

Allan, Corey; Maré, David C.; Hyslop, Dean R.

Working Paper

Monopsony in the New Zealand Labour Market: First Estimates from Administrative Data

IZA Discussion Papers, No. 17614

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Allan, Corey; Maré, David C.; Hyslop, Dean R. (2025) : Monopsony in the New Zealand Labour Market: First Estimates from Administrative Data, IZA Discussion Papers, No. 17614, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/314511>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 17614

**Monopsony in the New Zealand Labour
Market:
First Estimates from Administrative Data**

Corey Allan
David C. Maré
Dean R. Hyslop

JANUARY 2025

DISCUSSION PAPER SERIES

IZA DP No. 17614

Monopsony in the New Zealand Labour Market: First Estimates from Administrative Data

Corey Allan

NZ Ministry of Business, Innovation and Employment

David C. Maré

Motu Economic and Public Policy Research and IZA

Dean R. Hyslop

Motu Economic and Public Policy Research

JANUARY 2025

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Monopsony in the New Zealand Labour Market: First Estimates from Administrative Data

We examine employer monopsony power in the New Zealand private sector labour market. New Zealand has a small, geographically dispersed population, meaning that outside employment options for workers may be limited. However, New Zealand is generally considered to have a flexible labour market with large gross labour market flows. Using firm and individual level microdata from StatsNZ's Longitudinal Business Database (LBD) and Integrated Data Infrastructure (IDI), we estimate monopsony power based on separation elasticities, on the estimated marginal product-wage wedge, and by direct estimation of firm-level labour supply elasticities. Estimates based on separation elasticities and the marginal product-wage wedge are reasonably consistent, with an implied wage markdown of at most 25%, on average. Direct estimates of labour supply elasticities are sensitive to small changes in specification, highlighting the identification difficulties. Our estimates based on separation elasticities and marginal product-wage wedges are broadly consistent with recent international evidence. These results suggest the presence of employer monopsony power in New Zealand's private sector, although the extent of that power may be limited.

JEL Classification: J42, J63, M50, D20

Keywords: monopsony, wage setting, worker mobility

Corresponding author:

David C. Maré
Motu Economic and Public Policy Research
Level 1, 97 Cuba Street
PO Box 24390
Wellington 6142
New Zealand
E-mail: dave.mare@motu.org.nz

1. Introduction

Who sets wages? This question is central to understanding how labour markets operate and how a range of policies will impact on the labour market. As Card (2022) discusses, there is increasing recognition that most firms possess at least some limited degree of wage setting power, also termed monopsony power. A large and growing international literature attempts to measure the degree to which firms do possess wage setting power – the ability of firms to set wages below the value of what workers produce. Understanding the extent to which firms exercise such power, and over which workers, can also further our understanding of existing labour market disparities in New Zealand (e.g., gender and ethnic wage gaps).

Some features of New Zealand’s labour market make it a promising location in which to study monopsony power. New Zealand has a small, geographically dispersed population, suggesting many workers may have limited employment options outside of their current firm. In addition, the labour market is not covered by NZ’s main anti-trust legislation meaning that some anti-competitive practices that would be deemed illegal in product markets are not so in the labour market. Restraints of trade (non-compete) clauses in employment contracts, which can limit a person’s ability to find work at similar firms within a specified geographic area and time, may be enforceable in certain circumstances.¹

On the other hand, New Zealand has a dynamic labour market with high levels of participation and employment, and large gross labour market flows. In 2022, 27% of workers had been in their jobs for less than 1 year, 19% in the UK, 21% in Canada, and 21% across the OECD.² In general, New Zealand’s labour market settings are viewed as conducive to labour reallocation and productivity growth, suggesting healthy competition in the labour market (Barnes et al. 2013).

In this paper we examine the extent of monopsony power in the New Zealand labour market. To our knowledge, ours is the first attempt to provide direct evidence on the extent of employer monopsony in New Zealand. We compare the results from three approaches to estimating monopsony power. First, we use the separation elasticity approach based on models of dynamic monopsony (Manning 2003); second, we estimate mark-ups based on production function approach in the industrial organisation (IO) literature (Yeh et al. 2022; De Loecker and Warzynski 2012); and third, we directly estimate labour supply elasticities to the firm.

We implement the three approaches using data for a consistent set of firms sourced from StatsNZ’s Integrated Data Infrastructure (IDI) and Longitudinal Business Database (LBD) over the period 2002–2019. We utilise data on firm employment, wages, and worker movements to estimate separation elasticities and labour supply functions, while using the LBD’s production data to estimate production functions and implement the production approach.

We find evidence of monopsony power in New Zealand, with implied wage markdowns ranging from 6% to 25%. There is some overlap in estimates between the dynamic monopsony approach and the production approach, with a set of industries providing broadly consistent estimates across these two approaches. However, the correlation between the two sets of estimates is generally low

¹ Restraint of trade clauses are typically considered prima-facie unenforceable. However, if an employer attempted to enforce such a clause and this attempt found was brought before the Employment Relations Authority (ERA) or Employment Court, it is incumbent on the employer to prove that the restriction is reasonable to protect their business from damages and the restrictions imposed on the former employee are not overly restrictive. Whether such clauses are enforceable in the courts or not, the presence of such clauses in employment contracts is common, and research has shown the presence of clauses is enough to lower worker mobility (Starr et al. 2020).

² Source: OECD.stat, Table K1 Employment by job tenure intervals – frequency: Job tenure less than 1 year, for dependent employees. Accessed 5 April 2024.

(between 0.14 and 0.41). The imprecision of the estimates makes it difficult to draw strong conclusions on the consistency of the approaches at the industry level. We also find the direct estimates of the labour supply elasticity are sensitive to small changes in the estimation approach, highlighting the difficulty in identifying these parameters at the firm level.

The presence of monopsony power implies the existence of frictions in New Zealand's labour market that limit the ability of workers to move between firms. These frictions could relate to asymmetric or incomplete information (e.g., whether a restraint of trade clause can be enforced in the specific circumstances, what other similar firms are paying their workers, job vacancies not listed on common search platforms), search costs (e.g., the time and effort involved in searching and applying for jobs), or institutional factors (e.g., occupational licencing, employer-specific work visas, restraint of trade clauses or non-compete covenants). These frictions affect workers by potentially lowering wages and generally result in lower levels of employment than would be achieved in the absence of these frictions. These frictions also impact a firm's ability to recruit and expand and may slow the pace of labour reallocation across the economy. An upward sloping labour supply curve, meaning the firm must raise its wage offer to attract a new worker while retaining their current workers, means firms are constrained in their ability to expand. This could therefore reduce the rate of productivity growth as the most productive firms find it more difficult to find the workers necessary to expand, while less productive firms are able to retain their workers and market share.

This work is complementary to earlier research on bargaining power and rent sharing in New Zealand (Allan and Maré, 2021; 2022). While the earlier work focussed on the sharing of economic rents earned by firms in the form of high wages, the current research gives further insights into the importance of labour market frictions in generating economic rents.

It also sits alongside a growing body of literature examining labour market imperfections, bargaining power, and the distribution of economic success in New Zealand. For example, Bertram and Rosenberg (2022) consider the role of policies and institutions in the evolution of the share of total income earned by labour and suggest that economic rents have become a much larger share of total economic surplus since the 1980s economic reforms. Sin et al. (2022) examine drivers of the gender wage gap in New Zealand and find that approximately one-third of the within-firm wage gap is explained by differences in bargaining over firm-specific rents. Fabling and Grimes (2019) find that the productivity improvements associated with ultra-fast broadband (UFB) adoption are shared unequally with different incumbent workers. Males with a Science, Technology, Engineering and Math (STEM) qualification or other high-level qualification tend to benefit most, while women see little wage growth following UFB adoption. This body of work, including our previous papers, shows a clear role for wage setting power in New Zealand.

Issues of market power in the labour market are being actively looked at by policy agencies in several countries. Much of this work has focussed on legal barriers to worker mobility, especially restraint of trade or non-compete clauses. The Federal Trade Commission (FTC) in the US proposed a rule to ban noncompete clauses in employment contracts, which it claimed would increase workers' earnings by nearly US\$3 billion (Federal Trade Commission, 2023). In 2023 the Australian Government announced a review of competition policy, including looking at the impact of non-compete clauses on labour markets. Research conducted as part of the review finds at least 1 in 5 workers in Australia was bound by some form of non-compete clause (Andrews and Jarvis, 2023). In the UK, the Competition and Markets Authority (CMA) has released a report into the state of competition in their labour market finding that, while the degree of labour market power hasn't increased over time, there are significant and persistent differences in the extent of monopsony power across labour markets (Competition and Markets Authority 2024).

The rest of this paper is structured as follows. Section 2 provides an overview of monopsony models (2.1), the current empirical literature (2.2), and New Zealand's institutional setting (2.3). Section 3 provides an overview of the empirical models and estimation strategies we employ. Section 4

discusses the data used in our study (4.1 and 4.2), and some summary statistics related to monopsony power in New Zealand (4.3 and 4.4); while section 5 presents our main results. Section 6 concludes.

2. Background and literature review

In this section, we provide a brief overview of models of monopsony in the labour market and discuss the empirical evidence of the extent of monopsony power in labour markets. We then give a short description of the institutional settings relevant to New Zealand's labour market and questions of market power.

2.1. A brief overview of monopsony models

The typical textbook way of thinking about labour markets and wage setting is based on perfect competition. In this world, both workers and firms are homogeneous, and wages are set by the interaction of market supply and demand. Firms face an infinitely elastic labour supply curve and can hire as many workers as they choose at the prevailing market wage. This leads to the well-known condition that, in equilibrium, the wage will equal the marginal revenue product of labour.

As discussed in Manning (2011), one implication of a perfectly competitive model is that all firms will pay the same wage for a given type of worker within a specific labour market. If a firm were to pay even 1c less, all current workers would leave, and the firm would be unable to hire any new workers. If a firm were instead to offer even 1c more than the prevailing market wage, they would be inundated with job applicants.

We see a large degree of heterogeneity in wages within labour markets. A large fraction of this heterogeneity is due to differences among workers themselves, but a significant proportion of this variation is attributed to firm-level differences (e.g., Abowd et al. 1999; Card et al. 2018; Maré and Hyslop, 2006; Criscuolo et al. 2020; OECD 2021). Research also suggests that at least part of the rise in wage inequality over the last 30-40 years is driven by increasing dispersion in the firm-specific component of wages (e.g., Song et al. 2019; Card et al. 2013). Common experience further tells us that not all workers will leave in response to a wage cut, and while high-paying firms may find recruitment easier they are not swamped with job applications when they post a vacancy.

Monopsony models are one way to address the gap between textbook theory and observed experience. The wage setting process in monopsony models has a firm-specific component. The defining feature in monopsony models is driven by firms facing upward sloping labour supply curves with a finite labour supply elasticity.

There are several departures from the perfectly competitive model that generate an upward-sloping supply curve for individual firms. These departures relate to factors that restrict a worker's ability or willingness to switch firms. The simplest is relaxing the assumption of many homogeneous firms and instead consider the case of a single purchaser of labour in the labour market. This firm faces a standard profit maximisation problem but because this firm's labour supply curve is the market labour supply curve, to hire new workers to expand production it must offer higher wages to entice workers into their labour force. That is, wages are an increasing function for firm employment (L). The firm decision is to choose L to maximise profits.

To see this, consider a departure from the perfectly competitive model, in which a firm that is a wage setter, in the sense that the labour supply it faces is less than infinitely elastic, and it can affect the supply of labour by changing its wage offer. Then its labour use will depend on the wage it pays, and the firm's profit function can be expressed as:

$$\pi = f(L) - w(L)L \quad (1)$$

where π is firm profit, $f(L)$ is the revenue function, L is the quantity of labour, and $w(L)$ is the wage rate the firm must pay to employ L workers. The first order conditions for profit maximisation yield the following expression for wages:

$$w = MRP_L \frac{\eta}{1 + \eta} \quad (2)$$

where $\eta = \frac{dL}{dw} \cdot \frac{w}{L}$ is the elasticity of labour supply to the firm with respect to the wage, and $MRP_L = \frac{df}{dL}$ is the marginal revenue product of labour. In a perfectly competitive labour market, $\eta = \infty$, and $w = MRP_L$. With a positive, but finite, η , the term $\frac{\eta}{1 + \eta} < 1$, meaning that $w < MRP_L$.

In equation (2), the term $\frac{\eta}{1 + \eta}$ measures the ratio of the wage to the marginal product of labour and depends on the elasticity of the labour supply curve. For positive values of η , this ratio will be bounded above by 1 and will be less than 1 for finite positive values of η . For instance, if $\eta = 9$, the wage will equal 90% of the marginal product. If $\eta = 2$, the wage will equal 67% of the marginal product.

This shows that the elasticity of the firm-specific labour supply curve is critical for understanding the extent of monopsony power in the labour market. The elasticity of labour supply can be transformed into several different metrics that quantify the relationship between wages and marginal products. We focus on the wage markdown, measured as the relative gap of the wage to marginal product. Rearranging equation (2), this is expressed as:

$$\text{Markdown} = 1 - \frac{w}{MRP_L} = \frac{1}{1 + \eta}. \quad (3)$$

This wage markdown measures the fraction of marginal product that is not paid in wages. For example, a labour supply elasticity of 2 means that the wage markdown is 33%, whereas an elasticity of 9 implies a 10% wage markdown.

Other monopsony models capture alternative sources of market power that don't depend on market dominance. One such class of models allows workers to have heterogeneous preferences over the non-wage attributes of firms. When firms differ in their non-wage attributes and where workers have heterogeneous preferences over these attributes, firms will gain a degree of wage setting power even when they are small in relation to the overall market. Card et al. (2018) and Kline et al. (2019) are examples of these types of models. These heterogeneous preferences mean that a worker may be willing to accept a lower wage for a job at a firm for which they have a particularly strong preference. It also makes it harder for competing firms to expand as they have to offer higher wages to attract workers from firms for which they may have a stronger preference. Heterogeneous preferences effectively limit the number of viable outside employment options available to workers, meaning that firms face upward sloping labour supply curves and have a degree of wage setting power.

Market frictions are a further source of wage setting power. In situations where job search is costly, a departure from the perfectly competitive model, firms will gain a degree of monopsony power (dynamic monopsony). When job search is costly, both the worker and the firm gain from maintaining an ongoing employment relationship. Search frictions effectively limit the set of outside options facing a worker as it is prohibitively expensive for the worker to be aware of or move to the full set of job options available to them.

Manning (2003) outlines a simple version of a dynamic monopsony model, building on the search model of Burdett and Mortensen (1998). Here, firms face a dynamic labour supply curve that depends on separations and recruitment, which both depend on the wage:

$$L_t = (1 - s(w))L_{t-1} + R_t(w) \quad (4)$$

where $s(w)$ is the separation rate, which depends negatively on the wage; and $R_t(w)(= r(w) \cdot L_{t-1})$ describes the flow of recruits to the firm, which depends positively on the wage. In this model, the elasticity of labour supply can be written as a function of the wage elasticities of separations and of

recruitment. In steady state (i.e. where $L_t = L_{t-1}$), the inflow of new recruits offsets workers who leave the firm ($L = \frac{R(w)}{s(w)}$). The long-run elasticity of the labour supply curve (η) with respect to the wage is:

$$\eta = \eta^R - \eta^S \quad (5)$$

where η^R and η^S are the elasticities of recruitment and separations with respect to the wage, respectively. Manning (2003) shows that the recruitment elasticity can be written as a function of the separation elasticity in a model where a worker's decision to quit depends only on relative wage offers, so the separation elasticity (which is more readily observable) provides enough information to infer the degree of monopsony power. In the simplest version of the model, the relationship between the labour supply elasticity and the separation elasticity is given by:

$$\eta = -2\eta^S. \quad (6)$$

This dynamic monopsony model shows that, under the steady state assumption, it is not necessary to estimate the elasticity of the labour supply curve to understand monopsony power. Estimating the quit or recruitment elasticity is sufficient to identify the extent of monopsony power. Given that job separations are more readily observable to the researcher, much of the empirical literature on dynamic monopsony focusses on estimating η^S .

Langella and Manning (2021) combine the insights from preference-based static monopsony models (e.g., Card et al. 2018) with those from dynamic monopsony models (e.g., Manning 2003). In their model, idiosyncratic preferences effectively reduce the likelihood that a new job offer represents an improvement from the worker's perspective, even if the job offer comes with a greater wage than they earn at their current firm. The decision to change jobs then depends not only on relative wages, but also on relative non-wage job attributes and the relative importance that workers place on wage and non-wage attributes. In the version of the Burdett-Mortensen model used in Manning (2003), firms are assumed to be homogeneous in non-wage characteristics (or workers have no preferences over these attributes). Langella and Manning (2021) also emphasise the use of separation and/or recruitment elasticities to estimate the extent of monopsony power that firms possess.

Thus, a central element of monopsony models is a feature that restricts workers outside employment options. In simple models, this feature is the lack of competing firms, with a single firm being the sole purchaser of labour in the market. In more general models, factors such as heterogeneous preferences for workplace attributes or search/mobility costs serve to reduce the outside employment options. It also gives insight into how we might expect different labour market policies and institutions will affect monopsony power. In general, anything that increases a worker's outside options or reduces the costs of exercising those options, will reduce monopsony power.

2.2. Empirical evidence on the extent of monopsony power

There is a large and rapidly growing literature that seeks to estimate the extent of labour market monopsony power across a range of contexts. Several recent papers provide an overview the recent literature (e.g. Card 2022; Ashenfelter et al. 2022; Manning 2021; Sokolova and Sorensen 2021) while Manning (2003; 2011) surveys the earlier literature.

Sokolova and Sorensen (2021) provide a meta-analysis of the monopsony literature, looking at over 1,000 estimates of the labour supply elasticity from 53 studies. These estimates include direct estimates, from regressing employment (L) on wages (w), inverse estimates (from regressing w on L), and estimates derived from separation or recruitment elasticities. The median estimate across all studies examined is 1.7, while the median for those based on direct methods (L on w , separations, or recruitment) is around 1.4. In general, estimates of the inverse labour supply elasticity tend to be larger than those based on direct methods. This is in part driven by some studies that estimate inverse elasticities that are near zero, implying very large labour supply elasticities. Estimates based

on separation or recruitment elasticities tend to be higher than those from regressions of employment on wages. The median estimate from the literature surveyed implies a 37% markdown ($= \frac{1}{1+1.7}$) of wages below marginal product.

While the median estimate across all studies reviewed by Sokolova and Sorensen (2021) is 1.7, the authors find that studies published in top journals tend to find larger estimates of the labour supply elasticity and therefore less monopsony power. Their meta-analysis further reveals that studies with a more credible identification strategy find larger labour supply elasticities. They produce a set of ‘best practice’ estimates based on their meta-regression and these range from 6 to 10 (implying wage markdowns of between 9% and 14%), although these estimates have relatively wide confidence intervals.³ In the absence of a credible identification strategy, we would expect estimates of the labour supply elasticity (or separation elasticity) to be biased towards zero, overstating the degree of monopsony power that firms possess. The results of Sokolova and Sorensen (2021) suggest that many studies may be overstating the extent of monopsony power.

Card (2022) surveys a range of recent papers on the role of firms in wage setting, including papers estimating separation or recruitment elasticities as well as papers estimating the relationship between wages and firm productivity. Across the literature surveyed, the (implied) estimates of labour supply elasticities fall in the range of 3.5-5, higher than the median estimate but similar to the range of direct estimates from Sokolova and Sorensen (2021), implying a wage markdown of approximately 20%. Card et al. (2018) show that the range of rent-sharing elasticities common in that literature (0.05-0.15) are consistent with labour supply elasticities of approximately 4 in their simple monopsonistically competitive model of wage setting. A labour supply elasticity of four implies a wage markdown of 20%.

A range of studies find variation in the extent of labour market monopsony power. Sokolova and Sorensen (2021) document several papers that have explicitly examined the markets for nurses and teachers and generally find lower labour supply elasticities, indicating greater monopsony power in these specific labour markets (examples include Falch 2010; Matsudaira 2014; Ransom and Sims 2010; Staiger et al. 2010). This is not unexpected, given that employment of nurses or teachers tends to be concentrated in a small number of large employers (often government funded). Others have found that firms have greater monopsony power over low wage workers (e.g., Webber 2015; Langella and Manning 2021). Webber (2022) and Hirsch et al. (2018) find evidence that firms have greater monopsony power during economic downturns when unemployment is rising. Hirsch et al. (2010), Webber (2016), and Sánchez et al. (2022) show that firm-specific labour supply elasticities for women tend to be lower than for men, pointing to firms having greater monopsony power over their female employees, contributing to the gender wage gap. Hirsch et al. (2010) find that about one-third of the gender pay gap may be due to wage discrimination by monopsonistic employers, while Webber (2016) finds that 60% of the elasticity difference is explained by marriage and child penalties that affect women but not men.

Overall, the literature does point towards monopsony power being a pervasive feature of labour markets. Focussing on more recent studies and those with more credible identification strategies, a central estimate of the labour supply elasticity seems to be around 4 (implying a 20% wage markdown), although many studies have found lower elasticities, either overall or for specific segments of the labour market.

³ This excludes estimates based on the inverse supply elasticity and estimates of separation elasticities for males and females where the female share of employment is imputed. See Sokolova and Sorensen (2021), Table 5.

2.3. Institutional setting

We now discuss the institutional settings in New Zealand. There are features of New Zealand's labour market that make it a promising jurisdiction in which to study monopsony power in the labour market. First, New Zealand has a relatively small, geographically dispersed population. The outside employment options for some workers may be limited, perhaps more so than in other countries with similar populations, but greater geographic advantages (e.g., small countries in or near the European Union).

On the other hand, New Zealand is generally considered to have a flexible labour market. New Zealand's labour market settings are found to be conducive to GDP growth (Barnes et al. 2013) and there are relatively high rates of job starts, job ends, job creation, and job destruction. However, labour market flows have remained low post-GFC, returning to pre-GFC levels immediately prior to the COVID-19 pandemic (Maré 2022).⁴ New Zealand also tends to have one of the highest employment rates and labour force participation rates, and among the lowest unemployment rates in the OECD.

Union membership in New Zealand has declined dramatically over the past 30 years, as it has in several other OECD countries. In 2018, trade union membership was around 18% of employment, compared to the OECD average of 16%.⁵ Union membership is most prevalent in the government sectors of public administration and safety, education, and healthcare and social services. Private sector union membership is around 11%, compared with 35% in the public sectors, a pattern seen in other countries (Ryall and Blumenfeld, 2019).

This decline in union membership means that wage setting in New Zealand is highly decentralised and is typically done through agreements between individual firms and workers. Where collective bargaining does occur, the agreements tend to cover a single employer, rather than an entire sector or occupation.

Two pieces of legislation are important for understanding market power in the New Zealand labour market – the Employment Relations Act 2000 and the Commerce Act 1986. The Employment Relations Act provides the legal backdrop for relationships between employers and employees (or their representatives) as well as promoting the enforcement of employment standards. The Act also sets minimum employment standards where these are not covered by other legislation (e.g., trial periods, agreed hours of work, good faith and fair process, termination of employment agreements).⁶

Termination of employment relationships is typically straightforward from an employee's perspective, where the required notice period is stipulated in the employment agreement. Employees do enjoy a level of employment protection, meaning they are unlikely to have their employment terminated at short notice in the absence of any serious misconduct.⁷ At-will employment is rare in New Zealand. According to OECD data, in 2019 New Zealand had slightly lower

⁴ Maré (2022) gives a good overview employment and wage trends in New Zealand.

⁵ At the time of writing, 2018 is the most recent year for which comparable trade union membership statistics are available from the OECD. Source: OECD.Stat. Trade Union Dataset. Percent of employees who are union members. Accessed 5 April 2024.

⁶ Other relevant legislation that covers minimum employment standards include the Holidays Act 2003, the Equal Pay Act 1972, and the Minimum Wage Act 1983. The full list of legislation relevant to employment relationships can be found [here](#).

⁷ An exception to this is where an employee is on a 90-day trial, where an employee could have their employment terminated within 90 days of starting a job. 90-day trial clauses have been allowed in employment agreements for different groups of firms since 2009.

levels of employment protection than the OECD average. New Zealand's levels of employment protection are higher than in Australia and the UK, and lower than in Germany and France.⁸

The Commerce Act 1986 is New Zealand's anti-trust legislation, and its purpose is to 'promote competition in markets for the long-term benefit of consumers within New Zealand.' The Act established the Commerce Commission to enforce relevant competition legislation. It outlines a range of restrictive trade practices and how business acquisitions and mergers that may substantively affect the level of competition should be handled. Important for our purposes, none of the provisions in the Commerce Act are applicable to the labour market. That is, a range of practices that could inhibit competition, while illegal in product markets, may not be considered illegal under the Commerce Act (e.g., non-poaching agreements, restraint of trade clauses).

Several US based studies have examined the impacts of non-poaching agreements and restraint of trade clauses (termed non-compete contracts in the US). Kruger and Ashenfelter (2022) document the existence of explicit non-poaching agreements in franchise agreements, which prohibit one franchisee from hiring another's employee within a specified period. They find that such clauses were present in approximately half of franchise agreements. Balusubramanian et al. (2022) study the effects of restraints of trade/non-compete clauses on the mobility rates and wages of technology workers. They find that firm-to-firm mobility rates are reduced when and where such clauses are more enforceable. They also find evidence that wages are reduced for workers at different career stages, suggesting that restraint of trade clauses reduce outside employment options for workers and give firms greater (dynamic) monopsony power. Starr et al. (2021) find that 18% of American workers are bound by such clauses, while the UK Competition and Markets Authority (CMA) estimates that 26% of UK workers are bound by these clauses (Competition and Markets Authority 2024). Similar research for Australia finds around 1 in 5 workers are bound by such agreements (Andrews and Jarvis 2023). While these types of clauses are more prevalent among higher-wage workers, they are also increasingly appearing in employment contracts for low wage workers (see Andrews and Jarvis (2023) for Australia and Alves et al. (2024) for the UK).

While restraint of trade clauses are potentially enforceable in New Zealand. If an employer attempts to enforce such a clause, and this is challenged by the employee, it is up to the employer to prove that the enforcement of the clause is reasonable – that is, it protects the employer from undue damages while not overly burdening the employee or inhibiting product market competition.⁹ Following a recent high-profile Employment Relations Authority case involving the enforcement of a restraint of trade clause in a media personality's contract,¹⁰ legal commentators have outlined the strict conditions that a court would consider when deciding whether such a clause is enforceable (see Borrows 2022; Cowan 2021). A search of the New Zealand Legal Information Institute database for the terms 'restraint of trade' shows there have been 116 cases brought before the Employment Court involving a restraint of trade clause since 1992.¹¹ A similar search of Employment Relations Authority (ERA) decisions reveals 387 such cases since 2000.¹² We are currently unaware of any

⁸ Source: OECD.stat. Table U1: Strictness of employment protection – individual and collective dismissals (regular contracts), version 4. Accessed 5 April 2024.

⁹ Starr et al. (2020) find that the presence of non-compete clauses reduces worker mobility even when such clauses are unenforceable. Prescott & Starr (2022) show that workers tend to believe such clauses are enforceable even when they are not. Workers who believe these clauses are enforceable may be more likely to turn down better job opportunities due to the perceived threat of legal action by their employer.

¹⁰ O'Brien vs. Discovery NZ Limited, 2022 NZERA 15

¹¹ <http://www.nzlii.org/> - search terms "restraint of trade". Searched on 13/03/2023.

¹² <https://www.employment.govt.nz/elaw-search/> - search terms "restraint of trade". Searched on 13/03/2023. The Employment Court and ERA are two parts of the dispute resolution system relating to employment matters. Decisions by the ERA can be challenged in the Employment Court. Note that the number

analysis of these data or classification of the data into industry or occupation classifications that we could use to compare our results. This also covers only cases that have gone through some official legal process – we have no information on when or where restraints of trade clauses are more likely to be included in employment agreements, nor where even the threat of enforcement may be strong (whether enforcement of the clause may be legal).

of cases cited above makes no distinction about whether the employer or employee brought the case to the relevant authority, nor whether the Authority or Court found in favour of the employer or employee.

3. Estimation strategy

In this section, we outline our strategy for assessing the extent of monopsony power in the New Zealand labour market, which uses three approaches. First, we estimate the wage elasticity of separations to infer the labour supply elasticity. Second, we directly estimate elasticities using a labour supply function approach. Finally, we directly estimate the ratio of marginal revenue product of labour to the wage. We then compare estimates from these three approaches.

3.1. Estimating separation elasticities

Estimating supply (or demand) elasticities is an old econometric problem. When both demand and supply shocks are present, regressing price (in this case wages) on quantities (in this case workers) does not identify either supply or demand elasticities.

When estimating the labour supply elasticity, we are essentially asking how employment and wages respond to a demand shock. To trace out the supply curve, we need an excludable instrument that acts as a demand shifter. The same holds true for estimating a quit or recruitment elasticity as these are proximate determinants of the labour supply elasticity. We focus here on estimating the quit elasticity.

To illustrate the simultaneity bias when estimating the separation elasticity, consider the following log-linear relationship (Townsend 2023):

$$\log s = \eta^s \log w + \zeta \quad (7)$$

where s is the probability of separation, w is the wage, and ζ captures firm-level heterogeneity in non-wage factors that affect a worker's choice to leave a firm e.g., non-wage amenities.

From equation (2), we have that the profit-maximising wage rate is a mark down from marginal product. Following Manning (2003), assume that firms have revenue functions of the form:

$$y = \frac{A}{1-\theta} L^{1-\theta}. \quad (8)$$

Using the first-order condition for the optimal wage from equation (2):

$$w = AL^{-\theta} \left(\frac{\eta}{1+\eta} \right). \quad (9)$$

Replacing L with the steady state version of the dynamic labour supply function (equation (4), $L = \frac{R(w)}{s(w)}$), and taking logs yields:

$$\log w = \log A - \theta \log R + \theta \log s + \log \left(\frac{\eta}{1+\eta} \right). \quad (10)$$

In equilibrium, the log wage depends linearly on the log separation rate, and therefore on ζ . Therefore, OLS estimates of equation (10) will be biased if $\theta < 1$. When $\theta < 1$, ζ is positively correlated with the wage, leading to an upward bias in the OLS estimate. Since the expected relationship between the separation rate and wages is negative, this will lead us to underestimate both the separation elasticity and labour supply elasticity and therefore overstate the amount of monopsony power that firms possess (Townsend 2023).

An instrument is needed to isolate the exogenous variation in wages and trace out the labour supply curve. This instrument should be firm-specific and affect only the firm's labour demand, not labour supply. Essentially what we need is a measure of A , which is a productivity term. From the

productivity tables, we have measures of firm-level multi-factor productivity (MFP), and we use the lag of these as instruments for wages (Fabling and Maré 2015a; 2019).¹³

Our main estimating equation is:

$$\ln s_{jt} = \alpha_j + \eta^s \ln w_{jt} + X^s \beta^s + \lambda_t + \varepsilon_{jt} \quad (11)$$

where s_{jt} is the (monthly) separation rate for workers in firm j in year t (measured by the firm-level average monthly separation rate over the year), w_{jt} is the firm-level average wage, X^s is a vector of other controls, α_j is a firm fixed effect, λ_t are a set of year dummies, and ε_{jt} the error term. As we are interested in firm-level labour supply curves, it's important to control for firm heterogeneity. That is, we are interested in how separations from a firm respond when the firm changes its wage and the firm fixed effect means the estimates are based on within-firm variation. Omitting the firm fixed effect means we will be picking up the cross-sectional relationship between separations and wages which may reflect differences in employment practices across firms. This also means that we are, in effect, using the residual component of MFP as the instrument.

To assess the robustness of our results, we consider alternative measures of the probability of separation and the wage. First, we replace the average monthly separation rate, which counts all separations from the firm, with an average monthly job-to-job separation rate. That is, what is the probability a worker leaves the firm and is observed working at a different firm within 2 months. This allows us to test whether separations to employment are more responsive to wages than separations to non-employment. Manning (2003) shows that the overall separation elasticity can be written as a weighted average of the elasticity of separations to employment and non-employment. Langella and Manning (2021) consider what would happen to the overall level of monopsony power in the labour market in response to an increase in the job-offer arrival rate, and this depends on whether separations to employment or non-employment are more sensitive to wages. Second, we replace the firm-level average wage with the firm fixed effect and residual from a 2-way fixed effect model. Bassier et al. (2022) find that separation elasticities (and therefore labour supply elasticities) are larger when focussing on the firm-specific component of wages.

3.2. Direct estimation of labour supply elasticities

In our second approach, we directly estimate labour supply elasticities and compare those implied by the separation elasticity to those estimated directly. We use the same identification strategy as in identifying the separation elasticity. We use estimated MFP as a demand-shift instrument for the wage to trace out the labour supply curve. To estimate labour supply elasticities directly we estimate:

$$\ln L_{jt} = \psi_j + \eta \ln w_{jt} + X^L \beta^L + \tau_t + \omega_{jt} \quad (12)$$

where L_{jt} is employment in firm j in year t , ψ_j a firm fixed effect, X^L a vector of other controls, τ_t are year dummies, and ω_{jt} the error term. We estimate equation (12) using 2SLS, using lagged estimated MFP as an instrument for the log wage, as in equation (11).

3.3. Estimating markdowns

Most approaches in the literature focus on getting estimates of the labour supply elasticity, either directly or via separation or recruitment elasticities. This is informative of the ratio of wages to

¹³ We have two MFP estimates available to us in the productivity tables – one from a Cobb-Douglas production function and one from a translog function. We test the robustness of our results to using either estimate or both as our instrument. We use lags of MFP to reduce the likelihood of any shocks (e.g., changes to employment practices) that might simultaneously affect MFP, wages, or separations.

marginal product as shown by equation (2). A recent paper by Yeh et al. (2022) takes a different approach and seeks to estimate the ratio of marginal product to the wage directly, i.e. $\frac{MRP_L}{w}$.

Yeh et al. (2022) base their approach on that of De Loecker and Warzynski (2012). This approach uses the dual conditions of profit maximisation and cost minimisation to provide an expression for the wedge between an input's marginal product and its price that can be calculated using information on costs as well as production function parameters. Central to this approach is the existence of a “free” input, which can be adjusted instantaneously without adjustment costs, and is free from market imperfections in the input market. For any free input, h , the wedge between price and marginal revenue product can be calculated as:

$$\mu_h = \theta_h \left(\frac{c_h}{y} \right)^{-1} \quad (13)$$

where μ_h is the mark up in input h , θ_h is the output elasticity of input h , c_h is the total cost of input h , and y is revenue. For a free input, equation (13) measures product market power.

The estimator employed by Yeh et al. (2022) for the ratio of the marginal product of labour to the wage is of the double-ratio type:

$$\frac{MRP_L}{w} = \frac{\mu_L}{\mu_M} \quad (14)$$

where μ_M the markup on materials and both μ_L and μ_M are calculated using equation (13). The markup on materials is included to control for the fact that monopoly power in the product market will reduce demand for a labour via reduced production, which will result in a higher marginal product of labour. Failure to account for imperfect product market competition could lead the researcher to overstate the extent of monopsony power if part of the wedge between marginal product and the wage is due to product market power. The key assumption behind the markdown approach to estimating monopsony power is the presence of some input, other than labour, which is not subject to monopsony forces, faces no adjustment costs, and is chosen statically (i.e., input demand is chosen now to be used now).

Yeh et al. (2022) use materials as a free input, to separate product market monopoly from labour market monopsony. If the assumptions of the model are valid, then both μ_M and μ_L will be greater than 1. Yeh et al. (2022) argue that, even if materials input is subject to monopsony forces (which would also result in a wedge between the marginal product and price), their double-ratio estimator would tell us about the markdown on labour *relative* to the markdown on materials. In this case, the double ratio estimator could be less than 1 if materials are subject to greater monopsony power.

Costs and revenue are readily available in most firm-level datasets, so the implementation of these markup/markdown estimators relies on getting a good estimate of the output elasticities θ_h from a production function. There is a significant literature dedicated to the estimation of production functions (e.g., Olley and Pakes 1996; Levinsohn and Petrin 2003; Akerberg et al. 2015; Wooldridge 2009). The most significant econometric challenge in their estimation is the nature of MFP, which is unobservable to the econometrician but at least partially observable to the firm. Innovations to MFP can affect a firm's input choices and failure to account for this correlation between MFP and input choices will lead to biased estimates of the production function parameters and therefore biased estimates of markups/markdowns (Griliches and Mairesse 1998).

Several methods have been used to estimate production functions, accounting for the simultaneity between productivity innovations and input choices. Many of these fall under the general control function approach (e.g. Olley and Pakes 1996; Levinsohn and Petrin 2003; Akerberg et al. 2015; Wooldridge 2009; De Loecker and Warzynski 2012; Wooldridge 2015). These techniques use a variety

of timing assumptions along with a variable that responds to the productivity innovation to separate MFP from the input choice and recover the production function parameters.¹⁴

We closely follow the approach of Yeh et al. (2022) and estimate industry-specific gross-output translog production functions using the estimator of Akerberg et al. (2015) (ACF).¹⁵ The translog functional form nests other common functional forms of production functions (e.g. Cobb-Douglas, constant elasticity of substitution) and provides an approximation to any arbitrary production function. A translog functional form is also necessary for calculating firm-specific output elasticities and marginal products, meaning we can estimate the distribution of markdowns within industries. Other functional forms, such as Cobb-Douglas, has no within-industry variation in marginal products. Using a Cobb-Douglas functional form would therefore entail estimating the average markdown across firms within an industry.

In our main results, we estimate separate production functions for each of the 16 1-digit ANZSIC industries in our sample, with separate intercepts for different sub-industries. This is motivated by our relatively small sample sizes for more detailed industry categories and limited time series. All variables going into the estimation are measured relative to the industry average.

While the production approach has been widely used to estimate markups (and recently markdowns) (e.g. De Loecker and Warzynski 2012; De Loecker et al. 2020; Raval 2023; Foster et al. 2022), it is not without its critics. Bond et al. (2021) point out that, when using deflated revenue data (using industry-level deflators) rather than physical quantities (as is the case in many firm-level datasets), then one is not estimating the output elasticity but rather the revenue elasticity. They show that the ratio of an input's *revenue elasticity* to its revenue share should be identically one. They further suggest that identification of the production parameters is more difficult in situations involving market power without imposing additional restrictions on the demand side, or when inputs are used for purposes beyond producing output (e.g., marketing).

Hashemi et al. (2022) take the analysis from Bond et al. (2021) and show that, under the same conditions where Bond et al. (2021) find that markups cannot be identified using revenue data, that input market distortions can be identified. They show that the ratio of a firm's revenue elasticity to its revenue share recovers the input price wedge, rather than the output price wedge. A researcher doesn't need to compute the double-ratio markdown estimator in this framework, as the output price markup under these estimation conditions is not identified (i.e., using equation (13) rather than equation(14)). This result holds when inputs can be measured in physical quantities, which is the case with labour. We report both the double-ratio markdown estimate (as in Yeh et al. (2022)) and the single-ratio estimate suggested by Hashemi et al. (2022).

3.4. Converting the different estimates to a wage markdown

To aid interpretation, we convert all our estimates to a single metric, the wage markdown (equation 3) that measures the proportion of marginal product that is not paid as wages. By doing this we can compare our estimates from the different approaches in a consistent and easy to understand metric.

¹⁴ Olley and Pakes (1996) use investment as their proxy variable, while Levinsohn and Petrin (2003) use materials.

¹⁵ The main difference between our implementation ACF estimation and that used in Yeh et al. (2022) is the correction for measurement error in output used in the calculation of revenue shares. We do not implement this correction in our calculation of revenue shares.

4. Data

Our data are drawn from the rich individual- and firm-level information contained within StatsNZ's Integrated Data Infrastructure (IDI) and Longitudinal Business Database (LBD).¹⁶ Our population of interest is all employing firms and their workers. We identify firms using the Permanent Enterprise Number (PENT), which corrects for breaks in firm identifiers that may arise from factors such as changes of legal structure (e.g., moving from an individual proprietorship to a limited liability company) using the methodology of Fabling (2011).

Our key data on firm-level employment and wages are drawn from the IDI/LBD labour tables (Fabling and Maré, 2015b). These tables contain monthly job-level (person-firm pair) information on gross monthly earnings and estimated FTE labour input, that we aggregate to the firm-year level.

We link firm-level financial information from the LBD productivity tables (Fabling and Maré, 2015a; 2019; Fabling, 2021). These contain firm-year observations on gross output, intermediate expenditure, the value of capital services and, for a subset of private for-profit firms, estimated multi-factor productivity (MFP). We restrict attention to firms with at least 5 employees in the private for-profit sector that have an MFP estimate in the LBD productivity tables. This leaves us with a final sample of 257,445 annual observations on 39,114 firms over the period 2002-2019. While our sample includes approximately 10% of firms in the private-for-profit measured sector population, these firms account for over 80% of FTE employment in the population.

4.1. Separation and direct labour supply elasticity data

MEASURING EMPLOYMENT

Our measure of annual firm employment is the firm's average monthly headcount employment over the year. This measure treats monthly employment as zero in any month a firm doesn't have any employees. This information is sourced from the labour tables and is our key left hand side variable for our direct labour supply elasticity estimation. While an FTE-based measure of labour provides a better conceptual measure of the quantity of labour, the FTE measure available in the labour tables is endogenous to wages so is not used.¹⁷ Given the number of separations is inherently a headcount measure, we also use the headcount measure of employment to calculate separation rates.

MEASURING SEPARATIONS AND JOB-TO-JOB TRANSITIONS

First, a job spell is defined to end in a month, if the worker is not observed at the firm in the following two months.¹⁸ The labour tables contain flags for spell ends, corresponding to the final month of a job spell. To calculate the number of monthly separations, we simply sum the number of spell-ends within a firm-month.

¹⁶ See <https://www.stats.govt.nz/integrated-data/integrated-data-infrastructure/> for more information on the IDI. For more information on the LBD, see Fabling and Sanderson (2016).

¹⁷ While we do not directly observe an individual's FTE status in the administrative data, there are cases of obvious departures from full-time status where we can adjust. These include where a worker's earnings are less than full time (40 hours per week) minimum wage earnings, and earnings in the first and last month of an employment spell where it is likely a worker was present for only part of a month. FTE adjustments in these cases are based on earnings in the affected month relative to some reference, either the minimum wage or the worker's earnings in the intervening months. However, for many part time workers, it is not possible to distinguish high paid part time from low paid full-time earnings. As the FTE adjustment depends on wages, and wages are our key RHS variable, we do not use the FTE measure in our analysis.

¹⁸ This definition allows for one-month gaps in observed employment (e.g., due to taking unpaid leave). If a person "re-joins" a firm following the end of a job spell, this is interpreted as a new employment spell.

Second, we define a subset of these spell-ends as job-to-job transitions, if the worker moved to employment in another firm.¹⁹ In particular, for every job separation, we identify whether the individual has an associated spell start at a different firm in the three months following their last month at the previous firm.²⁰ The firm's monthly job-to-job separations is then the sum of these job-to-job separations in a month. The difference between the total number of separations and the total number of job-to-job separations is then the number of separations to non-employment.

From this monthly information, we calculate the average monthly number of separations and job-to-job separations within a 12-month period. We divide these by the annual measure of firm employment described in 4.1.1 above, to give our separation rate measures.

MEASURING WAGES

Our measure of wages is based on monthly wage and salary filings for tax purposes. For each firm-year, we sum the total monthly wages and salaries paid to workers who are not in their first or last month with the firm and are not obviously part-time.²¹ We calculate the average firm wage rate by dividing this sum by the average monthly number of workers who contribute to this sum. This is our best proxy for the underlying average firm wage rate and is standard in studies that use the Fabling-Mare tables (e.g. Maré et al. 2021; Allan & Maré, 2021).

One issue with using the average firm wage, rather than individual wage rates, is that changes in workforce composition, associated with hires and separations, may lead to changes in the average wage rate and induce a correlation between wages and separations in the absence of any monopsony power. We address this issue in two ways. First, we create a set of variables that describe the characteristics of a firm's workforce along the dimensions of age, gender, ethnicity, and urban/rural location. This helps to control for the confounding effects of changes in workforce composition.

The second way we deal with changes in firm composition is to follow the analysis of Bassier et al. (2022) and use the firm-specific components from a two-way fixed effect model of wages, introduced by Abowd et al. (1999):

$$\ln w_{ijt} = \alpha_i + \varphi_j + X_{it}\beta + \tau_t + e_{ijt} \quad (15)$$

where the worker fixed effect (α_i) is the time-invariant, transferable worker wage premium, the firm fixed effect (φ_j) is the time-invariant firm wage premium, $X_{it}\beta$ is a vector of time-varying individual characteristics, τ_t is a year-specific effect, and e_{ijt} the residual. Estimates of the components of a two-way fixed effect model are available in the Fabling-Mare tables, based on FTE earnings. Our measure of the firm specific component of wages is constructed as:

$$\tilde{w}_{jt}^F = \ln \bar{w}_{jt} - \bar{\alpha}_j - \bar{X}_{jt}\beta - \tau_t = \varphi_j + \bar{e}_{jt} \quad (16)$$

where bars denote firm averages. Including the residual provides within-firm variation that is independent of changes in workforce composition, which are captured by $\bar{\alpha}_j$ and $\bar{X}_{jt}\beta$.

¹⁹ The dynamic monopsony model introduced in Manning (2003) distinguishes between separations to employment and non-employment. Job-to-job separations is our proxy for separations to employment. Total separations include these, voluntary separations to non-employment, and involuntary job separations. Langella & Manning (2021) point out that the treatment of involuntary separations (e.g., layoffs, firings for disciplinary reasons) remain on the to-do list in monopsony models.

²⁰ Note, we require the new-firm is distinctly different from the firm the worker has left: a worker who has a 2-month gap in earnings and returns to the same firm in the third month is counted as a job-end but not a job-to-job transition. That is, we allow for at most two months of non-employment in defining job-to-job separations.

²¹ Workers are defined as obviously part-time if their monthly earnings are less than the equivalent of 40 hours per week at the minimum wage.

4.2. Production data for markdown estimation

The productivity data described in Fabling and Maré (2015a; 2019) combines data from survey and administrative sources into a harmonised dataset for the purposes of studying firm productivity.²² The dataset contains information on gross output, intermediate expenditure, the value of capital services, and total labour input. The labour input comes from the labour tables described above, with the addition of information on self-employed working proprietors (also in the set of labour tables) and is adjusted for obvious departures from full-time status to better approximate FTE labour input. The productivity tables cover approximately 70% of employing firms in the private-for-profit measured sector, accounting for about 80% of employment in the sector (Fabling and Maré, 2019).²³ These data have been used extensively to study issues related to firm productivity in New Zealand (Fabling, 2021; Fabling and Grimes, 2021; Chappell and Jaffe, 2018).

We use information on gross output, intermediate expenditure, total labour input, and the value of capital services. Total labour input is slightly different from the measure of employment described in 4.1.1 above, in that we utilise the FTE adjusted employment measure as this more closely approximates the quantity of labour used by a firm than a simple headcount measure. In addition, it includes labour input from any working proprietors. This more accurate measure of labour input is arguably important for obtaining reliable estimates of production function parameters and MFP.

MFP estimates are available for the subset of firms in the sample described at the beginning of this section. These estimates are derived from industry-specific gross output production functions, based on both Cobb-Douglas and translog production function specifications. These functions are estimated using OLS and contain year and firm fixed effects. MFP estimates are available for the subset of employing firms that are not either entering or exiting (i.e., not in their first or last year of operation). While this reduces the coverage in terms of the number of firms (particularly from excluding WP only firms which are a large fraction of the firm population), it has little effect on coverage in terms of output or inputs (Fabling 2021). We use these MFP estimates as instruments in the estimation of separation and labour supply elasticities.

SAMPLE SUMMARY STATISTICS

Table 1: Firm-level summary statistics

	Mean	Std. Deviation
Avg. wage	\$62,300	\$21,000
Gross output	\$14,312,000	\$151,640,000
Capital	\$1,732,000	\$20,226,000
Materials	\$8,223,000	\$124,316,000
Separation rate	6.3%	5.5%
Job-to-job rate	2.9%	3.1%
Employment	50	263
N	257,445	
N Firms	39,114	

Notes: Dollar values are expressed in 2018NZD. Number of observations and number of firms have been randomly rounded to base 3 for confidentiality purposes.

²² The survey data is from the Annual Enterprise Survey (AES) and the administrative data are annual tax financial statements (IR10).

²³ Working proprietor only firms are also included in the productivity tables, although the coverage is slightly lower at about 60% (Fabling & Maré, 2019)

Table 1 presents the means and standard deviations of our key variables. On average firms pay an annual wage of around \$62,000 and employs about 50 workers. An average firm produces roughly \$14 million of gross output, using \$8.2 million worth of materials and \$1.7 million worth of capital services. In the average month, 6.3% of the firm's workforce leaves the firm (separation rate) and just under half of the firm's separations are to another job (job-to-job separation rate of 2.9%).

4.3. Descriptive evidence of monopsony in New Zealand

We now begin to explore the extent of monopsony power in the New Zealand labour market using measures and methods suggested by Manning (2003) and Langella and Manning (2021). In what follows we utilise the monthly administrative job-level information from the labour tables developed by Fabling and Maré (2015b) as well as unit-record Household Labour Force Survey (HLFS) data.²⁴ We begin by looking at the fraction of new hires that come from a previous employer. A greater number of new hires that come from a previous employer suggests strong competition for existing workers. We then look at how the separation rate varies across the wage distribution, which provides some insight into how sensitive separations are to the wage and which segments of the labour market there may be more monopsony power. This analysis also provides the context in which to interpret our main results in section 5.

THE SHARE OF NEW HIRES FROM NON-EMPLOYMENT

First, the share of a firm's hires that don't come from another job is suggested as a macro indicator of monopsony power by Manning (2003), who shows that the share of recruits from non-employment is monotonically decreasing in the ratio of the job offer arrival rate to the job destruction rate. When the job offer arrival rate is high relative to the job destruction rate, more workers are hired from employment, implying greater competition for existing workers and less monopsony power for firms. As a result, workers' wages are closer to the competitive (marginal product) level and the wage markdown is reduced. Conversely, when the job offer arrival rate is low relative to the job destruction rate, there is less competition for existing workers as firms can expand employment by hiring unemployed workers (those whose jobs were destroyed in the previous period) without putting upward pressure on wages, and firms have greater monopsony power.

Figure 1 shows the proportion of hires that are from employment over the period January 2000 to December 2020.²⁵ This is derived from the labour tables. New hires are identified by the 'spell start' flags in the monthly job-level data.²⁶ We consider two characterisations of a new hire being from employment: first, if the worker had a spell-end at a different firm in the previous month; and second, if there was no more than a 3-month gap between their last firm's spell-end prior to the current firm's spell-start (i.e. at most 2 consecutive months of non-employment between spells).

Although the levels of the shares differ (by about 12 percentage points), both series follow similar trends. The share of new hires employed at a different firm the month prior was around 38% (50% for up to 3 months prior) prior to the GFC and this level was relatively constant over the period 2004-2008. These shares fell during the GFC, to 34% (45%). During this time, the hiring rate also fell as firms were shedding workers, and the unemployment rate increased. The share of new hires coming from other firms remained low from 2009 until 2014, after which it began to increase again as the

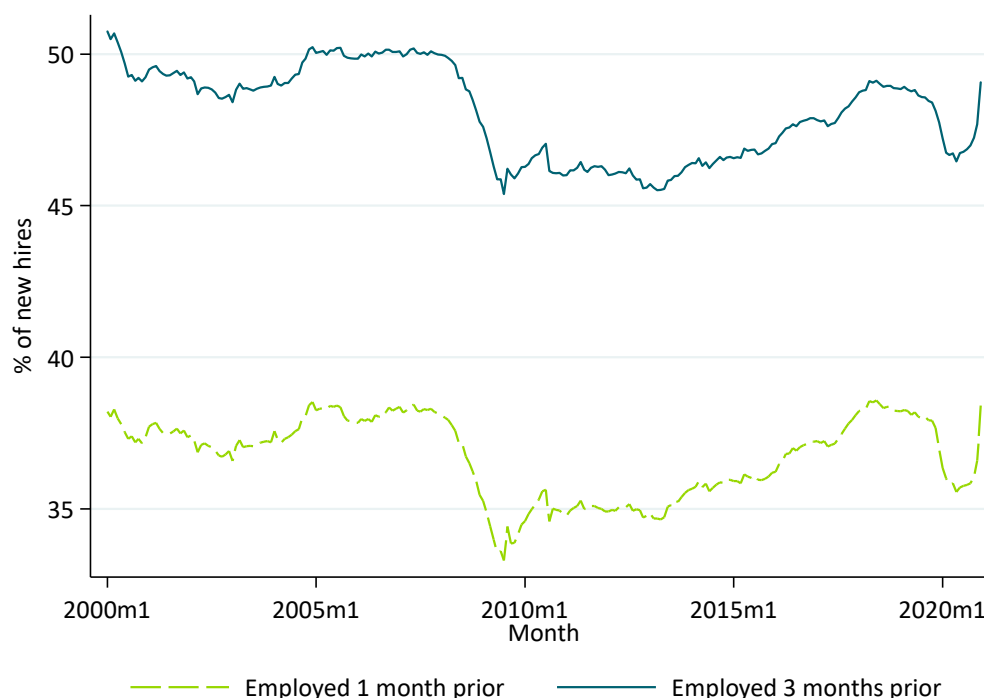
²⁴ The HLFS is a quarterly survey with an eight-quarter rolling panel – i.e., households are in the survey for up to eight quarters, with a fraction rotating in and out each quarter. Since 1997, an Income Supplement (IS) survey is administered to HLFS respondents in the June Quarter, which collects information on hourly and weekly earnings and other sources of income.

²⁵ The series plotted are 12-month symmetric moving averages of the underlying monthly series, to reduce the visual impact of the month-to-month volatility.

²⁶ New hires are on average 7% of employment each month in our data.

labour market began to improve and unemployment began to fall. In the 2 years prior to the COVID-19 pandemic, the share of new hires from previous employment had returned to pre-GFC levels. The dip at the end of the series reflects the strict lockdown implemented in March 2020 in response to the COVID-19 pandemic.

Figure 1: Share of new hires from previous employment 2000-2020



The share of new hires from previous employment found here are comparable to those reported in Manning (2003) for the US and UK, although there is no overlap in the time periods.²⁷ We wouldn't expect this share to be 100% given the entry of new workers into the labour market (e.g., from education and immigration), but ratios below 50% do seem low. This means that at least half of new hires were not in previous employment prior to starting their new job. Manning (2003) suggests this indicates significant monopsony power.

HOW SEPARATION AND QUIT RATES VARY WITH THE WAGE

Next, to get a better sense of the segments of the labour market where firms may possess greater monopsony power, we replicate the analysis of Langella and Manning (2021). The authors examine how separation and job-to-job transition rates vary with the wage level and conclude that firms have greater monopsony power over workers at the bottom of the wage distribution. They further examine how the share of quits to employment varies with the wage. This gives us an indication of whether quits to employment are more sensitive to the wage than quits to unemployment. If quits to employment are more sensitive, then when wages increase, the reduction in quits to employment will be greater than the reduction in quits to unemployment, share of quits to employment will be decreasing in the wage. If quits to unemployment are more sensitive, then the share will increase with the wage.

²⁷ Manning (2003) reports the fraction of new recruits from non-employment, which is the complement of what we report here. If a new hire is not from employment, then they must be from non-employment.

We begin by using data from the HLFS, as the analysis in Langella and Manning (2021) is done using household survey data for the UK and US. We focus on the June quarter surveys as information on hourly wages is collected only in the June quarter (via the income supplement). For individuals who are also interviewed in the September quarter, we take information on the labour force status and job tenure, which we use to infer job separations and job-to-job transitions between the June and September quarters. A worker has a job separation if they are employed in the June quarter and either not employed in the September quarter or if their reported tenure in the September quarter is less than the number of weeks between their June quarter and September quarter interview. The latter is our measure of a job-to-job transition.

To get a better idea of whether an individual is high or low paid relative to the labour market they are in, we follow Langella and Manning (2021) and use the residual from a standard earnings equation as our measure of the (log) wage. The earnings equation controls for a quadratic in age, gender, ethnicity, family type, highest qualification, whether the job is temporary, and 2-digit industry and occupation. We group individuals into bins based on the residualised wage, with each bin containing 5% of the (weighted) residualised wage distribution.

Figure 2 plots the marginal effects from a logit regression of whether an individual had a job separation on dummy variables for the bins, where the coefficients are the percentage difference in separation rates relative to the mean. A movement along the x-axis represents a percent change in wages and a movement along the y-axis a percent change in the probability of a separation (relative to the average probability), meaning the slope of this relationship can be interpreted as the separation elasticity. We see that the relationship between separation probabilities and the residualised wage is flat across approximately the lower half of the wage distribution. This shows the likelihood of leaving a job does not decline much with increases in the wage in the lower part of the wage distribution, indicating greater monopsony power. There is a strong, negative relationship between the residualised wage and separation probability for residualised wages 20%-60% higher than the average. The relationship then flattens out again at the top of the residualised wage distribution. These results are consistent with those in Langella and Manning (2021) – that separations are less sensitive to the wage at the extremes of the residualised wage distribution, suggesting firms have greater monopsony power over these workers. However, our estimates are relatively imprecise.

Figure 3 plots the results from a second logit regression, this time estimating the probability a worker has a job-to-job transition, conditional on having a job separation. This gives us information on whether separations to employment are more sensitive to wages than separations to non-employment. Although the imprecision of the estimates makes it difficult to say anything conclusively, the point estimates here are relatively flat across the distribution, suggesting quits to employment are just as sensitive to wages as quits to non-employment. The imprecision is largely due to our smaller sample size. Langella and Manning (2021) results for the US and UK suggest that quits to non-employment are more sensitive to the wage than quits to employment, although this result is not universal across studies (e.g., Bassier et al. 2022).

We next repeat the above exercise using administrative data. As this is based on the full employed population, it permits a much larger sample size, aiding the precision of our estimates. In addition, it will provide a test of whether we see similar patterns in the administrative data, which is the data we use to generate our main results.

Our measure of the residualised (log) FTE wage in the administrative data is derived from a 2-way fixed effect model of the form described in equation (15). We create two measures of the residualised wage. The first subtracts the individual-specific components of the 2-way fixed effect

model ($\alpha_i + X_{ijt}\beta$), and one which further subtracts the firm-specific component (φ_j).²⁸ The first more closely resembles the residualised wage from the HLFS, while the latter further controls for whether the firm is high or low paying in general. To be as consistent with the HLFS analysis as possible, we present results where the wage measure excludes only the individual-specific components, however the patterns are similar when we further exclude the firm-specific component of wages.

The data we use here is from the Fabling-Maré labour tables (Fabling and Maré, 2015b). We take monthly information from the June quarter (April, May, June) and, for workers who have a spell end during this period, look to see if they have a job-to-job transition in the three months following a job end. We take the spell end as the measure of a job separation. Note this is a slightly different definition of a job separation to the one used in the HLFS analysis. Here we use a monthly separation rate, for HLFS we used a quarterly separation rate.

Figure 4 repeats the analysis of Figure 2 and shows the results of a logit regression on the probability of a job separation across the residualised wage distribution from administrative data. Here we see a similar pattern to Figure 2 using HLFS data, but more precisely estimated. The slope of the curve is flatter at the bottom and top of the residualised wage distribution, and steeper between approximately -0.4 and 0.1. This is consistent with firms possessing greater monopsony power at the extremes of the wage distribution. The relatively flat part of the curve at the top of the wage distribution likely reflects that more highly paid workers place a greater weight on non-wage amenities when making labour supply decisions, which would be the case if these amenities are a normal good. Greater weight placed on non-wage aspects would lead to a lower wage elasticity of separations. At the bottom end, it more likely reflects limited outside employment options or a lack of accurate knowledge of the alternative employment options. Jäger et al. (2021) find that workers tend to anchor their expectations of pay at other firms at their current wage, meaning that low-paid workers at low paying firms will tend to underestimate wages elsewhere.

Figure 5 replicates Figure 3 using administrative data, looking at how the probability of a job-to-job separation varies with the wage. What is immediately clear is the improvement in precision from the increased sample size afforded by the administrative data. In general, the slope of the relationship is relatively flat, suggesting that, across most of the wage distribution, separations to employment are as sensitive to the wages as separations to non-employment. This is less true at the bottom end of the distribution, where the probability of a separation to employment is increasing in the wage. This indicates that separations to employment are less sensitive to the wage than separations to non-employment at the bottom of the wage distribution. This pattern differs somewhat from that in the UK and US found by Langella and Manning (2021). Their results suggest that quits to employment are less sensitive to the wage than quits to unemployment, and this relationship is strongest in the middle of the wage distribution. This implies that an increase in the job-offer arrival rate will increase the extent of monopsony power.

²⁸ The individual-level variables included in the $X\beta$ term are gender-specific quartics in age.

Figure 2: How the separation rate varies with the wage - HLFS data

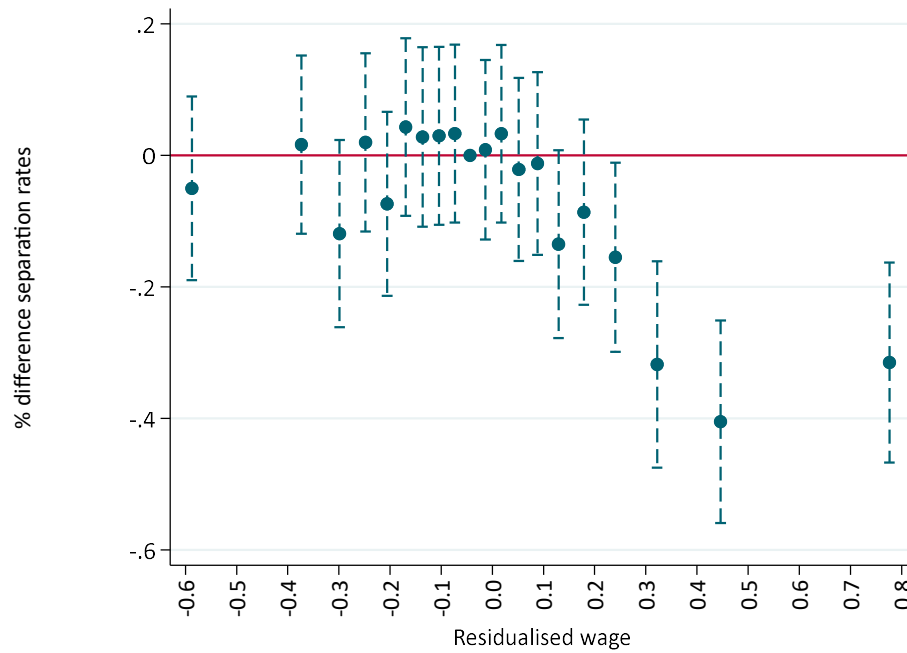


Figure 3: How the job-to-job separation rate varies with the wage – HLFS data

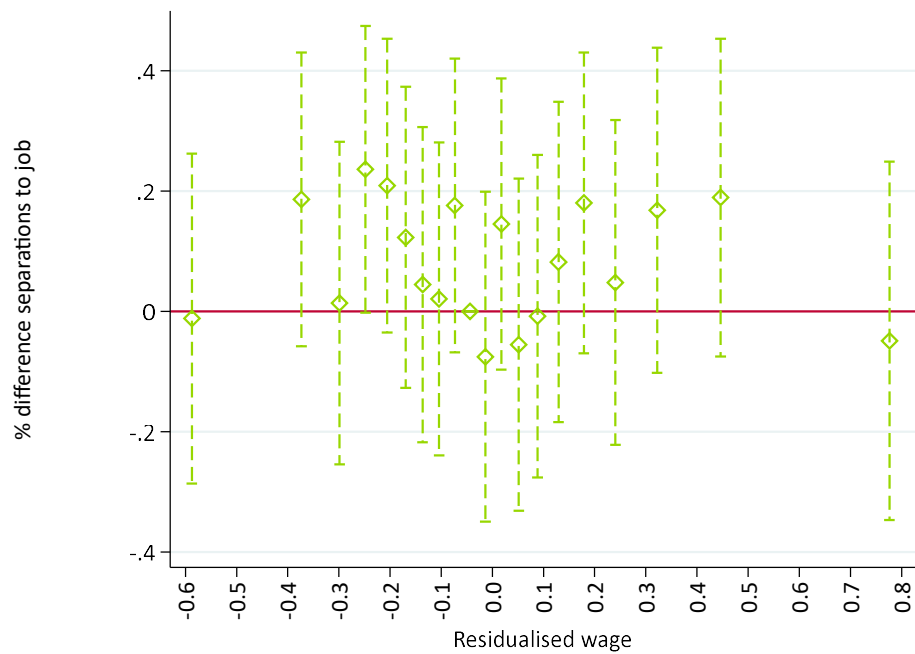


Figure 4: How the separation rate varies with the wage - administrative data

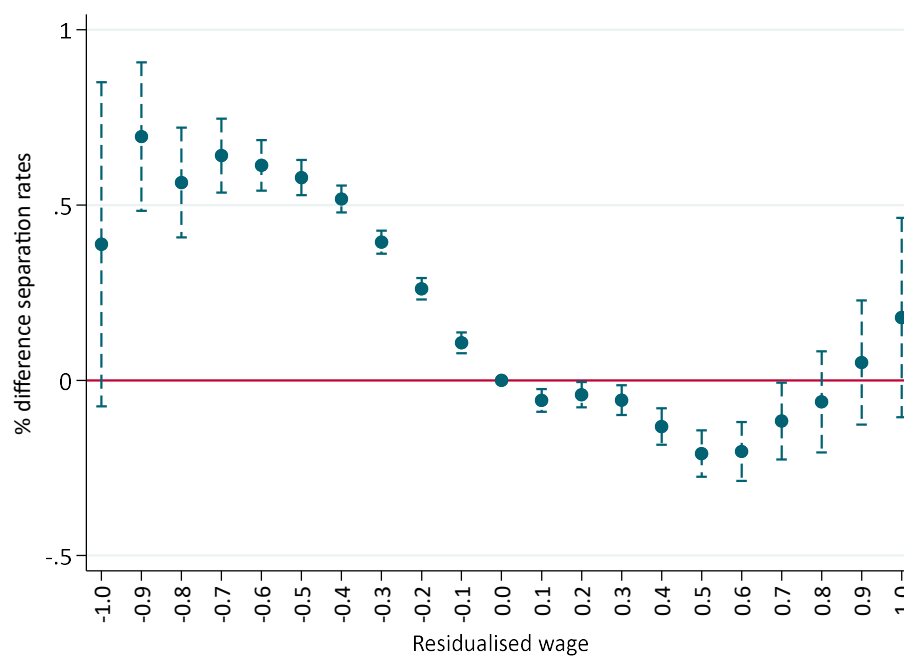
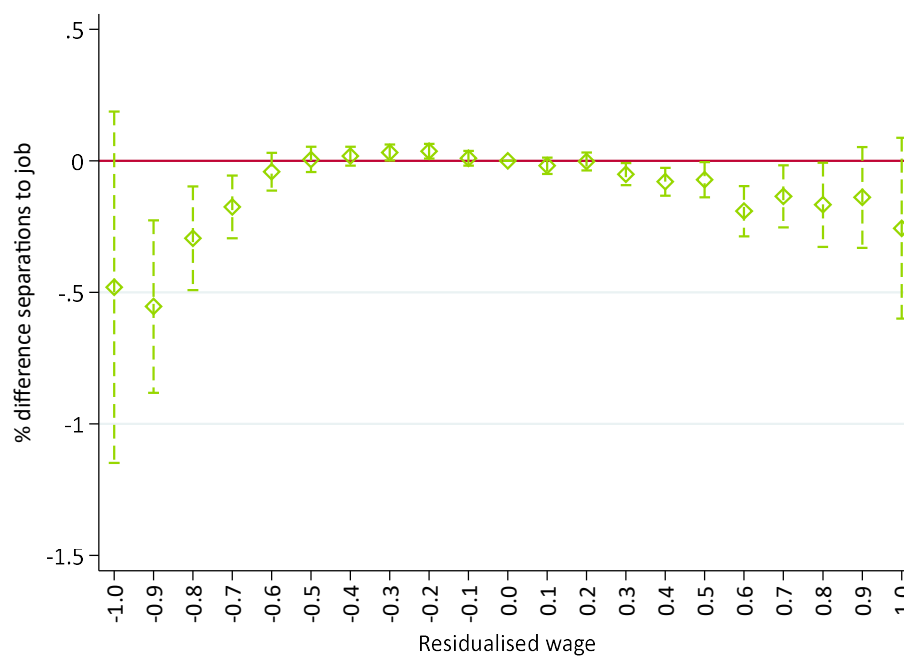


Figure 5: How the job-to-job separation rate varies with the wage – administrative data



These descriptive results are broadly consistent with those presented in Langella and Manning (2021). That is, a significant share of new hires come from non-employment, and that monopsony power appears strongest at the bottom of the labour market. However, the likelihood of a job-to-job transition relative to separation to non-employment is relatively constant throughout the wage distribution, suggesting that separations to employment are as sensitive to the wage as separations to non-employment. In the context of the Langella and Manning (2021) model, this suggests that any increases in the job-offer arrival rate is unlikely to change the extent of monopsony power that firms possess.

5. Results

We begin by presenting overall results for the extent of monopsony power in the NZ labour market using our various measures. Table 2 presents the results from estimating separation elasticities (η^S , using equation (11)), the ratio of marginal product to wage ($\frac{MPL}{w}$, using equations (12) and (13)), the implied labour supply elasticity (η), and the implied wage markdown ($\frac{1}{1+\eta}$, from equation (3)). The top panel reports unweighted results, while the bottom panel presents employment-weighted results.²⁹

Columns 1-4 report results based on the separation elasticity, estimated using equation (11).³⁰ In columns 1 and 2 total separations are used as the dependent variable, while in columns 3 and 4 job-to-job separations are the dependent variable. Columns 1 and 3 use the average firm wages as the wage measure, while columns 2 and 4 use the firm wage premium (AKM firm fixed effect plus residual).³¹

Comparing results across the weighted and unweighted results based on separation elasticities, we see that the implied labour supply elasticity is between 3 and 7, and the results are quite similar across the weighted and unweighted specifications. Results based on job-to-job separations tend to yield substantially higher estimated separation elasticities (and implied labour supply elasticities) with respect to wages than those based on total separations. As job-to-job separations are more likely to reflect voluntary decisions made on the part of the worker, we may expect such separations to be more sensitive to the wage. In contrast, total separations include workers who leave a firm because of exiting the labour market, redundancy, being fired, etc., and more likely driven by non-wage considerations or reflect a decision made by the firm rather than the worker.

We also see that separations (either total or job-to-job) are more sensitive to the firm wage premium than the average firm wage, a result consistent with Bassier et al. (2022). If part of the observed wage is some market return to the characteristics of workers that firms employ and not necessarily under the control of the firm, then we might expect that separations will be less sensitive to the average wage than the part of the wage that is more at firms' discretion.³² Replacing the average wage with the firm wage premium increases the separation elasticity by approximately 40%.

²⁹ For the unweighted results in the top panel, the separation elasticity estimates are from unweighted firm-level regressions, the firm-specific estimates $\frac{MPL}{w}$ are averaged across firms, and the implied markdown and labour supply elasticities are for the average firm. For the weighted results in the bottom panel, the separation elasticity estimates are from employment-weighted firm-level regressions, the firm-specific $\frac{MPL}{w}$ estimates are employment-weighted averages across firms, and the implied markdown and labour supply elasticities are for the firm that employs the average worker.

³⁰ Table A1 and Table A2 in Appendix A compare OLS and IV estimates of equation (11)

³¹ These results use lags of both Cobb-Douglas and translog MFP estimates as instruments. Results are robust to using either estimate in isolation.

³² This will depend on which individual characteristics correlate with both wages and the probability of a separation. For example, if more educated workers both earn higher wages and are more likely to move jobs, this would result in lower (less negative) estimates of separation elasticities and implied labour supply elasticities. On the other hand, age may be correlated with lower wages but higher turnover, which would yield higher (more negative) separation elasticity estimates. In practice, the direction of bias will depend on how different characteristics are correlated with both wages and separations, as well as the strength of these correlations. Our results here, and those of Bassier et al. (2022), suggest that high-wage individuals are more mobile.

Table 2: Estimates of monopsony power in New Zealand - separation elasticities and marginal product-wage ratios

	Separation elasticities				Estimates of $\frac{MPL}{w}$ (median)	
	(1)	(2)	(3)	(4)	(5)	(6)
	All separations, avg. wage	All separations, firm premium	J2J separations, avg. wage	J2J separations, firm premium	ACF, single ratio	ACF, double ratio
Unweighted						
η^s	-1.590*** (0.132)	-2.180*** (0.185)	-1.907*** (0.161)	-2.618*** (0.222)	---	---
$\frac{MPL}{w} = \frac{1 + \eta}{\eta}$	---	---	---	---	1.29	1.12
$\eta = -2 * \eta^s$	3.18	4.36	3.81	5.33	3.44	8.33
Markdown = $\frac{1}{1+\eta}$	23.9%	18.7%	20.8%	15.8%	22.5%	10.7%
Employment weighted						
η^s	-1.529*** (0.487)	-2.154*** (0.775)	-2.482*** (0.563)	-3.499*** (0.896)	---	---
$\frac{MPL}{w} = \frac{1 + \eta}{\eta}$	---	---	---	---	1.06	0.91
$\eta = -2 * \eta^s$	3.06	4.31	4.86	7.00	16.67	-11.11
Markdown = $\frac{1}{1+\eta}$	24.6%	18.8%	17.1%	12.5%	5.6%	-9.89%

Notes: Standard errors in brackets and are clustered at the firm level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively. Separation elasticities and IV FE production functions were estimated using the *ivreghdfe* command of Correia (2017) and Baum et al. (2010), while the ACF production functions used in the MPL/w calculations were estimated using the *prodest* command of Rovigatti and Mollisi (2018), all in Stata® 16. For the separation elasticity estimates, unweighted refers to estimates from unweighted firm-level regressions, whereas employment weighted refers to firm-level regressions where firm employment is used as the weight. All production function estimates are from unweighted firm-level regressions. From these, firm specific MPL/w ratios are calculated. Unweighted medians are simple medians of these firm specific MPL/w ratios across firms, whereas employment-weighted medians weight each firm by its employment in calculating the median.

The implied labour supply elasticities of between 3 and 7 suggest that the wage markdown is between 12% and 25%. Our estimates are similar to those reported in Card (2022) and at the higher end of the separation elasticity-based estimates reported in Sokolova and Sorensen (2021), although our estimates are within their range of best practice estimates from the literature. The headline estimates that we report in this paper are also those for the average worker. Recall from Figure 4 that the relationship between wages and the probability of separations is steepest in the middle of the wage distribution and flatter at the extremes. Our estimates reflect workers in the middle of the wage distribution and may therefore understate the degree of monopsony power firms have over low (or high) wage workers.

Using job-to-job separations, the implied markdown using the average wage measure is 21% in the unweighted results compared to 17% in the weighted results; while using the firm wage premium measure, the implied markdown is 16% unweighted and 12.5% weighted. This range of implied markdowns (12.5% – 21%) suggests that search frictions have a relatively even impact across the labour market and may be a source of monopsony power for both large and small firms.

Columns 5 and 6 in Table 2 report results from the production approach. Column 5 reports the median single-ratio estimate, and column 6 the median double-ratio estimate.³³ In the unweighted results, the median estimated ratio of marginal product to the wage is 1.29 using the single ratio, and 1.12 using the double-ratio estimate, showing that marginal products are between 12% and 21% larger than wages. This translates to estimated labour supply elasticities of between 3.4 and 8.3, and wage markdowns of between 10.7% and 22.5%. These estimates are broadly in line with the comparable unweighted separation elasticity-based estimates, where the implied markdowns range between 15% and 24%. The median estimate from Yeh et al. (2022) for US manufacturing is 1.36, or a 26.7% markdown.

The weighted production-based estimates, where we calculate an employment-weighted median from firm-level markdown estimates, are smaller than the corresponding unweighted estimates, implying less monopsony power. The single ratio estimate implies a wage markdown of 5.6%, while the double ratio estimate implies a wage markdown of -9.9%. Given the issues identified by Bond et al. (2021) and uncertainty over the validity of the assumption that materials are free from monopsony forces, our preferred estimates using the production approach are the single-ratio estimates.³⁴

That the weighted production-based results imply less monopsony power than the unweighted results may appear strange at first glance, particularly if we think monopsony power is positively correlated with firm size. We don't see a clear pattern in the weighted or unweighted results suggesting more or less monopsony power in the separation-based estimates. Estimates based on total separation suggest slightly more monopsony power in the weighted results, suggesting larger firms have marginally greater monopsony power. The opposite is true for estimates based on job-to-job transitions.

³³ In these columns, weighted and unweighted refer to the calculation of the median, not the estimation of the production functions. The unweighted median is the median markdown estimated across all firms in the sample, while the weighted median is estimated using employment weights.

³⁴ If the implied wage markdown ($\frac{1}{1+\eta}$ from equation (3)), is negative, this suggests that the markup on materials, used to control for product market power, is larger than the wage markdown. Yeh et al. (2022) suggests this may result from a failure of the assumption that materials are free from monopsony forces and means the markup on materials may reflect both product market and input market power. In these cases, the double-ratio estimate measures the relative amount of monopsony power that firms possess over workers and other materials suppliers. When the double ratio is less than 1, it suggests that firms have more price setting power in the materials input markets than in the labour market.

We further analyse the overall results in Table 2 by examining estimated markdowns by industry, and the estimates are shown in Table 3. The top panel shows unweighted estimates, while the bottom panel shows the weighted estimates. Columns 1 and 2 display separation-based estimates (using total separations), while columns 3 and 4 show the median single and double ratio estimates from the production approach. In the table, estimates outside the feasible range of 0-100% are in parentheses.

These comparisons are also shown graphically in Figure 6. Here the columns show separation-based estimates using the two wage variables – average observed wage and firm wage premium, while the rows show the unweighted and weighted results, respectively. In all graphs, we plot the single-ratio markdown estimate against the relevant separation-based markdown estimate. In the figures we censor low and high outliers at either 0 or 100, respectively. Correlation coefficients are shown in the top left corner in each figure. These are conditional on both the separation and production-based estimates being within the feasible (0 – 100) range.

Of the industries where the markdown estimates are in the feasible range, unweighted separation-based markdown estimates range between 2.5% in rental and real estates and 30% in construction. Production based estimates range between 5.2% in transport, postal and warehousing to 44% in professional and technical services. However, there are some industries where the estimated markdowns are negative. This occurs in in both approaches but is more prevalent in the in the production approach. Some of these industries are small and the estimates are likely poorly identified (e.g., mining, utilities). The negative results in the separation approach are the result of positive, but statistically insignificant estimates of the separation elasticity. In the production approach, given that the wage share of revenue is positive, it must come from negative estimates of the marginal product of labour, driven by the non-linear terms present in the translog production function.³⁵ The advantage of the translog function is that we can estimate firm-specific marginal products and then calculate firm-specific markdowns. However, estimates of the marginal products from the translog function can be volatile, particularly for firms with input factor ratios very different to the industry mean.

These issues are exacerbated in the weighted results. The production functions are estimated using mean-deviated variables, and the curvature of the production function is fitted around the mean. Some large firms may have input factor ratios that are far from the within-industry mean, meaning the estimates may be inappropriate for firms with extreme factor ratios, resulting in implausible estimates for the marginal products. Given that the overall production-based results are aggregations of firm-specific estimates based on industry-specific production functions, these large outlier firms have an outsized influence on the weighted results. This issue is exacerbated in the double-ratio estimates, where there are two sets of potentially poorly estimated marginal products. That the production approach suggests less monopsony power than the separation approach more likely reflects these identification and estimation challenges, rather than any structural difference in the two estimators.

Identification and estimation challenges notwithstanding, where industries do have reasonable estimates using both approaches, they are positively correlated. Among the small set of industries with reasonable estimates in the weighted results, the correlation is relatively strong. However, this is a small subset of industries and firms. Focussing on the unweighted results, industries where firms appear to have greater monopsony power include construction (E) and professional services (M), while firms in information media and telecommunications (I) and other services (S) appear to possess less monopsony power. Given the identification and estimation challenges involved in the production approach, our preferred set of estimates are those from the weighted separation approach.

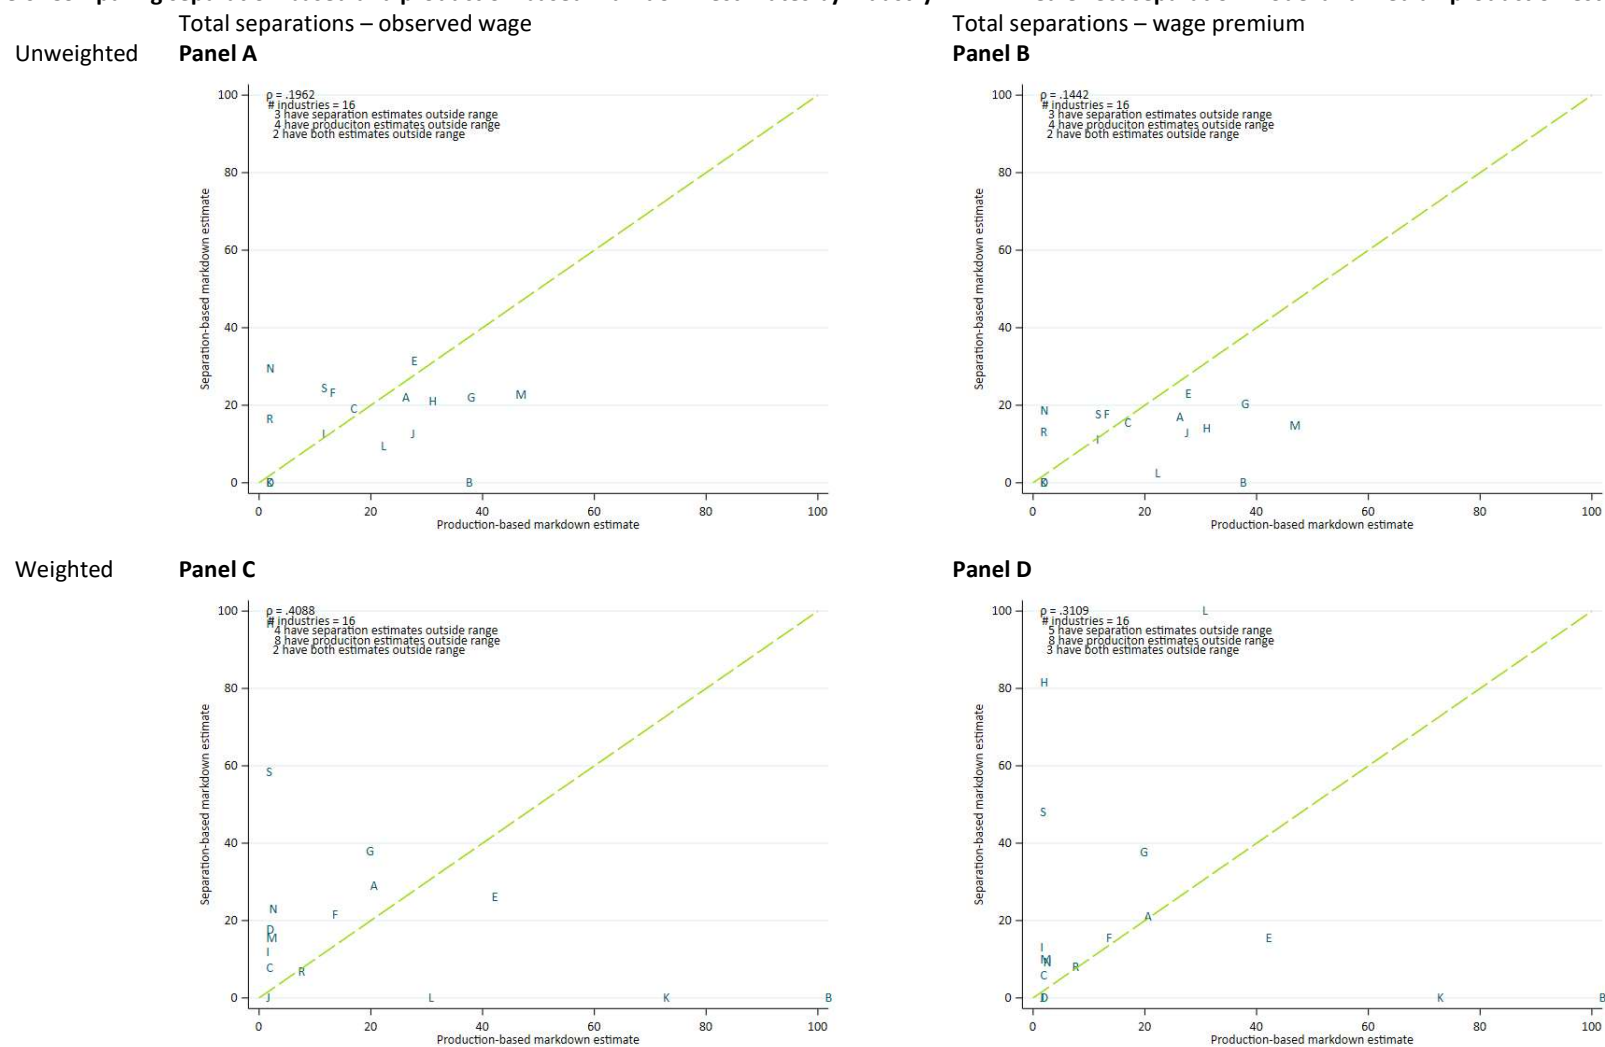
³⁵ These findings are consistent when using estimate for more detailed industry categories.

Table 3: Estimated wage markdowns by industry

		Total separations observed wage	Total separations firm wage premium	Production approach single ratio	Production approach double ratio
Unweighted					
A	Ag., forest and fish	22.1%**	16.9%**	24.3%	13.5%
B	Mining	(-21.5%)	(-9.4%)	35.7%	35.5%
C	Manufacturing	19.2%***	15.5%***	15.0%	9.7%
D	Utilities	(-29.1%)	(-6.7%)	(-36.6%)	(-69.2%)
E	Construction	31.3%***	22.9%***	25.9%	22.7%
F	Wholesale	23.3%***	17.7%***	11.4%	-32.9%
G	Retail	21.9%***	20.4%***	36.0%	24.0%
H	Hospitality	20.9%***	13.9%***	29.1%	26.1%
I	Trans., post. and warehousing	12.5%***	11.1%***	10.0%	5.2%
J	Info media and telecoms	12.6%	12.7%	25.9%	34.4%
K	Finance and insurance	(-109.9%)	(-204.1%)	(-32.0%)	(-92.5%)
L	Rental and real estate	9.5%	2.5%	20.5%	12.7%
M	Prof. and technical services	22.7%***	14.7%***	44.6%	23.3%
N	Admin. and support services	29.5%***	18.6%***	(-4.4%)	(-36.4%)
R	Arts and recreation	16.5%	13.1%	(-12.4%)	(-64.8%)
S	Other services	24.4%***	17.7%***	9.9%	(-30.2%)
Weighted					
A	Ag., forest and fish	29.0%***	21.1%***	18.6%	7.3%
B	Mining	(-21.2%)	(-8.1%)	(302.4%)	(509.6%)
C	Manufacturing	7.7%***	5.9%***	(-9.5%)	(-17.7%)
D	Utilities	17.6%*	(-5.3%)	(-35.3%)	(-115.2%)
E	Construction	26.0%*	15.5%*	40.3%	42.3%
F	Wholesale	21.5%**	15.4%**	11.8%	6.8%
G	Retail	37.9%	37.7%	17.8%	9.1%
H	Hospitality	96.7%	81.6%	(-213.3%)	(-249.8%)
I	Trans., post. and warehousing	11.8%*	13.0%	(-26.3%)	20.7%
J	Info media and telecoms	(-221.2%)	(-2.8%)	(-46.7%)	(-166.4%)
K	Finance and insurance	(-38.1%)	(-84.0%)	71.0%	(-36.4%)
L	Rental and real estate	(-97.8%)	(218.3%)	29.0%	51.0%
M	Prof. and technical services	15.4%***	10.0%***	(-29.4%)	(-89.9%)
N	Admin. and support services	22.9%***	9.2%**	0.5%	(-116.4%)
R	Arts and recreation	6.9%	7.9%	5.7%	(-54.3%)
S	Other services	58.4%	48.1%	(-33.6%)	(-99.1%)

Notes: estimates contained in parentheses are outside the feasible range of markdown estimates, 0-100%. ***, **, and * indicate that the underlying separation elasticity estimate is statistically significant at the 1%, 5%, and 10% level, respectively. We do not have standard errors for the production-based estimates.

Figure 6: Comparing separation-based and production-based markdown estimates by industry – firm-fixed effect separation model and median production estimates



Notes: Each figure compares separation-based and production-based markdown estimates by industry. Two sets of production-based estimates are used (weighted single ratio, unweighted single ratio), and four sets of separation-based estimates are used (total separations observed wage, total separations firm wage premium, both weighted and unweighted). Estimates outside the feasible range are censored at either 0 or 100, whichever is closer. The 45-degree line shows where separation-based and production-based estimates are equal.

Our estimates of direct labour supply elasticities are reported in Table A3 and Table A4 in Appendix A. These estimates are highly sensitive to the choice of weighting. Unweighted IV estimates produce labour supply elasticity estimates broadly consistent with estimates derived from separation elasticities. However, the weighted estimates are very different. Table A4 explores this issue by looking at the two MFP instruments in isolation. The sensitivity is most apparent when estimating an employment-weighted regression using MFP estimates from a translog production function. We believe the issues here stem from a failure of the exclusion restriction. While MFP is in principle a demand shifter, and doesn't affect the labour supply curve, its calculation as output less a combination of inputs (including labour), does generate a negative correlation between MFP and total labour which is independent of the wage. This relationship between MFP and firm size is stronger when MFP is calculated from translog production function and employment weights are used.

6. Conclusions

As far as we are aware, we provide the first estimates of the overall extent of monopsony power in New Zealand's private sector labour market. We estimate monopsony power using three distinct approaches. First, we estimate the wage elasticity of separations as suggested by models of dynamic monopsony. Second, we estimate firm labour supply curves directly. Third, we implement the production approach to estimating wage markdowns, following the approach of Yeh et al. (2022).

Using three alternative approaches we find evidence that monopsony power is present in New Zealand's labour market, with wage markdowns of up to 25%. There is a degree of overlap in the estimates from the separation elasticity and the production approach. While estimates from the separation approach tend to imply greater monopsony power than those based on the production approach, this is likely due to identification and estimation challenges in the estimation of production functions, rather than structural differences in the two approaches. Our estimates of wage markdowns are in line with recent international evidence surveyed by Card (2022) and are on the lower side of implied wage markdowns included in the meta-analysis of Sokolova and Sorensen, (2021). Our estimates are also in line with those from Townsend and Allan (2024), which estimates labour supply elasticities in a structural wage-posting model in New Zealand.

We have also explored variation in monopsony power across industries. These estimates are noisy and should be interpreted with some caution it appears however, that firms in the construction and professional services industries have greater monopsony power, while those in information media and other services have slightly less monopsony power. More work will need to be undertaken to get a better idea of the distribution of monopsony power across the labour market. Case studies of industries with different degrees of monopsony power may be useful in deepening our understanding of the factors contributing to differences in monopsony power across the economy.

We have estimated the extent of monopsony power possessed by the average firm or exerted over the average worker in New Zealand's private sector labour market. International evidence, and some descriptive evidence produced in this paper, suggest that monopsony may be greater at the low-wage end of the labour market (e.g., Langella and Manning, 2021; Webber, 2015). We have, however, excluded large sectors of the workforce where monopsony power may be more prevalent, such as the education and healthcare sectors. While New Zealand's dynamic labour market and flexible labour market settings may act as a counterweight to monopsony power, there may be pockets of the labour market where these counterweights are less effective. Future research should examine the distribution of monopsony power in different pockets of the labour market.

Our results imply the presence of important labour market frictions in New Zealand and the consequences of these frictions for workers. However, these results do not provide much guidance on the source of these frictions or whether they may be amenable to policy. Future work could usefully also seek to better understand the sources of monopsony power to better inform policy discussions around improving the functioning of New Zealand's labour market and improving outcomes for both workers and firms. The degree of monopsony power affects the extent to which the benefits and costs of labour market policies, and other policies that affect the labour market, accrue to employees or employers.

DISCLAIMER FOR MBIE

The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the author(s) and do not necessarily reflect the views of the Ministry of Business, Innovation and Employment or the New Zealand Government. The Ministry of Business, Innovation and Employment and the New Zealand Government take no responsibility for any errors or omissions in, or for the correctness of, the information contained here, or the results of any actions taken on the basis of information in this document. The paper is presented not as policy or advice, but with a view to inform and stimulate wider debate.

STATS NZ 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 and LBD 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.

ACKNOWLEDGEMENTS

We thank Wilbur Townsend, Michael Coelli, Jacques Poot, Tracy Mears, Natalie Labuschagne, and colleagues in MBIE's former Chief Economist Unit for comments on earlier drafts. This paper benefited from comments from participants at the 2022 New Zealand Association of Economists Conference, the 2023 Western Economic Association International Conference, attendees at a Productivity Commission Sausage Roll Seminar, and internal MBIE seminar participants.

7. References

- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67(2): 251-333.
- Akerberg, Daniel A, Kevin Caves, and Garth Frazer. 2015. "Identification properties of recent production function estimators." *Econometrica* 83 (6): 2411-2451.
- Allan, Corey, and David C Maré. 2022. "Who benefits from firm success? Heterogenous rent sharing in New Zealand." *MBIE Occasional Paper 22/02* (Ministry of Business, Innovation and Employment). <https://www.mbie.govt.nz/dmsdocument/20013-who-benefits-from-firm-success-heterogenous-rent-sharing-in-new-zealand-pdf>.
- Allan, Corey, and David C. Maré. 2021. "Do workers share in firm success? Pass-through estimates for New Zealand." *MBIE Occasional Paper 21/03* (Ministry of Business, Innovation and Employment). <https://www.mbie.govt.nz/dmsdocument/17166-do-workers-share-in-firm-success-pass-through-estimates-for-new-zealand>.
- Alves, Julian, Jason Greenberg, Yaxin Guo, Ravija Harjai, Bruno Serra, and John Van Reenen. 2024. "Labour market power: New evidence on Non-compete Agreements and the effects of M&A in the UK." *Centre for Economic Performance Discussion Paper 1976* (Centre for Economic Performance, LSE). <https://cep.lse.ac.uk/pubs/download/dp1976.pdf>.
- Andrews, Dan, and Bjorn Jarvis. 2023. "The ghosts of employers past: How prevalent are non-compete clauses in Australia?" *e61 Micro Note* (e61 Institute). <https://e61.in/the-ghosts-of-employers-past-how-prevalent-are-non-compete-clauses-in-australia/>.
- Ashenfelter, Orley C., David Card, Henry S. Farber, and Michael Ransom. 2022. "Monopsony in the labor market: New empirical results and new public policies." *Journal of Human Resources* 57 (S): S1-S10.
- Balusubramanian, Natarajan, Jin Woo Chang, Mariko Sakakibara, Jagadeesh Sivadasan, and Evan Starr. 2022. "Locked in? The enforceability of covenants not to compete and the careers of high-tech workers." *Journal of Human Resources* 57 (2): S349-S396.
- Barnes, Sebastian, Romain Bouis, Philippe Briard, Sean Dougherty, and Mehmet Eris. 2013. "The GDP impact of reform: A simple simulation framework." *OECD Economics Department Working Paper 834* (OECD Publishing). https://www.oecd-ilibrary.org/economics/the-gdp-impact-of-reform_5kgk9qjnhkmt-en.
- Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu. 2022. "Monopsony in movers: The elasticity of labor supply to firm wage policies." *Journal of Human Resources* 57 (S): s50-s86.
- Baum, C.F., M.E. Schaffer, and S Stillman. 2010. "ivreg2: Stata module for extended instrumental variables/2SLS, GMM, and AC/HAC, LIML and k-class regression." <http://ideas.repec.org/c/boc/bocode/s425401.html>.
- Bertram, Geoff, and Bill Rosenberg. 2022. "How policy and institutions affect the shares of wages, profit, and rent in New Zealand." *Paper presented at the 2022 NZAE Conference*. <https://geoffbertram.files.wordpress.com/2022/07/nzae-2022-session-1.1-bertram-rosenberg.pdf>.
- Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. 2021. "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data." *Journal of Monetary Economics* 121: 1-14.

- Borrows, Abi. 2022. *Show some restraint: Restraint of trade clauses in employment agreements*. Saunders and Co. Accessed 11 2022. <https://www.saunders.co.nz/news/show-some-restraint/>.
- Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage differentials, employer size, and unemployment." *International Economic Review* 68 (1): 257-273.
- Card, David. 2022. "Who set your wage?" *American Economic Review* 112 (4): 1075-90.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and labor market inequality: Evidence and some theory." *Journal of Labor Economics* 36: S13-S70.
- Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128 (3): 967-1015.
- Chappell, Nathan, and Adam Jaffe. 2018. "Intangible investment and firm performance." *Review of Industrial Organization* 52(4): 509-559.
- Competition and Markets Authority. 2024. *Competition and market power in UK labour markets*. London, UK: Competition and Markets Authority. <https://www.gov.uk/government/publications/competition-and-market-power-in-uk-labour-markets>.
- Correia, Sergio. 2017. "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects." <https://ideas.repec.org/c/boc/bocode/s457874.html>.
- Cowan, James. 2021. "Restraints of trade - "not worth the paper they are written on"?" *Employment News*, July, 12 ed. Accessed November 2022. <https://www.al.nz/restraints-of-trade-not-worth-the-paper-they-are-written-on/>.
- Criscuolo, Chiara et al. 2020. "Workforce composition, productivity and pay: The role of firms in wage inequality." *OECD Social, Employment and Migration Working Paper No. 241* (OECD Publishing). https://www.oecd-ilibrary.org/economics/workforce-composition-productivity-and-pay-the-role-of-firms-in-wage-inequality_52ab4e26-en.
- De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and firm-level export status." *American Economic Review* 102 (6): 2437-71.
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The rise of market power and the macroeconomic implications." *Quarterly Journal of Economics* 135 (2): 561-644.
- Fabling, Richard. 2011. "Keeping it together: Tracking firms in New Zealand's Longitudinal Business Database." *Motu Working Paper 11-01* (Motu Economic and Public Policy Research). <https://www.motu.nz/our-research/population-and-labour/longitudinal-household-survey/keeping-it-together-tracking-firms-in-new-zealands-longitudinal-business-database/>.
- Fabling, Richard. 2021. "Living on the edge: An anatomy of New Zealand's most productive firms." *Motu Working Paper 21-01* (Motu Economic and Public Policy Research). <https://www.motu.nz/our-research/productivity-and-innovation/firm-productivity-and-performance/living-on-the-edge-an-anatomy-of-new-zealands-most-productive-firms-2/>.
- Fabling, Richard, and Arthur Grimes. 2021. "Picking up speed: Does ultrafast broadband increase firm productivity?" *Information Economics and Policy* 57.
- Fabling, Richard, and Arthur Grimes. 2019. "Ultra-fast broadband, skill complementarities, gender and wages." In *Motu Working Paper 19-23*. Wellington: Motu Economic and Public Policy

- Research. <https://www.motu.nz/our-research/urban-and-regional/infrastructure/ultra-fast-broadband-skill-complementarities-gender-and-wages/>.
- Fabling, Richard, and David C. Maré. 2015b. "Addressing the absence of hours information in linked employer-employee data." *Motu Working Paper 15-17* (Motu Economic and Public Policy Research). <https://www.motu.nz/our-research/productivity-and-innovation/firm-productivity-and-performance/addressing-the-absence-of-hours-information-in-linked-employer-employee-data/>.
- Fabling, Richard, and David C. Maré. 2019. "Improved productivity measurement in New Zealand's Longitudinal Business Database." *Motu Working Paper 19-03* (Motu Economic and Public Policy Research). <https://www.motu.nz/our-research/productivity-and-innovation/firm-productivity-and-performance/improved-productivity-measurement-in-new-zealands-longitudinal-business-database/>.
- Fabling, Richard, and David C. Maré. 2015a. "Production function estimation using New Zealand's Longitudinal Business Database." *Motu Working Paper 15-15* (Motu Economic and Public Policy Research). <https://www.motu.nz/our-research/productivity-and-innovation/firm-productivity-and-performance/production-function-estimation-using-new-zealands-longitudinal-business-database/>.
- Fabling, Richard, and Lynda Sanderson. 2016. "A Rough Guide to New Zealand's Longitudinal Business Database (2nd edition)." *Motu Working Paper 16-03* (Motu Economic and Public Policy Research). <https://www.motu.nz/our-research/productivity-and-innovation/firm-productivity-and-performance/a-rough-guide-to-new-zealands-longitudinal-business-database-2nd-edition/>.
- Falch, Toberg. 2010. "The elasticity of labour supply at the establishment level." *Journal of Labor Economics* 28 (2): 237-266.
- Federal Trade Commission. 2023. "FTC Proposes Rule to Ban Noncompete Clauses, Which Hurt Workers and Harm Competition." <https://www.ftc.gov/news-events/news/press-releases/2023/01/ftc-proposes-rule-ban-noncompete-clauses-which-hurt-workers-harm-competition>.
- Foster, Lucia S., John C. Haltiwanger, and Cody Tuttle. 2022. "Rising markups or changing technology?" *NBER Working Paper 30491* (National Bureau of Economic Research). <https://www.nber.org/papers/w30491>.
- Griliches, Zvi, and J. Mairesse. 1998. "Production functions: The search for identification." In *Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial Symposium*, edited by Steinar Strøm, 169–203. Cambridge: Cambridge University Press.
- Hashemi, Arshia, Ivan Kirov, and James Traina. 2022. "The production approach to markup estimation often measures input distortions." *Economics Letters* 217: 110673.
- Hirsch, Boris, Elke J. Jahn, and Claus Schnabel. 2018. "Do employers have more monopsony power in slack labor markets?" *ILR Review* 71 (3): 676–704.
- Hirsch, Boris, Thorsten Schank, and Claus Schnabel. 2010. "Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany." *Journal of Labor Economics* 28 (2): 291–330. doi:10.1086/651208.
- Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer. 2024. "Worker beliefs about outside options." *Quarterly Journal of Economics* 139 (3): 1505-1556.

- Kleibergen, F., and R. Paap. 2006. "Generalized Reduced Rank Tests Using the Singular Value Decomposition." *Journal of Econometrics* 133 (1): 97-126.
- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2019. "Who Profits from Patents? Rent-Sharing at Innovative Firms." *Quarterly Journal of Economics* 134 (3): 1343-1404.
- Kruger, Alan B., and Orley Ashenfelter. 2022. "Theory and evidence on employer collusion in the franchise sector." *Journal of Human Resources* 57 (S): S324-S348.
- Langella, Monica, and Alan Manning. 2021. "Marshall Lecture 2020: The measure of monopsony." *Journal of the European Economic Association* 19 (6): 2929-2957.
- Levinsohn, James, and Amil Petrin. 2003. "Estimating production functions using inputs to control for unobservables." *Review of Economic Studies* 70: 317-341.
- Manning, Alan. 2011. *Imperfect competition in the labor market*. Vol. 4, in *Handbook of labor economics*, edited by David Card and Orley Ashenfelter, 973-1041. Elsevier.
- . 2003. *Monopsony in Motion*. Princeton: Princeton University Press.
- Manning, Alan. 2021. "Monopsony in the labor market: A review." *Industrial and Labor Relations Review* 74 (1): 3-26.
- Maré, David C. 2022. "The labor market in New Zealand 2000-2021." *IZA World of Labor* 427.
- Maré, David C., and Dean R. Hyslop. 2006. "Worker-firm heterogeneity and matching: An analysis using worker and firm fixed effects estimated from LEED." *Statistics New Zealand LEED Research Paper* (Statistics New Zealand).
- Maré, David C., Lynda Sanderson, and Richard Fabling. 2021. "Earnings and Employment in Foreign-Owned Firms." In *Labor Markets, Migration, and Mobility: Essays in Honor of Jaques Poot*, edited by William Cochrane, Michael P. Cameron and Omoniyi Alimi, 1-40. Singapore: Springer.
- Matsudaira, Jordan D. 2014. "Monopsony in the low-wage labor market? Evidence from minimum nurse staffing regulations." *Review of Economics and Statistics* 96 (1): 92-102.
- OECD. 2021. *The Role of Firms in Wage Inequality: Policy Lessons from a Large Scale Cross-Country Study*. Paris: OECD Publishing. https://www.oecd.org/en/publications/2021/12/the-role-of-firms-in-wage-inequality_101116ef.html.
- Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64: 1263-1297. <http://www.jstor.org/stable/2171831>.
- Prescott, JJ, and Evan Starr. 2024. "Subjective Beliefs about Contract Enforceability." *The Journal of Legal Studies* 53 (2): 435-488.
- Ransom, Michael R., and David P. Sims. 2010. "Estimating the firm's labor supply curve in a "new monopsony" framework: Schoolteachers in Missouri." *Journal of Labor Economics* 28 (2): 331-355.
- Raval, Devesh. 2023. "Testing the production approach to markup estimation." *Review of Economic Studies* 90 (5): 2592-2611.
- Rovigatti, Gabriele, and Vincenzo Mollisi. 2018. "Theory and practice of total-factor productivity estimation: The control function approach using Stata." *The Stata Journal* 18 (3): 618-662.
- Ryall, Sue, and Stephen Blumenfeld. 2019. *Unions and Union Membership in New Zealand - report on 2017 Survey*. Wellington: Centre for Labour, Employment and Work.

- Sánchez, Rafael, Javier Finot, and Mauricio G. Villena. 2022. "Gender wage gap and firm market power: Evidence from Chile." *Applied Economics* 54 (18): 2109-2121.
- Sin, Isabelle, Steven Stillman, and Richard Fabling. 2022. "What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, Bargaining and Discrimination." *Review of Economics and Statistics* 104 (4): 636-651.
- Sokolova, Anna, and Todd Sorensen. 2021. "Monopsony in labor markets: A meta-analysis." *Industrial and Labor Relations Review* 74 (1): 27-55.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2019. "Firming up inequality." *Quarterly Journal of Economics* 134 (1): 1-50.
- Staiger, Douglas O., Joanne Spetz, and Phibbs Ciaran. 2010. "Is there monopsony in the labor market? Evidence from a natural experiment." *Journal of Labor Economics* 28 (2): 211-236.
- Starr, Evan P., J.J. Prescott, and Norman D. Bhisara. 2021. "Noncompete Agreements in the US Labor Force." *The Journal of Law and Economics* 64 (1): 53-84.
- Starr, Evan, JJ Prescott, and Norman Bishara. 2020. "The Behavioural Effects of (Unenforceable) Contracts." *Journal of Law, Economics, and Organization* 36 (3): 633-687.
- Townsend, Wilbur. 2023. "Estimating separation elasticities when firms differ."
- Townsend, Wilbur, and Corey Allan. 2024. "How restricting migrants' job options affects both migrants and existing residents." <https://wilburtownsend.github.io/papers/visas.pdf>.
- Webber, Douglas A. 2015. "Firm market power and the earnings distribution." *Labour Economics* 35: 123-134.
- Webber, Douglas A. 2016. "Firm-level monopsony and the gender pay gap." *Industrial Relations: A Journal of Economy and Society* 55: 323-345.
- Webber, Douglas A. 2022. "Labor market competition and employment adjustment over the business cycle." *The Journal of Human Resources* 57 (S): S87-S110.
- Wooldridge, Jeffrey M. 2015. "Control Function Methods in Applied Econometrics." *Journal of Human Resources* 50 (2): 420-445.
- Wooldridge, Jeffrey M. 2009. "On estimating firm-level production functions using proxy variables to control for unobservables." *Economics Letters* 104: 112-114.
- Yeh, Che, Claudia Macaluso, and Brad Hershbein. 2022. "Monopsony in the US labor market." *American Economic Review* 112 (7): 2099-2138.

Appendix A – comparing OLS and IV estimates of separation elasticities and labour supply elasticities

Table A1: Separation elasticity estimates - comparing OLS and IV for total separations

	OLS – wage	OLS – premium	IV – wage	IV – premium
Unweighted				
$\hat{\eta}^S$	0.112*** (0.0163)	0.284*** (0.0173)	-1.590*** (0.132)	-2.180*** (0.185)
N	257,445	257,445	257,445	257,445
R2	0.724	0.725	-	-
Weak IV	-	-	296.9	341.7
OverID p-value	-	-	0.00737	0.0141
UnderID p-value	-	-	0.000	0.000
Employment weighted				
$\hat{\eta}^S$	0.0453 (0.0492)	0.369*** (0.0533)	-1.529*** (0.487)	-2.154*** (0.775)
N	257,445	257,445	257,445	257,445
R2	0.857	0.857	-	-
Weak IV	-	-	31.19	37.64
OverID p-value	-	-	0.471	0.611
UnderID p-value	-	-	0.000	0.000

Notes: Standard errors, shown in parentheses, are clustered at the firm level. All regressions include year dummies, firm fixed effects, and a set of workforce demographic controls (% of workforce by gender, ethnicity, age groups, and the fraction of the workforce in primary urban areas, secondary urban areas, and rural areas). Weak IV is the Kleibergen-Paap Wald rk F statistic (Kleibergen Paap, 2006). UnderID p-value is the p-value from the Kleibergen-Paap rk LM test for underidentification. OverID p-value is the p-value from Hansen's J test of overidentifying restrictions

Table A2: Separation elasticity estimates - comparing OLS and IV for job-to-job separations

	OLS – wage	OLS – premium	IV – wage	IV – premium
Unweighted				
$\hat{\eta}^S$	0.0368* (0.0192)	0.132*** (0.0205)	-1.907*** (0.161)	-2.618*** (0.222)
N	257,445	257,445	257,445	257,445
R2	0.652	0.652	-	-
Weak IV	-	-	296.9	341.7
OverID p-value	-	-	6.63e-06	2.44e-05
UnderID p-value	-	-	0.000	0.000
Employment weighted				
$\hat{\eta}^S$	-0.0945 (0.0611)	0.195** (0.0782)	-2.482*** (0.563)	-3.499*** (0.896)
N	257,445	257,445	257,445	257,445
R2	0.811	0.811	-	-
Weak IV	-	-	31.19	37.64
OverID p-value	-	-	0.257	0.428
UnderID p-value	-	-	0.000	0.000

Notes: See notes to Table A1.

Table A3: Direct estimates of labour supply elasticities

	OLS – wage	OLS – premium	IV – wage	IV – premium
	Unweighted			
$\hat{\eta}$	0.0298 (0.0211)	0.489*** (0.0156)	0.0688 (0.162)	1.585*** (0.207)
Implied markdown = $\frac{1}{1+\hat{\eta}}$	97%	67%	93%	38.7%
N	257,445	257,445	257,445	257,445
R ²	0.942	0.943	-	-
Weak IV	-	-	303.9	382.2
OverID p-value	-	-	0.000	0.000
UnderID p-value	-	-	0.000	0.000
	Employment weighted			
$\hat{\eta}$	-0.388*** (0.112)	0.316*** (0.0816)	-2.484*** (0.713)	-1.364 (1.025)
Implied markdown = $\frac{1}{1+\hat{\eta}}$	163%	76%	-67%	-274%
N	257,445	257,445	257,445	257,445
R ²	0.988	0.988	-	-
Weak IV	-	-	29.11	40.09
OverID p-value	-	-	0.000	0.000
UnderID p-value	-	-	0.000	0.000

Notes: See notes to Table A1.

Table A4: Direct estimates of labour supply elasticities - comparing instruments

	IV (CD)– wage	IV (CD)– premium	IV (TL)– wage	IV (TL)– premium
	Unweighted			
$\hat{\eta}$	3.320*** (0.229)	3.993*** (0.235)	0.605*** (0.165)	0.808*** (0.211)
Implied markdown = $\frac{1}{1+\hat{\eta}}$	23%	20%	62.3%	55.3%
N	257,445	257,445	257,445	257,445
Weak IV	562.2	733.9	604.9	604.9
OverID p-value	-	-	-	-
UnderID p-value	0.000	0.000	0.000	0.000
	Employment weighted			
$\hat{\eta}$	3.178*** (0.855)	4.036*** (0.974)	-1.174* (0.676)	-1.733* (1.052)
Implied markdown = $\frac{1}{1+\hat{\eta}}$	23.9%	19.9%	-574.7%	-136.4%
N	257,445	257,445	257,445	257,445
Weak IV	60.59	78.70	47.53	47.53
OverID p-value	-	-	-	-
UnderID p-value	-	-	-	-

Notes: See notes to Table A1. CD indicates that MFP derived from a Cobb-Douglas production function is used as the instrument. TL indicates MFP derived from a Translog production function is used as the instrument.