Essays in the application of quasi-experimental methods to linked administrative data

Lisa Meehan

A thesis submitted to Auckland University of Technology
in fulfillment of the requirement for the degree of
Doctor of Philosophy (PhD)

2024

Primary supervisor: Professor Gail Pacheco
Secondary supervisor: Associate Professor Peer Skov

Thesis Abstract

Each of the three papers in this thesis applies quasi-experimental methods to administrative data in order to make causal inferences about the effects of a particular event or policy.

Paper 1: Does citizenship improve migrant parents' integration and children's outcomes?: Evidence from a natural experiment

Does obtaining citizenship improve the integration of immigrants into the host country's society? Does it improve the outcomes of immigrants' children? Attempts to answer these questions are thwarted by potential self-selection: those who are more motivated to integrate are also more likely to naturalise. In order to make causal inferences, this paper exploits a natural experiment of New Zealand's removal of birthright citizenship. It finds that this policy change had no effect on family outmigration behaviour, parents' fertility and labour market outcomes, nor on children's health outcomes. These results contrast to existing evidence on Germany's introduction of birthright citizenship, which was found to have positive effects on parents' and children's integration outcomes. These different findings may reflect differences in the two countries' broader immigration policy settings, and the resulting differences in the characteristics of the countries' migrants. These differences also highlight that the New Zealand case is of potentially more relevance to Canada and the US, the only two western countries that retain unrestricted birthright citizenship, given greater similarities in these countries' immigration policy settings.

Paper 2: Workforce vaccine mandates: The effect on vaccine uptake and healthcare workers' labour market outcomes

As part of its COVID-19 policy response, the New Zealand government implemented vaccination mandates as a condition of ongoing employment for certain workers. This paper examines the effect of these mandates on vaccination uptake among mandated healthcare, education and corrections workers and on healthcare workers' labour market outcomes. This is enabled by New Zealand's linked population-wide administrative data, which includes a comprehensive national vaccination register linked to tax records to identify employment outcomes.

Overall, the results suggest that in the context of already-high vaccination rates, workforce vaccine mandates provided limited benefit in terms of increasing vaccination rates among mandated workers. Moreover, they negatively impacted healthcare workers' labour market outcomes, which may have had wider consequences in terms of exacerbating existing health workforce skills shortages.

Paper 3: The effect of a minor health shock on labour market outcomes: The case of concussions

The literature on health shocks finds that minor injuries have only short-term labour market impacts. However, mild traumatic brain injuries (mTBIs, commonly referred to as concussions) may be different as the medical literature highlights that they can have longer-term health and cognitive effects. Moreover, TBIs are one of the most common causes of disability globally, with the vast majority being mild. Thus, it is important to understand the impact of mTBIs on labour market outcomes.

We use administrative data on all medically-diagnosed mTBIs in New Zealand linked to monthly tax records to examine the labour market effects of a mTBI. We use a comparison group of those who suffer a mTBI at a later date to overcome potential endogeneity issues, and employ a doubly-robust difference-in-differences method. We find that suffering a mTBI has negative labour market effects. Rather than dissipating over time, these negative effects grow, representing a decrease in employment of 20 percentage points and earning losses of about a third after 48 months. Our results highlight the need for timely diagnosis and treatment to mitigate the effect of mTBIs to reduce economic and social costs.

Contents

Thesis Abstract i

Introduction 1

References 6

I Does citizenship improve migrant parents' integration and children's outcomes?: Evidence from a natural experiment 7

1 Introduction 10

2 Policy context 12
2.1 Background on NZ's immigration settings 12
2.2 Removal of birthright citizenship 14
2.3 The benefits of NZ citizenship 15

3 Existing evidence 18
3.1 Parental outcomes 20
3.2 Children's outcomes 22

4 Data 23
4.1 Population of interest 24
4.2 Treatment and control group definitions 26
4.3 Outcome measures 26
4.4 Other control variables 28

5 Methodology 30
5.1 Validity of identification strategy 30
5.2 Covariate balance 34

6 Results 38
6.1 Outmigration 38
6.2 Fertility 42
6.3 Labour market outcomes 43
6.4 Children's health outcomes 45

7 Robustness **50**

8 Discussion and conclusion **50**

References **53**

Appendices **57**

A Robustness: Varying treatment window graphs **58**

II Workforce vaccine mandates: The effect on vaccine uptake and healthcare workers' labour market outcomes **66**

1 Introduction **69**

2 Background and policy context **70**

3 Literature **74**

4 Data **79**

4.1 RQ1 vaccine uptake . . . 79

4.1.1 Creating a sample of employed individuals . . . 79

4.1.2 Identifying individuals subject to COVID-19 vaccination mandates 80

4.1.3 Identifying compliance with COVID-19 vaccination mandates . 82

4.2 RQ2 healthcare labour market outcomes . . . 84

4.2.1 Creating a sample of employed individuals . . . 84

4.2.2 Identifying individuals' vaccination status . . . 85

4.2.3 Identifying individuals' overseas spells . . . 85

4.2.4 Defining labour market outcomes . . . 86

4.3 RQ1 and RQ2 additional descriptive variables . . . 88

5 Method **88**

5.1 RQ1: Vaccine uptake . . . 88

5.2 RQ2: Healthcare worker labour market outcomes . . . 89

6 Results: RQ1 - Vaccine uptake **93**

6.1 Summary statistics . . . 93

6.2 Vaccination uptake over time . . . 100

7 Results: RQ2 - Healthcare workers labour market outcomes **103**

7.1 Tracking workers' outcomes over time . . . 103

7.2 Triple difference: Estimating the role of mandates in HCWs' labour market outcomes . . . 107
7.3 Robustness: Two-period triple difference . . . 109
7.4 Heterogeneity analysis . . . 109

8 Policy discussion **118**

9 Conclusion **119**

References **121**

Appendices **126**

A Difference-in-differences: Estimating the role of mandates in vaccination uptake **127**

B Two-period triple difference regression results **128**

III The effect of a minor health shock on labour market outcomes: The case of concussions **129**

1 Introduction **132**

2 Literature review **134**

3 Empirical strategy **136**
3.1 Model . . . 137
3.2 Identification . . . 139

4 Data and population of interest **139**

5 Results **143**
5.1 Monthly treatment effects . . . 144
5.2 Heterogeneity analysis . . . 147
5.3 Effects of mTBIs on accident compensation payments . . . 154

6 Conclusion **158**

References **160**

Conclusion **164**

List of Figures

I.1 Density of births for population of interest 33
I.2 Average number of days spent overseas over time 39
I.3 Dynamic difference-in-differences results: Days spent overseas 41
I.4 Cumulative number of children over time 43
I.5 Dynamic difference-in-differences results: Cumulative number of children 43
I.6 Employment rate over time, % 44
I.7 Dynamic difference-in-differences results: Employment rate 45
I.8 Annual earnings over time (2006Q1 NZ$) 46
I.9 Dynamic difference-in-differences results: Earnings 46
I.10 Accident claims over time 47
I.11 Dynamic difference-in-differences results: Accident claims 48
I.12 Hospital admissions over time 48
I.13 Dynamic difference-in-differences results: Hospital admissions 49
I.14 Prescriptions over time 49
I.15 Dynamic difference-in-differences results: Prescriptions 50
I.A.1 Dynamic DiD results by treatment window: Days spent overseas: Mothers 58
I.A.2 Dynamic DiD results by treatment window: Days spent overseas: Fathers 59
I.A.3 Dynamic DiD results by treatment window: Days spent overseas: Children 59
I.A.4 Dynamic DiD results by treatment window: Cumulative no. of children: Mothers 60
I.A.5 Dynamic DiD results by treatment window: Cumulative no. of children: Fathers 60
I.A.6 Dynamic DiD results by treatment window: Employment: Mothers . 61
I.A.7 Dynamic DiD results by treatment window: Employment: Fathers . . 61
I.A.8 Dynamic DiD results by treatment window: Earnings: Mothers 62
I.A.9 Dynamic DiD results by treatment window: Earnings: Fathers 62
I.A.10 Dynamic DiD results by treatment window: Accident claim rate . . . 63
I.A.11 Dynamic DiD results by treatment window: No. of accident claims . . 63
I.A.12 Dynamic DiD results by treatment window: Hospital admission rate . 64
I.A.13 Dynamic DiD results by treatment window: No. of hospital admissions 64
I.A.14 Dynamic DiD results by treatment window: Prescription rate 65
I.A.15 Dynamic DiD results by treatment window: No. of prescriptions . . . 65

II.1 COVID-19 policy response Stringency Index: NZ versus OECD 71
II.2 Share of population fully vaccinated: NZ versus OECD 72
II.3 Cumulative COVID cases and excess mortality: NZ versus OECD . . 91
II.4 Cumulative double-vaccination rate by industry 101
II.5 Tracking workers' labour market outcomes over time 104
II.6 Tracking workers' industry switching rates over time 107
II.7 Triple difference results: Change in outcome variable as percentage of the counterfactual 108
II.8 Gender: Triple difference results: Change in outcome variable as a percentage of the counterfactual 110
II.9 Age: Triple difference results: Change in outcome variable as a percentage of the counterfactual 112
II.10 Ethnicity: Triple difference results: Change in outcome variable as percentage of the counterfactual 113
II.11 NZ born: Triple difference results: Change in outcome as percentage of the counterfactual 114
II.12 NZ born: Triple difference results: Change in outcome as percentage of the counterfactual 115
II.13 Deprivation: Triple difference results: Change in outcome as percentage of the counterfactual 115
II.14 Income: Triple difference results: Change in employment as percentage of the counterfactual 116
II.15 Income: Triple difference results: Change in earnings as percentage of the counterfactual 117
III.1 Annual average of monthly earnings and employment, by year of treatment 142
III.2 Monthly effects of mTBI on earnings and employment 145
III.3 Monthly effects of mTBI on earnings and employment 152
III.4 Monthly effects of mTBI on ACC payments 155

List of Tables

I.1 Rights granted by NZ visa type 15
I.2 Derived versus actual citizenship indicator 26
I.3 Variable descriptions and sources 29
I.4 Covariate balance across treatment and comparison groups: Children's characteristics 36

I.5 Covariate balance across treatment and comparison groups: Parents' characteristics 37
I.6 Two-period difference-in-differences: Parents' outcomes 40
II.1 Key dates for COVID-19 vaccination mandates for health, corrections, and education industries 83
II.2 Descriptive statistics of workers in industries that barely, partially, or heavily faced COVID-19 vaccination mandates 93
II.3 District Health Board compositions of workers in industries that barely, partially, or heavily faced COVID-19 vaccination mandates 95
II.4 Descriptive statistics of workers in health industries that heavily faced COVID-19 vaccination mandates 96
II.5 Descriptive statistics of workers in corrections industries that heavily faced COVID-19 vaccination mandates 98
II.6 Descriptive statistics of workers in affected education industries that heavily faced COVID-19 vaccination mandates 99
II.A.1 Difference-in-differences estimating the effect of COVID-19 mandates on vaccination update 127
II.B.1 Two-period triple difference regressions 128
III.1 Effects of mTBI on earnings and employment 143
III.2 Heterogeneous effects of mTBI on earnings and employment 153
III.3 Effects of mTBI on accident compensation payments 154
III.4 Heterogeneous effects of mTBI on ACC payments 157

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Chapter 1 is sole authored. Chapter 2 is co-authored with Livvy Mitchell and Gail Pacheco (my primary supervisor). Chapter 3 is co-authored with Florian Fouquet, Gail Pacheco (my primary supervisor) and Alice Theadom.

I am the sole author of Chapter 1. I am the principal author of Chapters 2 and 3. For Chapter 2, I came up with the research idea, was the Principal Investigator and main grant proposal writer for the Ministry of Health grant that we received for this research (grant number PROP-20), designed the methodological approach in conjunction with Gail Pacheco, undertook the data analysis in conjunction with Livvy Mitchell and completed the main write-up of the paper (with additional writing contributions from Livvy Mitchell and Gail Pacheco). For Chapter 3, I contributed to the research idea along with Alice Theadom and Gail Pacheco, developed the methodological approach along with Florian Fouquet with advice from Gail Pacheco. I completed the main write-up of the paper, along with contributions by Florian Fouquet, and editing and expert advice from Gail Pacheco and Alice Theadom.

Signature:

Name: Lisa Meehan
Date: 28th March 2024

Primary supervisor

Signature:

Name: Professor Gail Pacheco
Date: 28th March 2024

Acknowledgements: Thank you to my supervisors, Professor Gail Pacheco and Associate Professor Peer Skov. Thanks also to my co-authors - it was a privilege collaborating with you. I would also like to thank Ministry of Health for funding the research project presented in Chapter 1, and Stats NZ for providing access to the data used in this thesis. All errors and omissions remain the sole responsibility of the authors.

Disclaimer: The results in this thesis are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI), which is managed by Stats NZ. For more information about the IDI please visit https://www.stats.govt.nz/integrated data.

The results are based in part on tax data supplied by Inland Revenue to Stats NZ under the Tax Administration Act 1994 for statistical purposes. Any discussion of the 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.

Introduction

This thesis applies quasi-experimental methods to linked administrative data to study questions related to the effect of policies or life events. This thesis demonstrates that using existing data collected for other purposes and linked together at the individual level combined with robust empirical methods can yield useful insights.

All three papers use New Zealand's (NZ) Integrated Data Infrastructure (IDI), which is a large research database managed by Stats NZ. The IDI brings together a host of disparate administrative and survey data and links it together at the individual level. The administrative data comes from a number of mostly government sources, such as the Ministry of Health, Ministry of Education, Inland Revenue, NZ Customs etc. This allows researchers to track individuals' health, education, labour market and numerous other outcomes over time. These data allow for the examination of the effect of a range of policy interventions and life events (such as health or other shocks).

In addition, all three papers apply difference-in-differences (DiD) methodology, particularly dynamic DiD, in order to draw causal inferences. Each paper uses a different take on a dynamic DiD approach. Paper 1 creates treatment and control groups based on a policy cut-off date. Paper 2 uses a triple difference estimator. Paper 3 uses a doubly-robust staggered DiD estimator.

The DiD methodology is one of the most commonly used identification strategies in empirical economics and has been crucial for evaluating the causal effects of a treatment or intervention where randomised controlled trials are not feasible (Currie et al., 2020). DiD has been popular in empirical work for about three decades, with seminal papers in using this methodology including Card (1990) and Card and Krueger (1994). The DiD methodology involves observing some units in time periods before and after they become treated, and is a particularly powerful framework for learning about causal effects. DiD identification strategies are, therefore, included among "natural experiment" methods in economics (Callaway, 2023). DiD typically focus on estimating the Average Treatment Effect on the Treated (ATT). The ATT is equivalent to the average path of outcomes experienced over time by the treated group relative to the average path of outcomes that the treated group would have experienced if they had not participated in the treatment. Thus, the key identification challenge in DiD applications is recovering this "counterfactual" path of outcomes for the treated group. The main identifying assumption for DiD is the parallel trends assumption. That is, in the absence of participating in the treatment, the counterfactual path of outcomes for the treated group is the same as the path of outcomes for the control group. Another assumption which has received more attention recently is treatment effect homogeneity - for example, differences in the timing of when units receive the treatment can lead to heterogeneous treatment effects.

The dominant approach to implementing the DiD methodology has been via the use of two-way fixed-effects (TWFE) regressions. The simple two-treatment group, two-time period version of this regression is:

$$Y_{it} = \theta_t + \sigma_i + \gamma D_{it} + v_{it} \tag{1}$$

where Y_{it} is the outcome of interest, θ_t is a time fixed effect, σ_i is a unit fixed effect, D_{it} is a dummy indicating whether or not unit i participated in the treatment in time period t and v_{it} are idiosyncratic, time-varying unobservables. If the parallel trends and treatment effect homogeneity assumptions are met, γ is the causal effect of participating in the treatment. This approach is taken in Paper 1 of this thesis.

In some instances when the parallel trend assumption is violated, a triple difference (DDD) approach may be useful. Like a standard DiD approach, the DDD strategy requires the parallel trend assumption to hold. However, the parallel trend in DDD is on a differential between two categories. Such a DDD strategy is employed in Paper 2, in examining the impact of COVID-19 vaccine mandates on the labour market outcomes of unvaccinated healthcare workers. One approach to estimating the effect would be to compare the outcomes of unvaccinated and vaccinated healthcare workers. This approach would have the advantage that both the treatment and comparison groups are in the same industry, and thus faced the same set of industry labour market conditions. However, there may be spillover effects since, for example, unvaccinated workers leaving the industry may increase health workforce skill shortages, thus improving the bargaining power and labour market outcomes of vaccinated healthcare workers. Alternatively, the exit of unvaccinated workers from the health industry could have potentially increased the pressure on vaccinated healthcare workers and expedited their exit from the health industry. An alternative comparison of unvaccinated healthcare workers and unvaccinated workers in industries that did not face vaccine mandates may also be problematic if the conditions of the two industries diverged for reasons other than the vaccine mandates. This is a realistic concern given that the COVID-19 pandemic itself may have impacted the health industry differently to other industries. For example, there may have been increased demand for healthcare workers, particularly relatively to other industries, many of which initially saw a reduction in demand due to COVID-19 policy measures such as lockdowns. Therefore, while each of these two DiD options may be problematic if used in isolation, a DDD approach which combines these two comparisons has the potential to yield unbiased estimates (Olden & Møen, 2022).

Another concern raised in recent DiD literature is that the standard TWFE approach is generally not robust to treatment effect heterogeneity (Borusyak et al., 2024; de Chaisemartin & D'Haultfœuille, 2020; Goodman-Bacon, 2021). A common situation where this may arise is if there are multiple time periods and variation in treatment timing. To address this issue, Callaway and Sant'Anna (2021) suggest an approach that estimates cohort-time-specific treatment effects, where a cohort is the group of units that is treated in the same period. With many cohorts and time periods, this approach results in many treatment effects, so the group-time average treatment effects, $ATT(g, t)$, can be aggregated in event-study parameters that provide average treatment effects at different periods after treatment, or an average treatment effect for the entire period. Paper 3 applies Callaway and Sant'Anna (2021)'s method to examine the effect of mild traumatic brain injury on labour market outcomes.

Turning our attention from methodological considerations to brief discussions of each paper, Paper 1 examines the effect of NZ's removal of birthright citizenship on the outcomes of migrant parents and their children. Attempts to analyse whether citizenship improves the integration of immigrants into the host country and their children's outcomes are complicated by potential self-selection: those who are more motivated to integrate are also more likely to naturalise. In order to make causal inferences, this paper uses NZ's removal of birthright citizenship as a natural experiment. The analysis is facilitated by the linked administrative data as it allows a number of relevant outcomes to be measured over time. This includes outmigration (as measured by NZ Customs border movement data), parents' fertility (as measured by Department of Internal Affairs birth registrations), parents' labour market outcomes (as measured by Inland Revenue tax data) and children's health outcomes (including injuries measured by Accident Compensation Corporation claims data, hospital admissions measured by Ministry of Health public and private hospital discharges data, and prescription information, as measured by Ministry of Health's pharmaceutical data). Whereas previous research in this area has used cross-sectional data and undertaken cohort analysis, the IDI provides the ability to track individuals over time. A DiD approach is taken whereby children of migrant parents affected by the policy (as identified by Ministry of Business, Innovation and Employment visa data) who were born just after the removal of birthright citizenship are compared with counterpart children who were born just before its removal. In contrast to existing international evidence, which is largely limited to Germany's introduction of birthright citizenship, I find that the removal of birthright citizenship had no effect on any of the parental and child outcomes examined.

Paper 2 also applies DiD methodology to data from the IDI. It examines the effect of workforce vaccine mandates on vaccine uptake and healthcare workers' labour market outcomes. As part of its COVID-19 policy response, the NZ government implemented vaccination mandates as a condition of ongoing employment for certain workers. This paper examines the effect of these mandates on vaccination uptake among mandated healthcare, education and corrections workers and on healthcare workers' labour market outcomes. It uses Inland Revenue employment data to identify mandated workers, and to track their employment and earnings outcomes. It compares vaccinated and unvaccinated workers, which is facilitated by the availability of a comprehensive national vaccination register. It finds that in the context of already-high vaccination rates, workforce vacccine mandates provided limited benefit in terms of increasing vaccination rates among mandated workers. Moreover, they negatively impacted healthcare workers' labour market outcomes, which may have exacerbated existing health workforce skills shortages.

Like the first two papers, Paper 3 applies DiD methodology to data from the IDI. It examines the effects of suffering a concussion on subsequent labour market outcomes. It uses Accident Compensation Corporation data to identify medically-diagnosed concussions. It uses Inland Revenue tax data to track individuals' employment and earnings outcomes. To overcome potential endogeneity issues, it creates a comparison group of those who suffer a concussion but at a later date, and employs a doubly-robust staggered DiD method. It finds that suffering a concussion has negative labour market effects. While the literature on health shocks generally finds that minor injuries have only short-term labour market impacts, it finds these effects do not dissipate over time, but rather grow.

Thus, this thesis demonstrates the utility of applying quasi-experimental methods to big data. While randomised control trials are the "gold standard" in terms of establishing causal effects, these are not always possible for practical and/or ethical reasons. However, the use of quasi-experimental methods to administrative data can provide a useful alternative in some situations.

References

Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*.

Callaway, B. (2023). Difference-in-differences for policy evaluation. In *Handbook of labor, human resources and population economics* (pp. 1–61). Springer Nature.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Card, D. (1990). The impact of the Mariel boatlift on the Miami labor market. *ILR Review*, *43*(2), 245–257.

Card, D., & Krueger, A. B. (1994). Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*, *84*(4), 772–793.

Currie, J., Kleven, H., & Zwiers, E. (2020). Technology and big data are dhanging economics: Mining text to track methods. *AEA Papers and Proceedings*, *110*, 42–48.

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–2996.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

Olden, A., & Møen, J. (2022). The triple difference estimator. *The Econometrics Journal*, *25*(3), 531–553.

Paper I

Does citizenship improve migrant parents' integration and children's outcomes?: Evidence from a natural experiment

Author: Lisa Meehan

Prelude

This paper applies quasi-experimental methods to linked administrative data. Specifically, it uses difference-in-differences (DiD) methodology applied to data from Stats NZ's Integrated Data Infrastructure. This prelude outlines the data and methods used.

In order to identify the population of interest, this paper uses Department of Internal Affairs birth records and Ministry of Business, Innovation and Employment visa data. It uses birth records to identify children born just before and just after birthright citizenship was removed. It uses visa data to identify children who were affected by the policy change - i.e. those with parents who were both on temporary visas at the time of their birth.

Outcomes are measured using a variety of IDI sources. Outmigration of parents and children are measured using NZ Customs border movements data. Parents' fertility outcomes are measured using Department of Internal Affairs birth registrations. Parents' labour market outcomes are measured using Inland Revenue tax data. Children's injury outcomes are measured using Accident Compensation Corporation claims data. Children's hospital admissions are measured using Ministry of Health public and private hospital discharges data. Children's prescription outcomes are measured using Ministry of Health's pharmaceuticals data.

Covariate information, such as ethnicity, child's gender, parents' age, deprivation status and so forth are taken from various IDI sources, such as the Stats NZ's personal details and address notifications tables, both of which collate information from across different IDI sources.

Dynamic DiD estimation is applied to these data. The treatment group consists of children affected by the policy who were born in the few months after the policy change and were, therefore, not entitled to birthright citizenship. The comparison group consists of a counterpart group of children born in the few months before the policy change.

Abstract

Does obtaining citizenship improve the integration of immigrants into the host country's society? Does it improve the outcomes of immigrants' children? Attempts to answer these questions are thwarted by potential self-selection: those who are more motivated to integrate are also more likely to naturalise. In order to make causal inferences, this paper exploits a natural experiment of New Zealand's removal of birthright citizenship. It finds that this policy change had no effect on family outmigration behaviour, parents' fertility and labour market outcomes, nor on children's health outcomes. These results contrast to existing evidence on Germany's introduction of birthright citizenship, which was found to have positive effects on parents' and children's integration outcomes. These different findings may reflect differences in the two countries' broader immigration policy settings, and the resulting differences in the characteristics of the countries' migrants. These differences also highlight that the New Zealand case is of potentially more relevance to Canada and the US, the only two western countries that retain unrestricted birthright citizenship, given greater similarities in these countries' immigration policy settings.

JEL: F22; J15; K37

Keywords: birthright citizenship; integration; outmigration; fertility; employment; earnings; health

1 Introduction

Does obtaining citizenship improve the integration of immigrants into the host country's society? Does it improve the outcomes of immigrants' children? Information on how the legal status of immigrants impacts their outcomes could inform the public debate and government policies in this area. However, empirical assessments of whether obtaining citizenship improves the outcomes of migrant families are inhibited by potential self-selection. Those who are more motivated to remain in the host country and integrate into its society are also more likely to naturalise, potentially leading to biased estimates of the impact of obtaining citizenship on any measure of integration. Thus, despite ongoing debates about citizenship policy, there is little evidence that isolates the causal effect of naturalisation.

Therefore, this paper uses a policy change in New Zealand (NZ) that removed birthright citizenship in 2006 to establish a causal relationship between citizenship status and the outcomes of migrant families. I exploit the differential treatment of those born in NZ just before the policy change, who were entitled to NZ citizenship regardless of their parents' legal status, and those born just after the policy change, who were only entitled to NZ citizenship if at least one of their parents had the legal right to reside in NZ indefinitely. Thus, this policy change provides a source of variation in the citizenship status of immigrant children that is independent of their parents' willingness to remain in NZ and integrate into its society.

The public debate for removing the automatic right to citizenship at birth often raises concerns about so-called 'birth tourism' and the use of 'anchor children' to exert parental rights to stay in the host country. For example, the removal of birthright citizenship in NZ appears to have been at least partly in reaction to some high-profile cases of parents who had overstayed their visas disputing deportation orders on the grounds that their NZ-born child/ren were NZ citizens (Geiringer, 2008; Sawyer, 2013).

Canada and the United States are the only two western countries that retain unrestricted birthright citizenship, and the public and political discourse is still active and contentious in these countries. Similar concerns about birth tourism and anchor children in Canada have prompted considerations to remove birthright citizenship in recent years,[1] and during his time as US President, Donald Trump expressed a desire to remove US birthright citizenship on several occasions.[2]

These debates highlight the potential costs of birthright citizenship. It may allow

[1] In Canada, unsuccessful parliamentary petitions to remove automatic citizenship entitlements to children born to two non-resident parents were made in 2012 and 2018.

[2] For example, in 2018, Trump said that he could end birthright citizenship with an executive order. However, the general consensus is that this is not the case as the right stems from the 14th amendment of the Constitution which states that "all persons born or naturalized in the United States, are subject to the jurisdiction therefore, are citizens of the United States".

people to gain the benefits of citizenship without the commensurate obligations. For example, people on tourist visas coming to the host country specifically to give birth with no intention of remaining in the country but to simply grant their children the right to live in the host country at some future point. Another case is immigrants who would not otherwise meet the criteria for residency being able to avoid deportation due to their child/ren's citizenship status, which appears to have been the main concern in NZ.

However, there appears to be less popular discourse about the potential costs of removing birthright citizenship in terms of reducing the incentives of migrants to integrate. Since returns to education and health investments are higher for citizen than non-citizen children, uncertainty about their ability to meet, and their child to meet, naturalisation requirements may reduce the incentives of migrant parents to invest in integration, leading to poorer labour market outcomes for themselves, and poorer education and health outcomes for their children.

This paper contributes to the limited causal evidence on the effects of migrants' legal status, particularly citizenship status and birthright citizenship. Indeed, almost all of the causal evidence on birthright citizenship relates to Germany's introduction of birthright citizenship in 2000. This existing research highlights that birthright citizenship reduces the likelihood of the parents of affected children leaving the country (Sajons, 2016), lowers parental fertility (Avitabile et al., 2014) and improves parents' social and cultural integration (Avitabile et al., 2013). For affected children, the research for Germany finds birthright citizenship improves health, educational and social and cultural integration outcomes (Avitabile et al., 2014; Felfe et al., 2020), although the effects are stronger among immigrant boys than girls (Dahl et al., 2022; Felfe et al., 2021).

However, it is questionable whether the German case is more broadly applicable. In particular, the NZ case is likely to be more relevant to the Canadian and US debates about birthright citizenship. Like Canada and the US, NZ is an anglophone country with high immigration rates. The overall immigration regimes of the countries are also similar in terms of having selective immigration policies and a low proportion of non-economic immigration (Behr & Fugger, 2020). Moreover, the specifics of NZ's birthright citizenship were more similar to the current policies of Canada and the US than Germany's birthright citizenship policy. Like the pre-2006 NZ system, children born in Canada or the US are automatically entitled to citizenship regardless of the legal status of their parents, with the only exception being the children of foreign diplomats. Germany's birthright citizenship is more restrictive, applying only to children who have at least one parent who has lived in Germany legally for eight or more years.

In addition, NZ's linked administrative data allows for improvements in the measurement of migrant outcomes. While previous research has relied on cohort-level analysis using repeated cross-sectional data, this paper undertakes individual-level analysis using population-wide longitudinal data. This allows, for example, actual outmigration to be observed rather than having to resort to a proxy measure as in Sajons (2016).

In contrast to the research on Germany, I find that the removal of birthright citizenship did not impact family outmigration patterns. It also did not effect parental labour market outcomes nor fertility. It also did not impact on children's health outcomes. These differences in findings compared with Germany may be due to the nature of NZ's migration system. Virtually all NZ migrants have a pathway to residency and the benefits of citizenship over residency are more limited than in Germany. This highlights that the effect of citizenship depends on the policy context.

In addition, the broader context may matter. For example, the fact that NZ has high rates of immigration and a large foreign-born population may mean migrants integrate relatively easily compared with Germany, resulting in less observed positive effects from the granting of citizenship.

Moreover, reflecting its selective immigration policy, the composition of NZ's immigrants is different than Germany's. For instance, NZ's foreign-born population has higher average education levels than its native-born population, whereas Germany's foreign-born population has lower average education levels. The benefits of citizenship are likely to be lower for more highly-skilled migrants, which may also explain why birthright citizenship has a positive effect in Germany but not in NZ.

2 Policy context

This section provides an overview of NZ's immigration policy generally, then provides more information on the 2006 removal of birthright citizenship and finally discusses the benefits of NZ citizenship.

2.1 Background on NZ's immigration settings

Following Canada's lead, NZ implemented a points-based immigration system in the 1990s which emphasised human capital attributes such as education and work experience. Over time, both countries have modified this approach to include more targeted selection based on labour market demand for specific skills and the attraction of international students (Akbari & MacDonald, 2014).

Migrants generally enter NZ on temporary visas initially, such as a work, student or visitor visas. Those who stay transition over time onto resident-class visas as they become eligible. However, some people can enter on resident-class visas immediately, such as refugees and people working in certain professions where there are acute skills shortages. The time it takes for those who are initially on temporary visas to be eligible for residency varies depending on initial visa type, work history, and other factors. However, in almost all cases, migrants have a pathway to residency if they meet certain criteria (relating to employment, health status etc.). That is, there are few genuine temporary or guest workers who do not have a pathway to residency.[3] Those on resident visas can then transition to permanent residency status, and finally, citizenship.[4]

Overall, NZ's approach to immigration is, therefore, similar to other anglophone OECD countries. Like Canada and the US, NZ has a restrictive immigration regime based on 1. a selective immigrant policy and 2. a low proportion of non-economic immigration (Behr & Fugger, 2020). The naturalisation pathway is also very similar in NZ, Canada and the US. All three countries generally require individuals to have been residents in the country for at least five years, been physically present in the country for a minimum amount of time during those five years, intend to continue living in the country, as well as meet other tax filing, good character, language and/or country knowledge criteria. In contrast, the German citizenship pathway is longer and more costly. To be eligible for naturalisation, a person generally has to have lived legally in Germany for at least eight years. In addition, one particular cost of German naturalisation that does not apply to becoming a NZ, Canadian or US citizen is the need to relinquish the citizenship of any other country since Germany does not generally allow naturalised citizens to hold dual citizenship. Relinquishing home-country citizenship involves losing access to practical benefits, such as the ability to move in/out of their home country freely, access to health and social services and so forth, and may also have large costs in terms of psychological and cultural attachments. In contrast, those who become naturalised NZ, Canadian and US citizens can retain their existing passports (provided they were issued by a country that allows dual citizenship). The higher costs of naturalisation in Germany may mean that the benefits of a birthright citizenship regime are greater in Germany than NZ, Canada or the US.

[3]The main exemption is recognised seasonal employer (RSE) workers, who are recruited only from eligible Pacific countries to work in agriculture and horticulture on a temporary basis.

[4]Throughout this paper, the current (post-2009) terms of 'resident visa' and 'permanent resident visa' are used for clarity. However, at the time of the birthright citizenship policy change in 2006, these were officially known as 'non-indefinite returning resident visas' and 'indefinite returning resident visas'.

2.2 Removal of birthright citizenship

Until January 2006, virtually all children born in NZ were automatically granted NZ citizenship. The only exemption was children of foreign diplomats. Thus, until 2006, NZ had a birthright citizenship policy that was in line with the current policies of Canada and the US.

The Citizenship Amendment Act 2005 removed birthright citizenship. Children born from 1 January 2006 onwards were only entitled to NZ citizenship if at least one of their parents had the right to reside indefinitely in NZ. In practice, this means that at least one of their parents is 1. a NZ citizen, 2. holds a resident-class visa (either residency or permanent residency), or 3. is a citizen of Australia, Tokelau, Niue or the Cook Islands (as citizens of these countries have the right to reside in NZ indefinitely without the need for a visa). Unlike the equivalent Australian and UK legislation, there is no provision for NZ-born children who are not entitled to citizenship to gain it after a long period of residency. For example, the Australia legislation allows those who were born in Australia but did not qualify for citizenship at birth to become citizens if they live in Australia for the first ten years of their life. Instead, the legal status of non-citizenship children is generally related to the legal status of their parents. For example, if a child is born in NZ to parents on temporary visas, the child will also be granted a temporary visa. As noted, virtually all documented migrants have a pathway to residency in NZ, and when at least one parent gains residency, their child will also gain residency.

The situation is less clear-cut for children born to undocumented migrants as they do not have a clear route to gaining legal status. One of the few routes available to them is via ministerial discretion as NZ's immigration law grants the Minister of Immigration absolute discretion to grant a visa of any type to someone unlawfully in NZ who would otherwise be liable for deportation. Another route is via a widespread amnesty scheme, however, the last time this option was provided was in 2000.

It should also be noted that there are relatively few undocumented migrants in NZ. Indeed, a key difference between NZ and most other countries is that almost all migrants enter NZ legally. As an island nation with no land borders that cannot easily be reached by sea on small or make-shift vessels, illegal border crossings are typically limited to cases of individuals crossing the border using forged travel documents. However, migrants can arrive in the country legally then fail to leave the country when their visas expire. Indeed, the term 'overstayers' is synonymous with undocumented migrants in NZ. It is estimated that there are about 14,000 overstayers in NZ (Heron & Barrow, 2023) out of a total population of about 5.2 million and a foreign-born population of about 1.5 million (based on Census 2018 population numbers). Thus, overstayers account for less than 0.3% of the total population and about 1% of the foreign-born population.

The share of undocumented migrants in Canada appears to be comparable to the NZ situation. Estimates of the number of undocumented migrants in Canada vary substantially, ranging from 20,000 to 500,000 (CIMM, 2022). These estimates equate to between 0.05% and 1.3% of Canada's total population and between 0.24% and 6% of the foreign-born population (based on Census 2021 population numbers), which is within the same range as the NZ estimates. However, the estimated share of undocumented migrants in the US population is much higher, at just over 11 million (Migration Policy Institute, 2020), they account for 3.3% of the total population and almost a quarter of the foreign-born population (based on US Census Bureau American Community Survey 2019 population numbers).

2.3 The benefits of NZ citizenship

In order to understand the potential effects of birthright citizenship, it is useful to consider the benefits of NZ citizenship. As mentioned, most migrants arrive in NZ on a temporary visa but have a pathway to transition to residency, although a few migrants are eligible for residency immediately. Therefore, I first discuss the benefits of a resident visa over a temporary visa, then the benefits of a permanent-resident visa over a resident visa, and finally the benefits of citizenship over permanent residency. Table I.1 provides a summary.

Table I.1: Rights granted by NZ visa type

Visa status	Work rights	Welfare benefits	Public healthcare	Public education	Indefinite NZ re-entry	Reside in Australia visa free
Visitor	None	No	No	No	No	No
Student	Restricted	No	No	No[1]	No	No
Work <2yr duration	Restricted	No	No	Partial[2]	No	No
Work ≥2yr duration	Restricted	No	Yes	Partial[2]	No	No
Residency	Unrestricted	Yes	Yes	Yes	No	No
Permanent residency	Unrestricted	Yes	Yes	Yes	Yes	No
Citizenship	Unrestricted	Yes	Yes	Yes	Yes	Yes

Notes: 1. With some exceptions, e.g. the child of a student visa holder who is enrolled in a PhD programme is entitled to free primary and secondary education. 2. The children of those on work visas are entitled to free primary and secondary school education, but are considered international students for the purposes of tertiary education.

Residency visas afford a number of advantages over temporary visa (most commonly work, student and visitor visas). First, in terms of work rights, there are no work restrictions for those on resident visas. In contrast, those with visitor visas are not entitled to work and those on student visas are entitled to work only 20 hours a week. Many of those on work visas are also restricted to working for a specific employer, and must go through a bureaucratic approvals process to change employers.

Moreover, employers are generally obliged to first search for NZ residents or citizens to fill vacancies before offering roles to those who hold, or could obtain, a temporary work visa. These restrictions mean that the administrative burden for an employer of hiring a resident is lower than hiring someone who is not. However, compared with many countries, including Germany, there are fewer restrictions on the types of roles that those on temporary visas with work rights can hold. Virtually all private sector roles (except for those in the sex industry) and most public sector roles (with some exceptions such as police and defence force roles as well as any roles requiring high-level security clearance) are open to those on temporary visas with work rights.

In terms of access to public services, unlike those on temporary visas, those on resident visas can access social services, such as welfare benefits. Those on work visas that entitle them to stay in NZ for two or more years can also access publicly-funded healthcare just as residents and citizens can. However, those on visitor, student and short-term work visas cannot access publicly-funded healthcare. It should also be noted that anyone who has an accident in NZ is entitled to medical treatment under the accident compensation scheme, regardless of their legal status. The children of those on work visas are entitled to free primary and secondary schooling, but are considered international students for tertiary education (and therefore, face much higher fees than domestic students and cannot access the government-run student loan scheme). The children of those on other types of temporary visas are generally considered international students for all educational levels. In addition, residents have the right to vote. When birthright citizenship was removed in 2006, there was no restrictions on foreign property ownership, although restrictions were introduced in 2018 meaning those on temporary visas generally cannot buy residential property.

After holding a resident visa for at least two years and demonstrating a commitment to NZ,[5] a migrant can then apply for a permanent resident visa. The only difference between residency and permanent residency is that resident visas have travel conditions that only allow the holder to re-enter NZ as a resident until a certain date, while permanent residency allows indefinite re-entry into NZ.

After residing in NZ on a resident-class visa for at least five years, individuals can apply to become citizens. In terms of citizenship versus permanent residency, there is one only one notable benefit for most people. NZ citizenship grants nearly unfettered access to Australia. NZ citizens are allowed to live and work in Australia indefinitely without the need for a visa provided they do not have serious criminal convictions. This is similar to Germany in the sense that German citizenship provides immigrants with the opportunity to live and work in other EU countries. However,

[5]There are several ways that migrants demonstrate a commitment to NZ, such as, spending the majority of their time in NZ; being NZ tax residents; having a business in NZ etc.

while Germany is a relatively high income country within the EU, Australia provides a higher-income option for NZ citizens, with GDP per capita that was approximately 24% higher than NZ's in 2022 (OECD, 2023b). Moreover, moving to Australia may be a more viable option in practice for a NZ citizenship than moving to another EU country for a German citizen since Australia and NZ are both anglophone countries and are culturally similar.

In virtually all other respects, NZ permanent residency provides the same rights as NZ citizenship.[6] Permanent residency, therefore, effectively provides citizenship in everything but name, which is particularly relevant for those whose home nation does not allow dual citizenship (such as China).

Thus, the benefits of NZ citizenship for migrants who wish to reside in NZ are smaller than the benefits of German citizenship. Holding a German passport has labour market benefits as there are several public sector roles that require one, such as teaching (Avitabile et al., 2014; Sajons, 2016). Moreover, some private sector professions, such as dentistry, medicine, pharmacy, law and architecture, are restricted to EU citizens (Avitabile et al., 2014).

In addition to the legal benefits inferred from citizenship, there may also be perceived benefits. For example, automatic granting of citizenship to their children may be perceived by migrant parents as a sign of goodwill on behalf of the host country, prompting them to be more willing to increase integration efforts. And, likewise, removal of it a sign of ill-will that prompts parents to decrease their integration efforts (Sajons & Clots-Figueras, 2014).

Another reason why German citizenship may have greater benefits than NZ citizenship is due to the different characteristics of the countries' migrants. As mentioned, NZ, along with Australia, Canada, the UK and the US, has a restrictive immigration regime based on 1. a selective immigrant policy and 2. a low proportion of non-economic immigration (Behr & Fugger, 2020). Reflecting this, NZ's foreign-born population tends to have higher education levels than its native-born population. In contrast, Germany's migrants have, on average, lower education levels than the native-born population. For example, 30% of Germany's native-born population aged 25-64 years has a tertiary qualification, compared with 25% of foreign-born adults aged 25-64. A third of NZ's native-born population has a tertiary qualification, compared with 53% of its foreign-born population (OECD, 2023a). It is likely that the benefits of citizenship are greater for more disadvantaged migrants with lower skill levels than for higher skilled, more advantaged ones who have better outside opportunities. In this

[6] Although not applicable to most people, an additional benefit of citizenship is that permanent residency can be revoked if an individual commits a serious crime. In contrast, naturalised citizens can only have their citizenship revoked if they have the citizenship of another country and voluntarily acted against the interests of NZ, or obtained their citizenship through fraud or misrepresentation.

respect, NZ is similar to Canada, where 56% of native-born adults aged 25-64 have a tertiary qualification versus 70% of foreign-born adults. In contrast, despite also having a restrictive immigration policy, a lower share of US foreign-born adults have a tertiary qualification compared with native-born adults (45% versus 51%) (OECD, 2023a).

3 Existing evidence

There are several studies which examine the relationship between gaining citizenship of the host country and migrant outcomes. For example, Bratsberg et al. (2002), Chiswick (1978), DeVoretz and Pivnenko (2005), Fougère and Safi (2009), and Steinhardt (2012) all examine naturalisation and labour market outcomes. However, these studies generally do not estimate causal relationships. Non-random selection into naturalisation means that those who are most motivated to integrate into the host country are also the most likely to be successful and are, therefore, also the most likely to naturalise.

However, some research uses plausibly exogenous variation in migration policies to study the consequences of a change in legal status. For example, Mazzolari (2009) examines the effects of five Latin American countries allowing dual citizenship on the labour market performance of immigrants in the US. Immigrants granted dual nationality rights were more likely to naturalise relative to immigrants from other Latin American countries, and experienced employment and earnings gains.

Another natural experiment is Switzerland's use of secret ballot referendums to decide on naturalisation applications until 2003. Leaflets describing the applicants were sent to local voters who then voted to accept or reject each individual applicant, with immigrants who gained a majority of "accept" votes receiving Swiss citizenship. This design allows the outcomes of migrants who received just over half "accept" votes and therefore received citizenship to be compared with those who received just under half "accept" votes so did not. Hainmueller et al. (2017) finds that naturalisation improves the long-term social integration of immigrants as measured by an integration scale that combines a variety of outcomes, including whether immigrants have plans to permanently stay in Switzerland, are members of a local social club, feel discriminated against, and read Swiss newspapers. The paper finds that these positive effects of naturalisation on social integration are sizable and persist for many years.

Another type of natural experiment involves amnesty programmes for undocumented migrants. However, these generally involve granting undocumented migrants legal status, such as work visas, which may provide them a pathway to citizenship rather than citizenship per se. Moreover, the effect of citizenship on undocumented migrants is not necessarily generalisable to migrants with the legal right to reside in the host country (e.g. those on a work visa). Indeed, it would be expected that the benefits of citizenship would be larger for undocumented migrants. Nonetheless, these studies find positive effects of undocumented immigrants being granted legal status. For example, Devillanova et al. (2018) finds that an amnesty program in Italy improved the employment outcomes of migrants. Several studies examining US amnesty programmes find that granting legal status to undocumented immigrants improves their earnings, but the effects on employment are unclear (e.g. Amuedo-Dorantes et al., 2007; Barcellos, 2010; Borjas & Tienda, 1993; Kaushal, 2006; Kossoudji & Cobb-Clark, 2002; Lozano & Sørensen, 2011; Pan, 2012).

Since examining birthright citizenship considers the effect of children's status on parental outcomes, another relevant strand of literature examines how children affect parental behaviour. For example, studies find that the gender of their child/ren impacts on parents' political choices, with those having daughters more likely to vote liberally (Owsald & Powdthavee, 2010; Washington, 2008). In addition, having a son can reduce the criminal convictions of young fathers relative to having a daughter (Dasgupta et al., 2022; Dustmann & Landersø, 2021), as well as improve educational and labour market outcomes (Dasgupta et al., 2022). Studies have also looked specifically at the impact of children on immigrant parents. Kuziemko (2014) documents that immigrant adults in California are less likely to learn English if there are English-speaking children in their household. Dustmann (2003) finds that children can affect the return migration of their parents, with the presence and number of children being associated with a decreased likelihood of return migration.

Finally, the most relevant literature to the current study examines Germany's introduction of birthright citizenship. This policy change meant that children born from 2000 onwards who had at least one parent who had lived legally in Germany for at least eight years was granted German citizenship at birth. Further, affected children were entitled to hold dual citizenship until the age of 23, at which time, they could choose which citizenship to retain. There was also a retrospective element to the policy change as children born between 1990 and 1999 who had at least one parent who met the eight-year residency criterion at the time of their birth could apply for German citizenship. The remainder of this section discusses the impact of this policy change along with some limited evidence on birthright citizenship from other countries.

3.1 Parental outcomes

Outmigration

It is theoretically ambiguous whether citizenship would increase or decrease the probability of outmigration (Sajons, 2016). Citizenship may entail benefits such as the availability of additional job opportunities, thereby decreasing the probability of outmigration. On the other hand, it lowers the cost of outmigration as it provides an unfettered ability to return to the host country at any point in the future (Sajons, 2016). Moreover, it may also increase outmigration to a third country (rather than to the home country) since, in the case of Germany, it opens up access to other EU countries. Likewise, NZ citizenship provides the ability to live and work in Australia without the need for a visa. Thus, the question of the effect of citizenship on outmigration is an empirical one.

Sajons (2016) appears to be the only paper which examines the impact of birthright citizenship on family outmigration. It uses German microcensus data, which is a representative 1% sample of all households in Germany to examine the effect of the introduction of birthright citizenship on outmigration by exploiting the differential treatment of birth cohorts around the policy enactment date. It finds that granting citizenship to immigrant children reduces the likelihood of outmigration.

However, the main limitation of Sajons (2016) is that family outmigration cannot be directly observed with the repeated cross-sectional data used. Instead, outmigration is estimated at the level of birth-year cohorts based on changes to a cohort's size. For example, outmigration between 2001 and 2006 for each birth-year cohort of children born between 1991 and 2002 is measured as the difference between the cohort size in 2001 and 2006, with an adjustment made for survey non-response. As Sajons (2016) discusses, this means that the estimates are vulnerable to sampling bias and measurement error. Given these issues, Sajons (2016) highlights the importance of further research on the impact of birthright citizenship whenever appropriate microdata becomes available.

Fertility

The introduction of birthright citizenship represents a positive shock to the returns of investing in a child's human capital, which may, therefore, impact on immigrant fertility. Avitabile et al. (2014) finds that Germany's introduction of birthright citizenship reduced immigrant fertility and improved children's outcomes (discussed further below), which is consistent with Becker's 'quality-quantity' fertility tradeoff (Becker & Lewis, 1973).

For their analysis of fertility, Avitabile et al. (2014) also use the German microcensus, and therefore relies on repeated cross-sectional data rather than longitudinal data that follows individuals over time. Moreover, the construction of the treatment and comparison groups opens the possibility that the two groups had different characteristics. Specifically, the treatment group comprised of parents affected by the reforms (i.e. non-citizens where at least one of the parents had lived in Germany for at least eight years), while the comparison groups comprised of 1. households where only one parent is a non-Germany citizen and 2. households where both parents are German citizens as these two groups should not be affected by the policy change. Indeed, Avitabile et al. (2014) finds that the pre-policy-change trends for the groups are not parallel. However, a robustness test using a placebo reform date suggests that the results are not driven by differential trends across the groups.

There is also some evidence on the effects of birthright citizenship on fertility for Ireland and the Dominican Republic. Ireland had birthright citizenship for just a five-year period between 1999 and 2004. Foad (2022) finds that non-British migrants living in Ireland when the policy was implemented increased their fertility in response. Foad (2022) also examines the effect on fertility of migrants arriving under the birthright citizenship regime. A common concern used to promote the end of birthright citizenship in Ireland was that it led to birth tourism. If this was the case, then the fertility rates of those arriving during the period where birthright citizenship existed should be higher than for those who arrived before it was implemented or after it was revoked. However, those arriving during the birthright citizenship period actually had lower fertility rates. This can be explained by a change in the composition of arriving migrants as those who arrived under the birthright citizenship regime also had higher levels of education and employment, which are both characteristics associated with lower fertility rates. Similar to Avitabile et al. (2014), a limitation of the Irish study is that it uses repeated cross-sectional data in the form of the Irish national census rather than longitudinal data that tracks individuals over time.

In 2010, the Dominican Republic changed its constitution to remove birthright citizenship for children born to undocumented immigrants. In 2013, the High Court revoked the citizenship of a large number of individuals by retrospectively applying the 2010 constitutional change. Amuedo-Dorantes et al. (2017) examines the impacts of these changes on Haitian immigrants and their descendants (to whom these policies were likely directed). It finds that, unlike the German case, the policy changes did not affect the fertility of Haitian households. This difference in findings is perhaps unsurprising given that, unlike Germany, the Dominican Republic is a developing country where access to health services and contraception by the Haitian population may be limited. As with the German and Irish studies, the study uses repeated cross-sectional data (in the form of Labour Force Surveys).

Labour market outcomes

Birthright citizenship could improve parental labour market outcomes if, for example, it increases migrant parents' investment in country-specific human capital through improved language proficiency, more frequent contact with natives and therefore better employment networks, and so forth. On the other hand, the effect could go the other way if parents invest more time in bringing up their citizen children and, therefore, invest less in their own labour market outcomes. They may also invest less in their own labour market outcomes if the introduction of birthright citizenship is considered a positive shock to the family's expected lifetime income. If part of the parents' motivation to work hard is to provide a better future for their child, this positive shock may decrease their motivation to increase their earnings.

Using German Microcensus data, Sajons (2019) finds that Germany's introduction of birthright citizenship did not have an impact on the employment rate of affected fathers, but affected mothers spent less time in employment and more time at home in their children's early years. In addition to being cross-sectional data, a disadvantage of the Microcensus is that it provides income rather than earnings information, and thus includes non-earned income such as welfare benefits, investment income etc. Furthermore, the income information is not precise as it is reported in bands rather than in euro amounts.

In terms of the Dominican Republic's removal of birthright citizenship, Amuedo-Dorantes et al. (2017) finds that the amendment increased the propensity of Haitian men to hold an informal sector job.

Social and cultural integration measures

Avitabile et al. (2013) examines whether the German reform had an effect on the integration of the social and cultural integration of immigrant parents. Using data from the German Socio-Economic Panel survey, it finds that the introduction of birthright citizenship increased the probability that foreign-born parents socialised with Germans and read German newspapers.

3.2 Children's outcomes

Avitabile et al. (2014) finds that Germany's introduction of birthright citizenship improved children's outcomes. In particular, the obesity gap between non-citizens' and citizens' children at preschool age dropped significantly for children born just after the introduction of birthright citizenship compared with those born immediately before it. In addition, Avitabile et al. (2014) finds that children affected by the reform had better non-cognitive development outcomes.

In terms of educational outcomes, Felfe et al. (2020) finds that Germany's introduction of birthright citizenship increased immigrant children's participation in non-compulsory preschool education and improved development outcomes measured at the end of the preschool period. It also led to immigrant children progressing faster through primary school and increased the likelihood of them attending the academic track of secondary school. By considering whether the observed changes were related to changes in official recommendations (e.g. paediatricians' recommendations of whether parents should delay their child's primary school start date, or teachers' recommendations of whether to choose the academic secondary school track) or not, the paper concludes that mechanism behind these results is that the introduction of birthright citizenship incentivised immigrant parents to provide their children with more similar educational opportunities to children from native families.

In terms of social and cultural outcomes, Felfe et al. (2021) examines whether Germany's introduction of birthright citizenship increased the propensity of immigrant youth to cooperate with their native-born peers. It finds that the policy led to increased cooperativeness towards natives among male immigrants, but not female immigrants.

Dahl et al. (2022) also finds gender differences, with the introduction of birthright citizenship in Germany improving the wellbeing, integration and schooling outcomes of immigrant boys. In contrast, it lowered measures of life satisfaction and self-esteem among immigrant girls, particularly Muslim girls. Moreover, Muslim girls granted birthright citizenship were less integrated into Germany society, as measured by self-reported social isolation and the likelihood of self-identifying as Germany. In terms of possible mechanisms, they find that immigrant Muslim parents invested less in their daughters' schooling and that these daughters have lower school achievement results if they were born after the reform. Parents were also less likely to speak German with their daughters. Their results suggest that immigrant girls were being encouraged by their parents to conform to a role within traditional culture, whereas boys were allowed to take advantage of the opportunities that come with citizenship.

For evidence outside of Germany, the Dominican Republic's removal of birthright citizenship led to a large reduction in school attendance by Haitian children and young people, with the share of those citing documentation barriers as the main reason for not being able to attend school increasing (Amuedo-Dorantes et al., 2017).

4 Data

This paper uses data from the Integrated Data Infrastructure, which is a large research database managed by Stats NZ. This rich database includes population-level administrative data from a range of sources. Records are anonymised and each person is assigned a unique identifier, allowing for individuals' data to be linked across data sources. Some of the main sources of data used are Department of Internal Affairs (DIA) birth records, NZ Customs border movements data, Ministry of Business, Innovation and Employment (MBIE) visa data, Inland Revenue (IR) income tax records, Accident Compensation Corporation (ACC) injury claims data, and Ministry of Health (MoH) hospital admissions and prescriptions data.

4.1 Population of interest

For the main specifications, I use DIA birth records to identify all births in NZ between July 2005 and June 2006 (i.e. six months before/after the policy change). Each child's birth record also includes the identifiers of their parents, allowing children and parents to be linked and both children's and parents' information to be gleaned from the IDI.

In terms of exclusions, in the case of multiple births, I include only one randomly chosen sibling. I also exclude children and their parents from the analysis if the child died before their 10$^{\text{th}}$ birthday since I follow parents' and children's outcomes for 10 years after a child's birth. To avoid a mother being in both the control and treatment groups, I exclude mothers who gave birth twice within a treatment window. These restrictions result in less than 0.7% of the sample being excluded.

To identify affected children, I use a variety of datasources to determine whether or not a child was eligible for birthright citizenship under the rules which applied from 2006 onwards. First, children are ineligible if the last visa both parents held before the child's birth was a temporary one, based on MBIE visa data. Second, children were identified as eligible if at least one parent was born in NZ (since all of these parents would have been born at a time when NZ granted automatic citizenship at birth), crossed the border using an NZ, Australian, Tokeluan or Niuean passport any time before the child's birth, or was granted a visa that allowed them to reside in NZ indefinitely before the child's birth. This information is based on DIA birth records, NZ Customs border movement data and MBIE visa data respectively. Note that there is no direct data on whether an individual is an NZ citizen in the IDI. Also note that information on border crossings using a foreign passport is not used to identify ineligible children due to the possibility that the parent actually held an NZ, Australian, Tokeluan or Niuean passport but used another passport (due to dual/multiple citizenships) when crossing the border.

These steps produce the following numbers. Between July 2005 and June 2006, over 59,000 children were born in NZ. Of these, only 858 (1.4%) had parents who were both on temporary visas, while over 57,000 had at least one parent who was entitled to reside in NZ indefinitely. Unfortunately, it was not possible to identify the citizenship eligibility status of a small share of children (just over 1,000 or approximately 1.8%). This could occur for several reasons. One example of a case where a child was eligible for birthright citizenship but it is not possible to see this in the data is when a child's parents were granted resident-class visas before 1997, the earliest that visa data are available in the IDI, and never subsequently crossed the border using an NZ, Australian, Tokeluan or Niuean passport. This could occur because they never left the country after border crossing data begins in 1997, or because they never gained NZ citizenship and, therefore, never held an NZ passport. An example of a case where a child is ineligible for birthright citizenship but it is not possible to see this via the steps described is if a child's parents entered NZ on temporary visas before 1997, or entered NZ at any time without a visa because they held a passport of a country NZ has a visa-waiver arrangement, and then failed to obtain a resident-class visa and remained in the country as 'overstayers'. However, overstayers who arrived on a temporary visa from 1997 onwards would be classified as being ineligible for birthright citizenship as the last visa they held before the birth of their child would be a temporary one.

Since a citizenship eligibility flag is available on the birth records of children born from 2006 onwards, when birthright citizenship ceased, I compare the derived eligibility variable described above with this birth record flag on citizenship eligibility for children in the population of interest who were born in 2006 (Table I.2). This comparison shows that the derived measure of birthright citizenship misclassifies just 0.4% (1.1%) of children who were eligible (ineligible) for birthright citizenship as being ineligible (eligible). Likewise, few children had unknown eligibility but were actually eligible (1.3%). However, a much higher share have unknown eligibility but were actually ineligible (20.9%). In essence, the analysis will be limited to those who could be identified as having both parents on a temporary visa at the time of their birth. While this excludes some children who were ineligible for citizenship but could not be classified as such, very few people who were eligible (ineligible) for citizenship are misclassified as being ineligible (eligible).

One limitation is that this research only examines those on temporary visas at the time of their child's birth, which is defined as those whose last granted visa before the child's birth was a temporary visa. Thus, it may capture some overstayers who stayed in the country beyond the expiration of their visitor, study or work visa. But, it would not capture overstayers who arrived in NZ from a visa-exempt country but then failed to leave once the exemption period was over. Given birthright citizenship would potentially be of greater value to a migrant who lacks legal status than a migrant on a temporary visa who will become eligible for a resident's visa at point in the future, the analysis may underestimate the true effect of birthright citizenship on migrants' outcomes. However, as previously mentioned, there are relatively few undocumented migrants in NZ and, therefore, any underestimate would likely be small.

Table I.2: Derived versus actual citizenship indicator

		DIA citizenship indicator		
		Eligible	Ineligible	Total
Derived citizenship indicator	Eligible	27,855	6	27,861
		(98.7%)	(1.1%)	(96.9%)
	Ineligible	12	417	429
		(0.04%)	(78.5%)	(1.5%)
	Unknown	369	111	477
		(1.3%)	(20.9%)	(1.7%)
	Total	28,233	531	28,767
		(100%)	(100%)	(100%)

Notes: This table presents the derived versus actual citizenship indicator for children born in January-June 2006 (i.e. children in the treatment group with a six-month treatment window). The actual citizenship indicator is only available on birth registrations after the removal of birthright citizenship after January 2006.

4.2 Treatment and control group definitions

The main specifications use a six-month treatment window whereby children (and parents of children) born between July 2005 and December 2005 form the control group, and those born between January 2006 and June 2006 form the treatment group. That is, those born six months on either side of the 1 January policy change. For reasons discussed in Section 5.1, for robustness, I also consider twelve, nine and three month treatment windows, and a six-month doughnut window which excludes the month immediately before and after the policy change (i.e. December 2005 and January 2006).

4.3 Outcome measures

I use several outcome measures: outmigration, parents' fertility, parents' labour market outcomes and children's health outcomes (Table I.3, Panel A).

To identify outmigration, I use Customs border movements data. This is a comprehensive database of all border movements into and out of NZ. As noted, given that NZ is an island nation with surrounding ocean that would be extremely difficult to traverse in a make-shift craft, unrecorded and/or illegal border crossings are relatively difficult and therefore less of an issue than in other countries. The border movements data allows the outmigration of mothers, fathers and children to be tracked separately. As far as I am aware, this is the only study in this area that uses actual border movement data. For example, Sajons and Clots-Figueras (2014) estimates outmigration based on the change in cohort size from repeated cross-sectional data. I measure how many days in each year mothers, fathers and children were overseas.

To measure parents' fertility, I use DIA birth records. I measure the cumulative number of children each parent had.

I use Inland Revenue (IR) tax data on wage and salary earnings to measure parents' employment and earnings outcomes. A parent is defined as employed if they have any wage/salary earnings in a given year. Earnings are deflated to 2006Q1 dollars using the consumer price index.

I examine three measures of children's health outcomes: injuries, hospital admissions and prescriptions. I use ACC injury claims data to measure children's injury outcomes. I use an indicator for whether or not they had any accepted injury claim in each of the 10 years following their birth, as well as the number of claims in each year.

ACC is a universal accident compensation system that covers all injuries that occur in NZ. Its universal nature and the fact that injury claims are submitted by treatment providers (to recover their treatment costs) rather than by patients should mitigate issues of underreporting and misreporting (Poland, 2018).

I also use MoH public and private hospital admissions data. I use an indicator of whether or not the child had at least one hospital admission within a given year, as well as a variable of the number of admissions in each year. Similarly, I use MoH prescription data to measure whether the child had any prescriptions within a given year, as well as a variable of the number of prescriptions in each year.

While it would be ideal to use pure health-outcome measures, all three health measures used are a combination of health-outcome and health-access measures. For injuries, ACC covers all accidents occurring in NZ and while some ACC treatment providers charge a patient co-payment, these co-payments are typically zero for children. Thus, cost is unlikely to be a barrier to access regardless of the child's legal status. However, there may still be differences in access due to other factors. For example, differences in parents' awareness of their child's ACC entitlements (although this is somewhat mitigated by the fact that treatment providers rather than the child's parents lodge ACC claims) and differences in the ability to access medical services (e.g. distance to the nearest provider).

For hospital admissions, the vast majority of admissions relate to public hospitals. Children in the control group are all citizens, and therefore entitled to free public hospital treatment. However, some children in the treatment group may not be entitled to free public hospital healthcare because of their legal status. Those whose parents were on work visas of two or more years' duration would be entitled to free public hospital care, but it would be user-pays for those with parents on other types of temporary visas, such as visitor, short-term work and student visas (although those on short-term work and student visas are typically required to hold private medical insurance as a condition of their visas). Similarly for prescriptions, during most of the time period, for those eligible for subsidised healthcare (i.e. everyone in the comparison group and some of those in the treatment group) most prescription medications were free for those under six years old and a flat fee of NZ$5 per prescription applied for those aged six and over.

Note that children's educational outcomes were not analysed since the first measured outcomes in the IDI are senior high school qualification results. At the time the analysis was undertaken, these were available up to and including the 2021 school year (the NZ school year runs from February to November). While a child born at the beginning of the analysis window (July 2005) would have been old enough to attempt the first set of high school qualifications (National Certificate in Educational Achievement (NCEA) Level 1) in 2021, some born at the end of the analysis window (June 2006) would not have yet attempted NCEA Level 1. Therefore, a comparison of educational achievement by the year 2021 would likely be misleading. However, this analysis will be possible as more years of data become available.

4.4 Other control variables

Other control variables are sourced from various IDI databases (Table I.3, Panel B). Age, gender and ethnicity was sourced from the derived personal details tables within the IDI. For those with multiple ethnicities, prioritised ethnicity was used, with the order of priorisation being: Māori, Pacific peoples, Asian, MELAA (Middle Eastern, Latin American or African) or other, and European.

For children who were ineligible for citizenship at birth, the nationality of the parents was sourced from Customs data, based on the passport used by the parent the last time they crossed the border before their child was born.

The region the child was living in at the time of their birth was based on address notifications information. In addition, the meshblock (the smallest geographic unit defined by Stats NZ, which is equivalent to a city block in urban areas) the child lived in at the time of birth was used to determine their deprivation index (see Salmond et al., 2007).

Table I.3: Variable descriptions and sources

Variable	Description	IDI source
	A. Outcome variables	
Outmigration	Number of days spent overseas	Customs border movements
Fertility	Cumulative number of children	DIA birth records
Employment	Positive wages/salary earnings	Inland Revenue
Earnings	Earnings from wages/salary (NZ$ 2006Q1)	Inland Revenue
Injuries	Child's number of accident claims	Accident Compensation Corporation
Hospital admissions	Child's number of hospital admissions	Ministry of Health
Prescriptions	Child's number of prescriptions	Ministry of Health
	B. Control variables	
Age	Parent's age when child born	Personal details
Ethnicity	Prioritised ethnicity	Personal details
Gender	Child's gender	Personal details
Nationality	Region of nationality based on parent's passport	Customs border movements
Residential region	Region the child lived in at time of birth	Address notifications
Deprivation	Deprivation index of meshblock child lived in at time of birth	Address notifications

5 Methodology

I employ dynamic difference-in-differences (DiDs) regressions of the following form to compare the treatment and control groups:

$$Y_{it} = \alpha + \beta Treat_i + \sum_{e\neq -4, e=-7}^{e=10} \delta_e.M_e + \sum_{e\neq -4, e=-7}^{e=10} \gamma_e.M_e.Treat_i + \theta X_{it} + \epsilon_{it} \qquad (2)$$

where Y_{it} is the outcome of interest, such as days spent overseas. $Treat$ is equal to 1 if the individual is in the treatment group (i.e. affected migrant children born between January and June 2006) and 0 if the individual is in the control group (i.e. affected children born between July and December 2005). M_e is a time dummy; for example, $M_{e=5} = 1$ if it is the fifth year after the child's birth and 0 otherwise. e ranges from 7 years before the child's birth to 10 years after the child's birth for parental outcomes (and from the year of the child's birth to 10 years after for children's outcomes). The base time period is $M_{e=-4}$ for parental outcomes (and $M_{e=0}$ for children's outcomes). X_{it} is a vector of control variables (see Panel B of Table I.3 for a list of control variables). The coefficient of interest is γ_e.

For robustness, I also estimate two-period DiDs:

$$Y_{it} = \alpha + \beta Treat_i + \delta Post_{it} + \gamma Treat_i \times Post_{it} + \theta X_{it} + \epsilon_{it} \qquad (3)$$

where $Post$ is equal to 1 if the time period is after the child's birth and 0 if it is before.

5.1 Validity of identification strategy

This sub-section discusses several potential threats to the identification strategy: adjustments to fertility behaviour, concurrent policy changes, changes to the timing of parents' residency visa applications, selective outmigration, changes in the composition of migrants and month-of-birth effects.

Fertility behaviour

The citizenship reform may have resulted in immigrant parents adjusting their fertility behaviour in at least two ways. First, parents may have changed the timing of birth to enable their child to benefit from the policy. However, this possibility seems unlikely. There were only 36.5 weeks between the legislation passing and it coming into effect, which is less than the period of a full-term pregnancy. While there was debate and media coverage about the possibility of the policy change before the Act was passed, there was uncertainty at the time as to whether the law change would pass since it required cross-party support under NZ's mixed-member proportional representation system given no one party held an outright majority. Furthermore, parents would have wanted the child to be born earlier in order for them to benefit from the more favourable birthright policy, which is more difficult to achieve than delaying conception.

There is still, however, a possibility that mothers with due dates very close to the policy-change cut-off could have chosen to give birth slightly earlier via induction or elective caesarean. However, the ability to do this is limited as it would be difficult to find a medical professional who would implement such a birth intervention for non-medical reasons, and also given there are restrictions in the choice of dates for inductions and elective caesareans in the public health system. It could, however, be plausible that women with private maternity care (who tend to be higher-income women) who were due very close to the policy-change date could have opted to be induced or have an elective c-section just before this date. Another factor which makes this even more unlikely, however, is that the policy change came into effect on 1 January. A low share of births occur between Christmas and New Year due to more medical staff being on holiday and, therefore, fewer inductions and elective caesareans being performed during this period (Stats NZ, 2018).

Table I.4 shows the share of children who were born via elective caesearean or whose mothers had their labour induced. There is no statistically significant difference in the shares for the treatment and comparison groups. Moreover, there is no statistically significant difference in average gestation length between the two groups. If mothers manipulated the timing of their birth, it would be expected that there would be more children in the comparison group who were born via elective caesearean or induced labour, and that they would have a shorter average gestation length.

For confidentiality reasons, the IDI only provides individuals' month and year of birth, not the exact date. Therefore, it is difficult to examine just those with due dates very close to the policy-change cut-off date. However, Figure I.1 investigates the distribution of births before and after the policy change. Panel A shows the monthly births to those who were ineligible for birthright citizenship under the new policy as a proportion of all births during the relevant calendar year for this group. For comparison, Panel B shows the same numbers but for those who were eligible for birthright citizenship under the new policy. In Panel A, there is no discernible pattern around the policy cut-off date, with the monthly variation being of a similar magnitude as in other years. In Panel B, there is less monthly variation due to a much larger number of children in the eligible group, but a comparison of the patterns for ineligible and eligible births does not suggest anything of concern.

Although it seems unlikely that those with due dates close to the cut-off manipulated their birth date, excluding births very close to the cut-off date may be prudent. Thus, as a robustness test, a doughnut sample that excludes births in December 2005 from the comparison group and January 2006 from the treatment group is used (see Section 7).

The second potential issue is that the removal of birthright citizenship may have increased the price of child quality and caused immigrant parents to increase their desired number of children, and reduce their investment in each child Avitabile et al. (akin to the findings for Germany of 2014). However, the fertility of affected parents is examined empirically, and no statistically significant difference in fertility is found (see Section 6).

Concurrent policy changes

If other policy changes that impacted on the population of interest happened at the same time as the birthright citizenship change then the analysis may not be able to separate out the effect of birthright citizenship from the effect of other policy changes. As far as I am aware, there were no relevant concurrent policy changes. However, paid parental leave was introduced in 2006 and was available to mothers who gave birth or were due to give birth on or after 1 July 2006. While the main treatment window, which includes children born six-months on either side of January 2006, ends in June 2006, it is likely that a few mothers whose children were born towards the end of this period but had a due date on or after 1 July were eligible for paid parental leave. While this would only impact a very small number of mothers in the treatment sample, as a robustness test, a three-month treatment window whereby the control group consists of children born between October and December 2005 and the treatment group consists of children born between January and March 2006, is also used (see Section 7).

Selective outmigration

Selective outmigration is a known issue in research examining immigrants' outcomes (Dustmann & Görlach, 2015). For example, those who are less successful at integrating into the host country's labour market may be more likely to leave the country. Thus, the labour market successes of migrants may be overestimated if only a non-random selection of migrants who remain in the country are examined.

In the current analysis, all migrants remain in the population of interest regardless of whether they leave the country or not, which, at least on the face of it, mitigates the selective outmigration issue. However, changes in outmigration patterns due to the policy change may still be an issue. For example, for labour market outcomes, having no earnings in NZ tax records could indicate that the individual is in NZ but

Figure I.1: Density of births for population of interest

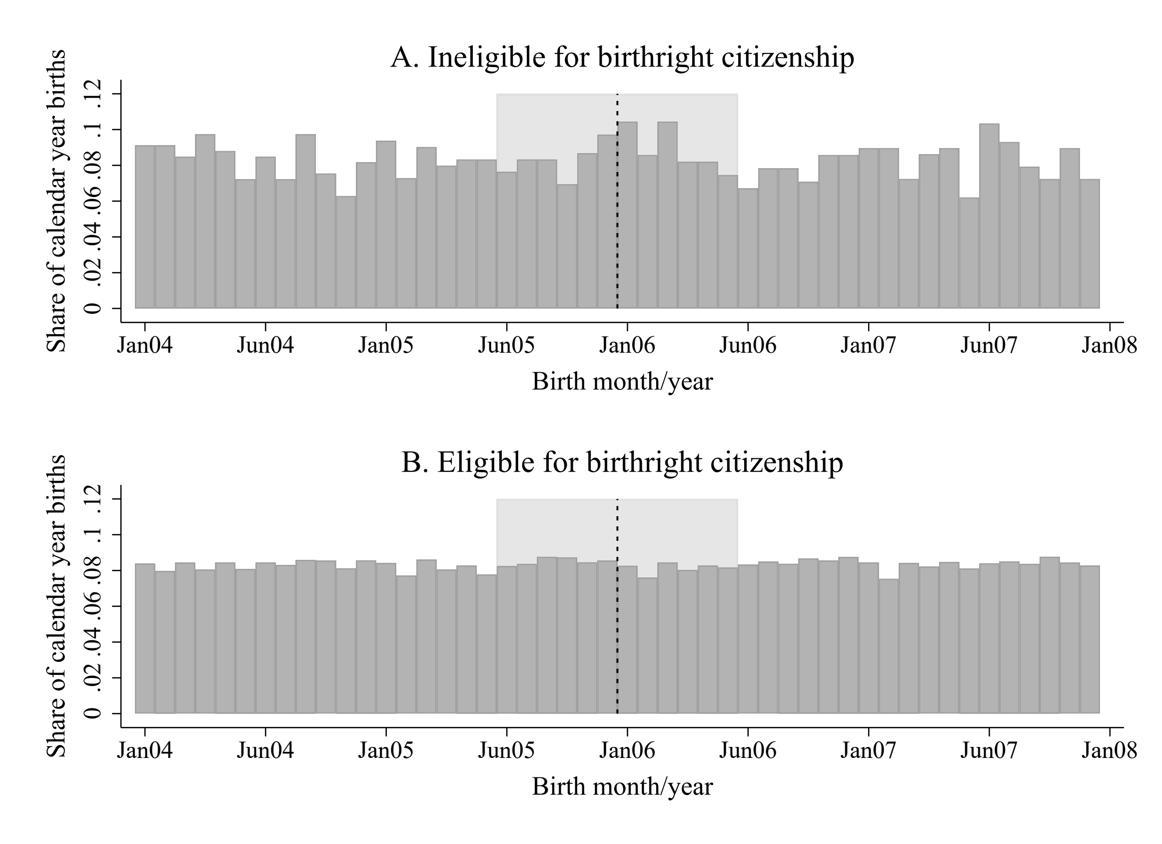

Notes: The light grey shaded area corresponds to births within the six-month treatment window.

has no earnings, or that they have no NZ earnings because they have left the country. Outmigration is, however, examined empirically. Reassuringly for the interpretation of the other integration measures, there is no evidence that the removal of birthright citizenship changed outmigration patterns (see Section 6.1).

Changes in the composition of migrants

The policy change could have changed the composition of the treatment group in other ways. One possibility is that parents who were eligible may have applied for, and been granted, residency sooner than they would have otherwise in order to ensure their child could become an NZ citizen at birth. However, this may not have been practically possible given residency visas can take some time to process. Furthermore, an examination of MBIE visa data does not show an increase in residency applications leading up to the policy change.

Another possibility is that the composition of migrants arriving in NZ changed. For example, the introduction of birthright citizenship in Ireland changed the characteristics of migrants - those who migrated after the policy change had, on average, higher education levels than those who migrated before. However, the analysis for

Ireland compared migrants who arrived up to five years after the policy change. In the current case, the treatment group is limited to children born six-months after the policy change. The short-time period post-change means that this is unlikely to be an issue. The longest time period between the legislation being passed and a child in the sample being born is just over 13 months. However, given the removal of birthright citizenship made NZ a potentially less attractive migration destination, it is possible that some families who had been planning to migrate to NZ decided not to.

As a robustness test, a three-month window is used to reduce the possibility that the composition of migrants changed over the treatment window. However, even for the longer six-month window, covariate balance tests suggest that the observable characteristics of the treatment and control groups are very similar, which further suggests migrant selection will not affect the results.

Month-of-birth effects

The comparison group includes children born in the second half of the year, while the treatment group includes those born in the first half of the year. This could lead to differences in the outcomes of the two groups due to seasonal effects. The potential impact of season-of-birth on children's later outcomes is well documented (Card, 1999). However, more recent research suggests the effect is mostly driven by differences in the characteristics of mothers who give birth at different times of year, with more births to younger, less educated, unmarried mothers occurring in winter (Buckles & Hungerman, 2013). Seasonal effects may also impact on parental labour market outcomes, at least in the short-term, depending on when a mother's intended return-to-work date falls relative to the Christmas/summer break (which coincide in NZ as a southern hemisphere country). This could impact mothers' exact return-to-work dates, and therefore, short-term maternal employment outcome measures.

As mentioned, the covariate balance tests suggest that the observable characteristics of the treatment and control groups are very similar, which allays concerns of a season-of-birth effect that is driven by differences in parental characteristics. Moreover, the preferred specification includes month-of-birth control variables. However, to further safeguard against this possibility, a 12-month treatment window is also used so that both the treatment and comparison groups include children born over an entire calendar year.

5.2 Covariate balance

Table I.4 compares the characteristics of children in the treatment and comparison groups. There are few differences in the general characteristics of the two groups. The only statistically significant differences are that the treatment group has a lower share of Asian children (significant at the 10% level) and a higher share of children residing in Canterbury (significant at the 10% level) or the rest of the South Island (significant at the 5% level). The magnitude of the differences are not large, and likely reflect that the number of children who are ineligible for birthright citizenship is actually quite small. As expected, if the treatment window is reduced to three months, there are fewer differences, with the only remaining difference being a higher share in the treatment group residing in Canterbury (significant at the 10% level) (not shown).

In terms of the child's birth characteristics, as mentioned, there are no statistically significant differences in average gestation length, or the share delivered via elective caesearean or whose mothers had their labour induced. There are also no differences in average birth weight or Apgar score (a standardised score out of 10 of infant health immediately after delivery).

Table I.5 examines the characteristics of parents. For mothers, the only statistically significant difference in general characteristics is the share of mothers whose region of nationality is Asia (significant at the 10% level). For fathers, the share whose region of nationality is Asia is also higher in the treatment than comparison group (significant at 5% level). There are also fewer fathers whose region of nationality is Europe (significant at the 10% level). There is also a higher share of fathers in the treatment group who are on 'other' temporary visa types (significant at the 5% level).

As with children's characteristics, if the treatment window is reduced to three months, as expected, any remaining differences between the treatment and comparison groups are further diminished (not shown). Only two statistically significant differences remain and both are only significant at the 10% level: the share of fathers who are of Asian nationality and the share of fathers who are on 'other' temporary visa types.

Table I.4: Covariate balance across treatment and comparison groups: Children's characteristics

	Comparison		Treatment		
	Count	Mean (s.d.)	Count	Mean (s.d.)	t-statistic
Child's general characteristics					
Female	426	0.47	429	0.51	-1.02
Ethnicity					
Pacific peoples	426	0.24	429	0.26	-0.66
Asian	426	0.41	429	0.35	1.84*
MELAA/Other	426	0.06	429	0.08	-0.90
European	426	0.30	429	0.32	-0.82
Region					
Auckland	423	0.57	420	0.54	1.24
Waikato	423	0.06	420	0.06	-0.16
Wellington	423	0.11	420	0.09	0.89
Rest of North Island	423	0.12	420	0.11	0.63
Canterbury	423	0.09	420	0.13	-1.81*
Rest of South Island	423	0.05	420	0.08	-1.99**
First born child for parent pairing	426	0.58	429	0.57	0.33
First born child for mother	426	0.55	429	0.55	0.32
First born child for father	426	0.56	429	0.55	0.25
Child's birth characteristics					
Gestation length (weeks)	426	39.13 (1.88)	426	39.34 (3.35)	-0.13
Birth weight (grams)	426	3378 (593)	426	3435 (601)	-1.01
Apgar score (0-10)	330	9.61 (0.70)	339	9.54 (0.76)	0.79
Part of multiple birth	426	s.	429	s.	0.07
Delivered via elective caesearean	411	0.10	396	0.08	-0.0057
Induced labour	402	0.17	390	0.18	1.15

Notes: This table presents t-test comparisons of the treatment and comparison groups using a six-month treatment window. Notation "s." means counts have been suppressed in accordance with Stats NZ confidentiality rules. Some characteristics have a smaller number of observations due to missing information. Asterix represents statistical significance at conventional levels, where * if $p < 0.10$, ** if $p < 0.05$, and *** if $p < 0.01$.

Table I.5: Covariate balance across treatment and comparison groups: Parents' characteristics

	Comparison		Treatment		
	Count	Mean (s.d.)	Count	Mean (s.d.)	t-statistic
Mother's characteristics					
Age at child's birth (years)	426	30.38 (5.16)	429	30.21 (5.22)	0.48
Ethnicity					
Pacific peoples	426	0.24	429	0.25	-0.58
Asian	426	0.41	429	0.36	1.49
MELAA/Other	426	0.06	429	0.08	-0.76
European	426	0.29	429	0.31	-0.61
Region of nationality					
Africa	411	0.10	423	0.09	0.82
Americas	411	0.07	423	0.07	-0.58
Asia	411	0.38	420	0.31	2.08**
Europe	411	0.18	423	0.23	-1.63
Oceania	411	0.27	423	0.30	-0.93
Visa type					
Work	411	0.63	420	0.67	-1.42
Student	411	0.06	420	0.05	0.36
Visitor	411	0.30	420	0.24	1.71
Other temporary visa	420	s.	429	0.02	-1.27
Years since arrival	426	2.14	429	2.48	-1.26
Father's characteristics					
Age at birth (years)	426	33.11 (5.93)	429	33.09 (5.90)	0.04
Ethnicity					
Pacific peoples	426	0.24	429	0.26	-0.58
Asian	426	0.40	429	0.32	2.22**
MELAA/Other	426	0.06	429	0.09	-1.25
Europe	426	0.30	429	0.33	-1.03
Region of nationality					
Africa	408	0.10	411	0.10	0.29
Americas	408	0.07	411	0.07	0.32
Asia	408	0.36	411	0.29	1.99**
European	408	0.18	411	0.24	-1.78*
Oceania	408	0.27	411	0.30	0.41
Visa type					
Work	405	0.73	411	0.75	-0.68
Student	405	0.10	411	0.07	1.56
Visitor	405	0.16	411	0.15	0.57
Other temporary visa	405	s.	411	0.04	-2.50**
Years since arrival	426	2.14	429	2.48	-1.26
Income and deprivation					
Mother's wage/salary earnings in year before birth (2006Q1$)	426	8,481 (16479)	429	8,296 (15079)	0.17
Father's wage/salary earnings in year before birth (2006Q1$)	426	25,901 (34886)	429	29,889 (37490)	-1.61
Parents' wage/salary earnings in year before birth (2006Q1$)	426	34,382 (43005)	429	38,186 (42335)	-1.30
Deprivation index	426	6.18 (2.73)	429	6.27 (2.82)	-0.37

Notes: This table presents t-test comparisons of the treatment and comparison groups using a six-month treatment window. Notation "s." means counts have been suppressed in accordance with Stats NZ confidentiality rules. Some characteristics have a smaller number of observations due to missing information. MELAA is Middle Eastern, Latin American and African ethnicity. Deprivation is measured on a 10-point scale with 1 being the least deprived and 10 being the most deprived. Asterix represents statistical significance at conventional levels, where * if $p < 0.10$, ** if $p < 0.05$, and *** if $p < 0.01$.

6 Results

This section presents results on family outmigration, parents' fertility and labour market outcomes, as well as children's health outcomes. The results use the base model specification of a six-month treatment window.

6.1 Outmigration

Figure I.2 presents descriptive trends on the number of days spent overseas in the seven years before and 10 years after the child's birth for the treatment and comparison groups. Seven years before the child's birth, mothers (Figure I.2A) and fathers (Figure I.2B) spent almost all their time overseas, with the average days spent overseas in that year being 336 for mothers and 333 for fathers. The fact that affected parents spent most of their time overseas seven years before the birth of the child is expected as if they had been living in NZ seven years before their child's birth, it is most likely that they would have transitioned from a temporary visa to a resident visa by the time their child was born and, therefore, not be part of the population of interest. Indeed, the descriptive statistics in Table I.5 show that, on average, parents had spent just over two years in NZ before their child's birth.

The amount of time that both mothers and fathers spent overseas decreases monotonically as the time to their child's birth gets closer. Both parents spend the lowest average number of days overseas during the year that their child was born (just over 50 days). The average number of days spent overseas then increases in the years after the child's birth but remains lower than in the early pre-birth years. By the 10th year after the child's birth, mothers spend an average of 155 days overseas, while fathers spend an average of 165 days overseas.

These descriptive graphs also suggest that there is little difference between the treatment group and comparison groups. The patterns in the days spent overseas over time are very similar for mothers (Figure I.2A), fathers (Figure I.2B) and children (Figure I.2C). Moreover, the pre-birth trends for mothers and fathers both suggest that the parallel trends assumption is met since the patterns for the treatment and comparison group are not only very close to being parallel, they are almost the same in terms of levels as well, with the treatment group spending just a little less time overseas than the comparison group in the years before the child's birth.

Figure I.2: Average number of days spent overseas over time

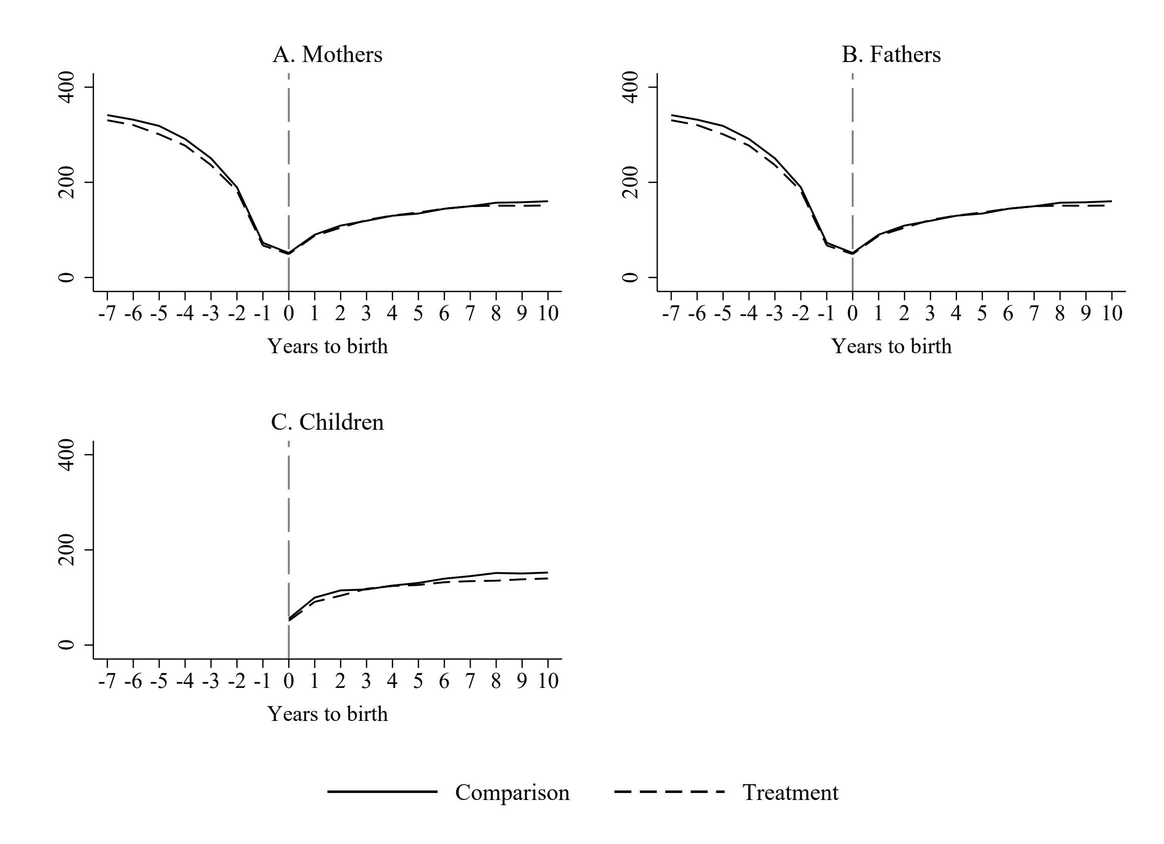

Panel A of Table I.6 presents two-period DiD results for the number of days spent overseas. Columns 1-3 show results for mothers, with Column 1 showing results from an estimation of Equation 3 with no controls, Column 2 showing the results with demographic controls (listed in Panel B of Table I.3) and Column 3 with demographic and month-of-birth controls. Columns 4-6 show the corresponding results for fathers. Results for children are not shown since there is no measure of the number of days they spent overseas before their birth.

The coefficient of interest, the interaction term between the treatment and post-period indicators, is not statistically significant in any specification for either mothers or fathers. That is, there is no evidence that the removal of birthright citizenship had an impact on outmigration patterns.

Table I.6: Two-period difference-in-differences: Parents' outcomes

	Mothers			Fathers		
	(1)	(2)	(3)	(4)	(5)	(6)
A. Days spent overseas						
Treatment indicator	-11.52*	-16.31**	-34.14**	-8.29	-10.80	-27.86*
	(6.70)	(6.60)	(14.22)	(6.76)	(6.65)	(14.29)
Post-period indicator	-128.92***	-129.78***	-129.78***	-115.86***	-116.08***	-116.08***
	(7.14)	(7.33)	(7.33)	(6.99)	(7.19)	(7.19)
Treatment * Post	9.06	10.23	10.23	5.33	4.66	4.66
	(10.05)	(10.25)	(10.25)	(9.91)	(10.23)	(10.23)
B. Cumulative number of children						
Treatment indicator	0.03	0.02	0.04	0.03	0.02	0.05
	(0.02)	(0.03)	(0.08)	(0.03)	(0.03)	(0.08)
Post-period indicator	0.59***	0.57***	0.57***	0.59***	0.58***	0.58***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Treatment * Post	0.01	0.01	0.01	0.01	0.02	0.02
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
C. Employment rate						
Treatment indicator	0.01	0.01	-0.01	0.06	0.05	0.05
	(0.02)	(0.02)	(0.04)	(0.02)	(0.02)	(0.04)
Post-period indicator	0.20***	0.20***	0.20***	0.29***	0.29***	0.29***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Treatment * Post	-0.01	-0.02	-0.02	-0.03	-0.03	-0.03
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
D. Earnings, 2006Q1 NZ$						
Treatment indicator	341	525	-726	1,981**	1,489*	815
	(435)	(461)	(1313)	(777)	(875)	(2345)
Post-period indicator	5,550***	5,658***	5,658***	16,119***	16,536***	16,536***
	(612)	(624)	(624)	(1240)	(1298)	(1298)
Treatment * Post	-65	-177	-177	367	132	132
	(920)	(944)	(944)	(1832)	(1892)	(1892)
Demographic controls	No	Yes	Yes	No	Yes	Yes
Month of birth controls	No	No	Yes	No	No	Yes
No. of observations	15,390	14,760	14,760	15,390	14,562	14,562
No. of individuals	855	820	820	855	809	809

Notes: This table presents the difference-in-differences regression results from Equation 3 for mothers and fathers, respectively. Standard errors are in parentheses. Asterix represents statistical significance at conventional levels, where * if $p < 0.10$, ** if $p < 0.05$, and *** if $p < 0.01$.

Figure I.3: Dynamic difference-in-differences results: Days spent overseas

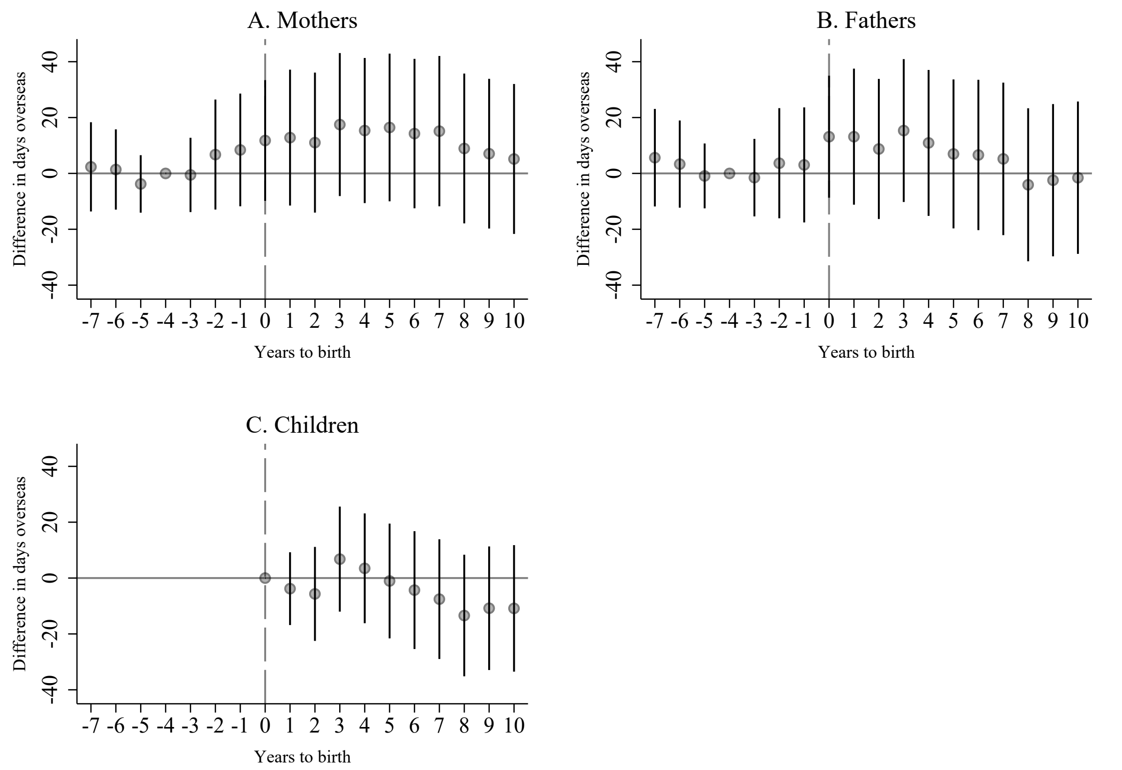

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. For Panels A and B, the base year is -4. For Panel C, the base year is 0.

To examine these results in more detail, dynamic DiD results are summarised graphically in Figure I.3. This presents the estimation of Equation 2 with demographic and month-of-birth controls included for mothers, fathers and children separately. Before the child's birth, there are no statistically significant differences in the number of days parents spent overseas, which, once again, highlights that the parallel trends assumption holds. After the child's birth, mothers of children in the treatment group spend more days overseas than mothers of children in the comparison group, but the difference is not statistically significant. Likewise, the number of days that fathers of children in the treatment group spent overseas is higher than the number of days that fathers of children in the comparison group spent overseas in some of the years after the child's birth, but none of these differences are statistically significant. For children, the difference between the number of days time spent overseas since their birth is sometimes lower and sometimes higher for the treatment versus the comparison group, and the difference is never statistically significant. These results contrast to Sajons (2016), which found that the introduction of birthright citizenship in Germany reduced the likely that the parents of immigrant children would leave the country.

It is unclear why the removal of birthright citizenship did not change the outmigration patterns of affected NZ migrant families. Although speculative, it could be a result of two opposing forces offsetting each other. The removal of birthright citizenship may have resulted in some parents staying in NZ longer in order to secure permanent residency or citizenship for themselves and their children, thus allowing their children to live in NZ at any stage in the future. It may have also meant that some left sooner due to the benefits to their children of staying in NZ being lower in the absence of birthright citizenship. These two effects may have offset each other. However, another possible explanation is that the removal of birthright citizenship did not materially change families' outmigration decisions. This explanation seems plausible in the case of NZ given that virtually all those on temporary visas have a pathway to residency and then citizenship. Moreover, NZ's immigrant population are relatively high skilled compared with the native-born population, in contrast to Germany. The benefits of citizenship are likely to be lower for higher-skilled migrants. Therefore, whether their child is granted birthright citizenship or not may not generally be a big factor in the decision making of NZ migrants.

6.2 Fertility

The international evidence on whether birthright citizenship impacts fertility is mixed. Evidence for Germany finds it reduces fertility among immigrant parents, while it is found to have no effect in the Dominican Republic and to increase fertility in Ireland. Thus, the effect of birthright citizenship on fertility seems to be context specific and, therefore, an empirical question.

Figure I.4 presents descriptive fertility trends for the treatment group versus the comparison group. Fertility is measured as the cumulative number of children over time and results are presented separately for mothers and fathers. These results suggest that the cumulative number of children for both mothers and fathers is slightly higher among the treatment group, although the magnitude of the difference is small and the trends over time for the two groups are very similar.

Examining this more formally with a two-period difference-in-differences analysis (Panel B of Table I.6), the coefficient of interest is very small and not statistically significant. Thus, there is no evidence that the removal of birthright citizenship impacted the fertility of immigrant parents. This is confirmed by the dynamic difference-in-difference results, with the cumulative number of children not being statistically significantly different between the treatment and comparison groups (Figure I.5).

Why do NZ's results differ from Germany's and Ireland's? Once again, it may be that the benefits of birthright citizenship are not as great in NZ due to NZ migrants being relatively well educated and/or because virtually all NZ migrants have a pathway to residency and then citizenship. Thus, birthright citizenship may not influence NZ migrants' behaviour.

Figure I.4: Cumulative number of children over time

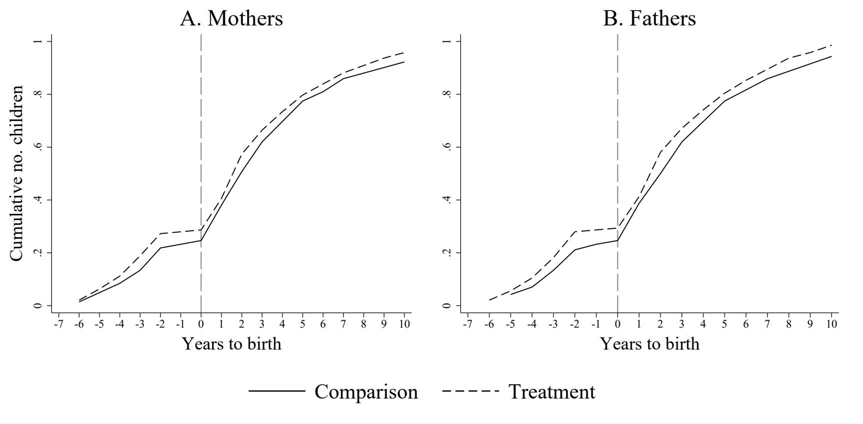

Figure I.5: Dynamic difference-in-differences results: Cumulative number of children

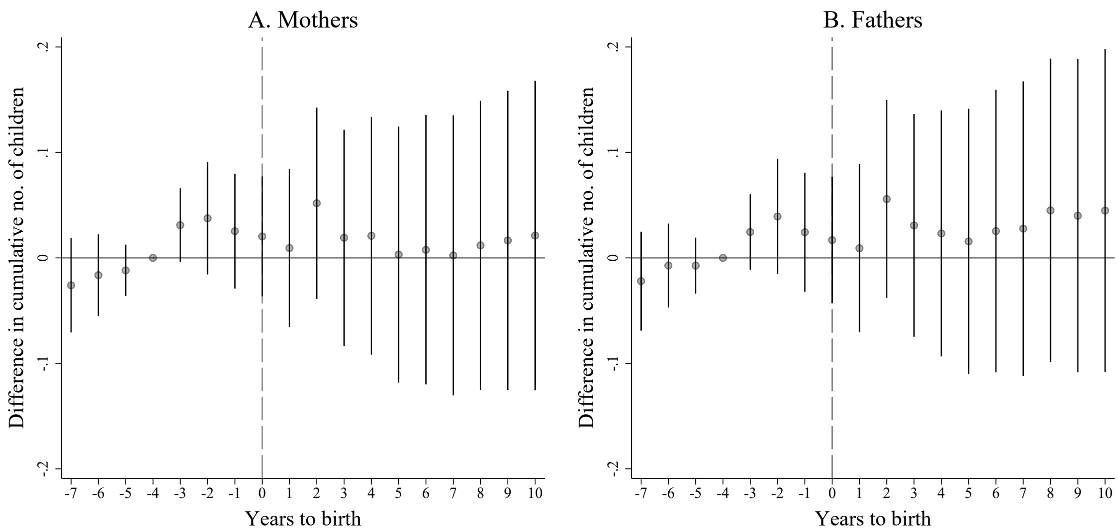

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

6.3 Labour market outcomes

The removal of birthright citizenship may have impacted parental employment and earnings. However, the direction of the effect is theoretically ambiguous. Removing birthright citizenship may reduce the incentives to invest in country-specific human capital and therefore worsen labour market outcomes. On the other hand, it may reduce payoffs from migrants' investments in their children, thus leading to less time spent at home with their children and more time spent in employment.

Employment

Descriptively, Figure I.6 suggests that the removal of birthright citizenship made no material difference to the employment rates of affected mothers and fathers. For mothers (Figure I.6A), seven years before the child's birth, employment rates are close to zero for both the treatment and control groups. This reflects that most of those in the sample were living overseas at that time (see Figure I.2). These employment rates steadily increase over time as more mothers arrive and begin to work in NZ. These migration patterns lead to the seemingly unusual finding that maternal employment rates peak close to the time of birth and do not fall much post-birth.

This lack of effect is confirmed by the two-period DiD results (Table I.6, Panel C). The coefficient of interest is small and insignificant for both mothers and fathers and for all specifications. The only significant coefficient is the post-period indicator, as both mothers and fathers have higher employment rates in this period, reflecting that more of them were in the country after their child was born than before it. The dynamic DiD results are also consistent with the two-period results, with no significant differences between the treatment and comparison groups in any time period for either mothers or fathers (Figure I.7).

Figure I.6: Employment rate over time, %

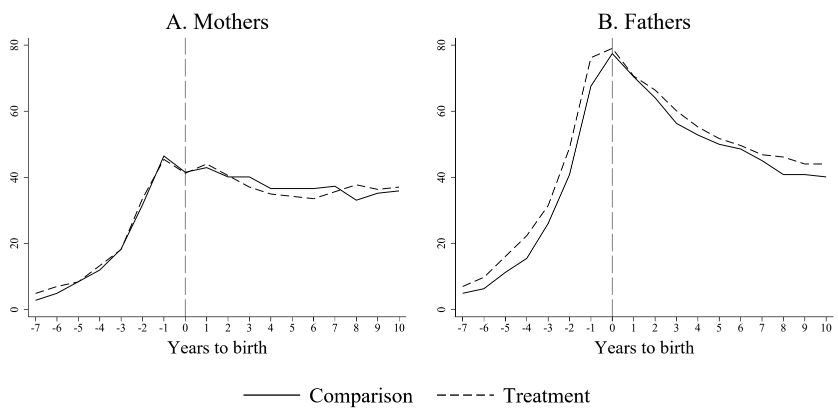

Figure I.7: Dynamic difference-in-differences results: Employment rate

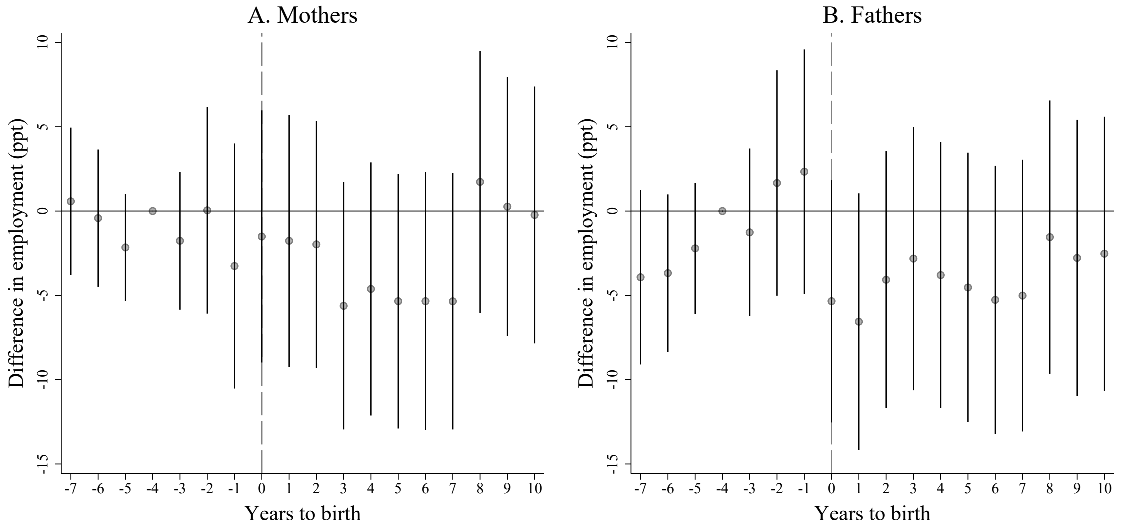

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Earnings

Figure I.8 shows the trends in earnings over time of mothers (Panel A) and fathers (Panel B) in the treatment and comparison groups. As with employment, the treatment and comparison groups have very similar earnings trajectories both before and after the child's birth. These similarities suggest that the removal of birthright citizenship made little difference to parents' labour market outcomes. As with employment, earnings are very low in the years before the child's birth, reflecting that most parents were not yet residing in NZ at that time. Earnings increase as the child's birth draws nearer and more of the parents are residing in NZ.

These descriptive trends are confirmed by two-period DiD results (Table I.6, Panel D). For all specifications for both mothers and fathers, the coefficient of interest is small in magnitude and not statistically significant. Dynamic DiD estimations confirm this result (Figure I.9), with no statistically significant differences between the earnings of the treatment and comparison groups in any time period for both mothers and fathers.

6.4 Children's health outcomes

The removal of birthright citizenship may have changed the incentives for parents to invest in their children. Differences in health outcomes could be one manifestation of this change in investment incentives. Thus, this section examines injuries, hospital admissions and prescription medications for the treatment versus comparison group.

Figure I.8: Annual earnings over time (2006Q1 NZ$)

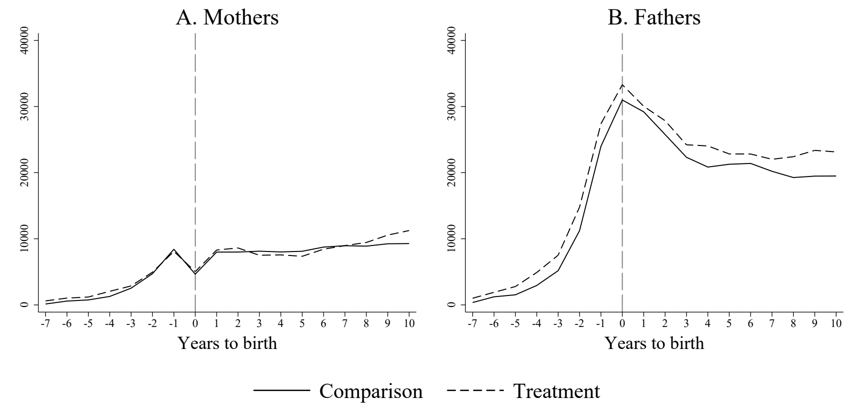

Figure I.9: Dynamic difference-in-differences results: Earnings

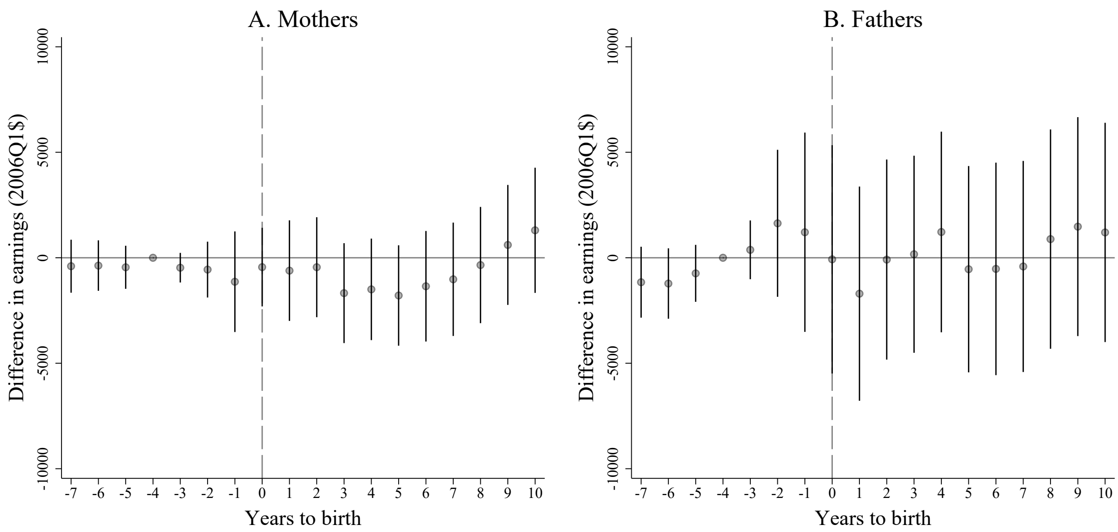

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4. Earnings measured in 2006Q1$.

Injuries

Injuries are measured using accident claims data. As discussed, this health measure is likely a combination of a health outcome measure and a measure of the propensity to seek medical treatment in the event of an accident. However, there are reasons to think that any differences between the treatment and comparison groups reflect differences in health outcomes. First, while existing research suggests there may be some differences in the propensity to seek medical treatment in the event of an injury by characteristics such as ethnicity and age (Poland, 2018), the treatment and comparison groups have very similar characteristics (see Section 4). In addition, the NZ accident compensation system is a universal, no-fault system, which covers all accidents which occur in NZ. Thus, the legal status of the child does not change their ability to access medical treatment that is covered by the accident compensation system. Moreover, even if parents of the treatment group's non-citizen children are unaware of that their child's injury is covered by the accident compensation system, claims for medical treatment are submitted by medical providers, which should further allay concerns about differential access.

Figure I.10 shows descriptively the share of children in the treatment and comparison groups with an accepted accident claim (Panel A) and the average number of claims (Panel B) over time. The time series are somewhat volatile, but there are no systematic differences between the treatment and comparison groups.

The dynamic DiD results (Figure I.11) confirm this lack of difference between the treatment and comparison groups, with no statistically significant differences between the groups in any of the years from their birth to ten years later.

Figure I.10: Accident claims over time

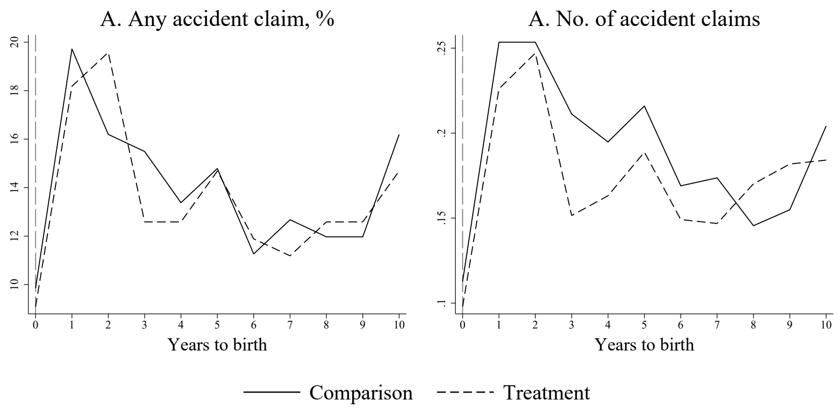

Hospital admissions

As with accidents, hospital admissions are likely a combination of a health outcome and health access measure. Unlike the case of accidents, there may be access differences between the treatment and comparison groups. As citizens, all members of the comparison group would be entitled to free public hospital care. However, only some of the members of the treatment group would be entitled to free care. Those whose parents were work visa holders who were eligible to remain in NZ for two years or more would be entitled to free hospital care. However, those whose parents were on a student or visitor visa, or a work visa of less than two years duration would not be eligible. Many, however, would have been required to hold medical insurance as part of their visa requirements.

Despite these potential differences in access to hospital care, there appears to be little difference in the share of children in the treatment and comparison groups who

Figure I.11: Dynamic difference-in-differences results: Accident claims

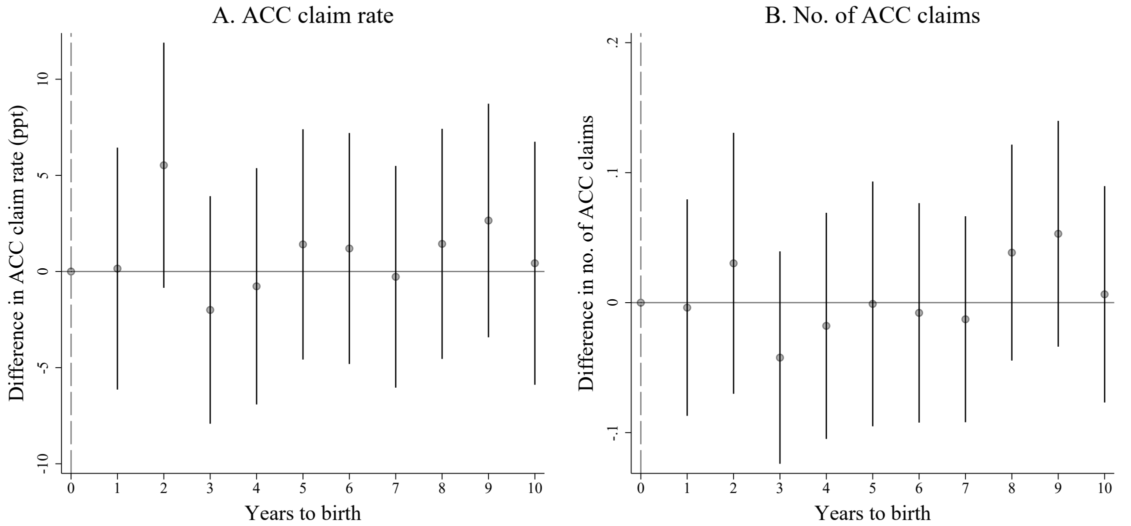

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

had a hospital admission (Figure I.12A), and also little difference in the average number of hospital admissions (Figure I.12B) over time. Almost all children in both groups had at least one hospital admission in the year they were born, which reflects that most children in NZ are born in a hospital. The incidence and number of hospital admissions is much lower in subsequent years.

DiD results confirm the descriptive trends, with no statistically significant differences in hospital admissions between the treatment and comparison groups (Figure I.13).

Figure I.12: Hospital admissions over time

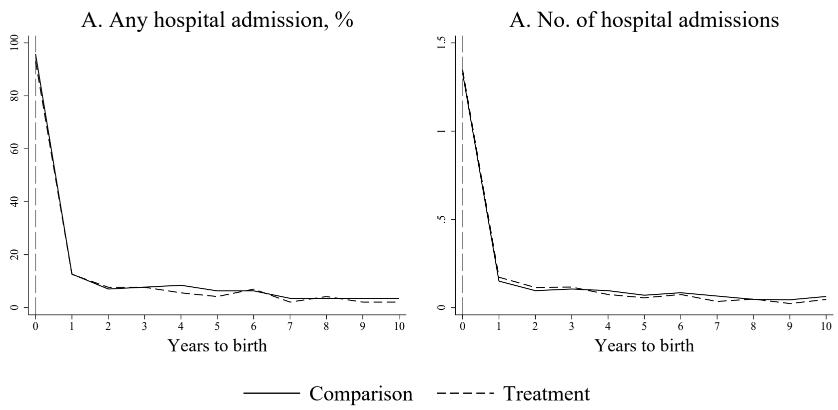

Figure I.13: Dynamic difference-in-differences results: Hospital admissions

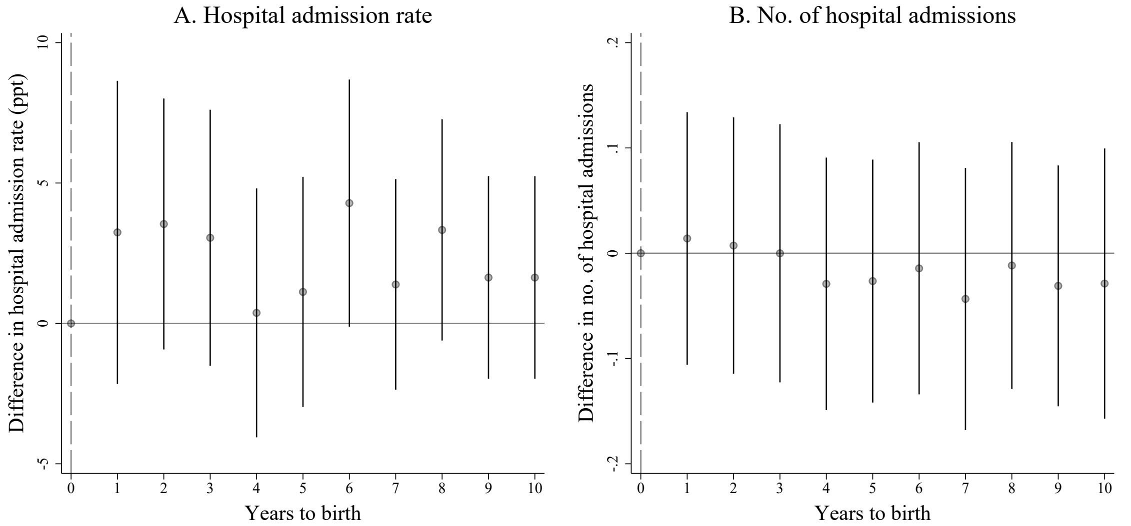

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Prescriptions

Like hospital care, subsidised prescriptions are available to all those in the comparison group, but to only some of those in the treatment group. Thus, prescription medication information may reflect differences in health access as well as health outcomes.

However, as with hospital admissions, the patterns of prescriptions over time are similar for the treatment and comparison groups. The descriptive trends suggest that the treatment group has a slightly lower share of the population who have at least one prescription and slightly lower average number of prescriptions (Figure I.14). However, DiD estimates find no statistically significant differences between the two groups (Figure I.15).

Figure I.14: Prescriptions over time

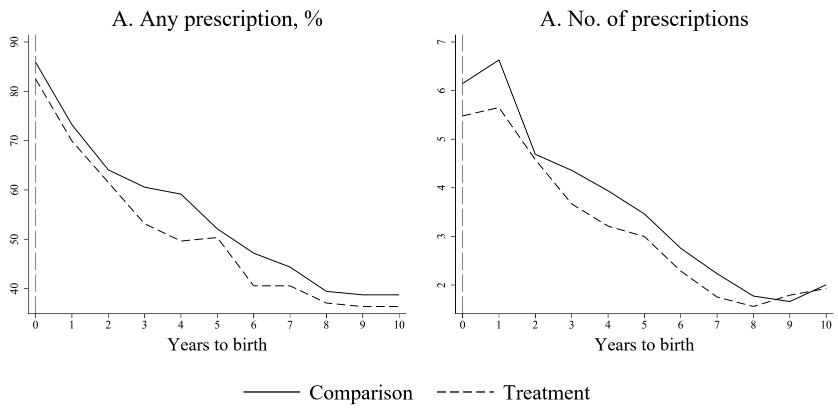

Figure I.15: Dynamic difference-in-differences results: Prescriptions

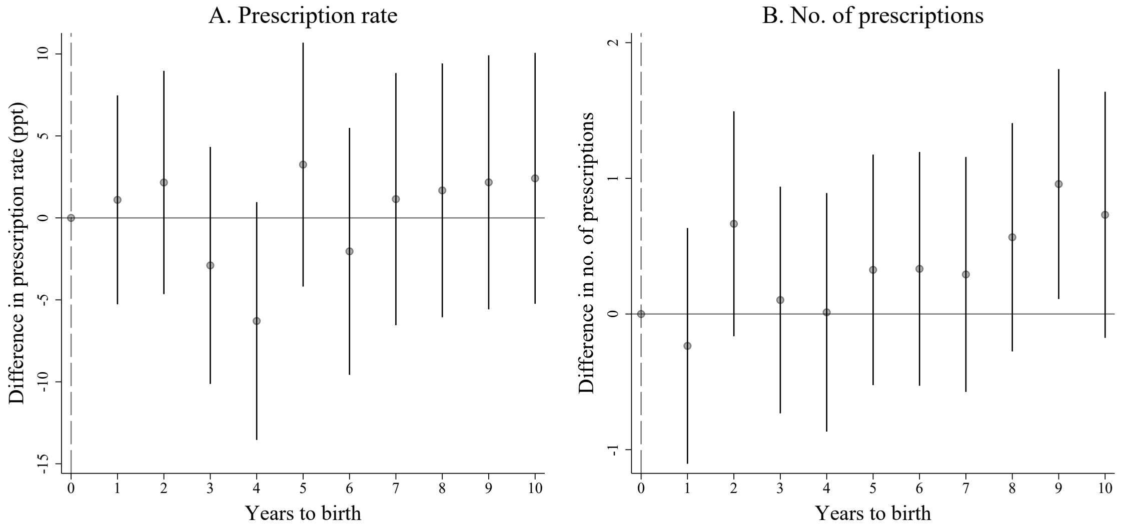

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

7 Robustness

As discussed in Section 5, for robustness, different treatment windows are used. A six-month treatment window with a one-month doughnut is examined to address concerns that the timing of births very close to the policy change date could have been manipulated. A 12-month treatment window is examined to address possible concerns about season-of-birth effects. Finally, a three-month window is examined to further increase the similarity between the treatment and control groups and to address the possible concern that migrants gained residency more quickly in response to the policy change, or that the policy change resulted in a change in the composition of migrants.

The results presented are generally robust to the treatment window used (see Appendix A). For all treatment windows except the 12-month window, the DiD results provide no evidence that the removal of birthright citizenship impacted on any of the outcomes examined. For the 12-month window, the employment rates of mothers in the treatment group were lower in the seven years following the child's birth, and statistically significant in years 2-7. Children's hospital admission rates were also higher and statistically significant for the 12-month window in the first two years following birth.

8 Discussion and conclusion

Does citizenship increase an immigrant's willingness to integrate into the host country's society, or is it the final step in the integration process without further consequences?

The answer to this question has potentially important implications for policy. If citizenship is an important catalyst that encourages migrants to integrate and invest in their own, and their children's, futures, then naturalisation should be made fairly accessible. However, if gaining citizenship does not have an independent effect on integration, or indeed, even reduces incentives to integrate, then a higher bar to receiving a host-country passport may be appropriate, with citizenship being viewed more as a reward for successfully completing the integration process.

However, it is difficult to make causal inferences about the effects of citizenship due to possible selection effects. Those who are more motivated to integrate are also more likely to naturalise. Thus, this paper exploits a natural experiment of NZ's removal of birthright citizenship in 2006 to examine the effects of citizenship on a range of parental and children's outcomes. As such, it adds to the evidence in this area which has, to date, largely be limited to Germany's introduction of birthright citizenship in 2000.

As an anglophone country with high immigration rates and a selective migration policy, the case of NZ is potentially more relevant to Canada and the US, the only two western countries which still retain unrestricted birthright citizenship, than the case of Germany. In addition, NZ's form of birthright citizenship was very similar to existing policy in Canada and the US in that anyone born in the country was entitled to citizenship regardless of their parents' legal status, with the only exception being children born to foreign diplomats. The form of birthright citizenship that Germany introduced is much more restrictive as it still requires at least one parent to have lived legally in Germany for at least eight years. Moreover, the pathways to citizenship in NZ, Canada and the US are very similar in terms of the process and length of time it takes to become a citizen, whereas the citizenship pathway in Germany tends to be longer and more costly, particularly as it generally requires naturalising individuals to relinquish their previous citizenship.

This paper finds that the removal of birthright citizenship did not impact family outmigration, parents' fertility or labour market outcomes, nor children's health outcomes. These results contrast with those for Germany, with several studies finding that the introduction of birthright citizenship improved parents' and children's outcomes across a range of dimensions. This difference in findings highlights that the benefits of citizenship are context specific. It is postulated that citizenship does not have a discernible impact in NZ because the benefits of citizenship are lower than in Germany. In NZ, virtually all migrants have a pathway to residency and citizenship - that is, there are very few true temporary or guest worker migrants. In addition, other residency options provide almost the same benefits as citizenship in NZ. Moreover, reflecting its selective migration policy, NZ's migrants have, on average, higher education levels than the native-born population, while in Germany, migrants have lower average education levels. It is likely that the benefits of citizenship are lower for more skilled migrants, which may also explain the difference in the findings for the two countries.

These results suggest that removing birthright citizenship in Canada and the US may not have much impact on the integration outcomes of immigrants. Like NZ, Canada relatively high-skilled migrant population and a low share of undocumented migrants. For the US, there are some key differences which may mean that the negative effects of removing birthright citizenship could be greater. First, the foreign-born population in the US are relatively low skilled - for example, they have lower average education levels than the native-born population. It is likely that citizenship has more benefits for lower-skilled migrants since it will open up more opportunities than in the case of higher-skilled migrants. Second, the share of undocumented migrants in the US is much higher than in NZ and Canada. Citizenship in general, and the automatic granting of citizenship to their children in particular, is likely to have greater benefits for undocumented migrants than documented ones. However, the NZ results do suggest that the costs of removing birthright citizenship may not be as great in the US as the German case suggests. While citizenship appears to be a catalyst to the integration process in German, it may merely be the final step in the integration process without further consequences for immigrants in countries such as NZ, Canada and the US.

References

Akbari, A. H., & MacDonald, M. (2014). Immigration policy in australia, canada, new zealand, and the united states: An overview of recent trends. *International Migration Review*, *48*(3), 801–822.

Amuedo-Dorantes, C., Bansak, C., & Raphael, S. (2007). Gender differences in the labor market: Impact of IRCA. *American Economic Review*, *97*(2), 412–416. https://doi.org/10.1257/aer.97.2.412

Amuedo-Dorantes, C., Gratereaux Hernández, C., & Pozo, S. (2017). *On the implications of immigration policy restricting citizenship: Evidence from the Dominican Republic* (IZA Discussion Paper No. 10602). https://www.iza.org/publications/dp/10602/on-the-implications-of-immigration-policy-restricting-citizenship-evidence-from-the-dominican-republic

Avitabile, C., Clots-Figueras, I., & Masella, P. (2013). The effect of birthright citizenship on parental integration outcomes. *The Journal of Law & Economics*, *56*(3), 777–810. https://doi.org/10.1086/673266

Avitabile, C., Clots-Figueras, I., & Masella, P. (2014). Citizenship, fertility, and parental investments. *American Economic Journal: Applied Economics*, *6*(4), 35–65.

Barcellos, S. H. (2010). *Legalization and the economic status of immigrants* (RAND Working Paper No. WR-754). https://doi.org/10.7249/WR754

Becker, G. S., & Lewis, H. G. (1973). On the interaction between the quantity and quality of children. *Journal of Political Economy*, *81*(2), S279–S288.

Behr, A., & Fugger, G. (2020). PISA performance of natives and immigrants: Selection versus efficiency. *Open Education Studies*, *2*(1), 9–36. https://doi.org/10.1515/edu-2020-0108

Borjas, G. J., & Tienda, M. (1993). The employment and wages of legalized immigrants. *The International Migration Review*, *27*(4), 712–747. https://doi.org/10.2307/2546910

Bratsberg, B., Ragan, J., James F., & Nasir, Z. M. (2002). The effect of naturalization on wage growth: A panel study of young male immigrants. *Journal of Labor Economics*, *20*(3), 568–597. https://doi.org/10.1086/339616

Buckles, K. S., & Hungerman, D. M. (2013). Season of birth and later outcomes: Old questions, new answers. *The Review of Economics and Statistics*, *95*(3), 711–724. Retrieved June 16, 2023, from https://www.jstor.org/stable/43554790

Card, D. (1999, January 1). The causal effect of education on earnings. In O. C. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (pp. 1801–1863, Vol. 3). Elsevier. https://doi.org/10.1016/S1573-4463(99)03011-4

Chiswick, B. R. (1978). The effect of Americanization on the earnings of foreign-born men. *Journal of Political Economy*, *86*(5), 897–921. Retrieved October 6, 2023, from https://www.jstor.org/stable/1828415

CIMM. (2022). *Minister's appearance before the Standing Committee on Citizenship and Immigration on supplementary estimates (c), 2021-2022 and main estimates, 2022-2023*. House of Commons Canada Standard Committee on Citizenship and Immigration. https://www.canada.ca/en/immigration-refugees-citizenship/corporate/transparency/committees/cimm-mar-03-2022.html

Dahl, G. B., Felfe, C., Frijters, P., & Rainer, H. (2022). Caught between cultures: Unintended consequences of improving opportunity for immigrant girls. *The Review of Economic Studies*, *89*(5), 2491–2528.

Dasgupta, K., Diegmann, A., Kirchmaier, T., & Plum, A. (2022). The gender reveal: The effect of sons on young fathers' criminal behavior and labor market activities. *Labour Economics*, *78*, 102224. https://doi.org/10.1016/j.labeco.2022.102224

Devillanova, C., Fasani, F., & Frattini, T. (2018). Employment of undocumented immigrants and the prospect of legal status: Evidence from an amnesty program. *ILR Review*, *71*(4), 853–881.

DeVoretz, D. J., & Pivnenko, S. (2005). The economic causes and consequences of Canadian citizenship. *Journal of International Migration and Integration*, *6*(3), 435–468. https://doi.org/10.1007/s12134-005-1021-6

Dustmann, C. (2003). Children and return migration. *Journal of Population Economics*, *16*(4), 815–830. https://doi.org/10.1007/s00148-003-0161-2

Dustmann, C., & Görlach, J.-S. (2015, January 1). Chapter 10 - selective out-migration and the estimation of immigrants' earnings profiles. In B. R. Chiswick & P. W. Miller (Eds.), *Handbook of the economics of international migration* (pp. 489–533, Vol. 1). North-Holland. https://doi.org/10.1016/B978-0-444-53764-5.00010-4

Dustmann, C., & Landersø, R. (2021). Child's gender, young fathers' crime, and spillover effects in criminal behavior. *Journal of Political Economy*, *129*(12), 3261–3301. https://doi.org/10.1086/716562

Felfe, C., Kocher, M. G., Rainer, H., Saurer, J., & Siedler, T. (2021). More opportunity, more cooperation? the behavioral effects of birthright citizenship on immigrant youth. *Journal of Public Economics*, *200*, 104448. https://doi.org/10.1016/j.jpubeco.2021.104448

Felfe, C., Rainer, H., & Saurer, J. (2020). Why birthright citizenship matters for immigrant children: Short- and long-run impacts on educational integration. *Journal of Labor Economics*, *38*(1), 143–182.

Foad, H. (2022). Birthright granted and revoked: The effects of Irish citizenship policy on migrant fertility. *AEA Papers and Proceedings*, *112*, 391–395. https://doi.org/10.1257/pandp.20221025

Fougère, D., & Safi, M. (2009). Naturalization and employment of immigrants in France (1968-1999) (A. F. Constant, M. Kahanec, & K. F. Zimmermann, Eds.) [Publisher: Emerald Group Publishing Limited]. *International Journal of Manpower*, *30*(1), 83–96. https://doi.org/10.1108/01437720910948410

Geiringer, C. (2008). Ding v Minister of Immigration; Ye v Minister of Immigration. https://www.wgtn.ac.nz/public-law/publications/working-papers/pdfs/VUW-NZCPL-001.pdf

Hainmueller, J., Hangartner, D., & Pietrantuono, G. (2017). Catalyst or crown: Does naturalization promote the long-term social integration of immigrants? *American Political Science Review*, *111*(2), 256–276. https://doi.org/10.1017/S0003055416000745

Heron, M., & Barrow, J. (2023). *A review of the processes and procedures around out of hours immigration compliance activity, and to identify and recommend potential changes to the process where required*. Michael Heron KC, Barrister. https://www.mbie.govt.nz/dmsdocument/26981-mhkc-inz-out-of-hours-final-report-29-june-2023

Kaushal, N. (2006). Amnesty programs and the labor market outcomes of undocumented workers. *Journal of Human Resources*, *XLI*(3), 631–647. https://doi.org/10.3368/jhr.XLI.3.631

Kossoudji, S. A., & Cobb-Clark, D. A. (2002). Coming out of the shadows: Learning about legal status and wages from the legalized population. *Journal of Labor Economics*, *20*(3), 598–628. https://doi.org/10.1086/339611

Kuziemko, I. (2014). Human capital spillovers in families: Do parents learn from or lean on their children? *Journal of Labor Economics*, *32*(4), 755–786. https://doi.org/10.1086/677231

Lozano, F., & Sørensen, T. A. (2011). *The labor market value to legal status* (IZA Discussion Paper Series No. 5492). https://repec.iza.org/dp5492.pdf

Mazzolari, F. (2009). Dual citizenship rights: Do they make more and richer citizens? *Demography*, *46*(1), 169–191. https://doi.org/10.1353/dem.0.0038

Migration Policy Institute. (2020). Profile of the unauthorized population: United states. https://www.migrationpolicy.org/data/unauthorized-immigrant-population/state/US

OECD. (2023a). Education and labour market outcomes of native- and foreign-born adults dataset. Retrieved September 16, 2023, from https://stats.oecd.org/Index.aspx?DataSetCode=EAG_MIGR#

OECD. (2023b). Level of GDP per capita dataset. Retrieved September 16, 2023, from stats.oecd.org

Owsald, A. J., & Powdthavee, N. (2010). Daughters and left-wing voting. *The Review of Economics and Statistics*, *92*(2), 213–227. Retrieved September 26, 2023, from https://www.jstor.org/stable/27867533

Pan, Y. (2012). The impact of legal status on immigrants' earnings and human capital: Evidence from the IRCA 1986. *Journal of Labor Research*, *33*(2), 119–142. https://doi.org/10.1007/s12122-012-9134-0

Poland, M. (2018). *The determinants of injury compensation claims in a universal claims environment* [Doctoral dissertation, University of Otago]. https://ourarchive.otago.ac.nz/bitstream/handle/10523/8497/PolandMichelleA2018PhD.pdf?sequence=1&isAllowed=y

Sajons, C. (2016). Does granting citizenship to immigrant children affect family outmigration? *Journal of Population Economics*, *29*(2), 395–420.

Sajons, C. (2019). Birthright citizenship and parental labor market integration. *Labour Economics*, *57*, 1–22. https://doi.org/10.1016/j.labeco.2019.01.001

Sajons, C., & Clots-Figueras, I. (2014). Birthright citizenship and education - do immigrant children need a passport to thrive? *VfS Annual Conference 2014 (Hamburg): Evidence-based Economic Policy*. Retrieved May 3, 2023, from https://ideas.repec.org//p/zbw/vfsc14/100470.html

Salmond, C., Crampton, P., & Atkinson, J. (2007). *NZDep2006 index of deprivation*. Department of Public Health, University of Otago. Wellington. https://www.otago.ac.nz/__data/assets/pdf_file/0027/317574/nzdep2006-index-of-deprivation-research-report-020348.pdf

Sawyer, C. (2013). The loss of birthright citizenship in new zealand. *Victoria University of Wellington Law Review*, *44*(3), 653–674. https://doi.org/10.3316/informit.859485620776206

Stats NZ. (2018). *Most common birthday table*. Stats NZ. https://www.stats.govt.nz/assets/Tools/Most-common-birthday-in-New-Zealand/most-common-birthdays-1980-2017.xlsx

Steinhardt, M. F. (2012). Does citizenship matter? the economic impact of naturalizations in germany. *Labour Economics*, *19*(6), 813–823. https://doi.org/10.1016/j.labeco.2012.09.001

Washington, E. L. (2008). Female socialization: How daughters affect their legislator fathers' voting on women's issues. *The American Economic Review*, *98*(1), 311–332. Retrieved September 26, 2023, from https://www.jstor.org/stable/29729973

Appendices

A Robustness: Varying treatment window graphs

Figure I.A.1: Dynamic DiD results by treatment window: Days spent overseas: Mothers

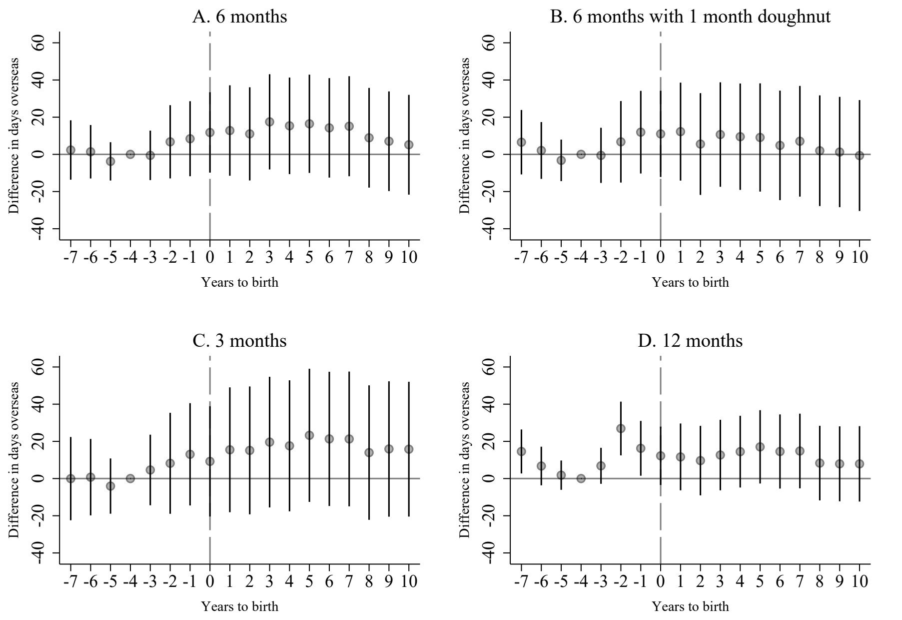

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Figure I.A.2: Dynamic DiD results by treatment window: Days spent overseas: Fathers

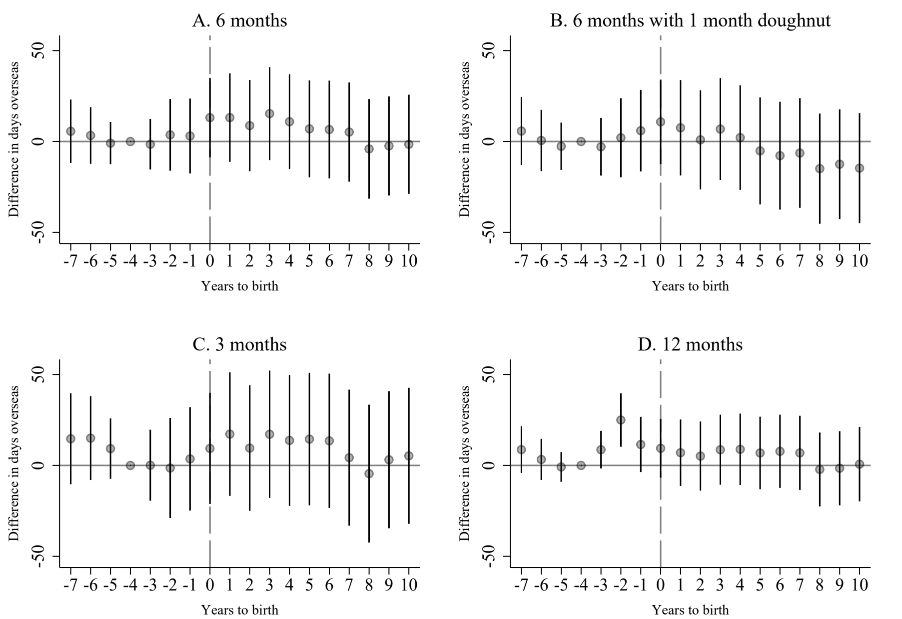

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Figure I.A.3: Dynamic DiD results by treatment window: Days spent overseas: Children

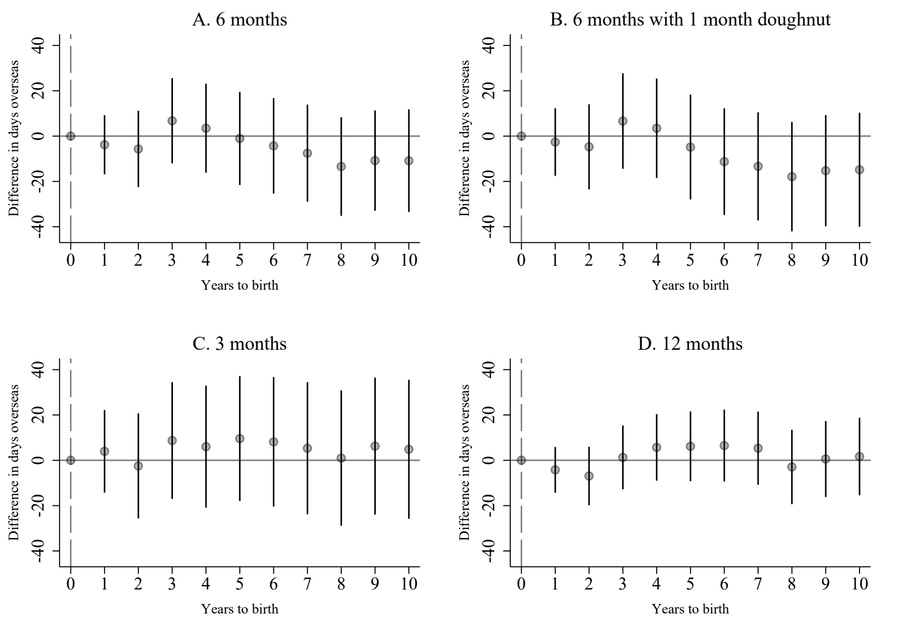

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Figure I.A.4: Dynamic DiD results by treatment window: Cumulative no. of children: Mothers

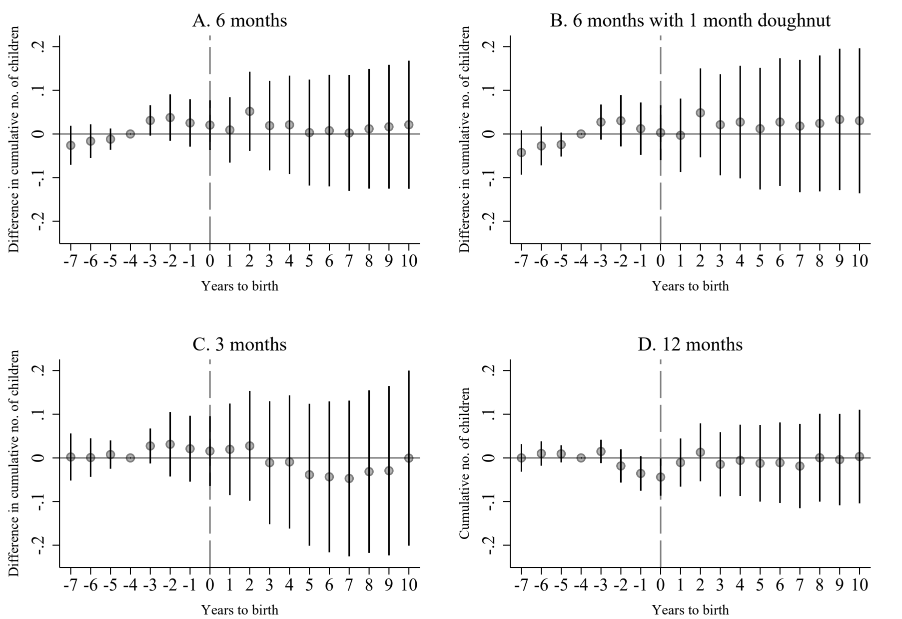

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Figure I.A.5: Dynamic DiD results by treatment window: Cumulative no. of children: Fathers

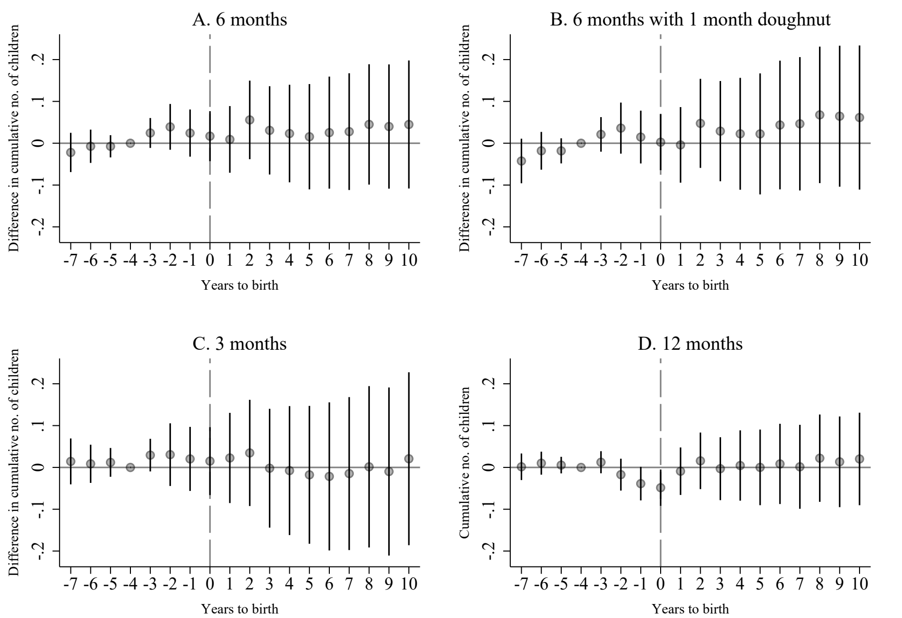

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Figure I.A.6: Dynamic DiD results by treatment window: Employment: Mothers

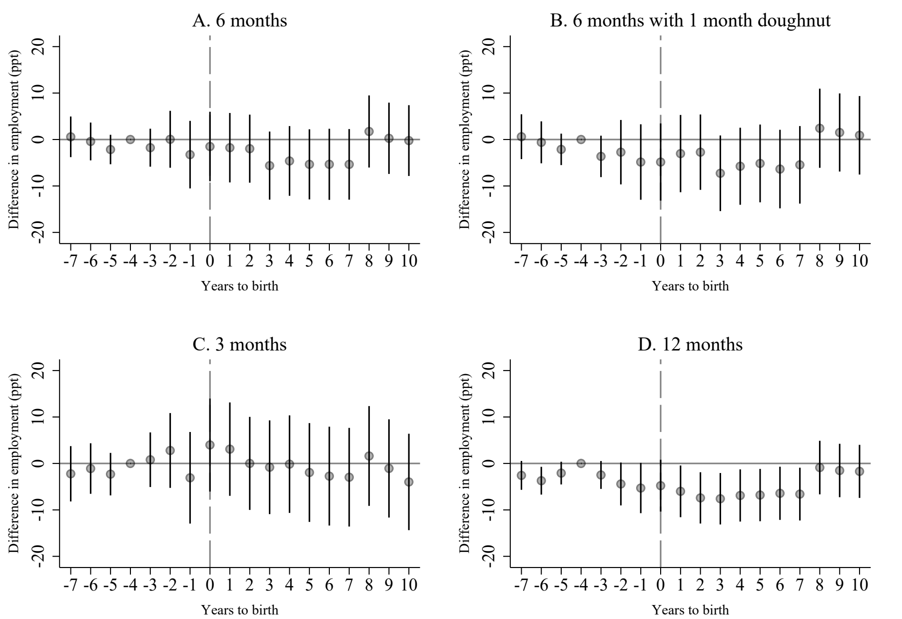

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Figure I.A.7: Dynamic DiD results by treatment window: Employment: Fathers

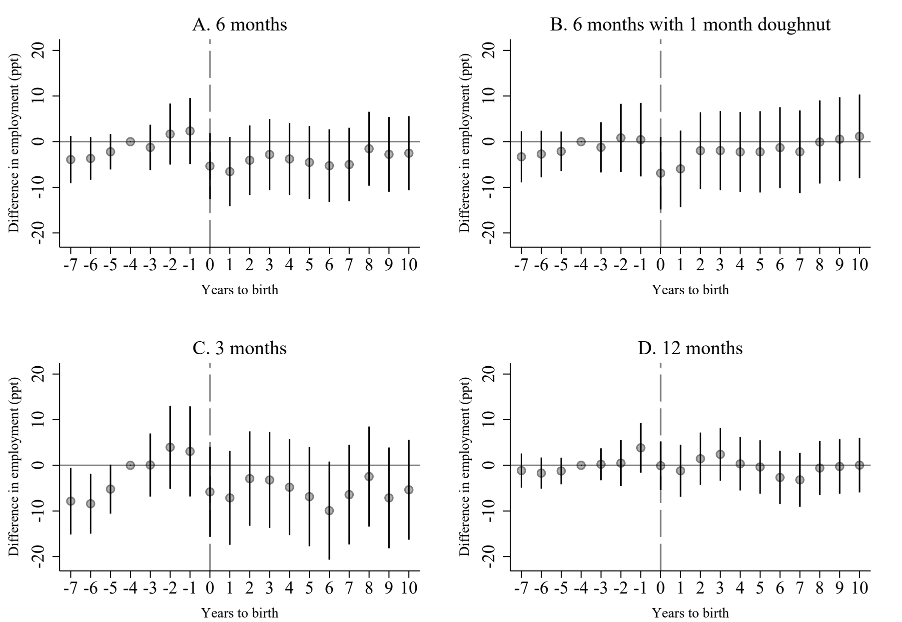

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4.

Figure I.A.8: Dynamic DiD results by treatment window: Earnings: Mothers

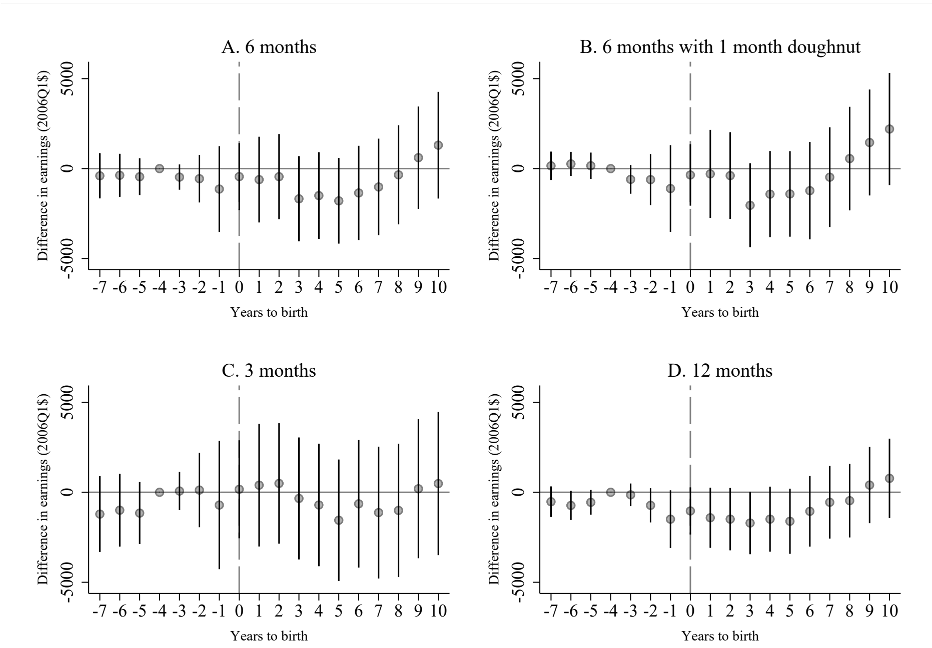

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4. Earnings measured in 2006Q1$.

Figure I.A.9: Dynamic DiD results by treatment window: Earnings: Fathers

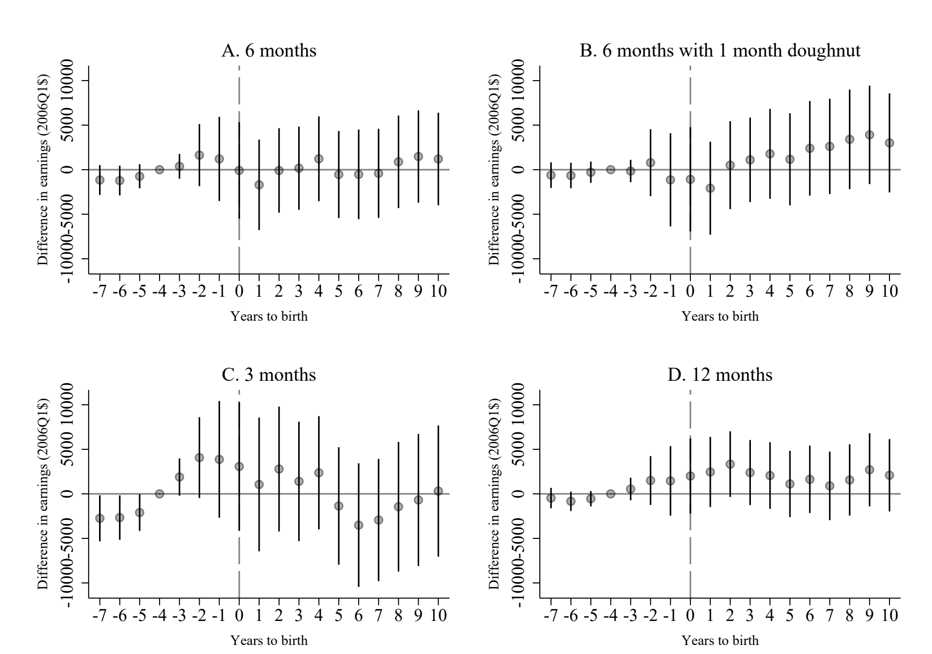

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is -4. Earnings measured in 2006Q1$.

Figure I.A.10: Dynamic DiD results by treatment window: Accident claim rate

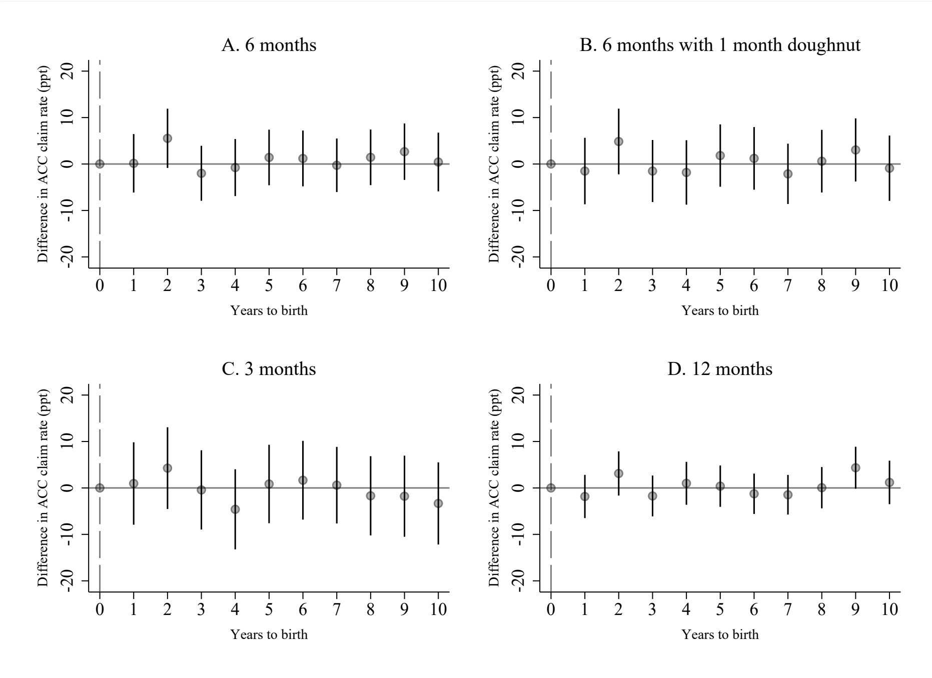

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Figure I.A.11: Dynamic DiD results by treatment window: No. of accident claims

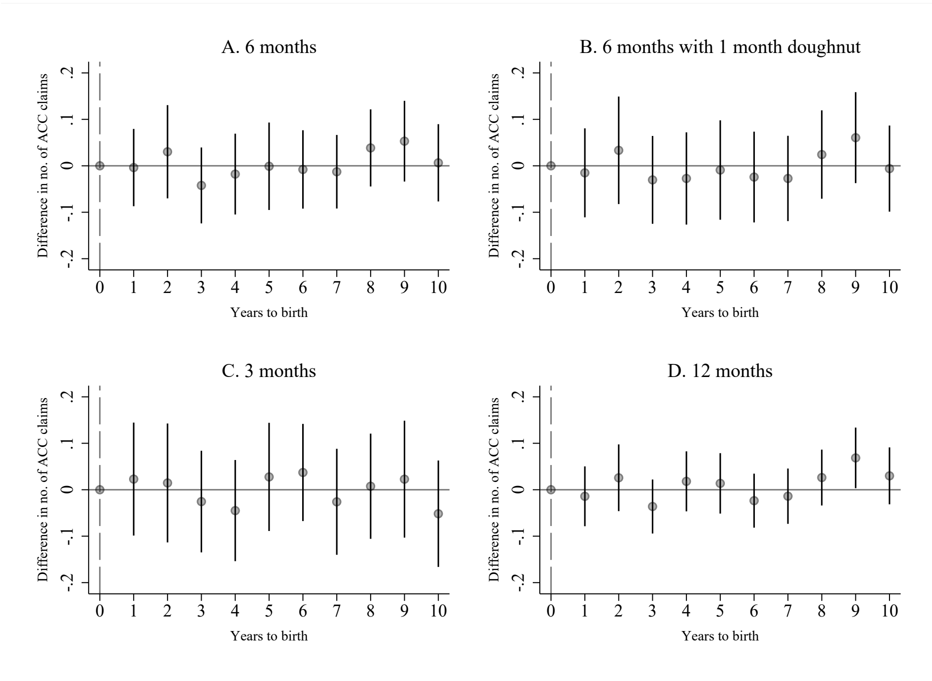

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Figure I.A.12: Dynamic DiD results by treatment window: Hospital admission rate

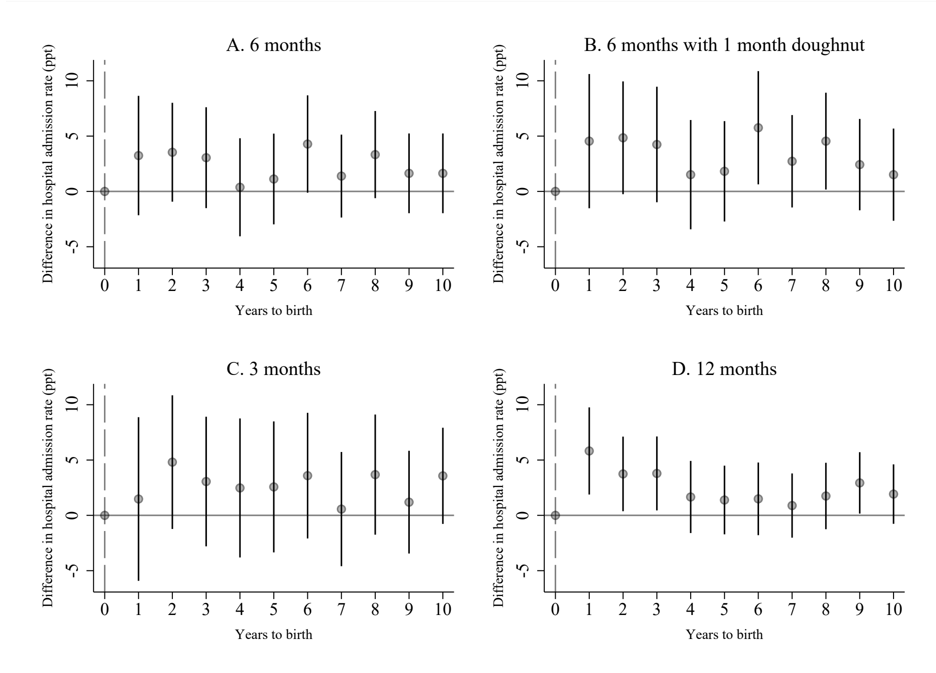

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Figure I.A.13: Dynamic DiD results by treatment window: No. of hospital admissions

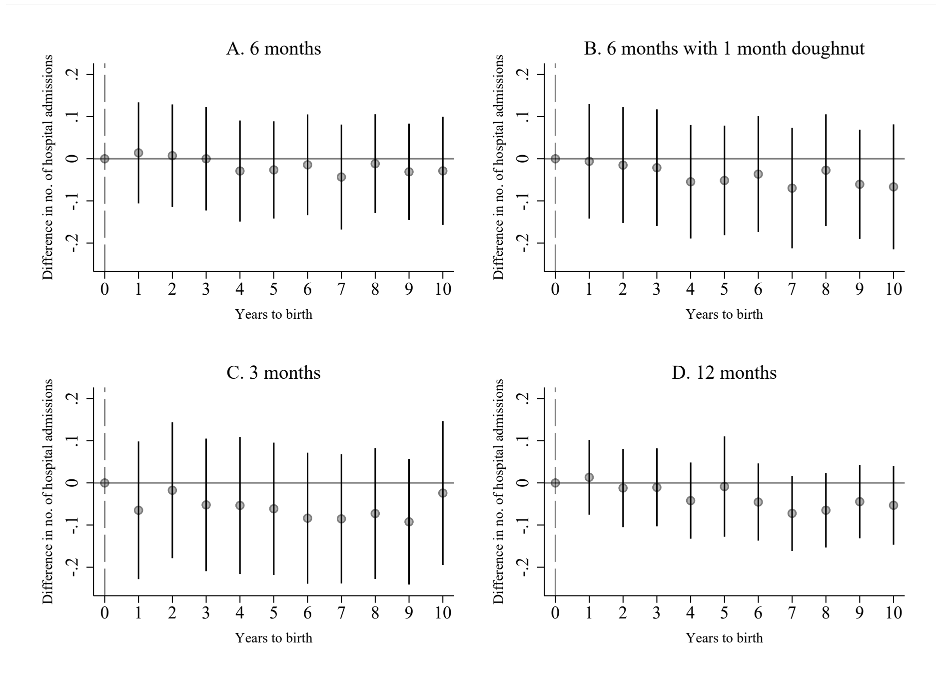

Notes: BPresents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Figure I.A.14: Dynamic DiD results by treatment window: Prescription rate

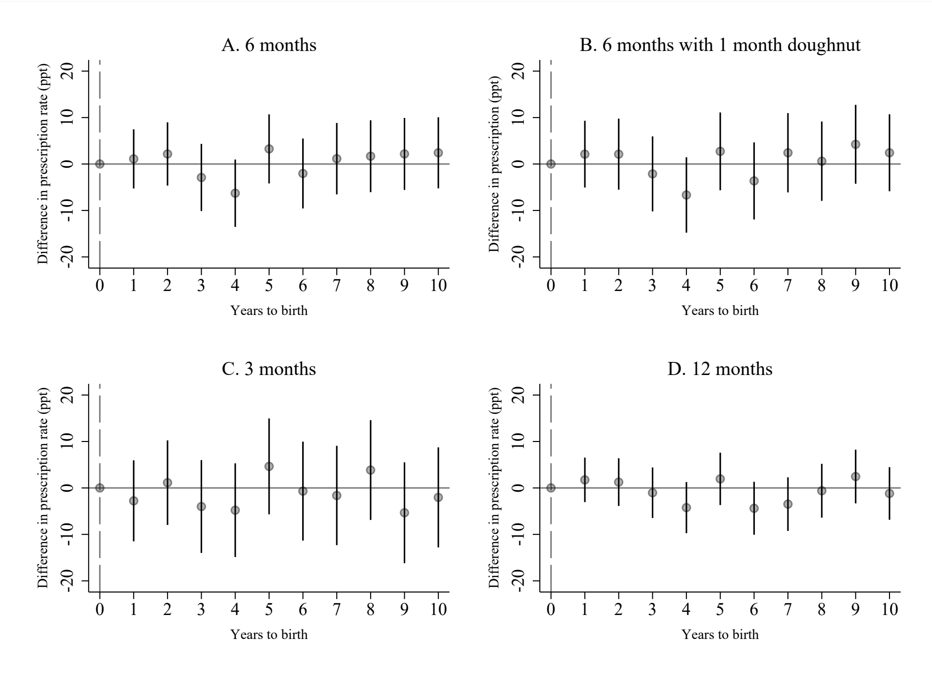

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Figure I.A.15: Dynamic DiD results by treatment window: No. of prescriptions

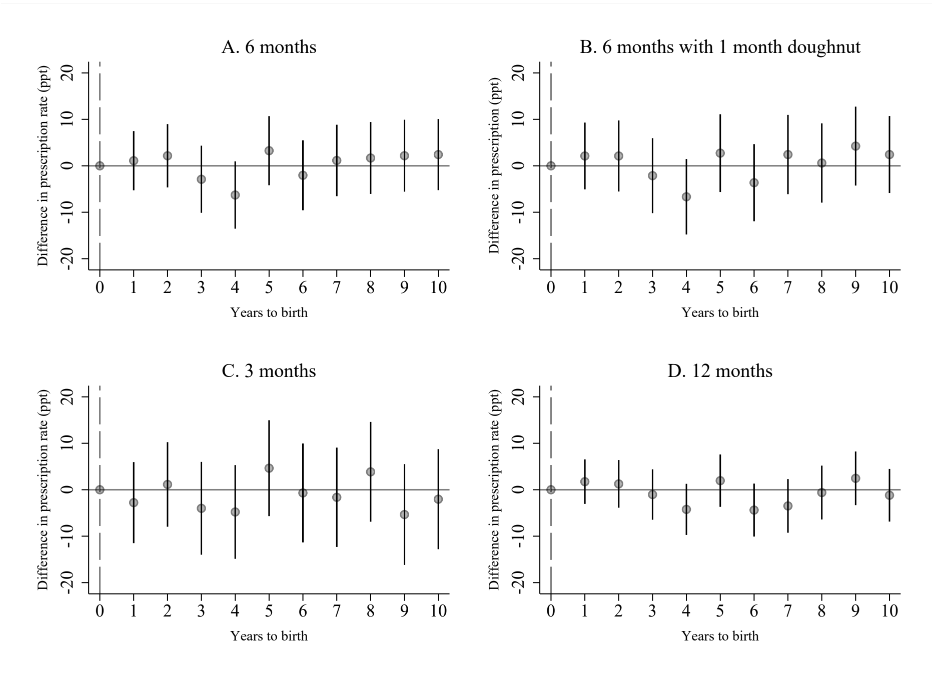

Notes: Presents the coefficient of interest (γ_e) from Equation 2 with a six-month treatment window and demographic and month-of-birth controls. Vertical bars represent 95% confidence intervals. The base year is 0.

Paper II

Workforce vaccine mandates: The effect on vaccine uptake and healthcare workers' labour market outcomes

Authors: Lisa Meehan, Livvy Mitchell and Gail Pacheco

Prelude

Like Paper 1, this paper uses difference-in-differences (DiD) estimation applied to data from Stats NZ's Integrated Data Infrastructure. While it uses some of the same data sources as Paper 1, including Inland Revenue tax data to measure labour market outcomes, some of the data used are different, such as the national immunisation register. Moreover, while Paper 1 used a policy cut-off date to create treatment and comparison groups, this paper uses a triple difference estimator to measure the effect of workforce vaccine mandates on healthcare workers' labour market outcomes.

This research is possible due to the linking of employment records from tax data and vaccination data from a comprehensive national immunisation register. In order to identify the population of interest and treatment and control groups, it uses data on individuals' employers and industry of employment from Inland Revenue tax data. In order to identify whether a worker is vaccinated, it uses data from Ministry of Health's national immunisation register. Workers' employment and earnings outcomes are measured using Inland Revenue tax data. Further covariate information is derived from various parts of the IDI, such as Stats NZ's personal details and address notifications tables, and Department of Internal Affairs birth register.

In order to examine the effect of workforce mandates on vaccination uptake, DiD estimations are used to compare mandated workers with those who were not mandated. The effect on healthcare workers' labour market outcomes is examined via triple difference estimators comparing unvaccinated and vaccinated healthcare workers with unvaccinated and vaccinated workers who did not face workforce mandates.

Abstract

As part of its COVID-19 policy response, the New Zealand government implemented vaccination mandates as a condition of ongoing employment for certain workers. This paper examines the effect of these mandates on vaccination uptake among mandated healthcare, education and corrections workers and on healthcare workers' labour market outcomes. This is enabled by New Zealand's linked population-wide administrative data, which includes a comprehensive national vaccination register linked to tax records to identify employment outcomes.

Overall, the results suggest that in the context of already-high vaccination rates, workforce vaccine mandates provided limited benefit in terms of increasing vaccination rates among mandated workers. Moreover, they negatively impacted healthcare workers' labour market outcomes, which may have had wider consequences in terms of exacerbating existing health workforce skills shortages.

JEL: C23, I12, I18
Keywords: COVID-19; vaccination; workforce mandate; employment; earnings

1 Introduction

As part of its COVID-19 policy response, the New Zealand (NZ) government implemented workforce vaccine mandates. These required certain types of workers, including health and disability, education, border and managed isolation, fire and emergency, police, defence and corrections staff, to be vaccinated in order to continue their employment. This paper examines the impact of these workforce vaccine mandates on the uptake of COVID-19 vaccinations among education, corrections and healthcare workers (HCWs)(RQ1: Research Question 1), and on the labour market outcomes (employment rates and earnings) of HCWs (RQ2: Research Question 2).

This paper adds to the very limited international evidence on the effect of workforce mandates on vaccination rates and labour market outcomes. Currently, there are only a handful of studies examining the effect of COVID-19 vaccine mandates on vaccine uptake, and most of these either study vaccine pass mandates (where proof of vaccination was needed to access non-essential services) as opposed to workforce mandates, or are limited to examining US nursing home staff mandates, which is only a subset of the wider health workforce. Moreover, there are few existing studies that use individual-level data, with almost all using state-/province-level data or nursing-home-level data. In addition, almost all of the existing studies using individual-level data involves self-reported survey information on vaccination status, rather than detailed administrative records. Indeed, there appears to be only one other study using individual-level administrative data to look at COVID-19 vaccine workforce mandates (namely Rubenstein et al., 2023, which examines the effect of New York City municipal employee mandates on vaccine uptake). There is even less evidence on the effect of COVID-19 vaccine mandates on HCW labour market outcomes, with only a couple of studies examining US nursing home staffing.

The NZ experience, therefore, offers a natural experiment of the effects of stringently applied and enforced nationwide workforce vaccine mandates on vaccine uptake and labour market outcomes. This analysis is enabled by the existence of a comprehensive, population-wide vaccination database that has details of the type of vaccine received, the number of the dose administered, and the exact date the dose was received. This database is linked to various other administrative data sources, including employment details from tax records, which allow HCWs' employment outcomes to be tracked. As such, the NZ experience offers a unique opportunity to examine the effect of vaccine mandates and provide an evidence base to inform their use in future pandemic planning.

Given the value of COVID-19 vaccines in preventing severe illness and death (Tenforde et al., 2022), NZ was one of a number of countries that either implemented, or attempted to implement, vaccine mandates. The international experience highlights how controversial these policies are, and the difficulties inherent in making trade-offs between public health considerations and the impingement on individual rights and the risk of eroding trust in government and scientific institutions. Indeed, mandates can entrench distrust and provoke reactance (a motivation to counter a threat to one's freedom) (Bardosh et al., 2022; Sprengholz et al., 2021, 2022). This can potentially strengthen anti-vaccine sentiment generally and reduce acceptance not just of COVID-19 vaccinations (Schmelz & Bowles, 2022), but also of other vaccines (Dubé et al., 2021). In addition, there were concerns that HCW mandates would further exacerbate staff shortages. This is reflected in the degree of opposition to these policies, which resulted in mandates being abandoned in some countries (e.g. the UK), and/or facing legal challenges in others (e.g. the US and NZ). These complex ethical considerations and the resulting level of controversy surrounding these policies further heightens the importance of having a sound evidence base on the their effectiveness.

A difference-in-differences (DiD) approach comparing vaccine uptake and labour market outcomes among workers subject to the mandates and those who were not subject to the mandates allows the effect of workforce vaccine mandates to be separated from other initiatives aimed at boosting vaccination uptake. For example, NZ also implemented population-wide initiatives such as vaccine passes, whereby proof of vaccination was needed to access non-essential businesses. This differentiation is important from a policy perspective since vaccine passes are a "softer" mandate which potentially restrict access to non-essential businesses while workforce mandates are a "harder" mandate which potentially prevent someone from earning a living in their chosen profession. This is also particularly relevant in the context of professions with skill shortages, such as healthcare, where workforce mandates can further contribute to these shortages and hinder timely delivery of health services.

2 Background and policy context

The first case of COVID-19 was reported in NZ on 28 February 2020. In response, the government implemented a zero-COVID elimination strategy. While the specific measures in place to achieve this changed over time, the main measures used included strict lockdowns, closing the border to foreign nationals and imposing a period of managed isolated for those entering the country.

Figure II.1 shows the stringency of NZ's policy response compared with the average for OECD countries. The spikes in the NZ series correspond to lockdowns, which involved the closure of non-essential businesses and services (including schools), strict restrictions on regional travel and the requirement to remain home except for essential travel (e.g. supermarket shopping, medical appointments etc.) or essential work (e.g. HCWs, supermarket workers etc.). These lockdowns occurred whenever cases of COVID-19 in the community were detected and were either nationwide or limited to specific regions where cases were detected. Although the specifics of what was permitted during lockdowns depended on the extent of community transmission, during the strictest lockdowns (officially known as Alert Level 4), NZ had the most stringent COVID-19 policy response in the world (Gibson, 2022b, 2022c). Figure II.1 also shows that NZ's Stringency Index remained high even when restrictions began to ease in other countries. This is because NZ pursued an elimination strategy for an extended period of time, with corresponding policy responses including lockdowns.

Figure II.1: COVID-19 policy response Stringency Index: NZ versus OECD

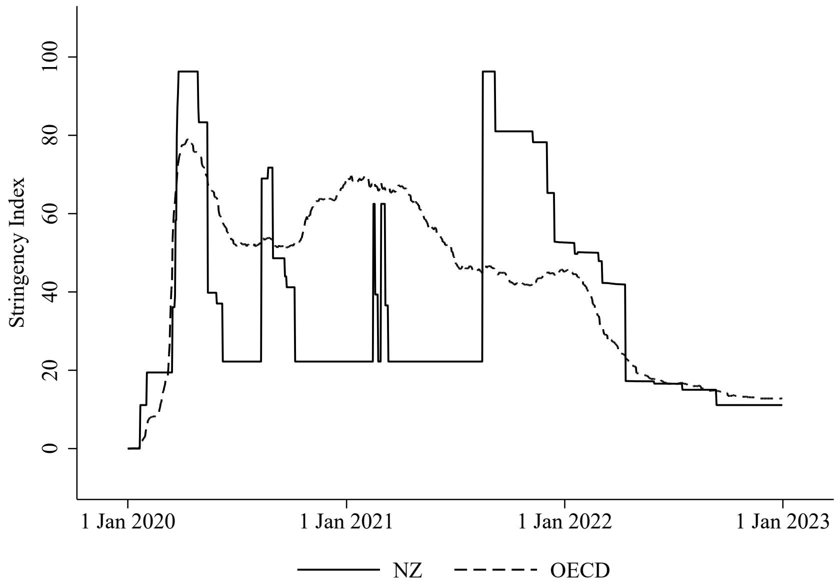

Note: OECD is a simple average of OECD countries with available data.
Source: Hale et al. (2021). Data accessed from https://github.com/OxCGRT/covid-policy-tracker on 27 January 2023

The first batch of COVID-19 vaccines arrived in NZ in February 2021. With the availability of vaccines, several additional vaccination-related policy measures were introduced. First, the government implemented a nationwide vaccine roll-out centered on the Pfizer vaccine. Due to supply issues, vaccines were initially offered to groups based on priority. Vaccinations were first available to vaccinators, and managed isolation workers and those they lived with, followed soon after by frontline HCWs. From March 2021, vaccine availability was extended to those most at risk of getting COVID-19 or developing serious illness as a result, including those aged 65 and over and those with underlying health conditions. Vaccines were then rolled-out to the general population in age groups. Those over 45 years were invited to get vaccinations from August 2021, those over 35 from September 2021 and everyone else was eligible from October

2021. However, when COVID-19 cases started to rise in the second-half of 2021, the roll-out proceeded slightly quicker than planned and vaccines were available to the entire population aged 12 and over from September 2021. Nonetheless, NZ's vaccine roll-out was initially slower than other countries, which is reflected in Figure II.2 showing the percentage of people who were fully vaccinated over time compared with the OECD average. NZ's resultant vaccination rate was, however, relatively high likely reflecting a reasonably high willingness to comply with government recommendations, in addition to initiatives to encourage vaccination such as vaccine passes to access non-essential services, advertising campaigns, vaccination rate targets to come out of lockdowns etc.[7]

Figure II.2: Share of population fully vaccinated: NZ versus OECD

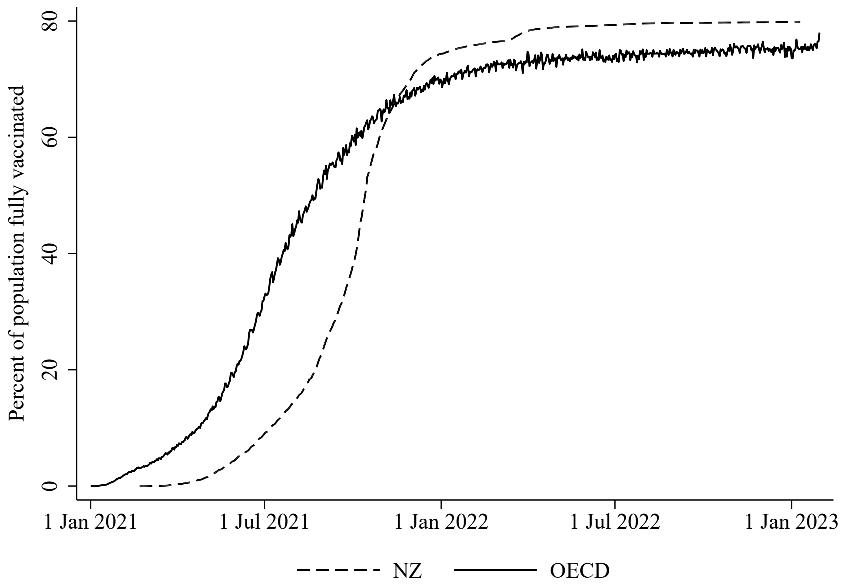

Notes: OECD is a simple average of OECD countries with available data.

Source: Our World in Data COVID-19 database. Accessed from https://github.com/owid/covid-19-data/tree/master/public/data on 1 February 2023

In April 2021, the government announced vaccinations would be mandatory for workers in managed isolation and quarantine (MIQ) facilities from 1 May 2021. In July 2021, mandatory vaccinations were extended to port and airport workers. These vaccine mandates involved a relatively small number of workers. In October 2021, mandates were extended to a large number of workers, including teachers, HCWs, corrections prison workers, frontline fire and emergency service workers, police and defence force personnel. This extension of the number of workers covered by the

[7]Due to data availability, Figure II.2 shows the share of the total population who were fully vaccinated rather than the share of the total eligible population, or total adult population. But, if anything, the use of total population rather than total eligible population as the denominator biases down NZ's vaccination rates since NZ's population is relatively young compared with the average OECD country.

mandates appears to be unexpected, with the government having earlier publicly ruled out the possibility of vaccine mandates, and with no media coverage of the possibility prior to the announcement in October. Deadlines were set for first and second vaccine doses, with staff who did not comply losing their jobs. Exemptions were only granted on medical grounds, and these had strict conditions and were administered centrally by the Ministry of Health. These were much narrower grounds for exemptions than, for example, the US, where COVID-19 vaccinations were mandated for HCWs employed by Medicare and/or Medicaid-accepting facilities, but exemptions were allowed on either medical or religious grounds (Rao et al., 2022).

Starting in October 2021, the government's COVID-19 response started to shift focus to a management rather than an elimination strategy, with vaccinations being a centrepiece of this strategy. Coming off the back of a nationwide lockdown in August to September 2021, and an ongoing lockdown in the largest city of Auckland, the government announced plans in October to implement a vaccine pass from early December 2021. This was a population-wide initiative, with anyone aged 12 or over requiring a vaccine pass to access public venues and non-essential businesses. Given this context, an important question is to what extent the workforce mandates increased vaccine uptake over-and-above other "softer" initiatives, particularly the population-wide vaccine passes.

While the stated purpose of the vaccine mandates according to the original COVID-19 Public Health Response (Specified Work Vaccinations) Order 2021 was to "prevent, and limit the risk of, the outbreak or spread of COVID-19" (NZ Government, 2021), the government later amended the Order's purpose to "ensure continuity of services that are essential for public safety, national defence, or crisis response". It is this latter purpose that was central to the legal challenge to the mandates and the court ruling that mandates for police and defence force personnel were unlawful (*Yardley v Minister for Workplace Relations and Safety*, 2022).

In December 2021, the government formally shifted away from an elimination strategy and a new protection framework for managing COVID-19 was introduced. Over time, many measures such as border closures, managed isolation, and vaccine passes were rolled back. Workforce mandates started to be removed from April 2022, with the last mandates, including for HCWs, ending in September 2022. Workers who had lost their jobs due to non-compliance with the mandates were not entitled to reinstatement once these mandates were removed.

3 Literature

The first research question of this paper - whether mandates increase vaccination rates - relates to a substantial literature on the economics of infectious diseases and vaccinations. Some studies in this area look at the relationship between the prevalence of an infectious disease and vaccination rates, and suggest that people are responsive to disease prevalence. For the US, Philipson (1996) finds that the prevalence of measles reduces the age at which the first measles vaccination occurs, and Oster (2018) and Schaller et al. (2017) find pertussis (whooping cough) outbreaks increase vaccination uptake. For Austria, Schober (2020) finds that measles outbreaks increase measles vaccine uptake. This literature also finds people respond to information. In particular, studies have found that MMR (measles, mumps, and rubella) vaccination uptake among children with highly educated mothers decreased in response to the controversial (and later retracted) study linking the MMR vaccine to autism (e.g. Anderberg et al., 2011).

This study is most related to quasi-experimental analyses in economics that examine vaccine mandates. Most existing evidence on the effectiveness of vaccine mandates relates to childhood immunisation as a condition for childcare or school entry. It mostly focuses on the US, which has a long history of using school-based mandatory vaccination laws to increase vaccination rates. For example, Carpenter and Lawler (2019) exploits the variation in the timing of mandate adoption across US states and applies a difference-in-differences methodology to the 2008-2013 waves of the National Immunization Survey-Teen. It finds strong evidence that Tdap (tetanus, diphtheria, and pertussis) vaccine mandates for middle school entry increase the uptake of the Tdap vaccine, and also have spill-over effects in raising vaccination rates of other, non-mandated vaccinations, such as the influenza vaccine. Abrevaya and Mulligan (2011) use data from the 1996-2006 National Immunization Survey (NIS) to examine daycare- and school-entry varicella (chickenpox) vaccine mandates in the US and find they increase immunisation rates. Lawler (2017) examines mandatory childcare-entry vaccinations versus non-binding recommendations to vaccine for hepatitis A and finds that recommendations increase vaccination rates among young children by 20 percentage points, while mandates increase rates by a further 8 percentage points. Moreover, recommendations only increase the probability that individuals will start the course of vaccinations, while mandates are effective at inducing them to complete the course. While the medical and public health literature examining mandates and childhood immunisations generally use data with less coverage, these also tend to find that school-entry mandates increase uptake (for a review, see Lee & Robinson, 2016).

There is limited existing evidence on the effectiveness of mandates in lifting vaccination rates in adults and outside of the US, and even less specifically focused on COVID-19 (Mello et al., 2022). Lindley et al. (2019) examines healthcare facility influenza vaccination mandates in the US, whereby employers implemented vaccines as a condition of employment, and assessed whether their effect differs depending on whether there are also state laws encouraging or mandating vaccinations. It finds that facility-level mandates increase influenza vaccination rates, with the increase being larger in states that have no or weaker laws. Carrera et al. (2021) compares US states that did and did not implement laws encouraging or mandating influenza vaccinations for hospital workers and finds that these laws reduce pneumonia and influenza mortality rates among the general population. Although it did not examine the first-order effect of whether it increased vaccination rates among hospital workers, the presence of the second-order effect of reducing mortality suggests a first-order effect also occurred.

In terms of examinations of COVID-19 vaccine mandates, there are only a handful of studies, and all but one of these either study vaccine pass mandates (where proof-of-vaccination was needed to access non-essential services) as opposed to workforce mandates, or are limited to examining US nursing home staff mandates, with no existing evidence for other countries and/or the wider healthcare workforce. In terms of the vaccine pass mandate literature, Karaivanov et al. (2022) exploits the variation in timing of these measures across Canadian provinces to apply a difference-in-differences approach. It finds that the announcement of a mandate led to a surge in new vaccinations (a more than 60 % increase in weekly first doses). It also undertakes time-series analysis for each province and for France, Italy and Germany which corroborates this finding. Another paper using a synthetic control model by comparing six countries that introduced vaccine passes with 19 control countries, finds that vaccine passes led to an increase in vaccinations (Mills & Rüttenauer, 2022).

There are also several studies which examine the effect of COVID-19 vaccine passes by creating a synthetic comparison group of countries without vaccine passes. Mills and Rüttenauer (2022) compares six countries that introduced vaccine passes with 19 control countries and finds that vaccine passes led to an increase in vaccinations (Mills & Rüttenauer, 2022). Similarly, Oliu-Barton et al. (2022) finds that the introduction of vaccine passes in France, Germany and Italy led to an increase in vaccine uptake and a decrease in hospitalisations and deaths. Comparing Lithuania, which required a vaccine pass to access certain businesses and events, with Poland, which introduced a vaccine pass but did not impose any such restrictions (it was used as a tool for international travel only), Walkowiak et al. (2021) finds that Lituania had markedly higher vaccination rates than Poland.

The only study to use a synthetic-comparison-group methodology to examine the use of workforce vaccine mandates that we are aware of is Cohn et al. (2022), which compares New York City (NYC) to other similar US counties to examine the effect of a policy package that included workforce vaccine mandates for municipal employees along with vaccination passes and vaccine incentive payments. It finds that these measures increased vaccination rates in NYC relative to other US counties. However, the effect of the vaccine mandates cannot be separated from the other policies to encourage vaccination.

In terms of the literature on US nursing home staff vaccine mandates, Syme et al. (2022) examines COVID-19 vaccine mandates with a test-out exemption for Mississippi nursing home staff. It finds that compared with surrounding states without mandates, the vaccination rates among Mississippi nursing home staff increased more, but that the gains were minimal. However, this study only covered nursing home staff (not HCWs more generally), and examined a less stringent policy that allowed workers to return negative COVID-19 tests twice a week instead of vaccinating. McGarry et al. (2022) also examines nursing home staff and finds that COVID-19 vaccine rates increased the most in states with a mandate and no test-out option, followed by states with a mandate and a test-out option, and least in states with no mandate. Plummer and Wempe (2023) finds that following the US Supreme Court's upholding of the federal COVID-19 vaccine mandate for HCWs in Medicare- and Medicaid-eligible nursing homes, vaccination rates among nursing home staff increased more in states that did not have state-level vaccine mandates than in states that did have mandates.

As far as we are aware, Rubenstein et al. (2023) is the only study that examines workforce COVID-19 vaccine mandates more broadly than just nursing home staff. It examines vaccine mandates for NYC municipal workers. An announcement was made in July 2021 that all NYC municipal employees would either need to be vaccinated against COVID-19 or return a negative test each week from September 13 2021. On 20 October 2021, it was announced that the test-out option would be removed from 1 November 2021, and all unvaccinated employees would be placed on unpaid leave from that date and eventually subject to termination. Comparing mandated municipal workers with all other working-age NYC residents, the study finds that the mandates did not increase vaccine uptake when the test-out option was available. However, uptake increased once the test-out option was removed. The comparison of NYC municipal workers and all other NYC working-age residents may, however, present issues since the characteristics of the two groups are not similar, with the comparison group having a much higher share of young people, men, Asians and Hispanics. Indeed, it is found that parallel trends does not hold as the vaccine uptake among the comparison group increased at a faster pace than the municipal employees group in the pre-treatment period. Moreover, a potentially larger issue with the apparent faster uptake of vaccinations after the test-out option was removed is that the population of municipal employees changed over time. Thus, as the paper acknowledges, the apparent faster increase in vaccinations among municipal employees could have been due to unvaccinated employees leaving their jobs and, therefore, no longer being counted in the treatment group, and because any new municipal employees would have been required to be vaccinated. Given unvaccinated workers were placed on unpaid leave and faced termination and any new employees needed to prove their vaccination status, it is difficult to know whether the stricter mandate without the test-out option led to an increase in vaccination uptake or merely changed the composition of the treatment group.

While staffing shortages are an important potential unintended consequence of COVID-19 vaccine mandates, there is even less evidence on worker labour market outcomes than on the effect of mandates on vaccine uptake. McGarry et al. (2022) finds no evidence of increased nursing home staffing shortages in states with COVID-19 vaccine mandates. However, this finding is based on self-reported facility-level nursing home staff shortages, which is subject to possible misrepresentation (Plummer & Wempe, 2023). Plummer and Wempe (2023), which examines the federal COVID-19 vaccine mandate, appears to be the most related to our study. It used nursing home payroll data to measure staffing levels (e.g. staff hours per resident per day) rather than self-reported staff shortage data due to concerns about possible misrepresentation. It finds that the mandates did not have a material impact on staffing levels. Once again, these studies are limited to nursing home staff rather than the wider healthcare workforce.

In terms of NZ evidence, we are not aware of any studies using quasi-experimental methods to examine the impact of vaccine mandates. However, there is a small economics literature examining the costs versus benefits of vaccine mandates. Lally (2021) undertakes cost-benefit analyses and finds that the costs of vaccine mandates for the general population are likely to far outweigh the benefits. However, because HCWs are more likely to come into frequent and close contact with sick people, the benefits of vaccine mandates *may* outweigh the costs for these workers (emphasis in original). Education workers are less likely to come into contact with people at high risk from COVID-19, and therefore, Lally (2021) finds that the costs of mandates outweigh the benefits for these workers.

Given the limited existing evidence on the impact of COVID-19 vaccine mandates on HCW vaccine uptake and workforce labour market outcomes, this paper makes a significant contribution to the literature. The existing quasi-experimental literature on mandates measures vaccination status using survey data for the most part, with some exceptions such as Karaivanov et al. (2022), which uses Canadian provincial data (rather than individual-level data), and the aforementioned studies on US nursing homes (McGarry et al., 2022; Plummer & Wempe, 2023), which use weekly data on staff vaccination rates at the nursing-home-facility level. The lack of studies using individual-level data stems from the fact that most existing analysis is for the US, where there is a lack of immunisation registries (Abrevaya & Mulligan, 2011). It seems that the only study to use individual-level data is Rubenstein et al. (2023), which examines NYC municipal employee mandates. However, it appears that the authors used aggregated information on individual vaccine records rather than having access to individual-level records as they state that they could not look at changes to municipal employee staffing due to only having access to aggregated data. In contrast, NZ's Ministry of Health has a comprehensive, population-wide vaccination database which includes information on all COVID-19 vaccine doses administered in NZ, as well as records of COVID-19 vaccines received overseas.[8] Anonymised individual-level vaccine records are available to researchers and linked to a rich set of data on individuals' characteristics via Stats NZ's Integrated Data Infrastructure (IDI), including employment records via tax data.

Moreover, improving the evidence base in this area is particularly important from a policy perspective given the problem of skill shortages in healthcare, and the possibility that mandates can further contribute to these shortages and hinder the timely delivery of essential health services. For NZ, this is even more pertinent given the stated purpose of the mandates was to ensure continuity of public services. Understanding the impact of mandates is also particularly relevant given the degree of controversy surrounding mandates. This is reflected in the qualitative literature, which suggests there is limited support for COVID-19 vaccine mandates among HCWs. For example, Woolf et al. (2022) finds that only 18% of surveyed UK HCWs favoured mandatory COVID-19 vaccination. Even more strikingly, a German survey found that few respondents were opposed to being vaccinated against COVID-19 if vaccinations were encouraged but voluntary (3.3%), but a much higher share were opposed to being vaccinated if vaccinations were mandatory (16.5%) (Schmelz & Bowles, 2022). This highlights the potential issue of reactance (Bardosh et al., 2022; Sprengholz et al., 2021, 2022), which can strengthen anti-vaccine sentiment generally (Schmelz & Bowles, 2022). In addition, there have been legal challenges to mandates. While the US Supreme Court upheld the federal vaccine mandate, these legal challenges were partially successful in NZ. The NZ courts ruled that the mandates were an unjustifiable limitation on the right to refuse medical treatment in the case of defence and police staff (although the ruling was made after the mandates for these workers had been lifted), but that

[8] During the period which the mandates were in place, NZ also required those entering the country to prove their COVID-19 vaccination status. Thus the COVID-19 vaccination register includes reliable information on vaccinations administrated overseas.

this limitation was justified in the case of HCWs and teachers. In the UK, vaccine mandates for NHS staff were set to come into force, but were abandoned amidst implementation issues (such as difficulties confirming the vaccination status of staff) and concerns about the loss of key staff (McKee & Schalkwyk, 2022).

Another consideration is that COVID-19 vaccinations, when the mandates were implemented and at the time of writing, reduced the severity of the illness but had limited efficacy in terms of reducing transmission (Gur-Arie et al., 2023). Thus, the positive externality argument for COVID-19 vaccine mandates was limited. Indeed, the NZ court ruling that vaccine mandates were unlawful in the case of police and defence force personnel highlighted that the Pfizer vaccine (which NZ's vaccine strategy was centered on) was not particularly effective at reducing transmission in the case of the main COVID-19 variants that existed at the time (i.e. Delta and Omnicron variants) (*Yardley v Minister for Workplace Relations and Safety*, 2022).

4 Data

We use population-wide linked administrative data from Stats NZ's Integrated Data Infrastructure (IDI). Our main data sources are the Inland Revenue Department's (IRD) Employer Monthly Schedule (EMS) data, which allows us to identify which individuals worked in sectors subject to COVID-19 vaccination mandates, and the Ministry of Health's COVID-19 vaccination register, which allows us to identify vaccination uptake and mandate compliance.

We first define our population of interest for evaluating the impact of vaccine mandates on vaccination uptake (RQ1), including how we identify those subject to the vaccine mandates and those who complied with the mandates. We then define the population of interest for evaluating the impact of vaccine mandates the labour market outcomes of HCWs (RQ2). Lastly, we define a range of demographic and socioeconomic descriptive variables used for both RQ1 (vaccine uptake) and RQ2 (HCW labour market outcomes) analyses.

4.1 RQ1 vaccine uptake

4.1.1 Creating a sample of employed individuals

To study the impact of COVID-19 vaccination mandates on vaccination uptake, we begin by identifying a cohort of individuals who were employed before the COVID-19 vaccines were largely available to the public (the first vaccine was administered in NZ on 19 February 2021, and vaccinations were initially limited to vaccinators and then MIQ staff and their families) and before any announcements of potential COVID-

19 vaccine mandates were made. Identifying a specific cohort avoids the issue of vaccination rates among mandated workers changing due to changes in the workforce composition over time (particularly given that unvaccinated workers were required to leave vaccine-mandated roles), which, as mentioned in the context of Rubenstein et al. (2023)'s analysis, would lead to an overestimate of the effect of mandates on vaccine uptake. Specifically, we use the personal details and IRD EMS data in the IDI to identify all working-aged people (aged 16-60) with positive earnings (in terms of receiving positive wages and salaries) in March 2021. March 2021 is chosen as the reference point not only because it is prior to widespread vaccine access and vaccination mandate announcements, but also because it is the end of the financial year in NZ. This gives 2,083,155 individual-job observations.

Next, to account for multiple job holdings, we observe each individual in their main job. A person's main job is defined as the job with the main ("M"-type) tax code in the EMS data. If there are multiple jobs with M codes, their main job is defined as the one with the highest earnings in the reference month of March 2021. This gives 1,946,859 individual observations.

The EMS data provide the Australia and New Zealand Standard Industry Classification 2006 (ANZSIC06) code relating to each employer-employee relationship, allowing us to identify the industry each individual is employed in. Since we require this industry information to identify individuals who were subject to COVID-19 vaccination mandates, we exclude those for whom their main job is missing an ANZSIC06 industry code. This leaves us with 1,941,942 individual observations with main job industry information.

We exclude those who died during our study period, leaving 1,940,370 individuals. Note that there were almost no COVID-related deaths in NZ so there are no selection issues caused by COVID deaths. Finally, we drop a very small number of individuals with dubious COVID-19 vaccination records (potentially driven by measurement error); for example, those with their first dose date after their second dose date, or those with missing first and second dose records but an 'additional dose' record. The resulting sample comprises 1,940,115 individuals.

4.1.2 Identifying individuals subject to COVID-19 vaccination mandates

Through examining iterations of the COVID-19 Public Health Response (Vaccinations) Order 2021 ('the Order'), and media releases on the the NZ Government's official website (beehive.co.nz) and Ministry of Health website, we are able to identify the announcement dates and commencement dates for each group of individuals who became subject to the COVID-19 vaccination mandates.

In Schedule 2 of the Order, "groups of affected persons" are defined in 10 'parts'.

These parts describe categories of work. For example, Part 7 is "Groups in relation to health and disability sector". We match the work description in each 'part' to a 7-digit ANZSIC06 industry code. We then categorise each of the 500+ 7-digit ANZSIC06 industries into one of the following three categories:

1. Industries barely covered by COVID-19 vaccination mandates
2. Industries partially covered by COVID-19 vaccination mandates
3. Industries heavily covered by COVID-19 vaccination mandates

Consequently, we observe which of the above categories each employed individual fell into, according to the ANZSIC06 industry code relating to their main job in March 2021.

The classification into 'barely', 'partially' and 'heavily' industries warrants further explanation. As the administrative data does not include information on an individual's occupation, we focus on industry of employment. This means that, for example, an administrator employed by a hospital would be classified as a health sector worker. In some cases, industries align well with the categories of work defined in the 'parts' of the Order, including in the case of HCWs since the definition was much broader than health practitioners (discussed below). In some cases, it only partially aligns. For example, MIQ workers were mandated by the Order under *Part 1 - Groups in relation to managed quarantine facilities*. However, MIQ facilities were hotels in NZ, and employment is therefore identified by the *H440000 Accommodation* ANZSIC06 industry code, which also includes all other hotel workers not employed at MIQ facilities. Thus, mandated hotel workers cannot be differentiated from non-mandated workers using industry classification codes, and we categorise *H440000 Accommodation* as a partially mandated industry.

As mentioned, fortunately the HCW category aligns well with Part 7 *Groups in relation to health and disability sector* which covered not only frontline HCWs, but also care workers (such as those in aged-care facilities), and workers whose role involved being within two metres of health practitioners or members of the public.[9] Thus, it aligns well with the ANZSIC06 industry code classification of *Q84 Hospitals*, *Q85 Medical and Other Health Care Services* and *Q85 Residential Care Services*. Note, however, that the use of ANZSIC06 industry codes means we cannot include workers who fall under the Order definition but are not employed in the health industry. For

[9]Specifically 1. Health practitioners; 2. Workers who carry out work where health services are provided to members of the public by one or more health practitioners and whose role involves being within 2 metres or less of a health practitioner or a member of the public for a period of 15 minutes or more; 3. Workers who are employed or engaged by certified providers and carry out work at the premises at which health care services are provided; 4. Care and support workers.

example, cleaners who are employed by a cleaning company but work in a hospital would be captured by the Order Part 7 definition but fall under ANZSIC06 *N7311 Building and Other Industrial Cleaning Services.*

Moreover, we will compare HCWs with workers in barely-mandated industries. These are industries such as *A Agriculture, Forestry and Fishing*, *E Construction* and *M Professional, Scientific and Technical Services* and all their associated sub-industries, where it is clear that government mandates did not apply generally. We refer to these as 'barely' mandated industries rather than non-mandated industries because there may have still been a few workers in some of these industries who were subject to mandates. For example, an IT worker employed by an IT firm and therefore falling under ANZSIC06 *M7000 Computer System Design and Related Services* but contracted to work in a hospital could have fallen under the definition of the Order, but is defined as belonging to a barely mandated industry. However, these small classification issues do not undermine the validity of the approach as the important point for our analysis is that the share of workers subject to mandates in 'heavily' mandated industries is much higher than the share in 'barely' mandated industries.

Another potential issue is that some employees were subject to employer-imposed, rather than government-imposed, vaccine mandates, which may downward bias our estimates. For example, most central and local governments, universities and even a few private businesses required staff to be vaccinated. We have classified employees in industries where employer-imposed mandates were common, such as central and local government, as partially mandated rather than barely-mandated industries to avoid them appearing in the comparison group for our DiD analysis. However, there may still be some employees in the barely-mandated comparison group who were subject to mandates. This issue will result in an underestimate of the effect of the mandates.

4.1.3 Identifying compliance with COVID-19 vaccination mandates

Our main interest is in HCWs, but as comparison points and to characterise compliance with the COVID-19 vaccination mandates, we also examine corrections prison and education workers. These two additional groups have clear links with ANZSIC06 industry codes. As mentioned, other workers covered by COVID-19 vaccination mandates are harder to identify via industry classifications because the Order did not mandate the whole industry.

Table II.1 details the mandate announcement dates and deadlines. In late January 2022, the government added a mandatory booster dose to the Order for these groups, but the deadline was less clear cut and we therefore do not examine this.[10]

[10] For HCWs, the booster dose deadline was the later of either 25 February 2022 or 183 days after the date of the second dose.

Table II.1: Key dates for COVID-19 vaccination mandates for health, corrections, and education industries

(1) Order	(2) ANZSIC06 industry	(3) Announcement date	(4) First dose deadline	(5) Second dose deadline
Part 7: Health and disability sector	Q84	11-Oct-21	15-Nov-21	1-Jan-22
	Q85	11-Oct-21	15-Nov-21	1-Jan-22
	Q86	11-Oct-21	15-Nov-21	1-Jan-22
Part 8: Corrections prisons	O771400	23-Oct-21	6-Nov-21	8-Dec-21
Part 9: Affected education services	P801	11-Oct-21	15-Nov-21	1-Jan-22
	P802	11-Oct-21	15-Nov-21	1-Jan-22
	P801	11-Oct-21	15-Nov-21	1-Jan-22

For individuals employed in industries under Parts 7, 8 and 9 of the Order, we use the Ministry of Health COVID-19 vaccination register to identify workers' compliance with the COVID-19 vaccination mandate. The vaccination register allows us observe if and when each individual received a first dose and second dose of an approved COVID-19 vaccine. By comparing actual vaccination dates with the mandated vaccination deadline dates, we categorise individuals into one of the following vaccination compliance categories:

1. Individuals who got vaccinated likely irrespective of vaccination mandates:
    - people who received their first dose of a COVID-19 vaccine before the vaccination mandate was announced, and who subsequently received a second dose before the mandated second dose deadline
2. Individuals who got vaccinated likely due to vaccination mandates:
    - people who received their first dose after the vaccination mandate was announced but before the mandated first dose deadline, and their second dose before the mandated second dose deadline
    - people who missed the mandated first dose deadline, but had both their first and second dose before the mandated second dose deadline
3. Individuals who were unvaccinated or uncompliant with vaccination mandates:
    - people who received no doses of a COVID-19 vaccine
    - people who received only one dose of a COVID-19 vaccine
    - people who received a first dose but had their second dose after the mandated second dose deadline

Anticipatory effects were considered but public information indicates it is unlikely to be an issue. Even if the mandates had been anticipated, vaccine-hesitant individuals were unlikely to get vaccinated until it was clear that this would be required as a condition of continued employment. However, there are several factors which suggest that they were not anticipated. The first vaccine mandates covering MIQ workers were announced in April 2021, and these were extended to port and airport workers in July 2021. However, we argue that these early vaccine mandates would not have induced an anticipatory effect among the vaccine hesitant for at least two reasons. First, in September 2020 the government had explicitly ruled out the possibility of COVID-19 vaccine mandates. The decision to mandate MIQ, port and airport workers could therefore be seen as a contradiction to earlier government media statements, and thus the prospect of vaccine mandates being applied more broadly was generally dismissed due to the government's initial sentiment against vaccine mandates. Second, the first set of vaccine mandates applied to a relatively small number of workers and occurred well before the HCW announcements, making it unlikely that the MIQ, port and airport mandates would have led HCWs to believe they would also be mandated.

This is supported by a Google news search for 'vaccine mandates health workers NZ' (and variants thereof), which reveals no media coverage on the possibility of HCW vaccine mandates before the government made the official announcement on 11 October 2021. In addition, a series of Ministry of Health-commissioned surveys about attitudes towards COVID-19 vaccines in NZ were carried out between December 2020 and October 2021. These included free-form responses and concerns about vaccine mandates were only reported in the last (October 2021) survey (HorizonResearch, 2021), which was the only one conducted after the mandate announcements. These pieces of evidence, in conjunction with the fact that the policy decisions at the time were being made quickly in a crisis-management mode suggests that HCW vaccine mandates were unlikely to have been widely anticipated.

4.2 RQ2 healthcare labour market outcomes

4.2.1 Creating a sample of employed individuals

To study the impact of COVID-19 vaccine mandates on health workforce labour market outcomes, we focus on two populations of workers: HCW and barely-mandated workers. As described in Section 5.2, we use the barely-mandated workforce as a comparison group to help isolate the mandate effect from the industry-specific and general pandemic effects on job separation rates.

We create the HCW and barely-mandated worker populations from the 2019 usually-resident population table in the IDI and observe all working aged (16-60 years) individuals with positive earnings (in terms of receiving positive wages and salaries) in March 2019 and who did not die throughout our study period. We use March 2019 as the last end-of-tax year before the COVID-19 pandemic began. To account for multiple job holdings (as with RQ1), we observe individuals in their main job in March 2019, defined first by the job with an M-type (main income source) tax code and then by the job that provides highest labour earnings. We then categorise individuals into industries using the 7-digit ANZSIC06 code associated with their main job in March 2019. As outlined in Section 4.1.2 for RQ1, we also link the 7-digit ANZSIC06 industry codes to the work description provided in the Order to identify whether a worker was heavily, partially, or barely subject to the COVID-19 vaccination mandates. This enables us to define a sample of 156,417 HCWs and 1,242,822 barely-mandated workers benchmarked in March 2019.

4.2.2 Identifying individuals' vaccination status

As described in Section 4.1.3, we identify a worker's vaccination status using the Ministry of Health's nationwide COVID-19 vaccination register in the IDI. For the HCW population, we define vaccination status as an indicator equal to 1 if the individual was double vaccinated and in compliance with the vaccination mandate and equal to 0 if the individual was unvaccinated or uncompliant with the mandate. We create this vaccination indicator consistent with the three vaccination compliance categories as defined for RQ1 in Section 4.1.3. Namely, we make the vaccination indicator equal to 1 if the HCW was vaccinated regardless of mandates or vaccinated likely due to the mandates and equal to 0 if the HCW was unvaccinated or uncompliant with the mandates.

For the barely-mandated worker population, we define vaccination status as an indicator equal to 1 if the individual was (at least) double-vaccinated and equal to 0 if the individual was unvaccinated or only received a single vaccination. While there were no mandate deadlines for the barely-mandated workers, we follow the same definitions as used for the HCW in Section 4.1.3 to make the two groups as comparable as possible. Therefore, we consider a worker to be double-vaccinated if they received at least two vaccinations as recorded in the vaccination register. Further, to be consistent with the HCW sample, we drop a very small number of barely-mandated workers who have dubious vaccination records.

4.2.3 Identifying individuals' overseas spells

It is important that we restrict our sample to workers who reside in NZ during our study period because vaccination records and earnings data could be misleading and/or missing for workers who spend considerable time overseas. Thus, we use the overseas spells border movement data to count the total number of days each worker spent outside NZ throughout the study period of March 2019 to November 2022. To identify workers who reside in NZ, we use Stats NZ's '12/16-month rule' [11] whereby they differentiate long-term migrants from short-term migrants (i.e., visitors) if the individual is in NZ for 12 out of 16 months in a given period. We adjust this 12/16 rule to our 45-month study period. This is equivalent to approximately 34/45 months, corresponding to about 1,020 days. Thus, we define a worker as residing in NZ if they were inside NZ for at least 1,020 days during our study period (approximately 1,350 days). Equivalently, a worker is defined as not residing in NZ if they were outside NZ for at least 330 days. This equates to 5.7% of HCWs and 1.1% of barely-mandated workers.

Overall, this leaves us with the following four groups for RQ2 analysis:

- Vaccinated HCW residing in NZ (144,087)
- Vaccinated barely-mandated workers residing in NZ (1,068,726)
- Unvaccinated HCW residing in NZ (12,330)
- Unvaccinated barely-mandated workers residing in NZ (174,099)

4.2.4 Defining labour market outcomes

We use IRD tax data in the IDI to obtain each worker's labour market information and examine HCWs' employment and earnings outcomes. We create unique individual-month observations by observing individuals in their main job each month following the same criteria as set out in Section 4.1. This results in a balanced monthly panel spanning 45 months from March 2019 to November 2022 for all HCW and barely-mandated workers. This equates to almost 63 million worker-month observations.

For both HCWs and barely-mandated workers, we define an *employment indicator* that equals 1 if the individual received positive wages and salaries in a given month, and 0 otherwise. We also create a monthly *labour earnings* variable showing the wages and salaries earned from the individual's main job per month.

We then define two additional industry-specific employment indicators. The first is a *same industry indicator*. For HCW, this indicator equals 1 if the individual is employed in the health industry and equals 0 otherwise (i.e. if the person is employed in a different industry or not employed). For barely-mandated workers, this indicator

[11] As detailed on Stats NZ's 'Migration Data Transformation' project webpage, https://www.stats.govt.nz/about-us/what-wedo/current-projects/migration-data-transformation-project/ (accessed on 3 May 2022) and Stats NZ (2017).

equals 1 if the individual is employed in the same 1-digit ANZSIC06 industry that they were in March 2019 (the '1-digit reference industry'),[12] and equals 0 otherwise.

The second industry-specific employment indicator is a *same industry indicator conditional on employment*. For HCW, this indicator equals 1 if the individual is employed in the health industry and equals 0 if the individual is employed but not in the health industry. For barely-mandated workers, this indicator equals 1 if the individual is still employed in their 1-digit reference industry and equals 0 if the individual is employed in a different 1-digit ANZSIC06 industry. These two same-industry indicators provide us with a means to examine the extent to which the mandates impacted on the exit of unvaccinated HCWs from the health industry.

We also examine rates of industry switching descriptively to examine whether workers exiting the health industry were being replaced, in order to gauge whether job separations caused by the mandates may have contributed to the industry's worker shortages. However, this measure of job accessions is of secondary concern since if there were excess workers leaving the health industry due to the mandates, it is difficult to see how these gaps could have been filled by job accessions. Most healthcare jobs are skilled roles that require high levels of training and experience, making it difficult to fill gaps from within the domestic labour market in the short-term. Historically, like other developed countries, NZ has filled these immediate gaps with inward migration. However, during the pandemic, offshore recruitment of both migrants and returning New Zealanders was very limited due to NZ's border restrictions,[13] and, at this time, NZ also had more restrictive migration conditions for HCWs than competing anglophone countries (such as Australia and Canada).[14]

For job separations, we examine whether HCWs move out of the industry, either into employment in another industry or out of employment entirely, by defining a *left health sector* indicator that equals 1 if a HCW is no longer employed in the health industry and 0 if they remain employed in this industry. For barely-mandated workers, this indicator equals 1 if the individual is no longer employed in their 1-digit reference industry, and 0 if they remain employed in their reference industry.

For job accessions, we examine whether barely-mandated workers move into the health sector by defining a *joined health sector* indicator that is equals 1 if a barely-

[12]The use of a 1-digit reference industry roughly matches the level of the health industry, which is the 1-digit industry *Q Health Care and Social Assistance* less social assistance services.

[13]The Ministry of Health was allocated priority access to 300 MIQ rooms a month for critical HCWs from November 2021, approximately 20 months after the NZ border restrictions were implemented. In addition, between November 2021 and February 2022, only 147 of the 900 allocated places were used.

[14]For example, during the period being investigated, Australia and Canada offered migrant nurses, midwives and doctors residency visas immediately, while in NZ, until December 2022, nurses midwives and some doctors (depending on their specialty) were only eligible for temporary work visas and had to wait at least two years before becoming eligible to apply for residency.

mandated worker becomes employed in the health industry, and 0 otherwise. For HCWs, this indicator equals 1 if they become employed in a barely-mandated industry. As discussed, our main interest is in job separations, and our population of interest is less well suited to measuring accessions as this measure does not consider NZ residents who moved from not being in employment to being HCWs, those who move from overseas to employment as a HCW, and those who move from working in partially-mandated industries to being HCWs.

4.3 RQ1 and RQ2 additional descriptive variables

We link several administrative datasets in the IDI to obtain demographic and socioeconomic information for all individuals in our RQ1 and RQ2 samples.

We use the personal details table to obtain an individual's age, sex and ethnicity.[15] To define an individual's migrant status, we use the Department of Internal Affairs births register to identify if the individual was born in NZ or not.

We use Stats NZ's derived address notification dataset to identify the meshblock associated with each individual's last known residential address. We use the meshblock code to identify each individual's residential District Health Board (DHB) area, and to identify the level of socioeconomic deprivation in the area as measured by the NZ Deprivation Index (NZDep) 2018.[16]

5 Method

5.1 RQ1: Vaccine uptake

To estimate the extent to which the COVID-19 vaccination mandates increased vaccination uptake, we employ a difference-in-differences estimation strategy. By comparing the vaccination uptake of HCWs (treatment group) with the vaccination uptake of individuals employed in barely-mandated industries (comparison group), we can isolate the effects of the mandates from the general increase in vaccinations due to other population-wide initiatives, such as vaccine passports. Specifically, we estimate the following model set out by Equation 4:

$$Y_{it} = \alpha + \beta HCW_i + \gamma Post + \delta(HCW_i.Post) + \epsilon_{it} \tag{4}$$

[15] We use Stats NZ's prioritisation rules to create mutually exclusive ethnicity categories prioritised as follows: Māori; Pacific peoples; Asian; Middle Eastern, Latin American, or African (MELAA); Other; European.

[16] The NZDep is a measure of socioeconomic deprivation based on the meshblock a person lives in, with a meshblock being roughly equivalent to a city block. The Index ranges 1-10, with 1 being the least deprived areas and 10 being the most deprived areas.

where Y is an indicator of vaccination uptake, equal to 1 if the individual is double-vaccinated and 0 otherwise.[17] T represents treatment status, where $HCW = 1$ for HCWs and $HCW = 0$ for individuals who work in barely-mandated industries. We observe individuals in two time periods as indicated by $Post$, where $Post = 0$ indicates the time period before the vaccine mandate was announced (i.e. the pre-announcement time period) and $Post = 1$ indicates the time period after the vaccine mandate was announced (i.e. the post-announcement time period). ϵ_{it} denotes the error term. To account for the possibility of serial and intra-group correlation, robust standard errors clustered at the level of the individual are used (Bertrand et al., 2004).

The coefficient of interest δ reveals the treatment effect of COVID-19 vaccination mandates on vaccination uptake. β is the treatment-group specific effect that accounts for permanent differences between the average vaccination uptake of the treatment group compared to the comparison groups (e.g. to account for the fact that mandated health workers may have permanently higher vaccination rates than workers in the comparison group). γ is the time trend common to the treatment and comparison groups.

5.2 RQ2: Healthcare worker labour market outcomes

As discussed, we examine the effect of mandates on several labour market outcomes for HCWs, including employment, earnings and job accession and separation rates. Note that while the data can reveal whether an unvaccinated health worker left employment, the data does not tell us why. It may be that some workers left voluntarily for non-mandate reasons, such as retirement, family pressures, or a career change, etc. It also may be due to the COVID-19 vaccination mandates. Therefore, to estimate the effect of the vaccine mandates on labour market outcomes for HCW, we use a triple difference-in-differences (DDD) analysis comparing unvaccinated and vaccinated HCWs with unvaccinated and vaccinated barely-mandated workers over time. Since labour market outcomes can be measured on a monthly basis, we use dynamic DiD estimates, although we also conduct robustness checks using two-period DiDs.

In terms of the treatment and comparison groups, we consider two potential options. One involves comparing the labour market outcomes of unvaccinated versus vaccinated HCWs, and another involves comparing the outcomes of unvaccinated HCWs versus unvaccinated barely-mandated workers. Both these options have potential advantages and issues, and we thus instead employ a triple difference method that incorporates both comparisons on the basis that the difference between two biased DiD estimates is potentially unbiased as long as the bias is the same in both

[17] As a robustness test, we also define vaccination uptake by only the first dose, and it did not qualitatively change the results.

estimators (Olden & Møen, 2022).

In terms of the comparison of unvaccinated HCWs with vaccinated HCWs, this has the advantage that both groups are in the same industry, and thus faced the same set of industry labour market conditions. However, there may be spillover effects since, for example, unvaccinated HCWs leaving the healthcare industry may increase skills shortages, improving the bargaining power and labour market outcomes of vaccinated HCWs. Alternatively, it could have potentially increased the pressure on vaccinated HCWs and expedited their exit from the health industry.

Furthermore, a comparison of unvaccinated HCWs with unvaccinated barely mandated workers may also be problematic if the conditions in the two industries diverged for reasons other than the vaccine mandates. This is a realistic concern given the pandemic itself may have impacted the health industry differently to other industries. For example, there may have been increased demand for HCWs, particularly relative to other industries, many of which intially saw a reduction in demand due to COVID-19 lockdowns (although, after this initial reduction, the labour market was buoyant and skills shortages became an issue in many industries). This high demand for HCWs could have potentially reduced (increased) job separation (accession) rates in the post-pandemic time period regardless of the imposition of mandates. On the other hand, the pandemic may have increased (decreased) job separation (accession) rates by creating a more stressful work environment for HCWs in a way that was not experienced by other industries, thus leading to higher HCW job separation rates regardless of the vaccine mandates.

However, this potential issue was likely to be less problematic in the case of NZ relative to other countries as NZ experienced few COVID-19 cases until March 2022 (Figure II.3, Panel A) so did not have the same issues of COVID-19 cases straining the health system. Indeed, NZ's excess mortality - the cumulative difference between the reported number of deaths since 1 January 2020 and the projected number of deaths for the same period based on previous years, shown in Figure II.3, Panel B - was actually negative while the OECD average was large and positive. This reflects that COVID-19 containment measures also greatly reduced the number of cases of, and deaths associated with, influenza and other respiratory illnesses, and the lockdown reduced accident-related deaths (e.g. by heavily reducing traffic volumes) (Kung et al., 2021).

Nevertheless, HCW stress was still an issue in NZ, due to factors such as staff shortages which were exacerbated by border closures. In addition, staff illness and isolation requirements also contributed to staff shortages. From March 2020, anyone who had potentially came into contact with a COVID-19-positive case, including causal contact such as visiting a supermarket within the same window of time, had to isolate

Figure II.3: Cumulative COVID cases and excess mortality: NZ versus OECD

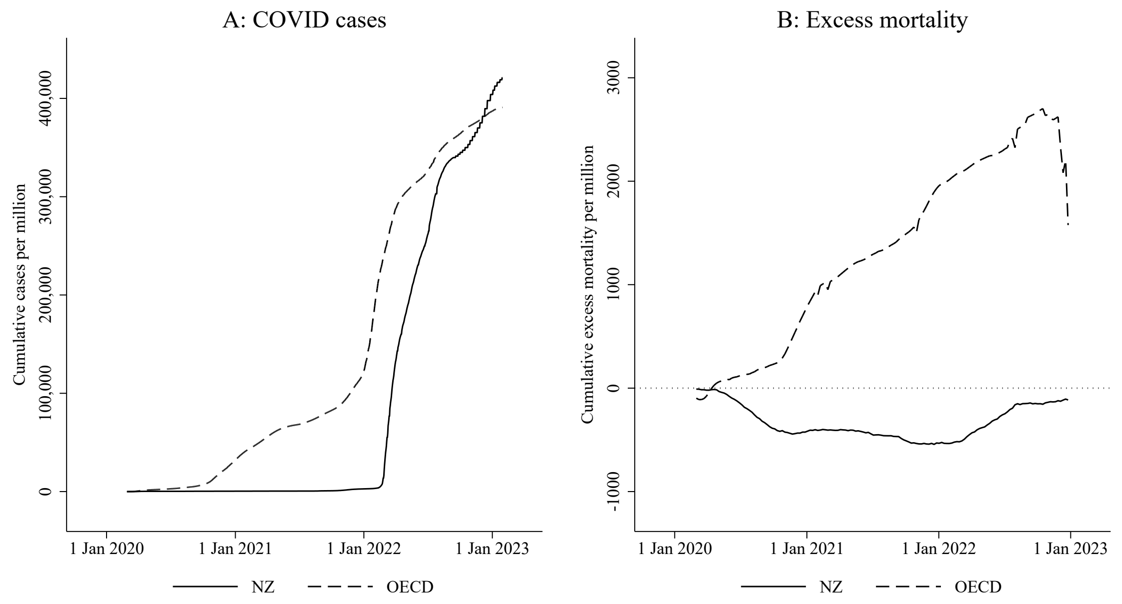

Notes: OECD is a simple average of OECD countries with available data.
Source: Our World in Data COVID-19 database. Accessed from https://github.com/owid/covid-19-data/tree/master/public/data on 1 February 2023.

at home for 14 days, which led to a large number of people isolating despite the low COVID-19 case numbers. While the rules around who had to isolate and for how long eased over time, up until September 2022, it was still necessary for those with COVID-19 and their household contacts to isolate for 7 days. Between September 2022 and August 2023, only those with COVID-19 had to isolate. While these isolation rules created staffing shortages across many industries, not just healthcare, given there was a higher probability of a HCW coming into contact with someone with COVID-19, this issue may have impacted HCWs more. Thus, to account for the possibilities of spillover effects and differential industry effects, we estimate implement a DDD approach.

This DDD approach takes the form:

$$\begin{aligned} Y_{sit} = & \alpha + \beta_0 UV_i + \beta_1 HCW_i + \beta_2 (UV_i.HCW_i) + \gamma_0 Post + \gamma_1 (UV_i.Post) \\ & + \gamma_2 (HCW_i.Post) + \gamma_3 (UV_i.HCW_i.Post) + \delta X_{sit} + \epsilon_{sit} \end{aligned} \tag{5}$$

where Y_{sit} refers to the outcome of interest for individual i in sector s at time t. In terms of the treatment groups, $UV_i = 1$ for those who are unvaccinated (i.e. those who did not meet the requirements of the relevant vaccine mandate by the stated deadline) and $UV_i = 0$ for those who are vaccinated (i.e. did meet the requirements of the relevant vaccine mandate). $HCW_i = 1$ for HCWs and $HCW = 0$ for barely-mandated workers. X_{sit} is a vector of control variables, such as age, gender, ethnicity and socioeconomic status and robust standard errors are clustered at the individual level.

This approach provides four different estimators of the effect of vaccine mandates on labour market outcomes. The sum of the estimators $\hat{\gamma}_0$, $\hat{\gamma}_1$, $\hat{\gamma}_2$, $\hat{\gamma}_3$ provides an "undifferenced" estimate of the effect on unvaccinated HCWs. To the extent that there may be other factors that may have affected the labour market outcomes of both unvaccinated barely-mandated and HCWs that were unrelated to the mandates, we can difference out a common "unvaccinated" effect by using $\hat{\gamma}_2 + \hat{\gamma}_3$. A third estimator ($\hat{\gamma}_1 + \hat{\gamma}_3$) uses the difference between unvaccinated HCWs and vaccinated HCWs to remove any common "HCWs" effect that both the unvaccinated and vaccinated share. Finally, a "triple difference estimator", $\hat{\gamma}_3$, differences out both vaccination status and sector effects and is therefore our coefficient of interest. That is, $\hat{\gamma}_3$ is the effect of mandates on unvaccinated HCWs relative to vaccinated HCWs and unvaccinated barely-mandated workers.

Because we have labour market outcomes on a monthly basis, as our preferred estimation method, we employ a dynamic DDD of the form:

$$\begin{aligned} Y_{sit} = &\alpha + \beta_0 UV_i + \beta_1 HCW_i + +\beta_2(UV_i.HCW_i) + \sum_{e\neq -3, e=-15}^{13} \gamma_0.M_e \\ &+ \sum_{e\neq -3, e=-15}^{13} \gamma_1(UV_i.M_e) + \sum_{e\neq -3, e=-15}^{13} \gamma_2(HCW_i.M_e) \\ &+ \sum_{e\neq -3, e=-15}^{13} \gamma_3(UV_i.HCW_i.M_e) + \delta X_{sit} + \epsilon_{sit} \end{aligned} \tag{6}$$

where M_e are event time indicators, where $M_e = 0$ is the vaccine mandate announcement in October 2021. That is, we track the outcome on a monthly basis from July 2020 to November 2022.

In general, the treatment effect is expressed as a percentage of the counterfactual (P_e) to provide a comparable sense of the magnitude of the effect. Using the example of employment:

$$P_e = \frac{\gamma_3}{E[\tilde{Y_{sit}}|t]}.100 \tag{7}$$

where P_e is the ratio of the parameter of interest (γ_3) from Equation 6 at event time t to the predicted employment outcome conditional on event time t, multiplied by 100.

6 Results: RQ1 - Vaccine uptake

This section presents results for RQ1 on vaccination uptake. It first presents summary statistics by COVID-19 vaccination mandate categories as well as vaccination compliance categories. It then examines vaccination uptake over time for different groups of mandated workers as well as barely-mandated workers. Finally, it presents DiD results.

6.1 Summary statistics

By COVID-19 vaccination mandate categories

Table II.2 presents descriptive statistics of the March 2021 worker cohort categorised by whether they were in barely-mandated, partially-mandated or heavily-mandated industries (as defined in Section 4.1.2). The largest group is the barely mandated category, which comprises nearly 1.3 million workers. The heavily mandated group has just over 460,000 workers and the partially mandated group has nearly 190,000 workers.

Table II.2: Descriptive statistics of workers in industries that barely, partially, or heavily faced COVID-19 vaccination mandates

(1) Characteristic	(2) Barely mandated	(3) Partially mandated	(4) Heavily mandated
Number of workers	1,289,007	189,501	461,604
Had at least one vaccine (%)	92.18	94.96	95.00
Demographic			
Age (years)	37.52	39.58	37.96
Female (%)	39.00	58.63	71.31
Gender unknown (%)	0.06	0.05	0.06
European (%)	53.81	55.97	53.54
Māori (%)	15.21	15.11	14.76
Pacific (%)	7.64	7.41	5.69
Asian (%)	19.04	16.89	22.06
MELAA/Other (%)	3.90	4.39	3.63
Ethnicity unknown (%)	0.40	0.23	0.32
NZ born (%)	62.44	61.92	59.29
Socioeconomic			
NZ Deprivation Index 2018	5.51	5.30	5.43
Monthly income in March 2021 ($)	6,522	7,127	5,445
Monthly earnings in March 2021 ($)	6,438	7,050	5,349
Monthly earnings from main job in March 2021 ($)	6,330	6,941	5,224

Notes: This table presents demographic and socioeconomic characteristics of all individuals who were employed in March 2021, categorised by whether their industry barely faced COVID-19 vaccination mandates, partially faced COVID-19 vaccination mandates, or heavily faced COVID-19 vaccination mandates (as defined in Section 4.1.2). Percentages may not always add up to 100 due to rounding.

The average age of workers in each group is fairly similar, with the barely-mandated group and the heavily-mandated being around 37-38 years old, on average, while the partially-mandate group is nearly 40 years old, on average. About 95% of both the partially-mandated and heavily-mandated groups received at least one COVID-19 vaccination, while this is only 92% for the barely-mandated group.

Females make up a much larger percentage of the heavily mandated group (71.3%) compared to the partially-mandated group (58.6%) and the barely mandated group (39%). This is unsurprising since the heavily-mandated industries are mostly female-dominated health and education industries, while the barely-mandated industries include those such as construction, which are more male-dominated.

Just over half of all three groups are European, about 15% are Māori, and about 7% are Pacific peoples. This suggests that the composition of Māori and Pacific peoples' in each COVID-19 mandate industry group is approximately representative of the composition of Māori and Pacific peoples' in the general New Zealand population. Approximately 60% of each group were born in NZ. The average deprivation level is fairly similar across the three COVID-19 mandate groups, sitting at around 5.4.

Across all income and earnings measures, workers in partially-mandated industries earn more than those in barely-mandated or heavily-mandated industries. Those in mandated industries received an average total income of $7,127 in March 2021, while those in barely mandated and heavily mandated industries received $6,521 and $5,445, respectively. These figures are similar when looking at total monthly earnings received in March 2021. Further, across all groups, approximately 98% of total monthly earnings received in March 2021 were from the individual's main job, suggesting few workers received additional earnings from secondary jobs.

Table II.3 shows the percentage of each COVID-19 mandate industry group that reside in each DHB region as at March 2021. The DHB composition of each group is fairly similar. A notable exception is that 15% of workers in the partially-mandated group reside in the Capital and Coast DHB, whereas this statistic is only about 6% and 7% for the barely-mandated and heavily-mandated groups, respectively. This is likely because the large majority of central-government employees reside in Wellington and public sector jobs were more likely to be subject to COVID-19 vaccination mandates.

By COVID-19 vaccination mandate compliance categories

For the mandated industries of interest (health, corrections, and education), we provide summary statistics by COVID-19 vaccination mandate compliance categories. This splits each group into three categories (as defined in Section 4.1.3): those who were likelyvaccinated regardless of mandates, those who were vaccinated potentially due to mandates, and those who did not comply with the mandates.

Table II.3: District Health Board compositions of workers in industries that barely, partially, or heavily faced COVID-19 vaccination mandates

(1) Characteristic	(2) Barely mandated	(3) Partially mandated	(4) Heavily mandated
Number of workers	1,289,007	189,501	461,604
District Health Boards			
Northland	2.68	2.65	3.43
Waitemata	12.99	10.37	12.78
Auckland	11.10	10.70	10.33
Counties Manukau	12.72	8.98	10.12
Waikato	8.12	7.04	8.62
Lakes	2.08	2.29	2.34
Bay of Plenty	4.99	3.44	4.82
Tairawhiti	0.99	0.62	1.01
Taranaki	2.34	1.67	2.34
Hawke's Bay	3.57	2.63	3.44
Whanganui	1.10	0.90	1.49
Mid-Central	3.22	4.01	4.13
Hutt Valley	2.96	5.34	3.35
Capital and Coast	5.81	15.50	7.32
Wairarapa	0.84	0.73	0.93
Nelson/Marlborough	3.23	2.26	2.97
West Coast	0.58	0.61	0.52
Canterbury	12.21	10.50	11.94
South Canterbury	1.30	0.73	1.07
Southern	6.85	8.78	6.85
Area outside DHB	.s	.s	.s
DHB unknown	0.32	0.23	0.18

Notes: This table shows the percentage of workers that reside in each District Health Board as at their last recorded address on March 2021, categorised by whether the individual's main job industry barely faced Covid-19 vaccination mandates, partially faced COVID-19 vaccination mandates, or heavily faced Covid-19 vaccination mandates. Percentages may not always add up to 100 due to rounding. Notation ".s" means counts have been suppressed in accordance with Stats NZ confidentiality rules.

Part 7 - Healthcare workers

Table II.4 presents the demographic and socioeconomic characteristics of HCWs who were subject to COVID-19 vaccination mandates under Part 7 of the Order. This group comprises 171,486 workers, of which 89.2% (152,937) were vaccinated before the mandate was announced and hence were likely vaccinated regardless of the mandate. About 5.5% (9,426) were vaccinated within the mandate announcement and vaccination deadlines and hence could have potentially vaccinated due to the mandate. The remaining 5.3% (9,123) were unvaccinated or uncompliant. Of those in the unvaccinated or uncompliant group, 21% received at least one COVID-19 vaccine.

This reasonable share of partially vaccinated uncompliant workers raises the possibility that some were at least somewhat open to being vaccinated. We do not know why they did not receive a second dose, but it may have been due to factors such as experiencing an adverse effect from the first dose.

HCWs who were vaccinated before the mandate was announced and those who did not comply with the mandate were 41 years old on average, while those who may have gotten vaccinated potentially due to the mandates were slightly younger, at 38 years old on average. The large majority of all three groups were female, ranging from 80-84 %. Those who were born in NZ make up a larger share of those who could have potentially been vaccinated due to the mandate than the other two compliance groups.

Table II.4: Descriptive statistics of workers in health industries that heavily faced COVID-19 vaccination mandates

(1) Characteristic	(2) Vaccinated regardless of mandate	(3) Vaccinated potentially due to mandate	(4) Unvaccinated or uncompliant with mandate
Number of workers	152,937	9,426	9,123
Had at least one vaccine (%)	100.00	100.00	20.91
Demographic			
Age (years)	41.07	38.04	41.46
Female (%)	80.80	84.15	83.76
Gender unknown (%)	.s	.s	.s
European (%)	52.33	47.77	53.47
Māori (%)	11.07	26.96	18.28
Pacific (%)	5.87	9.17	7.04
Asian (%)	26.96	12.00	14.21
MELAA/Other (%)	3.58	4.07	3.98
Ethnicity unknown (%)	.s	.s	.s
NZ-born (%)	52.19	69.35	55.48
Socioeconomic			
NZ Deprivation Index 2018	5.42	6.44	5.85
Monthly income in March 2021 ($)	6822	4935	5368
Monthly earnings in March 2021 ($)	6736	4742	5195
Monthly earnings from main job in March 2021 ($)	6563	4626	5060

Notes: This table presents demographic and socioeconomic characteristics of all individuals who were employed in health industries in March 2021, categorised by vaccination mandate compliance behaviour. Percentages may not always add up to 100 due to rounding. Notation ".s" means counts have been suppressed in accordance with Stats NZ confidentiality rules.

In terms of differences by ethnicity, European HCWs were about as likely to be in each of the three compliance groups. Asian HCWs were more likely to be in the group that would have been vaccinated regardless of the mandate. Māori and Pacific HCWs were more likely to be in the group that were potentially vaccinated due to the mandate compared with the other two compliance groups.

While our data cannot shed light on the reasons behind these ethnicity patterns, it could potentially raise issues of the coercive nature of the mandates further eroding trust in public institutions among these workers. This is a particular issue in relation to Māori HCWs given the historical legacies of colonisation affecting trust among the Māori population. This is reflected in, for example, NZ's General Social Survey, which shows that 44% of Māori rated their trust in parliament as low compared with 29% of the total population, and 47% of Māori feeling that the public had little to no influence on government decision making, versus 37% of the total population (Stats NZ, 2018). Moreover, qualitative research involving Māori, Pacific and disability communities undertaken as part of the government's equity review of the COVID-19 response highlights that these groups felt that mandates further disadvantaged them (Paipa et al., 2023). Qualitiative research on NZ HCWs specifically also highlights that Māori HCWs felt that the mandates were another measure that was being imposed upon them and contributed to their sense of a loss of control (Dewar et al., 2024).

There are also clear differences by socioeconomic status. Health workers who likely would have vaccinated regardless of mandates and workers who did not comply with mandates have an average deprivation level of approximately 5.5, while those who could have potentially gotten vaccinated due to mandates have a higher average index of nearly 6.5 (where a higher index indicates a higher level of deprivation). Similarly, health workers who could have been vaccinated potentially due to the mandates have lower monthly income and monthly earnings than those in the other two groups. This may reflect that health workers with lower socioeconomic status experienced stronger financial imperatives and therefore were more likely to comply with the mandate to avoid losing their jobs.

Part 8 - Corrections workers

Table II.5 presents the demographic and socioeconomic characteristics of workers in corrections industries who were subject to a COVID-19 vaccination mandate under Part 8 of the Order. This group comprises 8,937 workers, of which 90.1% (8,055) were vaccinated before any mandate announcement and hence were likely vaccinated regardless of the mandate, 5.5% (492) were vaccinated after the mandate announcement but before the mandate deadlines and hence could have potentially been vaccinated due to the mandates, and the remaining 4.4% (390) did not comply with the mandate. These statistics are fairly similar to the vaccination compliance shown by health workers in Table II.4. Of the corrections workers who did not comply with the mandate, 9% received at least one vaccine dose.

Corrections workers in the first and third compliance groups were 43 years old on average, while those who may have gotten vaccinated due to the mandates are slightly younger, at 40 years old on average. Again, Māori and Pacific Peoples are overrepresented among those who were vaccinated potentially due to the mandate compared to those who likely would have vaccinated regardless of the mandate and those who did not comply with the mandate.

Unlike the results for health workers, corrections workers have a more balanced gender split, where those who were vaccinated before or after the mandate announcement have approximately a 50:50 female-to-male ratio, while it was about 60:40 for those who did not comply with the mandate. This could reflect that women are more likely to be secondary income earners within families, and, therefore, could have been less concerned about losing their jobs by not complying. It could also reflect that women are more likely to be COVID-19 vaccine hesitant (Toshkov, 2023) and more likely to experience adverse effects from the vaccine (Duijster et al., 2023; Green et al., 2022).

A similar story is evident when looking at the socioeconomic status variables across vaccination mandate compliance categories as with HCWs. Corrections workers who may have gotten vaccinated due to the mandates have a higher deprivation score and lower monthly earnings and monthly income than workers in the other two compliance categories. However, unlike HCWs, corrections workers who were not compliant look more similar to group two (those who may have gotten vaccinated due to the mandates) than group one (those who would have likely gotten vaccinated regardless of the mandate) in terms of socioeconomic status indicators.

Table II.5: Descriptive statistics of workers in corrections industries that heavily faced COVID-19 vaccination mandates

(1) Characteristic	(2) Vaccinated regardless of mandate	(3) Vaccinated potentially due to mandate	(4) Unvaccinated or uncompliant with mandate
Number of workers	8,055	492	390
Had at least one vaccine (%)	100.00	100.00	9.23
Demographic			
Age (years)	43.23	40.02	42.85
Female (%)	50.54	48.78	57.69
Gender unknown (%)	.s	.s	.s
European (%)	52.03	42.68	48.46
Māori (%)	19.03	25.61	20.77
Pacific (%)	12.33	20.73	14.62
Asian (%)	12.33	6.71	10.00
MELAA/Other (%)	4.21	4.27	3.85
Ethnicity unknown (%)	.s	.s	.s
NZ-born (%)	60.86	63.41	58.46
Socioeconomic			
NZ Deprivation Index 2018	5.60	6.43	6.16
Monthly income in March 2021 ($)	9349	8444	8305
Monthly earnings in March 2021 ($)	9327	8414	8279
Monthly earnings from main job in March 2021 ($)	9264	8371	8231

Notes: This table presents demographic and socioeconomic characteristics of all individuals who were employed in corrections industries in March 2021, categorised by vaccination mandate compliance behaviour. Percentages may not always add up to 100 due to rounding. Notation ".s" means counts have been suppressed in accordance with Stats NZ confidentiality rules.

Part 9 - Education workers

Table II.6 presents the demographic and socioeconomic characteristics of workers in education industries who were subject to a COVID-19 vaccination mandate under Part 9 of the Order. This group comprises 122,397 workers, of which 83.4% (102,084) were vaccinated before any mandate announcement and hence were likely vaccinated regardless of the mandate, 10% (12,285) were vaccinated within the mandate announcement and vaccination deadlines and hence may have potentially been vaccinated due to the mandate, and the remaining 6.6% (8,028) did not comply with the mandate. Thus, education industries have a lower percentage of workers that were vaccinated regardless of mandates compared with health and corrections industries, and a higher percentage of workers who potentially got vaccinated due to the mandates. While this could reflect differences in the degree of vaccine hesitancy, it may, however, reflect that, unlike health and corrections workers, education workers did not have early access to the vaccine (discussed more below). As with health and corrections workers, a reasonable minority (24%) of education workers in the non-compliant group received at least one vaccine dose.

Table II.6: Descriptive statistics of workers in affected education industries that heavily faced COVID-19 vaccination mandates

(1) Characteristic	(2) Vaccinated regardless of mandate	(3) Vaccinated potentially due to mandate	(4) Unvaccinated or uncompliant with mandate
Number of workers	102,084	12,285	8,028
Had at least one vaccine (%)	100.00	100.00	23.65
Demographic			
Age (years)	42.09	36.39	40.48
Female (%)	83.02	84.69	86.06
Gender unknown (%)	.s	.s	.s
European (%)	66.57	50.79	57.47
Māori (%)	14.41	34.68	25.15
Pacific (%)	5.04	7.28	6.24
Asian (%)	10.76	4.64	6.54
MELAA/Other (%)	3.19	2.59	3.21
Ethnicity unknown (%)	.s	.s	.s
NZ-born (%)	69.38	81.39	68.95
Socioeconomic			
NZ Deprivation Index 2018	5.07	6.37	5.90
Monthly income in March 2021 ($)	4989	4167	4005
Monthly earnings in March 2021 ($)	4918	3954	3798
Monthly earnings from main job in March 2021 ($)	4829	3867	3712

Notes: This table presents demographic and socioeconomic characteristics of all individuals who were employed in education industries in March 2021, categorised by vaccination mandate compliance behaviour. Percentages may not always add up to 100 due to rounding. Notation ".s" means counts have been suppressed in accordance with Stats NZ confidentiality rules.

There is a bit more variation across compliance categories in age among education workers than the health and corrections workers. Education workers who would have likely vaccinated regardless of mandates are 42 years old on average, with those who did not comply were slightly younger, at 40 years old on average. However, those who were potentially vaccinated due to the mandates were younger, at 36 years old on average. Like the health industries, the large majority of all three groups are female, ranging 83-86%.

European and Asian education workers are more likely to have been vaccinated regardless of mandates, compared to the other two groups. Like the health and corrections workers, Māori and Pacific education workers are more likely to have potentially been vaccinated due to the mandates compared to groups one and three. Over 80% of education workers who were potentially vaccinated due to mandates were born in NZ, while this statistic is nearly 70% for the other two groups.

Like health and corrections workers, education workers who were vaccinated potentially due to the mandate have a higher deprivation score on average, compared with those that were vaccinated regardless of mandates or those who did not comply. However, unlike the health and corrections workers, education workers who were not compliant had the lowest monthly income and monthly earnings measures, compared to groups one and two.

6.2 Vaccination uptake over time

Figure II.4 shows the cumulative share of health, education and corrections workers who had received two vaccine doses over time, compared with workers in barely-mandated industries II.4. The vertical lines represent the mandate announcement date (11 October 2021 for all three mandated groups), the first dose deadline for HCW and education workers (15 November 2021, whereas it was slightly earlier on 6 November 2021 for corrections workers) and the second dose deadline for HCW and education workers (1 January 2022, whereas it was 8 December 2021 for corrections workers) respectively.

HCWs had the earliest and fastest rate of vaccine uptake, with a fast initial uptake once the nationwide vaccination roll-out began, followed by a gradual increase to a 'steady state' vaccination rate. This pattern is perhaps unsurprising given they were one of the earliest groups to gain access to the vaccine as part of the roll-out strategy to manage the initially limited vaccine supply. Recall that HCWs had ready access to the vaccine from March 2021, whereas the vaccine availability was rolled-out by age group categories for the general population, with it being widely accessible to all those aged 12 and over from September 2021. Moreover, we would expect health

Figure II.4: Cumulative double-vaccination rate by industry

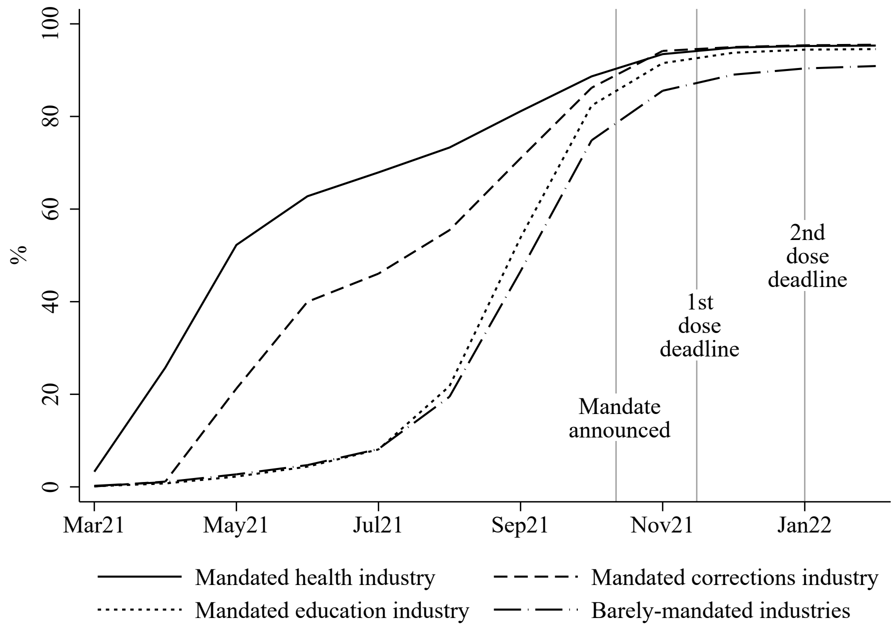

workers to have a higher propensity to vaccinate regardless of the mandates than the general population. By October 2021, when the vaccine mandate was announced, HCWs' double-vaccination rate had already reached just over 89%, and this increased gradually after the announcement to level off at about 95%. Visual inspection suggests there was not a discontinuous jump in the vaccination rate after the announcement of mandates, and the increase was part of an ongoing but slowing upwards trajectory.

Uptake among correction workers followed a similar trajectory as health workers, but the fast initial uptake began later from May 2021. This likely reflects that they also had early access to the vaccination, but the roll-out for this group started later, in May 2021. The slowing uptake in June likely reflects that the roll-out for this group was suspended in June 2021 in order to manage limited vaccine stocks, before resuming again in July 2021. Similar to health workers, there does not appear to be a discontinuity in the vaccination rate following the mandate announcement. By October, the double-vaccination rate had reached 86% before levelling off at about 95%.

The pattern for education sector workers is different, and is more similar to workers in the barely-mandated industries. Education workers did not have early access to the vaccine, which is likely reflected in the slow initial uptake followed by a sizable increase around August 2021, when the vaccination became more widely available to older age groups, followed by a larger increase around September 2021, when the vaccine became available to every aged 12 and over. Despite this later access to the vaccine, education workers had reached a double-vaccination rate of 82% by October 2021, with the rate levelling off at about 95%, which is very similar to the vaccination rate among health and corrections workers. The comparison with education and corrections workers suggests that the vaccination uptake over time, and particularly the pattern before the vaccine mandates were announced, is strongly linked to the availability of the vaccine rather than anticipatory effects.

Comparing HCWs with barely-mandated workers in Figure II.4 does, however, reveal an issue for DiD analysis: the parallel trends assumption is violated. HCWs' early access to the vaccine resulted in much faster uptake among HCWs than barely-mandated workers in the pre-treatment period. Indeed, the results for a DiD estimate as detailed in Equation 4 provide a negative δ coefficient, suggesting the mandates actually decreased vaccine uptake among health workers (see Appendix A). However, this is in line with Figure II.4, as the vaccine rate among HCWs was considerably higher than among barely-mandated workers by the time the mandates were announced in October 2021, and thus had less room to increase after the announcement. In contrast, many barely-mandated workers would have only gained access to the vaccine in September 2021, and thus, their vaccination rates were still on a stronger upwards trajectory.

While it is difficult to overcome this parallel trends issue for this research question,[18] provided there were no anticipatory effects, the fact that the HCWs' double-vaccination rate had already reached just over 89% before the mandate was announced, and that it levelled off at about 95% suggests that the mandate would have, at the very most, increased vaccination rates by six percentage points among HCWs. Moreover, given HCWs' vaccination rates were still on an upwards trajectory when the mandates were announced, it is likely that the effect of the mandates would have been less than this upper bound six percentage points. In addition, the lack of a discontinuous jump in vaccine rates around the time of the announcement is telling. This contrasts with international research examining vaccine passes (rather than vaccine mandates) for France, Italy and, to a lesser extent, Germany (Oliu-Barton et al., 2022), as well as Lithuania (Walkowiak et al., 2021). In these cases, there was a jump in vaccine rates after the vaccine pass was announced. Furthermore, vaccine rates in these countries were much lower before the announcement (less than 65% had received one dose at the time of the announcements in all of these countries), providing more room for vaccine passes in these countries to potentially increase vaccination rates than in the case of vaccine mandates in NZ.

[18]We considered the use of other comparison groups, but due to the early access to vaccines for health and corrections workers and a likely higher propensity to vaccine among all three groups of workers regardless of the mandates, it is difficult to overcome the parallel trends issue.

7 Results: RQ2 - Healthcare workers labour market outcomes

We now turn to our second research question: did the vaccine mandates impact the labour market outcomes of unvaccinated healthcare workers? In particular, did they increase healthcare worker job separation rates? These questions are important in terms of workers' outcomes. They are also important in the context of ongoing HCW shortages experienced not only in NZ, but many countries.

7.1 Tracking workers' outcomes over time

As described in Section 4.2, we track four groups of workers over time. We track HCWs who complied with the vaccine mandate (144,087 'vaccinated HCWs') and those who did not (12,330 'unvaccinated HCWs'). We also track barely-mandated workers who were vaccinated within the vaccine mandate timelines (although they were not subject to the mandates) (1,068,726 'vaccinated barely-mandated workers'), and barely-mandated workers who were not vaccinated within the mandate timelines (174,099 'unvaccinated barely-mandated workers').

Panel A of Figure II.5 shows employment rates over time. By construction, these are 100% in March 2019 as we defined our population of interest at this date (see Section 4.2). The employment patterns are very different between vaccinated and unvaccinated HCWs even before the mandates were announced, with unvaccinated HCWs having lower employment rates over time. This is perhaps unsurprising given we know from Section 6.1 that vaccinated and unvaccinated have different characteristics - for example, unvaccinated workers tend to have lower earnings and therefore may have lower labour market attachment. Moreover, these differences seem to be related to vaccination status rather than industry of employment as the employment rate patterns for vaccinated HCWs and vaccinated barely-mandated workers, as well as unvaccinated HCWs and unvaccinated barely-mandated workers, are similar.

The employment rates of unvaccinated HCWs fell slightly faster than that of unvaccinated barely-mandated workers in the months leading up to the vaccine mandate announcement. This could be indicative of anticipatory effects. However, other evidence suggests that anticipatory effects were unlikely (e.g. a lack of media discussion of the possibility of mandates, as discussed in Section 4.1.3) and, in any case, it seems unlikely that individuals would leave their roles before they had to if they did not have another role to go to, unless there were extenuating circumstances. Indeed, a more likely explanation is that unvaccinated HCWs may have felt pressure and elevated workplace stress due to their vaccination status that was not felt to the same extent by unvaccinated barely-mandated workers. This conjecture is supported by qualitiative research which highlights that unvaccinated workers experienced ostracism at work and resulting high levels of workplace stress (Dewar et al., 2024). In any case, this pre-announcement employment effect is minimal.

Figure II.5: Tracking workers' labour market outcomes over time

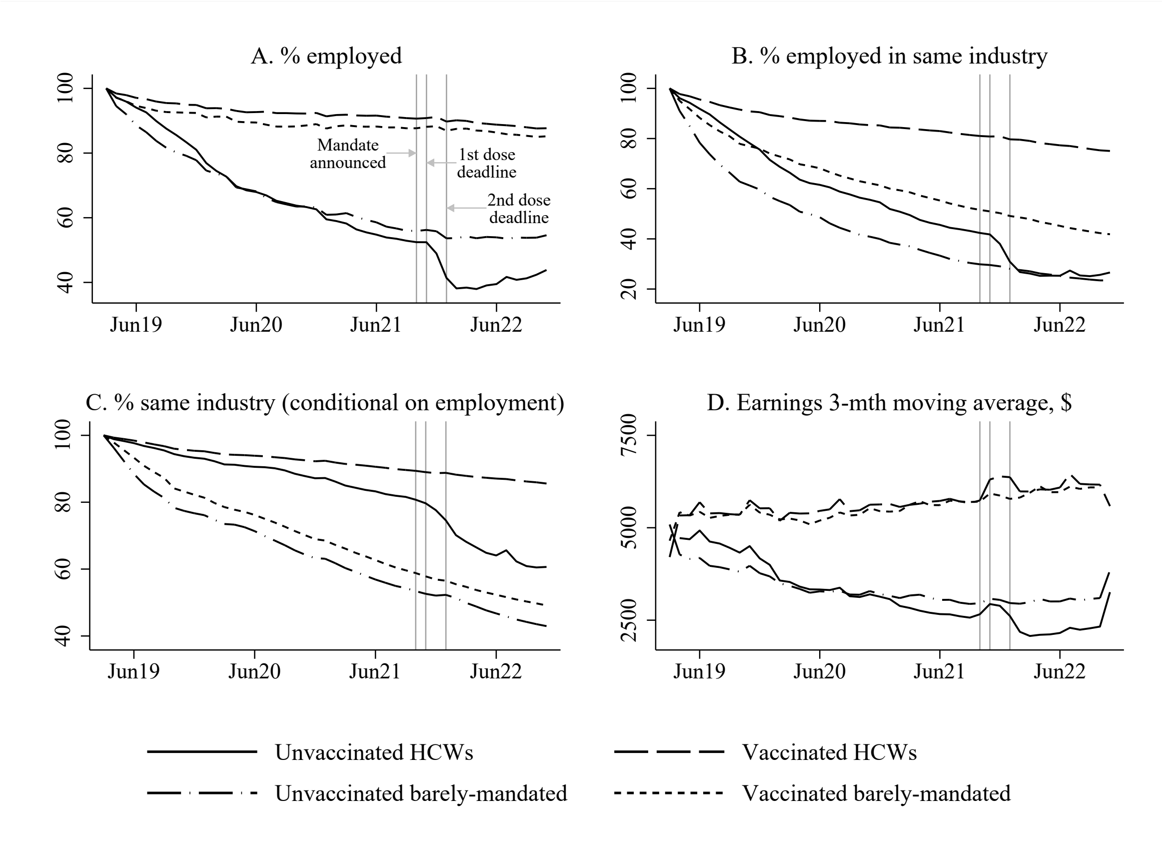

Notes: The vertical lines at October 2021, November 2021 and January 2022 represent the vaccine mandate announcement date, first dose compliance deadline and second dose compliance deadline respectively.

Little happened to the employment of unvaccinated HCWs immediately after the mandate was announced in October 2021. However, from November 2021, when the first dose requirement came into effect, the employment rate of unvaccinated HCWs dropped noticeably relative to those of unvaccinated barely-mandated workers. From early 2022, the employment rate of unvaccinated HCWs started to recover gradually, presumably as unvaccinated workers began new jobs in other non-mandated industries. In addition, the HCW mandate was lifted in late September 2022, towards the end of the time period examined, which would have allowed unvaccinated former HCWs to take up positions in the health industry again. However, as noted, there was no obligation for employers to reinstate them into their previous roles and, indeed, qualitative research suggests some unvaccinated HCWs had difficulties finding employment in the health industry even after the mandates were lifted (Dewar et al., 2024).

We now examine employment rates within the same industry (Figure II.5, Panel B). For HCWs, this measures whether the individual remained employed within the healthcare industry. For barely-mandated workers, this is whether the individual was employed within the same 1-digit ANZSIC industry (as explained in Section 4.2). Vaccinated HCWs had the highest propensity to remain employed within the same industry, while unvaccinated barely-mandated workers had the lowest. Of most relevance is that there was a distinct drop in same-industry employment for unvaccinated HCWs following the first dose deadline that is not observed for any of the three other groups. Same-industry employment for unvaccinated HCWs did increase slightly in the last few months of the series, possibly reflecting the lifting of the mandate in late September 2022. However, the slight increase began before the mandate was lifted and may reflect that some HCWs were being redeployed within the health industry to roles that had no contact with health practitioners or the general public, and, therefore, were not covered by the mandate. The fact that not all HCWs were subject to the vaccine mandate (see Section 4.1) also explains why employment within the healthcare industry for unvaccinated workers does not fall to 0%.

Part of the reason for the lower same-industry employment among unvaccinated workers shown in Panel B of Figure II.5 could be the lower employment rates among these workers (as shown in Panel A of Figure II.5). Therefore Panel C shows same-industry employment conditional on being employed. The smaller gap between unvaccinated HCWs (barely-mandated workers) and vaccinated HCWs (barely-mandated workers) highlights that some of the differences in Panel B are due to employment rate differences. However, the same general pattern of a large post-mandate drop in health-industry employment among unvaccinated HCWs that is not observed for the other three groups of workers is evident. Once again, the fact that unvaccinated HCWs experienced a faster fall in same-industry employment conditional on being employed than the vaccinated HCWs in the months leading up to the mandate announcement could signal anticipatory effects, but more likely reflects that there was pressure on unvaccinated HCWs due to their vaccination status even before the mandates were announced, leading to somewhat elevated job separation rates among this group.

Panel D of Figure II.5 tracks earnings, taking a three-period moving average to reduce fluctuations due to cyclical seasonality. Vaccinated HCWs and barely-mandated workers have similar earnings trends, which are higher than those of unvaccinated workers, which is consistent with their higher employment rates. Unvaccinated HCWs workers experienced an initial increase in earnings after the mandate was announced, followed by a sharp drop in earnings after the mandate announcement that was not experienced by the other groups of workers. The initial increase in earnings among unvaccinated HCWs likely reflects that many were receiving final pay cheques before leaving their jobs, which tend to be larger than a typical pay cheque due to factors such as the payment of outstanding holiday pay. However, average earnings of unvaccinated HCWs began to recover towards the end of the period, which likely reflects a combination of workers finding alternative employment and the mandate being lifted in late September 2022 so unvaccinated former HCWs could potentially return to the healthcare industry.

These descriptive trend graphs also provide further rationale for the use of a triple difference approach. This approach requires only one parallel trend to hold. Since the pairing of worker groups is arbitrary and mathematically equivalent, it does not matter which pairing the parallel trend holds for (Olden & Møen, 2022). That is, it can hold for unvaccinated HCWs and vaccinated HCWs, or unvaccinated HCWs and unvaccinated barely-mandated workers. The descriptive graphs show that the assumption of least one parallel trend holding is generally plausible, and it varies whether this is met through the unvaccinated HCWs and vaccinated HCWs comparison, or the unvaccinated HCWs and unvaccinated barely-mandated workers comparison. Since we are focussing on dynamic DDDs, the existence of pre-trends will be examined more systematically below in Section 7.2.

While Figure II.5 examines descriptively the main outcome variables that we will undertake DDD analysis on, it focuses on HCWs propensity to remain in the health industry given the mandate. However, it is possible that even if the mandate resulted in worse labour market outcomes, including a lower rate of employment among unvaccinated HCWs in general, and in healthcare specifically, the roles left vacant by unvaccinated workers could have been filled via new entrants into the health industry. As discussed, given the large extent of shortages of HCWs not just in NZ but globally, that many health roles require years of training and experience, and that NZ's borders were largely closed during the period of the vaccine mandate (thus limiting off-shore recruitment), it seems unlikely that new entrants into the industry would have offset the loss of unvaccinated workers.

However, to examine this descriptively, Figure II.6 looks at the rate at which vaccinated and unvaccinated barely-mandated workers switched to working in the health industry (conditional on employment), and the rate at which vaccinated and unvaccinated HCWs switched to barely-mandated industries. Prior to the mandates, unvaccinated HCWs had a somewhat higher propensity to leave the health industry and begin work in a barely-mandated industry than vaccinated HCWs. However, this difference increased markedly following the announcement of the vaccine mandate. There is little worker movement from barely-mandated industries to the health industry overall, and only a small amount of movement of unvaccinated barely-mandated workers to the health industry. Moreover, after the mandate announcement, this small amount of movement for unvaccinated barely-mandated workers slowed even further, as jobs in the health industry became largely closed off to them.

Figure II.6: Tracking workers' industry switching rates over time

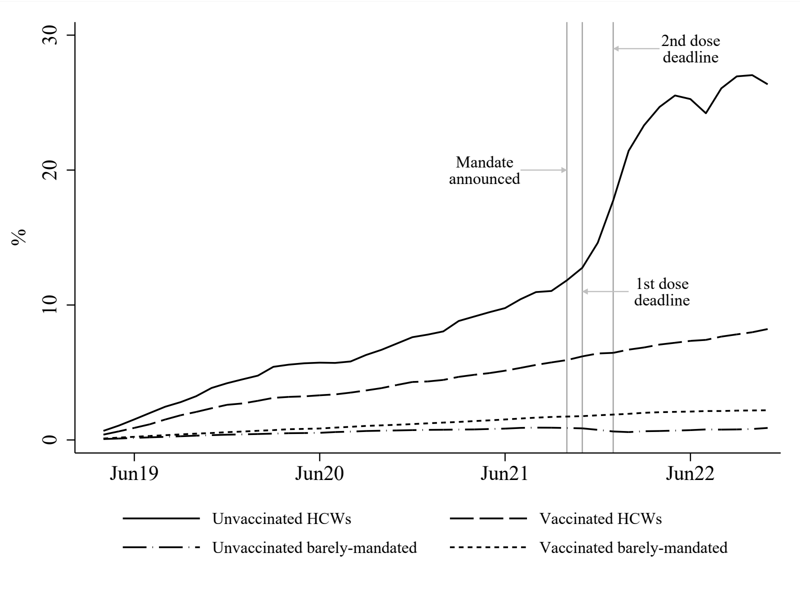

Notes: The vertical lines at October 2021, November 2021 and January 2022 represent the vaccine mandate announcement date, first dose compliance deadline and second dose compliance deadline respectively.

7.2 Triple difference: Estimating the role of mandates in HCWs' labour market outcomes

We now formally test the effect of mandates on the outcomes discussed descriptively in Section 7.1 above. Due to computational limitations, all estimates in this section are based on a 10% random sample of the population of interest equating to 3,891,945 individual-month observations (a balanced panel of 134,205 individuals over 29 months).

Figure II.7 plots the coefficient of interest (the triple interaction coefficient, γ_3, from Equation 6) as a percentage of the counterfactual. Month zero is October 2021, when the mandate was announced, as indicated by the first vertical line. The two additional vertical lines indicate the deadlines by which mandated workers were required to have one and two vaccine doses (November 2021 and January 2022 respectively).

Panel A of Figure II.7 presents employment rates. There are some significant differences in the employment rate of unvaccinated HCWs and the comparison group in the months before the mandate was announced. Moreover, because there was a general downward pre-trend the post-mandate difference may be somewhat overestimated, although the magnitude of these pre-trends are relatively small and are not statistically significant in the six months before the mandate announcement. Reassuringly, analysis for separate socioeconomic groups presented in Section 7.4 exhibit fewer pre-trends and the same general effects in the post-mandate period.

By the second month after the mandate was announced and the first month after the first dose deadline (December 2021), the employment rate of unvaccinated HCWs was lower than the comparison group. These post-announcement differences are statistically and economically significant. The employment rate is more than 17 percentage points lower (not shown) or 15% of the counterfactual employment rate in some months.

Figure II.7: Triple difference results: Change in outcome variable as percentage of the counterfactual

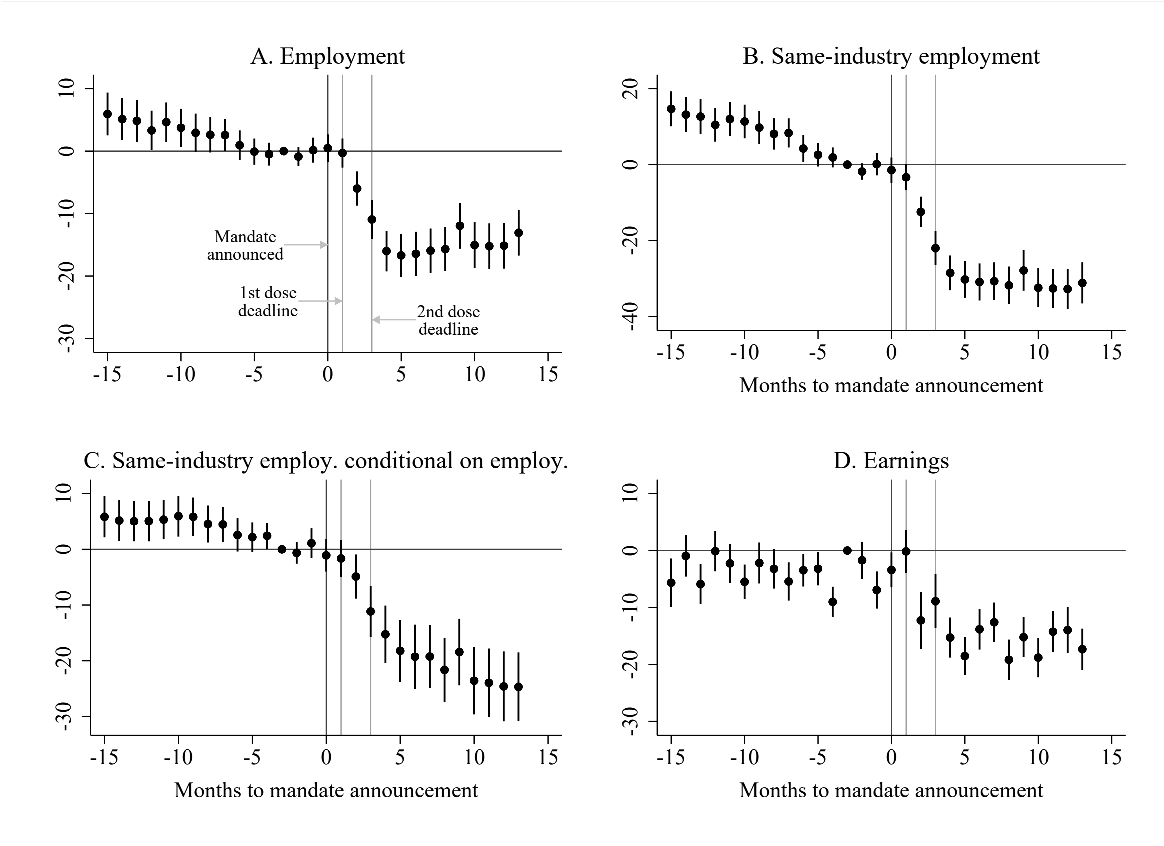

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

The pattern for same-industry employment is similar to that of overall employment, with unvaccinated HCWs much less likely to remain employed in the health industry after the mandate was announced (Figure II.7, Panel B). Unsurprisingly, the magnitude of the difference is larger than in the case of the overall employment rate, with the same-industry employment rate being up to 22 percentage points lower, or 33% of the counterfactual employment rate. The results for same-industry employment conditional on being employed are similar, although the magnitude of the effect is smaller (up to 17 percentage points or 25% of the counterfactual).

The earnings of unvaccinated HCWs is sometimes lower and statistically significant relative to the comparison group in the pre-announcement period. However, there is a noticeable drop in the earnings difference from the first month after the first dose deadline (December 2021). The earnings differential is statistically significant and negative in all months after the mandate deadline, and of economically significant magnitude (up to about $1,476 lower earnings in a month or 19% of the counterfactual earnings).

In summary, these results suggest that the vaccine mandates had a negative effect on overall employment rates, rates of employment within the health industry and workers' earnings.

7.3 Robustness: Two-period triple difference

As a robustness check, we also estimate two-period triple difference regressions (presented in Appendix B). The estimates are consistent with those from the dynamic triple difference regressions. The vaccine mandate is estimated to result in an employment rate which is 14 percentage points lower or -14% of the counterfactual, a same-industry employment rate that is 21 percentage points lower or 27% of the counterfactual, and a same-industry employment rate conditional on employment that is 13 percentage points lower or -17% of the counterfactual. Earnings are estimated to be about $700 lower or -11% of the counterfactual.

7.4 Heterogeneity analysis

The HCW vaccine mandate may have impacted different types of workers differently. For example, perhaps older unvaccinated HCWs were more likely to fall out of employment than younger unvaccinated HCWs, either because they had more difficulty transitioning to alternative employment, or because they were more likely to enter early retirement. To explore possible heterogeneity in the impact of the mandates, this section undertakes separate DDD analysis for different groups by gender, age, ethnicity, birthplace, deprivation level and income quartile. For brevity, we present only results for employment and earnings.

Gender

As mentioned, the pre-trends in the sub-group analysis tend to be smaller than in the overall analysis. Figure II.8 shows that for men, there are no statistically significant differences in employment before the mandate announcement. For women, the pre-trends are small and not statistically significant in the six months before the mandate

Figure II.8: Gender: Triple difference results: Change in outcome variable as a percentage of the counterfactual

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

announcement. For both men and women, significant differences emerge one month after the first dose deadline. The magnitude of the effect of the mandates is larger for men than women. For men, the employment rate is up to 20 percentage points lower, or -19% of the counterfactual. For women, the employment rate is up to 14 percentage points lower, or up to -14% of the counterfactual. This is somewhat unexpected given women are more likely to be secondary income earners and, therefore, presumably less attached to the labour market. However, it could be that unvaccinated female HCWs were more able to transition to alternative employment than unvaccinated male HCWs.

For earnings, the effect of the mandate is clearer for women than men. Unvaccinated male HCWs tend to have lower earnings than the comparison group even in many of the months prior to the mandate announcement, although the post-announcement differences are larger and more of them are statistically significant. For women, there are few pre-announcement periods with statistically significant differences, and statistically significant differences in earnings emerge one month after the first dose deadline.

However, consistent with the employment results, the change in post-announcement earnings is larger for men than women. For men, earnings are up to $2,636 lower, or -26% of the counterfactual. For women, earnings are up to $1,177 lower, or -19% of the counterfactual.

Age

Estimating separate results for younger (aged 20-39) and older (aged 40-60) workers shows that there are no statistically significant pre-announcement differences for younger workers, and only a few for older workers (Figure II.9). Employment rates for both younger and older workers fell after the mandate announcement. In terms of the magnitude of the effect, younger workers have a somewhat smaller fall in employment (up to -16 percentage points or -16% of the counterfactual) than older workers (up to -19 percentage points or -18% of the counterfactual). In addition, the employment rates of younger workers recover more over time, perhaps indicating they were better able to transition into alternative employment, or that older workers had lower labour market attachment (e.g. more likely to enter early retirement).

In terms of earnings (Figure II.9), while there are some statistically significant pre-announcement differences, there is a marked drop in earnings from one month after the first dose deadline for both younger and older workers. Once again, the effect was somewhat larger for older workers (up to -$1,791 or -20% of the counterfactual) than younger workers (up to -$1,384 or -18% of the counterfactual).

Ethnicity

Examining separate results for European, Māori and Pacific workers (Figure II.10), for employment rates, there are no statistically differences for any of these ethnic groups prior to the mandate announcement. After the mandate announcement, unvaccinated European and Māori HCWs experienced a drop in employment relative to the comparison groups. The magnitude of the effects of the mandates on employment among unvaccinated European and Māori HCWs is similar (a fall in employment of up to 22 percentage points for Europeans and 20 percentage points for Māori, equating to about -20% of the counterfactual for both). The Pacific worker results are different, with no statistically significant effects on employment of the mandates.

It is unclear why the unvaccinated Pacific HCWs did not experience a statistically significant drop in employment rates. Same-industry employment results (not shown) have a similar pattern of large, negative effects for unvaccinated European and Māori HCWs but no significant effects for Pacific HCWs. This suggests that the Pacific employment rate results were driven by unvaccinated HCWs remaining in the health

Figure II.9: Age: Triple difference results: Change in outcome variable as a percentage of the counterfactual

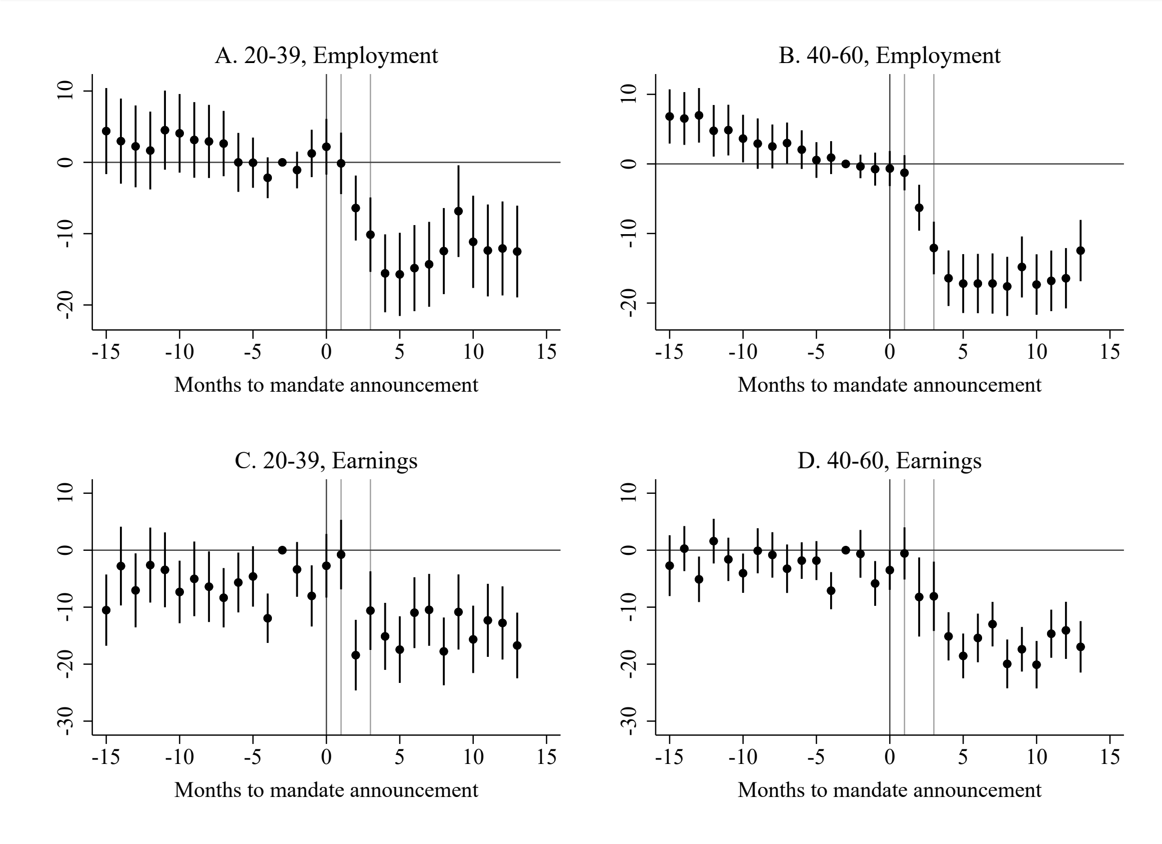

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

industry rather than having higher transition rates to other non-mandated industries. Thus, it could be that Pacific HCWs were more likely to work in roles that were not covered by the mandate (that is, they were not health practitioners and did not have roles that involved being within two metres of a health practitioner or member of the public).

In terms of earnings, the mandates had a negative effect on the earnings of European and Māori unvaccinated HCWs, with the effect for these two groups being of similar magnitude (up to -\$1,802 or 21% of the counterfactual for Europeans and -\$1,526 or -23% for Māori). In line with the employment results, there is much less of a clear drop in earnings after the mandate announcement for unvaccinated Pacific HCWs, and none of the differences are statistically significant.

Figure II.10: Ethnicity: Triple difference results: Change in outcome variable as percentage of the counterfactual

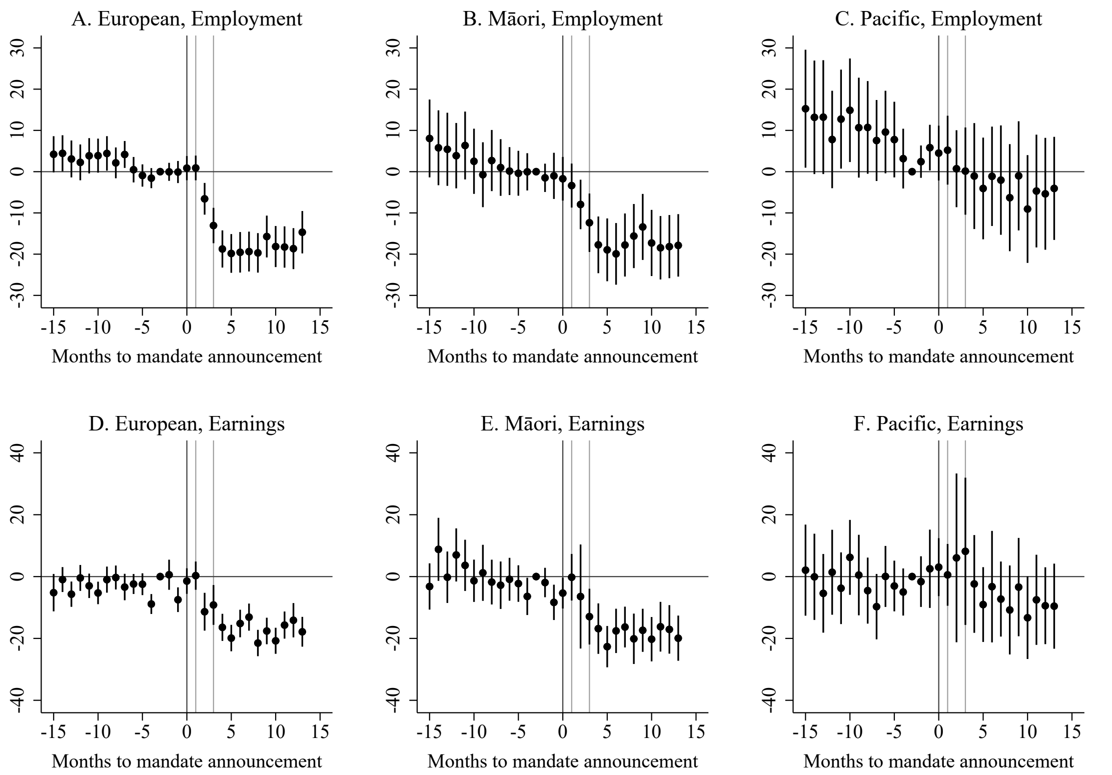

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

Born in NZ

Comparing those who were born in NZ with those who were not, Figure II.11 shows that both groups have a large drop in employment after the mandate was announced. The magnitude of the drop in employment is larger among unvaccinated HCWs who were born in NZ (-21 percentage points or -20% of the counterfactual) than those who were born overseas (-13 percentage points or -13% of the counterfactual). This larger effect for those born overseas is also true of same-industry employment (not shown). Thus, it appears that foreign-born HCWs more readily transitioned to employment in non-mandated industries, but also that they may have had roles within the health industry that were less likely to be covered by the mandate.

In terms of earnings, the effect of the mandate is clearer for unvaccinated NZ-born HCWs than those born overseas due to some pre-announcement negative earnings differences for those born overseas. However, both experience a drop in earnings after the mandate announcement. For unvaccinated NZ-born HCWs, the mandate resulted

Figure II.11: NZ born: Triple difference results: Change in outcome as percentage of the counterfactual

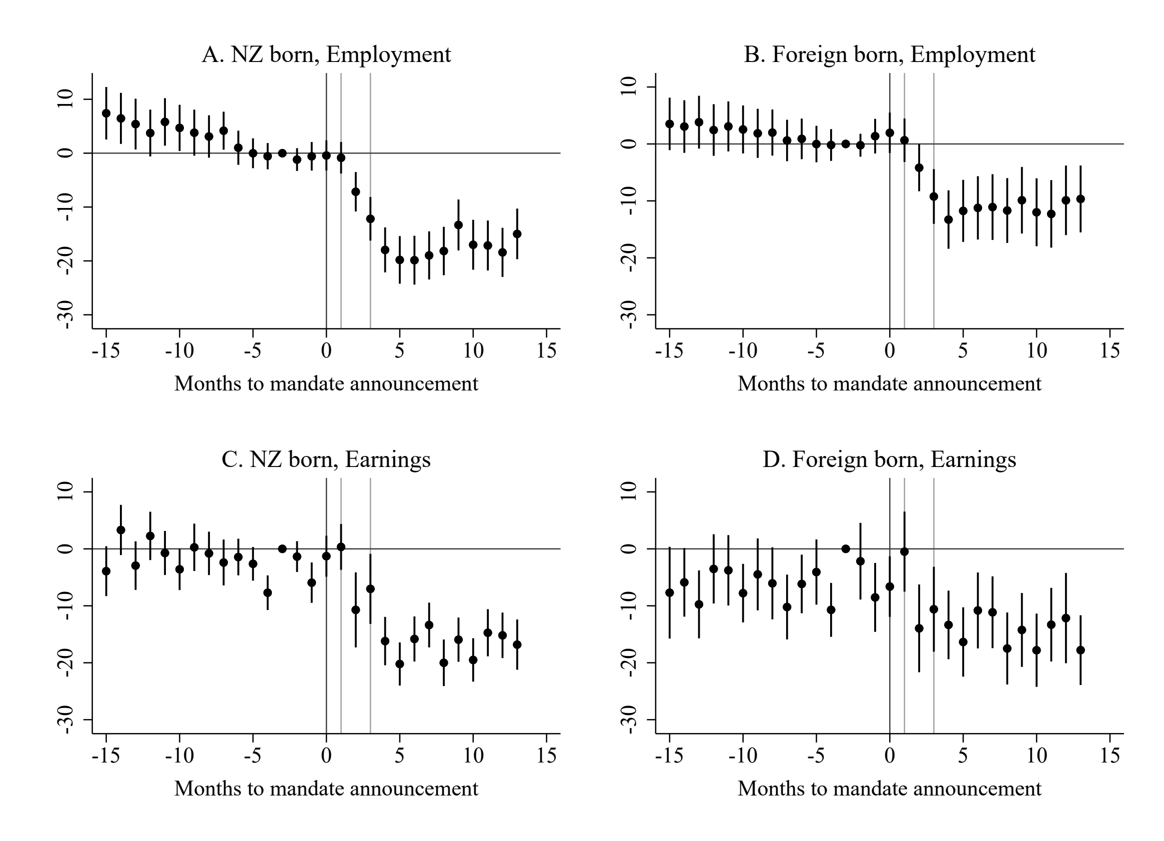

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

in an earnings drop of up to $1,593 (or -20% of the counterfactual). For unvaccinated foreign-born HCWs, the mandate resulted in an earnings drop of $1,411 (or -18% of the counterfactual).

Deprivation index

For both medium-to-high and low deprivation unvaccinated HCWs, the mandate resulted in a drop in employment. The effect is larger for those who had low measured deprivation (-16 percentage points or -16% of the counterfactual) than those with medium-to-high levels of deprivation (-21 percentage points or -20% of the counterfactual). The results for same-industry employment are of very similar magnitude for both groups (not shown). This suggests that those with low levels of deprivation were less likely to transition to alternative employment rather than being less likely to work in a health industry role that was not covered by the mandate. This may be because those with higher levels of deprivation had to find alternative employment

Figure II.12: NZ born: Triple difference results: Change in outcome as percentage of the counterfactual

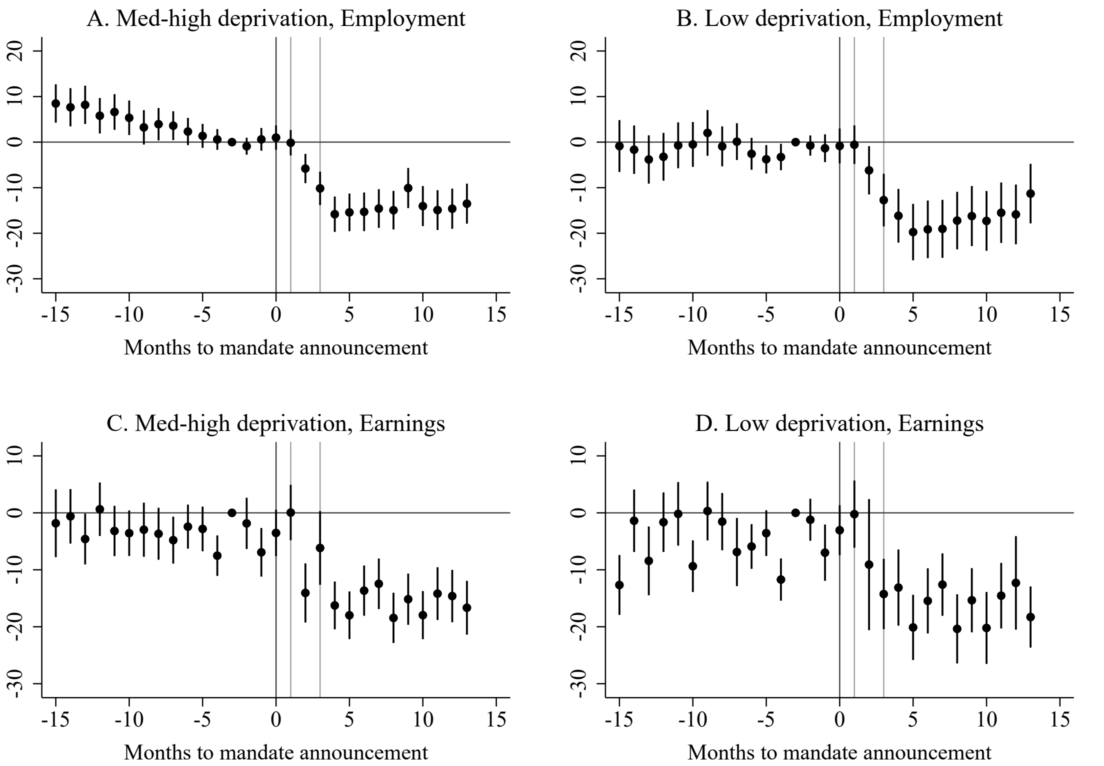

Figure II.13: Deprivation: Triple difference results: Change in outcome as percentage of the counterfactual

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

whereas those with low levels of deprivation were able to choose not to do so and exit employment instead.

Unvaccinated HCWs from both deprivation groups also experienced a drop in earnings after the mandate announcement. The effect was also similar in magnitude, with a \$1,294 drop (-18% of the counterfactual) for medium-to-high deprivation unvaccinated HCWs, and a \$1,929 drop for low deprivation unvaccinated HCWs (-20% of the counterfactual).

Income quartiles

We now undertake separate analysis by income quartiles. For all four quartiles, the mandate had a negative employment effect. However, the magnitude of the effects and the patterns over time differ. For unvaccinated HCWs in the lowest two income quartiles, the employment effect is larger than for the higher quartiles (-19% for quartiles 1 and 2 versus -15% and -16% of the counterfactual for quartiles 3 and 4). In addition, the employment rates of quartiles 3 and 4 begin to recover over time, while this is not observed for quartiles 1 and 2. This could reflect that those on higher incomes can more readily transition to roles in non-mandated industries. This would accord with literature that finds that workers with lower incomes and lower skills and/or qualifications are less resilient to adverse events, such as recessions (e.g., Shibata, 2021) and health shocks (e.g., García-Gómez et al., 2013).

Figure II.14: Income: Triple difference results: Change in employment as percentage of the counterfactual

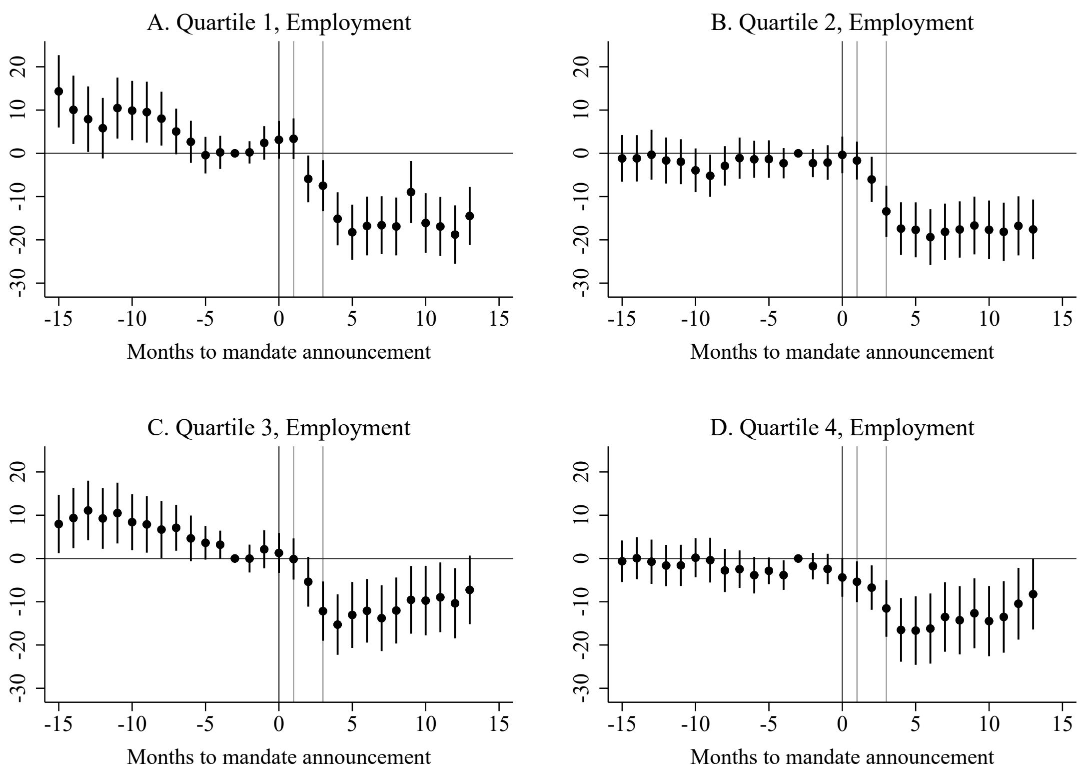

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

In terms of earnings, the vaccine mandate had a negative effect on the earnings of unvaccinated HCWs in all four income quartiles. The magnitude of the effect in terms of the dollar fall in earnings increases as income increases (from -\$737 for income quartile 1 to -\$3,358 for income quartile 4). However, this is due to differences in income levels across the quartiles as the effects in terms of a comparison with the counterfactual are of similar magnitude (-19% for quartiles 1 and 3, -21% for quartile 2 and -23% for quartile 4).

Figure II.15: Income: Triple difference results: Change in earnings as percentage of the counterfactual

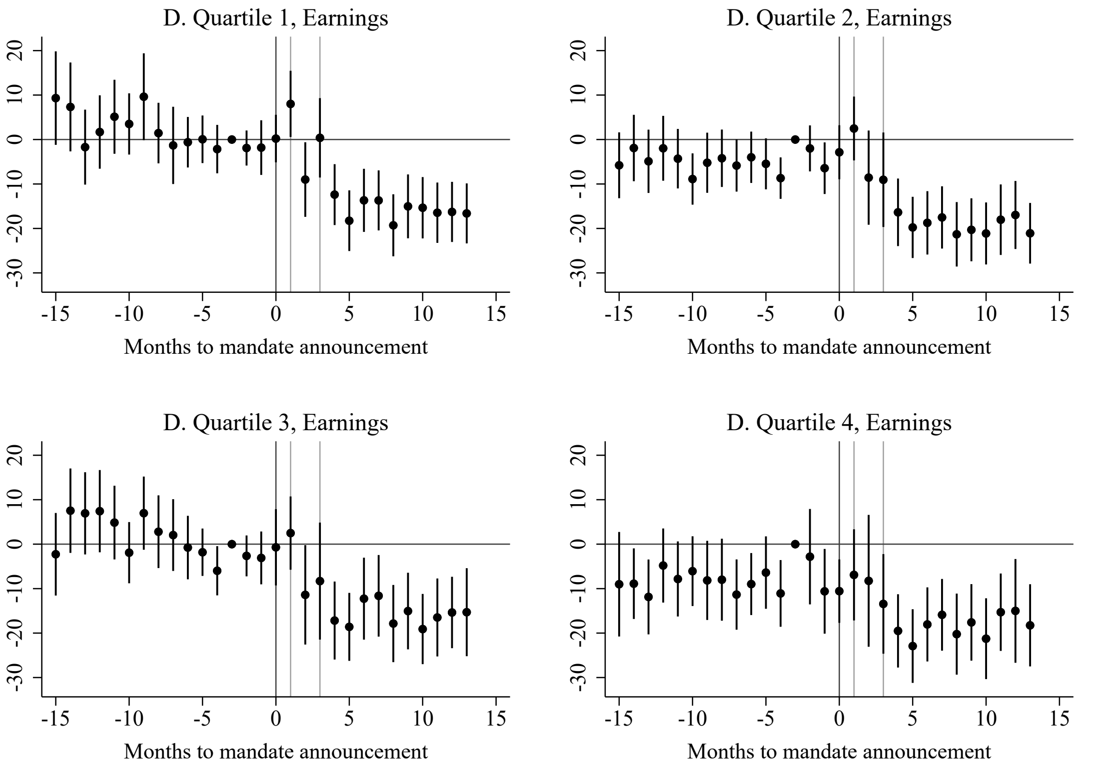

Notes: Estimates of the triple-interaction coefficient from Equation 6 with socioeconomic controls, as a percentage of the counterfactual. The vertical lines at time 0, time 1 and time 3 represent the vaccine mandate announcement date and 1st and 2nd dose compliance deadline respectively. The vertical bars around the point estimates are the 95% confidence intervals.

8 Policy discussion

Our findings suggest that vaccine mandates did little to increase the uptake of COVID-19 vaccinations given that uptake was already high in general, and particularly high among workers covered by the mandates. Moreover, they had a negative effect on the labour market outcomes of unvaccinated HCWs, which not only had consequences for the individuals involved, but also likely contributed to ongoing skills shortages in the health industry.[19] Therefore, while the stated purpose of the mandates was to ensure continuity of essential public services, they are more likely to have hindered rather than helped this goal. For future pandemic planning, this suggests that vaccine mandates should be used judiciously.

While beyond the scope of our analysis, another consideration is whether mandates could have the unintended consequence of crowding out vaccination willingness, with potential spillover effects to other vaccinations (Dubé et al., 2021). For example, the experimental economics literature finds that people are averse to control, with agents exerting more effort when the principal implements a high-trust, low-control system and less effort under a low-trust, high-control system (e.g., Burdin et al., 2018; Ziegelmeyer et al., 2012).

In terms of COVID-specific evidence, a representative panel survey in Germany found that mandates "substantially increase opposition to vaccination" (Schmelz & Bowles, 2022, p.1). This survey, with three waves starting in May 2021 (when the double-vaccination rate in Germany was less than 7%), found that few respondents were consistently opposed to being vaccinated if vaccinations were encouraged but remained voluntary (3.3%). However, a much higher share were consistently opposed to being vaccinated if vaccinations were mandatory (16.5%) (Schmelz & Bowles, 2022). Furthermore, the opposition to voluntary vaccinations was more transient - many of those who opposed in one survey wave changed their view in support of voluntary vaccination in later waves. In contrast, the opposition to mandatory vaccinations was more stable - the majority of those who opposed in one wave remained opposed in later waves. In addition, those opposed to mandated vaccinations had similar demographic and socioeconomic characteristics as the overall German population. However, what differentiated them was their level of trust in public institutions, their beliefs about vaccine efficacy, and whether they viewed mandated vaccinations as a restriction on their freedom. Thus, vaccine mandates may negatively impact people's sense of civic

[19] It would have been useful to additionally examine the impact on patient outcomes of the mandate policy. This could have been possible by analysing whether facility-level outcomes differed by the share of workers who left a particular facility due to the mandates. Unfortunately, an individual's place of employment and the health facility where a patient received treatment cannot be linked in the IDI. This is because the tax data assigns workplaces an employer ID, and health data assigns treatment facility IDs, but there is no way to link the two.

duty and the feel-good factor associated with "doing the right thing". Gibson (2022a) also highlights that an unexpected cost of the COVID-19 pandemic may be erosion of public confidence in all vaccines, which is partly driven by inflated claims about COVID-19 vaccine efficacy creating unrealistic expectations among the public about what the vaccines could achieve. Gibson (2023) further argues that public misunderstanding of COVID-19 vaccine trials may have contributed to NZ's adoption of vaccine mandates, despite the costs of doing so outweighing the benefits (Lally, 2021).

Another aspect to consider in terms of policy inferences is the NZ context and environment when the mandates were introduced. Whether vaccine mandates will be a useful tool going forward will largely depend on whether voluntary compliance will be as high in future pandemics as it was at the time of the mandate introduction in 2021. The high voluntary vaccination rate was likely driven by a combination of factors, including a generally strong sense of civic duty and high levels of trust in the government, as well as other "softer" policies to encourage vaccination, such as vaccine passes to access non-essential buisnesses/services, vaccination rate targets to end lockdowns, mass vaccine events and so forth. Our analysis suggests that when voluntary vaccination rates are high, the benefits of mandates are limited, and are likely outweighed by the spillover costs in terms of worsening health workforce shortages. However, it remains an open question as to whether these results would hold in circumstances where voluntary vaccination rates are lower. Indeed, international research on the introduction of vaccine passes in jurisdictions where vaccination rates were much lower than at the time of NZ's introduction of vaccine mandates suggests that such measures do contribute to increasing vaccination rates.

9 Conclusion

This paper examines the effect of workforce vaccine mandates on vaccination uptake and healthcare workers' (HCWs') labour market outcomes. We use linked population-wide administrative data from New Zealand, which includes a comprehensive national vaccination register linked to tax records to identify employment outcomes.

We employ a difference-in-differences approach to isolate the effects of workforce vaccination mandates from the effects of the NZ government's population-wide initiatives to boost vaccination rates, particularly vaccine passes to access non-essential businesses/services. However, no comparison group could be found where the parallel trends assumption held due to HCWs' early access to vaccinations. However, vaccination rates were already very high among mandated workers when the mandates were announced, leaving little room for vaccination rates to increase. Moreover, unlike international studies examining vaccine passes, there is no discontinuous jump in vaccination rates following the mandate announcement.

We additionally apply a dynamic triple difference approach (DDD) to examine healthcare workers' labour market outcomes, comparing unvaccinated HCWs with vaccinated HCWs and vaccinated and unvaccinated workers in industries that were not covered by workforce mandates. We find that the mandates negatively impacted on unvaccinated workers' overall employment rates, their rates of employment within the health industry and their earnings. While some groups, such as higher-income workers, saw some recovery in their labour market outcomes over time, the negative effects persisted for most groups of workers throughout the 13-month post-announcement period.

Overall, the results suggest that in the context of already-high vaccination rates, workforce vaccine mandates may not have provided much benefit in terms of increasing vaccination rates among mandated workers. Moreover, they came at a cost in terms of HCWs' labour market outcomes, which may have had wider negative consequences in terms of the supply of healthcare workers in an area where skills shortages were already an issue.

References

Abrevaya, J., & Mulligan, K. (2011). Effectiveness of state-level vaccination mandates: Evidence from the varicella vaccine. *Journal of Health Economics*, *30*(5), 966–976.

Anderberg, D., Chevalier, A., & Wadsworth, J. (2011). Anatomy of a health scare: Education, income and the MMR controversy in the UK. *Journal of Health Economics*, *30*(3), 515–530.

Bardosh, K., de Figueiredo, A., Gur-Arie, R., Jamrozik, E., Doidge, J., Lemmens, T., Keshavjee, S., Graham, J. E., & Baral, S. (2022). The unintended consequences of COVID-19 vaccine policy: Why mandates, passports and restrictions may cause more harm than good. *BMJ Global Health*, *7*(5), e008684. https://doi.org/10.1136/bmjgh-2022-008684

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, *119*(1), 249–275.

Burdin, G., Halliday, S., & Landini, F. (2018). The hidden benefits of abstaining from control. *Journal of Economic Behavior & Organization*, *147*, 1–12.

Carpenter, C. S., & Lawler, E. C. (2019). Direct and spillover effects of middle school vaccination requirements. *American Economic Journal: Economic Policy*, *11*(1), 95–125.

Carrera, M., Lawler, E. C., & White, C. (2021). Population mortality and laws encouraging influenza vaccination for hospital workers. *Annals of Internal Medicine*, *174*(4), 444–452.

Cohn, E., Chimowitz, M., Long, T., Varma, J. K., & Chokshi, D. A. (2022). The effect of a proof-of-vaccination requirement, incentive payments, and employer-based mandates on COVID-19 vaccination rates in New York City: A synthetic-control analysis. *The Lancet Public Health*, *7*(9), e754–e762.

Dewar, J., Barbarich-Unasa, T. W., Pacheco, G., Meehan, L., & Wilson, D. (2024). *Hidden behind a cloak of silence and exclusion: A qualitiative study of health professionals and mandated COVID-19 vaccinations (forthcoming).*

Dubé, È., Ward, J. K., Verger, P., & MacDonald, N. E. (2021). Vaccine hesitancy, acceptance, and anti-vaccination: Trends and future prospects for public health. *Annual Review of Public Health*, *42*, 175–191.

Duijster, J. W., Lieber, T., Pacelli, S., Van Balveren, L., Ruijs, L. S., Raethke, M., Kant, A., & Van Hunsel, F. (2023). Sex-disaggregated outcomes of adverse events after COVID-19 vaccination: A dutch cohort study and review of the literature. *Frontiers in Immunology*, *14*.

García-Gómez, P., van Kippersluis, H., O'Donnell, O., & van Doorslaer, E. (2013). Long term and spillover effects of health shocks on employment and income. *Journal of Human Resources*, *48*(4), 873–909.

Gibson, J. (2022a). Widespread public misunderstanding of pivotal trials for COVID-19 vaccines may damage public confidence in all vaccines. *Frontiers in Public Health*, *10*.

Gibson, J. (2022b). Government mandated lockdowns do not reduce COVID-19 deaths: Implications for evaluating the stringent New Zealand response. *New Zealand Economic Papers*, *56*(1), 17–28.

Gibson, J. (2022c). Hard, not early: Putting the New Zealand COVID-19 response in context. *New Zealand Economic Papers*, *56*(1), 1–8.

Gibson, J. (2023). Public misunderstanding of pivotal COVID-19 vaccine trials may contribute to New Zealand's adoption of a costly and economically inefficient vaccine mandate. *New Zealand Economic Papers*, *57*(1), 31–40.

Green, M. S., Peer, V., Magid, A., Hagani, N., Anis, E., & Nitzan, D. (2022). Gender differences in adverse events following the pfizer-BioNTech COVID-19 vaccine. *Vaccines*, *10*(2), 233.

Gur-Arie, R., Hutler, B., & Bernstein, J. (2023). The ethics of COVID-19 vaccine mandates for healthcare workers: Public health and clinical perspectives. *Bioethics*, *37*(4), 331–342. https://doi.org/10.1111/bioe.13141

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., & Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 government response tracker). *Nature Human Behaviour*, *5*(4), 529–538.

HorizonResearch. (2021). *COVID-19 vaccine: 28 october - 9 november 2021*. Ministry of Health. https://www.health.govt.nz/covid-19-novel-coronavirus/covid-19-vaccines/covid-19-vaccine-information-health-professionals/covid-19-vaccine-strategy-planning-insights/covid-19-vaccine-research-insights

Karaivanov, A., Kim, D., Lu, S. E., & Shigeoka, H. (2022). COVID-19 vaccination mandates and vaccine uptake. *Nature Human Behaviour*, 1–10.

Kung, S., Doppen, M., Black, M., Hills, T., & Kearns, N. (2021). Reduced mortality in New Zealand during the COVID-19 pandemic. *The Lancet*, *397*(10268), 25.

Lally, M. (2021). A cost-benefit anlaysis of covid-19 vaccine mandates. https://www.wgtn.ac.nz/__data/assets/pdf_file/0011/1983899/r-covid-vaccine-mandates.pdf

Lawler, E. C. (2017). Effectiveness of vaccination recommendations versus mandates: Evidence from the hepatitis a vaccine. *Journal of Health Economics*, *52*, 45–62.

Lee, C., & Robinson, J. L. (2016). Systematic review of the effect of immunization mandates on uptake of routine childhood immunizations. *Journal of Infection*, *72*(6), 659–666.

Lindley, M. C., Mu, Y., Hoss, A., Pepin, D., Kalayil, E. J., van Santen, K. L., Edwards, J. R., & Pollock, D. A. (2019). Association of state laws with influenza vaccination of hospital personnel. *American Journal of Preventive Medicine*, *56*(6), e177–e183.

McGarry, B. E., Gandhi, A. D., Syme, M., Berry, S. D., White, E. M., & Grabowski, D. C. (2022). Association of state COVID-19 vaccine mandates with staff vaccination coverage and staffing shortages in US nursing homes. *JAMA Health Forum*, *3*(7), e222363.

McKee, M., & Schalkwyk, M. C. I. v. (2022). England's U turn on covid-19 vaccine mandate for NHS staff. *BMJ*, *376*, o353.

Mello, M. M., Opel, D. J., Benjamin, R. M., Callaghan, T., DiResta, R., Elharake, J. A., Flowers, L. C., Galvani, A. P., Salmon, D. A., Schwartz, J. L., Brewer, N. T., Buttenheim, A. M., Carpiano, R. M., Clinton, C., Hotez, P. J., Lakshmanan, R., Maldonado, Y. A., Omer, S. B., Sharfstein, J. M., & Caplan, A. (2022). Effectiveness of vaccination mandates in improving uptake of COVID-19 vaccines in the USA. *The Lancet*, *400*(10351), 535–538.

Mills, M. C., & Rüttenauer, T. (2022). The effect of mandatory COVID-19 certificates on vaccine uptake: Synthetic-control modelling of six countries. *The Lancet Public Health*, *7*(1), e15–e22.

Olden, A., & Møen, J. (2022). The triple difference estimator. *The Econometrics Journal*, *25*(3), 531–553.

Oliu-Barton, M., Pradelski, B. S. R., Woloszko, N., Guetta-Jeanrenaud, L., Aghion, P., Artus, P., Fontanet, A., Martin, P., & Wolff, G. B. (2022). The effect of COVID certificates on vaccine uptake, health outcomes, and the economy. *Nature Communications*, *13*(1), 3942.

Oster, E. (2018). Does disease cause vaccination? disease outbreaks and vaccination response. *Journal of Health Economics*, *57*, 90–101.

Paipa, K., Hayward, S., Hamilton, K., & Leaoasavaii, D. (2023). *Review of the equity response to COVID-19: Final report for ministry of health*. Ministry of Health. https://www.health.govt.nz/publication/te-rau-ora-equity-review

Philipson, T. (1996). Private vaccination and public health: An empirical examination for u.s. measles. *The Journal of Human Resources*, *31*(3), 611–630.

Plummer, E., & Wempe, W. F. (2023). Evidence on the effects of the federal COVID-19 vaccine mandate on nursing home staffing levels. *Journal of the American Medical Directors Association*.

Rao, R., Koehler, A., Beckett, K., & Sengupta, S. (2022). COVID-19 vaccine mandates for healthcare professionals in the united states. *Vaccines*, *10*(9), 1425.

Rubenstein, B. L., Amiel, P. J., Ternier, A., Helmy, H., Lim, S., Chokshi, D. A., & Zucker, J. R. (2023). Increases in COVID-19 vaccination among NYC municipal employees after implementation of vaccination requirements. *Health Affairs*, *42*(3), 357–365.

Schaller, J., Schulkind, L., & Maghakian Shapiro, T. (2017). *The effects of perceived disease risk and access costs on infant immunization* (NBER Working Paper No. No. 23923). National Bureau of Economic Research. https://doi.org/10.3386/w23923

Schmelz, K., & Bowles, S. (2022). Opposition to voluntary and mandated COVID-19 vaccination as a dynamic process: Evidence and policy implications of changing beliefs. *Proceedings of the National Academy of Sciences*, *119*(13), e2118721119.

Schober, T. (2020). Effects of a measles outbreak on vaccination uptake. *Economics & Human Biology*, *38*, 100871.

Shibata, I. (2021). The distributional impact of recessions: The global financial crisis and the COVID-19 pandemic recession. *Journal of Economics and Business*, *115*, 105971.

Sprengholz, P., Betsch, C., & Böhm, R. (2021). Reactance revisited: Consequences of mandatory and scarce vaccination in the case of COVID-19. *Applied Psychology: Health and Well-Being*, *13*(4), 986–995.

Sprengholz, P., Felgendreff, L., Böhm, R., & Betsch, C. (2022). Vaccination policy reactance: Predictors, consequences, and countermeasures. *Journal of Health Psychology*, *27*(6), 1394–1407.

Stats NZ. (2018). *Kiwis perceive high political trust but low influence*. Retrieved February 14, 2023, from https://www.stats.govt.nz/news/kiwis-perceive-high-political-trust-but-low-influence/

Syme, M. L., Gouskova, N., & Berry, S. D. (2022). COVID-19 vaccine uptake among nursing home staff via statewide policy: The Mississippi vaccinate or test out policy. *American Journal of Public Health*, *112*(5), 762–765.

Tenforde, M. W., Self, W. H., Gaglani, M., Ginde, A. A., Douin, D. J., Talbot, H. K., Casey, J. D., Mohr, N. M., Zepeski, A., McNeal, T., Ghamande, S., Gibbs, K. W., Files, D. C., Hager, D. N., Shehu, A., Prekker, M. E., Frosch, A. E., Gong, M. N., Mohamed, A., ... IVY Network. (2022). Effectiveness of mRNA vaccination in preventing COVID-19-associated invasive mechanical ventilation and death - United tates, march 2021-january 2022. *MMWR. Morbidity and mortality weekly report*, *71*(12), 459–465.

Toshkov, D. (2023). Explaining the gender gap in COVID-19 vaccination attitudes. *European Journal of Public Health*, *33*(3), 490–495.

Walkowiak, M. P., Walkowiak, J. B., & Walkowiak, D. (2021). COVID-19 passport as a factor determining the success of national vaccination campaigns: Does it work? the case of lithuania vs. poland. *Vaccines*, *9*(12), 1498.

Woolf, K., Gogoi, M., Martin, C. A., Papineni, P., Lagrata, S., Nellums, L. B., McManus, I. C., Guyatt, A. L., Melbourne, C., Bryant, L., Gupta, A., John, C., Carr, S., Tobin, M. D., Simpson, S., Gregary, B., Aujayeb, A., Zingwe, S., Reza, R., . . . Pareek, M. (2022). Healthcare workers' views on mandatory SARS-CoV-2 vaccination in the UK: A cross-sectional, mixed-methods analysis from the UK-REACH study. *eClinicalMedicine*, *46*, 101346.

Yardley v Minister for Workplace Relations and Safety. https://www.courtsofnz.govt.nz/assets/Uploads/2022-NZHC-291.pdf

Ziegelmeyer, A., Schmelz, K., & Ploner, M. (2012). Hidden costs of control: Four repetitions and an extension. *Experimental Economics*, *15*(2), 323–340.

Appendices

A Difference-in-differences: Estimating the role of mandates in vaccination uptake

We estimate two-period DiD regressions to evaluate the extent to which the COVID-19 vaccination mandates increased vaccination uptake across HCWs, correction workers, and education workers, separately. The comparison group comprises workers in barely-mandated industries. Table II.A.1 presents the results. As mentioned in Section 5.1, the coefficient of interest in Equation 4 is δ (column (5) of Table II.A.1).

Table II.A.1 reveals a negative and significant effect of the COVID-19 vaccination mandates on vaccination uptake across all three mandated groups. For HCWs, the COVID-19 vaccination mandate led to a 28.3% decrease in vaccination uptake relative to barely-mandated workers. For corrections workers, the decrease was 12.5%, while for education workers, the decrease was 2.3%. All effects are highly significant and are robust to the inclusion of demographic and socioeconomic controls.

Table II.A.1: Difference-in-differences estimating the effect of COVID-19 mandates on vaccination update

(1) Treatment Group	(2) Comparison Group	(3) Treatment Indicator	(4) Post-period Indicator	(5) Treatment * Post
HCWs	Barely mandated workers	0.345*** (0.001)	0.419*** (0.001)	-0.283*** (0.001)
Corrections Workers	Barely mandated workers	0.170*** (0.004)	0.284*** (0.000)	-0.125*** (0.006)
Education Workers	Barely mandated workers	0.073*** (0.001)	0.419*** (0.001)	-0.023*** (0.002)

Notes: This table presents the difference-in-differences regression results from Equation 4 for HCWs, corrections workers, and education workers, respectively. Column 3 presents the coefficient on the treatment indicator. Column 4 presents the coefficients on the post-period indicator. Column 5 presents the coefficient on the interaction between the treatment and post-period indicators. Standard errors are in parentheses. Asterix represents statistical significance at conventional levels, where * if $p < 0.10$, ** if $p < 0.05$, and *** if $p < 0.01$.

As discussed in Section 6.2, these negative estimates are a consequence of the vaccination time trends presented in Figure II.4. Since mandated workers, particularly HCWs, had early access to the vaccine and, thus, a much faster vaccination uptake in the pre-treatment period, there was little room for mandate workers to increase their vaccination uptake after the mandates were announced, whereas barely-mandated workers had more room to increase uptake in the post-period. That is, the results are a consequence of the parallel trends assumption being violated.

B Two-period triple difference regression results

Table II.B.1: Two-period triple difference regressions

	(1) Employment (Δ employ. rate)	(2) Same-industry employment (Δ employ. rate)	(3) Same-industry employment conditional on employment (Δ employ. rate)	(4) Earnings (Δ $)
Unvaccinated	-0.112***	-0.083***	-0.023***	-830.45***
	(0.003)	(0.004)	(0.005)	(25.71)
HCW	0.041***	0.231***	0.221***	612.62***
	(0.002)	(0.003)	(0.003)	(33.73)
Unvaccinated * HCW	-0.024*	-0.073***	-0.019	-369.90***
	(0.013)	(0.015)	(0.013)	(97.29)
Post-period	-0.012***	-0.126***	-0.133***	453.10***
	(0.001)	(0.001)	(0.001)	(6.30)
Unvaccinated * Post	-0.039***	0.004	0.003	-382.02***
	(0.003)	(0.003)	(0.003)	(16.71)
HCW * Post-period	-0.011***	0.068***	0.093***	29.32*
	(0.002)	(0.002)	(0.002)	(17.25)
Unvaccinated * HCW * Post-period	-0.140***	-0.210***	-0.130***	-700.32***
	(0.012)	(0.013)	(0.014)	(73.58)
Socioeconomic controls	Yes	Yes	Yes	Yes
Individual-month observations	3,891,945	3,891,945	3,891,945	3,891,945

Notes: This table presents the two-period difference-in-differences regression results from Equation 5. Results for the outcome of interest of employment, same-industry employment, same-industry employment conditional on employment and earnings are presented in Columns 1, 2, 3 and 4 respectively. Standard errors are in parentheses. Asterix represents statistical significance at conventional levels, where * if $p < 0.10$, ** if $p < 0.05$, and *** if $p < 0.01$.

Paper III

The effect of a minor health shock on labour market outcomes: The case of concussions

Authors: Florian Fouquet, Lisa Meehan, Gail Pacheco and Alice Theadom

Prelude

Like the first two papers, this paper uses difference-in-differences (DiD) estimation applied to data from Stats NZ's Integrated Data Infrastructure. Also like the first two papers, it uses Inland Revenue tax data to measure labour market outcomes. while Paper 1 used a policy cut-off date to create treatment and comparison groups and Paper 2 uses a triple difference estimator, this paper uses a doubly-robust staggered DiD estimator.

This paper uses Accident Compensation Corporation data to identify the population of interest. Accident Compensation Corporation data provides information on medically-diagnosed mild traumatic brain injuries (mTBIs, commonly referred to as concussions). These data include not only hospital-treated mTBIs, but also those treated by primary care physicians. As mentioned, employment and earnings outcomes are measured using Inland Revenue tax data. Individual characteristics, such as age, gender and ethnicity come from Stats NZ's personal details table, which collates this information from across different IDI sources.

To address potential endogeneity, we use a comparison group of those who experienced the same mTBI shock but at a later date. One concern is that the individuals in our sample are not all treated at the same time, and the treatment effect may vary over time. To address this possibility, we apply a doubly-robust staggered DiD estimator.

Abstract

The literature on health shocks finds that minor injuries have only short-term labour market impacts. However, mild traumatic brain injuries (mTBIs, commonly referred to as concussions) may be different as the medical literature highlights that they can have longer-term health and cognitive effects. Moreover, TBIs are one of the most common causes of disability globally, with the vast majority being mild. Thus, it is important to understand the impact of mTBIs on labour market outcomes.

We use administrative data on all medically-diagnosed mTBIs in New Zealand linked to monthly tax records to examine the labour market effects of a mTBI. We use a comparison group of those who suffer a mTBI at a later date to overcome potential endogeneity issues, and employ a doubly-robust difference-in-differences method. We find that suffering a mTBI has negative labour market effects. Rather than dissipating over time, these negative effects grow, representing a decrease in employment of 20 percentage points and earning losses of about a third after 48 months. Our results highlight the need for timely diagnosis and treatment to mitigate the effect of mTBIs to reduce economic and social costs.

JEL: I10, I14, J01, J31

Keywords: health shock; traumatic brain injury; employment; earnings

1 Introduction

There is a growing economics literature on the impact of health shocks on labour market outcomes. Previous studies have shown negative effects of such shocks on both employment and income (e.g., García Gómez & López Nicolás, 2006; García-Gómez et al., 2013; Lenhart, 2019). However, there is limited evidence on the effects of mild traumatic brain injuries (mTBIs, commonly referred to as concussions) on subsequent labour market outcomes, despite the increasing awareness of the potential health and cognitive consequences of TBIs and the high incidence of such injuries. Moreover, while there are studies on mTBI in the health literature, very few examine the longer-term effects of mTBI on labour market outcomes, and virtually none account for potential endogeneity and attempt to establish causal links.

A traumatic brain injury (TBI) is defined in the International Classification of Diseases (ICD-10) as "*a traumatically induced structural injury or physiological disruption of brain function as a result of external force that is indicated by new onset or worsening of at least one of the following clinical signs immediately following the event: any alteration in mental status (e.g., confusion, disorientation, slowed thinking, etc.); any loss of memory for events immediately before or after the injury; any period of loss of, or a decreased level of, consciousness, observed or self-reported*". While there is increasing awareness of the impact of repeated concussions among athletes in contact sports, most mTBIs are caused by falls during everyday activities, with less than 30% occurring while playing sports (Accident Compensation Corporation, 2022). Thus, the risk of mTBIs is widespread.

In New Zealand, more than 36,000 people experience a TBI each year, with 95% of these cases being considered mild (Feigin et al., 2013). Worldwide, TBI is one of the most common causes of disability and death in adults and the leading cause of disability in children and young adults (Hyder et al., 2007; Langlois et al., 2006). Further highlighting the widespread risk of mTBI, in the US, the incidence of mTBI is more than five times higher than that of medium-severity TBIs and 32 times higher than that of high-severity TBIs (Miller et al., 2021). TBIs also carry significant costs - in terms of healthcare spending alone, the total estimated healthcare spending attributable to non-fatal TBI in the US was more than $40 billion in 2016, with the costs relating to mTBI being greater than more severe cases of TBI due to the higher incidence rate (Miller et al., 2021).

While apparently minor in nature, mTBIs may have persistent symptoms and long-term effects (Dean & Sterr, 2013). The exact mechanisms involved are complex and are not yet fully understood. However, van der Horn et al. (2020) summarises the existing evidence, highlighting the importance of the interplay between physiological and psychological processes. This evidence suggests that in the early stages, the physiological issues (e.g. cell injury and inflammation) and the acute stress response leads to neural network dysfunction. Patients often report symptoms such as fatigue, headaches, poor concentration and unstable moods. As time passes, psychological processes become more influential, and factors such as coping styles, personality, emotional regulation and the extent of other life demands come into play. Long-lasting symptoms can emerge, potentially leading to negative impacts on social integration, educational and labour market outcomes, and even to increases in antisocial behaviour (Theadom et al., 2023; Wehman et al., 2017; Williams et al., 2015). Moreover, while public awareness is growing of the issues caused by repeated mTBIs in contact sports, just one mTBI can have negative effects, further highlighting the potential implications for the general population (Theadom et al., 2023).

Given the persistent health and cognitive outcomes for mTBI cases, it is likely that labour market effects will differ from other forms of minor or temporary health shocks. Existing evidence suggests that temporary health issues (specifically those that last three months or less) do not have long-term effects on employment and income (Pelkowski & Berger, 2004), and that effects on labour market outcomes generally increase as the severity of the health shock increases (Crichton et al., 2011). Given this evidence, we could expect mTBI to have limited effects on future employment or earnings. However, mTBIs are, by their nature, potentially different from other minor injuries and may have longer-lasting effects.

Therefore, this paper examines (1) if mTBI has an effect on future employment and earnings; (2) if any effect is limited to the short term or if it remains significant after a longer time period; and (3) how the effects of mTBIs differ by age, gender, ethnicity or occupational skill level. In addition, we also explore the extent to which the accident compensation system offsets earning losses.

To this end, we use a staggered difference-in-differences framework (Callaway & Sant'Anna, 2021) that allows us to study the monthly effects of the health shock on earnings and employment, while taking into account the fact that not all individuals experience the shock at the same time. Because individuals who experienced a mTBI might have specific unobservable characteristics that affect both their risk of TBI and their labour outcomes, we follow Fadlon and Nielsen (2019) and construct counterfactuals for individuals suffering from a mTBI using individuals who experience one in the future. This is facilitated by New Zealand's linked administrative data, which includes Accident Compensation Corporation (ACC) injury data linked to Inland Revenue (IR) income data. These data allow us to estimate the monthly treatment effects on both earnings and employment up to 48 months after the shock.

These data provide advantages over existing studies. First, ACC data covers all medically-diagnosed mTBIs whereas other studies are typically reliant on hospital data, which is more limited in terms of identifying TBIs, particularly mild ones (Graff et al., 2019). Indeed, only about a fifth of individuals in our sample were treated at a hospital, with the majority being treated by a primary care physician. Moreover, IR income data are available on a monthly rather than an annual basis, allowing for the short- and medium-term effects of mTBIs to be better observed.

In addition to these data advantages, this paper offers methodological and policy contributions. First, there are few existing studies that examine the medium-term effects of mTBI on labour market outcomes (Graff et al., 2019; Theadom et al., 2017). Given that mTBI can, in contrast to most other minor physical injuries, have long-lasting effects, it is important to examine longer-term outcomes.

There also appears to be only one existing paper which analyses this issue using quasi-experimental methods to account for the potential endogeneity of suffering a TBI and labour market outcomes. Specifically, Fallesen and Campos (2020) use those who suffer TBIs at a later date as a comparison group, as our analysis also does. Relative to that paper, we make a number of further contributions to the literature. First, as mentioned, we use all medically-diagnosed mTBIs, which has advantages over the Fallesen and Campos (2020)'s use of Danish hospital and emergency room data to identify those who suffer a TBI. In addition, we use monthly rather than annual earnings. As our results will highlight, the use of monthly data provides a much more nuanced story about the effect of TBIs on employment and earnings, which also reinforces clinical observations. Moreover, Fallesen and Campos (2020) applies a two-way staggered DiD method, which may not be adequate to identify an average treatment effect when effects are heterogeneous (de Chaisemartin & D'Haultfœuille, 2020; Goodman-Bacon, 2021). Thus, we apply Callaway and Sant'Anna (2021)'s doubly-robust staggered difference-in-differences approach. We believe that there are just a handful of studies to date to apply this approach, and to the best of our knowledge our paper is the first to apply it to examine the impacts of a health shock.

We also make a contribution from a policy perspective. Given the high prevalence of mTBI and the fact that mTBI can, unlike most other minor physical injuries, have long-lasting effects, we add to the understanding of the potential costs to individuals in terms of their future earnings, and to society in terms of lost productivity and the potential increased burden on health and social welfare systems.

2 Literature review

This study relates to the economics literature on health shocks and subsequent labour market outcomes, as well as the health literature on the consequences of TBIs.

In terms of the economics literature on health shocks, numerous studies explore the links between health and labour market outcomes, whether investigating how unemployment and income can affect health or how health conditions alter labour market trajectories. There is previous evidence of a negative effect of job loss and unemployment on health (Eliason & Storrie, 2009; Sullivan & von Wachter, 2009). There is also a strong positive relationship between income and health, known as the "income gradient of health" (Case et al., 2002). However, these results may suffer from endogeneity and reverse causality issues. Therefore, more recent studies explored the reverse relationship, *i.e.* the effect of health on labour market outcomes.

Empirical studies investigating these effects examine various health shocks such as road accidents (Dano, 2005; Halla & Zweimüller, 2013) or sudden hospitalisations (García-Gómez et al., 2013; Lindeboom et al., 2016), as well as self-assessed health indicators (Contoyannis & Rice, 2001; Lenhart, 2019; Riphahn, 1999) and disability (Lechner & Vazquez-Alvarez, 2011). In general, these studies find negative effects of adverse health shocks on both employment and earnings (García-Gómez et al., 2013; Halla & Zweimüller, 2013). In general, studies have also found that whether the negative labour market effects persist over time depends on the severity of the injury. For example, Crichton et al. (2011) shows that more severe health shocks have stronger effects on labour market outcomes, and Pelkowski and Berger (2004) finds that temporary illnesses do not have any significant long-term effects on employment and income.

Findings also highlight differences between demographic groups in terms of the impact of health shocks. Riphahn (1999) and Pelkowski and Berger (2004) find a larger effect on income for women, but a stronger effect on employment for men. Similar effects on employment are found in Dano (2005), Lenhart (2019) and Zucchelli et al. (2010), and on earnings in Crichton et al. (2011). However, Contoyannis and Rice (2001) and Dano (2005) find that only men experience long-term effects on earnings. Many studies also find stronger effects on employment for older workers (Crichton et al., 2011; García-Gómez et al., 2013; Halla & Zweimüller, 2013), often related to the higher severity of the accidents and a lower ability of this population to recover from them. Part of the larger negative effect on employment with age can also be attributed to the opportunity of early retirement (Zucchelli et al., 2010). On the contrary, young workers seem to experience higher and more persistent income penalties (Halla & Zweimüller, 2013), maybe because of a negative effect on productivity but lower effects on employment. Finally, individuals with lower income prior to the shock suffer larger detrimental effects (Crichton et al., 2011; Dano, 2005; García-Gómez et al., 2013; Riphahn, 1999).

The institutional context also has an influence on labour market outcomes after health shocks. For example, Lechner and Vazquez-Alvarez (2011) show that while health shocks reduce the probability of being employed and decrease labour earnings, negative effects on income are (at least partially) compensated by disability benefits in countries with more protective social security systems. García-Gómez (2011) also show that the negative effect on employment can be reduced in the presence of quotas for disabled workers.

To our knowledge, there is no existing research in the economics literature specifically examining TBIs and subsequent labour market outcomes. However, there is health research on this topic. For example, using Danish hospital data, Graff et al. (2019) finds a strong negative effect on employment five years after the shock and Fallesen and Campos (2020) finds large income penalties. Using New Zealand data on 245 adults from the Brain Injury Incidence and Outcomes New Zealand in the Community longitudinal study, Theadom et al. (2017) shows that more than 15% of the individuals had exited the labour force four years after experiencing a mTBI, and those who remained in employment suffered work limitations and productivity losses.

An issue in this literature is the ability to attribute the outcomes to the TBI given that there may be a correlation between generally unobservable individual characteristics, such as risk preferences, and both the likelihood of suffering a TBI and labour market outcomes. As far as we are aware, there is only one existing study which addresses this issue using quasi-experimental methods. Fallesen and Campos (2020) uses a comparison group of those who suffer a mTBI at a future time point, following Fadlon and Nielsen (2019)'s approach. It finds that suffering a mTBI reduces average annual salary by 4.2%, mostly due to lower employment rates among those who suffered from a concussion. As discussed, we use a similar methodology, but are able to include all medically-diagnosed mTBIs, including those treated by primary healthcare providers and use monthly, rather than annual, employment and income data, allowing for a clearer differentiation between short-run and more persistent effects. Moreover, we use the method of Callaway and Sant'Anna (2021) to address the potential issue of bias in the application of standard two-way fixed effects regressions to staggered difference-in-differences analysis.

3 Empirical strategy

For causal estimates of the effects of mTBI on labour market outcomes, we need counterfactual outcomes of what would have happened to these individuals if they had not suffered from a mTBI. A simple comparison with those who have not suffered a mTBI would not enable any observed effect to be attributed to the mTBI. For instance, the (unobserved) characteristics of those who suffer from a mTBI are likely to be different from those who did not suffer from one (e.g. more likely to engage in risk-taking behaviour), and these characteristics are also likely to be correlated with labour market outcomes. This could lead to selection bias when comparing those who did and did not suffer from a mTBI and an overestimation of the mTBI's effects. To address this concern, we use a quasi-experimental design and construct counterfactuals using individuals who experienced the same mTBI shock but in the future (as per Fadlon & Nielsen, 2019). Therefore, our control group is not composed of individuals who will never be treated, but instead, of individuals who are not treated yet at each point in time.

Focusing on mTBIs is not an arbitrary choice. Indeed, this kind of accident is likely to be sudden and unpredictable, ensuring the randomness of the timing that we need in our strategy. Moreover, mTBIs are a very common injury and a major cause of disability and death according to the World Health Organization (WHO). They are, therefore, of particular interest in terms of public health policies and have consequences for other socio-economic outcomes. In addition, the high incidence of mTBI provides us with a large sample size which is likely to improve the robustness of our results.

One concern with our quasi-experimental design is that the individuals in our sample are not all treated at the same time, and that treatment effects are likely to vary over time. Therefore, results from a standard two-way fixed effects (TWFE) regression (i.e. a regression including both individual and time fixed effects) may be biased (Callaway & Sant'Anna, 2021; de Chaisemartin & D'Haultfœuille, 2020; Sun & Abraham, 2021) and cannot necessarily be interpreted as causal effects.[20] To address this issue, we rely on the doubly-robust difference-in-differences estimator proposed by Callaway and Sant'Anna (2021).

3.1 Model

We estimate the following staggered difference-in-differences (DiD) model:

$$Y_i^{g,t} = \alpha_1^{g,t} + \alpha_2^{g,t}.G_{i,g} + \alpha_3^{g,t}.1\{T = t\} + \beta^{g,t}.(G_{i,g} \times 1\{T = t\}) + \gamma.X_i + \epsilon_i^{g,t} \qquad (8)$$

where g denotes the groups, each group corresponding to all individuals starting to be treated at time G. $Y_i^{g,t}$ denotes the labour earnings of individual i in group g at time t, α_2 and α_3 are respectively group and time fixed effects, and X_i denotes the individual (time-invariant) controls. To differentiate between the intensive and

[20]For our analysis, TWFE estimations provide substantively the same results as the presented Callaway and Sant'Anna (2021) appraoch, although with small but consistent pre-treatment effects in the months immediately before the shock.

extensive margins, we also estimate the effects of mTBI on employment, using the same specification as in Equation 8, but with $Y_i^{g,t}$ the employment observed in individual i in group g at time t. Under limited anticipation and homogeneous treatment effects assumptions, the average treatment effects $ATT^{g,t}$ are given by $\beta^{g,t}$. However, they are not obtained through the standard $ATT^{g,t} = \mathbb{E}[Y_t(g) - Y_t(0)|G_g = 1]$, but are instead re-weighted using propensity scores. Following Callaway and Sant'Anna (2021), the average treatment effect of group g at time t can be written:

$$ATT(g;t;\delta) = \left[\left(\frac{G_g}{\mathbb{E}[G_g]} - \frac{\frac{p_{g,t+\delta}(X)(1-D_{t+\delta})(1-G_g)}{1-p_{g,t+\delta}(X)}}{\mathbb{E}\left[\frac{p_{g,t+\delta}(X)(1-D_{t+\delta})(1-G_g)}{1-p_{g,t+\delta}(X)}\right]} \right) \left(Y_t - Y_{g-\delta-1} - m^{ny}_{g,t,\delta}(X)\right) \right] \tag{9}$$

where $m^{ny}_{g,t,\delta}(X) = \mathbb{E}[Y_t - Y_{g-\delta-1}|X, D_{t+g}, G_g = 0]$ is the population outcome regression for the "not-yet treated" group at time $t + g$. Here, δ is the known duration of the anticipation period, and $p_{g,t+\delta}(X)$ is the probability of being first treated at time g conditional on covariates X and on either belonging to group g, so that $G_g = 1$, or belonging to the "not-yet treated" group at time $t + \delta$, so that $(1 - D_{t+\delta})(1 - G_g) = 1$.

Since we are interested in the monthly average effects of the mTBI and their dynamics, rather than the group-time ATT, we want to recover at each date the average effect of being treated for the group(s) that have been treated for exactly e time periods. This effect is given by:

$$\theta(e) = \sum_{g=1}^{\tau} 1\{g + e \leq \tau\} ATT(g, g + e) P(G = g|G + e \leq \tau) \tag{10}$$

where τ is the date of treatment of the last group.

We also compute the overall average treatment effects, which are analogous to the ATT in a standard two-period difference-in-differences model, and can be written:

$$\theta_\Delta = \frac{1}{\Delta + 1} \sum_{e=0}^{\Delta} \theta(e) \tag{11}$$

where Δ is the duration of the post-period. We compute these estimates for different values of Δ from 12 to 48 months. With $\Delta = 12$, we obtain the average treatment effect of the mTBI on labour market outcomes in the first year following the shock; with $\Delta = 24$, the average treatment effect of the mTBI in the first two years following the shock, etc.

3.2 Identification

This model is identified under five assumptions: (1) irreversibility of treatment, (2) limited anticipation, (3) random sampling, (4) parallel trends and (5) overlap (Callaway & Sant'Anna, 2021). Conditions (1) and (2) are verified by the sudden nature of the health shock we study. TBIs, by definition, cannot be unexperienced, and *a priori* cannot be anticipated. The random sampling assumption (3) is likely to be verified given the nature of our data. Indeed, we are able to observe all mTBI which have been treated by health services during the period (see Section 4 for more details).

The parallel trends assumption (4) is standard in DiD frameworks, although in this model, it should hold after conditioning on covariates. This allows for potential group-specific trends that would bias the estimates to be accounted for. Because we use not-yet treated individuals as counterfactuals to control for unobserved heterogeneity (Fadlon & Nielsen, 2019), and given the nature of the treatment, there is no reason to believe that individuals who experienced the TBI earlier are different from those who experienced it later. Therefore, this assumption is likely to hold even unconditionally (see the trends in Figure III.1). Nevertheless, we control for a set of demographic characteristics that could affect the outcomes. As will be shown, the average ATT observed for the pre-period is sometimes statistically significant, but economically negligible (less than NZ$5, see Table III.1), as well as the monthly ATT (Figure III.2), which confirms the validity of this assumption.

The overlap assumption (5) states that a positive share of the sample starts being treated at each time period g, and that, for each t and g, there is at least a small probability of not being treated, ensuring common support between the treated group and their counterfactuals. In other words, the actual date of treatment has to be uncorrelated with the probability of being treated. We use not-yet treated individuals as the control group to limit the presence of unobservables that would affect the probability of treatment and, given the suddenness of TBIs, the timing of treatment for this population can be considered as good as random, ensuring the overlap assumption is met.

4 Data and population of interest

The Integrated Data Infrastructure (IDI) is a large research database managed by Stats NZ. It holds population-wide linked micro-data about life events, including health events, income and other labour market information, from various government agencies and Stats NZ surveys.[21]

[21] See https://www.stats.govt.nz/integrated-data/integrated-data-infrastructure/ for more detailed information about the IDI.

We use multiple IDI data sources to conduct our analysis. We use information provided by the Accident Compensation Corporation (ACC) to identify individuals who experienced a mTBI between January 2012 and June 2022. For each of these individuals, we retrieve their wages and salaries for each month between January 2010 and December 2019 from the Inland Revenue (IR) data. Individual characteristics, such as age, gender and ethnicity come from Stats NZ's personal details table, which collates this information from across different IDI data sources.

NZ's ACC system is unique internationally and has some advantages over hospital and/or emergency room, or worker compensation data that has been used in previous studies. ACC is a compulsory, universal, no-fault compensation system that encompasses all accidents that occur in New Zealand. Compensation covers medical treatment costs and, for those in employment, income compensation of up to 80% of their pre-injury earnings for as long as the injury impacts on their ability to work. Treatment claims are lodged by medical providers rather than the injured individual, mitigating underreporting issues, and meaning coverage extends beyond hospital and emergency room treatment to, for example, primary health services. While unreported mTBIs are likely to still be an issue, the inclusion of mTBIs treated by primary health services means that ACC data are likely to capture more mTBIs than hospital data, particularly those of milder severity.

Our sample is comprised of individuals who experienced a mTBI between January 2012 and June 2022, identified using ICD-9 diagnosis code of 850, ICD-10 of S06.0 and/or ACC READ code of S6 (which is an ACC-specific code used to indicate the diagnosis of a concussion by a medical practitioner). Ideally, we would restrict attention to concussions with a short period or no loss of consciousness. While the ICD-9, ICD-10 and ACC READ codes include sub-categories of concussions which allow for these to be coded according to how long the loss of consciousness was, medical practitioners generally do not code the diagnosis beyond the higher-level 'concussion' classification. However, consistent with previous research (Feigin et al., 2013), ACC notes that the vast majority (95%) of TBIs are mild (ACC, 2017). Our data also indicates that the majority are initially treated by a primary care physician (79%) or other non-specialist medical practitioner, such as a nurse (83%), which further suggests that the majority are mild in nature.

We restrict our sample to individuals aged at least 25 years old at the beginning of the pre-treatment period (i.e. 24 months before they experience the mTBI), and younger than 65 at the time of the mTBI. We exclude from our analysis individuals who had experienced a prior TBI of any severity from 1994 (when ACC records begin) onwards. We also exclude individuals who experience multiple TBI during the study's period. This restriction means that any observed labour market effects can be more clearly attributed to the initial health shock, rather than the cumulative effect of the initial mTBI and any subsequent mTBIs suffered during the study period, particularly as health literature suggests that suffering a mTBI puts an individual at greater risk of suffering subsequent mTBI (Gils et al., 2020). We are therefore examining the effect of just one medically-diagnosed mTBI on labour market outcomes. Further, we exclude individuals who die before 2020 and those who lived overseas during the period (identified via Customs border movement data).

Most individuals in our sample are included in both the control and treatment groups, but at different points in time. This reduces the risk of unobserved differences between the treatment and control ("not-yet treated") groups. However, individuals who suffered the shock after December 2019 are only in the control ("not-yet treated") group. Although we observe the employment and earnings from January 2010, nobody in the sample is treated before January 2012, ensuring that we observe a pre-period of at least 24 months for each individual.

We include controls for gender, age at the time of the mTBI, ethnicity and fixed effects for region of residence at the time of the shock. Indeed, the inclusion of individual covariates should further reduce any remaining risk of unobservable differences between the treated and not-yet treated groups, which would result in important biases of our estimates if not taken into account (Callaway & Sant'Anna, 2021; Heckman et al., 1997). To avoid any reverse effect, we only include time-invariant demographics and pre-treatment covariates (Wooldridge, 2005).

We create a balanced panel of labour market outcomes for individuals in our population of interest. We measure employment and earnings outcome variables using IR monthly tax records. Monthly earnings are expressed in New Zealand dollars (NZ$), deflated to the prices of the first quarter of 2010 (2010Q1) using the consumer price index. One limitation of the data is that we cannot observe hours of employment. Therefore, we cannot assess how much of any change in the intensive margin is due to reduced hours versus reduced earnings per hour.

Overall, there are 35,301 individuals in the population of interest. Almost half (49.2%) are female and the average age is 43.5 years. In terms of ethnicity, 62.7% are European, 17.0% Māori, 5.3% Pacific peoples, 10.8% Asian, 2.3% Middle Eastern, Latin American or African, and 2.0% identify with other ethnicities. In terms of regional distribution, 28.3% live in Auckland, 13.7% live in Canterbury, 10.0% in Waikato, 9.6% in Wellington, 7.0% in the Bay of Plenty and the remaining 31.4% are spread across 11 smaller regions. About a third were employed in high-skilled occupations (Australian and NZ Standard Classification of Occupations (ANZSCO) skill levels 1 or 2) and two-thirds were employed in low-to-middle skilled occupations (ANZSCO skill levels 3, 4 or 5).

Figure III.1: Annual average of monthly earnings and employment, by year of treatment

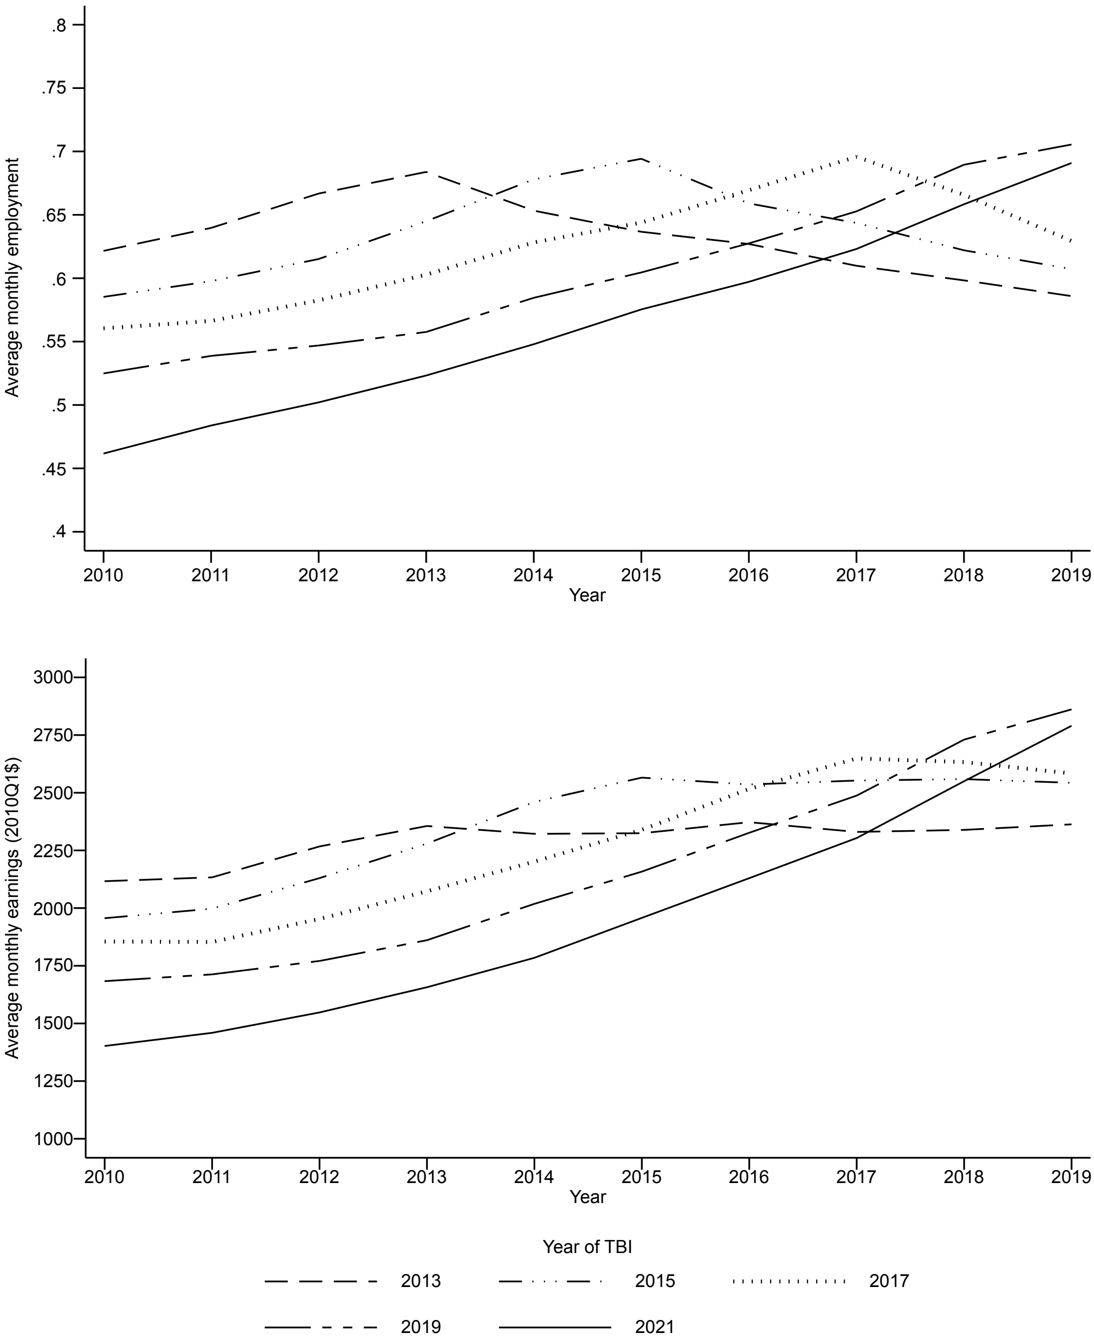

Note: These figures display the annual average of monthly earnings and employment for individuals experiencing a mTBI in years 2013, 2015, 2017, 2019 and 2021. For sake of clarity and since the pattern is similar, we do not report the average earnings and employment for individuals experiencing a mTBI in years 2012, 2014, 2016, 2018, 2020 and 2022.
Source: Authors' calculations using IDI.

Figure III.1 displays the trends in earnings and employment for the different cohorts (aggregated by year of treatment). Regardless of the year of treatment, we observe that both outcomes are increasing before the mTBI, supporting the parallel trends assumption. The absolute differences in values between cohorts might be related to demographic variation between cohorts (especially in terms of age), which argues for the inclusion of the individual covariates in our model. After the shock, we show a clear decrease in employment, while wages stagnate rather than decrease. This suggests that we should find a stronger effect of mTBI on employment than on earnings.

5 Results

We analyse the effects of suffering from a mTBI on future earnings and employment. Table III.1 presents the overall average treatment effects for different post-shock horizons (θ_Δ in Equation 11). We find a clear negative effect of the health shock on employment and earnings (Columns 1a and 1b). In the first year following the mTBI, the average monthly employment effect is -6.2 percentage points. In the four years following the shock, the average employment effect is -13.1 percentage points. The average monthly earnings penalty is around NZ$250 (NZ$500) in the first year (the first four years) after the shock, which is about 10% (20%) of pre-treatment average earnings.

Table III.1: Effects of mTBI on earnings and employment

	(1) Whole sample		**(2) Employed at time e**	**(3) Always employed**
	(a) Monthly earnings	(b) Employment	Monthly earnings	Monthly earnings
Pre-effect	4.63***	0.00***	-0.52	-1.39
	(0.55)	(0.00)	(0.79)	(0.95)
<u>*ATT*</u>				
$\Delta = 12$	-255.58***	-0.06***	-108.72***	-67.62***
	(9.11)	(0.00)	(8.84)	(17.75)
$\Delta = 24$	-346.55***	-0.09***	-114.66***	-67.38***
	(10.72)	(0.00)	(9.67)	(18.33)
$\Delta = 36$	-428.09***	-0.11***	-132.22***	-75.43***
	(12.33)	(0.00)	(10.80)	(19.46)
$\Delta = 48$	-505.74***	-0.13***	-155.88***	-89.23***
	(13.97)	(0.00)	(12.10)	(20.95)
N individuals	35,301	35,301	35,301	5,454

Note: Staggered difference-in-differences estimates are obtained using Equation 8. The effects reported in this table are calculated using Equation 11. Effects in Columns (1a) and (1b) are calculated for all individuals in the sample. Effects in Column (2) are calculated conditional on being employed at time e. Effects in column (3) are calculated on the subsample of individuals employed in all of the 120 periods of our data. Standard errors are clustered at the individual level. *** indicates significance at the 1% level, ** at the 5% level and * at the 10% level.
Source: Authors' calculations using IDI.

Since we are interested in distinguishing between the intensive and extensive margins, we also estimate the effects on earnings conditional on employment. Column 2 in Table III.1 reports the θ_{Δ} obtained by including only individuals in employment at time e in the calculation of $\theta(e)$. By comparison, column 3 reports the results obtained by restricting our sample to individuals who were employed in all of the 120 periods of our data (every month between January 2010 and December 2019). These results confirm the negative effect of mTBI on earnings. However, the magnitude of the penalties are lower than for the whole sample. Conditional on being employed at time e (always being employed), the penalty is around NZ\$110 (NZ\$70) in the year following the TBI, against NZ\$250 for the whole sample. Moreover, looking at longer time horizons, the average monthly penalty seems to increase slower for individuals who remain in employment. Therefore, most of the effect on earnings seems to be driven by a drop in employment. In order to confirm this, we next analyse how these labour market effects evolve on a monthly basis.

5.1 Monthly treatment effects

Figure III.2 displays the average monthly effects calculated using Equation 10. We find a negative and significant effect of the health shock on both employment and earnings (Figures III.2a and III.2b) starting immediately after the mTBI. Employment falls by three percentage points and earnings fall by NZ\$230 in the first month after the shock. The wage penalty remains reasonably stable during the first six months following the shock, before increasing strongly and continuously to reach NZ\$330 a month after one year and more than NZ\$800 after four years (Figure III.2b).

Therefore, although these injuries are considered to be minor and temporary in nature, mTBIs appear to have strong and persistent effects on earnings. This contradicts findings by Pelkowski and Berger (2004), which found that temporary illnesses generally have limited impact on labour market outcomes. However, it is in line with the medical literature about TBIs, which evidences longer-term (and sometimes hidden) health and cognitive effects, as well as employment effects, of mTBI (Dean & Sterr, 2013; Ribbers, 2007; Theadom et al., 2017).

Figure III.2: Monthly effects of mTBI on earnings and employment

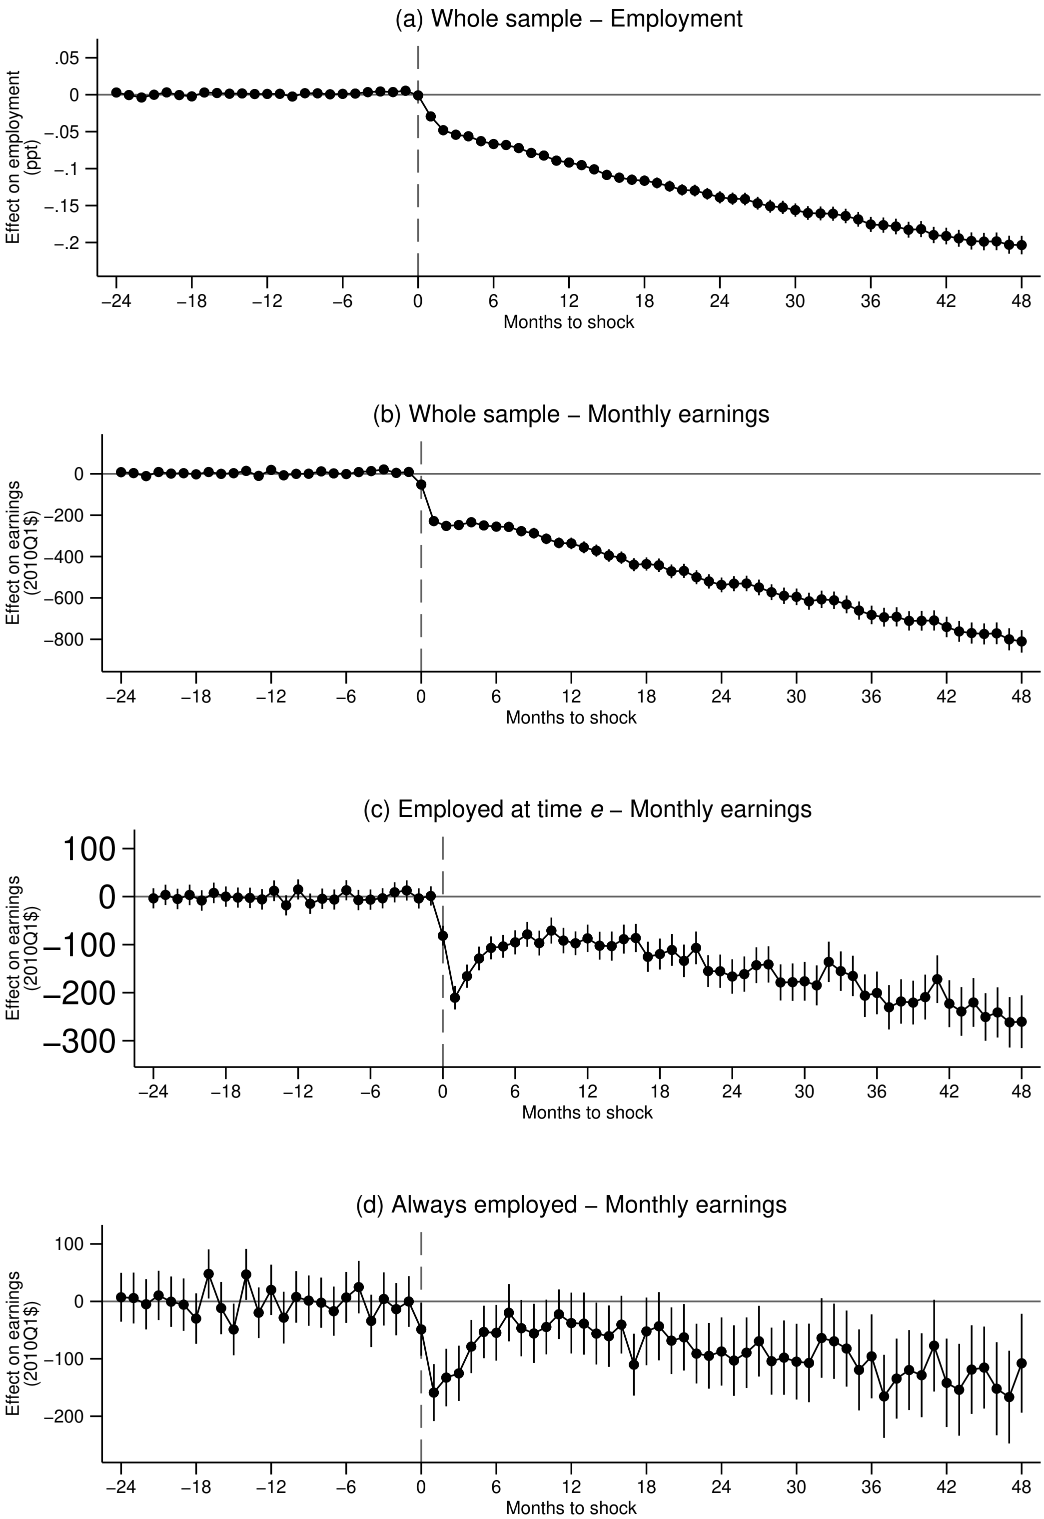

Note: These figures display the average monthly effects calculated using Equation 10 along with the 95% confidence intervals. Effects in (a) and (b) are calculated for all individuals in the sample. Effects in (c) are calculated conditional on being employed at time e. Effects in (d) are calculated on the subsample of individuals employed in all of the 120 periods of our data. Standard errors are clustered at the individual level.
Source: Authors' calculations using IDI.

The results for earnings conditional on employment (Figures III.2c and III.2d) confirm that the negative effect on earnings is mostly (but not entirely) due to a decrease in employment. This is illustrated by the muted effects on earnings in Figure III.2c (conditional on being employed at time e) and Figure III.2d (conditional on being always employed) relative to the impact on earnings evident for the full sample in Figure III.2b. This finding is consistent with previous evidence about brain conditions, such as strokes (Tanaka, 2021). The effect of mTBI on employment is smaller than the effect of strokes, with a decrease in employment of around 30% after four years while Tanaka (2021) reports a 55% decrease, but remains much higher than for other diseases, such as cancer (Heinesen et al., 2018). This might be related to cognitive difficulties resulting from brain trauma, that do not arise following other health shocks (Theadom et al., 2017). Figure III.2c (Figure III.2d) also shows that for those employed at time e (in each time period), earnings initially drop after the injury, then recover somewhat, before gradually falling again. Thus, wages of people who experienced a mTBI but remain more attached to the labour market, start stagnating in the longer term.

These results confirm that the strong longer-term effects on earnings for the full sample are mostly driven by transitions out of employment. For individuals able to remain in employment, there seem to be smaller effects of mTBI, resulting in only small decreases in either hours of work or productivity and thus lower wage penalties. These smaller effects are likely to be due to temporary part-time employment during the recovery period.

Another noteworthy feature of Figure III.2c is that for those employed at time e, earnings initially drop after the injury, then recover somewhat, before gradually falling again. We observe a similar pattern for individuals employed in each time period (Figure III.2d), with a sudden decrease in earnings right after the TBI, then a temporary recovery and finally a slow decrease after 18 to 24 months.This suggests that the wages of people who experienced a mTBI but remain more attached to the labour market and manage to stay employed, start stagnating in the longer term.

While there are few studies examining the longer-term effects of mTBI on labour market outcomes, our results accord with the limited research in this area. In particular, previous health research has found that while many of those who suffer mTBIs return to their previous work duties within a short timeframe, they often struggle to meet the demands of their employment in the longer term and may go on to exit employment or reduce their work intensity (Theadom et al., 2017). Individuals who suffer from a mTBI report having difficulties with losing their train of thought, concentration and sticking to a routine. A reasonable proportion (about one in six) report having to make changes to their working style to assist them in managing persistent symptoms at work (Theadom et al., 2017). While this reduced productivity upon returning to work may be manageable in the short term, over the medium term, individuals may be less able to manage their persistent symptoms and their employers are likely to become less accommodating over time. Indeed, individuals may return to work before they are fully recovered but they, and their employers, may have the expectation that they will gradually recover and return to their pre-injury levels of productivity. As it becomes evident that they are unable to undertake their pre-injury tasks with a similar level of efficiency, they may then begin to either exit employment, reduce their hours, or take on less demanding roles. This further highlights how mTBI may be different from other forms of minor injuries in terms of having longer-lasting effects.

An important caveat of our analysis is that our focus is on medically-diagnosed mTBI. As mentioned, this is an improvement on existing research that only examines hospital-treated mTBI. However, it must be recognised that there are likely to also be undiagnosed and untreated mTBIs in addition to the medically-diagnosed mTBI that we examine. The expected labour market impacts if the analysis included undiagnosed mTBIs is not clear. The impacts may be smaller on average if we assume that the undiagnosed mTBIs are less severe in nature. The impacts could also be larger on average if the diagnosed mTBIs receive more appropriate medical guidance and treatment relative to the undiagnosed mTBIs and thus have relatively better labour market outcomes in the longer term.

Also note that the persistent earnings losses also reflect that the estimate is relative to the comparison group of those who suffer TBI at a later date. The earnings of the comparison group continue to grow, thus contributing to the relative earning losses of the treatment group over time. Thus, even in cases where the earnings of individuals stagnate, rather than decrease in absolute terms, they will still be falling behind the comparison group whose earnings continue to grow.

5.2 Heterogeneity analysis

We estimate the heterogeneous effects on earnings and employment across gender, age and occupational skill level[22] at the time of the shock, and whether the individual suffered from a mTBI only or mTBI and other injuries (such as a limb fracture, for example). Note that information on occupation, and therefore occupational skill level, is only collected at the time of the injury. We, therefore, cannot assess whether changes in earnings over time are related to changes in occupation and an associated lowering of the required occupational skill level. Results (θ_Δ) obtained through Equation 11

[22] We use the Australian and New Zealand Standard Classification of Occupations (ANZSCO) classification to define our two skill groups. High skilled corresponds to ANZSCO skill levels 1 and 2, low to middle skilled corresponds to ANZSCO skill levels 3, 4 and 5.

are reported in Table III.2 for the different demographic subgroups.

Results by gender (Panel 1 of Table III.2) show that the penalty is slightly higher for men than for women. Regarding income, this might partly reflect that women earn lower wages on average than men. The year before they experienced the mTBI, women earned on average NZ$2,130 a month, versus NZ$2,790 a month for men. The average monthly penalty during the first year after the shock thus corresponds to a relative loss of 9% for women and 11% for men. However, while employment is similar for both genders the year before the shock (67.6% for women, 67.9% for men), men suffer from slightly higher employment penalties (around one percentage point higher) both in the short and longer terms. This could be because women have less consistent labour market attachment in general than men due to, for example, being more likely to spend time out of the labour force to raise children, thus leading to smaller differences between the treatment and comparison groups.

There is a strong heterogeneity in the effect of mTBI depending on the age of the individuals (Panel 2 of Table III.2). Wage penalties are higher for younger workers (under 40) than for older ones (aged 40+). In the first year following the shock, the average monthly income loss is around NZ$300 for younger workers, compared with NZ$230 for older ones. This absolute magnitude difference is also more stark in percentage terms given older workers have higher average pre-shock earnings than younger workers. This difference grows over time, with an income loss of around NZ$650 in the fourth year after the mTBI for younger workers and NZ$420 for older workers. Our results thus contradict those of Charles (2003), who found stronger negative effects of health shocks (transition to disability) on the earnings of older workers, which he associated with a destruction of the accumulated human capital and a shorter time horizon to recover from the shock. Our results show an even higher gap between age groups for employment, with a eight percentage point decrease in employment in the first year for younger workers compared with a five percentage point decrease for older workers. This employment effect grows to 18 percentage points in the first four years after the shock for younger workers, versus 10 percentage points for older workers.

Similar to the case of gender, the differences between younger and older workers might be explained by differences in the employment and earnings trajectories of the comparison group. For example, the earnings of younger workers tend to grow more strongly than those of older workers, who are more likely to have already entered their prime earning years. Thus, younger workers who experience a mTBI may lose an increasing amount of ground against the comparison group of other younger workers whose earnings are rising at a quicker pace.

It may also be that younger workers have a lower opportunity cost of exiting employment, at least in the short-term, due to lower average pay rates and potentially having lower financial responsibilities (e.g. less likely to have dependents, a mortgage etc.). A counterargument to this possibility is that the lifetime opportunity cost of time out of the workforce is likely to be higher for younger workers due to the greater potential that this will have a scarring effect. However, individuals may discount the future heavily or be myopic and therefore not fully factor in this consideration.

Another possible explanation is that younger people have lower labour market attachment in general, and lower attachment to their employer in particular. For example, older workers are likely to have higher longer job tenure than their younger counterparts. Thus, in the event of an accident that impacts their ability to undertake their job, older workers may have more support from their employer to return to work (e.g. more options to take on lighter duties for a time). Younger people may also be in more precarious work situations. For example, they may be less likely to be on permanent contracts and more likely to be on fixed-term or casual contracts. Therefore, more than an effect on the ability to keep one's job, mTBIs might have a negative effect on the probability of finding a subsequent job, whether because of reduced productivity (Theadom et al., 2017) or because of a scarring effect of having experienced such a health shock or having experienced a period of non-employment.[23] Unfortunately, our data do not contain information on whether individuals are employed on fixed-term or permanent contracts, preventing us from verifying this possibility.

In terms of occupational skill level, mTBIs are more detrimental for workers in high-skilled jobs than for those in low-to-middle-skilled ones (Panel 3 of Table III.2). Wage penalties in the longer term are about twice as high for the former in comparison with the latter (respectively -NZ$330 against -NZ$220 in the first year, -NZ$750 against -NZ$380 with a four years time horizon). Similar to what is observed by gender, the smaller effect on wages for low-to-middle skilled workers may (at least partly) be explained by lower pre-shock earnings on average. Indeed, with respective average earnings of NZ$3,300 and NZ$2,100 the year before the shock, the relative wage losses are similar at around 10% for both groups in the year following the mTBI. Nevertheless, in the longer term, the relative penalties increase more for high-skilled workers (around 36% against 29% after 48 months). Contrary to what we observe for age, this could be related to a destruction of human capital (Charles, 2003). Our results differ from the economics of health shocks literature which finds that those with lower income prior to the shock suffer larger detrimental effects (Crichton et al., 2011; Dano, 2005; García-Gómez et al., 2013; Riphahn, 1999). Similar results are found by Heinesen and

[23]To the best of our knowledge, there is no existing work about such an effect due to TBI. However, there is previous evidence of adverse effects of other health conditions on future employability, see e.g. Ameri et al. (2018), Drydakis (2010), and Hipes et al. (2016).

Kolodziejczyk (2013) for cancers, with individuals with lower education being more at risk of exiting employment. However, closer to our results, these authors also highlight a larger increase of the effect in the long term for more educated workers. Besides, Heinesen et al. (2018) show that the labour market effects of health shocks are highly dependent on the skill content of the jobs. In particular, they show that cancers have more negative effects in jobs requiring physical or manual skills. In our case, however, this could explain why we observe stronger effects of mTBI in high-skilled occupations. Indeed, higher skilled occupations are often more reliant on cognitive functions, which are the most affected by TBIs (Dean & Sterr, 2013). Indeed, our results align with those of Theadom et al. (2017), which finds that productivity losses following a mTBI are greater for mental and interpersonal tasks than for physical tasks.

It is expected that those who received the initial treatment for their mTBI in a hospital would have a more severe injury on average than those who were initially treated in a non-hospital setting (e.g. a primary healthcare practice) and, therefore, more negative employment and earnings outcomes. However, there is little difference in the employment consequences of the hospital-treated and non-hospital-treated (Panel 4 of Table III.2). Employment fell by seven percentage points over the first 12 months for the hospital-treated group, and six percentage points for the non-hospital-treated group. These employment losses had grown to 13 percentage points and 11 percentage points over the 48 months since the injury for the hospital and non-hospital treated respectively. The hospital-treated had larger earnings losses, although the magnitude of the differences are not large. The hospital-treated had earnings losses of about NZ$300 a month over the first 12 months, increasing to just over NZ$500 over the first 48 month. The non-hospital-treated had earnings losses of about NZ$250 over the first 12 months, increasing to NZ$500 over the first 48 months.

These smaller-than-expected differences may be because the difference in injury severity between hospital- and non-hospital-treated injuries is not great. For example, if hospital-treated mTBI is largely due to the timing of the injury - for example, if it is mainly those who were injured after hours or on the weekend when primary healthcare practices are generally closed. However, it could also be the case that the characteristics of those who are treated in the hospital are different. In particular, hospital treatment for accidents is free in New Zealand, while primary healthcare treatment is in general partially subsidised, with patients being required to make a co-payment towards their treatment costs. Moreover, some individuals do not have a primary care physician and may, therefore, have little choice but to go to a hospital emergency department. Since differences in primary healthcare access tend to be related to socioeconomic status, it is perhaps unsurprising that those who receive hospital treatment are less likely to be in high-skilled occupations (25% versus 36%

for those receiving non-hospital treatment). They are also a little less likely to be employed in the year before the mTBI (60% were employed for at least one month of the year, versus 63% for those receiving non-hospital treatment), and had lower earnings (NZ$2,014 average monthly earnings in the year before the mTBI, versus NZ$2,220 for those receiving non-hospital treatment). Another possible explanation is that those who are initially treated in a hospital setting have more severe mTBIs, but are provided with more specialised treatment and better follow-up care and guidance, thus leading to similar outcomes as those treated in non-hospital settings.

The results for those who suffered a mTBI only versus those who suffered a mTBI as well as other injuries as a result of the accident show that those who suffer a mTBI plus other injuries have higher earnings losses and falls in employment initially than the mTBI-only group, but these gaps decrease over time. For example, monthly earnings losses are on average 55% higher for the group with mTBI and other injuries compared with the mTBI-only group in the first 12 months, but over 48 months, this difference is just 8%. If examined on a monthly basis (Figure III.3), there are no statistically significant differences in earnings (employment) between the mTBI-only and multiple injury group after 9 (10) months. This convergence of effects between the two groups suggests that the initial shock is larger for those with multiple injuries, but over time, individuals recover from their other injuries (such as fractured limbs, for example), and the remaining longer-lasting effects are largely due to the mTBI. This result is consistent with the health shocks literature that temporary, minor injuries do not have a lasting impact on labour market outcomes, and further highlights that mTBIs are different and tend to have a lasting impact.

Figure III.3: Monthly effects of mTBI on earnings and employment

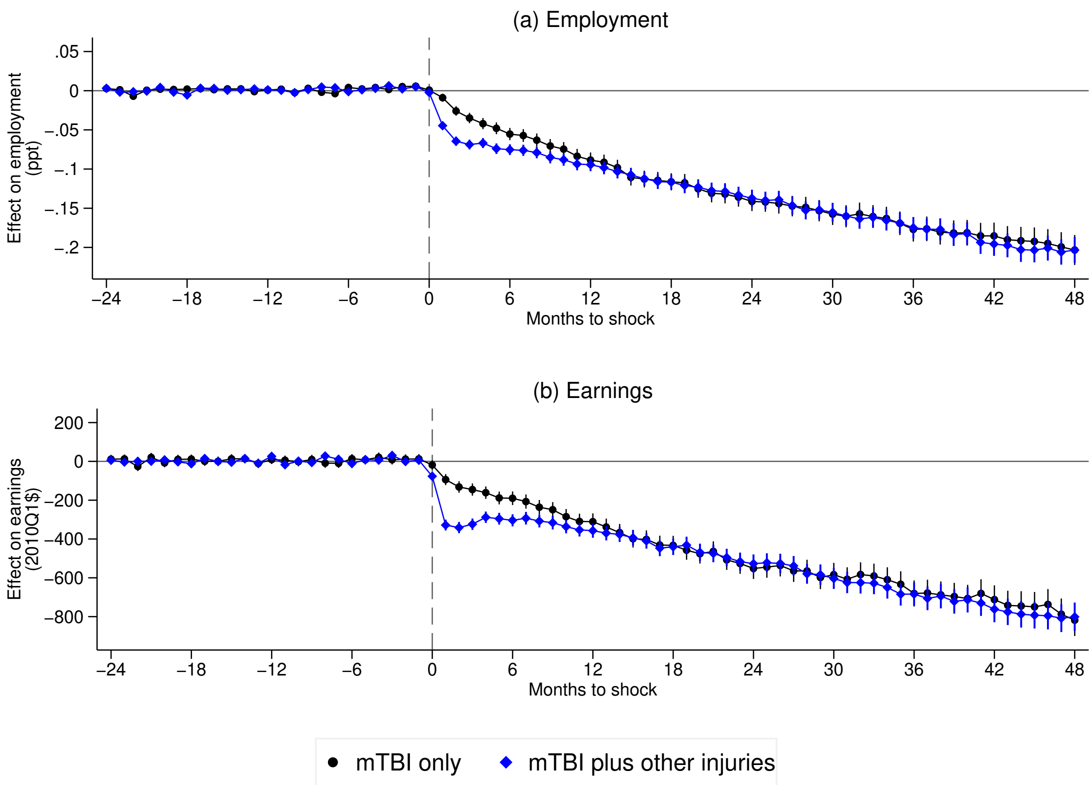

Note: These figures display the average monthly effects calculated using Equation 10 along with the 95% confidence intervals. Effects in (a) and (b) are calculated for all individuals in the sample. Effects in (c) are calculated conditional on being employed at time e. Effects in (d) are calculated on the subsample of individuals employed in all of the 120 periods of our data. Standard errors are clustered at the individual level.
Source: Authors' calculations using IDI.

Table III.2: Heterogeneous effects of mTBI on earnings and employment

Gender	**(1) Women** ($N = 17,385$)		**(2) Men** ($N = 17,916$)	
	(a) Monthly earnings	(b) Employment	(a) Monthly earnings	(b) Employment
Pre-effect	3.47***	0.00***	5.73***	0.00***
	(0.70)	(0.00)	(0.84)	(0.00)
ATT				
$\Delta = 12$	-192.47***	-0.05***	-313.65***	-0.07***
	(11.43)	(0.00)	(14.03)	(0.00)
$\Delta = 24$	-278.79***	-0.08***	-408.56***	-0.09***
	(13.60)	(0.00)	(16.36)	(0.00)
$\Delta = 36$	-356.29***	-0.11***	-493.84***	-0.12***
	(15.90)	(0.00)	(18.62)	(0.00)
$\Delta = 48$	-424.41***	-0.13***	-579.63***	-0.13***
	(18.12)	(0.01)	(20.96)	(0.00)
Age	**(3) Under 40** ($N = 14,952$)		**(4) 40+** ($N = 20,349$)	
	(a) Monthly earnings	(b) Employment	(a) Monthly earnings	(b) Employment
Pre-effect	9.65***	0.00***	1.11	0.00***
	(0.93)	(0.00)	(0.68)	(0.00)
ATT				
$\Delta = 12$	-299.52***	-0.08***	-226.11***	-0.05***
	(15.30)	(0.00)	(11.26)	(0.00)
$\Delta = 24$	-437.26***	-0.12***	-287.00***	-0.07***
	(18.26)	(0.00)	(13.10)	(0.00)
$\Delta = 36$	-548.33***	-0.15***	-350.21***	-0.08***
	(21.07)	(0.00)	(15.07)	(0.00)
$\Delta = 48$	-646.44***	-0.18***	-415.74***	-0.10***
	(23.88)	(0.01)	(17.12)	(0.00)
Skill level	**(5) Low- to mid-skilled** ($N = 23,505$)		**(6) High-skilled** ($N = 11,796$)	
	(a) Monthly earnings	(b) Employment	(a) Monthly earnings	(b) Employment
Pre-effect	3.90***	0.00***	6.65***	0.00***
	(0.63)	(0.00)	(1.07)	(0.00)
ATT				
$\Delta = 12$	-217.62***	-0.06***	-328.78***	-0.07***
	(10.60)	(0.00)	(17.36)	(0.00)
$\Delta = 24$	-274.18***	-0.08***	-486.88***	-0.11***
	(12.43)	(0.00)	(20.50)	(0.00)
$\Delta = 36$	-325.85***	-0.10***	-626.25***	-0.14***
	(14.15)	(0.00)	(24.03)	(0.01)
$\Delta = 48$	-380.95***	-0.11***	-748.34***	-0.16***
	(15.98)	(0.00)	(27.41)	(0.01)
Hospital	**(7) Hospital-treated** ($N = 7,890$)		**(8) Non-hospital treated** ($N = 27,411$)	
	(a) Monthly earnings	(b) Employment	(a) Monthly earnings	(b) Employment
Pre-effect	6.62***	0.00***	4.12***	0.00***
	(1.15)	(0.00)	(0.63)	(0.00)
ATT				
$\Delta = 12$	-291.13***	-0.07***	-245.40***	-0.06***
	(20.10)	(0.00)	(10.22)	(0.00)
$\Delta = 24$	-365.73***	-0.09***	-340.10***	-0.09***
	(23.26)	(0.01)	(12.06)	(0.00)
$\Delta = 36$	-437.03***	-0.11***	-423.25***	-0.11***
	(26.53)	(0.01)	(13.91)	(0.00)
$\Delta = 48$	-512.16***	-0.13***	-500.01***	-0.11***
	(29.81)	(0.01)	(15.78)	(0.00)
Multiple injuries	**(9) mTBI only** ($N = 14,567$)		**(10) mTBI plus other injuries** ($N = 20,734$)	
	(a) Monthly earnings	(b) Employment	(a) Monthly earnings	(b) Employment
Pre-effect	5.44***	0.00***	4.02***	0.00***
	(0.84)	(0.00)	(0.73)	(0.00)
ATT				
$\Delta = 12$	-194.31***	-0.05***	-301.48***	-0.07***
	(13.55)	(0.00)	(12.28)	(0.00)
$\Delta = 24$	-315.13***	-0.08***	-370.78***	-0.09***
	(16.23)	(0.00)	(14.27)	(0.00)
$\Delta = 36$	-404.73***	-0.11***	-446.42***	-0.11***
	(18.68)	(0.00)	(16.43)	(0.00)
$\Delta = 48$	-484.00***	-0.13***	-522.53***	-0.13***
	(20.96)	(0.00)	(18.74)	(0.00)

Note: Staggered difference-in-differences estimates are obtained using Equation 8. The effects reported in this table are calculated using Equation 11. Standard errors are clustered at the individual level. *** indicates significance at the 1% level, ** at the 5% level and * at the 10% level. Source: Authors' calculations using IDI.

5.3 Effects of mTBIs on accident compensation payments

In addition to the effects of mTBIs on labour market outcomes, we also examine to what extent the ACC system can mitigate the earning losses we observe. To this end, we estimate the effects of mTBIs on accident compensation payments using the same methodology. Table III.3 shows the accident compensation payment results, alongside the monthly earnings results from Table III.1 for reference.

As expected, individuals who experienced a mTBI receive accident compensation payments following the shock, the average monthly treatment effect for over the first 12 months after the mTBI being around NZ$100. However, this amounts to only about 40% of earning losses. As the earning losses grow over time, the accident compensation payments actually shrink in relative terms, reaching just 11% of average earning pre-injury losses by the fourth year.

Table III.3: Effects of mTBI on accident compensation payments

	(1) Monthly ACC payments	(2) Monthly earnings	(3) -ACC/Earning losses
Pre-effect	-0.42***	4.63***	0.09
	(0.12)	(0.55)	
<u>*ATT*</u>			
$\Delta = 12$	101.92***	-255.58***	0.40
	(3.79)	(9.11)	
$\Delta = 24$	77.80***	-346.55***	0.22
	(3.62)	(10.72)	
$\Delta = 36$	63.75***	-428.09***	0.15
	(3.64)	(12.33)	
$\Delta = 48$	54.98***	-505.74***	0.11
	(3.84)	(13.97)	
N individuals	35,301	35,301	35,301

Note: Staggered difference-in-differences estimates are obtained using Equation 8. The effects reported in this table are calculated using Equation 11. Effects are calculated for all individuals in the sample. Standard errors are clustered at the individual level. For comparison, earnings effects from Table III.1 are repeated here.
*** indicates significance at the 1% level, ** at the 5% level and * at the 10% level.
Source: Authors' calculations using IDI.

To examine this in more detail, Figure III.4 provides monthly estimates. In the first month following the mTBI, the average amount received is NZ$220, which covers almost entirely the loss in labour earnings (-NZ$230, see Figure III.2b).[24] After that initial payment, the compensation gradually decreases, not exceeding NZ$100 after one year and NZ$50 after two years. Given that, at the same time, labour earnings penalties are strongly and continuously increasing, individuals who experienced a mTBI suffer strong income losses. The average monthly ACC payment over the four

[24]ACC will pay up to 80% of pre-injury earnings as compensation. It is, therefore, unclear why the average ACC payments equate to more than 80% of earnings losses in the first month after the injury. However, it may be due to payment timing. For example, if there is an initial delay in the payment of ACC in the month of the accident that is then made up in the next month, or because ACC payments are made weekly while most employees are paid fortnightly, which may mean payments are more likely to fall in the earlier month than wage/salary earnings.

years is NZ$55, while the average wage loss is almost ten times higher (NZ$505, see Table III.1).

Figure III.4: Monthly effects of mTBI on ACC payments

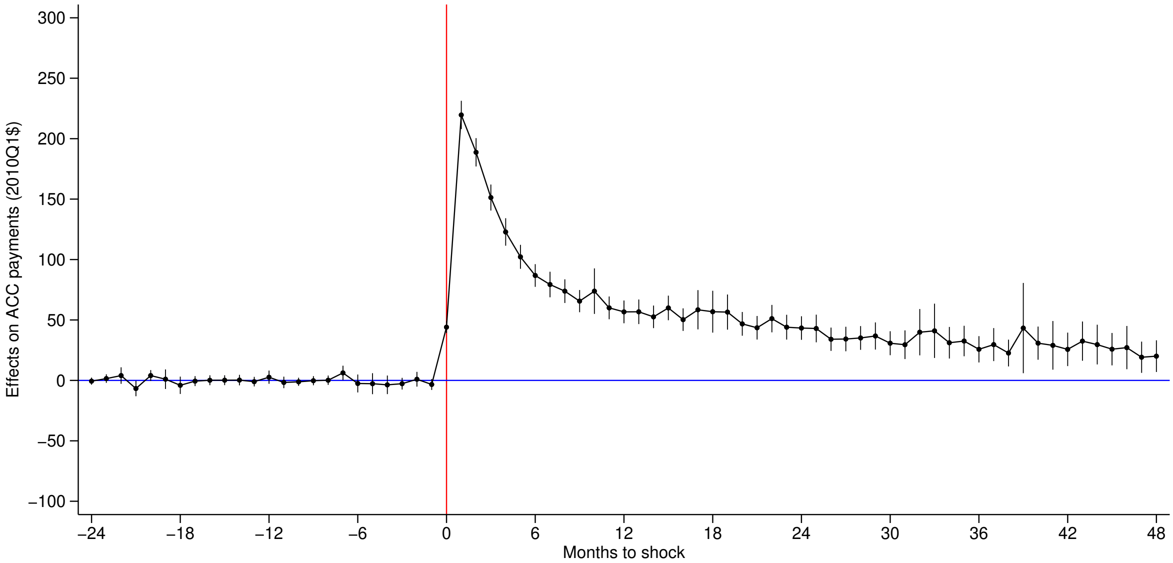

Note: This figure displays the average monthly effects calculated using Equation 10 along with the 95% confidence intervals. Standard errors are clustered at the individual level.
Source: Authors' calculations using IDI.

Heterogeneity analysis

We observe similar patterns in accident compensation payments for all the demographic subgroups (Table III.4). The main differences concern the amount received due to the fact that the ACC system calculates the compensation payments based on previous earnings. Therefore, ACC payments are slightly higher for men, those in high-skilled occupations and older workers. However, in terms of ACC payments as a proportion of earning losses, there are some differences particularly for younger versus older workers. ACC payments account for just over a quarter of earning losses for those under 40 years old in the first 12 months after the mTBI, falling to 7% over a 48 month period. ACC payments as a share of earning losses are almost twice as high for those aged 40 and over, accounting for just over half of earning losses over the first 12 months after the mTBI, falling to 15% over a 48 month period. Although we cannot say definitively why this is the case, it could be because ACC payments are based on pre-injury earnings, and for younger workers, the earnings of the comparison group continue to go up as their careers progress and they enter their prime earning years, while ACC payments are only adjusted to increase in line with the Labour Cost Index. In contrast, for those aged 40 and over, earnings of the comparison group are rising at a slower pace due to these individuals already being in their prime earning years and, therefore, experiencing slower earnings growth. However, there may also be other explanations. For example, ACC may more strongly encourage younger mTBI sufferers to return to work given their potential lifetime earnings compensation liability is higher, or because they are perceived to be more able to recover from their injury and return to work.

There are also some differences between those treated in a hospital versus those treated in a non-hospital setting such as a primary healthcare practice. Those treated in a hospital have higher monthly ACC payments in both absolute terms, and as a share of lost earnings. The difference in absolute terms could be related to injury severity if those treated in hospitals have a more severe TBI on average. For example, someone with a more severe injury may not be able to work at all and would, therefore, receive the maximum amount of income compensation. Someone with with a less severe injury may be able to work but only at reduced hours and, therefore, may receive only partial income compensation. However, this would not explain the difference in terms of the share of lost earnings. A potential reason for the difference in terms of share of lost earnings could be that those treated in a hospital were better able to access ACC services, including income compensation. Perhaps, for example, because they were more likely to have been treated by medical practitioners who were more specialised in TBI and gave additional guidance.

Table III.4: Heterogeneous effects of mTBI on ACC payments

Gender	**(1) Women** ($N = 17,386$)		**(2) Men** ($N = 17,915$)	
	(a) Monthly ACC payments	(b) ACC/ Earning losses	(a) Monthly ACC payments	(b) ACC/ Earning losses
Pre-effect	-0.21	0.06	-0.61***	0.11
	(0.13)		(0.19)	
ATT				
$\Delta = 12$	78.75***	0.41	123.81***	0.39
	(4.56)		(5.97)	
$\Delta = 24$	57.98***	0.21	96.43***	0.24
	(4.34)		(5.72)	
$\Delta = 36$	46.19***	0.13	80.09***	0.16
	(4.39)		(5.72)	
$\Delta = 48$	40.13***	0.09	68.78***	0.12
	(4.81)		(5.90)	
Age	**(3) Under 40** ($N = 14,950$)		**(4) 40+** ($N = 20,351$)	
	(a) Monthly ACC payments	(b) ACC/ Earning losses	(a) Monthly ACC payments	(b) ACC/ Earning losses
Pre-effect	-0.36**	0.04	-0.46***	0.42
	(0.16)		(0.16)	
ATT				
$\Delta = 12$	80.68***	0.27	116.06***	0.51
	(4.89)		(5.41)	
$\Delta = 24$	59.49***	0.14	89.91***	0.31
	(4.63)		(5.19)	
$\Delta = 36$	49.91***	0.09	72.67***	0.21
	(4.70)		(5.19)	
$\Delta = 48$	44.24***	0.07	61.76***	0.15
	(5.06)		(5.43)	
Skill level	**(5) Low- to mid-skilled** ($N = 23,505$)		**(6) High-skilled** ($N = 11,796$)	
	(a) Monthly ACC payments	(b) ACC/ Earning losses	(a) Monthly ACC payments	(b) ACC/ Earning losses
Pre-effect	-0.43***	0.11	-0.38**	0.06
	(0.15)		(0.18)	
ATT				
$\Delta = 12$	89.63***	0.41	127.16***	0.39
	(4.39)		(7.26)	
$\Delta = 24$	69.58***	0.25	94.76***	0.19
	(4.21)		(6.93)	
$\Delta = 36$	56.92***	0.17	77.86***	0.12
	(4.22)		(7.01)	
$\Delta = 48$	49.42***	0.13	66.41***	0.09
	(4.40)		(7.55)	
Hospital	(7) Hospital-treated ($N = 7,890$)		(8) Non-hospital treated ($N = 27,411$)	
	(a) Monthly ACC payments	(b) ACC/ Earning losses	(a) Monthly ACC payments	(b) ACC/ Earning losses
Pre-effect	-0.84***	0.13	-0.31***	0.08
	(0.26)		(0.13)	
ATT				
$\Delta = 12$	146.71***	0.50	89.09***	0.36
	(8.88)		(4.16)	
$\Delta = 24$	110.54***	0.30	68.33***	0.20
	(8.10)		(4.05)	
$\Delta = 36$	88.92***	0.20	56.43***	0.13
	(7.84)		(4.12)	
$\Delta = 48$	76.23***	(0.15)	48.68***	(0.10)
	(7.84)		(4.41)	

Note: Staggered difference-in-differences estimates are obtained using Equation 8. The effects reported in this table are calculated using Equation 11. Standard errors are clustered at the individual level. *** indicates significance at the 1% level, ** at the 5% level and * at the 10% level. Source: Authors' calculations using IDI.

6 Conclusion

This study investigates the effects of mild traumatic brain injuries (mTBIs) on future labour market outcomes. It is important to understand the effects of mTBIs as they are a common injury with a high incidence rate. Indeed, the health literature on mTBIs is growing, with a recognition that our understanding of brain injuries and their effects is lagging behind our understanding of other physical injuries. Despite the growing body of economic research on health shocks, this is the first study to examine TBIs specifically. Past economics literature highlights that minor / temporary injuries have only a short-term effect on labour market outcomes. However, the health literature suggests mTBI could have more lasting effects despite their seemingly minor nature due to the possibility of ongoing cognitive impairment. This paper therefore fills a clear gap in the literature by examining the causal effect of suffering a mTBI on subsequent employment and earnings outcomes.

We use population-wide administrative data on all medically-diagnosed mTBIs linked to employment and earnings data from tax records. To account for possible endogeneity, we construct comparison groups of those who suffer from a mTBI but at a future date and apply a doubly-robust staggered difference-in-differences estimator. We examine the effect of suffering just one mTBI by excluding those who had previously been diagnosed with a TBI or who suffered a subsequent TBI (of any severity).

We find that individuals who experience a mTBI suffer adverse effects on employment and large earning losses. A large part of the effect is at the extensive margin (i.e. an employment effect), although the intensive margin is also important in explaining the earning losses. Indeed, for individuals who manage to stay employed after a mTBI, we observe an initial drop in earnings, followed by a recovery, before earnings gradually decrease again in the medium term. This is in line with the medical literature on mTBIs, which finds that many individuals return to work after an initial recovery period and may manage adequately for a time. However, the expectation that they will continue to gradually improve until they reach their pre-injury level of productivity may not eventuate, resulting in them performing at a lower level for a prolonged period of time. Their employers are also likely to become less understanding of these lower productivity levels as time passes. These persistent symptoms then leads to them leaving their roles and/or failing to progress in their careers as they would have otherwise done so, resulting in poor medium-term labour market outcomes.

There is some difference in the magnitude of the adverse effects across groups. The largest difference is between younger and older workers, with younger workers experiencing greater adverse effects on employment and earnings. Additional smaller differences are evident by gender and occupational level. The negative labour market effects are larger for men than women, and for those in high-skilled occupations relative to those in low-to-middle skilled occupations.

We also investigate the accident compensation payments following a mTBI. We find that earning losses are largely offset by accident compensation payments immediately following the shock. However, while earnings continue to fall over the four years after the mTBI, accident compensation payments decrease over time. This raises questions about whether the current income replacement scheme recognises and provides adequate assistance for the potential long-term consequences of mTBIs.

Overall, our results show that, despite being classified as minor injuries, suffering a mTBIs can have important and persistent detrimental effects on individuals' labour market outcomes. Our findings highlight the need for timely diagnosis and treatment to mitigate the effects of mTBIs and reduce the burden on the individual, in terms of not only health costs, but also economic and social costs experienced in the labour market. Early intervention would reduce the likelihood of these negative effects, particularly longer-term effects, from occurring. Given the richness of linked administrative data used in this study, future work could explore the broader effects of mTBI. This includes the potential impact on educational outcomes of young people; other health outcomes (such as mental health); and relationship outcomes, for example, through an examination of family formation/dissolution patterns.

References

ACC. (2017). *Traumatic brain injury strategy and action plan (2017-2021)*. Accident Compensation Corporation. Wellington, New Zealand. Retrieved August 22, 2023, from https://www.acc.co.nz/assets/provider/1bf15d391c/tbi-strategy-action-plan.pdf

Accident Compensation Corporation. (2022). Concussion / TBI dataset. Retrieved August 6, 2023, from https://catalogue.data.govt.nz/dataset/acc-concussion-tbi-data/resource/bd83401b-459b-4c51-9f15-d6b3b8c0c6b4

Ameri, M., Schur, L., Adya, M., Bentley, F. S., McKay, P., & Kruse, D. (2018). The disability employment puzzle: A field experiment on employer hiring behavior. *ILR Review*, *71*(2), 329–364. https://doi.org/10.1177/0019793917717474

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Case, A., Lubotsky, D., & Paxson, C. (2002). Economic status and health in childhood: The origins of the gradient. *American Economic Review*, *92*(5), 1308–1334. https://doi.org/10.1257/000282802762024520

Charles, K. K. (2003). The longitudinal structure of earnings losses among work-limited disabled workers. *The Journal of Human Resources*, *38*(3), 618–646. https://doi.org/10.2307/1558770

Contoyannis, P., & Rice, N. (2001). The impact of health on wages: Evidence from the British Household Panel Survey. *Empirical Economics*, *26*(4), 599–622. https://doi.org/10.1007/s001810000073

Crichton, S., Stillman, S., & Hyslop, D. (2011). Returning to work from injury: Longitudinal evidence on employment and earnings. *ILR Review*, *64*(4), 765–785. https://doi.org/10.1177/001979391106400407

Dano, A. M. (2005). Road injuries and long-run effects on income and employment. *Health Economics*, *14*(9), 955–970. https://doi.org/10.1002/hec.1045

Dean, P., & Sterr, A. (2013). Long-term effects of mild traumatic brain injury on cognitive performance. *Frontiers in Human Neuroscience*, *7*.

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–2996.

Drydakis, N. (2010). Labour discrimination as a symptom of HIV: Experimental evaluation: The Greek case. *Journal of Industrial Relations*, *52*(2), 201–217. https://doi.org/10.1177/0022185609359445

Eliason, M., & Storrie, D. (2009). Does job loss shorten life? *Journal of Human Resources*, *44*(2), 277–302. https://doi.org/10.3368/jhr.44.2.277

Fadlon, I., & Nielsen, T. H. (2019). Family health behaviors. *American Economic Review*, *109*(9), 3162–3191. https://doi.org/10.1257/aer.20171993

Fallesen, P., & Campos, B. (2020). Effect of concussion on salary and employment: A population-based event time study using a quasi-experimental design. *BMJ Open*, *10*(10), e038161. https://doi.org/10.1136/bmjopen-2020-038161

Feigin, V. L., Theadom, A., Barker-Collo, S., Starkey, N. J., McPherson, K., Kahan, M., Dowell, A., Brown, P., Parag, V., Kydd, R., Jones, K., Jones, A., & Ameratunga, S. (2013). Incidence of traumatic brain injury in New Zealand: A population-based study. *The Lancet Neurology*, *12*(1), 53–64. https://doi.org/10.1016/S1474-4422(12)70262-4

García Gómez, P., & López Nicolás, A. (2006). Health shocks, employment and income in the Spanish labour market. *Health Economics*, *15*(9), 997–1009. https://doi.org/10.1002/hec.1151

García-Gómez, P. (2011). Institutions, health shocks and labour market outcomes across Europe. *Journal of Health Economics*, *30*(1), 200–213. https://doi.org/10.1016/j.jhealeco.2010.11.003

García-Gómez, P., Kippersluis, H. v., O'Donnell, O., & Doorslaer, E. v. (2013). Long-term and spillover effects of health shocks on employment and income. *Journal of Human Resources*, *48*(4), 873–909. https://doi.org/10.1353/jhr.2013.0031

Gils, A. v., Stone, J., Welch, K., Davidson, L. R., Kerslake, D., Caesar, D., McWhirter, L., & Carson, A. (2020). Management of mild traumatic brain injury. *Practical Neurology*, *20*(3), 213–221. https://doi.org/10.1136/practneurol-2018-002087

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, *225*(2), 254–277.

Graff, H. J., Siersma, V., Møller, A., Kragstrup, J., Andersen, L. L., Egerod, I., & Rytter, H. M. (2019). Labour market attachment after mild traumatic brain injury: Nationwide cohort study with 5-year register follow-up in denmark. *BMJ Open*, *9*(4), e026104. https://doi.org/10.1136/bmjopen-2018-026104

Halla, M., & Zweimüller, M. (2013). The effect of health on earnings: Quasi-experimental evidence from commuting accidents. *Labour Economics*, *24*, 23–38. https://doi.org/10.1016/j.labeco.2013.04.006

Heckman, J. J., Ichimura, H., & Todd, P. E. (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The Review of Economic Studies*, *64*(4), 605–654. https://doi.org/10.2307/2971733

Heinesen, E., Imai, S., & Maruyama, S. (2018). Employment, job skills and occupational mobility of cancer survivors. *Journal of Health Economics*, *58*, 151–175.

Heinesen, E., & Kolodziejczyk, C. (2013). Effects of breast and colorectal cancer on labour market outcomes—average effects and educational gradients. *Journal of Health Economics*, *32*(6), 1028–1042.

Hipes, C., Lucas, J., Phelan, J. C., & White, R. C. (2016). The stigma of mental illness in the labor market. *Social Science Research*, *56*, 16–25. https://doi.org/10.1016/j.ssresearch.2015.12.001

Hyder, A. A., Wunderlich, C. A., Puvanachandra, P., Gururaj, G., & Kobusingye, O. C. (2007). The impact of traumatic brain injuries: A global perspective. *NeuroRehabilitation*, *22*(5), 341–353.

Langlois, J. A., Rutland-Brown, W., & Wald, M. M. (2006). The epidemiology and impact of traumatic brain injury: A brief overview. *The Journal of Head Trauma Rehabilitation*, *21*(5), 375–378. https://doi.org/10.1097/00001199-200609000-00001

Lechner, M., & Vazquez-Alvarez, R. (2011). The effect of disability on labour market outcomes in germany. *Applied Economics*, *43*(4), 389–412. https://doi.org/10.1080/00036840802599974

Lenhart, O. (2019). The effects of health shocks on labor market outcomes: Evidence from UK panel data. *The European Journal of Health Economics*, *20*(1), 83–98. https://doi.org/10.1007/s10198-018-0985-z

Lindeboom, M., Llena-Nozal, A., & van der Klaauw, B. (2016). Health shocks, disability and work. *Labour Economics*, *43*, 186–200. https://doi.org/10.1016/j.labeco.2016.06.010

Miller, G. F., DePadilla, L., & Xu, L. (2021). Costs of nonfatal traumatic brain injury in the United States, 2016. *Medical Care*, *59*(5), 451. https://doi.org/10.1097/MLR.0000000000001511

Pelkowski, J. M., & Berger, M. C. (2004). The impact of health on employment, wages, and hours worked over the life cycle. *The Quarterly Review of Economics and Finance*, *44*(1), 102–121. https://doi.org/10.1016/j.qref.2003.08.002

Ribbers, G. M. (2007). Traumatic brain injury rehabilitation in the netherlands: Dilemmas and challenges. *The Journal of Head Trauma Rehabilitation*, *22*(4), 234. https://doi.org/10.1097/01.HTR.0000281839.07968.32

Riphahn, R. T. (1999). Income and employment effects of health shocks a test case for the German welfare state. *Journal of Population Economics*, *12*(3), 363–389. https://doi.org/10.1007/s001480050104

Sullivan, D., & von Wachter, T. (2009). Job displacement and mortality: An analysis using administrative data. *The Quarterly Journal of Economics*, *124*(3), 1265–1306. https://doi.org/10.1162/qjec.2009.124.3.1265

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

Tanaka, A. (2021). The effects of sudden health reductions on labor market outcomes: Evidence from incidence of stroke. *Health Economics*, *30*(6), 1480–1497.

Theadom, A., Barker-Collo, S., Jones, K., Kahan, M., Te Ao, B., McPherson, K., Starkey, N., Feigin, V., Feigin, V., Theadom, A., Barker-Collo, S., McPherson, K., Kydd, R., Barber, P. A., Parag, V., Brown, P., Starkey, N., Dowell, A., Kahan, M., … Te Ao, B. (2017). Work limitations 4 years after mild traumatic brain injury: A cohort study. *Archives of Physical Medicine and Rehabilitation*, *98*(8), 1560–1566. https://doi.org/10.1016/j.apmr.2017.01.010

Theadom, A., Meehan, L., McCallum, S., & Pacheco, G. (2023). Mild traumatic brain injury increases engagement in criminal behaviour 10 years later: A case-control study. *Frontiers in Psychiatry*, *14*, 1154707. https://doi.org/10.3389/fpsyt.2023.1154707

van der Horn, H. J., Out, M. L., de Koning, M. E., Mayer, A. R., Spikman, J. M., Sommer, I. E., & van der Naalt, J. (2020). An integrated perspective linking physiological and psychological consequences of mild traumatic brain injury. *Journal of Neurology*, *267*(9), 2497–2506. https://doi.org/10.1007/s00415-019-09335-8

Wehman, P. H., Targett, P. S., & Avellone, L. E. (2017). Educational and vocational issues in traumatic brain injury. *Physical Medicine and Rehabilitation Clinics of North America*, *28*(2), 351–362. https://doi.org/10.1016/j.pmr.2016.12.010

Williams, W. H., McAuliffe, K. A., Cohen, M. H., Parsonage, M., Ramsbotham, J., & General The Lord David. (2015). Traumatic brain injury and juvenile offending: Complex causal links offer multiple targets to reduce crime. *The Journal of Head Trauma Rehabilitation*, *30*(2), 69–74. https://doi.org/10.1097/HTR.0000000000000134

Wooldridge, J. (2005). Violating ignorability of treatment by controlling for too many factors. *Econometric Theory*, *21*(5), 1026–1028. https://doi.org/10.1017/S0266466605050516

Zucchelli, E., Jones, A. M., Rice, N., & Harris, A. (2010). The effects of health shocks on labour market exits: Evidence from the HILDA survey. *Australian Journal of Labour Economics*, *13*(2), 191–218. https://doi.org/10.3316/informit.496905059127698

Conclusion

This thesis applies quasi-experimental methods to linked administrative data to make causal inferences about two policy issues and a health shock. All three papers use dynamic difference-in-differences methodology applied to Stats NZ's Integrated Data Infrastructure (IDI).

Paper 1 examines the effect of citizenship on parents' integration outcomes and children's health outcomes. Does citizenship increase an immigrant's willingness to integrate into the host country's society, or is it the final step in the integration process without further consequences? The answer to this question has potentially important implications for policy. If citizenship is an important catalyst that encourages migrants to integrate and invest in their own, and their children's, futures, then naturalisation should be made fairly accessible. However, if gaining citizenship does not have an independent effect on integration, or indeed, even reduces incentives to integrate, then a higher bar to receiving a host-country passport may be appropriate, with citizenship being viewed more as a reward for successfully completing the integration process.

However, it is difficult to make causal inferences about the effects of citizenship due to possible selection effects. Those who are more motivated to integrate are also more likely to naturalise. Thus, this paper exploits a natural experiment of NZ's removal of birthright citizenship in 2006 to examine the effects of citizenship on a range of parental and children's outcomes. As such, it adds to the evidence in this area which has, to date, largely be limited to Germany's introduction of birthright citizenship in 2000.

As an anglophone country with high immigration rates and a selective migration policy, the case of NZ is potentially more relevant to Canada and the US, the only two western countries which still retain unrestricted birthright citizenship, than the case of Germany. In addition, NZ's form of birthright citizenship was very similar to existing policy in Canada and the US in that anyone born in the country was entitled to citizenship regardless of their parents' legal status, with the only exception being children born to foreign diplomats. The form of birthright citizenship that Germany introduced is much more restrictive as it still requires at least one parent to have lived legally in Germany for at least eight years. Moreover, the pathways to citizenship in NZ, Canada and the US are very similar in terms of the process and length of time it takes to become a citizen, whereas the citizenship pathway in Germany tends to be longer and more costly, particularly as it generally requires naturalising individuals to relinquish their previous citizenship.

This paper finds that the removal of birthright citizenship did not impact family outmigration, parents' fertility or labour market outcomes, nor children's health outcomes. These results contrast with those for Germany, with several studies finding that the introduction of birthright citizenship improved parents' and children's outcomes across a range of dimensions. This difference in findings highlights that the benefits of citizenship are context specific. It is postulated that citizenship does not have a discernible impact in NZ because the benefits of citizenship are lower than in Germany. In NZ, virtually all migrants have a pathway to residency and citizenship - that is, there are very few true temporary or guest worker migrants. In addition, other residency options provide almost the same benefits as citizenship in NZ. Moreover, reflecting its selective migration policy, NZ's migrants have, on average, higher education levels than the native-born population, while in Germany, migrants have lower average education levels. It is likely that the benefits of citizenship are lower for more skilled migrants, which may also explain the difference in the findings for the two countries.

These results suggest that removing birthright citizenship in Canada and the US may not have much impact on the integration outcomes of immigrants. The NZ case is particularly relevant to Canada, which also has a relatively high-skilled migrant population and a low share of undocumented migrants. For the US, there are some key differences which may mean that the negative effects of removing birthright citizenship could be greater. First, the foreign-born population in the US are relatively low skilled - for example, they have lower average education levels than the native-born population. It is likely that citizenship has more benefits for lower-skilled migrants since it will open up more opportunities than in the case of higher-skilled migrants. Second, the share of undocumented migrants in the US is much higher than in NZ and Canada. Citizenship in general, and the automatic granting of citizenship to their children in particular, is likely to have greater benefits for undocumented migrants than documented ones. However, the NZ results do suggest that the costs of removing birthright citizenship may not be as great in the US as the German case suggests. While citizenship appears to be a catalyst to the integration process in Germany, it may merely be the final step in the integration process without further consequences for immigrants in countries such as NZ, Canada and the US.

More generally, NZ is an interesting case study of an anglophone country with high migration rates where migration policy changes in a range of areas have been implemented over the years. Along with the availability of high-quality data on visa applications and border movements linked to other datasets, such as health and tax information, this opens numerous possibilities in terms of analysing the impact of different migration policy settings on the outcomes of immigrants and their children.

Paper 2 examines the effect of workforce vaccine mandates on vaccination uptake and healthcare workers' (HCWs') labour market outcomes. We use linked population-wide administrative data from New Zealand, which includes a comprehensive national vaccination register linked to tax records to identify employment outcomes.

We employ a difference-in-differences approach to isolate the effects of workforce vaccination mandates from the effects of the NZ government's population-wide initiatives to boost vaccination rates, particularly vaccine passes to access non-essential businesses/services. Unfortunately, no comparison group could be found where the parallel trends assumption held due to HCWs' early access to vaccinations. However, vaccination rates were already very high among mandated workers when the mandates were announced, leaving little room for vaccination rates to increase. Moreover, unlike international studies examining vaccine passes, there is no discontinuous jump in vaccination rates following the mandate announcement.

We additionally apply a dynamic triple difference approach (DDD) to examine healthcare workers' labour market outcomes, comparing unvaccinated HCWs with vaccinated HCWs and vaccinated and unvaccinated workers in industries that were not covered by workforce mandates. We find that the mandates negatively impacted on unvaccinated workers' overall employment rates, their rates of employment within the health industry and their earnings. While some groups, such as higher-income workers, saw some recovery in their labour market outcomes over time, the negative effects persisted for most groups of workers throughout the 13-month post-announcement period.

Overall, the results suggest that in the context of already-high vaccination rates, workforce vaccine mandates may not have provided much benefit in terms of increasing vaccination rates among mandated workers. Moreover, they came at a cost in terms of HCWs' labour market outcomes, which may have had wider negative consequences in terms of the supply of healthcare workers in an area where skills shortages were already an issue.

Given ongoing HCW worker shortages in NZ (and globally), there is potential to further explore the health workforce within the IDI. Although the absence of administrative data on the individual's role/occupation within the health sector is a potential limitation, this information could be proxied for in some contexts via a combination of census and educational qualifications data. This analysis could include consideration of recruitment into the industry, in terms of the transition from the education system to the health workforce, as well as via immigration given the high share of migrant workers in this sector. There is also potential to examine the labour market trajectories of those within the health workforce to understand retention and the destinations (e.g. future employment, outmigration etc.) of those leaving the sector.

Paper 3 investigates the effects of mild traumatic brain injuries (mTBIs) on future labour market outcomes. It is important to understand the effects of mTBIs as they are a common injury with a high incidence rate. Indeed, the health literature on mTBIs is growing, with a recognition that our understanding of brain injuries and their effects is lagging behind our understanding of other physical injuries. Despite the growing body of economic research on health shocks, this is the first study to examine TBIs specifically. Past economics literature highlights that minor / temporary injuries have only a short-term effect on labour market outcomes. However, the health literature suggests mTBI could have more lasting effects despite their seemingly minor nature due to the possibility of ongoing cognitive impairment. This paper therefore fills a clear gap in the literature by examining the causal effect of suffering a mTBI on subsequent employment and earnings outcomes.

We use population-wide administrative data on all medically-diagnosed mTBIs linked to employment and earnings data from tax records. To account for possible endogeneity, we construct comparison groups of those who suffer from a mTBI but at a future date and apply a doubly-robust staggered difference-in-differences estimator. We examine the effect of suffering just one mTBI by excluding those who had previously been diagnosed with a TBI or who suffered a subsequent TBI (of any severity).

We find that individuals who experience a mTBI suffer adverse effects on employment and large earning losses. A large part of the effect is at the extensive margin (i.e. an employment effect), although the intensive margin is also important in explaining the earning losses. Indeed, for individuals who manage to stay employed after a mTBI, we observe an initial drop in earnings, followed by a recovery, before earnings gradually decrease again in the medium term. This is in line with the medical literature on mTBIs, which finds that many individuals return to work after an initial recovery period and may manage adequately for a time. However, the expectation that they will continue to gradually improve until they reach their pre-injury level of productivity may not eventuate, resulting in them performing at a lower level for a prolonged period of time. Their employers are also likely to become less understanding of these lower productivity levels as time passes. These persistent symptoms then leads to them leaving their roles and/or failing to progress in their careers as they would have otherwise done so, resulting in poor medium-term labour market outcomes.

There is some difference in the magnitude of the adverse effects across groups. The largest difference is between younger and older workers, with younger workers experiencing greater adverse effects on employment and earnings. Additional smaller differences are evident by gender and occupational level. The negative labour market effects are larger for men than women, and for those in high-skilled occupations relative to those in low-to-middle skilled occupations.

We also investigate the accident compensation payments following a mTBI. We find that earning losses are largely offset by accident compensation payments immediately following the shock. However, while earnings continue to fall over the four years after the mTBI, accident compensation payments decrease over time. This raises questions about whether the current income replacement scheme recognises and provides adequate assistance for the potential long-term consequences of mTBIs.

Overall, our results show that, despite being classified as minor injuries, suffering a mTBIs can have important and persistent detrimental effects on individuals' labour market outcomes. Our findings highlight the need for timely diagnosis and treatment to mitigate the effects of mTBIs and reduce the burden on the individual, in terms of not only health costs, but also economic and social costs experienced in the labour market. Early intervention would reduce the likelihood of these negative effects, particularly longer-term effects, from occurring. Given the richness of linked administrative data used in this study, future work could explore the broader effects of mTBI. This includes the potential impact on educational outcomes of young people; other health outcomes (such as mental health); and relationship outcomes, for example, through an examination of family formation/dissolution patterns.

There is growing awareness of the potential impacts of mTBIs and the IDI offers an opportunity to further explore the effect of mTBIs on a range of other outcomes. This includes further work into the impact on criminal activity, mental health outcomes and relationship formation and dissolution patterns. In terms of the impact of health shocks more generally, the consequences of these shocks for affected individuals' family members, such as their children, is also a potential area for further research.

In summary, all three papers demonstrate the utility of applying quasi-experimental methods to linked administrative data. While randomised control trials are the "gold standard" in terms of establishing causal effects, these are not always possible for practical and/or ethical reasons. However, the use of quasi-experimental methods to administrative data can provide a useful alternative in some situations and yield valuable insights.