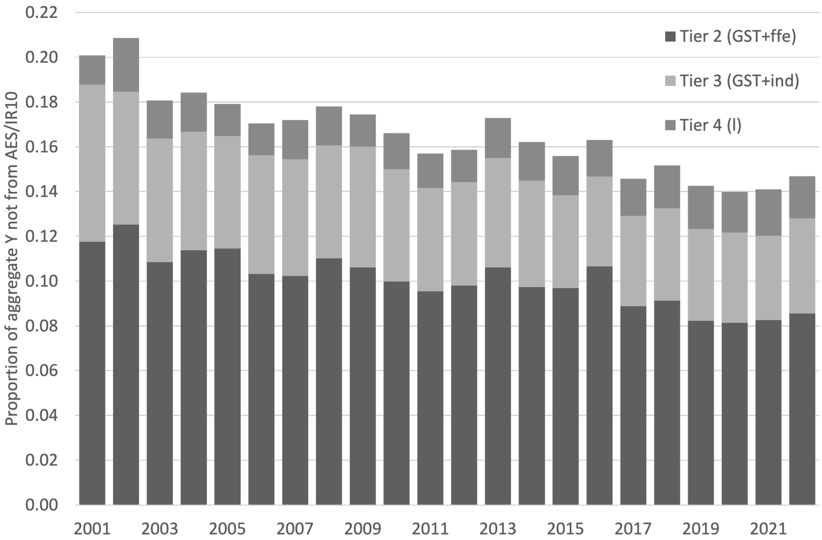

Intermediate consumption(M)

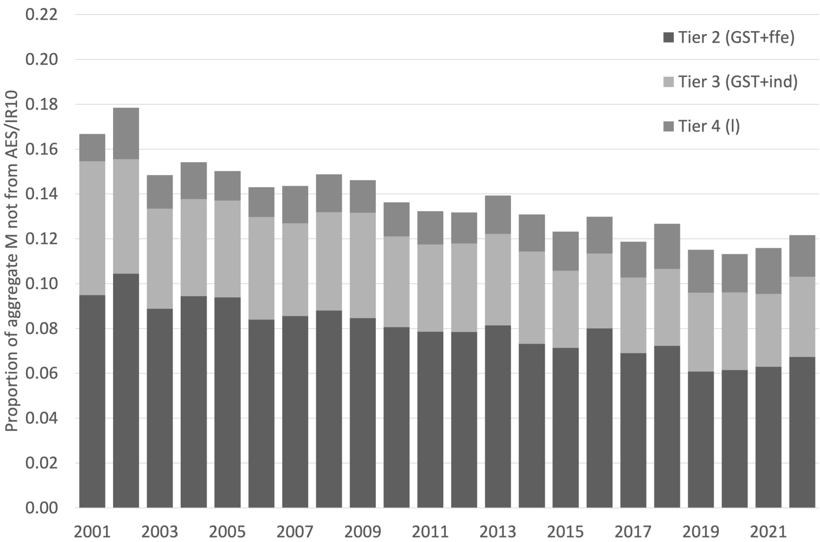

Capital services (K)

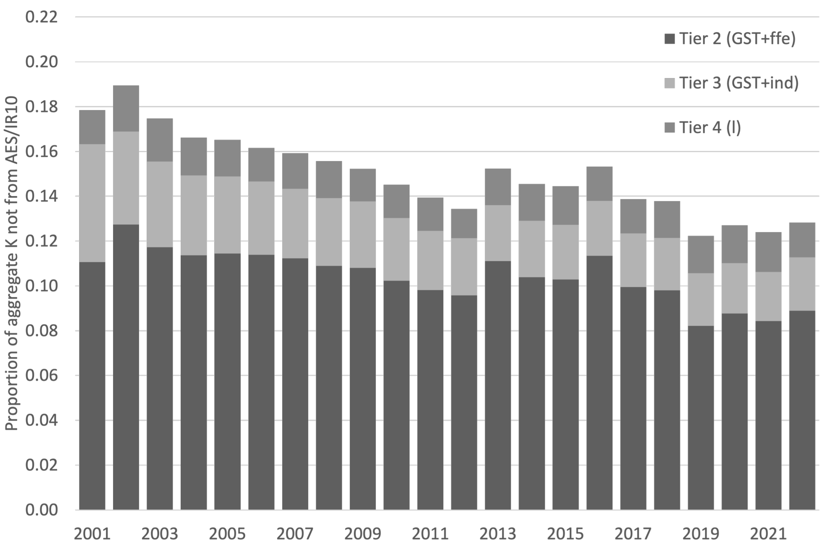

Figure 8: Mean change in labour input (l) by tier and t

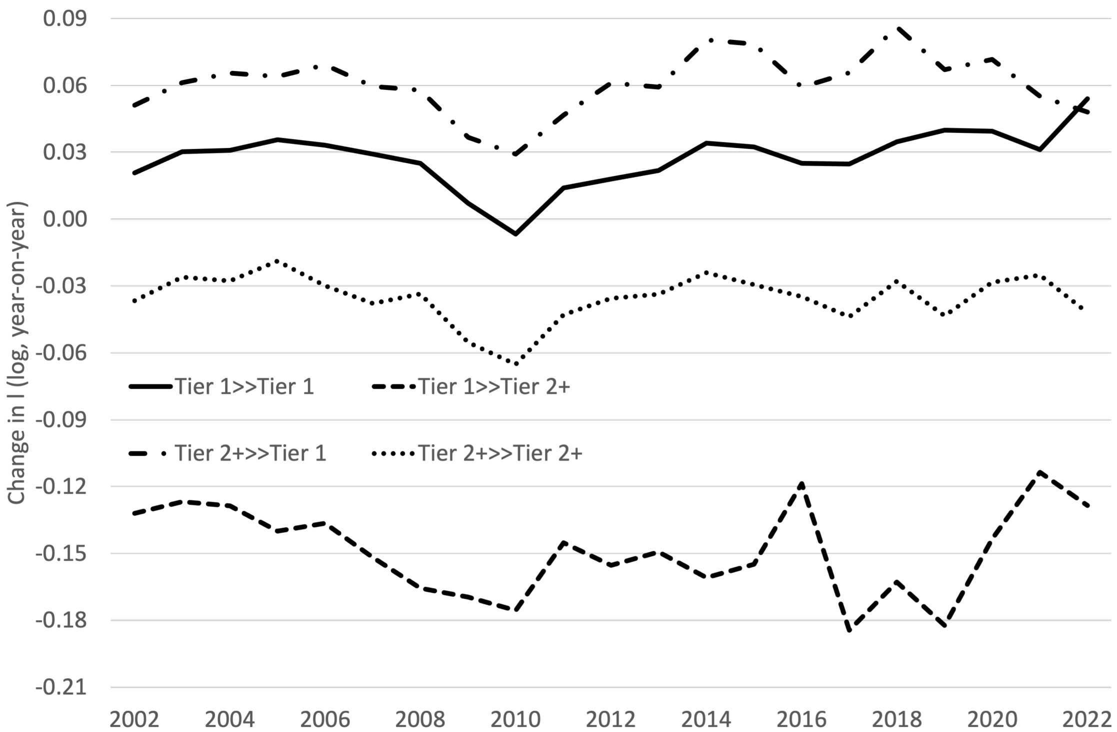

Figure 9: Mean change in productivity ratios by tier and t

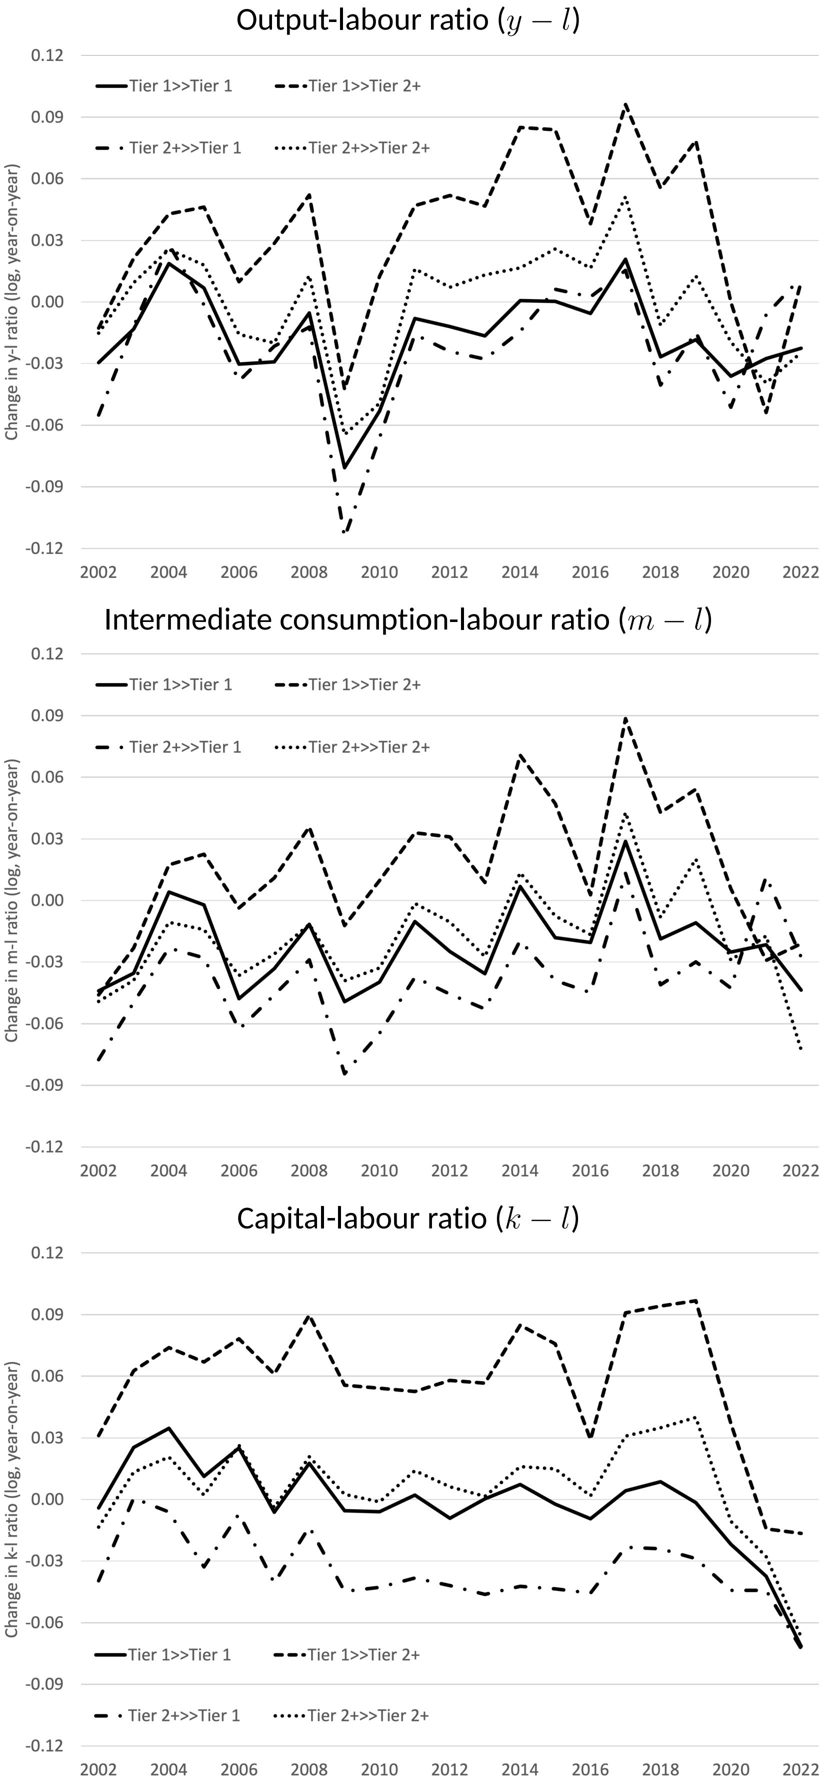

Figure 10: Distribution of change in productivity ratios by tier

Output-labour ratio $(y - l)$

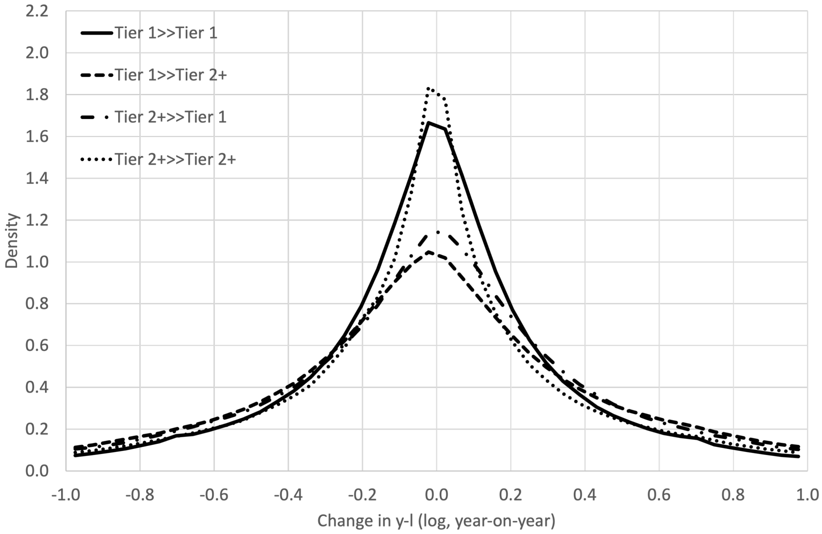

Intermediate consumption-labour ratio $(m - l)$

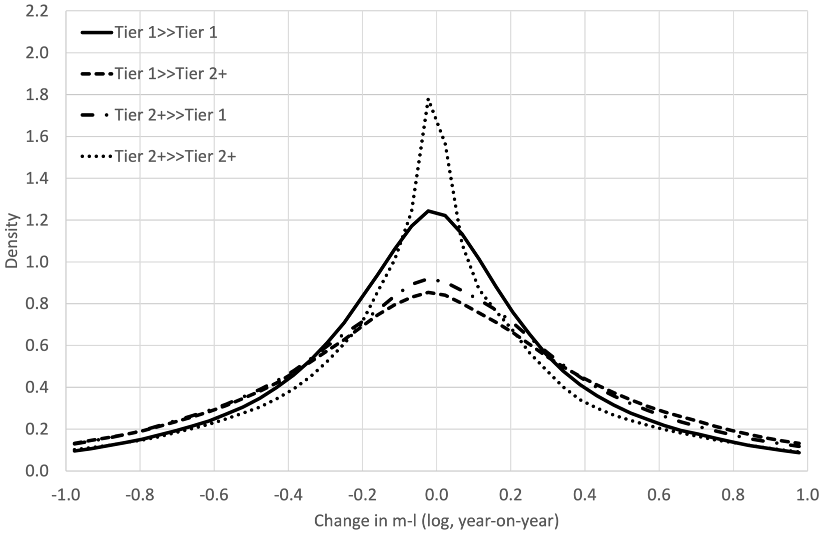

Capital-labour ratio $(k - l)$

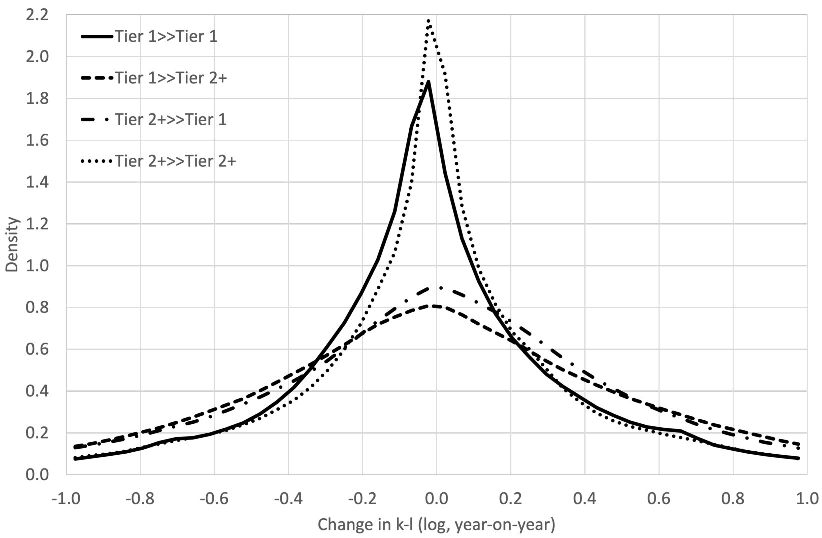

Kernel density (Stata default parameters) truncated for confidentiality and clarity.

Figure 11: Firm-level standard deviation of productivity ratios by tier and l

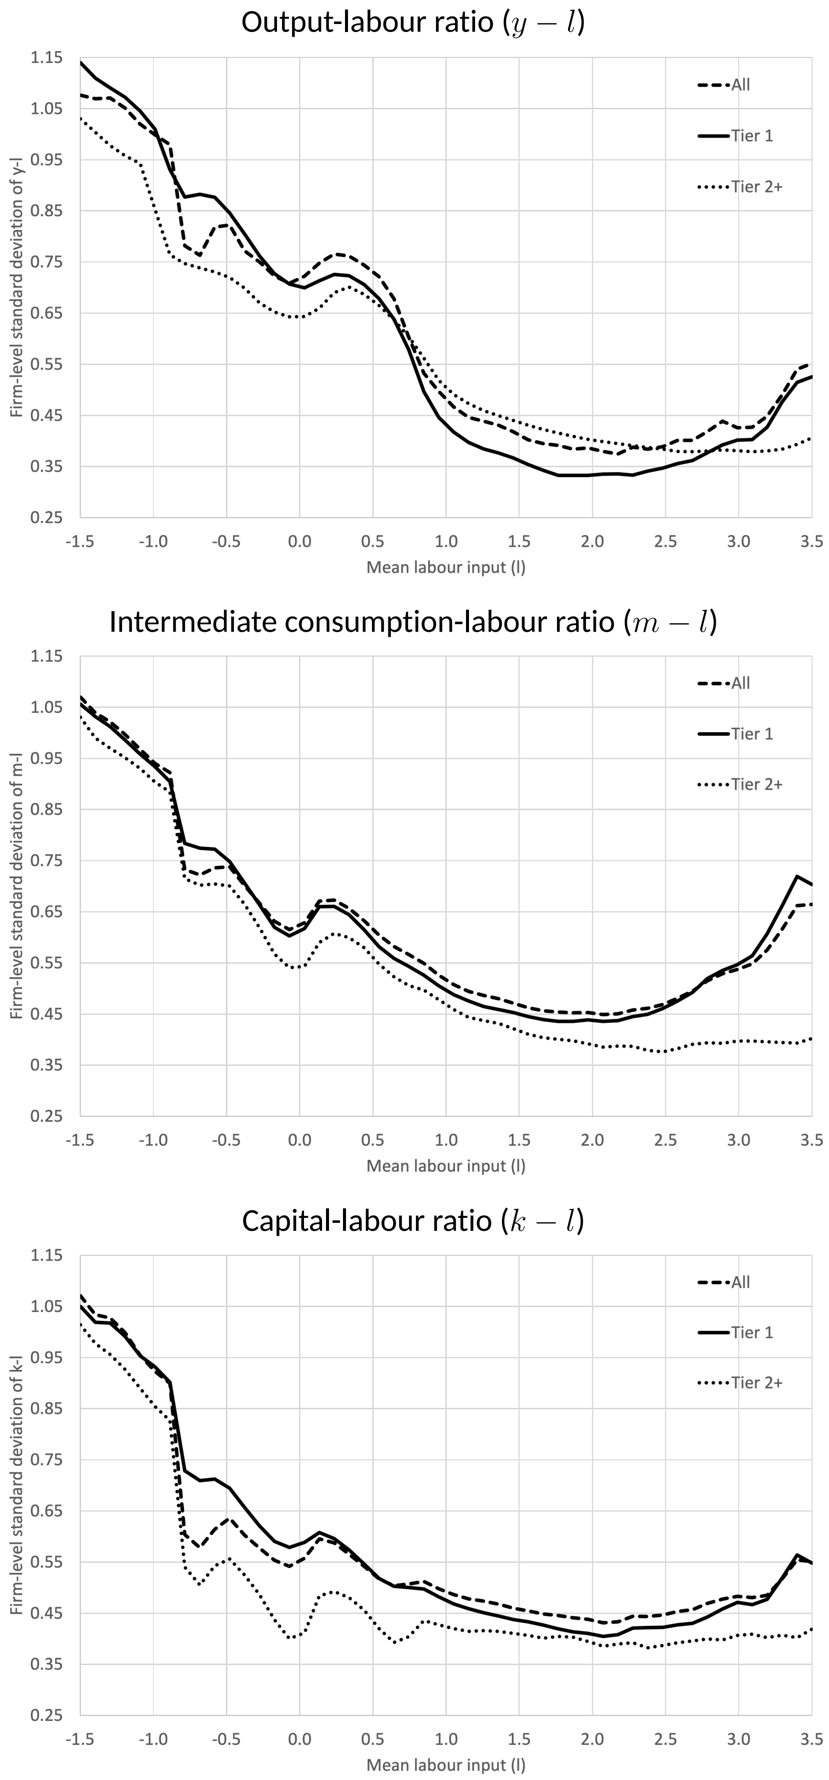

Kernel-weighted local polynomial regression (Stata default options) with mean l truncated for confidentiality and clarity. Analysis restricted to firms with multiple observations within/outside tier 1. Varying the minimum firm observations count has minimal effect on the shape of the curves.

Appendix A – Productivity innovations since Fabling & Maré (2019)

AES/IR10 quality tests (2021)
Fabling (2021b) identified quality issues with the IR10 balance sheet component. As a consequence, additional tests are now present to spot and remove firm-year observations where the fixed asset schedule is not itemised. At the same time we identify and remove a small number of outlier AES observations that were previously resulting in firms being dropped from the productivity data (due to subsequent cleaning steps).

Plant locations (2022)
Employing plant (PBN) location (latest meshblock instance) has been added to the monthly employment table, which allows researchers to easily identify the presence of firms across regions. This change is largely a convenience feature that prevents the need to link through to the Business Register (BR) and determine the current meshblock instance. Working proprietor labour input is not allocated to PBNs, so the BR is still needed to identify non-employer firm locations.

GST processing (2022)
GST data processing now accounts for industry-level seasonal variation, building on the (now defunct) Business Activity Indicator approach implemented by Stats NZ. These changes apportion two- and six-monthly GST returns to the monthly level (pent_month_GST_IDI_202310) then aggregate to the firms' permanent balance date year (pent_year_GST_IDI_202310). GST filing patterns are used to identify missing filling frequency information. Monthly filers within each industry determine the relative proportion of two-monthly filer output to be assigned to each month within the return period. The combined monthly dataset for monthly and two-monthly filers is then used to determine the relative proportions for six-monthly filers. We now retain GST returns where sales and purchases are both zero. These returns were previously dropped.

Right-of-use accounting changes (2024)
The introduction in 2020 of new accounting rules for large firms (expenses of at least $2 million) caused a substantial increase in reported aggregate depreciation expenses, which is a component of K in the productivity dataset. In response to the accounting changes, Stats NZ introduced additional AES questions, allowing them (and us) to unwind these changes for AES respondents. We additionally use these data to identify AES-IR10 firms whose IR10s are also likely to be affected by these changes. We use the firms AES responses to estimate the size of the adjustment required to IR10 components. IR10-only firms that adopt the right-of-use rules will have inconsistent filing of K over time, but we expect this group to be small because: AES is targeted at the firms that are likely adopters; and testing suggests it is difficult to identify clear cases of right-of-use adoption in the IR10-only sample.

Appendix B – Estimated coefficients

Table 10: Estimated coefficients for equation (1) with pooled industry

	y		m		k	
	(1)	(2)	(3)	(4)	(5)	(6)
s_{it}	0.8316**	0.6833**	0.2476**	0.2680**	0.1585**	0.1980**
	[0.0018]	[0.0023]	[0.0014]	[0.0013]	[0.0017]	[0.0013]
p_{it}	0.1175**	0.1813**	0.6852**	0.5364**	0.5241**	0.3450**
	[0.0016]	[0.0020]	[0.0014]	[0.0015]	[0.0017]	[0.0014]
$\delta(P=0) \times s_{it}$	-0.0513**	0.0075	0.2253**	0.1868**	-0.0014	0.0615**
	[0.0138]	[0.0142]	[0.0262]	[0.0200]	[0.0196]	[0.0138]
$\delta(S=0) \times p_{it}$	0.4422**	0.2447**	-0.0623**	-0.0474**	-0.2010**	-0.0963**
	[0.0183]	[0.0165]	[0.0076]	[0.0066]	[0.0077]	[0.0056]
$\delta(P=0)$	1.5939**	1.3106**	1.4858**	1.2297**	3.9829**	1.7097**
	[0.1474]	[0.1506]	[0.2727]	[0.2075]	[0.2000]	[0.1438]
$\delta(S=0)$	-0.8759**	0.3711*	2.7304**	2.6626**	3.5728**	2.5657**
	[0.1756]	[0.1570]	[0.0725]	[0.0626]	[0.0759]	[0.0546]
Adjusted R^2	0.760	0.425	0.823	0.465	0.656	0.256
Controls for: year	Y	Y	Y	Y	Y	Y
4-digit industry	Y	N	Y	N	Y	N
Firm fixed effects	N	Y	N	Y	N	Y

N(observations) is 4,293,390. Coefficients not estimated separately by t. (Unreported) OLS specifications with only year controls have adjusted R^2 of 0.731 (y), 0.766 (m), 0.507 (k). Robust standard errors clustered on firm (*;** indicates significance at the 5%;1% level).

Table 11: Estimated coefficients for equation (3) with pooled industry

	y		m		k	
	(1)	(2)	(3)	(4)	(5)	(6)
l_t	0.4679**	0.4893**	0.6105**	0.4527**	0.5144**	0.3357**
	[0.0078]	[0.0059]	[0.0059]	[0.0043]	[0.0053]	[0.0037]
$l_t \times \delta(\mathrm{WP} > 0 \wedge \mathrm{FTE} > 0)$	0.5502**	0.2700**	0.3630**	0.1981**	0.3642**	0.2406**
	[0.0080]	[0.0058]	[0.0063]	[0.0045]	[0.0057]	[0.0039]
$l_t \times \delta(\mathrm{WP} = 0)$	0.3748**	0.0705**	0.1636**	-0.0368**	0.1753**	0.0392**
	[0.0077]	[0.0058]	[0.0060]	[0.0042]	[0.0054]	[0.0036]
$l_t \times \delta(L_{t-1} = 0)$	-0.1303**	-0.003	-0.2035**	-0.0261**	-0.1946**	-0.0370**
	[0.0044]	[0.0039]	[0.0039]	[0.0031]	[0.0037]	[0.0028]
$l_t \times \delta(L_{t+1} = 0)$	-0.2563**	-0.0865**	-0.3345**	-0.1213**	-0.3058**	-0.1279**
	[0.0060]	[0.0059]	[0.0051]	[0.0047]	[0.0045]	[0.0035]
$\delta(\mathrm{WP} > 0 \wedge \mathrm{FTE} > 0)$	0.8157**	0.3276**	0.7599**	0.2996**	0.3584**	0.1376**
	[0.0046]	[0.0039]	[0.0045]	[0.0035]	[0.0042]	[0.0030]
$\delta(\mathrm{WP} = 0)$	1.3027**	0.7094**	1.3891**	0.7606**	1.0218**	0.6305**
	[0.0050]	[0.0059]	[0.0049]	[0.0050]	[0.0046]	[0.0047]
$\delta(L_{t-1} = 0)$	-0.1419**	-0.2821**	-0.0734**	-0.2221**	-0.2035**	-0.3181**
	[0.0043]	[0.0036]	[0.0035]	[0.0028]	[0.0032]	[0.0025]
$\delta(L_{t+1} = 0)$	-0.4286**	-0.3485**	-0.2922**	-0.2482**	-0.2631**	-0.2887**
	[0.0057]	[0.0051]	[0.0047]	[0.0040]	[0.0041]	[0.0032]
Adjusted R^2	0.456	0.131	0.517	0.154	0.535	0.151
Controls for: year	Y	Y	Y	Y	Y	Y
4-digit industry	Y	N	Y	N	Y	N
Firm fixed effects	N	Y	N	Y	N	Y

N(observations) is 4,506,405. Coefficients not estimated separately by t. Reference groups are WP-only firms (ie, $\mathrm{FTE} = 0$) and ongoing firms (ie, $L_{t-1} > 0 \wedge L_{t+1} > 0$). (Unreported) OLS specifications with only year controls have adjusted R^2 of 0.387 (y), 0.442 (m), 0.356 (k). Robust standard errors clustered on firm (*;** indicates significance at the 5%;1% level).

Figure 12: Estimated coefficients over t (pooled industry & firm fixed effects)

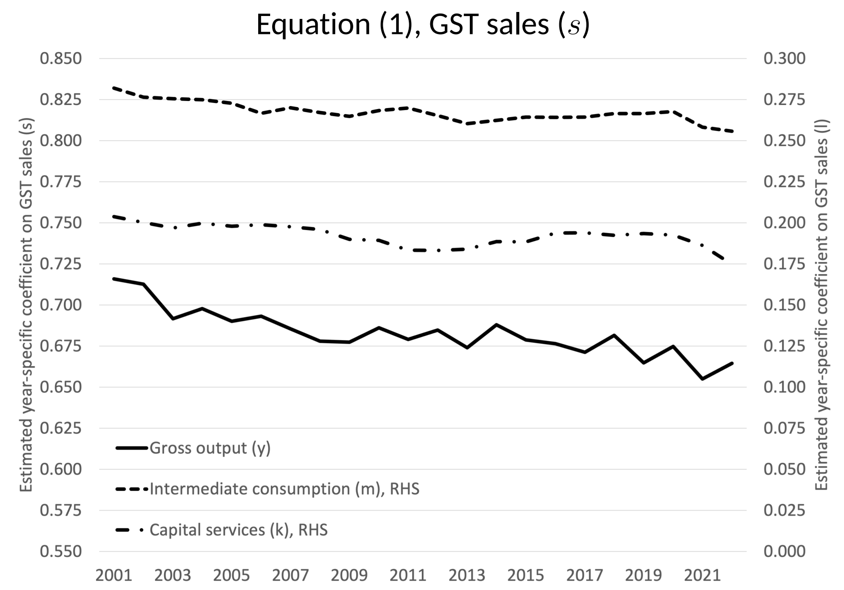

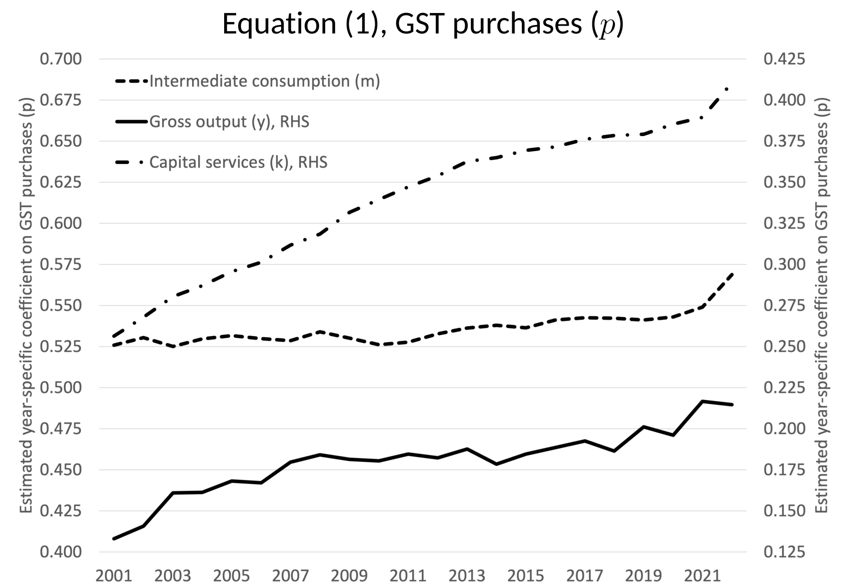

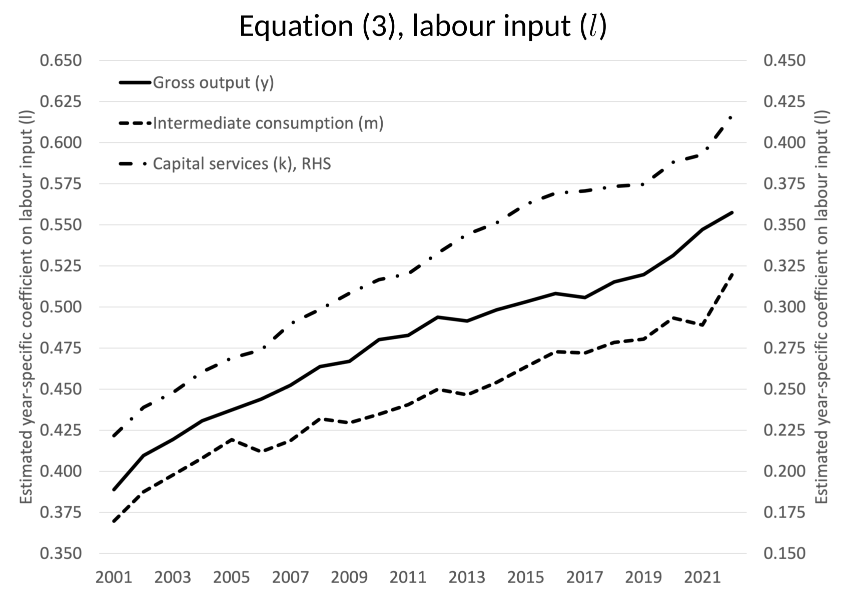

Motu

economic & public policy research

for other Motu working papers: www.motu.nz