e61 Research Note | Number 9 | 18 June 2026

IT'S ABOUT IP AND NOT IT: KNOWLEDGE CAPITAL AND THE AUSTRALIA-US PRODUCTIVITY GAP

Gianni La Cava

Executive summary

Market-sector labour productivity in Australia has fallen behind the United States, particularly over the past decade. Productivity is now around 12 per cent higher in the US than in Australia. Almost all of the divergence is concentrated in two sectors —manufacturing and information — where US firms have built much larger stocks of intellectual property (IP) capital, particularly in research and development (R&D) and software.

The productivity gap partly reflects lower rates of IP investment in Australia. But the evidence in this note suggests a more important explanation: the limited diffusion of knowledge capital to Australian firms and industries. US-owned subsidiaries operating in Australia invest in IP at close to US rates and are significantly more productive than Australian-owned firms in the same sectors. Frontier US technology is already present in Australia, both through cloud and AI infrastructure and the local operations of multinational firms, but this has not closed the productivity gap for domestic firms.

This pattern is not confined to manufacturing and information. Across a wide range of industries, sectors where Australian firms have fallen behind in IP capital intensity also tend to be those where the productivity gap has widened. Controlling for within-industry differences in IP capital accounts for much of the aggregate divergence over the 21st century.

Australia's productivity challenge is not just about increasing IP investment, but about improving the ability of domestic firms to adopt and apply existing technologies. This points to the importance of management capability, skills, organisational capacity and financing systems that support firms scaling into adoption. Labour market and competition settings that facilitate knowledge diffusion are also likely to matter. Overall, the evidence suggests that Australia's productivity problem is as much about how knowledge is used and diffused as where it is created.

1. Australia's 21st century productivity gap

It is well established that productivity growth in Australia has slowed over recent decades (Andrews et al 2022, Duretto et al 2021). Part of this slowdown reflects a shift in economic activity towards the non-market sector, including education and care (Maltman and Rankin 2024). But even within the market sector, Australia's productivity performance has diverged from that of the United States.

A key feature of this divergence is its recency. Australian and US market-sector labour productivity (real output per hour worked) tracked each other closely for much of the 21st century. Since around 2016, however, US productivity has grown about 12 percentage points more than Australia's, opening up effectively the entire post-2000 gap in less than a decade (Figure 1). This gap is evident in both labour productivity and total factor productivity. Understanding its causes matters for policymakers: closing the gap would deliver substantial gains in Australian living standards.

Several hypotheses have been proposed, including differences in industry composition (Dolman et al., 2007), capital investment (Stanford, 2025), business dynamism (OECD, 2021), global integration (Kirchner, 2020) and the environment for innovation (IMF, 2021). This research note incorporates elements of each of these hypotheses in that it documents a strong empirical regularity: industries where Australian firms have fallen behind the United States in IP capital intensity are also those where the productivity gap has widened. The pattern holds across a wide range of sectors and not just in the most knowledge-intensive industries.

I thank Kevin Fox, Luci Ellis, Ewan Rankin, Colin Burns and Lachlan Vass for helpful comments and suggestions. Replication files are available upon request.

Figure 1: Market Sector Labour Productivity

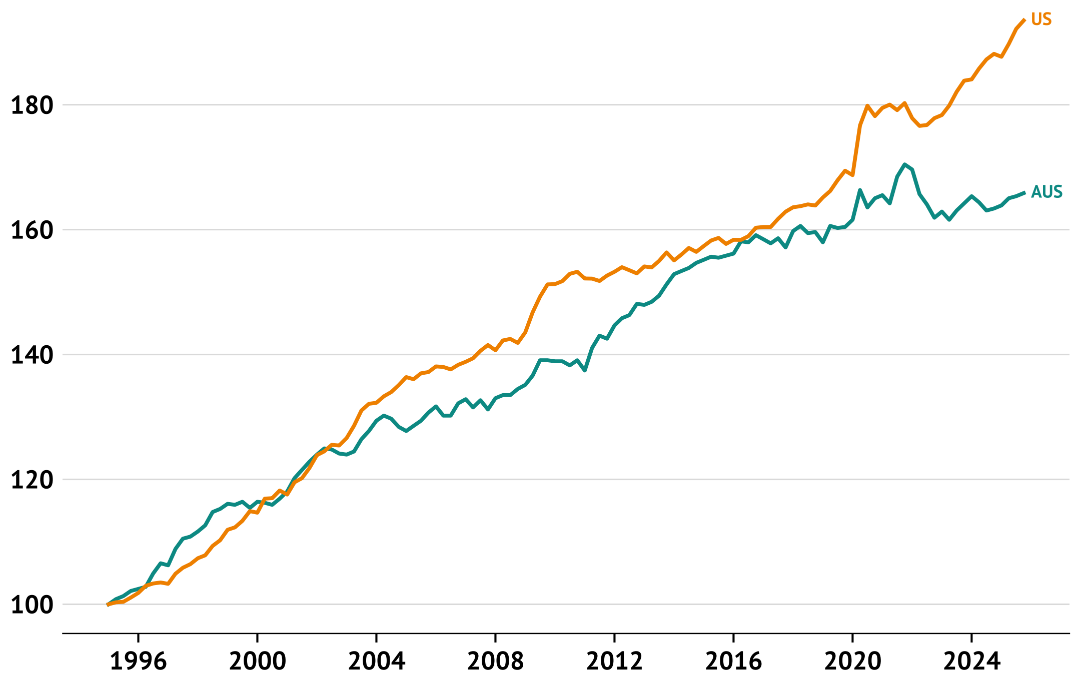

Sources: ABS; BLS; e61 Institute.

2. From Solow's paradox to knowledge capital

Nobel laureate Robert Solow famously quipped in 1987, 'You can see the computer age everywhere except in productivity statistics.' This remark, often dubbed the Solow Paradox, highlights that investing in information technology hardware is not enough to guarantee higher productivity. Decades later, Australia's experience echoes this point. Despite broadly adopting IT systems and spending substantially on machinery and equipment, productivity gains have been modest.

This research note suggests that what matters for productivity is not physical IT capital alone — computers and telecommunication equipment — but the broader stock of *intangible or knowledge capital*: the ideas, software, organisational know-how, and innovations that drive new products and processes (Romer, 1994). The distinction matters because knowledge capital, unlike a machine, is infinitely usable (or 'non-rival'). A piece of software can be used across an entire firm or industry at near-zero marginal cost, generating returns that compound over time.

Knowledge capital includes products that Australian businesses use every day: Microsoft Azure and Amazon Web Services in cloud computing; Salesforce and Workday in business software; Oracle and SAP in enterprise resource planning; Pfizer and Merck in pharmaceutical R&D; Nvidia in GPU compute and the CUDA software stack that underpins most modern AI systems. Each of these is intangible capital created in the United States that, once built, can be rented or licensed at near-zero marginal cost to firms anywhere in the world. The productivity payoff therefore accrues not only to the firms that fund the investment but to the much wider set of firms that adopt it.

The United States' superior productivity performance has coincided with a widening gap in IP investment intensity. American businesses invest a substantially larger share of output in R&D and software — the two largest components of IP investment (Figure 2). By contrast, Australian firms have historically spent more on tangible assets such as non-residential construction. Total business investment as a share of GDP has been broadly similar in the two countries, but the composition has diverged: the US has shifted toward intangibles, while Australia's investment remains tilted toward physical assets.

It is important to recognise that part of Australia's tilt toward physical capital reflects features of its economic structure that are difficult to change. For instance, Australian industry is largely dominated by mining and construction which in turn reflects Australia's high natural resource endowments and relatively fast population growth. The analysis in this note will account for some of these differences in economic structure.

Figure 2: Market Sector Non-Residential Business Investment (% of market-sector GVA)

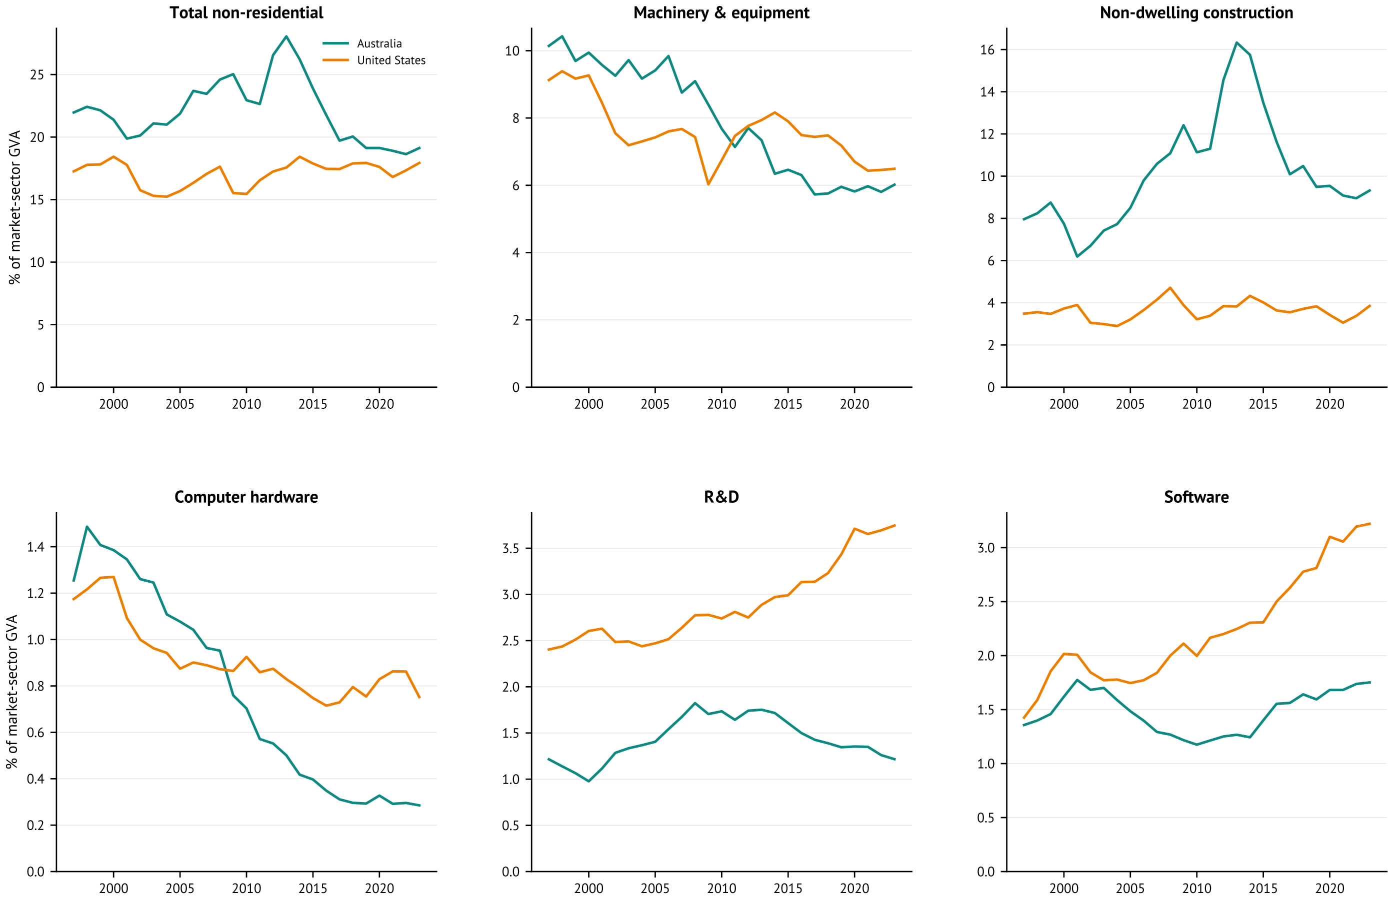

Notes: Investment and GVA are summed across the 15 market-sector divisions (ANZSIC / NAICS 1-digit, excluding real estate, public administration, education and health) for each country.
Sources: ABS; BEA; e61 Institute.

3. The productivity gap — it's about IP and not IT

The aggregate productivity gap between Australia and the US is largely accounted for by two sectors — manufacturing and information (Figure 3, top row). Manufacturing is a broad sector that includes everything from food processing to semiconductor fabrication. Information encompasses IT services, media, telecommunications and data-related services. Together, these sectors account for around 30 per cent of US market-sector output but explain nearly all of the cross-country divergence in labour productivity growth since 2000. Excluding them, the productivity trajectories of the two countries are broadly similar — Australia has, in fact, slightly outpaced the United States.[1]

What sets these sectors apart is their IP investment intensity (Figure 3, bottom row). US firms in manufacturing and information invest nearly 20 per cent of sectoral gross value added in IP products, compared with around 5 per cent in Australia — a fourfold gap. In the rest of the market sector, IP investment rates have historically been broadly similar across the two countries, although the US has accelerated relative to Australia over the past decade.

The same concentration is evident in the *stock* of IP capital built up over recent decades.[2] Manufacturing and information account for over 60 per cent of the US market-sector IP capital stock, reflecting decades of accumulated investment in R&D, software and other intangibles (Figure 4). In Australia, the same sectors account for around a quarter of the IP capital stock, with a much larger share in professional services and finance. The United States has built its knowledge capital in sectors

1 See Appendix A for the underlying productivity trends across individual industries for the two countries.
2 The Australian national accounts classify mineral exploration as intellectual property, following the SNA 2008 framework. The US Bureau of Economic Analysis classifies it as structures. To ensure cross-country consistency, mineral exploration is excluded from IP throughout this note. Including it would increase Australia's measured IP stock in mining but would not materially change the comparison in other sectors.

Figure 3: Labour Productivity and IP Investment by Market Sector

Sources: ABS; BEA; BLS; e61 Institute.

that have grown fastest globally over the past 25 years; Australia's knowledge capital is more dispersed and more oriented toward the domestic economy.

Three questions follow from this pattern. First, does the link between IP capital and productivity hold *outside* of manufacturing and information? Second, is the US productivity advantage in these sectors due to a handful of dominant 'superstar' firms? And third, does the gap reflect Australian firms' lack of access to frontier US technology, or their inability to absorb it? The third question turns out to have a striking answer, previewed here and developed in Section 8: US-owned subsidiaries operating in Australian manufacturing and information are around 1.5–2x as productive as Australian-owned firms in the same sectors. Frontier US technology is therefore already here. The productivity gap seems to be more a story about what Australian-owned firms do with it.

Figure 4: Composition of Market Sector IP Capital Stock (2020--2023 average)

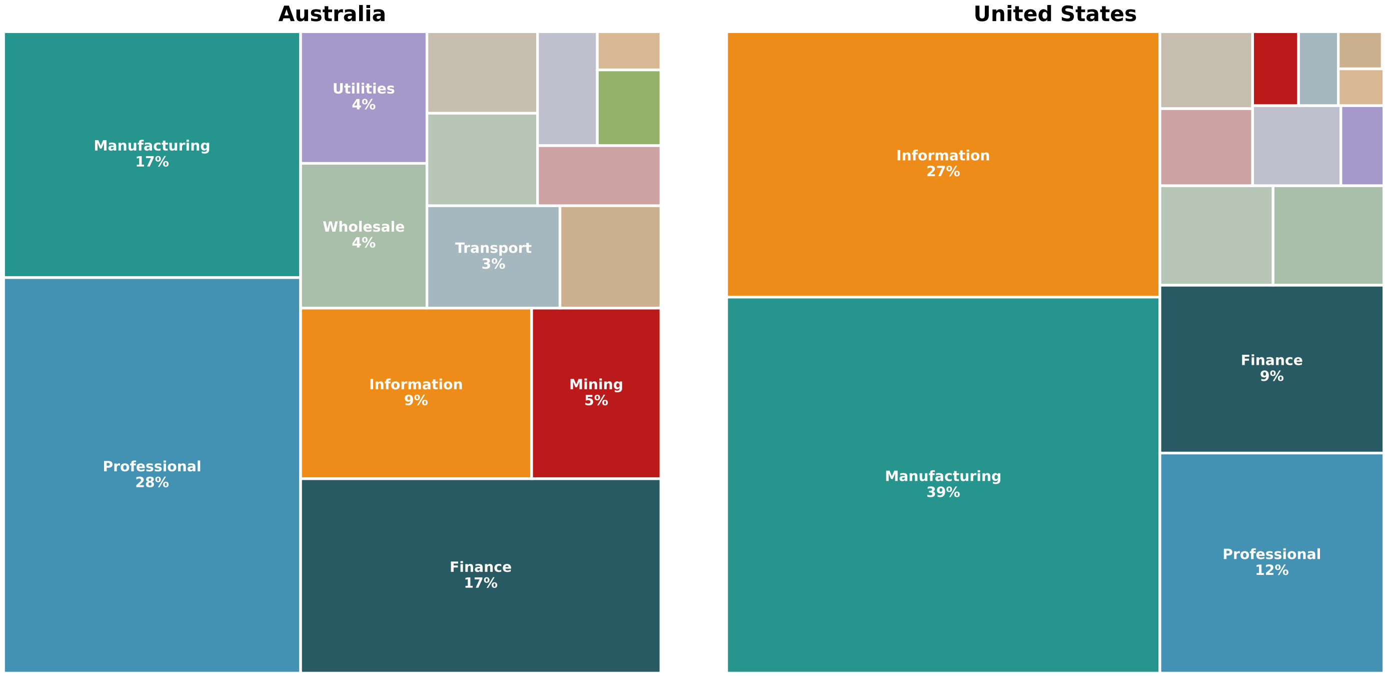

Sources: ABS; BEA; e61 Institute.

Box 1: Is Australia a laggard or the US a leader?

An obvious question is whether the productivity and IP gaps reflect Australian underperformance, US outperformance, or both. International benchmarking against other OECD economies suggests the answer depends on the sector. The right benchmark for that question is the *level* of sectoral labour productivity (PPP-adjusted value added per hour worked): fast growth from a low base does not make a country a productivity leader, and slow growth from a frontier level does not make a country a laggard.

Manufacturing: Australia is the laggard. In level terms, Australian manufacturing labour productivity sits well below the OECD median of advanced economies, while the US sits near the top of the high-income pack. The growth pattern is consistent with that level picture: Australia has gradually fallen further behind the high-income median, partly reflecting its orientation toward resource-related manufacturing, which typically operates at lower productivity levels. (A few catch-up economies in eastern Europe have grown faster than the US, but they remain well below US and Australian levels.) The investment patterns mirror the level story: the US is a global leader in manufacturing IP intensity, while Australia ranks near the bottom of the OECD (Figure 5, right panel).

Information: the US is the outlier. In the information sector, the US is at the productivity level frontier (Figure 6, left panel): US value added per hour is around USD 167 in PPP terms, around 50 per cent above the next-best high-income economies (Belgium, Switzerland, Sweden, Norway). Australia sits roughly at the high-income OECD median on the ANZSIC division J definition — though on the broader ISIC scope used by other OECD countries (which includes IT services and computer programming, sectors that ANZSIC places in division M), AUS ranks lower. Either way, the US lead over Australia is large and unmatched by any other peer economy. The US level advantage reflects its uniquely high rate of IP investment (Figure 6, right panel). Australia is mid-pack on information-sector IP intensity.

Overall, Australia's gap with the United States reflects both a laggard and leader story. In manufacturing, Australia sits below the high-income OECD pack on productivity levels and IP intensity. In information, the US is at the global frontier and Australia is below the median. In both cases, where Australia ranks on IP investment intensity maps closely onto where it ranks on productivity.

Figure 5: OECD: Manufacturing Sector Productivity Level and IP Investment Intensity

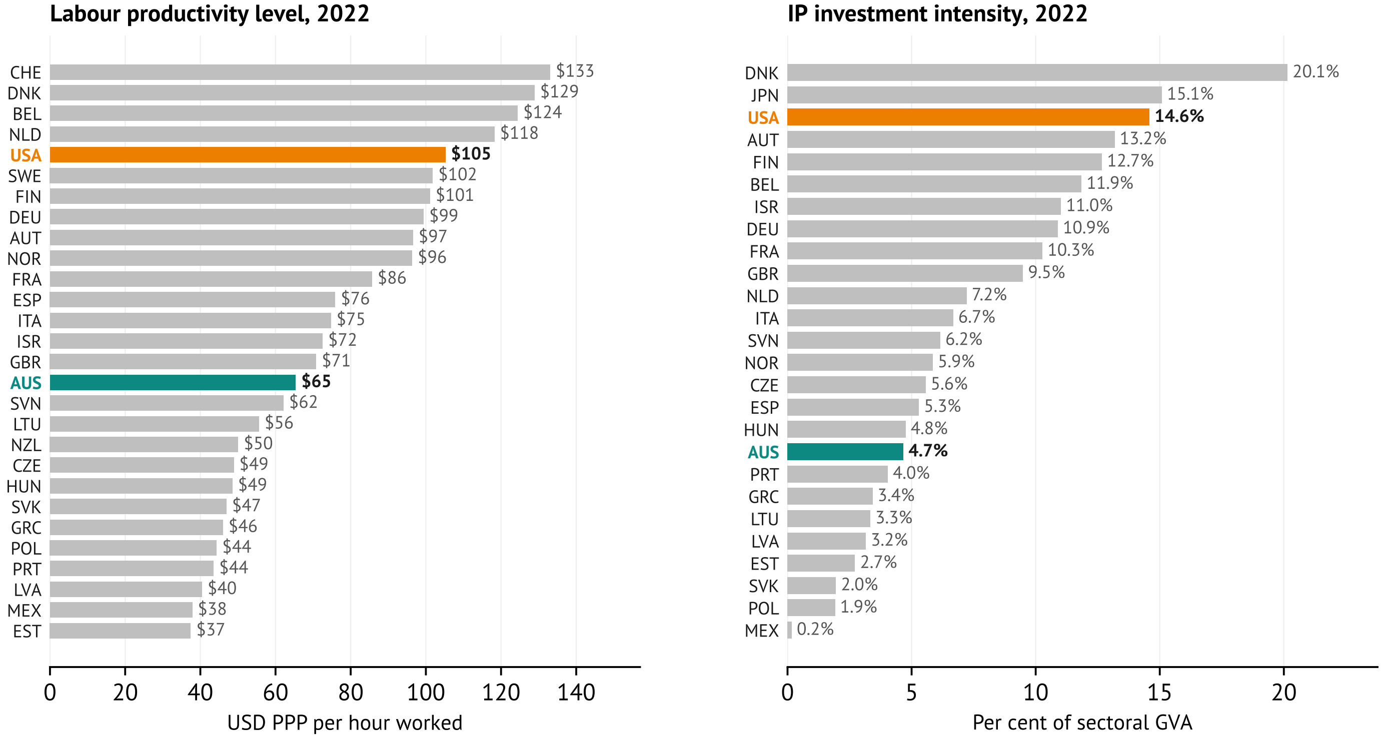

Notes: Left panel shows the level of labour productivity in 2022 (USD per hour worked, PPP-adjusted) ranked across OECD countries with available data, with Australia and the United States highlighted. Ireland, Luxembourg, Chile and Iceland are excluded from the level comparison because their measured value added per hour is distorted by multinational tax artefacts (IRL, LUX) or small-economy data-quality issues (CHL, ISL). Right panel shows IP investment as a share of sectoral gross value added; the AUS value is computed from ABS GFCF (5204064) on the ANZSIC scope, with mineral exploration excluded for cross-country consistency, while other countries are from OECD STAN.
Sources: ABS; BEA; BLS; OECD; e61 Institute.

4. The importance of knowledge investment

Section 3 showed that the productivity gap is concentrated in the same sectors where the IP investment gap is largest. But a sector-level correlation could reflect many things — industry structure, comparative advantage, or the fact that the US is home to a handful of dominant technology firms. This section explores whether the relationship between IP capital and productivity holds *within* industries, and whether it extends beyond manufacturing and information. If it does, then the IP–productivity link is a general phenomenon and not an artefact of a few outlier sectors. This would mean that closing Australia's IP gap matters for productivity everywhere, not just in the two sectors that dominate the aggregate.

The discussion so far has focused on labour productivity, but the same broad pattern holds for total factor productivity (TFP), which abstracts from differences in labour and capital inputs and isolates the residual contribution of technology, business organisation and intangibles. Industry-level TFP trends for Australia and the United States are shown in Appendix A — the gaps are concentrated in the same sectors that drive the labour productivity gap, with the largest divergences in manufacturing and information.

To pin down the link between IP capital and productivity, industry-level regressions are estimated (see Box 1 for the framework and data). The industry regressions control for all time-invariant differences between sectors, including structural factors like Australia's resource endowment or the US's larger domestic market. What remains is within-industry variation: changes over time in how far Australia has fallen behind (or caught up with) the United States in both IP capital and productivity.

Figure 6: OECD: Information Sector Productivity Level and IP Investment Intensity

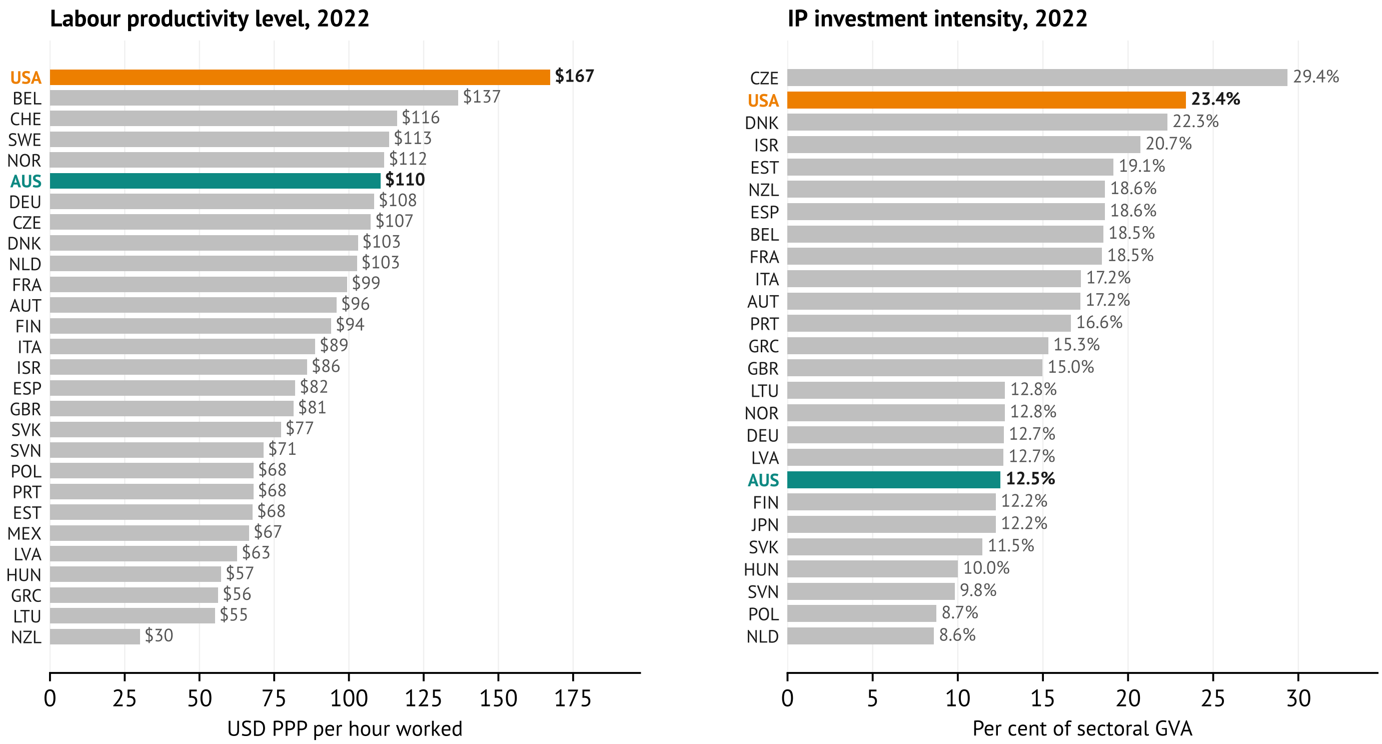

Notes: Left panel shows the level of labour productivity in 2022 (USD per hour worked, PPP-adjusted) ranked across OECD countries with available data, with Australia and the United States highlighted. The Australian value is based on ANZSIC division J (information media and telecommunications), which is narrower than the ISIC Rev.4 division J definition used by other OECD countries: computer programming and IT consultancy services (ISIC 62) fall under ANZSIC division M rather than division J. The narrower ANZSIC scope for AUS therefore likely overstates Australia's relative position in the cross-country comparison; the broader ISIC scope used by other countries would place AUS lower in the ranking. Ireland, Luxembourg, Chile and Iceland are excluded from the level comparison for the reasons noted in Figure 5. Right panel shows IP investment as a share of sectoral GVA; for comparability, the AUS series in the right panel is sourced from OECD (SNA Table 8 IPP investment over STAN nominal value added) on the ISIC scope, rather than from the ABS ANZSIC measure used elsewhere in the note.
Sources: ABS; BEA; BLS; OECD; e61 Institute.

Box 2: Framework and industry-level data

Framework. The regression analysis is based on the endogenous growth literature (Aghion & Howitt, 1992, 1997) and empirical work linking R&D to productivity at the firm and industry levels (Griffith et al., 2004; Ugur et al., 2016). Countries are indexed by $j = \{AUS, US\}$, industries by $i = 1, \ldots, N$, and financial years by t. Value-added (Y) in each industry is given by a neoclassical production function:

$$Y^j_{it} = A^j_{it} F^j(K^j_{it}, L^j_{it}) \quad (1)$$

where K is the total capital stock — including both tangible (machinery, equipment, structures) and intangible (intellectual property products) capital — L is labour, and A is total factor productivity (TFP). TFP is assumed to depend on the stock of IP capital (G) and other factors (X) that vary across industries, countries, and time:

$$A^j_{it} = \Psi(G^j_{it}, X^j_{it}) \quad (2)$$

Taking logs of the ratio of Australian to US TFP yields the relative TFP level (or 'productivity gap') for industry i in year t:

$$\hat{a}_{it} = \eta \hat{g}_{it} + \gamma \hat{X}_{it} + \mu_{it} \quad (3)$$

where $\hat{a}_{it} = \ln(A^{AUS}_{it}/A^{US}_{it})$ is the productivity gap, and $\hat{g}_{it} = \ln(G^{AUS}_{it}/G^{US}_{it})$ is the relative IP capital stock (the 'IP capital gap'). The vector $\hat{X}_{it}$ includes controls such as relative human capital and trade exposure. Writing the relationship in log-ratio form imposes a common elasticity η across the two countries: a given proportional change in IP capital is assumed to yield the same proportional change in productivity in Australia as in the United States. If the true returns to IP differ — for example, because Australian firms have weaker absorptive capacity — the estimated coefficient should be interpreted as a weighted average of the two countries' true elasticities, with the weighting determined by the data. The error term μ_{it} is decomposed into industry fixed effects (α_i), decade-specific time effects (λ_t), and an idiosyncratic shock (ϵ_{it}):

$$\mu_{it} = \alpha_i + \lambda_t + \epsilon_{it} \tag{4}$$

The key parameters of interest are η, which measures how differences in IP capital are associated with differences in productivity, and λ_t, which capture the decade-specific time trends in the productivity gap. Because IP capital is already embedded in K, the coefficient η should be interpreted as the productivity return to IP capital *above and beyond* its direct factor share — that is, the spillovers and non-rival returns that arise from the diffusion of knowledge across firms, industries and the broader economy.[a]

Data. The model is estimated using harmonised industry-level data for Australia (Australian Bureau of Statistics) and the United States (Bureau of Economic Analysis and Bureau of Labor Statistics). The dataset includes for each industry and year:

- **Output:** gross value-added and gross output, in both nominal and real terms.
- **Employment:** hours worked and headcount.
- **Capital:** net capital stock and gross investment, in both nominal and real terms, split into machinery and equipment, non-dwelling construction, and intellectual property products (IPP). IPP is further decomposed into R&D, software and artistic originals. Mineral exploration is excluded from IP to ensure consistency between Australian and US definitions.

The sample includes all industries in the market sector, covering the 1-digit ANZSIC divisions in Australia except education, health and public administration. Real estate is also excluded due to differences in coverage. The US industries are based on the NAICS classification and are broadly equivalent at the 1-digit level. The sample covers FY1996/97 to 2023/24, giving more than 500 industry-year observations for each country.

Constructing the TFP gap. The model requires an industry-year estimate of the TFP gap.[b] I construct this index using a translog production framework with smoothed labour shares,[c] which compares Australian and US output, capital, and labour input for each industry:

$$\hat{a}_{it} = \hat{y}_{it} - \bar{\sigma}_i \hat{l}_{it} - (1 - \bar{\sigma}_i)\hat{k}_{it} \tag{5}$$

where $\hat{y}_{it} = \ln(Y_{it}^{AUS}/Y_{it}^{US})$, $\hat{k}_{it} = \ln(K_{it}^{AUS}/K_{it}^{US})$ and $\hat{l}_{it} = \ln(L_{it}^{AUS}/L_{it}^{US})$ are the cross-country gaps in real output, real capital, and hours worked. The variable $\bar{\sigma}_i$ is the average labour share in Australia and the US for each industry over the sample period. The capital stock used to construct the TFP gap includes all forms of capital, so η is identified from variation in the *stock* of IP capital beyond its direct factor share.

Differences in price and output measurement, particularly in ICT-intensive sectors, may bias the US productivity advantage upward: the BEA applies hedonic quality adjustments to a wider set of ICT categories than the ABS does, which inflates measured real output growth in the US relative to Australia. The broad patterns documented in this note are unlikely to be overturned by these adjustments, but their magnitude may be affected. The analysis also does not adjust for cross-country differences in productivity based on purchasing power parity, as the focus here is on productivity as a supply-side concept; PPP-adjusted estimates are reported in Appendix D and make little difference to the long-run gap.

a Data limitations make it difficult to include estimates of industry-level human capital and international trade that are comparable on a cross-country basis, so the analysis in this note assumes that $\gamma = 0$. However, proxies for both industry-level labour quality and international trade exposure are included in the robustness tests in the Appendix and make little difference to the key results.
b The published estimates of industry productivity differ between the US and Australia. ABS labour productivity is typically based on gross value-added while the BLS estimates for the US are based on gross output. The BEA industry estimates of gross value-added are combined with the BLS estimates of hours worked to make the cross-country estimates more comparable.
c Following (Caves et al., 1982) and (Harrigan, 1997).

The results show a consistent relationship: industries where Australian firms have less IP capital than their American counterparts tend to have a larger productivity gap (Figure 7). Table 1 presents the regression estimates. The baseline spec-

Figure 7: US--AUS Gap in Total Factor Productivity and IP Capital Intensity by Industry

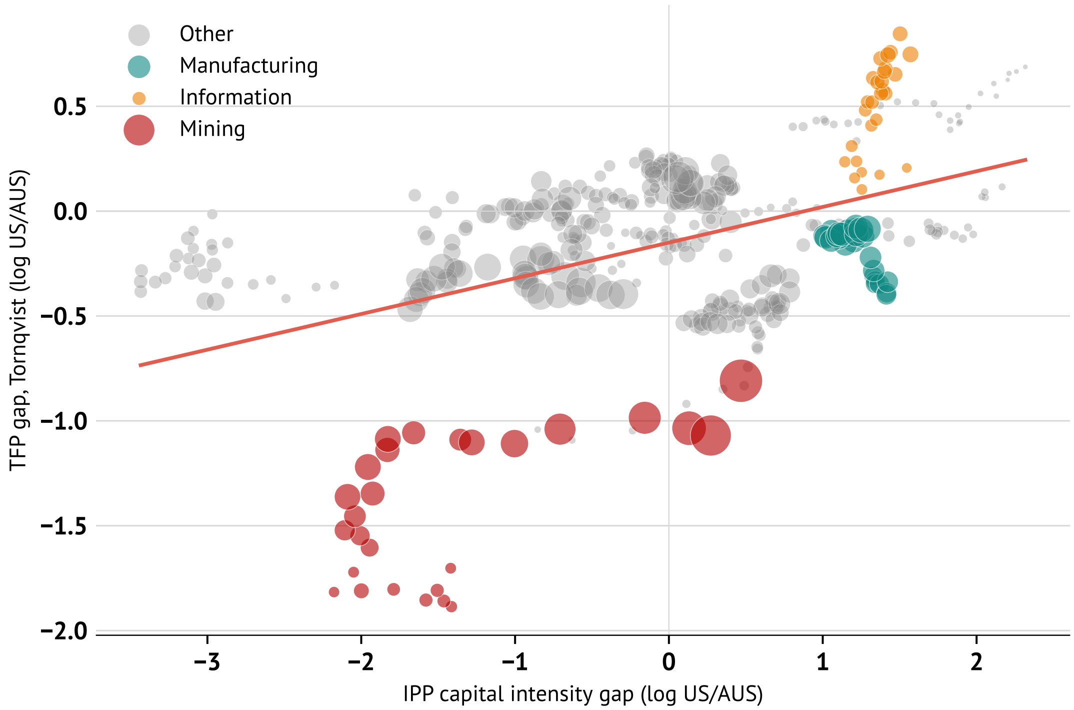

Notes: Each point is an industry–year observation. The y-axis is the TFP gap (log US/AUS) and the x-axis is the IPP capital intensity gap (log US/AUS). Bubble size is proportional to Australian nominal GVA. The solid line is the OLS fit.

ification implies that a 1 per cent widening of Australia's IP capital gap relative to the United States is associated with a 0.17–0.19 per cent widening of the TFP gap.

Critically, this relationship is not driven by manufacturing and information. Columns 3–4 and 7–8 exclude both sectors from the sample, and the estimated coefficient is essentially unchanged. Differences in IP capital intensity are correlated with productivity differences across a wide range of industries — not just the two that account for most of the aggregate gap. The aggregate productivity gap may be concentrated, but the underlying IP–productivity relationship is general.

The relationship is weaker in specifications that exclude Australia's mining sector: the GVA-weighted coefficient falls from around 0.17 to 0.04 (see Appendix C). Part of this sensitivity is mechanical rather than substantive.[3]

The estimated decade fixed effects provide a measure of how much the aggregate productivity gap has widened, and how much of that widening is associated with IP. Under the GVA-weighted specification, the within-industry TFP gap widened by roughly 10 percentage points between the late 1990s and the 2020s. Controlling for IP capital reduces the unexplained widening to about 4 percentage points (Figure 8). The growing divergence in IP capital is associated with more than half of the within-industry productivity widening.

These estimates are associations, not causal effects. Reverse causality is a potential concern: sectors that experience faster productivity growth generate higher profits, which funds more IP investment. Omitted factors are also plausibly correlated with IP intensity, such as human capital, management quality, and trade openness. However, proxies for these variables are included in robustness tests in Appendix C and do not overturn the main results.

Despite these caveats, the pattern is consistent. The OLS specification is positive and significant across every sample restriction considered, and aligns with a large international literature linking intangible investment to productivity (Byrne et al., 2016; Corrado et al., 2009; Haskel & Westlake, 2018).

3 Australia's mining sector is large by GVA and gets heavy weight in the WLS; its measured IP capital is also artificially low because mineral exploration — which the ABS classifies as IPP but the BEA does not — has been excluded for cross-country comparability. Including mineral exploration symmetrically would narrow the measured AUS–US IP gap in mining and raise the WLS coefficient. The equally-weighted OLS result is not sensitive to mining and remains positive and significant across all sample restrictions, so the qualitative finding is robust; the magnitude under GVA weighting should be read as partly reflecting both Australia's industrial structure and the mineral-exploration measurement convention.

Table 1: Industry-Level Regressions: TFP Gap and IP Capital Gap

	OLS				WLS (GVA-weighted)			
	(1) All	(2) All	(3) Excl. C, J	(4) Excl. C, J	(5) All	(6) All	(7) Excl. C, J	(8) Excl. C, J
IP capital gap ($\hat{\eta}$)		0.192***		0.187***		0.172***		0.184***
		(0.048)		(0.056)		(0.042)		(0.045)
2000s	0.012	0.012	−0.003	−0.006	0.009	0.011	−0.028	−0.028
	(0.038)	(0.033)	(0.042)	(0.035)	(0.037)	(0.038)	(0.037)	(0.037)
2010s	0.065	0.101	0.024	0.062	0.052	0.080	−0.011	0.016
	(0.076)	(0.072)	(0.081)	(0.077)	(0.082)	(0.084)	(0.089)	(0.089)
2020s	0.086	0.071	0.036	0.019	0.102	0.042	0.040	−0.036
	(0.089)	(0.071)	(0.094)	(0.070)	(0.107)	(0.089)	(0.121)	(0.088)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	405	405	351	351	405	405	351	351
Industries	15	15	13	13	15	15	13	13

Notes: Dependent variable is the Törnqvist TFP gap (log US/AUS). The IP capital gap is the log ratio of US to Australian intellectual property capital stock per hour worked. Intellectual property is defined as R&D, software and artistic originals; mineral exploration is excluded to ensure consistency with US definitions. All specifications include industry fixed effects (within-industry demeaning). Columns (1)–(4) are estimated by OLS; columns (5)–(8) by weighted least squares with weights proportional to average Australian industry gross value added. Standard errors (in parentheses) are clustered by industry using $t(G-1)$ critical values. Base period: 1997–1999. Columns (3)–(4) and (7)–(8) exclude manufacturing (ANZSIC C) and information (ANZSIC J).
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Sources: ABS; BEA; BLS; e61 Institute.

Figure 8: Average US--AUS Total Factor Productivity Gap by Decade

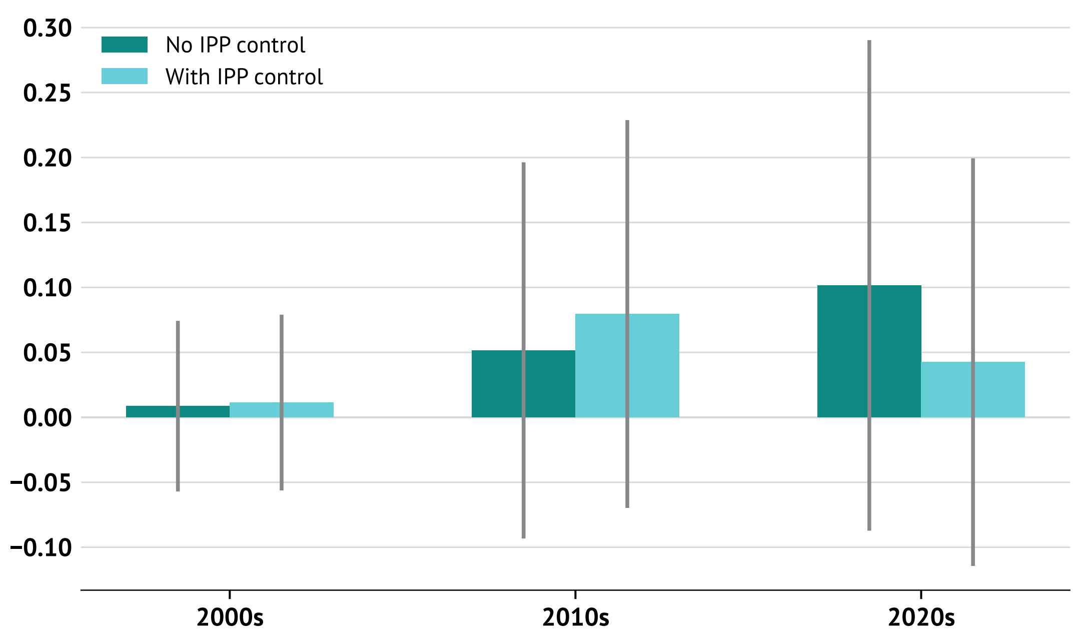

Notes: Decade fixed effects from a WLS regression of the TFP gap on industry fixed effects, with and without the IP capital gap. Weights are proportional to average Australian industry GVA. Base period is 1997–1999. Standard errors clustered at the industry level.

The regression estimates imply that a one per cent widening of the IP capital gap is associated with a 0.17–0.19 per cent widening of the TFP gap. Because the dependent variable is TFP — constructed by netting out all factor inputs including IP capital — the coefficient should be interpreted as the productivity return to IP capital *above and beyond* its direct contribution as a factor input. In standard growth-accounting calculations, the direct factor share of intangible capital is typically also in the 0.10–0.20 range. Adding the two together implies a total social return to IP capital in the order of 30–40 cents per dollar invested, broadly consistent with estimates of R&D spillovers in the international literature (Bloom et al., 2013). These returns are economically large, and reflect the non-rival and cumulative nature of knowledge capital relative to tangible investment.

5. Superstars and scale

Is the US productivity advantage in manufacturing and information driven by a handful of dominant firms, or is it broad-based? To answer this question, we need to turn to alternative sources of data — namely, industry-level market concentration statistics and publicly listed company-level data.

Information: consistent with a meaningful superstar role

There appears to be a meaningful role for large 'superstar' firms in the US information sector, though the evidence for this is indirect. US Census data show that the top four firms in information account for around a third of sector revenue and explain roughly 44 per cent of measured productivity growth within the sector between 2017 and 2022, suggesting that firm-level concentration may be an important part of the story over that short period.

Publicly listed company data also suggest that the US–Australia gap in information-sector productivity growth cumulated to around 168 percentage points between 2000 and 2022. Over the same period, the composition of the largest US firms in the sector changed dramatically: IBM and legacy telecoms were displaced at the top by Alphabet and Microsoft, whose combined revenue grew from near zero to almost half a trillion dollars. The rise of the US information sector is closely associated with the rise of a small number of globally dominant firms, while Australia has no obvious equivalent global leaders.

These estimates should be interpreted as indicative rather than definitive. They are based on nominal revenue-per-worker measures and do not map perfectly to the national-accounts concepts used elsewhere in the note. They are best viewed as illustrating how concentrated productivity gains can be within IP-intensive sectors, rather than as a precise decomposition of the Australia–US productivity gap.

Manufacturing: broad-based, not superstar-driven

Manufacturing is a different story. The US outperformed Australia by less than 2 percentage points in nominal productivity growth between 2017 and 2022, and the top four US manufacturers explain around a third of sector-wide productivity growth — too little, relative to the narrow headline gap, to constitute a meaningful superstar channel. Over the longer 2000–2022 horizon, the cumulative productivity gap reached 32 percentage points in nominal terms, but top-firm concentration was broadly stable at 15–19 per cent of listed-company revenue.

The US manufacturing advantage is also broad-based across sub-industries (Appendix B, Figure B.1). US firms have pulled ahead in machinery and equipment, chemicals and petroleum products, and a broad 'other' category covering clothing, footwear and furniture, while Australian manufacturers have grown faster in food, beverages and several mining-related branches. Part of the aggregate US advantage reflects the larger weight of computer-related manufacturing, but this is a compositional story and not necessarily a firm-level superstar story.

6. R&D concentration and the diffusion challenge

A striking feature of the Australia–US gap is not just the *level* of R&D spending but its *distribution* across firms. Business R&D is heavily concentrated in the largest US firms and has become more so over the past decade: the share of R&D accounted for by the largest US firms has risen from around 2 to 3 per cent of GDP (Figure 9). A handful of US firms — Alphabet, Microsoft, Apple and Meta — each spend more on R&D than Australia's entire business sector combined. In Australia, the opposite trend

Figure 9: R&D Spending by Firm Size (% of GDP)

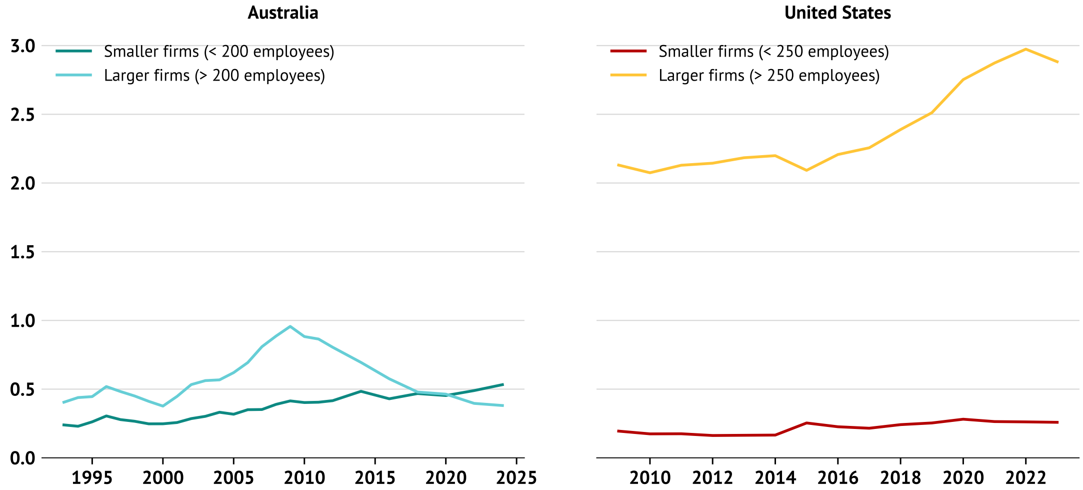

Sources: ABS; BEA; NCSES; e61 Institute.

has taken hold. R&D has become *less* concentrated over time, with small and medium-sized firms (fewer than 200 employees) now accounting for the majority of business R&D.[4]

The Australian tilt toward smaller R&D performers is at least partly a policy-induced pattern. The R&D Tax Incentive (RDTI) provides a more generous refundable offset to firms with aggregated turnover below A$20 million, which skews the effective subsidy toward smaller claimants. The result is an R&D system that spreads support across a broad base of small firms rather than concentrating it on firms with the scale to build and deploy knowledge capital at the frontier.

Whether this design is well-matched to Australia's productivity problem depends on what that problem actually is. The evidence in this note suggests it is largely one of *diffusion* — absorbing knowledge developed elsewhere, particularly in the United States, and redeploying it across the domestic economy. Diffusion is not a passive process. It requires absorptive capacity: specialist management, engineering and research teams, deep links with frontier partners, and the organisational capability to integrate new technology into existing processes. These capabilities are typically easier to build at scale. Large firms can sustain dedicated R&D divisions, in-house technical talent and long-horizon investment; small firms with a handful of researchers may innovate at the margin but are less well-placed to absorb and redeploy a complex technology developed by a US frontier firm.

The international comparison reinforces the point. Small economies that have built globally competitive IP-intensive industries — Taiwan in semiconductors, Sweden in telecommunications and pharmaceuticals, Israel in cyber and software — have done so largely through a small number of firms operating at scale rather than through a broad base of small R&D performers. Australia's own exceptions — Cochlear, CSL and ResMed in medical devices and biotechnology, Atlassian and WiseTech in enterprise software — illustrate the same point from a domestic base.

None of this is an argument against R&D support for small firms; they remain an important source of new firm formation and experimentation. But it does raise the question of whether a system concentrated on small-firm R&D is well-calibrated to the absorptive-capacity requirements of a diffusion-led productivity strategy. A contemporary parallel is taking shape in artificial intelligence: frontier AI is being developed by a small set of very large US firms, and Australia's ability to absorb and redeploy those capabilities will depend less on whether local firms can match the frontier and more on whether the domestic innovation system supports the kind of firms that can deploy those capabilities at scale.

7. Technology transfer from the United States

The productivity gap could be narrowed by making it easier for Australian firms to adopt and adapt ideas developed abroad. Rather than trying to replicate frontier innovation domestically, Australian firms may benefit from absorbing and applying

4 The concentration of R&D spending is even more lopsided in the US information and manufacturing sectors. Large US firms (more than 250 employees) account for more than 90 per cent of R&D in both sectors, compared with 43 and 55 per cent respectively in Australia.

technologies, business methods and know-how generated in the US. This diffusion channel could be as important for productivity as home-grown innovation, particularly in sectors where Australian firms lag behind international best practice.

Box 3: Adoption in practice: three examples

Cloud computing (Australian firms adopting US IP). Australia's IP imports from the US have roughly quadrupled as a share of GDP over the past decade (Figure 10), driven heavily by Australian firms' use of US-built cloud and AI infrastructure. Commonwealth Bank, Telstra, the major airlines and most Australian-headquartered startups now run their core compute on Amazon Web Services, Microsoft Azure or Google Cloud. The capital expenditure that built these platforms — and the R&D that keeps them at the frontier — shows up in the US national accounts, not Australia's; what shows up here is the productivity payoff for the Australian firms that use them. The same hyperscalers' Sydney and Melbourne data-centre announcements bring physical infrastructure onshore, but the underlying IP capital — the chips, software and model weights — remains with the parent firms in the US.

Domino's Australia (a US technology stack delivered through an Australian operator). Domino's Pizza Enterprises holds the master franchise for the Domino's brand across Australia, New Zealand and parts of Europe, and is one of the most automated quick-service operations in the country. The online ordering platform, store-fulfilment systems, GPS-tracked delivery and predictive store management all run on technology developed by the US parent and adapted locally. Sales per store and labour hours per delivery improved sharply through the mid-2010s as the digital ordering platform was rolled out. The case illustrates that the relevant constraint is rarely access to the technology — the underlying IP is licensable — but the management and execution capability needed to deploy and maintain it at scale.

Costco vs Coles and Woolworths (a US affiliate operating alongside Australian incumbents). Costco entered Australian retail in 2009 and now operates a small number of warehouse-format stores. Reported sales per labour hour are several times those of the major Australian supermarkets despite Costco operating under the same wage and regulatory framework. The difference is overwhelmingly a story of organisational design — limited SKU range, large pack sizes, scan-and-go checkout, US-developed inventory management — rather than a different technology endowment. Costco's productivity advantage is also evidence of the limits of pure technology transfer: the Australian incumbents have access to most of the same tools and have not converged.

The common thread is that the binding constraint is rarely access to frontier IP. It is the organisational capability, scale and management quality required to deploy that IP profitably — the absorptive-capacity story that the rest of this section quantifies.

Recent data point to a sharp rise in Australia's purchases of intellectual property from abroad. IP imports as a share of GDP have more than doubled over the past decade to around 0.66 per cent of GDP in FY2024–25 (Figure 10), driven almost entirely by transactions with US firms.[5] The US alone now accounts for roughly 0.5 per cent of GDP, up from 0.1 per cent a decade earlier. For context, total domestic business R&D spending is about 1 per cent of GDP, so IP imports from the US alone are now about half of Australia's total business knowledge investment.

A substantial share of these IP imports likely reflect Australian firms' consumption of US cloud computing and AI services, including Microsoft Azure, Amazon Web Services and Google Cloud. Recent announcements of large-scale data centre investments in Australia by US hyperscalers support this hypothesis: the physical infrastructure sits on Australian soil, but the underlying IP capital — the chips, software and model weights — remains with the parent firms in the US. These patterns suggest that the US will be an important external source of commercial knowledge capital for Australia for the foreseeable future.

Much of the inflow of US technology is also consumed on Australian soil by US firms themselves. Majority-owned affiliates of US multinationals account for around 6 per cent of Australian market-sector value added and 3 per cent of market-sector employment. Their presence is heavily concentrated in the two sectors where the productivity gap is largest: US affiliates produce roughly 16 per cent of value added in both Australian manufacturing and information, compared with around 4 per cent across the rest of the market sector.[6]

5 IP imports are estimated based on the ABS import category called 'charges for the use of intellectual property n.i.e'. A broader measure that includes data processing, computer hardware and software consultancy shows a similar rising trend.

6 Affiliate shares are calculated from BEA *Activities of U.S. Multinational Enterprises* (majority-owned foreign affiliates, Australia) combined with ABS industry value added, employment and business R&D statistics for matching years. BEA affiliate value added reflects transfer-pricing and cost-allocation conventions that do not align perfectly with ABS national accounts, so the ownership decomposition should be read as indicative of orders of magnitude rather than as precise shares.

Figure 10: Intellectual Property Product Imports (% of GDP)

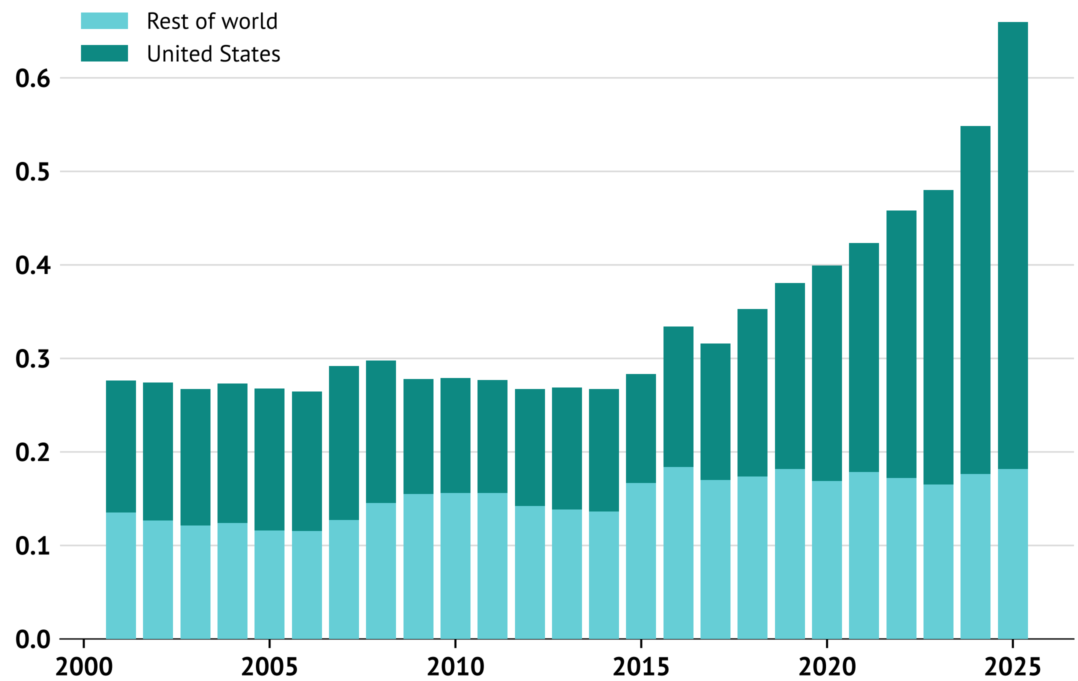

Sources: ABS; e61 Institute.

US affiliates perform approximately 8 per cent of all Australian business R&D, despite accounting for only around 6 per cent of market-sector value added. Put differently, about one in every twelve dollars of business R&D undertaken in Australia occurs inside a US multinational subsidiary. To the extent that the resulting intellectual property is owned, controlled, or commercially exploited by the US parent, this means that a non-trivial share of Australia's measured business R&D is tied to foreign ownership. This share is likely larger once imported IP is included (Figure 10).

Taken together, the two channels of US-linked knowledge investment are substantial. IP imports from the US amount to around 0.5 per cent of GDP, while R&D performed in Australia by US affiliates accounts for a further 0.1 per cent of GDP. In combination, these flows are equivalent to close to three-quarters of the R&D performed by Australian-owned firms. This suggests that the commonly cited figure of around 1 per cent of GDP for domestic business R&D overstates Australia's indigenous commercial knowledge base. A sizeable share of the knowledge investment supporting Australian productivity is either developed abroad or undertaken domestically within US-owned subsidiaries.

A simple ownership decomposition illustrates both the scale and the limits of this channel (Figure 11). In manufacturing, US-owned affiliates operating in Australia have recorded a value added per worker at or above their US parent companies' average in both the 2010s and 2020s, and roughly twice the level of Australian-owned firms in the same sector. In information, US affiliates operate at around half the US aggregate level — reflecting the distribution and cloud-infrastructure character of their local activity rather than frontier product development — but still around 30 per cent above Australian-owned firms on average over the 2020s. In manufacturing in particular, stripping the US-owned slice out of the Australian aggregate widens the domestic-only productivity gap by close to 10 log points. The adjustment is smaller in information — US affiliates there are less productive than the local aggregate in the 2010s — but the broader point holds: the level of Australian productivity in these sectors is noticeably lifted by the presence of US-owned firms.

These patterns point to a particular reading of Australia's productivity problem. Frontier US technology is already physically located in the two sectors where the gap is largest, through the local operations of major US firms. But more than a decade of that presence has not narrowed the productivity gap for Australian-owned firms in the same sectors. If simple co-location with frontier firms were sufficient, the gap should already be smaller than the evidence in Sections 3–5 suggests. The weight of the policy problem therefore shifts from technology access toward *absorptive capacity* — the management, skills and organisational capabilities that determine whether domestic firms can adapt and apply the knowledge flowing in from their American-owned neighbours.

Figure 11: Labour Productivity by Ownership Cohort, Decadal Averages

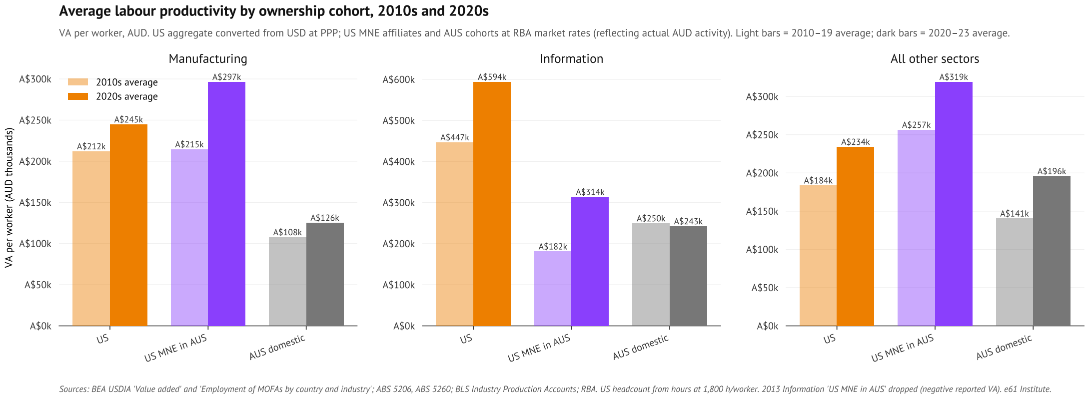

Notes: Average value added per worker for the 2010s (2010–2019) and 2020s (2020–2023). Values expressed in AUD thousands. The US aggregate is converted from USD at PPP; US MNE affiliates in Australia and Australian cohorts are converted at RBA market exchange rates, reflecting the activity actually taking place in AUD. US employment is imputed from BLS hours worked at 1,800 hours per worker. The 'Australia, domestic-owned' cohort is constructed by subtracting US majority-owned affiliate value added and employment from the Australian industry totals. 'All other sectors' aggregates mining, wholesale trade, retail trade, finance and insurance, and professional services. The 2013 observation for US MNE affiliates in information is dropped from the 2010s average because BEA reported a negative value added figure for that year.
Sources: ABS; BEA; BLS; e61 Institute; RBA.

8. Policy implications

Australia's productivity strategy will likely rely more on *diffusion* — adopting and adapting knowledge capital produced elsewhere — than on producing frontier IP at scale. Frontier IP-intensive investment is concentrated in a small set of US firms, Australia has not produced firms at equivalent scale, and Australian firms are already importing substantial IP from the United States. The weight of policy therefore shifts from fostering frontier innovation toward lowering the barriers to diffusion and building the capacity to absorb it.

Financing that is tilted toward small firms

The supply of risk capital is the most important barrier to knowledge investment. R&D is hard to collateralise, generates returns only after long lags and is poorly suited to traditional bank finance. The US has built a deep equity-financing ecosystem — public markets, venture capital, and cash-generative incumbents that self-finance further R&D — that funds knowledge-heavy firms at scale.

Australia's system is thinner and structurally tilted toward small firms. Cumulative venture capital raised by Australian startups since 2000 is around US$34 billion, against roughly US$2.2 trillion in the US — a ratio of about 65:1, against a GDP ratio closer to 14:1.[7] The R&D Tax Incentive reinforces the small-firm tilt through a more generous refundable offset below a A$20 million turnover threshold (Section 6).

That design is reasonable if the objective is to encourage entry into innovation. It is harder to defend if the objective is diffusion, which puts a premium on firms large enough to absorb and redeploy complex technologies from abroad. A system that makes support meaningfully less generous at the scaling phase works against that objective. Ignaszak et al. (2025) find that, per dollar spent, it is *young* rather than *small* firms that deliver the largest boost to aggregate growth — firms with the potential to scale, not firms that happen to be small today. Rebalancing the RDTI toward growth potential rather than firm size, without removing support for new entrants, is a natural place to start.

Absorptive capacity is the binding constraint

US multinationals have operated at scale in Australian manufacturing and information for decades (Section 7), and IP imports from the US have roughly doubled as a share of GDP. If technology access were the binding constraint, the domestic gap would already be smaller. The constraint is more plausibly *absorptive capacity* — the management, skills and organisational capabilities that let firms adapt and apply new technology.

7 VC investment data from Dealroom.co.

The relevant policy levers sit mostly outside the R&D tax system: skilled migration and post-study work rights, non-compete enforcement, STEM education and management training, and competition settings that let innovative firms grow. Frameworks that support spillovers from US affiliates — labour mobility, supplier-development programs, collaborative R&D with domestic firms and universities — are a relatively high-leverage use of public resources. And because R&D in the sectors where the gap is largest is labour- rather than capital-intensive, the availability of researchers, engineers and technicians bears more directly on closing the gap than further capex incentives.

The pattern across industries is consistent: where Australian firms fall behind in knowledge capital, they fall behind in productivity. Closing the gap will depend less on producing Australian equivalents of US frontier firms than on ensuring that domestic firms can adopt and scale the technologies those frontier firms create.

References

Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, *60*(2), 323–351. https://doi.org/None

Aghion, P., & Howitt, P. (1997, December). *Endogenous growth theory* (1st ed., Vol. 1). The MIT Press. https://doi.org/None

Bloom, N., Schankerman, M., & Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, *81*(4), 1347–1393.

Byrne, D. M., Fernald, J. G., & Reinsdorf, M. B. (2016). Does the united states have a productivity slowdown or a measurement problem? *Brookings Papers on Economic Activity*, *47*(1), 109–182. https://doi.org/10.1353/eca.2016.0008

Caves, D. W., Christensen, L. R., & Diewert, W. E. (1982). Multilateral comparisons of output, input, and productivity using superlative index numbers. *The Economic Journal*, *92*(365), 73–86. Retrieved November 16, 2025, from http://www.jstor.org/stable/2232257

Corrado, C. A., Hulten, C. R., & Sichel, D. E. (2009). Intangible capital and u.s. economic growth. *Review of Income and Wealth*, *55*(3), 661–685. https://doi.org/10.1111/j.1475-4991.2009.00343.x

Dolman, B., Parham, D., & Zheng, S. (2007). *Can australia match US productivity performance?* (Staff Working Paper). Australian Treasury.

Griffith, R., Redding, S., & Reenen, J. V. (2004). Mapping the two faces of r&d: Productivity growth in a panel of oecd industries. *The Review of Economics and Statistics*, *86*(4), 883–895. https://doi.org/None

Harrigan, J. (1997). Technology, factor supplies, and international specialization: Estimating the neoclassical model. *The American Economic Review*, *87*(4), 475–494. Retrieved November 16, 2025, from http://www.jstor.org/stable/2951359

Haskel, J., & Westlake, S. (2018). *Capitalism without capital: The rise of the intangible economy*. Princeton University Press.

Ignaszak, M., Robbins, D., & Sedláček, P. (2025). *The bang for the buck: Aggregate impact of firm-level R&D incentives* (Working Paper No. No. 062). The Productivity Institute.

IMF. (2021). *Australia: 2021 Article IV consultation* (Country Report No. No. 2021/256). International Monetary Fund.

Kirchner, S. (2020). *Failure to converge: The Australia–US productivity gap in long-run perspective* (tech. rep.). United States Studies Centre, University of Sydney.

OECD. (2021). *OECD economic surveys: Australia 2021* (tech. rep.). OECD Publishing. Paris.

Romer, P. M. (1994). The origins of endogenous growth. *Journal of Economic Perspectives*, *8*(1), 3–22. https://doi.org/10.1257/jep.8.1.3

Stanford, J. (2025). *Investing in Australia's future: The decline of non-mining business investment* (tech. rep.). Australia Institute, Centre for Future Work.

Ugur, M., Trushin, E., Solomon, E., & Guidi, F. (2016). R&d and productivity in oecd firms and industries: A hierarchical meta-regression analysis. *Research Policy*, *45*(10), 2069–2086. https://doi.org/10.1016/j.respol.2018.07.002

APPENDIX PRODUCTIVITY AND CAPITAL INTENSITY BY INDUSTRY

Over the past 30 years, Australian firms have seen faster rates of total factor productivity growth than US firms in industries such as finance, wholesale trade and construction (Figure A.1). But productivity among US firms has grown more rapidly in industries such as manufacturing, information and mining.

While the focus of the industry-level regression analysis in the note is on TFP growth, similar industry patterns hold for labour productivity (Figure A.2). The most notable exception is the mining industry — labour productivity growth has been similar across the two countries. A similar, if less extreme, divergence shows up in transport, postal and warehousing: Australian labour productivity grew roughly 36 per cent between 1997 and 2023 compared with around 10 per cent in the United States, but TFP growth in Australian transport was only about 5 per cent against 12 per cent in the US. The wedge between LP and TFP growth in Australian transport reflects substantial capital deepening — a large increase in capital per worker tied to the mining investment boom and associated logistics infrastructure — rather than a data issue in the underlying series.

The divergent outcomes between the US and Australia are mirrored by their different investment patterns. Australian businesses invest as much or more capital as US firms overall (Figure A.3). However, where they put that investment has differed. Australian firms historically devoted a larger share of investment to tangible assets like machinery, equipment, and non-residential construction, as well as to information technology hardware. In contrast, US firms invest more in intangibles across most industries — particularly in R&D (Figure A.4) and software (Figure A.5). Even within sectors like manufacturing and information services, Australian businesses have tended to spend relatively more on non-IP assets (physical plant, equipment, IT hardware), while US businesses spent more on R&D and other IP.

The gap in IP investment between Australia and the US across industries is mostly explained by differences in the intensity of R&D spending (Figure A.4). The gap in software investment intensity is typically much smaller (Figure A.5).

Figure A.1: Total Factor Productivity by Industry

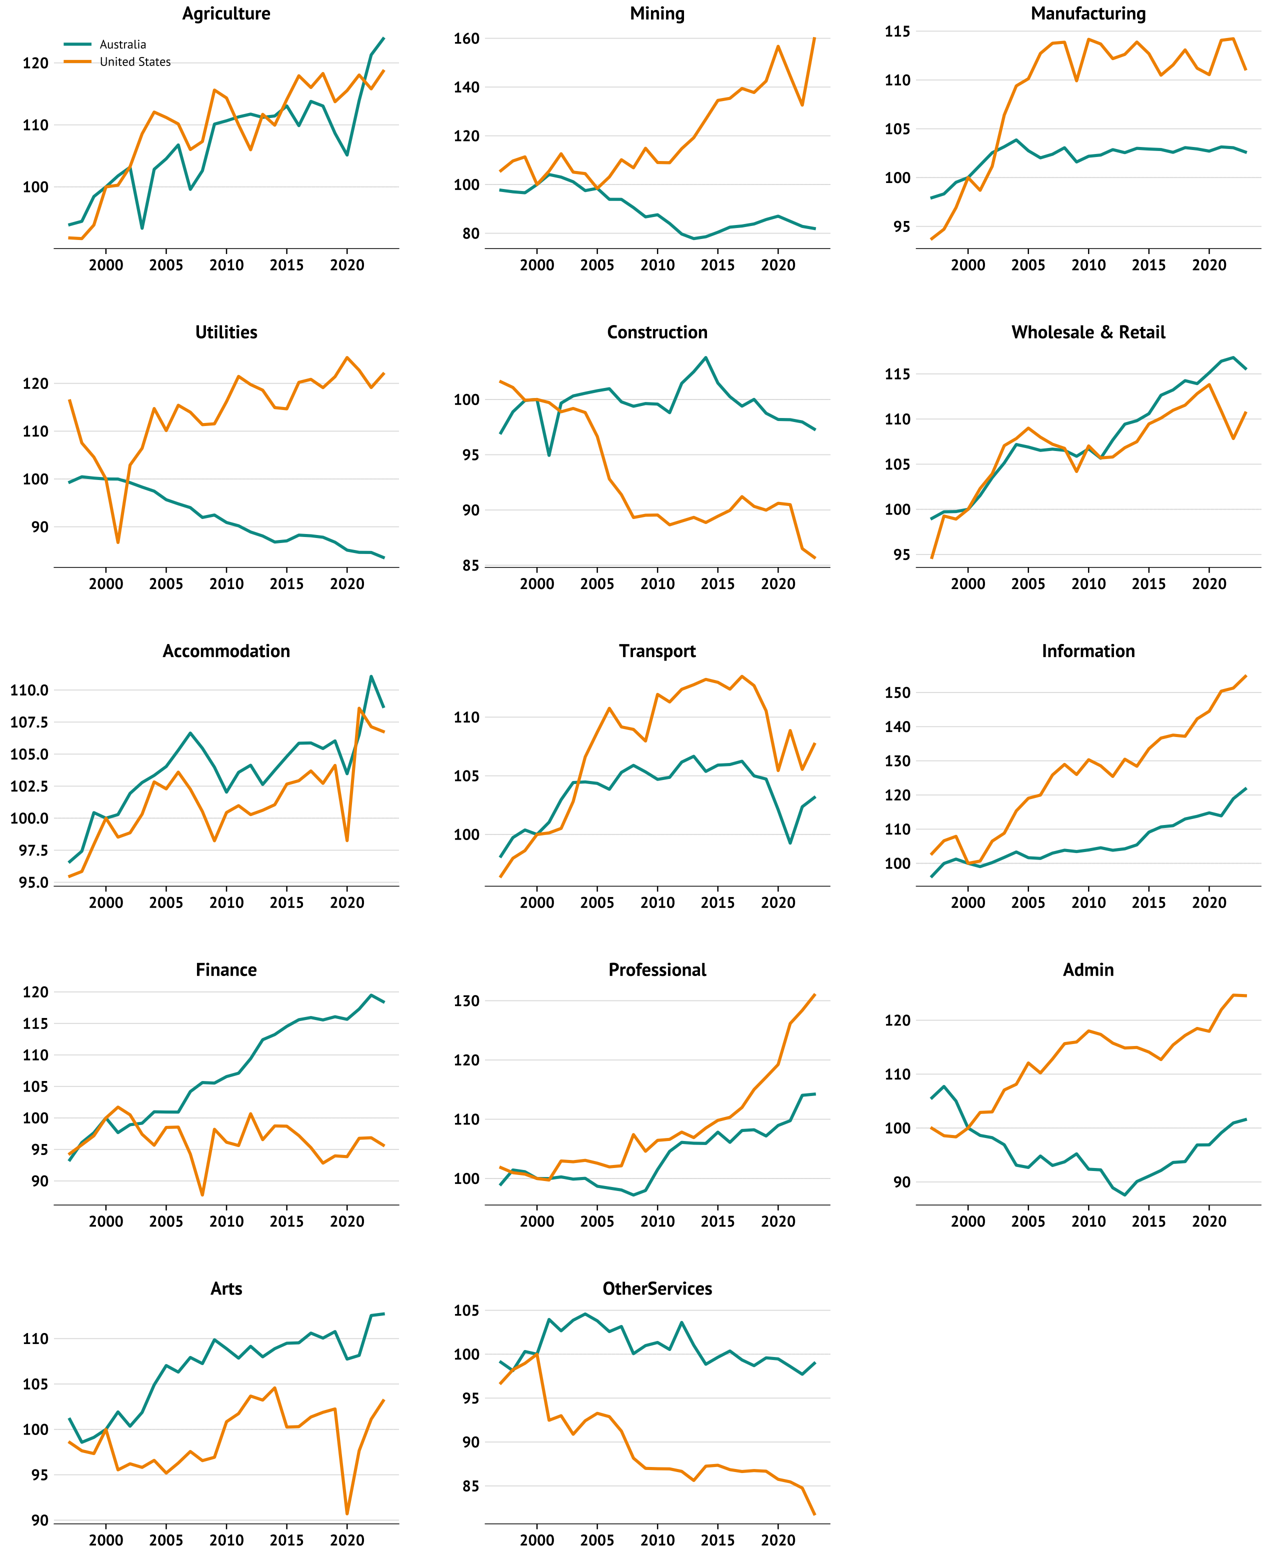

Sources: ABS; BEA; BLS; e61 Institute.

Figure A.2: Labour Productivity by Industry

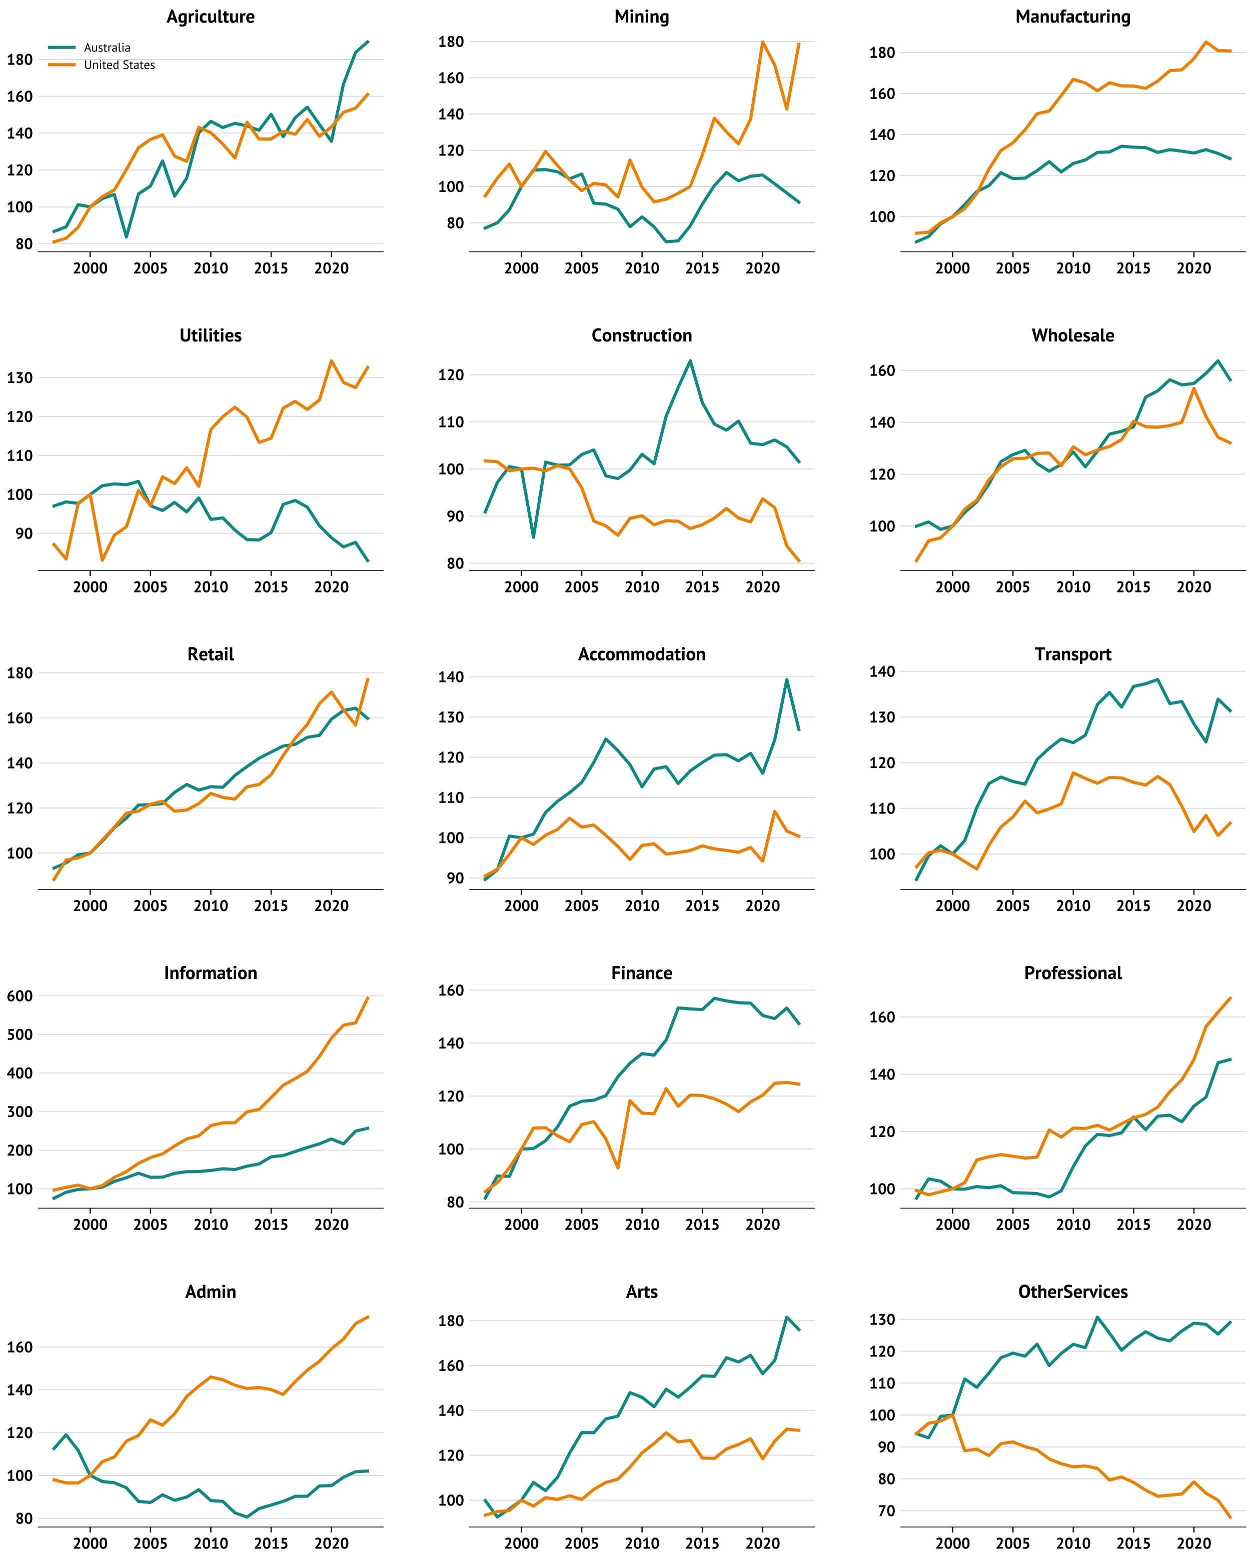

Notes: Real GVA per hour worked, 2000 = 100. *Sources:* ABS; BEA; BLS; e61 Institute.

Figure A.3: Non-residential Investment (% of industry GVA)

Sources: ABS; BEA; BLS; e61 Institute.

Figure A.4: R&D Investment (% of industry GVA)

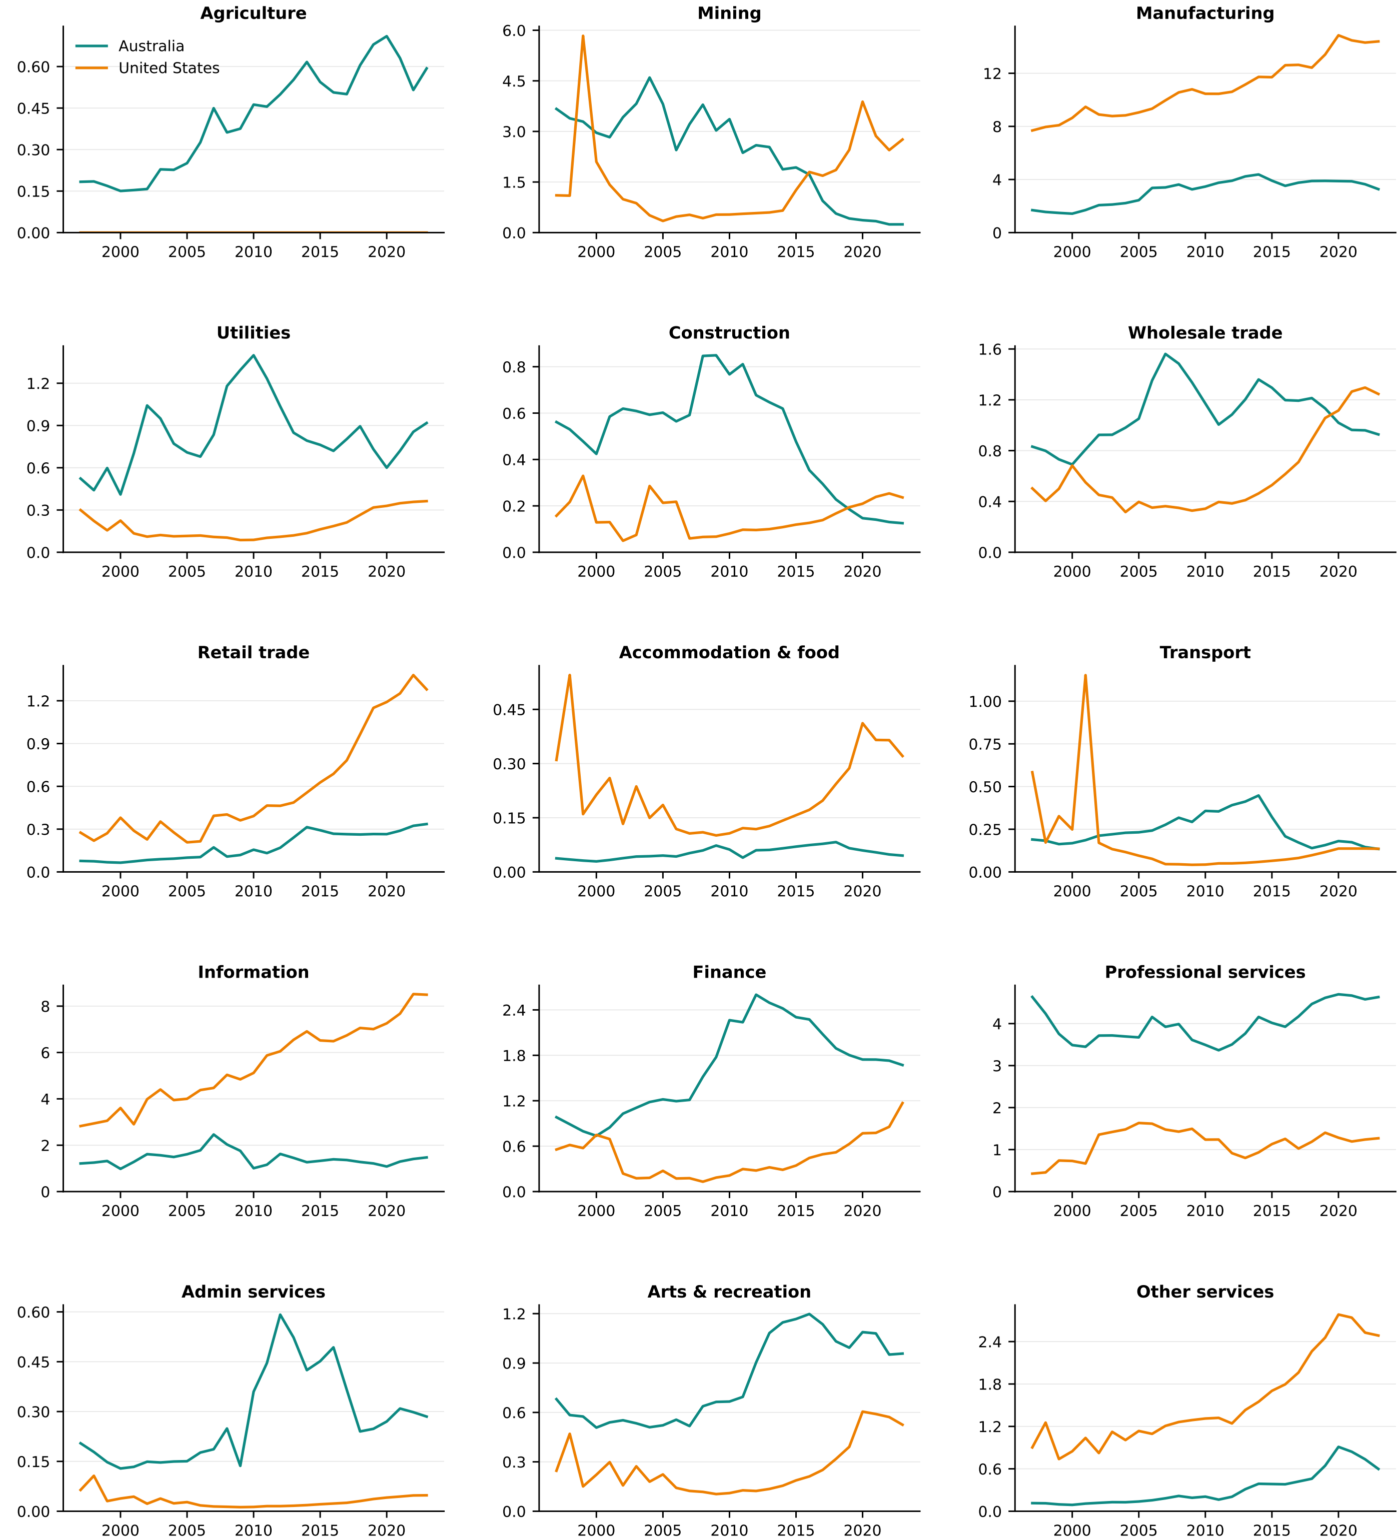

Sources: ABS; BEA; BLS; e61 Institute.

Figure A.5: Software Investment (% of industry GVA)

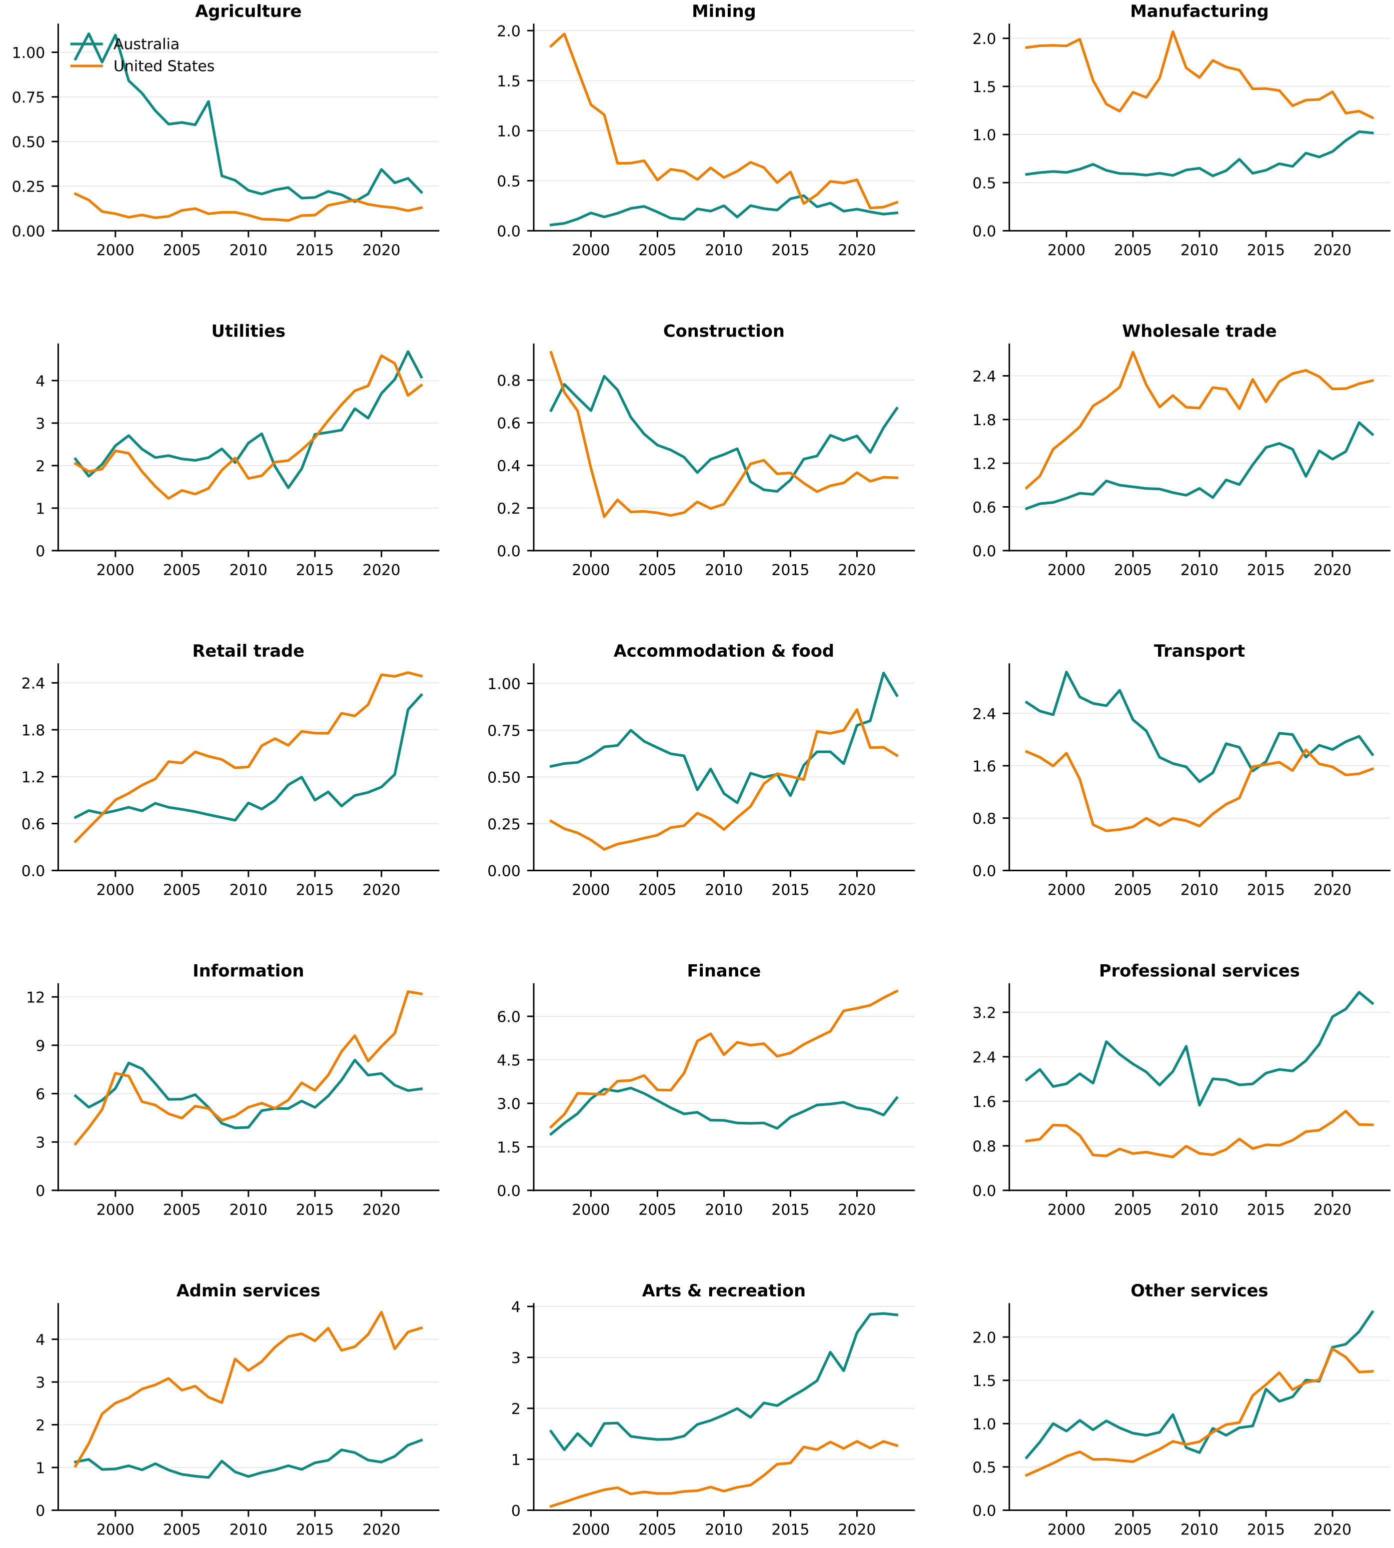

Sources: ABS; BEA; BLS; e61 Institute.

APPENDIX SUB-INDUSTRY ANALYSIS OF THE MANUFACTURING AND INFORMATION SECTORS

A more disaggregated breakdown of the manufacturing and information sectors can be undertaken because the Australian and American sub-industries are broadly similar at the 2-digit level. However, due to data limitations, the analysis here is based on labour productivity using real gross value-added and headcount.

Figure B.1: Manufacturing Sector Productivity by Sub-industry

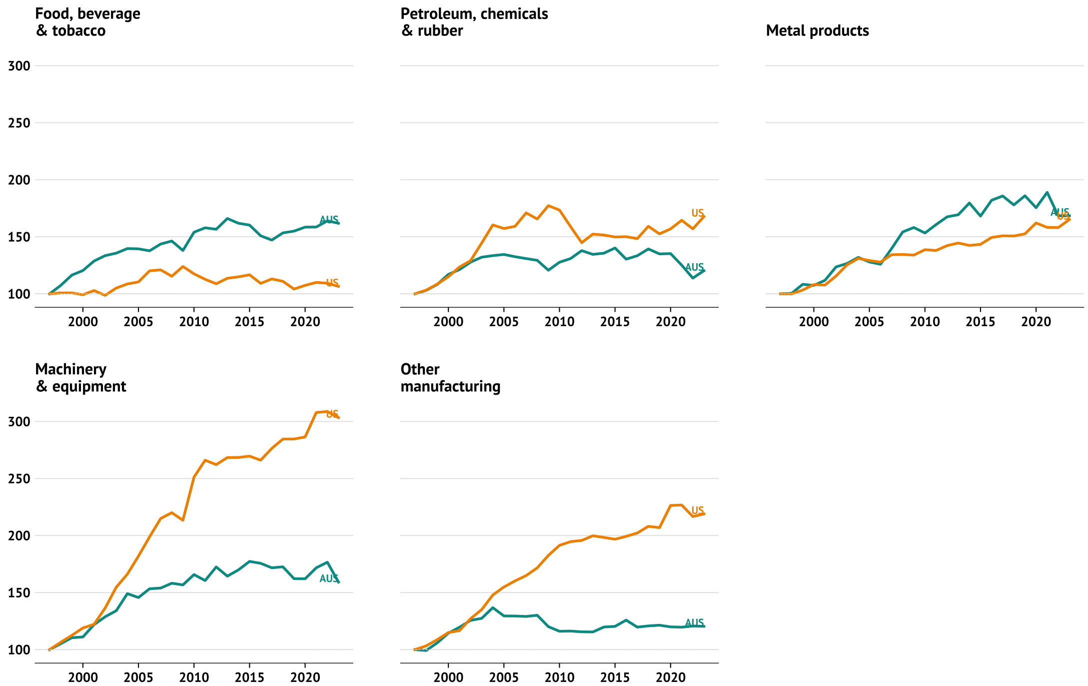

Sources: ABS; BEA; BLS; e61 Institute.

The analysis of the manufacturing sub-industries shows that US firms have not pulled ahead of Australian firms in every part of the sector (Figure B.1). Since the early 2000s, American manufacturers have recorded stronger labour productivity growth in areas like machinery and equipment, petrol, coal, chemical and rubber products, and a broad “other” category that covers clothing, footwear and furniture. Australian manufacturers, by contrast, have grown faster in food, beverages and tobacco, along with several mining-related manufacturing branches, including metals. The overall US manufacturing sector also looks stronger because computer-related manufacturing makes up a much larger share of US manufacturing than it does in Australia. This sub-industry is led by large US companies such as Apple, Dell, HP and Intel.

The breakdown of sub-industries within the information sector shows that US telecommunications firms have seen faster productivity growth than Australian firms since the mid-2000s (Figure B.2). The productivity gap in the ‘other information and media services’ sub-industry is also larger and widening over time, which likely reflects the growing weight of software, digital media, cloud and platform-based activities in the US relative to Australia.

Figure B.2: Information Sector Productivity by Sub-industry

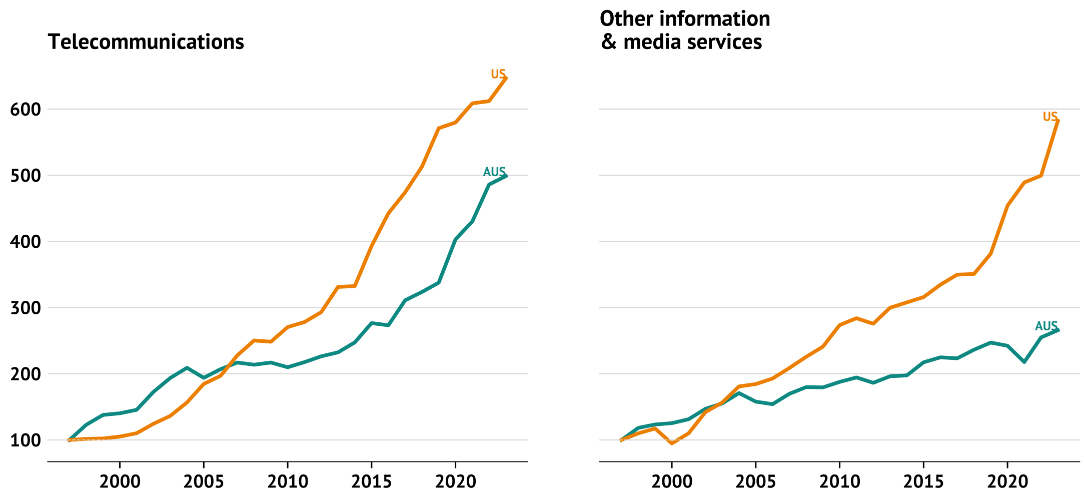

Notes: US 'Telecommunications' includes broadcasting.
Sources: ABS; BEA; BLS; e61 Institute.

This appendix reports additional robustness checks for the main regression in Table 1. Three modifications are considered: (i) replacing decade dummies with year fixed effects, which absorb a finer set of common shocks; (ii) excluding finance (ANZSIC division K), where output measurement based on financial intermediation services indirectly measured (FISIM) is notoriously imperfect; and (iii) excluding mining (ANZSIC division B), where mineral exploration is excluded from IP capital by construction and where the productivity-IP relationship may therefore be hard to identify cleanly.

The headline result — a positive and economically meaningful coefficient on the IP capital gap — is robust to year fixed effects and to excluding finance. The result also survives in the equally-weighted OLS specification when mining is excluded. However, the GVA-weighted specification is sensitive to the inclusion of mining: dropping mining reduces the WLS coefficient from around 0.17 to about 0.04. This reflects the fact that Australia's mining sector is large by gross value added but records very low IP capital intensity (mineral exploration is excluded), so it contributes substantially to the estimated relationship under GVA weighting. The OLS specification, which weights all industries equally, is less affected.

Trade exposure and human capital controls

The baseline regression sets $\gamma = 0$ — it omits industry-level controls beyond the IP capital gap and decade fixed effects — because comparable cross-country measures are not readily available. This section relaxes that restriction using two proxies.

Trade exposure is constructed from the OECD Trade in Value Added (TiVA) database, 2025 edition, at the ISIC Rev. 4 1-digit level. For each industry–country–year, trade exposure is defined as gross exports plus gross imports divided by gross output. The regressor $\hat{X}_{it}^{trade}$ is then the log ratio of US to Australian trade exposure. TiVA coverage is 2000–2022, which reduces the sample from 405 to 345 industry-year observations. The ratio captures both export orientation and import competition in a single index and is consistent across countries because TiVA is built from harmonised inter-country input-output tables.

Human capital is proxied by the ratio of quality-adjusted hours to total hours in each industry-country-year. For Australia, both series are published by the ABS in Table 9 of the *Estimates of Industry Multifactor Productivity* release (Cat. 5260.0.55.002) at the ANZSIC 1-digit level. For the United States, the equivalent ratio is constructed from the BLS Integrated Industry Productivity Account by dividing the *Labor Input* (composition-adjusted) index by the *Labor Hours* index at the detailed-NAICS level, then aggregating to ANZSIC-compatible 1-digit divisions using labour-compensation weights. Both countries' series are re-indexed

Table C.1: Robustness: TFP Gap and IP Capital Gap

	OLS		WLS (GVA-weighted)			
Sample	$\hat{\eta}$	SE	$\hat{\eta}$	SE	N	Industries
Panel A: Decade time effects (baseline specification)						
All market	0.192***	(0.048)	0.172***	(0.042)	405	15
Excl. C, J	0.187***	(0.055)	0.184***	(0.044)	351	13
Excl. K (Finance)	0.192***	(0.046)	0.164***	(0.036)	378	14
Excl. B (Mining)	0.137*	(0.067)	0.036	(0.052)	378	14
Excl. K, B	0.142*	(0.068)	0.038	(0.059)	351	13
Excl. C, J, K, B	0.120	(0.070)	0.036	(0.047)	297	11
Panel B: Annual time effects (year fixed effects)						
All market	0.210***	(0.053)	0.186***	(0.045)	405	15
Excl. C, J	0.201***	(0.064)	0.194***	(0.052)	351	13
Excl. K, B	0.159*	(0.077)	0.046	(0.070)	351	13

Notes: Each row reports the coefficient on the IP capital gap ($\hat{\eta}$) from a within-industry regression of the Törnqvist TFP gap on the IP capital gap and time effects. Panel A uses decade dummies (2000s, 2010s, 2020s); Panel B replaces these with year fixed effects. Standard errors are clustered by industry. The samples progressively exclude manufacturing (C), information (J), finance (K) and mining (B). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Sources: ABS; BEA; BLS; e61 Institute.

Table C.2: Robustness: Trade Exposure and Human Capital Controls

	OLS		WLS (GVA-weighted)			
Specification	$\hat{\eta}$	SE	$\hat{\eta}$	SE	N	Industries
Baseline (no controls)	0.192^{***}	(0.048)	0.172^{***}	(0.042)	405	15
+ Trade exposure	0.174^{***}	(0.038)	0.131^{***}	(0.023)	345	15
+ Human capital	0.198^{***}	(0.046)	0.168^{***}	(0.037)	405	15
+ Trade & human capital	0.177^{***}	(0.035)	0.123^{***}	(0.024)	345	15

Notes: Each row reports the coefficient on the IP capital gap ($\hat{\eta}$) from a within-industry regression of the Törnqvist TFP gap on the IP capital gap, decade dummies and the indicated controls. Trade exposure is the log US/AUS ratio of (gross exports + gross imports) / gross output at ISIC 1-digit level from OECD TiVA 2025 ed., covering 2000–2022. Human capital is the log US/AUS ratio of labour-quality indexes (BLS Integrated Industry Productivity Account labour-input over labour-hours indexes; ABS 5260.0.55.002 quality-adjusted hours over hours), re-indexed to 2017 = 1 in each country. Standard errors are clustered by industry. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Sources: ABS; BEA; BLS; OECD; e61 Institute.

to 2017 = 1, so the cross-country gap $\hat{X}_{it}^{hc}$ captures relative *growth* in labour quality since 2017 rather than the level of that quality.[8]

Adding these controls leaves the IP capital coefficient largely unchanged (Table C.2). The OLS estimate moves from 0.19 in the baseline to 0.18 with the trade control, 0.20 with human capital, and 0.18 with both; the WLS estimate moves from 0.17 to 0.12–0.13 once trade is included, a roughly 25 per cent reduction but still statistically and economically significant. The coefficients on trade exposure and human capital themselves are statistically insignificant in every specification, and the point estimate on the human-capital gap is imprecisely determined because the underlying series show only small differential movements between the two countries over the sample period (see the discussion of levels versus trends in the footnote).

Taken together, controlling for trade exposure and human-capital proxies does not overturn the main finding that IP capital differences are a robust correlate of the productivity gap.

8 Interpreting the human-capital control as a differential-trend proxy is a limitation: ABS and BLS both publish labour-quality indexes rather than levels, and differences in methodology (ABS decomposes by age × sex × education × hours; BLS by age × sex × education) make cross-country level comparisons unreliable. The coefficient should therefore be read as the impact of relative *changes* in workforce composition, not the level of labour quality.

APPENDIX LABOUR PRODUCTIVITY ADJUSTED FOR PURCHASING POWER PARITY

An adjustment for purchasing power parity between Australia and the United States makes little difference to the aggregate productivity gap when estimated over a 30-year period (Figure D.1). The PPP adjustment leads to larger cyclical fluctuations in the productivity gap between the two countries, which mostly reflects real exchange rate adjustments during the Australian mining boom and bust episode in the early part of the 21st century.

Figure D.1: PPP-adjusted Market Sector Labour Productivity

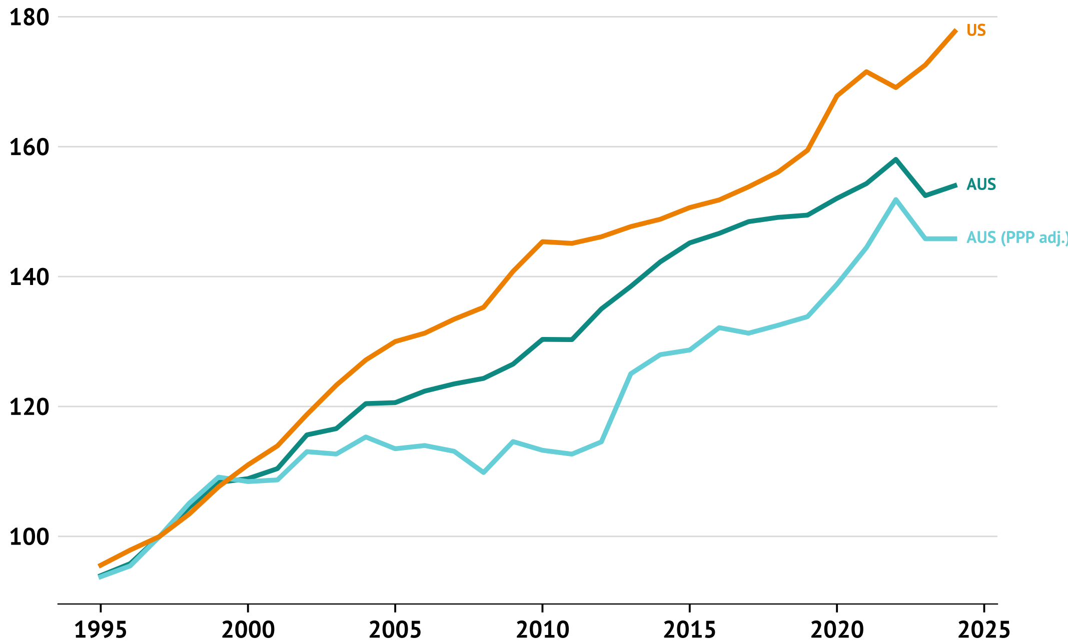

Sources: ABS; BEA; BLS; OECD; e61 Institute.