Empirical Essays in Intergenerational Mobility and Early Childhood Human Capital

Nathan Deutscher

A thesis submitted for the degree of
Doctor of Philosophy of
The Australian National University
December 2018

Declaration

This thesis does not contain material that has been accepted for the award of any other degree or diploma in any University or other tertiary institution. The publication status and extent of my contribution to the individual chapters is outlined below:

Chapter 1: N. Deutscher and B. Mazumder. Intergenerational mobility in Australia: national and regional estimates using administrative data. (working paper only)

I took sole responsibility for the data curation and analysis, and shared the write up of the paper with my co-author, Bhashkar Mazumder.

Chapter 2: N. Deutscher. Place, jobs, peers and the importance of the teenage years: exposure effects and intergenerational mobility.
(1^{st} round revise and resubmit at *AEJ: Applied Economics*)

Chapter 3: N. Deutscher. What drives second generation success? The role of education, culture and social context.
(submitted to *Journal of Population Economics.*)

Chapter 4: N. Deutscher and R. Breunig. Baby bonuses: natural experiments in cash transfers, birth timing and child outcomes
(published in the *Economic Record*, March 2018)

I developed the research questions, design and took sole responsibility for data curation and analysis. I wrote the first draft of the paper, and my co-author, Robert Breunig, then redrafted it for submission.

Other than where noted above, this thesis does not contain any material previously published or written by another person, except where due reference is made in the text.

Nathan Deutscher

December 2018

Acknowledgements

As with any endeavour, this thesis reflects not only my abilities and efforts, but those of many others, and a good dose of luck.

I am greatly indebted to my supervisory panel — Professor Robert Breunig, Dr Bhashkar Mazumder and Associate Professor Tue Gorgens. Bob is the approachable, enthusiastic and clear-eyed supervisor every PhD student needs — I sincerely hope that this thesis does not contain our last joint production! Bhash has been remarkably generous and insightful in his contributions, and has opened doors I would have thought closed. And Tue's pointers on the econometrics have never missed the mark.

The Sir Roland Wilson Foundation and the Department of the Treasury provided generous financial support, but also an environment that allowed me to learn quickly, and take risks. Early visits to international workshops and experts exposed me to the frontier of knowledge, while a career to return to allowed me to take risks in my research. At key points my Treasury mentor Jenny Wilkinson, and fellow scholars, have provided valuable support.

The Australian Taxation Office has been central to the success of this work. Their early in-principle support was central to my decision to embark on the PhD. And since those discussions in late 2014, there has been much work in building a world-class intergenerational dataset. Their expertise, advice, support and professionalism has been personally invaluable. But it has also placed this research agenda on a solid footing — while this thesis is the first use of tax data in studying intergenerational mobility in Australia, I am confident it will not be the last.

This thesis also rests on the support of numerous other data custodians, with none of the papers drawing on off-the-shelf data. Both the level and trend that I've seen in the availability, ease of access and scope of data available to researchers give me confidence in the future of empirical research in Australia.

In the chapters that follow, I show that even in the land of the ‘fair go’, family matters for individual success. Long since leaving home, I still have my parents to thank for kindling my early love of learning, work ethic and grit. The relatively isolated experience of a PhD has often left my partner Hilary as my main confidant and counsel, through numerous ups and downs. I could not ask for a better person to turn to; her questions, perspective and prodding have kept me on track. Finally, thanks to my two boys, Marius and Pippin — your delight in life is infectious.

Abstract

This thesis contributes to the vast literature on intergenerational mobility and human capital formation. It consists of four papers presenting new facts on intergenerational mobility in Australia, and fresh insights into its measurement and underlying causal mechanisms. I use a variety of empirical techniques, and datasets ranging from millions to mere hundreds of observations.

In Chapter 2, I introduce a new Australian intergenerational dataset. I use this data to present a brief overview of intergenerational income mobility in Australia, at both a national and regional scale. Special attention is paid to a variety of measurement issues known to bias measures of mobility. Australia has a high degree of intergenerational mobility but still displays significant regional variations.

In Chapter 3, I examine the causes of regional variation in intergenerational mobility. Is it driven by sorting across regions or does it reflect a causal effect of place? I apply the methodology introduced in Chetty and Hendren (2018a), exploiting variation in the age at which children move to identify a causal effect of exposure to place. I find that place matters, and it matters most in the teenage years. I identify two mechanisms that may explain a causal effect of place in the teenage years — entry to local labour markets and peer effects.

In Chapter 4, I explore the intergenerational mobility of migrants to Australia. The second generation in some migrant communities is remarkably successful, even conditional on outcomes for the first generation. I present a new decomposition of intergenerational income mobility, and use it to highlight the pivotal role of education in second generation success. I provide evidence of a role for culture and the context of migration in shaping the educational achievements and aspirations of the second generation.

In Chapter 5, I examine the causal influence of family income itself on child outcomes. I use the sharp date-of-birth eligibility cut-off for the Australian Baby Bonus to examine the effects of a modest boost to family income at birth. While the income boost is enough to induce parents to shift births, resulting in changes in infant health outcomes, it has no appreciable effect on Year 3 test scores.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Equality of opportunity is central to many conceptions of a just society, and a perennial touchstone in public policy debates. Accordingly, there is a vast related literature in economics, perhaps most notably on intergenerational mobility — the extent to which economic outcomes persist from one generation to the next (Solon (1999); Black and Devereux (2011)) — and early human capital formation (Currie and Almond (2011)). This thesis contributes to this literature. I provide the most complete picture of intergenerational mobility in Australia to date and insights new to the international literature on a variety of causal mechanisms — including causal place and peer effects and the potential influence of less tangible factors such as culture and social context.

In Chapter 2, co-authored with Bhashkar Mazumder, I provide new estimates of intergenerational mobility in Australia. These are the first estimates of standard measures of mobility using Australian administrative data. Until recently, estimates of intergenerational income mobility in Australia have been indirect, imputing parent income based on occupation (Leigh (2007); Mendolia and Siminski (2016)). More recent direct estimates have made it clearer that Australia is more mobile than the United States, but have remained relatively imprecise (Murray et al. (forthcoming)). This thesis provides a much more precise and complete overview of intergenerational mobility in Australia. Australia emerges as a relatively mobile country by international standards, with much less persistence in outcomes than the United States and even much of Europe. While less pronounced than in the United States, there is nonetheless notable geographic variation in intergenerational mobility.

In Chapter 3, I consider the extent to which this geographic variation in intergenerational mobility in Australia reflects an underlying causal effect of place. I apply the methodology introduced in Chetty and Hendren (2018a), exploiting variation in the age at which children move to identify the effect of an additional year exposed to place at varying points in the lifecycle. Despite Australia being a much more mobile country, with less geographic variation in intergenerational mobility, I find evidence that the geographic variation that does exist is mostly causal — just as it is in the United States. By moving to a new place at birth, a child could expect to pick up about 70% of the gap in outcomes between their origin and destination.

In a new contribution, I show the causal effect of exposure to place varies over childhood — teenagers are more sensitive to place than young children.[1] This finding need not contradict, but rather complements, the vast literature on the importance of early childhood (e.g. Case et al. (2005); Cunha and Heckman (2007); Currie and Almond (2011)). It simply suggests that the mechanisms driving a causal effect of the places considered here differ from those that matter in early childhood. I then shine a light on two such mechanisms that may explain the causal effect of place and the sensitivity of the teenage years — local labour market effects and peer effects.

A relatively prosaic explanation for place effects is that they may be driven by local labour market entry. In Page and Solon (2003a), for example, most of the correlation in the adult earnings of neighbouring boys in the United States is explained by the large earnings premium in urban areas and the high correlation between childhood and adult urbanicity. Indeed, the results from Chapter 2 suggest local labour market conditions may drive some of the observed variation in intergenerational mobility in Australia. Using a variety of approaches, I estimate that around 15-55% of the causal effect of exposure to place is explained by the local labour market entry. In short, an additional year exposed to a place in childhood lifts the probability an individual ends up working there, and earning any

[1]This same pattern has since been observed in more recent work in the United States (Chetty, Friedman, Hendren, Jones and Porter (2018)).

associated earnings premium.

Finally, peer effects are a potential candidate for some of the residual effect of place, given the sensitivity of the teenage years and the large psychological literature on the increasing importance of peers at this time (Brown and Larson (2009)). While there is a large literature on the estimation of peer effects in education (Sacerdote (2011)), studies on the long-run influence of neighbourhood peers are rare. I identify peer effects using cross-cohort variation in the mean parent income ranks of permanent postcode residents. In short, if an individual just happens to be born into a particularly rich or poor cohort for their postcode, do they do any better or worse? I find a 10 percentile rank point increase in the mean parent income rank of an individual's postcode-cohort peers is associated with a 0.2-0.3 percentile point increase in their own income rank at age 24 — around a fifth of the influence of parents at this age. These peer effects are open to a variety of interpretations: an individual's peers' parents may matter directly, as role models or job contacts; or they may matter only insofar as they proxy for the behaviours and aspirations found in the peers themselves. Nonetheless, the results suggest that social interactions within neighbourhoods are a potential explanation for the causal effect of place.

In Chapter 4, I examine differences in intergenerational mobility across a different set of subpopulations — new migrants to Australia. My focus is on explaining why some migrant communities do better or worse in the second generation than might be expected based on first generation outcomes. To this end, I present a new decomposition of the residual from the canonical intergenerational income mobility regression. I show differences in education mobility, rather than differences in the return to education, account for most of the differences between migrant communities in second generation incomes conditioned on those of the first.

What then explains differences in education mobility across migrant communities? I present evidence in favour of a role for both culture and social context. I explore the role of culture using the 'epidemiological approach' (Fernández (2011)), exploiting the fact that migrants may bring elements of their culture with them, but leave the source country institutions and economy behind. To the extent that source country outcomes predict outcomes in the destination, it plausibly reflects a causal influence of culture. I find that second generation migrants from countries that outperform on PISA tend to end up with more education and, even more clearly, outperform on similar tests in Australia. Even more notably, the context of migration matters — migrant communities whose fathers were poorer, conditional on their education, tend to get more education in the second generation. This effect appears to be driven by family aspirations rather than child ability, and points to the potential importance of social context in determining educational aspirations and attainment. These effects hold when controlling for detailed family background and school fixed effects, and when looking within rather than between major source regions.

Finally, in Chapter 5, co-authored with Robert Breunig, I return to considering the population at large and the role of family income, though this time in a causal framework. We exploit the introduction of the Australian Baby Bonus — a cash transfer of AUD3,000 to parents of babies born on or after 1 July 2004. Past research has highlighted the unprecedented birth shifting that accompanied this policy (Gans and Leigh (2009)), leading to differences in birth weight and gestation length. We use a difference-in-differences design, excluding the period affected by birth shifting, to test for lasting effects of additional family income at birth on child outcomes, as measured by test scores in Year 3. While the treatment is small, the full universe of test scores provides sufficient power to detect effect sizes at the upper end of those in the existing literature (as surveyed in Currie and Almond (2011)). Despite the payment's clear value to parents, there is no evidence it improved test scores, and only modest evidence of an effect for those from more disadvantaged backgrounds.

As a whole, these papers paint a picture of intergenerational mobility in Australia that is both encouraging and challenging. Encouraging, in that Australia maintains relatively high levels of intergenerational mobility, both at a national level and across most regions. But challenging too, in that some of the mechanisms suggested in this thesis — local labour markets, peers and role models, culture and context — may be particularly challenging to redress.

Chapter 2

Intergenerational mobility in Australia: national and regional estimates using administrative data

Abstract

We produce the first estimates of intergenerational mobility in Australia using tax data covering over a million individuals born between 1978 and 1982. We find that the intergenerational elasticity in total income is 0.185, and that the rank-rank slope is 0.215. These are among the lowest estimates for advanced economies. We show that there is both substantial upward mobility from the bottom of the income distribution and substantial downward mobility from the top. We also produce the first regional estimates of intergenerational mobility for Australia. While mobility is rapid throughout most of the country there is meaningful dispersion — with the mining boom in particular driving strong upward mobility over the period observed.

This chapter was co-authored with Bhashkar Mazumder. Nathan Deutscher was supported by a Sir Roland Wilson scholarship, funded by the Australian Treasury and the Australian National University. We would like to thank Julia Neville, Thomas Abhayaratna, Bruce Bastian,

2.1 Introduction

In recent decades, as economic forces such as changes in technology and rising globalization has impacted economies throughout the world, policy makers are increasingly concerned about both inequality and intergenerational mobility (OECD (2018)). Understanding the degree to which children have the opportunity to succeed irrespective of their parents' economic circumstances is now a central question facing virtually all advanced economies.

Thus far, however, there is only limited evidence on the degree of intergenerational income mobility for Australia. Most existing studies lack data on income for long periods in both generations and researchers have imputed parental income using methods that have often required strong assumptions.[1] There is also no evidence on geographic differences in intergenerational mobility within the country. National estimates for Australia are interesting in their own right and, when compared with other countries, may also shed light on the value of certain institutional features of the country and how these features may enhance or impede mobility. Regional differences within Australia may also provide insights as to how differences in industry composition, economic growth or other unique aspects of different regions may matter for mobility. For example, Bütikofer et al. (2018) highlight the importance of resource booms in Norway in promoting intergenerational mobility.

A key obstacle to producing definitive mobility estimates for Australia has been the lack of large-scale nationally representative administrative panel data linking parents and children. Population-wide data would enable researchers to produce more precise national estimates as well as estimates for smaller geographic

Matt Power and Julia Rymasz for their invaluable expertise, advice and support accessing the intergenerational data. This research uses data from the Australian Taxation Office. All findings, opinions and conclusions are those of the authors and do not necessarily represent the views of the Australian Government or any of its agencies. The views expressed here do not reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve system or the Australian Treasury. The research plan was approved by the Australian National University Human Ethics Committee, protocol number 2017/832.

[1] Studies to date include Leigh (2007); Mendolia and Siminski (2016); Murray et al. (forthcoming). The last of these provides the first estimates of intergenerational income mobility for Australia where both parent and child incomes are observed directly. The earlier studies imputed parental income based on parental occupation. A recent study using administrative data is Cobb-Clark et al. (2017), though the focus here is on intergenerational welfare dependency rather than income mobility per se.

regions. Administrative data with long enough panels also holds promise for more effectively addressing the variety of measurement challenges that is endemic to the intergenerational mobility literature (e.g. Solon (1992); Mazumder (2005); Haider and Solon (2006)).

We address these challenges by producing the first estimates of intergenerational mobility based on Australian tax data. We find that intergenerational persistence in income is quite low and that intergenerational mobility is consequently quite high. We estimate the intergenerational elasticity (IGE) for Australia to be 0.185 and the rank-rank slope to be 0.215. By way of contrast estimates for the United States put the IGE at 0.5 or higher and the rank-rank slope at 0.4 or higher (Mazumder (2016)). Our estimates for Australia are comparable to Nordic countries such as Denmark and Norway where estimates of these parameters are typically around 0.2 or lower. These findings suggest that Australia is among the most intergenerationally mobile countries in the world.

We also use other rank-based measures of intergenerational mobility that provide measures of directional mobility. We find that the expected rank of individuals whose parents were at the 25^{th} percentile is the 45^{th} percentile. Those who started at the 75^{th} percentile could expect to land at the 56^{th} percentile.[2] This suggests that there is substantial upward and downward mobility. In the United States there are only about 53 cities out of 381 that have higher rates of upward mobility than Australia taken as a whole.

Our qualitative results do not change when adjusting for potential sources of bias, including nonlinearities in intergenerational relationships, missing or incorrect parent-child links, missing income observations and measurement of income over too short a period or too early or late in life. Consistent with recent research abroad and in Australia, rank-based measures appear less sensitive to many of these concerns (Nybom and Stuhler (2017); Mazumder (2016); Murray et al. (forthcoming)). Our most conservative estimates lift our estimated IGE to 0.210 (from 0.185) and our rank-rank correlation to 0.232 (from 0.215).

Finally, we estimate standard measures of intergenerational mobility across Australian regions. While there is significantly less dispersion in mobility across

[2] If there were "perfect" mobility such that everyone's rank was randomly distributed, then the expected rank would be at the 50^{th} percentile.

Australian regions relative to the United States, meaningful differences nonetheless emerge — both within the country and within individual cities such as Sydney and Melbourne. Perhaps most notable is the influence of the resources boom, which lifted the expected rank outcomes of children born in resource-rich regions, but did so relatively consistently across the income spectrum. Over this period the mining regions thus typically have higher expected ranks for children, conditional on parent income, but are no more mobile when looking at rank-rank slopes. This highlights the conceptual differences between different measures of intergenerational mobility.

2.2 Data

We use a new intergenerational dataset drawn from Australian federal income tax returns from 1991 to 2015. The Australian Taxation Office (ATO) has produced the data as an extension of its existing research files, the ATO Longitudinal Income Files.[3] Family links primarily come from linking children to adults living at the same address when the child registers for a Tax File Number (TFN): a unique personal identifier issued by the federal government. The algorithm for linking is also informed by a subset of families in which children are directly claimed as dependents on tax returns.[4]

The universe for our baseline sample includes 1.1 million individuals born in Australia between 1 July 1978 and 30 June 1982 who registered for a TFN and remained resident in Australia through 2015. Of these 90% are linked to parents. This is comparable to the matching rate attained by Chetty et al. (2014) for their core sample of children born from 1980-82. Our baseline sample closely mirrors population benchmarks for family structure, median parental age at birth and family size (Table 3.1). Compared to the population, our sample contains a slightly higher share of two parent families and a slightly lower share of families headed by single mothers. Our sample is also slightly skewed towards smaller families.

Our primary measure of income is individual total pretax income. This is

[3]For further information on these files, see https://alife-research.app/.

[4]Precise details on the linking procedure and other features of the data can be found in Deutscher (2018).

Table 2.1: Family characteristics in sample and population

Birth cohort	Full sample 1978-1982	Population Various
Family structure (%)		
Couple	84	81
Lone mother	11	15
Lone father	5	4
Median parental age at birth (years)		
Mother	27	26
Father	29	29
Family size(%)		
1	13	8
2	38	38
3	30	34
4	13	15
5	4	4
6	2	1
7 or more	1	1
Mean family size	2.7	2.8
Number of children	1,136,900	1,100,000
Number of children linked to parents	1,025,800	NA
Number of families	792,900	835,800

Notes: Population estimates are based on: Family Characteristics Survey 1997, Australian Bureau of Statistics (2017b) (family structure, 1973-82 birth cohorts); Births, Australian Bureau of Statistics (2010) (median parental age at birth, 1978-82 birth cohorts); and the 1991 Census, Australian Bureau of Statistics (1991) (family size, 1978-82 birth cohorts).

the most commonly used income measure across Australian Bureau of Statistics household surveys and is commonly used in the literature (e.g. Chetty et al. (2014)). In years where an individual has filed a tax return, this is their reported total income or loss. In years where an individual has not filed a tax return, it is the sum of individual salary and wages reported by employers, and taxable allowances, benefits and pensions reported by government welfare agencies, where available. Those who have no return or 3^{rd} party information are recorded as having zero total pretax income. This income measure includes labor and capital income, and taxable government payments such as unemployment and study benefits. It is prior to any tax deductions or offsets. Income variables are measured in 2015 dollars, adjusted for inflation using the headline consumer price index published by the Australian Bureau of Statistics (Australian Bureau of Statistics (2017c)). Below, we also discuss how different income measures relate to different concepts of mobility — for example, mobility in latent earnings potential or consumption possibilities.

In our baseline analysis, we calculate parental household income as the average of the combined annual income of the parent(s) over eleven years from 1991 to 2001. This choice of years balances the benefits of comparability with international studies of intergenerational mobility within a country (such as Chetty et al. (2014)), where shorter time periods are used, with the importance of averaging over many years to generate a better proxy for lifetime income (Mazumder (2005, 2016)). Child household income is defined similarly as the average of the combined annual income of the child and their most recently reported spouse (as at 2015) over the five years from 2011 to 2015. We examine the sensitivity of our national estimates to both of these choices.

Finally, for the purposes of defining geography for our regional estimates of mobility, children are assigned to the first geographic location associated with their primary parent. These locations either arise from a geocoded address or a residential postcode for the parent. In both cases we assign children to the associated Statistical Area 4 (SA4), as defined by the Australian Bureau of Statistics, 2011. These SA4s delineate broad labor markets, and are the closest Australian analogue to the commuting zones of Chetty et al. (2014).

2.3 Methodology

In this section we describe the various measures we use for our national and regional estimates.

2.3.1 Intergenerational elasticity

The intergenerational elasticity (IGE) has been the most commonly used measure of intergenerational mobility in economics. The IGE characterizes the rate of intergenerational persistence in a particular outcome (measured in logs such as log income) and one minus the IGE can be viewed as a gauge of intergenerational mobility. The IGE is the estimate of β obtained from the following regression:

$$y_{1i} = \alpha + \beta y_{0i} + \varepsilon_i \tag{2.1}$$

where y_{1i} is the log of income in the child's generation and y_{0i} is the log of income in the parents' generation.[5] The estimate of β provides a measure of intergenerational persistence in log income and $1 - \beta$ can be used as a measure of mobility. A value of 0.2, for example, suggests that if the difference in income between two families is 10 per cent, in the parent generation, then on average, approximately 2 per cent of this gap would be expected in the income of the children's generation.[6] An IGE of 0.2 would be indicative of a low degree of persistence and a fairly high degree of mobility compared to an IGE of 0.6.

2.3.2 Intergenerational correlation

In contrast to the IGE, the intergenerational correlation (IGC), or Pearson correlation, is a measure of positional mobility and indicates the degree to which a child can expect to occupy the same position in the income distribution. The IGC in income, for example, abstracts from any changes in the variance in income across generations. Formally, the IGE is equal to the IGC times the ratio of the standard deviation of log income in the child's generation to the standard deviation of log income in the parents' generation:

$$IGE = IGC \frac{\sigma_{y_1}}{\sigma_{y_0}} \tag{2.2}$$

[5]Often the regression will include age controls but few other covariates since β is not given a causal interpretation but rather reflects all factors correlated with parent income. We include financial year of birth dummies to control flexibly for the age of the child when income is measured.

[6]Since the data is measured in logs, the difference in log income approximates the percentage difference in income.

2.3.3 Rank-based measures

A closely related measure to the IGC is the rank-rank slope or Spearman correlation, which is obtained from the following regression:

$$r_{1i} = \alpha + \rho r_{0i} + \varepsilon_i \tag{2.3}$$

where r_{1i} and r_{0i} now represent the percentile rank in income in each respective generation. In this case, ρ provides an estimate of persistence in rank position and $1 - \rho$ provides a measure of positional mobility. In addition to estimates of rank persistence, following Chetty et al. (2014) we also use the rank-rank regression framework to calculate expected ranks at the 25^{th} and 75^{th} percentiles.[7] These statistics are useful for thinking about 'directional' mobility. For example, if the expected rank of individuals whose parents were at the 25^{th} percentile is the 40^{th} percentile then this would suggest average upward mobility of about 15 percentiles.

A key advantage of using the rank-based measures is that when using ranks based on the national income distribution, they can be used to make "apples to apples" comparisons of various subgroups of the population. Most notably for our purposes, we can compare regions within Australia to one another and be confident that our intergenerational rank mobility estimates mean the same thing in all places.

Finally, we also produce a matrix of transition probabilities across quintiles of the income distribution. This approach has also commonly been used to summarize intergenerational income mobility. Transition probabilities also provides measures of directional mobility as well as showing how mobility may differ at different points of the income distribution.

2.4 National results

We begin by showing some descriptive figures. Figure 2.1 plots, for each percentile in the parent income distribution, the mean log of total income of adult children

[7]We present these two percentiles purely for descriptive purposes. Given the rank-rank slope and the intercept (or expected rank at any given percentile) it is of course possible to calculate the expected rank at any percentiles of interest.

against the mean log of total income of their parents. What is immediately evident is that the relationship is nonlinear with a much flatter slope at lower and higher levels of parent income and a steeper slope at the middle of the distribution. This pattern is similar to that seen in Canada (Corak and Heisz (1999)) and the United States (Chetty et al. (2014)), where the IGE is also highest in the interior of the income distribution.[8] The patterns in the data shown in Figure 1 suggest that in addition to estimating the IGE which characterizes the entire distribution in one summary statistic, it is also useful to consider how mobility might differ at different points in the distribution, a point we return to when discussing nonlinearities and estimating transition probabilities.

Figure 2.1: Intergenerational income mobility – log of income

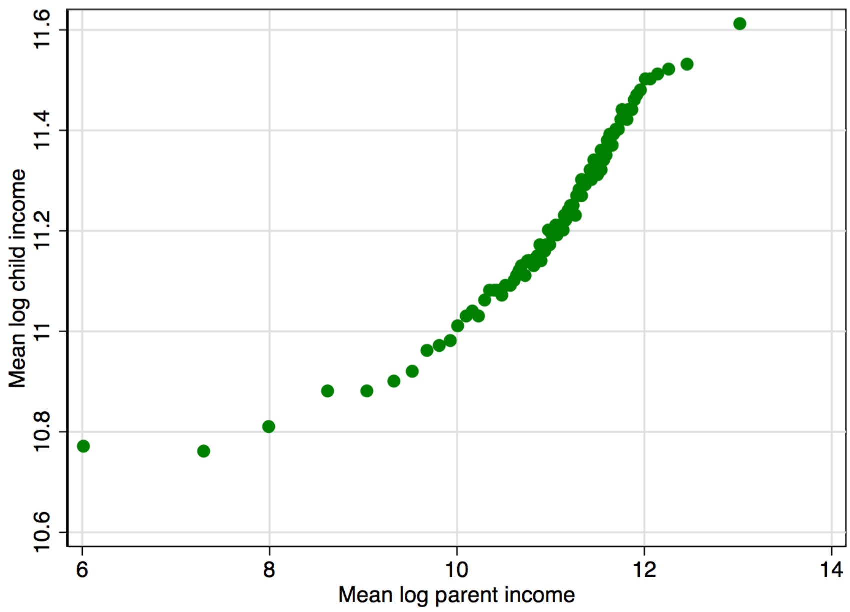

Notes: Chart plots the mean log child and parent total household income for each percentile of the parent income distribution.

Figure 2.2 plots intergenerational income mobility in terms of income ranks. Working with income ranks allows us to characterize upward mobility and down-

[8]A similar pattern is observed through in Sweden, outside the top percentile, where the IGE again increases Björklund et al. (2012). A similar, though less pronounced kick up in the top percentile is also apparent in Figure 1.

ward mobility in a more readily interpretable way. In addition, the slope of this line, the rank-rank slope, provides an additional measure of intergenerational persistence. Figure 2 also compares the rank-rank relationship in Australia to those in United States and Sweden — two advanced economies known to have vastly different experiences of mobility.

Consistent with the existing literature, the intergenerational relationship is more linear in ranks than it is in log incomes (Chetty et al. (2014)), making it easier to summarize in a single statistic. That said, there are interesting nonlinearities. In Australia, the rank-rank slope is much flatter between the 15^{th} and 95^{th} percentile than it is below the 15^{th} percentile and above the 95^{th} percentile. A greater slope and hence greater persistence in outcomes towards the bottom of the parent income distribution is in fact a feature across all three countries. This may point to different mechanisms underlying entrenched disadvantage. Similarly, Sweden has a pronounced increase in the slope at the top of the distribution, which may again point to different mechanisms, such as intergenerational wealth transfers (Björklund et al. (2012)). A comparison of the three countries suggest that expected ranks are highest for the bottom 10 per cent and top 1 per cent in Sweden, for the lower-middle income earners in Australia and for the upper-middle income earners in the United States. While the mapping of ranks to actual living standards and social standing will differ across all three countries, this highlights some nuances that can be lost in single summary statistics.

In Table 2 we present our baseline national estimates of intergenerational income mobility. In Panel A we consider four different measures of income: wages and salary, all private income (including investment income), total income (including government transfers) and disposable income (total income income minus taxes). These different measures of income capture different concerns. Wage and salary income may be viewed as most indicative of earnings ability as it presumably reflects productivity and excludes passive sources of income. Private or market income includes passive sources of income as an additional mechanism by which advantage may be passed from one generation to the next. Finally total and disposable income will result in measures accounting for redistribution first through transfers (total income), and then through the tax system (disposable income). Disposable income might also be most reflective of consumption possibilities and

Figure 2.2: Intergenerational income mobility – income ranks

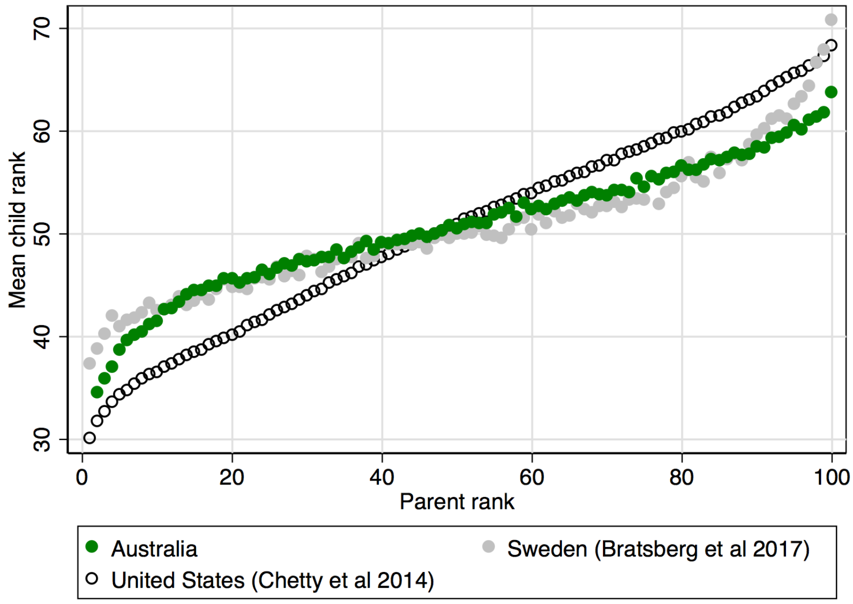

Notes: Chart plots the mean child and parent total household income rank for each percentile bin of the parent income distribution.

therefore of welfare.

In column (1) we find that the estimates of the IGE range from 0.107 for wages and salary to as high as 0.192 for private income. Our estimate for total income is 0.185. The estimates for the IGC (Pearson correlation) shown in column (2) are consistently lower, ranging from 0.114 to 0.159. Given the mechanical relationship between the IGE and IGC, this implies that the variance in the income measures is higher in the child generation than in the parent generation. The IGC abstracts from this increase in inequality, while the IGE captures it and is higher as a consequence. Finally, the estimates of the rank-rank slope (Spearman correlation) presented in column (3) range from 0.186 to 0.222. Our baseline estimate for the rank-rank slope in total income is 0.215. Across both the IGE and Spearman correlation, persistence is lowest when looking purely at wages and salaries, highest when looking at all private income and then progressively less for total and then disposable income.[9] This aligns with the typical effects of capital income, transfers

[9]The elasticity estimates for wages and salaries may appear strikingly low, particularly in

Table 2.2: National measures of intergenerational income mobility

	IGE	Pearson correlation	Rank-based		
			Spearman correlation	$E[r_{1i} \vert r_{0i} = 25]$	$E[r_{1i} \vert r_{0i} = 75]$
Panel A: Income definition					
Wages	0.107	0.114	0.186	45.8	55.1
	(0.001)	(0.001)	(0.001)	(0.0)	(0.0)
Private	0.192	0.150	0.222	44.8	55.9
	(0.002)	(0.001)	(0.001)	(0.0)	(0.0)
Total	0.185	0.159	0.215	45.0	55.8
	(0.001)	(0.001)	(0.001)	(0.0)	(0.0)
Disposable	0.175	0.148	0.211	45.1	55.7
	(0.001)	(0.001)	(0.001)	(0.0)	(0.0)
Panel B: Household, own or spousal income					
Women, household	0.181	0.156	0.211	46.5	57.0
	(0.002)	(0.002)	(0.001)	(0.1)	(0.1)
Women, own	0.166	0.149	0.174	37.4	46.1
	(0.002)	(0.002)	(0.001)	(0.0)	(0.1)
Women, spouse	0.117	0.136	0.126	55.7	62.0
	(0.002)	(0.002)	(0.002)	(0.1)	(0.1)
Men household	0.188	0.161	0.217	43.7	54.5
	(0.002)	(0.002)	(0.001)	(0.1)	(0.1)
Men, own	0.181	0.159	0.209	53.4	63.8
	(0.002)	(0.002)	(0.001)	(0.1)	(0.1)
Men, spouse	0.117	0.138	0.100	40.0	45.0
	(0.003)	(0.002)	(0.001)	(0.0)	(0.0)

Notes: Presents estimates of five different measures of intergenerational persistence for different income definitions and units of observation. The default is to estimate using total income at a household level for the full sample. In Panel A we vary income only: wages income is the self- or third-party reported individual salary and wages; private income is total income minus self- or third-party reported government payments; total income is as defined in the text; and disposable income is taxable income minus gross tax. In Panel B we split the sample into women and men (based on child gender) and vary whether the child's adult household income, own income or spouse income is the outcome of interest.

comparison to the international literature on earnings mobility. They are likely underestimates of true persistence in earnings since labour income, for the self-employed in particular, may appear under a number of different tax return labels. As a rough correction, if we restrict attention to and taxes on static measures of inequality. These various measures however, all tell a fairly consistent story that intergenerational persistence in Australia is quite low with estimates typically around 0.2 or lower. By way of contrast, estimates of intergenerational persistence in total income in the United States are typically around 0.5 or higher (Mazumder (2016)).

In columns (4) and (5) we present estimates of the expected rank at the 25^{th} and 75^{th} percentiles to illustrate the movements in income ranks implied by the high level of intergenerational mobility. We find that upward mobility is quite high as children whose parents were at the 25^{th} percentile can expect to rise nearly to the median at the 45^{th} percentile. Similarly those born into the 75^{th} percentile can expect to fall to the 56^{th} percentile. Naturally, behind these average experiences there is a diversity of outcomes, which we will later explore through the probabilities of transitioning between given quintiles of the income distribution.

In Panel B we consider how the estimates differ by the gender of the child, and whether it is the child themselves, their spouse or their combined outcomes that are considered.[10] This allows us to consider the role of household formation in driving the results. For both women and men, measures of persistence based on individual income are modestly lower when looking at their individual income, rather than their household income. This gap is slightly larger for women, pointing to a slightly greater role for household formation in driving their observed household outcomes. Nonetheless, the differences between the genders are mostly modest. While women have lower expected ranks in the income distribution, the relationships between their income levels and position in the income distribution and those of their parents are similar to those of men. There is only a slightly stronger connection between parent and child outcomes for men versus women, and the connection between parent and child spouse outcomes are fairly similar. This finding contrasts with the United States literature, which has found more substantial differences,

those children and parents for whom wages and salaries constituted at least 80 per cent of their total income, the intergenerational elasticity rises from 0.107 to 0.131. A more comprehensive measure of earnings mobility is outside the scope of this paper, but could potentially draw on the net income from working measure that underpinned the (now abolished) Mature Age Workers Tax Offset, a targeted Australian earned income tax credit discussed in some detail by Breunig and Carter (2018).

[10] Those without a spouse are coded as having zero spouse income, which means they are dropped in the IGE and Pearson correlations, but included in the rank specifications.

with lower persistence in outcomes for women, and driven more through household formation, though these differences tend to be smaller in studies examining more recent birth cohorts (Chadwick and Solon (2002); Chetty et al. (2014); Mitnik et al. (2015)).[11]

Finally, for robustness, we also estimated the IGE and the other parameters in total income using several other approaches shown in Appendix Table 1. First, given the nonlinearities at the tails of the income distribution, we produced estimates just using the middle 80 percentiles (10^{th} to 90^{th} percentile). As expected based on Figures 1 and 2, this boosts our estimate of the IGE slightly higher to 0.241 and lowers our estimate of the rank-rank correlation to 0.181. We also produce estimates where we use inverse probability weights to account for differences in the probability that a child is successfully linked to their parents.[12] This has a relatively modest effect on our results producing an IGE of 0.191 (up from 0.185) and a rank-rank correlation of 0.217 (up from 0.215). Finally, we produce a set of estimates using only the links we are most confident about, and weighting on the probability that a child has one of these high quality parent links.[13] Incorrect parent-child links may be expected to bias down our mobility estimates, but again the changes are relatively modest and in varying directions, with the IGE rising to 0.195 but the rank-rank correlation falling to 0.208.

As noted earlier, one disadvantage of summary statistics of intergenerational income mobility is the loss of finer detail about underlying movements. A world in which all children born into the 25^{th} percentile end up at the 45^{th} percentile is very different from one where this is simply the average across outcomes that span the full income distribution. A common way to capture this nuance is to

[11] The most recent and precise estimates of mobility to date in Australia (Murray et al. (forthcoming)) do not estimate mobility by gender due to sample size constraints.

[12] To do this we first calculate the percentile income ranks for all children, including those not linked to parents. We then calculate for each percentile bin the proportion of children who are linked to parents. The inverse of this provides the weight for the subsequent regressions. More complex approaches are possible, accounting for a wider set of potential covariates (for example, sex and location). However, this risks a false sense of precision given the inability to weight on the (unobserved) joint distribution of child and parent incomes, and given this limitation we have not conducted further weighting exercises.

[13] The ATO data includes a variable that captures the quality of the parent-child link on the interval [0,1]. We include only those links for which this is 0.95 or greater, which drops around 10 per cent of the sample.

examine transition probabilities. Table 3 presents the probability a child born into a given quintile of the parent income distribution transitions to each quintile in the child adult income distribution. While the most common outcome is that a child stays in the same quintile they were born into, there are large proportions moving both up, and down. The transition probability of 12.3 per cent from the bottom quintile to the top quintile again marks Australia out as among the most mobile of the advanced economies. An Australian child born into the bottom quintile is over 60 per cent more likely to reach the top quintile than a child born in the United States (where the transition probability is 7.5 per cent). Further, while Denmark has more mobility than Australia as measured by a rank-rank slope of 0.180 relative to 0.215, they have less upward mobility on this measure with a probability of transition from the bottom to the top quintile of 11.7 per cent (Boserup et al. (2013)).

Table 2.3: Intergenerational transition matrix

		Parent quintile				
		1	2	3	4	5
Child quintile	5	12.3	15.9	18.6	22.5	**30.7**
	4	15.5	19.0	21.1	**22.6**	21.9
	3	18.5	21.1	**21.5**	21.0	18.0
	2	22.7	**22.3**	20.9	18.7	15.4
	1	**31.0**	21.8	18.0	15.3	14.0

Notes: Shows the per cent frequency with which a child with parents in a given income quintile (column) ends up in given income quintile (row) themselves. The main diagonal is shaded grey, with figures in bold.

2.5 Sensitivity to the age and window of income observation

As the intergenerational literature has evolved, increasing attention has been paid to the importance of the length of time and ages over which incomes are observed. Intergenerational mobility can be greatly overstated if measured over too short a period, or too early or late in the lifecycle (Solon (1992); Mazumder (2005); Haider and Solon (2006)).

In Figures 3a and 3b we show the influence of the age and window over which parent and child incomes are measured for the IGE and the rank-rank correlation. Each series is centered on a given year of observation, with the length of the window of observation increasing along the x axis. By centering on a given year of observation, we fix the average age at which child or parent income is measured, allowing us to consider the influence of age and window width separately. A number of observations can be made.

First, persistence is higher when income is measured in mid-to-late working life. For children, measuring income in a window centered later in life (in their 30s in this case) yields notably higher measures of persistence. For parents, measuring income in a window centered on their late 40s to early 50s produces the highest measures of persistence. This latter result contrasts slightly with international findings that have suggested measuring earnings at slightly earlier ages. For example, Mazumder (2005) notes that from the late 40s incomes tend to become more volatile, introducing a potential downward bias to measures of persistence. However, such findings are typically for much earlier birth cohorts, and the results here may be consistent with the extension of working lives. In addition, unlike most detailed examinations of attenuation and life-cycle biases, we present results based on both parents' incomes — not just fathers. The incomes of women later in life may well be a better reflection of the endowments passed on to their children than those earlier in their working lives when child care may limit their labor force participation.

Second, persistence is higher when income is observed over a longer period of time. While this is particularly true for parents, it also holds for children, pointing to the existence of non-classical measurement error (Nybom and Stuhler (2017)).[14] It remains common for mobility studies to concern themselves primarily with the length of time over which parent incomes are observed, and the age at which child incomes are observed. The results in Figures 3a and 3b suggest the influence of length and age should be examined thoroughly for both generations.

Finally, comparing the two measures reveals some interesting differences. As has been noted in past work, the rank-rank correlation is less sensitive to the age

[14] As classical measurement error in the left hand side variable would not bias the estimated coefficients.

Figure 3a: Lifecycle and attenuation biases in the IGE

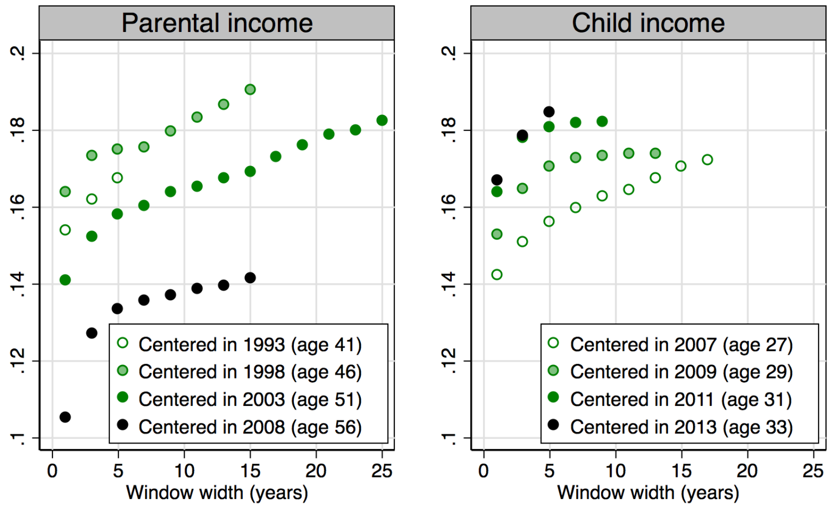

Figure 3b: Lifecycle and attenuation biases in the rank-rank correlation

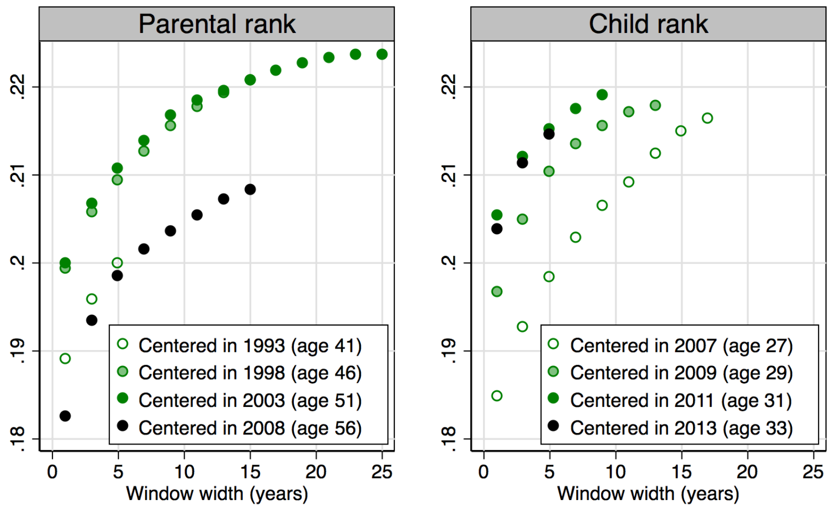

Notes: Presents estimates of the IGE and rank-rank correlation, varying the center year and width (in years) over which either parent or child log incomes or income ranks are observed. The corresponding average ages of the parents and children are shown in brackets. The default is to measure parent incomes over 11 years centered in 1996 and child incomes over 5 years centered in 2013.

and window over which incomes are measured (Mazumder (2016); Nybom and Stuhler (2017)). However, when faced with the trade-off between age of measurement and window width, the rank-rank correlation tends to rise more with window width — the rank-rank correlation is greatest when parent income is measured over the full 25 years, even though the average parental age will be 63 at the end of this window. Similarly, the rank-rank correlation is greatest when child income is measured over 9 years, even though this includes incomes observed in their 20s. On the other hand, the intergenerational elasticity estimates are greatest for somewhat shorter windows that are centered closer to the middle of parent and child working lives.

A remaining concern regarding our national estimates may be that while the individual potential biases may be small, their cumulative effect may be more significant. In Appendix Table 2, we present a set of 'conservative' estimates, where we measure parent income over the full 25 years (centered on an average age of 51) and child income over 9 years (centered on an average age of 31) as well as weighting for the probability of inclusion in the sample and, in the second row, using only the highest quality links. We also repeat this exercise with the slightly smaller windows, and ages closer to the mid-life, that appear better suited to the IGE. The IGE estimates range from 0.188 to 0.210 (compared to 0.185 in Table 2), while the rank-rank correlation ranges from 0.220 to 0.232 (compared to 0.215 in Table 2). Once again, these estimates place Australia among the most mobile advanced economies.

It is worth noting that our estimated levels of persistence are somewhat less than the current benchmark estimates for Australia. Murray et al. (forthcoming) estimate an intergenerational elasticity of 0.28 (s.e. 0.05) and a rank-rank correlation of 0.27 (s.e. 0.05) using the Household Income and Labour Dynamics in Australia (HILDA) survey. They note their elasticity estimate in particular is likely biased down as parent income is only observed over five years. These are the first direct estimates of intergenerational mobility in Australia and improve upon earlier estimates for which parent income had to be imputed (Leigh (2007); Mendolia and Siminski (2016)). Nonetheless, the estimates remain imprecise and subject to potentially complex attrition biases. For example, it is plausible that more mobile children — be it upwardly or downwardly mobile — are more likely to be lost in surveys, which would bias upwards the estimated persistence. For these reasons, at a minimum, we favor using our estimates when comparing Australia to other countries' mobility estimates based on administrative data.[15] It would be useful for future research to better explore the differences between estimates produced using survey and administrative data in Australia.

2.6 Regional estimates

We now present the first regional, rather than purely national, estimates of intergenerational mobility for Australia. Variations in intergenerational mobility within nations are increasingly used in the broader literature as a means of shining a light on potential mechanisms underlying intergenerational mobility — for example, in the work of Chetty et al. (2014) and Davis and Mazumder (2018) for the United States; Güell et al. (2018) for Italy and Bütikofer et al. (2018) for Norway.

In this section we focus on two measures of intergenerational mobility — the rank-rank slope and the expected rank at the 25^{th} percentile. Appendix Table 3 presents these estimates for all 87 Australian regions, alongside the sample size, estimates of the intergenerational elasticity and transition probability from the bottom to the top quintile. In future work, we will consider a still broader set of mobility measures.

Figure 5 presents three charts characterising intergenerational mobility for regions in Australia (shown in green) and the United States (shown in black, and based on Chetty et al. (2014)). The central chart shows a scatter plot of the expected rank of a child born to parents at the 25^{th} percentile of the income distribution against the rank-rank correlation in the region. These two measures are positively correlated in both countries. In the histograms we show the dispersion in the individual measures. While there is a notable dispersion in the Australian estimates, it is much less than that seen in the United States. For example, a child born to parents at the 25^{th} percentile in a mobile Australian region (at the 90^{th} percentile of regions ranked by mobility) can expect to end up only 8 percentile rank points higher than if they were born in an immobile Australian region (at the 10^{th} percentile of regions). For the United States, the gap in expected outcomes for a poor child between high and low mobility regions is nearly double this, at 15

[15]We are not aware of any systematic comparison of mobility estimates derived from survey data with estimates derived from administrative data.

Figure 2.3: Distribution of intergenerational mobility measures across regions in Australia and the United States

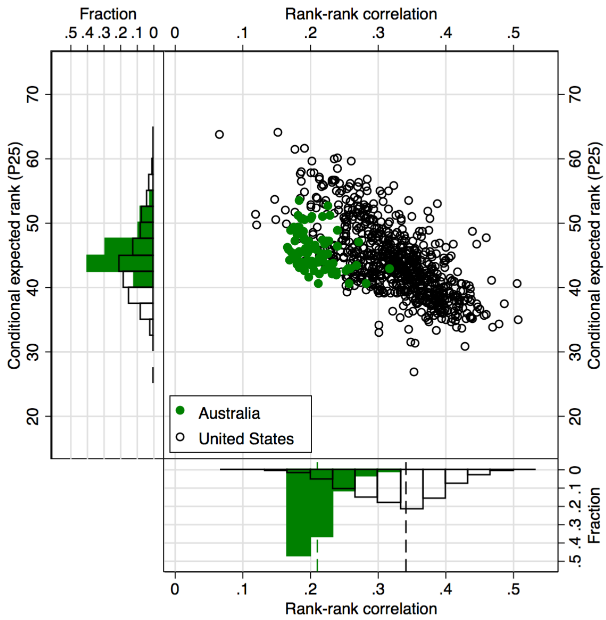

Notes: Presents estimates of the expected rank, conditional on being born into the 25^{th} percentile of the national parent income distribution, against the rank-rank correlation for similarly-sized regions in Australia (87 regions) and the United States (741 regions). A scatter plot of the joint distribution and histograms of the individual distributions of the two intergenerational mobility measures are shown.

percentile rank points.

The fact that there is less variation in mobility within Australian than within the United States is perhaps unsurprising for two reasons. First, Australia is a more centralised federation than the United States, with less geographic variation in policies that might influence mobility. Second, Australia's high level of mobility may mean more regions are near upper bounds on mobility determined by factors such as levels of assortative mating and skills transmission. This latter possibility is the histogram of the rank-rank slopes accompanying Figure 5, where there appears to be a missing left tail of regions where the rank-rank slope is in the low 0.10s.

Finally, Figure 5 maps intergenerational mobility for Australia and its two largest cities — Sydney and Melbourne. Once again we can see variation in the estimates, within the nation and the cities, though there appears to be less dispersion in mobility within Melbourne. An interesting set of contrasts between the two maps are the high expected ranks for children from poor families in some regional areas of Queensland and Western Australia that have fairly unexceptional rank-rank mobility.[16] This likely reflects the influence of the mining boom driving strong local labour markets in these regions over the period of observation. Norwegian research has shown that resource shocks can improve outcomes for children for poor families (Bütikofer et al. (2018)). A similar mechanism, but potentially more transient, may well be at play here.

Figure 2.4: Maps of intergenerational mobility within Australia and its two largest cities

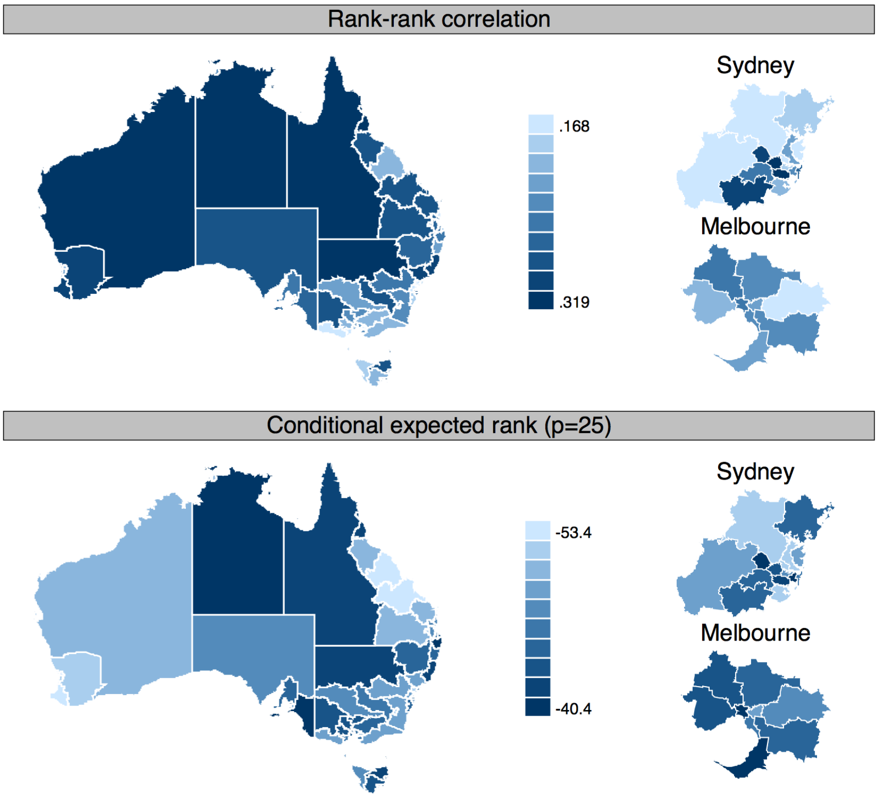

Notes: Presents estimates of the of the expected rank, conditional on being born into the 25th percentile of the national parent income distribution and the rank-rank correlation across Australian regions, including within the two largest cities, Sydney and Melbourne.

What then will we be able to learn about intergenerational mobility from variations across Australian regions? A threshold question is the extent to which

[16] Apparent in the South Western and North Eastern corners of Australia.

these variations arise from the influences of the regions themselves versus simple differences in the types of families that live there. Deutscher (2018) replicates and extends the work of Chetty and Hendren (2018a) in the Australian setting, finding that most of the differences in mobility across Australian regions is indeed causal. A child moving at birth between two Australian regions can expect to pick up around 70 per cent of the gap between the outcomes of permanent residents of those regions who have the same birth cohort and parent income rank. Further analysis of the regional estimates presented here may thus shed helpful light on the mechanisms underlying intergenerational mobility.

2.7 Conclusions

Intergenerational mobility is of key interest to policy-makers in Australia and beyond — it motivates many public policies, and has been thrown into sharp relief in the face of rising inequality. A new Australian evidence base aids domestic policy-makers; and it adds to the international literature by presenting a detailed picture of intergenerational mobility in a country with different demographics, institutions and economic circumstances.

We present the most precise and comprehensive set of estimates of intergenerational income mobility for Australia to date. Australia emerges as one of the most mobile of the advanced economies — not far off Nordic levels of mobility despite a much smaller welfare state. Having a parent 10 percentile rank points higher in the income distribution increases your expected rank by a little over 2 percentile rank points. This relationship is a little more pronounced at the bottom 15 per cent of the parent income distribution, possibly reflecting different transmission mechanisms behind entrenched disadvantage.

Finally, we examine differences in mobility across regions in Australia. While Australia has much less dispersion in mobility relative to the United States, differences still emerge, both within the country and individual cities. Both macroeconomic and more finely-grained factors appear to be in play. For example, the mining boom appears to have lifted the expected ranks of children in resource-rich states. But even within individual cities, such as Sydney, regions with measures of mobility at either end of the Australian experience sit alongside one another.

Since the turn of the century, the intergenerational mobility literature has devoted considerable attention to understanding the causal mechanisms that may underlie persistence in economic outcomes (Black and Devereux (2011)). This paper presents a new Australian evidence base, with a set of estimates robust to traditional measurement concerns. In doing so, it provides a further example of substantial intergenerational mobility outside the Nordic countries, and provides a new source of variation to exploit in studies seeking to understand the mechanisms behind intergenerational mobility.

2.A Additional tables

Table 2.4: National measures of intergenerational income mobility — robustness

	IGE	Pearson correlation	Rank-based		
			Spearman correlation	$E[r_{1i} \| r_{0i} = 25]$	$E[r_{1i} \| r_{0i} = 75]$
Middle 80%	0.241	0.195	0.181	46.2	55.2
	(0.002)	(0.002)	(0.001)	(0.0)	(0.0)
	747,000	747,000	772,600	772,600	772,600
Weighted	0.191	0.165	0.217	43.8	54.6
	(0.001)	(0.001)	(0.001)	(0.0)	(0.0)
	900,700	900,700	965,700	965,700	965,700
Highest quality links	0.195	0.161	0.208	44.0	54.4
	(0.002)	(0.001)	(0.001)	(0.0)	(0.0)
	807,700	807,700	855,100	855,100	855,100

Notes: Presents estimates of five different measures of intergenerational persistence varying the sample only: having first calculated income ranks, log income and normalized log income, we: restrict estimation to children from the middle 80% of the parent income distribution; weight children by the inverse of the probability that a child at the same percentile rank in the child income distribution is linked to parents; and restrict estimation to children with the highest quality links, where the primary parent's predicted probability of being a parent is at least 0.95, and weight children by the inverse of the probability that a child at the same percentile rank in the child income distribution is linked to such parents.

Table 2.5: National measures of intergenerational income mobility — conservative

	IGE	Pearson correlation	Rank-based		
			Spearman correlation	$E[r_{1i}\|r_{0i} = 25]$	$E[r_{1i}\|r_{0i} = 75]$
Panel A: Conservative rank-rank windows					
Weighted	0.188	0.173	0.232	43.4	55.0
	(0.001)	(0.001)	(0.001)	(0.0)	(0.0)
Highest quality links	0.205	0.180	0.225	42.5	53.7
	(0.002)	(0.001)	(0.001)	(0.0)	(0.0)
Panel A: Conservative IGE windows					
Weighted	0.198	0.169	0.227	43.5	54.9
	(0.001)	(0.001)	(0.001)	(0.0)	(0.0)
Highest quality links	0.210	0.172	0.220	42.5	53.5
	(0.002)	(0.001)	(0.001)	(0.0)	(0.0)

Notes: Presents estimates of five different measures of intergenerational persistence varying both the sample and the window over which income is observed. Panel A presents estimates based on windows of observation of: 25 years for parents, centered in 2003 and implying an average age of 51 years; and 9 years for children, centered at in 2011 and implying an average age of 31 years. Panel B presents estimates based on windows of observation of: 15 years for parents, centered in 1998 and implying an average age of 46 years; and 5 years for children, centered at in 2013 and implying an average age of 33 years.

Table 2.6: Regional estimates of intergenerational income mobility

SA4 code	SA4 name	Number of children	Mobility metric			
			IGE	Spearman correlation	$E[r_{1i} \mid r_{0i} = 25]$	$P[r_{1i} > 80 \mid r_{0i} \leq 20]$
101	Capital Region	8,800	0.173	0.198	45.3	10.9
			(0.015)	(0.011)	(0.4)	
102	Central Coast	13,800	0.164	0.182	43.8	11.0
			(0.011)	(0.009)	(0.3)	
103	Central West	10,500	0.196	0.208	46.4	12.3
			(0.013)	(0.010)	(0.4)	
104	Coffs Harbour - Grafton	6,800	0.160	0.190	43.6	10.4
			(0.016)	(0.013)	(0.4)	
105	Far West and Orana	6,600	0.210	0.260	42.7	10.7
			(0.017)	(0.012)	(0.4)	
106	Hunter Valley exc Newcastle	12,200	0.205	0.242	46.3	13.2
			(0.014)	(0.009)	(0.4)	
107	Illawarra	14,600	0.182	0.219	45.1	13.1
			(0.012)	(0.008)	(0.3)	
108	Mid North Coast	9,600	0.178	0.234	42.6	9.2
			(0.013)	(0.011)	(0.3)	
109	Murray	6,000	0.176	0.194	45.2	10.3
			(0.017)	(0.013)	(0.4)	
110	New England and North West	9,800	0.182	0.219	43.9	9.2
			(0.013)	(0.010)	(0.3)	
111	Newcastle and Lake Macquarie	17,400	0.174	0.210	45.7	11.6
			(0.011)	(0.007)	(0.3)	
112	Richmond - Tweed	10,700	0.167	0.206	42.4	9.7
			(0.012)	(0.010)	(0.3)	
113	Riverina	8,700	0.208	0.220	44.3	8.6
			(0.015)	(0.011)	(0.4)	
114	Southern Highlands and Shoalhaven	6,200	0.150	0.182	43.8	12.6
			(0.016)	(0.013)	(0.4)	
115	Sydney - Baulkham Hills and Hawkesbury	10,200	0.173	0.177	49.2	18.3
			(0.013)	(0.010)	(0.5)	
116	Sydney - Blacktown	14,000	0.191	0.240	42.3	9.7
			(0.012)	(0.008)	(0.3)	
117	Sydney - City and Inner South	5,500	0.173	0.199	41.4	11.7
			(0.017)	(0.013)	(0.4)	
118	Sydney - Eastern Suburbs	7,300	0.189	0.216	44.2	17.3
			(0.014)	(0.012)	(0.5)	
119	Sydney - Inner South West	21,700	0.206	0.257	42.7	12.9
			(0.009)	(0.006)	(0.2)	
120	Sydney - Inner West	8,200	0.178	0.178	45.1	14.4
			(0.014)	(0.011)	(0.4)	
121	Sydney - North Sydney and Hornsby	13,800	0.175	0.193	47.3	19.7
			(0.011)	(0.009)	(0.5)	
122	Sydney - Northern Beaches	9,900	0.151	0.168	46.1	16.1
			(0.014)	(0.010)	(0.5)	
123	Sydney - Outer South West	13,900	0.211	0.235	43.6	11.5
			(0.013)	(0.008)	(0.3)	
124	Sydney - Outer West and Blue Mountains	17,500	0.148	0.169	45.8	13.9
			(0.011)	(0.008)	(0.3)	
125	Sydney - Parramatta	15,900	0.229	0.270	43.3	11.2
			(0.011)	(0.007)	(0.3)	
126	Sydney - Ryde	6,300	0.162	0.172	48.8	19.5
			(0.017)	(0.013)	(0.6)	
127	Sydney - South West	14,300	0.143	0.211	43.8	12.0
			(0.010)	(0.008)	(0.3)	
128	Sydney - Sutherland	11,300	0.196	0.188	48.5	16.7
			(0.014)	(0.010)	(0.5)	
201	Ballarat	7,500	0.175	0.187	43.4	9.2
			(0.016)	(0.012)	(0.4)	
202	Bendigo	7,200	0.154	0.201	43.4	9.2
			(0.015)	(0.012)	(0.4)	
203	Geelong	12,600	0.147	0.171	44.1	9.9
			(0.012)	(0.009)	(0.3)	
204	Hume	8,600	0.168	0.188	43.8	9.4
			(0.015)	(0.011)	(0.4)	
205	Latrobe - Gippsland	14,100	0.147	0.184	45.5	12.6
			(0.011)	(0.009)	(0.3)	
206	Melbourne - Inner	11,600	0.176	0.213	40.4	9.9
			(0.012)	(0.009)	(0.3)	
207	Melbourne - Inner East	14,500	0.181	0.194	45.9	16.9
			(0.011)	(0.008)	(0.4)	
208	Melbourne - Inner South	13,700	0.177	0.197	44.1	13.0
			(0.011)	(0.008)	(0.4)	
209	Melbourne - North East	21,100	0.178	0.196	43.9	11.2
			(0.010)	(0.007)	(0.3)	
210	Melbourne - North West	14,700	0.188	0.205	43.1	10.8
			(0.011)	(0.008)	(0.3)	
211	Melbourne - Outer East	28,100	0.163	0.175	45.1	12.0
			(0.009)	(0.006)	(0.2)	
212	Melbourne - South East	26,200	0.155	0.195	43.6	11.6
			(0.008)	(0.006)	(0.2)	
213	Melbourne - West	22,500	0.141	0.189	43.0	9.9
			(0.009)	(0.007)	(0.2)	

Table 2.7: Regional estimates of intergenerational income mobility (cont.)

SA4 code	SA4 name	Number of children	Mobility metric			
			IGE	Spearman correlation	$E[r_{1i} \mid r_{0i} = 25]$	$P[r_{1i} > 80 \mid r_{0i} \leq 20]$
214	Mornington Peninsula	12,200	0.186	0.192	42.4	9.9
			(0.013)	(0.009)	(0.3)	
215	North West	8,000	0.195	0.229	43.5	8.6
			(0.017)	(0.011)	(0.4)	
216	Shepparton	7,100	0.170	0.197	43.4	8.6
			(0.016)	(0.012)	(0.4)	
217	Warrnambool and South West	7,100	0.164	0.174	45.3	9.8
			(0.017)	(0.012)	(0.4)	
301	Brisbane - East	9,000	0.143	0.179	47.1	11.8
			(0.015)	(0.011)	(0.4)	
302	Brisbane - North	8,000	0.203	0.220	47.3	15.0
			(0.017)	(0.012)	(0.4)	
303	Brisbane - South	11,800	0.179	0.222	46.6	14.0
			(0.012)	(0.009)	(0.4)	
304	Brisbane - West	7,700	0.166	0.185	49.2	18.7
			(0.016)	(0.012)	(0.6)	
305	Brisbane Inner City	5,300	0.183	0.201	45.8	15.6
			(0.018)	(0.014)	(0.6)	
306	Cairns	9,500	0.213	0.254	42.5	11.4
			(0.012)	(0.010)	(0.4)	
307	Darling Downs - Maranoa	5,700	0.165	0.229	47.1	12.4
			(0.016)	(0.014)	(0.4)	
308	Fitzroy	10,800	0.200	0.220	50.9	18.8
			(0.015)	(0.010)	(0.4)	
309	Gold Coast	13,800	0.163	0.182	43.4	11.2
			(0.011)	(0.009)	(0.3)	
310	Ipswich	11,300	0.162	0.216	44.6	10.8
			(0.013)	(0.010)	(0.3)	
311	Logan - Beaudesert	15,000	0.174	0.218	44.0	9.9
			(0.011)	(0.008)	(0.3)	
312	Mackay	8,200	0.172	0.184	53.4	20.3
			(0.015)	(0.011)	(0.5)	
313	Moreton Bay - North	8,500	0.157	0.214	44.4	9.9
			(0.014)	(0.011)	(0.3)	
314	Moreton Bay - South	7,100	0.177	0.182	48.1	13.6
			(0.020)	(0.013)	(0.5)	
315	Queensland - Outback	3,800	0.270	0.319	42.7	11.3
			(0.021)	(0.016)	(0.6)	
316	Sunshine Coast	10,600	0.143	0.175	44.6	12.3
			(0.014)	(0.011)	(0.3)	
317	Toowoomba	7,300	0.161	0.191	47.7	13.4
			(0.016)	(0.012)	(0.4)	
318	Townsville	10,300	0.210	0.225	46.9	14.1
			(0.015)	(0.010)	(0.4)	
319	Wide Bay	12,300	0.176	0.226	47.0	12.4
			(0.011)	(0.010)	(0.3)	
401	Adelaide - Central and Hills	11,400	0.180	0.187	45.3	12.9
			(0.013)	(0.009)	(0.4)	
402	Adelaide - North	19,700	0.184	0.233	42.0	8.2
			(0.009)	(0.007)	(0.2)	
403	Adelaide - South	17,700	0.177	0.204	43.9	10.7
			(0.010)	(0.008)	(0.3)	
404	Adelaide - West	9,000	0.174	0.240	41.8	9.5
			(0.013)	(0.010)	(0.3)	
405	Barossa - Yorke - Mid North	4,800	0.145	0.213	43.7	8.0
			(0.018)	(0.015)	(0.5)	
406	South Australia - Outback	4,800	0.224	0.227	45.1	11.4
			(0.026)	(0.014)	(0.5)	
407	South Australia - South East	7,600	0.191	0.215	42.0	7.6
			(0.017)	(0.012)	(0.4)	
501	Bunbury	6,800	0.190	0.230	51.0	17.2
			(0.016)	(0.013)	(0.5)	
502	Mandurah	2,500	0.154	0.227	52.5	21.8
			(0.023)	(0.022)	(0.8)	
503	Perth - Inner	5,000	0.159	0.183	51.1	22.5
			(0.016)	(0.014)	(0.7)	
504	Perth - North East	8,700	0.158	0.188	50.6	20.2
			(0.015)	(0.012)	(0.4)	
505	Perth - North West	20,000	0.182	0.202	50.6	19.3
			(0.009)	(0.007)	(0.3)	
506	Perth - South East	19,000	0.175	0.192	50.4	19.0
			(0.010)	(0.008)	(0.3)	
507	Perth - South West	13,000	0.181	0.204	50.9	18.5
			(0.011)	(0.009)	(0.4)	
508	Western Australia - Outback	9,700	0.258	0.273	46.9	15.5
			(0.015)	(0.011)	(0.5)	
509	Western Australia - Wheat Belt	6,200	0.204	0.241	48.7	16.9
			(0.018)	(0.014)	(0.5)	
601	Hobart	11,800	0.229	0.259	40.5	8.1
			(0.015)	(0.009)	(0.3)	
602	Launceston and North East	7,700	0.161	0.222	42.6	8.2
			(0.014)	(0.011)	(0.4)	

Table 2.8: Regional estimates of intergenerational income mobility (cont.)

SA4 code	SA4 name	Number of children	Mobility metric			
			IGE	Spearman correlation	$E[r_{1i} \mid r_{0i} = 25]$	$P[r_{1i} > 80 \mid r_{0i} \leq 20]$
603	South East	1,400	0.192	0.184	43.0	6.8
			(0.036)	(0.029)	(0.8)	
604	West and North West	7,300	0.158	0.181	44.6	9.6
			(0.016)	(0.012)	(0.4)	
701	Darwin	5,300	0.181	0.199	47.0	16.4
			(0.019)	(0.013)	(0.6)	
702	Northern Territory - Outback	2,500	0.267	0.283	40.5	7.2
			(0.027)	(0.019)	(0.8)	
801	Australian Capital Territory	18,000	0.184	0.191	48.5	16.7
			(0.011)	(0.008)	(0.4)	

Notes: Presents estimates of intergenerational mobility for those born in Australia in the 1978-82 financial years. Parent household total pretax incomes are measured from 1991-2001, while the total household incomes of the adult children are measured from 2011-15. Sample sizes rounded to the nearest 100 and standard errors in parentheses.

2.B Robustness of key measures to treatment of missing values

In this Appendix, we examine the sensitivity of our national measures of intergenerational mobility to the treatment of missing values — years in which child or parent incomes are not observed. Recent work has noted the sensitivity of the intergenerational elasticity to such assumptions, citing it as partial justification for adopting new measures of intergenerational mobility. For example, Chetty et al. (2014) promotes the rank-rank slope while Mitnik et al. (2015) proposes a new elasticity measure, which we also present here.[17] There are many potential ways to treat missing values in income data. We consider the following:

- Imputing \$1 to all missing values
- Imputing \$1,000 to all missing values
- Imputing \$10,000 to all missing values
- Dropping annual missing values
- Dropping lifetime missing values

Recent concerns have been in the context of missing values for child income, so we begin by applying these transformations to missing values in the child and child's spouse income histories (Panel A). However, we go on to apply the same transformations to missing values in the parents' income histories (Panel B) and both child and parents histories (Panel C). In all cases we apply the same treatments to negative income, though negative incomes are sufficiently rare that this does not influence the conclusions drawn.

Table 2.9 presents the results from this exercise. As expected, the intergenerational elasticity and correlation are much more sensitive to the treatment of zeroes than the rank based mobility measures, across all panels. However, when only concerned with child household missing values, this sensitivity is greatly reduced

[17] This measure can be and is estimated here as the coefficient from a Poisson regression of child income on parent log income, with birth cohort dummies.

when imputations are restricted to more plausible values. For example, it is difficult to imagine a situation in which it is appropriate to impute an income of $1 — reasonable imputed values for earnings capacity, subsistence income or similar would likely be much higher. Strikingly, Panels B and C show that the elasticity estimates are even more sensitive when missing values in parent income are also treated. This is true even for the IGE measure proposed by Mitnik et al. (2015), which is robust to the treatment of missing values in the child generation.

Table 2.9 provides an important caveat on the IGE and Pearson correlation measures. That said, the range of values apparent still mark Australia out as a particularly mobile advanced economy. The exercises also highlight the potential importance of missing values in parent income histories, as well as those in child income histories. An assessment of the most appropriate treatment of these missing values is well beyond the scope of this paper. Such an exercise would need to include an assessment of the underlying processes generating missing values. For example, missing values arising from unemployment, caring responsibilities or emigration may differ substantially in the information they carry about expected lifetime incomes.

Table 2.9: National measures of intergenerational income mobility

	IGE	IGE Mitnik	Pearson correlation	Rank-based		
				Spearman correlation	$E[r_{1i}\|r_{0i}=25]$	$E[r_{1i}\|r_{0i}=75]$
Panel A: Child missing values						
Impute $1	0.26	0.17	0.12	0.22	45.0	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Impute $1k	0.20	0.17	0.17	0.22	45.0	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Impute $10k	0.17	0.16	0.18	0.22	45.0	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Drop annual	0.17	0.16	0.17	0.22	44.9	55.8
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Drop lifetime	0.19	0.16	0.16	0.22	44.9	55.8
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Panel B: Parent missing values						
Impute $1	0.07	0.07	0.13	0.22	45.0	55.8
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Impute $1k	0.17	0.15	0.15	0.22	45.0	55.8
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Impute $10k	0.26	0.22	0.16	0.21	45.1	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Drop annual	0.22	0.20	0.14	0.18	46.4	55.5
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Drop lifetime	0.19	0.17	0.16	0.20	46.0	55.8
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Panel C: Child and parent missing values						
Impute $1	0.12	0.07	0.16	0.22	45.0	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Impute $1k	0.19	0.15	0.23	0.22	45.0	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Impute $10k	0.24	0.21	0.22	0.21	45.0	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Drop annual	0.21	0.19	0.15	0.19	46.1	55.7
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)
Drop lifetime	0.19	0.16	0.16	0.20	45.8	55.9
	(0.00)	(0.00)	(0.00)	(0.00)	(0.0)	(0.0)

Notes: Presents estimates of intergenerational mobility for those born in Australia in the 1978-82 financial years. Parent household total pretax incomes are measured from 1991-2001, while the total household incomes of the adult children are measured from 2011-15. Each row applies a different treatment to missing values, either imputing a value to them, dropping them or treating them as zeroes and dropping only lifetime zeroes. Standard errors in parentheses.

Chapter 3

Place, jobs, peers and the teenage years: exposure effects and intergenerational mobility

Abstract

I show that where a child grows up has a causal effect on their adult income, but that place matters most in the teenage years. I use variation in the age at which Australian children move to identify this pattern of place exposure effects. I explore two potential explanations. First, this pattern is partly explained by the fact that spending more years in a place in adolescence lifts the probability of entering the associated local labor market and earning any corresponding wage premium. Second, I identify long-lasting peer effects using cross-cohort variation in peer parental income among permanent postcode residents.

The author has been supported by a Sir Roland Wilson scholarship, funded by the Australian Treasury and the Australian National University. The helpful suggestions and advice of Robert Breunig, Lorenzo Cappellari, Tue Gorgens, Nathan Hendren, Bhash Mazumder, Pat Sharkey and Jan Stuhler have been greatly appreciated. I would like to thank Julia Neville, Thomas Abhayaratna, Bruce Bastian, Matt Power and Julia Rymasz for their invaluable expertise, advice and support accessing the intergenerational data. This research uses data from the Australian Taxation Office. All findings, opinions and conclusions are those of the author and do not necessarily represent the views of the Australian Government or any of its agencies. This chapter uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Survey was initiated and is funded by the Australian Government Depart-

3.1 Introduction

A new wave of studies highlights that where a child grows up influences their adult income, education and social outcomes (Chetty and Hendren (2018a); Chetty et al. (2016); Chyn (2018); Damm and Dustmann (2014)). Previously, correlations in the outcomes of neighboring children sat alongside more tenuous evidence for an underlying causal relationship. Yet knowing when and why place matters is just as important as knowing if place matters. How much could we gain by intervening earlier in childhood?

I find that where an Australian child grows up has a causal effect on their adult outcomes, but place matters most in the teenage years. I show this using de-identified Australian intergenerational tax data. I explore two explanations for the observed place effects — the fact that where you grow up influences where you end up (local labor market conditions) and who you grow up with (peer effects). This chapter provides further evidence for causal place effects, but highlights the importance of the teenage years. It also suggests a role for mechanisms – local labor markets and peers – that may be more challenging to redress than simply the quality of local institutions and policy settings.

I begin with the empirical framework introduced in the landmark study by Chetty and Hendren (2018a). I use variation in the *age* at which children move within Australia to identify the causal effect of exposure to place. This methodology avoids the selection bias inherent in simple comparisons between those choosing to live in different places or move between them. The key identifying assumption is that selection effects do not vary with the age at which children move.

The ability to observe children moving from infancy onwards, and hence detect when place matters most, is the key advantage of the Australian data. Before age 11, the outcomes of children who move converge to those of the permanent residents of their destination at a rate of around 1.1% for each year spent in the destination. After age 11, this rate of convergence is around 4.2%. These find-

ment of Social Services (DSS) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this chapter, however, are those of the author and should not be attributed to the Australian Government, DSS or the Melbourne Institute. The research plan was approved by the Australian National University Human Research Ethics Committee, protocol number 2017/832.

ings have implications for the interpretation and use of the estimates from the influential work of Chetty and Hendren (2018a,b). As those authors note, their estimates by necessity relied on incomplete child histories, capturing moves from age nine onwards.[1] This chapter's findings suggest that extrapolating back from their estimates could overstate the effect of place.[2] In the Australian setting, a child moving at birth to a place with one percentile higher incomes for permanent residents would receive around 70% of that effect themselves. Extrapolating back the exposure effects observed from age nine onwards yields a higher (biased) estimate of around 100%. On the other hand, the relative effects of specific places identified in Chetty and Hendren (2018b) could well capture the period in which location matters most — providing more comfort that they would be robust to complete childhood histories.

The pronounced sensitivity of the teenage years complements the existing human capital literature, which has long identified the importance of early childhood (e.g. Case et al. (2005); Cunha and Heckman (2007); Currie and Almond (2011)). For example, shocks to fetal health and targeted policy interventions in early childhood have been found to have long-lasting consequences. This chapter finds that the causal effects of the places considered here are most potent later in childhood — perhaps reflecting differences between these mechanisms versus those most frequently examined in prior literature. The second part of the chapter examines two such mechanisms that may result in a causal effect of place and explain the importance of adolescence.

First, place can matter because where you grow up is often where you end up working. Controlling for a child's location in adulthood, interacted with their cohort, halves the estimated exposure effects. This may overstate the role of local labor markets as it controls for an endogenous outcome — place effects may operate in part by giving children the skills required in the stronger labour markets. To provide a clearer test, I examine whether the outcomes of those who move better

[1] Chetty and Hendren (2018a) use outcomes at age 24 for the estimation of the average causal effects of place, which limits them to moves from age nine. They use outcomes at age 26 for the estimation of the causal effects of specific places, which limits them to moves from age 11. For earlier outcomes, moves at earlier ages were captured.

[2] Chetty and Hendren are clear about the assumptions inherent in such extrapolations of average or specific causal effects of place. They also frequently take a conservative approach, for example by extrapolating over less than the full period over which exposure to place may matter.

reflect those who grew up in a place, or those who end up there — and hence subject to the associated local labor market but not necessarily the childhood environment. These exogenous controls also lower the estimated exposure effects. Depending on the specification, anywhere between 15-55% of the effect of exposure to place in the teenage years can be explained by local labor markets, with the least restrictive specification consistent with the midpoint of this range.[3]

Second, place can matter because of who you grow up with. I identify positive and statistically significant peer effects among permanent residents of postcodes using cross-cohort variation in peer parental income rank. A 10 percentile rank point increase in the mean parent income rank of your postcode-cohort peers is associated with a 0.2-0.3 point increase in your own household income rank at age 24. In contrast, in placebo tests that replace an individual's peers with those from the financial years either side, the estimated peer effects are small and statistically insignificant. Finally, the peer effects appear to be driven by interactions between same-sex peers, which makes some nonsocial explanations far less likely. Given the large psychological literature on the increasing importance of peers during the teenage years (Brown and Larson (2009)), such peer effects may play a role in explaining the observed pattern of exposure effects. While I find that simple controls for peers cannot explain the observed effects of place, this may reflect the limitations of working with indirect measures of peers.

A final contribution of this chapter is in validating and extending the research design of Chetty and Hendren (2018a). I show their approach can be applied in Australia — a country with less than a tenth of the population of the United States, and less geographic variation in intergenerational mobility. I am also able to replicate many of their validation exercises, with largely comforting results. The observed exposure effects remain under a range of specifications and sample selection criteria, and when comparing same-sex siblings. The approach may thus be able to yield fresh insights from further applications in a wide variety of settings.

[3]These tests for a role for local labor market conditions go beyond those in Chetty and Hendren (2018a). In their work, the described controls for a child's location in adulthood are time invariant and thus capture only persistent differences in local labor market conditions. The exogenous controls included in this chapter, based on those ending up in a location, are also new. That said, we might expect a stronger role for local labor market conditions in determining Australian place effects for a number of reasons, which are discussed in more detail in Section 3.6.1

3.2 Related literature

Many studies have documented correlations in the economic outcomes of children growing up in the same neighborhood (Black and Devereux (2011); Sharkey and Faber (2014)). However, the extent to which this reflected sorting by families into neighborhoods remained an open question. Early experimental studies typically found little evidence of neighborhoods affecting economic outcomes (Katz et al. (2001); Oreopoulos (2003); Ludwig et al. (2008)). Using family fixed effects, Aaronson (1998) found evidence for a causal role of neighborhoods, but small sample sizes were a constraint on this and other early quasi-experimental work.

In their seminal work, Chetty and Hendren (2018a) introduce a new and ingenious approach — using variation in the age at which children move locations — and apply it to administrative data in the United States to identify compelling causal effects of place. Crucially, these effects vary according to the time *exposed* to the location. A separate paper revisits the Moving to Opportunity experiment and also finds that location matters according to time exposed, with beneficial outcomes for those moving to a lower-poverty neighborhood before the teenage years (Chetty et al. (2016)).

A recurring theme in Chetty and Hendren (2018a) and Chetty et al. (2016) is that these exposure effects appear relatively constant — the outcomes of children who move to a better neighborhood improve *linearly* in proportion to the amount of time they spend growing up there. However, as Chetty and Hendren note, the observed linearity is based on children moving at ages nine and above, and extrapolating from this is "a strong assumption that should be evaluated in future work".[4] Depending on how their estimates are extrapolated, a child moving at birth to a new location in the United States could be expected to pick up anywhere from 56% to 92% of the difference in permanent residents' expected outcomes between their origin and destination.[5] Further, their causal estimates for specific places are based on moves made from age 11 onwards, and thus potentially miss differences, if any, in the causal effects of places in early childhood.

There is no reason to expect that each year of childhood in a better location will

[4] See pages 4 and 53 of Chetty and Hendren (2018a).

[5] Depending on whether their 4% annual rate of convergence is applied only to the 14 years that they observe it over, from age 9 to 23 years, or to the full 23 years.

be equally important.[6] Indeed, the literature most closely related in methodology generally finds this is not the case. Age at migration to a country often matters most in the teenage years, with negligible or marginal gains to English language proficiency, years of schooling, earnings, height and cognitive skills from an extra year in the destination country prior to adolescence (e.g. Bleakley and Chin (2004); Van den Berg et al. (2014); Basu (2018)).[7] Exposure to place within a country can also matter more within the teenage years for outcomes such as teen parenthood (Wodtke (2013)). Indeed, such patterns are also apparent in Chetty and Hendren (2018a) when college attendance and teen birth are the outcomes considered.[8]

This chapter adds to the literature by showing that exposure to place, at least on the scale considered here, matters most in the teenage years. Note this need not contradict the large literature on the sensitivity of early childhood (e.g. Case et al. (2005); Cunha and Heckman (2007); Currie and Almond (2011)) — it may simply be the large neighborhoods examined here mask substantial variation in what matters in the early years, while exhibiting variation in what matters most during the teenage years. The shocks entailed by moving may simply differ to those examined in the literature on early childhood.

So why might place matter? An important observation in the economic literature on neighborhood effects has been that where you grow up might matter because it influences the local labor market you end up in. In Page and Solon (2003a), most of the correlation in the adult earnings of neighboring boys in the United States is explained by the large earnings premium in urban areas and the high correlation between childhood and adult urbanicity.[9] A potential role for local labor markets has also been found in studies of neighborhood effects in Norway

[6] I occasionally refer to locations with higher predicted total household income ranks as having 'better' outcomes or as 'better' locations. This is a convenient shorthand and is not intended as a definitive normative statement, as many features of place, well beyond those considered here, will be relevant for individual wellbeing.

[7] This is more typically framed in terms of there being a critical (pre-teen) age, beyond which one's ability to catch up to those born in the country is progressively more limited.

[8] Chetty and Hendren (2018a) note on teen birth: "the gradient is especially steep [exposure effects are especially high] between ages 13 and 18, suggesting that a child's neighborhood environment during adolescence may play a particularly important role in determining teen birth outcomes" (page 52). This chapter suggests the same could be said for broader economic outcomes.

[9] The same authors find the same is true of neighboring girls (Page and Solon (2003b)).

and Australia (Raaum et al. (2006); Overman (2002)). I find that up to half the exposure effects identified in this chapter can be explained by local labor market conditions.[10]

Another explanation for a causal role of place is the influence of peers. Peer relationships become more salient and complex from adolescence (Brown and Larson (2009)), and potentially explain part of the sensitivity of the teenage years. Peer effects have received a lot of attention in the education literature (Sacerdote (2011)), reflecting both their potential role in school settings and the availability of data with clearly identified cohorts. School-based studies have also gone well beyond student achievement, to broader behaviors with potentially long-lasting consequences, including the decision to drop out (e.g. Gaviria and Raphael (2001)). Explicit studies of long-run effects are rare, but one notable exception is Black et al. (2013), where cross-cohort variation is used to study the causal effects of teenage peers on long-run outcomes in Norway.[11] I similarly use cross-cohort variation to demonstrate the presence of statistically significant, large and lasting peer effects in Australia.

3.3 Data

I use de-identified Australian Taxation Office (ATO) intergenerational data drawn from federal income tax returns from the 1991 to 2015 financial years. Parent-child links primarily come from linking children to adults living at the same address when the child registers for a Tax File Number (TFN): a unique personal identifier issued by the federal government. The links are also informed by a short period in which children were claimed on tax returns. Children are linked to a 'primary parent', the adult considered most likely to be a parent based on shared address, shared

[10]If the influence of an individual's initial local labor market fades with time, then this may also help explain Danish findings that neighborhoods account for a large share of the variance in permanent earnings between siblings early in the working life, but a negligible share beyond age 30 (Bingley et al. (2016)).

[11]Black et al. (2013) find girls' outcomes benefit from having a higher proportion of female peers, while boys' outcomes are harmed. They find little evidence that peers' mothers' matter, but peers' fathers' income seems to matter for boys. Their study differs in ways that make it difficult to directly compare the results. Most notably, they define peers as those who were in the ninth grade in the same year and school — a more direct but also much narrower definition.

last name and age at the child's birth.[12] Children are also linked to the primary parent's first reported spouse over the period 1991-2015.[13] More details on the construction of the data is in Appendix 3.E.[14]

3.3.1 Sample definitions

I initially restrict attention to a sample of those who: 1) were born in Australia between the 1978 and 1991 financial years; and 2) remained resident in Australia through to 2015. Earlier cohorts cannot be reliably restricted to those born in Australia and later cohorts are too young to observe at age 24, the main outcome measure in this chapter. The resulting families closely mirror population benchmarks for family structure, median parental age at birth and family size (Table 3.1). There are slightly more children in couple families and fewer in lone mother families, likely reflecting re-partnering of lone mothers that is captured in our panel but not in the population cross-section. For the 1978-1991 birth cohorts considered here, I have parent links for around 3.1 million children, or 92% of those in the data.[15]

I split the full sample in two based on residential history — creating permanent resident and mover subsamples. I use Statistical Area 4 (SA4), as defined by the Australian Bureau of Statistics (ABS), as my unit of geography. There are 107 SA4 covering Australia, typically with populations in the range of 100,000-500,000.

[12]Intergenerational links based on shared residential addresses are behind the Canadian data used in a number of influential studies (Corak and Heisz (1999); Oreopoulos (2003); Corak and Piraino (2011)). However, the Australian data has the notable advantage of drawing on earlier registration addresses and additional information such as claims for children on tax returns. The data most closely resembles that used by Chetty et al. (2014) in its quality and coverage.

[13]Provided the youngest child in the family is no older than 24 at the time the parent reported the spouse. Questions on spouse details have always appeared on tax returns over this period, but answers have not always been required. Linking to the earliest reported spouse rather than the spouse at a fixed point in time or in childhood attempts to cater for the resulting underreporting of spouses. Nonetheless, a restriction based on the age of the youngest child is used to avoid linking to spouses that were not present during childhood.

[14]The intergenerational data was constructed by ATO staff. The author wrote code which was submitted to the ATO to be run internally on the resulting datasets to produce the results presented in this chapter.

[15]This is on par with the 91% link rate achieved by Chetty et al. (2014) for the 1980-1988 cohorts used in Chetty and Hendren (2018a) — see Online Appendix Table 1 in Chetty et al. (2014).

Table 3.1: Family characteristics in sample and population

Birth cohort	Full sample 1978-1991	Population Various
Family structure (%)		
Couple	86	80
Lone mother	9	16
Lone father	5	4
Median parental age at birth (years)		
Mother	27	27
Father	30	30
Family size(%)		
1	12	13
2	38	40
3	30	30
4	13	12
5	4	3
6	2	1
7 or more	1	1
Mean family size	2.7	2.6
Number of children	3,376,800	3,185,400
Number of children linked to parents	3,108,000	NA
Number of families	1,834,300	1,772,300

Notes: Population estimates are based on: Family Characteristics Survey 2003, Australian Bureau of Statistics (2017b) (family structure, 1979-93 birth cohorts); Births, Australian Bureau of Statistics (2010) (median parental age at birth, 1978-91 birth cohorts); and the 1991 Census, Australian Bureau of Statistics (1991) (family size, 1978-91 birth cohorts).

The choice of SA4 is motivated by the methodology — these units strike the best balance between having a large permanent resident population (and thus more precise predicted outcomes for those who move) and having a large population of moves (and thus a larger sample).[16] Permanent residents are those whose primary parent files from only one SA4 from 1991 through to the year the child turned 35. Movers are those whose primary parents file from multiple SA4 — I focus on those

[16]The next unit up in the main structure of the Australian Statistical Geography Standard is the state/territory level — Australia has only eight of these and only 27% of the moves in the sample cross state/territory lines. The next unit down in the structure is the Statistical Area 3 level, with populations in the range of 30,000-130,000 — too small to generate precise predicted outcomes based on permanent residents.

moving once (with parents filing from two distinct SA4).

I limit the influence of measurement error in the presence and timing of moves in three ways. First, I consider only moves between SA4 where the primary parent filed in at least two years from both. Second, I consider only moves where the parent's first year filing from the destination immediately follows their last year filing from the origin — this ensures the child's age at move is known to within a year. Third, I restrict attention to those moving at least 15 kilometers, thus dropping short moves that just happen to cross SA4 boundaries.[17]

To investigate peer effects, I begin by restricting attention to those growing up in the same postcode. I create a set of permanent postcode residents whose primary parent files from only one postcode from 1991 through to the year the child turned 20. Large neighborhoods are no longer required by the methodology, and this is closer to the common conception of a neighborhood and the scale on which social interactions take place. It also more closely resembles the existing literature, where peer effects are commonly examined in the context of relatively small groups (e.g. workplaces, grades or classrooms). Nonetheless, I check for differences in estimated peer effects by postcode size.[18]

In all analyses, individuals with strictly negative parent or child adult total pre-tax family income are dropped, as negative income is typically associated with high wealth and hence a poor indicator of actual economic wellbeing.

3.3.2 Variable definitions

Individual income is defined as total pre-tax income.[19] In years where an individual has filed a tax return, this is their reported total income or loss. In years where

[17]In Chetty and Hendren (2018a) only moves of 100 miles or more are considered. A less restrictive condition is used here as Australia's population is highly concentrated in a small number of major cities — around a third of Australians live in Sydney or Melbourne. Moves within these cities are thus of policy interest and empirically important. In the unrestricted mover subsample, 44.3% of moves are within the same city while 26.3% are within Sydney or Melbourne.

[18]The postcode identifiers are randomized in the data, and are thus unable to be aggregated to higher level geographies to examine the effect of neighborhood size.

[19]Individual income is clearly identified in Australian tax returns as the individual is the primary unit of taxation. While family-based income tests exist for some taxes and many family benefits, income is reported on an individual basis.

an individual has not filed a tax return, it is the sum of individual salary and wages reported by employers through annual pay-as-you-go (PAYG) payment summaries, and taxable government allowances, benefits and pensions. This income measure includes labor and capital income, and taxable government payments such as unemployment and study benefits. It is prior to any tax deductions or offsets. Income variables are measured in 2015 dollars, adjusted for inflation using the headline consumer price index published by the Australian Bureau of Statistics (2017c).

Parent income

Parent family income is the sum of the income of the identified parent(s) in each year, averaged from 1991 to 2001 inclusive. This window is chosen to strike a balance between potential attenuation biases arising from measuring parental income over too short a period versus too late in life when income is typically more volatile (Mazumder (2005)).

Child income

Child adult family income is the income of the child and their most recently reported spouse. I follow Chetty and Hendren (2018a) in measuring child income at age 24 or, more precisely, in the financial year in which they turn 24. The intent here is not to proxy for a child's lifetime income, but rather to gauge whether, how and why place might have a causal effect on later outcomes. As such, measuring outcomes at this age is not as obviously problematic as it would be in more typical intergenerational mobility studies, where life-cycle bias has received significant attention (for a recent examination, see Nybom and Stuhler (2017)). Nonetheless, key results are also examined with outcomes measured at later ages for a subset of individuals.

Location

In constructing the de-identified file, parents were assigned to the residential postcode listed on their tax return in each year.[20] Parental postcode was mapped

[20]For non-filing parents the postcode was recorded as missing. Postcodes were interpolated across periods of missing observations — if a parent lists the same postcode either side of such

to higher-level geographies using the 2011 Australian Statistical Geography Standard released by the Australian Bureau of Statistics (2011). In particular, the ABS postcode correspondence was used to map postcodes to Statistical Areas; a hierarchy broadly designed to range from local labor markets (SA4) to local communities (SA2).[21] Move distances were calculated based on the longitude and latitude centroids from the ABS postal areas corresponding to the origin and destination postcodes.

3.3.3 Sample comparisons and summary statistics

Table 3.2 presents key features of the sample compared against Chetty and Hendren (2018a). The key advantages of the Australian data are the longer panel of tax returns and longer span of birth cohorts linked to parents. This allows outcomes to be observed for children moving from age two (rather than from age nine). Despite the longer panel, sample sizes are significantly smaller, reflecting Australia's population. The samples of permanent residents and 1-time movers are 1,683,800 and 313,900.

A challenge to identifying causal effects of place in Australia is that these effects are likely smaller than in the United States. Australia is a relatively centralized federation, with the federal government controlling the individual tax and transfer systems and exerting significant influence in health and education policy. Potentially reflecting this, there is less variation in outcomes across Australian regions. Panel C in Table 3.2 shows the distribution across regions of the expected household income rank of a child born into the 25^{th} percentile of the income distribution. This is the 'absolute mobility' measure reported by Chetty et al. (2014). There is less regional dispersion in these expected ranks in Australia — the difference between the regions at the 90^{th} versus the 10^{th} percentile is only 7.8 rank points in Australia versus 14.7 in the United States.

Finally, the geographic units used in this chapter are smaller on average, and

a gap they were assumed to have been in the one location the whole time.

[21]The ABS postcode correspondence assigns postcodes to areas, roughly in proportion to population. A postcode may sit entirely within an area or be split across multiple areas. Postcodes are assigned to the area containing the largest split. Around 80% of postcodes sit within the one SA4. While some individuals will be misallocated, this misallocation is consistent by postcode and hence simply a caveat on the definition of the geographical units under consideration.

much more uniform in size (Panel D). This reflects the differing treatment of major cities. In the United States, major cities such as New York and Chicago are represented by a single commuting zone, whereas here they are represented by several SA4. In the Australian context of a smaller and more concentrated population this is an advantage, as within-city variations can be exploited. However, smaller geographic units also limit my ability to conduct some validation exercises.

In Appendix Table 3.8 summary statistics are presented for the permanent resident and 1-time mover samples. Moving children tend to be from slightly higher income families and are more often from couple families; they also tend to end up with higher incomes themselves. However, these differences are small and mask substantial variation.[22] In Appendix Table 3.9 summary statistics are presented for the origin and destination of the 1-time movers — there is a slight tendency to move to places with lower parent income ranks, fewer permanent residents and worse predicted outcomes for children.

3.4 Empirical framework

3.4.1 Estimation of causal place effects

I use the identification strategy introduced in Chetty and Hendren (2018a). This exploits variation in the *age* of children when their primary parent moves to identify the causal effect of *exposure* to place.

In the first step, I use the sample of permanent residents to estimate the relationship between parent and child household income ranks (p_i and y_i) in each SA4. The following parsimonious linear specification is used:

$$y_i = \alpha_{ls} + \beta_{ls} p_i + \varepsilon_i \tag{3.1}$$

where l and s denote the child's SA4 location and financial year of birth cohort respectively. Rank-based measures are less sensitive to attenuation and life-cycle bias — a point made in a number of studies, but particularly clear in the lifetime income histories available in Swedish data (Nybom and Stuhler (2017)). In

[22] In the United States, the moving families have higher mean but lower median incomes, and their children end up with lower mean and median incomes. See Table 1 of Chetty and Hendren (2018a).

Table 3.2: Comparison of data with Chetty and Hendren (2018a)

	Deutscher (2018)	Chetty and Hendren (2018a)
Panel A: Sample size and time span		
Birth cohorts	1978-1991 (14 years)	1980-1988 (9 years)
Permanent residents	1,683,800	19,499,662
1-time movers	313,900	1,553,021
Panel B: Location and income information		
Data range	1991-2015 (25 years)	1996-2012 (17 years)
Potential range of age at move	1-39 years	9-32 years
Analysis range of age at move	2-34 years	9-30 years
Panel C: Regional distribution of expected rank of child born into bottom half		
Mean region	47.1	43.9
p10 region	44.0	37.3
Median region	46.6	43.3
p90 region	51.8	52.0
Panel D: Regional distribution of average annual cohort size		
Mean region	2,200	4,596
p10 region	900	251
Median region	1,700	1,512
p90 region	3,300	10,358

Notes: Panels A and B are based on this chapter and Chetty and Hendren (2018a). Range of age at move assumes outcomes are measured at age 24. The distributions of child rank outcomes and children are calculated from the author's calculations and Chetty et al. (2014) (online data 5) for the 1978-82 and 1980-82 birth cohorts respectively.

many countries, the linear specification is also a reasonable approximation to the relationship between parent and child household income ranks (Bratberg et al. (2017)). To illustrate the suitability of the specification, Appendix Figure 3.6 plots the mean child household income rank at age 24 by parent household income decile for the 18 largest SA4 birth cohorts for 1978: both the general linearity of the relationship and relatively small deviations from the national relationship are striking. The fitted model from equation 3.1 then generates two predicted ranks for each child in the sample of 1-time movers — denoted $\bar{y}_{ops}$ and $\bar{y}_{dps}$ — based on two counterfactuals, namely that they were permanent residents of their origin and destination respectively.

In the second step, I use the sample of 1-time movers to estimate plausibly causal place effects. Specifically, I estimate a child's eventual household income rank as a function of their predicted origin outcome ($\bar{y}_{pos}$) and the 'shock' to this coming from their destination ($\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$), interacted with their age at move m. In the most general model, I estimate:

$$y_i = \alpha_{qosm} + \sum_{s=1978}^{1990} I(s_i = s)(\kappa_s \Delta_{odps}) + \sum_{m=2}^{34} I(m_i = m)(b_m \Delta_{odps}) + \varepsilon_i \tag{3.2}$$

where α_{qosm} is a set of fixed effects for the child's parent income decile q, origin o, cohort s and age at move m. The coefficients b_m capture the expected increase in rank associated with moving at age m to a destination with a 1 percentile rank higher predicted outcome than your origin for permanent residents. The coefficients κ_s allow for measurement error arising from potential mis-measurement of the child's origin — some children may have been born somewhere other than their parent's first recorded location, and this is more of a concern with the earlier cohorts.[23] Similar estimates of b_m are obtained from a more parsimonious and less computationally burdensome model:

$$\begin{aligned} y_i = &\sum_{s=1978}^{1991} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}) + \sum_{m=2}^{34} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p) \\ &+ \sum_{s=1978}^{1990} I(s_i = s)(\kappa_s \Delta_{odps}) + \sum_{m=2}^{34} I(m_i = m)(b_m \Delta_{odps}) + \varepsilon_i \end{aligned} \tag{3.3}$$

where the first sum captures origin effects, the second disruption effects, the third cohort controls and the fourth the coefficients of interest.

The b_m coefficients capture both the causal effect of moving at age m (β_m) and a selection effect (δ_m). The selection effect captures the idea that parents who move may have other attributes that affect child outcomes, for better or worse, such as differing levels of wealth, job or relationship security. Indeed, as shown in Table 3.8 movers are more likely to come from high-income and couple families. To identify causal effects, I follow Chetty and Hendren (2018a) in making the additional strong assumption that selection effects do not vary with a child's age at move ($\delta_m = \delta$ for all m). This will be discussed in further detail later.

[23] Since our data begins in 1991, it is only for the most recent cohort that we have location over the child's full childhood.

However, with this assumption, the causal effect of place can be identified. First, the selection effect δ can be identified from b_m where m is greater than the age at which the outcome is measured. Second, the causal effect β_m of moving at age m can be identified by subtracting the selection effect δ from b_m. The causal effect of an additional year of exposure at age m can be identified as $\gamma_m := b_m - b_{m-1}$. If these exposure effects are a constant $\gamma_{\tilde{m}}$ over age groups $\tilde{m} \in M$ then we can further parametrize equation 3.3 as:

$$\begin{aligned} y_i = & \sum_{s=1978}^{1991} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}) + \sum_{m=2}^{34} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) \\ & + \sum_{s=1978}^{1990} I(s_i = s)(\kappa_s \Delta_{odps}) + \sum_{\tilde{m} \in M} (\delta_{\tilde{m}} + \gamma_{\tilde{m}} e_{\tilde{m}}) \Delta_{odps} + \varepsilon_i \end{aligned} \tag{3.4}$$

where $e_{\tilde{m}}$ is the number of years in age group $\tilde{m}$ that the child was exposed to the destination for. I will mostly use the age groups $\{2, ..., 11\}$, $\{12, ...24\}$ and $\{25, ..., 34\}$.[24]

Finally, the above models fall into the general class of two-step estimation, where some regressors in the model of interest (one of equations 3.2-3.4) are generated from an earlier model (equation 3.1). A resulting econometric concern is that not accounting for the uncertainty in the generated regressors can lead to inconsistent standard errors (Pagan (1984)). I address this in two ways. First, I use the estimated standard error on the Δ_{odps} terms to restrict the sample to moves with more precisely estimated regressors. Specifically, I consider moves where the standard error on Δ_{odps} is less than 2 percentile rank points. This is a more direct means of ensuring precision than the population restriction used in Chetty and Hendren (2018a), where the baseline analysis considers only commuting zones with populations of 250,000 or more in 2000. Second, Murphy and Topel (1985) provide a formula for calculating asymptotically correct standard errors in fairly general settings, and these Murphy-Topel standard errors are presented where possible throughout this chapter.[25] Appendix 3.D describes in more detail the calculation

[24] For the post-outcome age group, we should have $\delta_{\tilde{m}} = \delta$ identifying the selection effect. I set $\delta_{\tilde{m}} = \delta_{\tilde{m}'}$ for the pre-outcome age groups, allowing only the exposure effect $\gamma_{\tilde{m}}$ to vary by age.

[25] Bootstrapping standard errors was another option but was not computationally feasible. Murphy-Topel standard errors cannot be calculated where the model of interest has regressors from multiple earlier models. In this case naive standard errors are report — this is mentioned

of Murphy-Topel standard errors, and shows key findings to be robust to varying the precision-based sample restriction.

3.4.2 Estimation of causal peer effects

Where a child grows up may matter simply because it determines who a child grows up with. A child's peers may influence their accumulation of skills, behaviors and aspirations — all feeding into eventual outcomes. There are many models of peer effects, ranging from roles for an entire peer group's mean ability, to the effect of individual 'bad apples' or 'shining lights' (Sacerdote (2011)). The adults in a child's neighborhood may also matter. Again, many mechanisms could drive this. For example, adults may serve as role models or facilitate job market search. This fits into the broader class of human capital externalities, which has a long history in the mobility literature.[26] I begin with a broad conception of peer effects, that includes all these potential pathways.[27]

Compelling causal identification of peer effects is notoriously difficult (for a recent discussion, see Angrist (2014)). Common shocks to a peer group — such as particularly good teachers or local labor markets — will generate spurious correlation in outcomes and the appearance of peer effects where none may exist. A common way to address this is to focus on ex ante variation in peer characteristics, and this is the approach taken here. Indeed, Angrist and Pischke (2009) suggest this is the 'best shot' at a causal investigation of peer effects; an approach that uses 'some measure of peer quality that predates the outcome variable and is therefore unaffected by common shocks' (page 196).

I exploit variation in the mean parent income ranks of financial-year-of-birth cohorts within a postcode.[28] Mean parent income rank will almost certainly vary

in the table notes where applicable. The differences between naive and Murphy-Topel standard errors in the regressions in this chapter tend to be relatively small.

[26]For example, Borjas (1992) shows the outcomes of a generation to relate not only to their parent's outcomes but also to the average outcomes of their parents' ethnic group.

[27]Manski (2000) describes three reasons members of the same group may end up with similar outcomes: 1) *endogenous interactions* whereby individual behavior varies with group behavior; 2) *contextual interactions* whereby individual behavior varies with group background; and 3) *correlated effects* whereby individuals behave similarly simply due to similar individual or institutional characteristics. I identify the combined effect of (1) and (2), rather than (1) alone.

[28]Financial year of birth was chosen as the time span of interest as it was readily available.

in a systemic way *across* postcodes and be correlated with unobservable factors influencing child outcomes. Similarly, there may well be trends in mean parental income rank correlated with trends in unobservable factors. However, idiosyncratic differences in birth timing will also generate plausibly exogenous differences between peer groups within a postcode. This approach is similar to that taken by Black et al. (2013) where cross-cohort variations in the ninth grade peer groups of Norwegians is exploited to identify peer effects. In this case, I estimate:

$$y_i = \alpha + \beta p_i + \eta c_{zs(i)} + \zeta \bar{c}_{zs(i)} + \varepsilon_i \tag{3.5}$$

where $c_{zs(i)}$ is the mean parent income rank in postcode z, birth cohort s, and $\bar{c}_{zs(i)}$ is the 3-, 5- or 7-year moving average of the same. The moving average controls for factors that may influence child outcomes, and be correlated with mean parent income rank within the postcode. I use leave-one-out means to exclude variation in peer group parent rank driven by an individual's own parents. In keeping with the earlier investigation of the causal effects of place, y_i is child household income rank at age 24 and p_i is their parent household income rank. Equation 3.5 is simply the canonical intergenerational mobility regression for the rank-rank (Spearman) correlation β, with two added terms — the peer mean parent income rank, with associated peer effect η, and the moving average of the same as a control.[29] As apparent in Table 3.3, there is significant variation in peer mean parent income ranks, both across postcode-year peer groups and across years within postcodes, the latter being the source of variation exploited here.

A potential criticism of the above is that families may be more (less) likely to leave a postcode if their child ends up in a particularly poor (rich) cohort. This selection process could result in unobserved differences between the families observed in the poor or rich peer groups. However, moving is costly and disruptive, so this may be less of an issue than endogenous school choice is in the more typical setting when peer groups are defined by school grades or classes. Nonetheless, I reestimate Equation 3.5 with family fixed effects to test the robustness of the

However, school entry cut offs also mean it is a good conceptual choice, as most Australian children born in the same financial year of birth end up in the same grade in school. The same is not true of calendar years.

[29]I also estimated specifications that combined these two terms, adding only the deviation in peer mean parent income rank $(c_{zs(i)} - \bar{c}_{zs(i)})$ to the regression, with similar results.

Table 3.3: Distribution of peer group size and mean parent income rank

	Mean	p10	p50	p90	N
Number of peers	134	22	99	288	1,339,300
Peer mean parent rank	52	38	51	66	1,334,500
Peer mean parent rank...					
minus 3-year ma	-.01	-3.2	.01	3.2	1,143,600
minus 5-year ma	-.02	-3.3	-.02	3.2	954,500
minus 7-year ma	.01	-3.4	.02	3.4	767,000

Notes: For the sample of permanent postcode residents, shows distributional statistics for the size and mean parent income rank of the peer groups defined by shared permanent postcode and financial year of birth. The moving averages are taken within the postcode, across the 3, 5 or 7 financial year of birth cohorts centred around the individual's financial year of birth.

results to fixed family unobservables.

3.5 The causal effect of exposure to place

I now turn to the results. In this section, I discuss the estimated causal effects of exposure to place in Australia, including the heightened importance of place in the teenage years. In the following section, I examine potential explanations for these patterns, including the role of local labor markets and peer effects.

I begin by examining Figure 3.1, which shows the estimated coefficients b_m from Equation 3.3. These represent the expected boost to an individual's household income rank at age 24 associated with a move at age m to a place with 1 percentile rank higher outcomes for permanent residents. Three patterns are of interest — the evidence of selection effects, positive exposure effects and the more pronounced sensitivity of the teenage years.

First, there is evidence of positive but age-invariant selection. Children do better at age 24 even if their parents only later move to a place with better outcomes. Comfortingly this selection effect appears invariant to the child's age at move: if I fit a line to the estimated regression coefficients b_m for $m \in \{25, ..., 34\}$ I get a slope coefficient of -0.00097 (s.e. 0.0085).[30] The mean selection effect is

[30]Selection effects are not considered for $m = 24$ as the child's outcome is potentially affected by moves at this age. Outcomes are measured in the income year the child turns 24. Age at move is determined by how old the child turns in the income year for which the primary parent

0.308, slightly higher than the 0.226 found for the United States in Chetty and Hendren (2018a). This is not surprising, as if Australia's lower regional variation in mobility compared to United States reflects less of a causal role of place this is exactly what one would expect.

Second, there is evidence of positive exposure effects — the benefits of a move to a place with better outcomes are greater the earlier one moves. Third, the exposure effects ($\gamma_m = b_m - b_{m+1}$) appear to vary systematically by age, with larger exposure effects in the teenage years. All up, a child whose parent moves at their birth would be expected to pick up about 70% of the difference in predicted outcomes between their origin and destination.[31]

In Appendix Table 3.10, I provide exposure effect estimates and model fit statistics for several versions of the parametrized model in equation 3.4. In column (1), I assume constant exposure effects as in Chetty and Hendren (2018a). In columns (2)-(8) I allow a kink, with exposure effects that are allowed to vary between early and late childhood. Consistent with the visual impression left by Figure 3.1, the best fitting parametrized model allows a kink at age 11.[32] An additional year in a place with 1 percentile rank point higher outcomes is associated with a gain at age 24 of 0.011 (s.e. 0.007) rank points before age 11 and 0.042 (s.e. 0.003) rank points after age 11; the p-value on a test of equivalence of the early and late childhood exposure effects is 0.00068. This model is the baseline model for the analysis that follows.

The ability to measure early childhood exposure effects is the primary advantage of the Australian data, and the finding of smaller place effects in these years is a key contribution of this chapter. It also seems unlikely this finding will be

files a return from a new location. While the parent will file the return after that income year, the move may well have occurred during it and affected child outcomes.

[31]The fact that most of the differences in outcomes between places in Australia appears to be causal does not imply that place explains a large portion of the variation in outcomes between individuals. As already noted, the geographic variation in outcomes in Australia is relatively modest, and it may be that Danish findings that neighborhoods account for a large share of the variance in permanent earnings between siblings early in the working life, but a negligible share beyond age 30 (Bingley et al. (2016)) would also replicate here, though this is beyond the scope of this chapter.

[32]This model is preferred to the other potential kink points on all the model selection criteria. It is also preferred to a model with constant exposure effects on all but the BIC, which places the highest penalty on the additional parameter.

Figure 3.1: Place exposure effect estimates for child income rank in adulthood

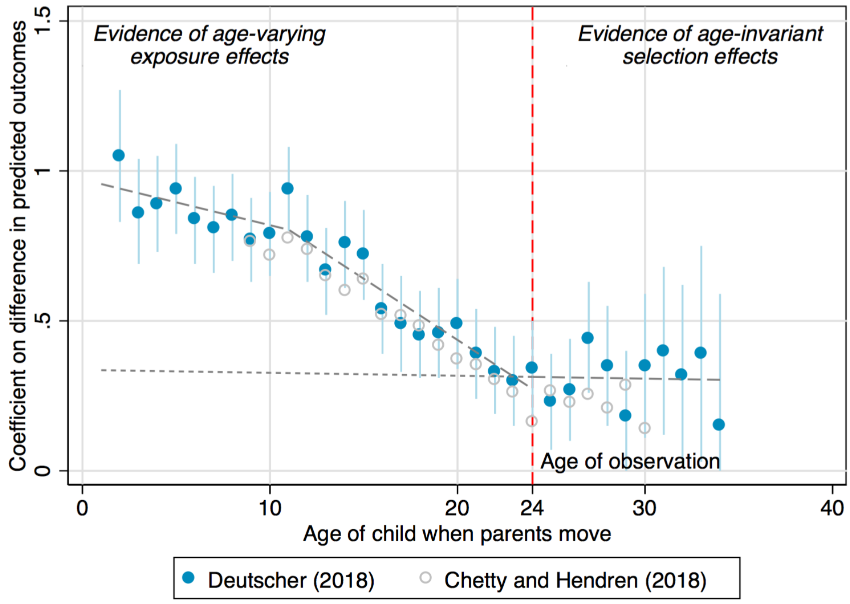

Notes: Estimated coefficients b_m from equation 3.3, with 95% confidence intervals. The b_m capture the expected boost to an individual's household income rank at age 24 from moving at age m to a place with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their age at parent move m with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. This replicates Panel B, Figure IV from Chetty and Hendren (2018a). The point estimates from that paper are also shown.

unique to Australia. As noted earlier, the age-at-migration literature has generally found additional years of exposure to the host country matter most in the teenage years for skills such as language acquisition. Further, there are hints in Chetty and Hendren (2018a) that a similar pattern may be observed in the United States once more data becomes available. For example, as the authors themselves have noted, the teenage years seem more important than earlier years in determining teen births and college attendance.[33]

[33] Also, in Figures IV and V in Chetty and Hendren (2018a) the coefficients b_m for $m \in \{9, ..., 12\}$ do not generally display as strong a linear relationship as the later years. These coefficients are also likely not as precisely estimated given they are based on only the later cohorts for which moves at this age can be observed.

The sensitivity of the teenage years is not driven by any particular subpopulation. In Appendix Table 3.11, exposure effects are estimated separately for various subpopulations — across all specifications the exposure effects are higher in late childhood. There is, however, some notable heterogeneity. Late childhood exposure effects are significantly larger for boys than for girls — 0.049 (s.e. 0.005) versus 0.032 (s.e. 0.005). Moves by individuals from poorer families or to better destinations are nearer to having constant exposure effects, though the teenage years still appear more sensitive.

Finally, the sensitivity of the teenage years is not confined to place exposure effects — the fixed costs associated with moving also appear larger in these years.[34] Appendix Figure 3.7 shows the expected household income rank at age 24 of a child born in 1991, with parent income rank of 50, and moving at age m between an origin and destination where their predicted outcome is also an income rank of 50 (implying $\Delta_{odps} = 0$). Positive selection effects are apparent — children exceed their predicted ranks at age 24 even if their parents only move later — but comfortingly these also appear age-invariant.[35] If these selection effects can also be extrapolated back then the chart suggests a fixed cost of moving between places. These fixed costs are also larger during the teenage years: they are largest at age 17 — the age at which students enter their final year of schooling and prepare for university entrance exams. The fixed costs of moving are relatively large, and may well outweigh the benefits. If we subtract the assumed selection effects, then in the sample of moves in the data, only 15.6% of moves (23.9% of moves to places with better outcomes) carry an expected benefit to the child based on their age at move m, parent rank p and the difference in outcomes of permanent residents Δ_{odps}.[36]

[34]'Fixed' in the sense that they do not vary with the difference in predicted outcomes between the origin and destination (Δ_{odps}).

[35]This positive selection into moving is consistent with Appendix Table 3.8, where movers were found to come from slightly higher income families, on average.

[36]This discussion is intended as illustrative only, as I don't present arguments in favor of a causal interpretation of these fixed costs. It is simply a useful caveat to bear in mind — the regressions here do not imply that moving to a better place is invariably associated with better expected outcomes. Further, fixed costs of moving are not without precedent. In an altogether different setting Chetty et al. (2016) find slightly negative effects for those moving in adolescence as part of the Moving to Opportunity experiment (but positive effects for those moving when young), which they suggest may reflect disruption effects.

3.5.1 Validation exercises

A caveat on these results is the strong underlying assumption that selection effects are age-invariant. There appear to be good grounds to doubt this assumption — both intuitively and based on the data. Appendix Figure 3.8 shows moves in late childhood are much less common and slightly skewed towards lower income and sole parent families. While parent income rank is controlled for in the regressions, differences in other background traits may lead to biased estimates of the causal effect of place.

In Appendix 3.C I replicate where possible and discuss in detail the validation exercises conducted in Chetty and Hendren (2018a). The results are generally comforting. Using a more general set of controls or a later age of observation does not alter the conclusions above. Adding family fixed effects results in only a very modest reduction in the estimated exposure effects, and the sensitivity of the teenage years remains apparent. Finally, an event study illustrates the best predictor of a mover's outcomes comes from looking at their cohort, rather than neighboring cohorts.

There are two cases where the validation exercises fail. In particular, the exposure effects for moves in the top decile of relative postcode outflows are significantly attenuated. The idea behind this exercise is that these moves seem more likely to have been forced — due to mass layoffs or natural disasters — rather than chosen. This is a potential caveat on the external validity of the results, but would also be consistent with the finding of Chyn (2018) that treatment effects may fundamentally differ between those choosing to move versus those forced to move. Finally, the mover's probabilities of falling into the top or bottom deciles are better predicted by the expected ranks of the permanent residents (rather than expected probabilities of the same outcomes), but this could readily be due to much smaller sample sizes, as discussed further in the Appendix.

In sum, the tests provide good evidence in support of a causal interpretation of the estimates in the Australian setting. Any omitted variable giving rise to the exposure effects must operate within the family in proportion to the time exposed to the destination, and replicate the outcomes of an individual's particular birth cohort. However, these tests are also consistent with a number of potential explanations for a causal effect of place — including potential roles for local market conditions and peer effects as considered in the following section.

3.6 Why the teenage years?

3.6.1 The role of local labor market conditions

Where a child grows up might matter simply because it relates to where they end up working. If a child ends up in a stronger local labor market they will likely do better even without their location having had any effect on their development.

If local labor markets drive the observed effects, then why would time spent in a location matter? Why would we see an exposure effect pattern? Figure 3.2 provides an explanation — a child's end location looks like a function of their exposure (through their primary parent) to locations growing up. For those moving once, the earlier a child's parent moves, the more likely the child is to end up in the destination by age 24. This is particularly the case for moves after 15, and there is no age at which a parent's move creates a large, obvious discontinuity in the child's probability of ending up in the destination location. A similar pattern is also apparent if controlling for family fixed effects (Appendix Figure 3.9). And for those moving twice, the more years spent in the initial destination the more likely the child is to end up there (Appendix Figure 3.10). There are intuitive explanations for a child's location being a function of their exposure to locations while growing up. A longer period of exposure in childhood might generate stronger ties to the location through social ties and preference formation. It may also indicate stronger exogenous ties to the location, such as the presence of other family or community members.[37]

[37]Mis-measurement of child location is also a possibility given older children are less likely to follow their parents. As noted in Chetty and Hendren (2018a) this possibility of children not moving with their parents means the exposure effects estimated on the basis of parent moves should be interpreted as an intention-to-treat effect rather than an effect of treatment on the treated. I do not attempt to calculate the latter, as the Australian tax data does not reveal whether children moved with their parents. However, Appendix Figure 3.11 shows that in a similarly constructed 1-time mover sample from the Household Income and Labour Dynamics in Australia (HILDA) survey all those aged 16 or less when their parent moved followed their parent. Even among these individuals, the child's end location is a function of their exposure to the destination. This suggests mis-measurement cannot be the primary explanation of the

Figure 3.2: Moving once: % of children in a given location at age 24, by age at parent move

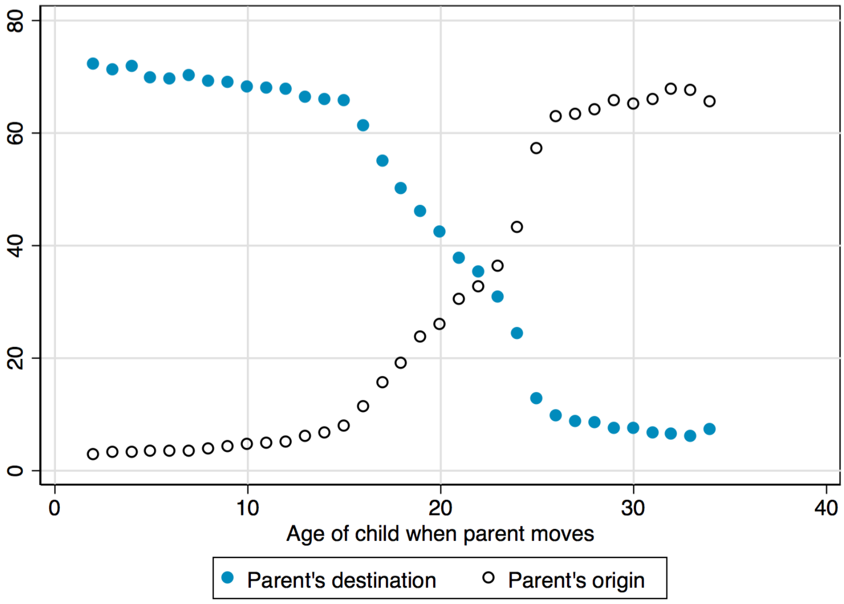

Notes: Based on the 1-time movers sample. For children of a parent who moves once, shows the % those in a given location at age 24 — the parent's origin or destination — by age at move. Those in a different location are not shown for clarity, while those without a known location are excluded. The child's location is known for all those lodging a tax return in the year they turn 24 with a valid postcode.

These charts refute some arguments against a role for local labor market conditions. Figure 3.2 shows that the pattern of exposure effects observed in Figure 3.1, with no obvious discontinuity in place effects just prior to labor market entry, could in fact still occur even if local labor markets explained all the observed place effects. Similarly, even if exposure to an interim destination matters for outcomes (as found by Chetty and Hendren (2018a) for those whose parents move multiple times), then this too is consistent with effects entirely driven by local labor markets.

In Table 3.4, I present estimated exposure effects with a variety of different controls for a child's end location. These are 'bad' control variables, in that they control for an endogenous outcome — an individual's location may indeed influence their income, but income may also influence location. Nonetheless, this exercise is

pattern observed in Figure 3.2.

useful as an initial investigation. In column (2) I add controls for a child's location at age 24. This partially attenuates the estimated exposure coefficient, which falls by around 40% from 0.042 to 0.026.[38] Since we would expect some attenuation anyway, given we are controlling for an endogenous outcome, it is tempting to read this as a lower bound for the exposure effect net of local labor market effects. However, local labor markets have cycles, and will generally differ in their strength over time. A more comprehensive set of local labor market controls would have the same richness as the place effects, with differences by cohort and parent rank.

Allowing local labor market controls to vary with time is important. In column (3) I interact child location with child cohort and the exposure effect estimate is now more than halved relative to the baseline, falling to 0.020. In column (4) I interact this with parent rank as well, but with negligible further effect. Table 3.4 suggests local labor market conditions could play a significant role in generating exposure effects. However, the question still remains as to how much this attenuation is simply a side-effect of controlling for an endogenous outcome.

An alternative and more compelling test of the importance of local labor market conditions is a placebo test along the lines of those used in validating the research design. If local labor markets explain the effects observed then the outcomes of children who end up residing in a location should be a better predictor of movers' outcomes than the outcomes of those who grew up there. To test this, I consider the subset of the permanent resident children who file a tax return from a valid postcode. This subsample can then be used to generate two sets of predictions for the movers — one based on where the child "grew up" (their primary parent's location) and one based on where they "ended up" (their filing location).

The choice of specification is a key consideration in conducting this exercise. I begin with the more parsimonious linear kink specification in equation (3.4), before moving to less restrictive specifications. Earlier, I chose age 11 as the kink point separating early and late childhood. However, this was chosen to maximize

[38]In Chetty and Hendren (2018a) the estimated exposure effect falls by only around 20%, from 0.040 to 0.031 (Table II). This is actually based on a specification that controls for child location in 2012, rather than at age 24, interacted with cohort. This could lead to either lower or greater attenuation relative to what would be observed if the specifications in column (2) and (3) were replicated in the US data. In practice however, the resulting coefficients are at most slightly smaller than the 0.031, and remain above those observed in the Australian data (personal communication with Nathan Hendren).

Table 3.4: Exposure effects with local labor market controls

	Baseline	With child location controls		
	(1)	(2)	(3)	(4)
Early exposure	0.011	0.011	0.009	0.010
	(0.007)	(0.006)	(0.006)	(0.006)
Late exposure	0.042	0.026	0.020	0.020
	(0.003)	(0.003)	(0.003)	(0.003)
Post-outcome exposure	0.008	0.015	0.009	0.011
	(0.013)	(0.012)	(0.012)	(0.012)
Controls				
child location ...		X	X	X
... X child cohort ...			X	X
... X parental rank				X
N	264,500	264,500	264,500	264,500

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure with various controls for child location. These represent the expected boost to an individual's household income rank at age 24 associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination in each stage of life $e_{\tilde{m}}$ with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. Child location controls cover: none (1); child location fixed effects (2); child location-cohort fixed effects (3); and child location-cohort fixed effects and interactions with parental rank (4). Child location is simply the corresponding SA4 for those filing a tax return with a valid postcode. Murphy-Topel standard errors are in parentheses.

the ability of those growing up in a place to explain the outcomes of those moving places. A kink at age 11 may thus tilt this test in favor of those growing up in a place over those ending up there. On the other hand, age 15 is a very clear kink point in the relationship between where you grow up and where you end up (Figure 3.2). Using this kink point may thus tilt the test in the other direction, maximizing the predictive content of the regressors based on those ending up in a place. I initially conduct this test using two specifications, one with a kink point at age 11 and one with a kink point at age 15, before further relaxing these constraints.

The results of this exercise suggest a clearer role for local labor markets. In

Table 3.5 I present the estimated late childhood exposure effects for a mover's rank and probability of making the top decile. Estimates are shown using, first separately and then simultaneously, the predictions based on either where the children of permanent residents grew up ("Growers") or where they ended up ("Enders"). The separate regressions produce exposure effect estimates that are similar to those from the baseline specification. When estimated simultaneously the results are highly sensitive to the specification. With a kink at age 11, the exposure effect based on the Growers falls, but only by a modest 15% when the outcome is mean income rank. If I instead locate the kink at age 15 then the predictions based on the Enders dominate, with the exposure effect based on the Growers falling by 55% when the outcome is mean income rank. These regressions leave open the possibility that where you grow up matters mostly because it influences where you end up, and that the endogenous controls in Table 3.4 are not, in fact, overstating the role of local labor markets.

Finally, I further relax the specification, returning to one that is more agnostic as to how exposure effects may vary over childhood. Specifically, I modify equation 3.3 to also include the predicted outcomes, and differences in predicted outcomes between origin and destination, based on the 'Enders'.[39] The resulting coefficients are shown in Figure 3.3. While not definitive, a few useful observations can be made. At almost all ages there is predictive power in the experiences of both those growing up and those ending up in a place. The predictive content of both $\Delta_{odps}^{Growers}$ and Δ_{odps}^{Enders} tends to fall with age, particularly in the teenage years. The precise fall in predictive content is quite sensitive to how it is measured. If we look simply at the difference in the average coefficients before age 11 and after age 24, then the coefficients on the Growers fall by around 2.7% a year, while those on the

[39]The precise specification is:

$$\begin{aligned} y_i = &\sum_{m=2}^{34} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p) + \sum_{s=1978}^{1990} I(s_i = s)(\kappa_s \Delta_{odps}^{Growers}) \\ &+ \sum_{s=1978}^{1991} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}^{Growers} + \alpha_s^3 \bar{y}_{pos}^{Enders}) \\ &+ \sum_{m=2}^{34} I(m_i = m)(b_m^{Growers} \Delta_{odps}^{Growers} + b_m^{Enders} \Delta_{odps}^{Enders}) + \varepsilon_i \end{aligned} \tag{3.6}$$

Table 3.5: Exposure effects using predicted outcomes of permanent versus eventual residents

	Mean rank			Top 10%		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Kink at age 11						
Growers	0.037		0.031	0.017		0.010
	(0.004)		(0.008)	(0.003)		(0.003)
Enders		0.029	0.005		0.017	0.012
		(0.003)	(0.007)		(0.002)	(0.003)
Panel B: Kink at age 15						
Growers	0.044		0.020	0.024		0.013
	(0.006)		(0.012)	(0.004)		(0.005)
Enders		0.039	0.024		0.023	0.018
		(0.005)	(0.010)		(0.004)	(0.004)
N	221,700	215,600	209,800	210,500	206,800	176,200

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ from equation 3.4 for late childhood for different outcomes and based on different predictions for origin and destination outcomes. The outcomes of those whose parents move once in their childhood provide the dependent variables. In columns (1)-(3) the outcome is the individual's household income rank at age 24, in columns (4)-(6) it is the event that this rank is in the top decile. As in Chetty and Hendren (2018a), I add the square of parent income rank to equation 3.1 before estimating the predicted probabilities of being in the top decile. When the outcome is being in the top decile, rather than mean rank, I use the precision threshold of $\Delta_{odps} < 10$ in order to restrict consideration to a similarly sized subsample of moves with the most precisely estimated differences between origin and destination. In all regressions the predicted outcomes for the origin and the destination, which provide a counterfactual for the moving individual, are based on a sample of those whose parents are permanent residents of a location, and who file a tax return in the year they turn 24 with a valid postcode (movers are also required to do this). In columns (1) and (4) the predicted outcomes for the origin and destination are based on the sample of individuals who grew up there, based on the location of their primary parent. In columns (2) and (5) the predicted outcomes for the origin and destination are based on the sample of individuals who file a tax return from there. In both these cases the specification is as in equation 3.4. In columns (3) and (6) both sets of generated regressors — those based on where individuals grew up and those based on where they ended up — are included in the regression. Standard errors are in parentheses (naive for columns (3) and (6) and Murphy-Topel otherwise).

Enders fall by around 1.4% a year. This would suggest where an individual ends up explains around a third of the total causal effect of place, a story between the two extremes emerging from Table 3.5.

These tests for a role for local labor market conditions go beyond those provided in Chetty and Hendren (2018a). In theory, these results could reflect specific features of the Australian economy over this period, such as the mining boom resulting from Chinese demand for Australian resources.[40] The mining boom is

[40]The mining boom saw employment in the industry increase well over three-fold from 2002

Figure 3.3: Place exposure effect estimates for child income rank in adulthood, based either on permanent residents who grew up or ended up in a location

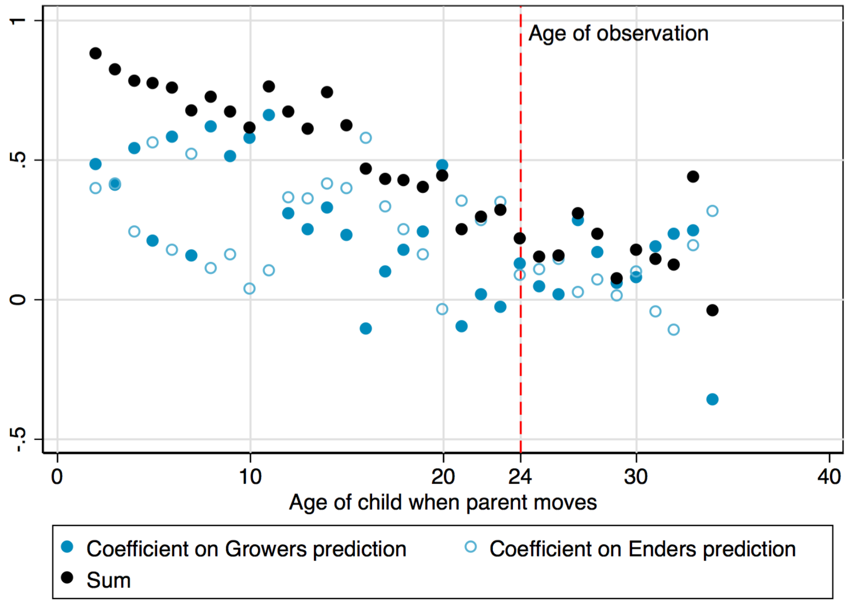

Notes: Estimated coefficients $b_m^{Growers}$ and b_m^{Enders} from equation 3.6, and their sum. The b_m capture the expected boost to an individual's household income rank at age 24 from moving at age m to a place with 1 percentile rank higher expected outcomes for permanent residents — either those who grew up in the locations or those who ended up there. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their age at parent move m with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o, calculated separately based on the permanent residents who grew up or ended up in the locations. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with $\Delta_{odps}^{Growers}$ to capture potential mis-measurement of the origin. Due to the high degree of correlation between $\Delta_{odps}^{Growers}$ and Δ_{odps}^{Enders} this last term is included based on the growers only. The results are not meaningfully different if this term is excluded all together.

certainly an important factor in recent experiences of intergenerational mobility in Australia: Appendix Figure 3.12 shows 'absolute mobility' in mining states has followed economic conditions relatively closely. However, in Appendix Table 3.12 I show this chapter's key findings remain when restricting to moves within the non-mining states. Even for those moving within the non-mining states, exposure effects are larger in the teenage years and significantly attenuated when including

to 2012 and was concentrated in two states — Western Australia and Queensland. Together these states accounted for around a third of Australian employment but two-thirds of mining employment over the period of observation (Australian Bureau of Statistics (2017d)).

local labor market controls. The same is true when restricting to moves within major metropolitan areas (Appendix Table 3.13).

The relative importance of local labor markets in Australia may simply reflect a lesser role played by other factors. As noted in Section 3.3.3, Australia has much less geographic variation in intergenerational mobility than the United States. This is perhaps unsurprising, as Australia is a much more centralized federation, with less geographic variation in policy settings. As a consequence, even if local labor markets are no more important in absolute terms than in the United States, they could account for a greater share of the causal effects of place observed here.

3.6.2 The role of peers

The results above suggest place effects and the sensitivity of the teenage years is partly — but not wholly — driven by local labor market conditions. A prime candidate for some of the residual effect is a role for peer effects. The psychological literature has long noted that the transition from early childhood to adolescence is marked by the increasing salience and complexity of peer relationships (Brown and Larson (2009)).

The Australian data shows clear evidence of peer effects. To identify peer effects I regress an individual's household income rank on their parents' rank and their peers' mean parent rank. Peers are those sharing a postcode and a financial year of birth. I include the moving average of peer mean parent ranks for that postcode as a control. Intuitively, if a child just happens to be born into a wealthier or poorer birth cohort for their postcode, do they do any better or worse? Children could be influenced either directly by their peers' parents — as potential role models, for example — or indirectly — insofar as they proxy for the abilities and behaviors of the peers themselves.

In Table 3.6 I display the coefficients on parent rank and peer parent rank. Across a range of specifications, an increase of 10 percentile ranks in the peer parent rank is estimated to increase a child's household income rank at age 24 by between 0.2 and 0.3 percentile rank points. The influence of peers is around a fifth that of the influence of parents at that age — a large yet plausible effect size. The addition of family fixed effects has little effect on the point estimates, though their statistical significance is lost. While the estimate is notably smaller with the shortest window width, this would be consistent with individuals having some peers drawn from the neighboring birth cohorts. In this case, having a higher mean peer parent rank would be expected to be more beneficial when it is conditional on the wider moving averages that place less weight on cohorts that may also exert some influence on outcomes.[41] Given this possibility, I use a window width of 7 years in the remaining analysis.

Table 3.6: Parent and peer influences on household income rank at age 24

	(1)	(2)	(3)	(4)	(5)	(6)
Parent rank	0.129	0.129	0.129			
	(0.002)	(0.002)	(0.002)			
Peer parent rank	0.017	0.028	0.025	0.019	0.028	0.030
	(0.008)	(0.009)	(0.011)	(0.020)	(0.023)	(0.027)
Specification						
Window width	3	5	7	3	5	7
Family fixed effects				X	X	X
N	1,126,200	939,700	754,900	1,126,200	939,700	754,900

Notes: coefficients from equation 3.5 — the regression of a child's household income rank at age 24 on: their parent household income rank; their peers mean parent rank; and the 3-, 5- or 7-year moving average of peer mean parent rank. Peers are defined by postcode and financial year of birth and exclude the individual in question. The moving average of the peer mean parent rank is taken across adjacent birth cohorts in each postcode. A peer's primary parent must have been a permanent resident of the postcode — not filing from outside it — from 1991 to the year in which the child turned 20. Robust standard errors, clustered by postcode, are in parentheses.

As a further test of the results, I conduct a placebo test, looking for any effect of the peers born in the years either side of an individual's own cohort.[42] In particular, I re-estimate the baseline specification from column (3) of Table 3.6 as if an individual's financial year of birth was $s + l$ rather than s, where $l \in \{-6, ..., 6\}$. Figure 3.4 presents the resulting peer effect estimates. The largest and only statistically significant estimate is that on an individual's own cohort. This suggests the results are not driven by inadequate controls for trends within postcodes that may generate spurious correlation between mean peer parental rank

[41] Black et al. (2013) also note this possibility, stating that they view estimates of theirs based on a 3-year moving average as a "lower bound on true peer effects."

[42] I thank Nathan Hendren for suggesting a test along these lines.

and individual outcomes.

Figure 3.4: Peer effect estimates: placebo test

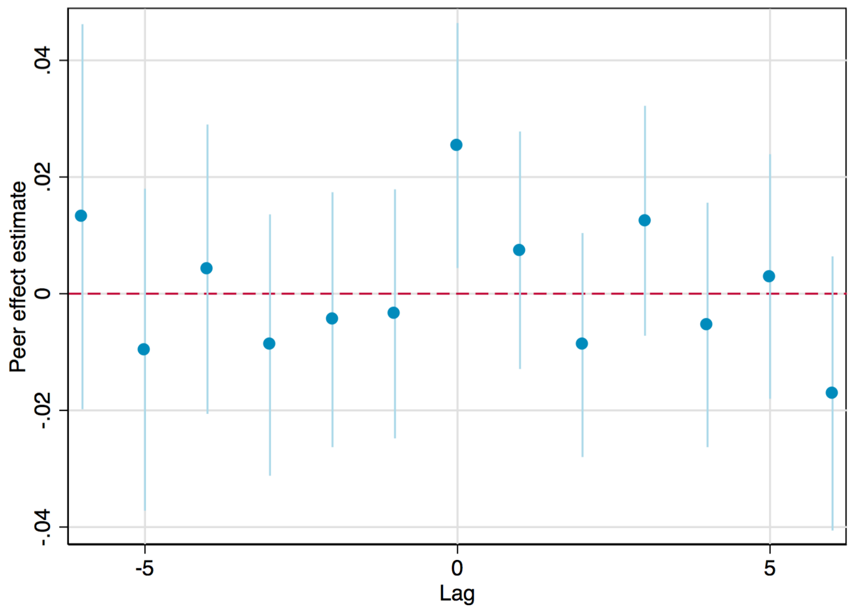

Notes: Based on permanent postcode residents. Shows the coefficients (and 95% confidence intervals) from a regression of household income rank at age 24 on own parent household income rank and the mean parent household income rank of peers (defined by shared permanent postcode and a financial year of birth that is shifted by a lag l relative to the individual's own). A 7-year moving average of the mean parent rank of peers is included as a control, in line with the specification in column (3) of Table 3.6.

Finally, the identified peer effects are not driven by particularly small or large postcodes. While those in smaller postcodes are perhaps more likely to interact with one another, they may also be a less complete picture of an individual's social network. In Appendix Figure 3.13 I present the estimated peer effects restricting attention to those in progressively smaller peer groups. The estimates are less precise, but range between 0.017 and 0.026 and all include the baseline estimates from Table 3.6 within their confidence intervals.

Peer effects do not appear to fade with time, and rather tend to increase with age. Figure 3.5 shows the coefficients on parent and peer parent rank lifting the age at which child outcomes are observed from 16 to 30, using the specification in column (3) of Table 3.6.[43] In this case, peer effects measured at age 30 are slightly

[43]Appendix Table 3.14 replicates Table 3.6 in full with outcomes measured at age 30.

larger than those measured at age 24: though they remain around a fifth of the parent effect. This suggests the peer effects are lasting; they may even be subject to a modest version of the life-cycle bias effects apparent in the estimated coefficients on parent rank and well known in the intergenerational mobility literature (Black and Devereux (2011)).

Figure 3.5: Influence of own and peer parents on household income rank at various ages

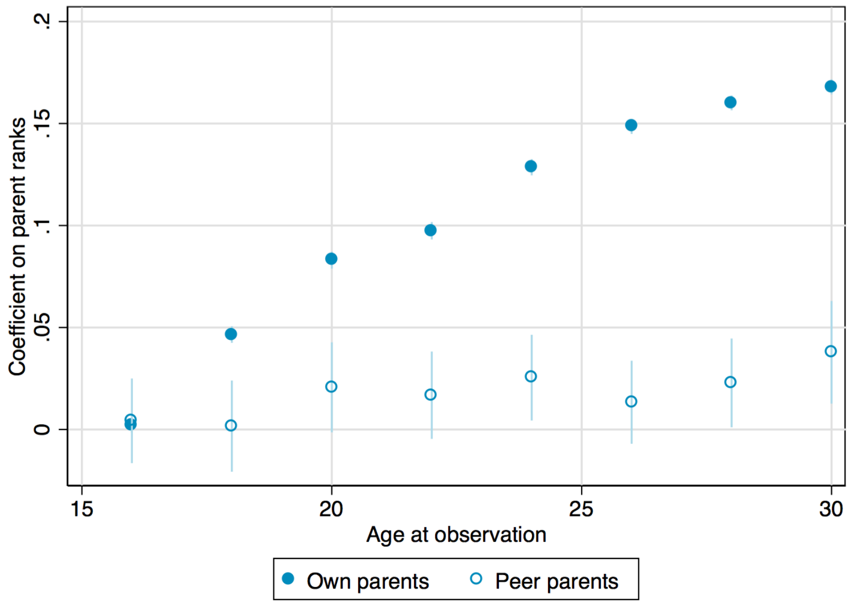

Notes: Based on permanent postcode residents. Shows the coefficients (and 95% confidence intervals) from a regression of household income rank at various ages on own parent household income rank and the mean parent household income rank of peers (defined by shared permanent postcode and financial year of birth). A 7-year moving average of the mean parent rank of peers is included as a control, in line with the specification in column (3) of Table 3.6.

Why might peers matter?

A strength and a weakness of the peer effects identified here is that they may reflect many underlying mechanisms. These include the influences of the peers themselves, and potentially their parents, arising from social interactions. But they may also include nonsocial externalities — for example, an individual may benefit indirectly if their peers' parents lobby for better teachers, more resources or more opportunities for their peer group cohort.

To better understand the likely mechanisms at play, I reestimate equation 3.5 separately for men and women, and with separate controls for female and male peers. Table 3.7 presents the estimated peer effects. For both women and men, it is the rank of their same-sex peers' parents that matters most. The results rule out some otherwise plausible explanations for the peer effects. A simple "boy/girl next door" explanation, whereby peers matter because some individuals marry their peers, would require the opposite-sex peers to matter most. Further, it appears less likely that nonsocial externalities drive the results, as they would then have to operate in a gender specific manner only — for example, richer parents only lobbying for teachers, resources or opportunities that specifically benefit their child's gender.

Finally, data limitations make it difficult to provide a compelling test of whether peers are more important at particular points in life, such as the teenage years. As a suggestive test, I examined how estimates varied if the definition of the peer group was expanded to include those moving into a postcode-cohort area before the cohort hit the teenage years. However, insufficient variation between these competing peer group measures meant it was not possible to draw firm conclusions. An added difficulty is the absence of direct data on peer relationships — it could be that peer relationships are formed in early childhood but only influence outcomes from adolescence onwards.

It is also difficult, given the data limitations, to nest peer effects in the model of exposure-to-place effects in a compelling way. For example, it is straightforward to include controls for the peer shock that movers experience.[44] However, these fail to explain the patterns observed — the late childhood exposure-to-place effect remains steady at 0.04, while the exposure-to-peers effects are small and statistically insignificant. It could be that peers do not, in fact, explain any of the causal effects of place examined earlier in the chapter. But another possibil-

[44]Specifically, I construct controls for the mean parent income rank of the permanent residents of the origin ($\bar{p}_{os}$) and the difference between this and the same for the destination ($\Delta_{ods} := \bar{p}_{ds} - \bar{p}_{os}$). I then re-estimate equation 3.4, including controls for $\bar{p}_{os}$ and Δ_{ods} that are identical to those for $\bar{y}_{pos}$ and Δ_{odps}. This allows for an exposure-to-peers effect. Whether these measures are based on the origin and destination SA4, or on the origin and destination postcodes, does not materially change the results.

Table 3.7: Parent and peer influences on household income rank at age 24 — by individual and peer sex

	(1)	(2)	(3)	(4)
Panel A: Men				
Parent rank	0.133	0.134	0.133	0.134
	(0.002)	(0.002)	(0.002)	(0.002)
Peer parent rank for...				
...all peers	0.028			
	(0.014)			
...male peers		0.019		0.024
		(0.010)		(0.011)
...female peers			0.001	0.001
			(0.010)	(0.011)
N	388,100	383,900	381,900	380,300
Panel B: Women				
Parent rank	0.123	0.123	0.123	0.123
	(0.002)	(0.002)	(0.002)	(0.002)
Peer parent rank for...				
...all peers	0.023			
	(0.015)			
...male peers		0.005		0.007
		(0.011)		(0.012)
...female peers			0.019	0.017
			(0.011)	(0.011)
N	366,800	361,700	362,200	359,600

Notes: coefficients from equation 3.5 — the regression of a child's household income rank at age 24 on: their parent household income rank; their peers mean parent rank; and the 3-, 5- or 7-year moving average of peer mean parent rank — separately by individual and peer sex. Peers are defined by sex, postcode and financial year of birth and exclude the individual in question. The moving average of the peer mean parent rank is taken across adjacent birth cohorts in each postcode. A peer's primary parent must have been a permanent resident of the postcode — not filing from outside it — from 1991 to the year in which the child turned 20. Robust standard errors, clustered by postcode, are in parentheses.

ity is that these simple controls fail to capture the true peers of movers, or their influence on movers' outcomes, any better than the predicted outcomes of the permanent residents do.[45] In the absence of direct measures of peer relationships,

[45]The predicted outcomes of the permanent residents in theory include the influence of the full set of peers that lived in that location, even if only temporarily. The controls for peers included here are based only on the permanent residents. A time-varying measure of peers was explored but is made difficult by the gaps in the location histories inherent in the tax data.

their influence on movers seems likely to remain embedded in the black box of the exposure-to-place effect observed here.

Further research on when peers influence outcomes and if they can explain the causal effect of place would be valuable. Also of interest is whether the effects identified here (for permanent residents) operate through the peers themselves or their parents. For example, in a recent contribution Bell et al. (2017) find striking relationships between an individual's patent rate in a specific technology class and the patent rates of not only their father, but their father's colleagues and their neighbors — such findings seem suggestive of a place for role model effects.

3.7 Conclusion

The seminal work of Chetty and Hendren (2018a) provided fresh evidence that where a child grows up matters for their later life outcomes. In this chapter I explore when and why place matters — critical questions for those seeking to redress the inequalities arising from causal place effects. I find place matters most in the teenage years, with a clear role for local labor markets, and evidence for peer effects at a more localised scale.

Exposure-to-place effects are largest in the teenage years, and generally small and not statistically significant in early childhood. This is consistent with age-at-migration studies finding the benefits to language acquisition from migrating a year earlier are generally largest in adolescence as well. It also accords with the findings of Chetty and Hendren (2018a) for earlier life outcomes, such as teenage births and college attendance. As such, it seems unlikely this finding is particular to Australia. This finding does not suggest that early childhood is unimportant. Rather, what matters in early childhood may be factors — such as the family or more localized environmental influences — where most variation is within rather than between the large neighborhoods examined here.

Any explanations for causal place effects should also seek to explain the sensitivity of the teenage years. Two possibilities are explored in this chapter — that where a child grows up matters because it influences where they end up (local labor market conditions) and who they grow up with (peer effects).

Local labor market conditions can partly explain the observed patterns, based on tests with endogenous and exogenous controls. When using exogenous controls based on those ending up in a place (rather than those growing up there), anywhere from 15-55% of the effect of exposure to place in late childhood is explained away. This is not driven by idiosyncratic factors, such as the mining boom which coincided with the period over which child outcomes are observed. Rather, the results point to a more general potential for local labor markets to generate exposure effect patterns — as where a child ends up is itself a function of their exposure to places growing up.

Sizable peer effects are also present among the permanent residents of Australian postcodes. Given the psychological literature on the heightened importance of peer relationships in adolescence (Brown and Larson (2009)), these may provide an explanation for some of the residual influence of place. While simple controls for peers fail to explain the exposure-to-place effects, this may reflect a need for better data on actual peer relationships. Further promising directions for research may include more detailed examinations of when peers matter, whether peers' parents matter directly (as role models or job contacts, say) or indirectly (as a measure of peer 'quality') and looking beyond the linear-in-means model of peer effects examined here.

Finally, this chapter serves as a validation of the research design introduced in Chetty and Hendren (2018a), and its broader applicability in settings with smaller populations and less geographic variation in intergenerational mobility. Future applications of this design in settings capturing child outcomes even later in life, and with the potential for more accurate measures of peer relationships, may yield further insights into why place matters.

3.A Additional charts

Figure 3.6: National and regional relationships between parent and child income ranks: permanent residents born in 1978

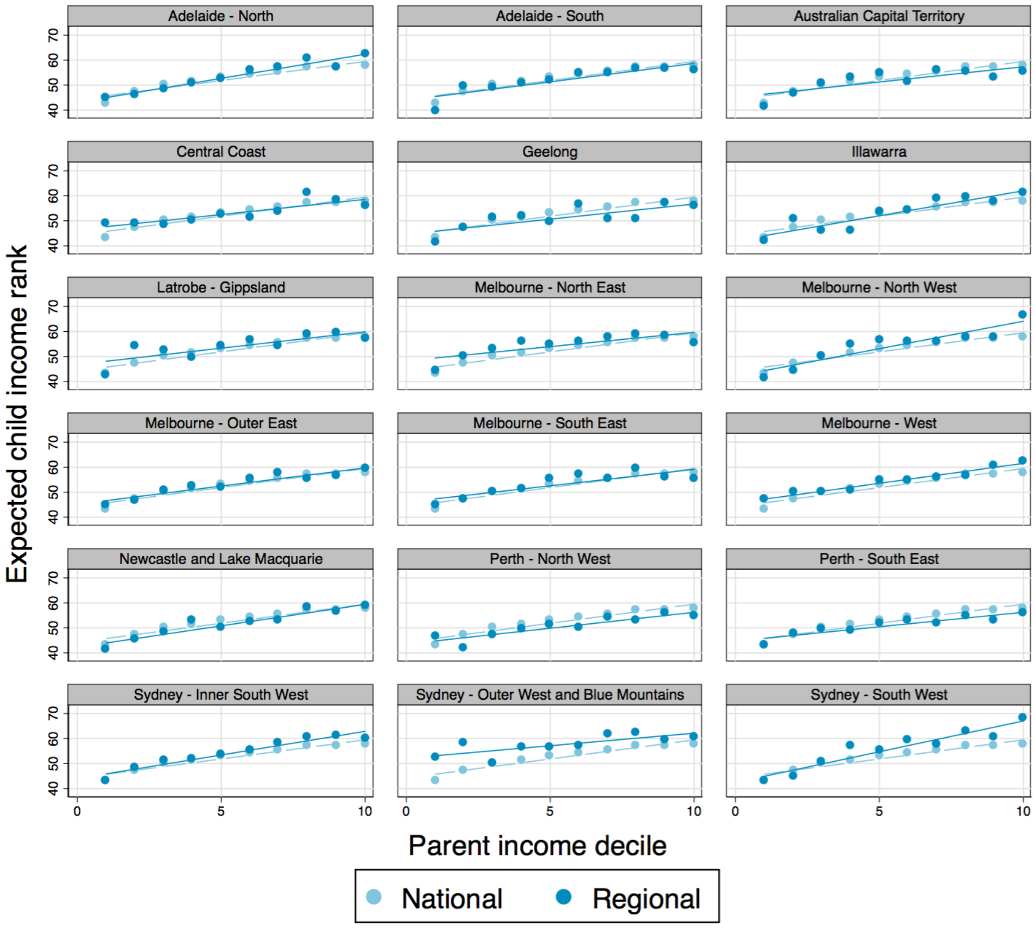

Notes: Based on the sample of permanent residents. Chart illustrates the mean household total income rank at age 24, by parent income decile, for children born in 1978 and in one of the 18 largest SA4.

Figure 3.7: Expected outcome of child: born in 1991, to parents with median income, and moving between places where similar permanent residents end up with median income

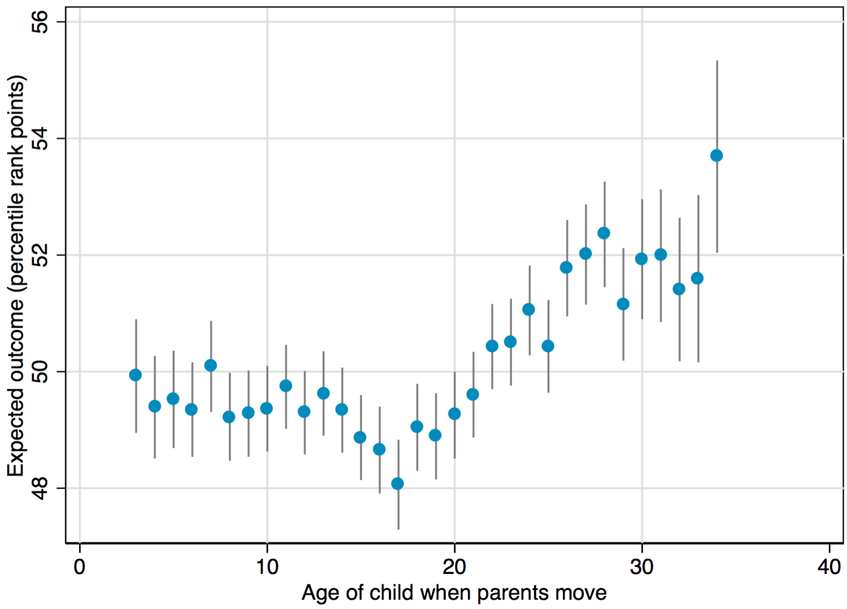

Notes: Estimated linear combination $\alpha^1_{1991} + 50\alpha^2_{1991} + \zeta^1_m + 50\zeta^2_m$ of coefficients from equation 3.3. This captures the expected household income rank at age 24 for a child: born in 1991; with parents at the 50^{th} percentile of the income distribution; and moving at age m between an origin and destination where their predicted outcome based on permanent residents is also the 50^{th} percentile. The full regression regresses the adult ranks y_i of those whose parents move once in their childhood on the interaction of their age at parent move m with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin.

Figure 3.8: Family characteristics by age at move: 1991 cohort

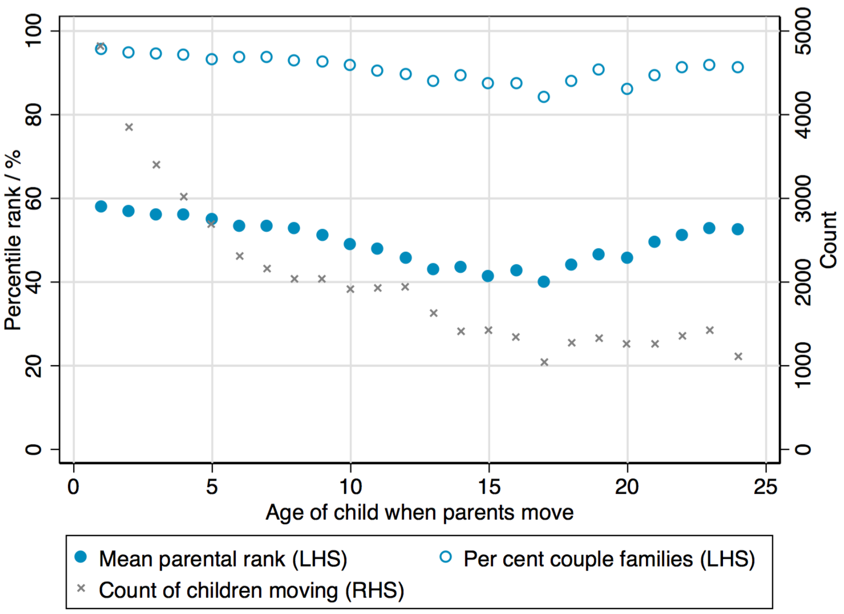

Notes: For the individuals born in the 1991 financial year whose parents move once, shows the mean parent rank, proportion in couple families and sample size by the individual's age at move.

Figure 3.9: Moving once: predicted probability of child being in destination at age 24, by age at parent move

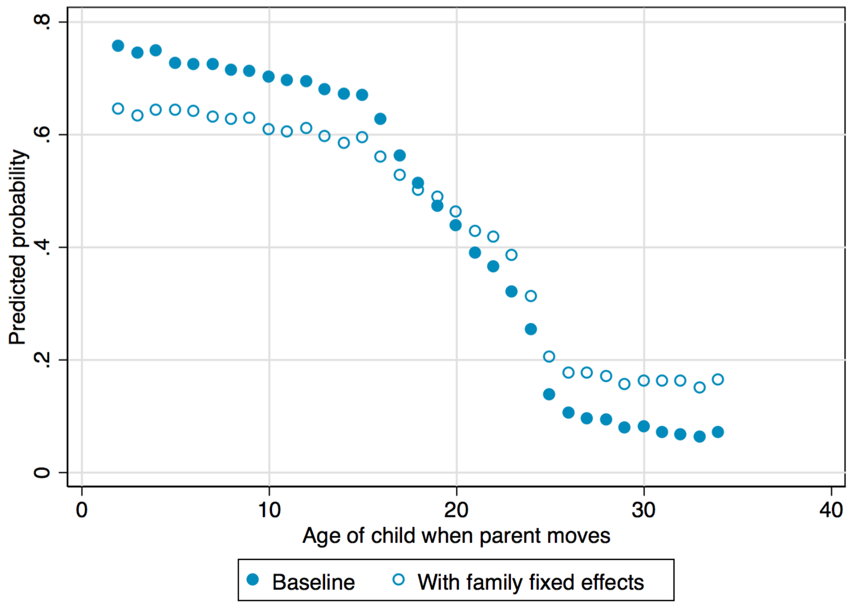

Notes: Based on the 1-time movers sample, restricted to those filing a tax return from a known location at age 24. The figure shows the predicted probability a child is filing from their parent's destination at age 24, by their age at move. Predicted probabilities are generated first from a linear regression of the binary outcome that a child is in their parents destination on age at move and birth cohort dummies (baseline), then from the same regression allowing for family fixed effects.

Figure 3.10: Moving twice: % of children in first destination at age 24, by years spent there

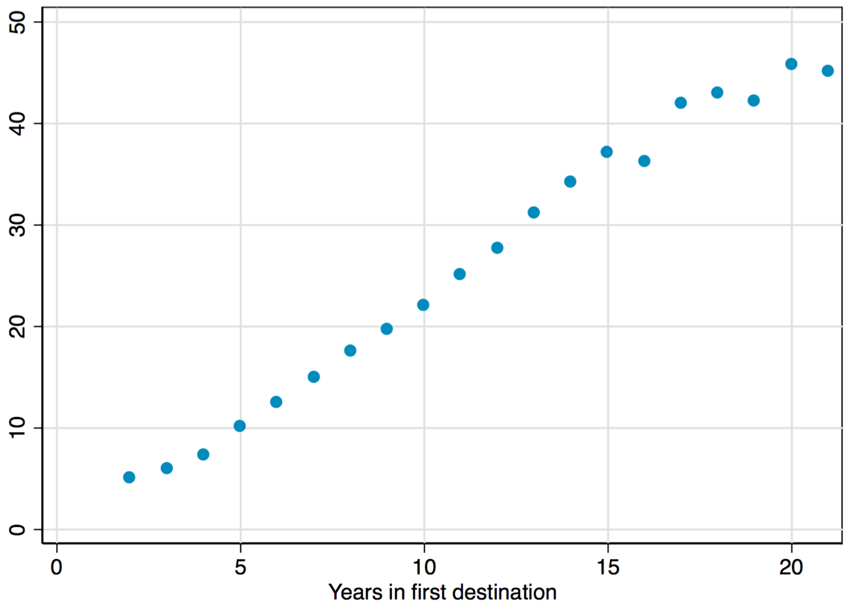

Notes: Based on a 2-time movers sample, defined analogously to the 1-time movers sample. For children of a parent who moves twice, shows the % of those in the first destination at age 24 by the time the parent spent in the first destination. The child's location is known for those lodging a tax return in the year they turn 24 with a valid postcode.

Figure 3.11: Moving once, survey data: % of children in destination at age 24, by age at parent move

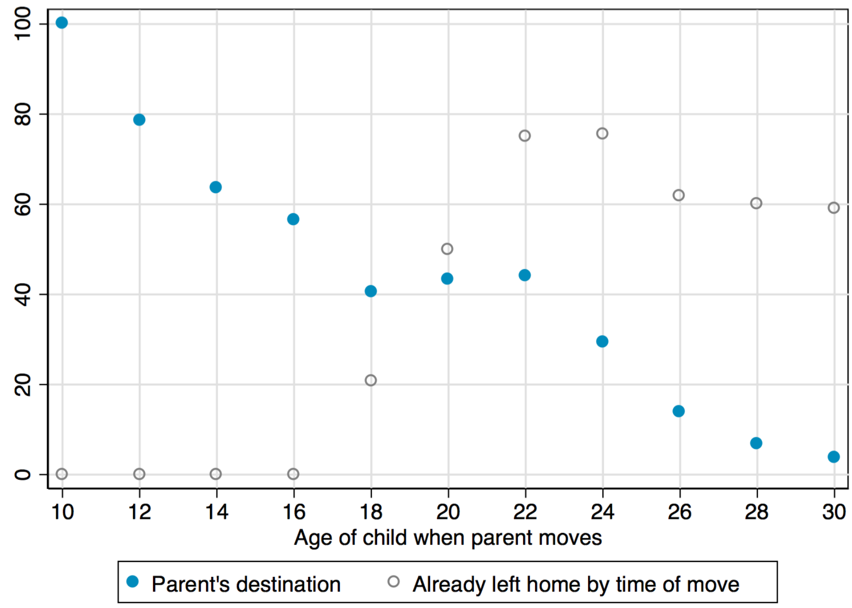

Notes: Replicates the 1-time mover sample as closely as possible using the Household Income and Labour Dynamics in Australia (HILDA) survey. I link respondents to their biological mother (or father, if the mother is missing). I restrict attention to respondents whose linked parent lived in two SA4 from 2001 (the first year of HILDA) to the year in which the respondent turned 35. Once again one SA4 episode must immediately precede the other. The chart shows the per cent share of respondents living, at age 24, in the parent's destination by their age at move. Also shown are the per cent share who had already left home by the time their parent moved.

Figure 3.12: Intergenerational mobility and economic conditions: deviation from national average in resource-rich states

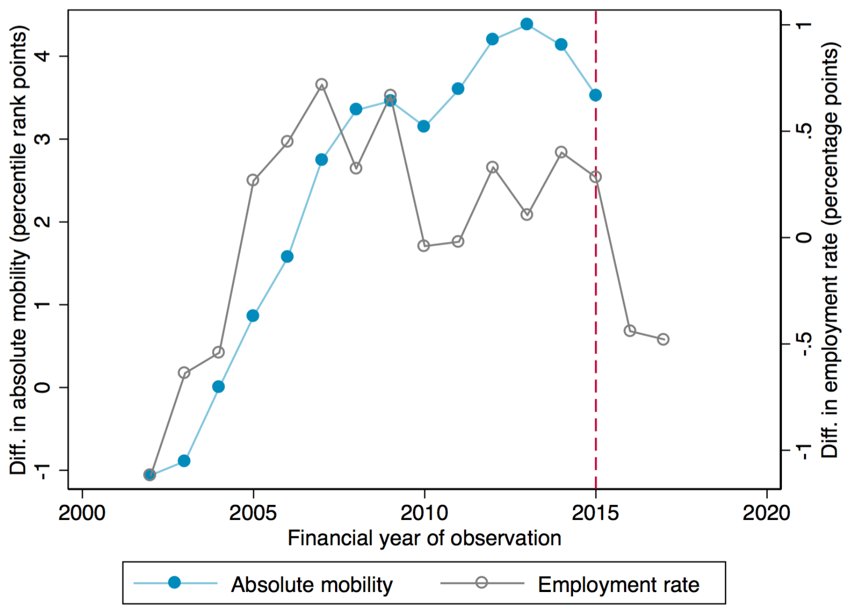

Notes: Illustrates the deviation from the national average in an individual's expected household income rank, given their parents were permanent residents of the resource-rich states of Queensland and Western Australia and were at the 25^{th} percentile of the national parent income distribution. This is shown alongside the deviation of those states' unemployment rates from the national average, from Australian Bureau of Statistics (2017d).

Figure 3.13: Effect of peer parents on household income rank at age 24: by peer group size

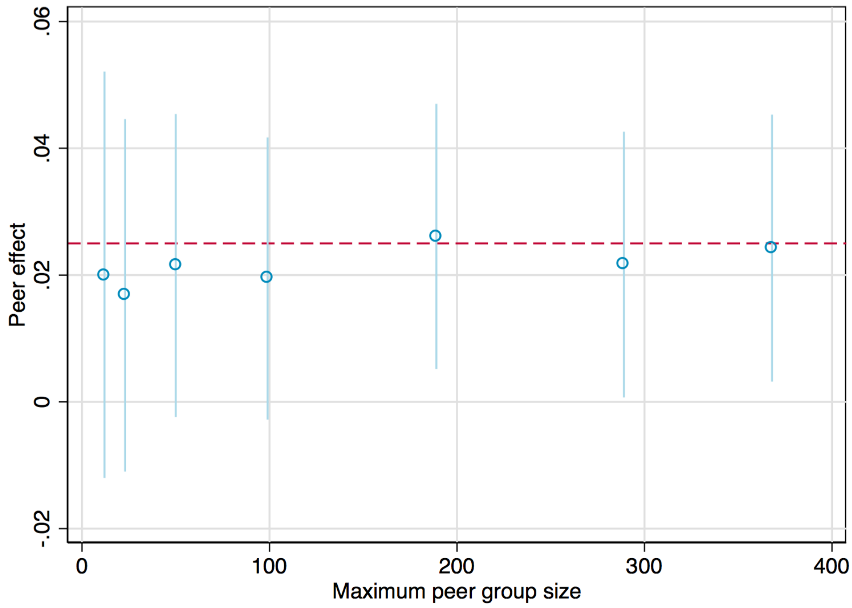

Notes: Based on permanent postcode residents. Shows the peer effect estimates (and 95% confidence intervals) from a regression of household income rank at age 24 on own parent household income rank and the mean parent household income rank of peers (defined by shared permanent postcode and financial year of birth). A 7-year moving average of the mean parent rank of peers is included as a control, in line with the specification in column (3) of Table 3.6. This regression is run for individuals in progressively smaller cohorts, with the chosen thresholds corresponding to the 5^{th}, 10^{th}, 25^{th}, 50^{th}, 75^{th}, 90^{th} and 95^{th} percentiles of cohort size in the sample. The dashed line indicates the baseline estimate on the full sample from column (3) of Table 3.6.

3.B Additional tables

Table 3.8: Summary statistics for permanent residents and one-time movers

	Permanent residents			1-time movers		
	Mean	Std. dev.	Median	Mean	Std. dev.	Median
Panel A: Family background						
Parent income ($)	79,300	71,600	72,600	86,100	80,700	77,700
Parent income rank	50.7	28.5	51	54.8	27.3	56
Indicator, in a couple family	0.87	0.33	1	0.91	0.29	1
Family size	2.7	1.2	3	2.7	1.2	2
Panel B: Outcomes						
Child income ($)	61,800	46,400	54,100	62,600	45,300	54,800
Child rank	52.2	28.4	53	52.8	28.6	54
N	1,683,800			313,900		

Notes: The full sample consists of those children born between 1978-91, remaining resident in Australia through to 2015 and linked to parents. The permanent residents are those children whose primary parent files from only one SA4 from 1991 through to the year the child turned 35. The 1-time movers are those whose primary parent filed from two SA4 from 1991 through to the year the child turned 35, filed from each at least twice, began filing in the destination the year after they ceased filing in the origin, and moved at least 15 kilometres (based on postcode centroids). Parent income is the average household total pre-tax income from 1991-2001 in 2015 dollars. Child income is the household total pre-tax income in the year the child turns 24. Ranks are calculated separately for each birth cohort.

Table 3.9: Difference between destination and origin: 1-time mover subsample

	Mean	Std. dev.	Median
Mean permanent resident parent rank	-1.06	9.98	-.83
Number of permanent residents	-95	1,000	-57
Predicted child rank	-.083	5.09	-.088
N		313,900	

Notes: Shows differences in the characteristics of the 1-time movers destinations and origins. These characteristics of place are based on the permanent residents. The difference in the predicted child rank is simply the difference in predicted values for a child in birth cohort s and with parent income rank p for a permanent resident of the origin o versus the destination d, that is $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$. The difference in the mean permanent resident parent rank is the difference in the means of the parent ranks p of the permanent residents in the same cohort s in the origin o and destination d, that is $\bar{p}_{os} - \bar{p}_{ds}$.

Table 3.10: Exposure effect estimates and model fit statistics: by model specification

	Linear	Piecewise linear with kink at age...						
		10	11	12	13	14	15	16
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	0.033							
	(0.002)							
Early		0.010	0.011	0.015	0.018	0.020	0.023	0.026
		(0.008)	(0.007)	(0.006)	(0.005)	(0.005)	(0.004)	(0.004)
Late		0.039	0.042	0.043	0.043	0.045	0.045	0.045
		(0.003)	(0.003)	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)
Post-outcome	0.008	0.008	0.008	0.008	0.009	0.009	0.009	0.009
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
$(R^2 - R^2_{max})10^6$	-10	-2	0	-1	-2	-2	-4	-6
$(aR^2 - aR^2_{max})10^6$	-9	-2	0	-1	-2	-2	-4	-6
$AIC_{min} - AIC$	-10	-3	0	-1	-3	-3	-5	-7
$BIC_{min} - BIC$	0	-4	-1	-2	-3	-3	-6	-8
N	264,500	264,500	264,500	264,500	264,500	264,500	264,500	264,500

Notes: Exposure effect estimates and model fit statistics for competing models of exposure effects — a constant exposure effects model as in Chetty and Hendren (2018a) and a piecewise linear model with the kink at varying ages. Model fit statistics are transformed as described to aid readability — higher values indicate better fits. Statistics are estimated from equation 3.4 for early ($m \in \{2, ..., k\}$), late ($m \in \{k, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure for varying values of the kink k (columns (2)-(8)) or assuming constant exposure effects in early and late childhood (column (1)). The coefficients represent the expected boost to an individual's household income rank associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. Murphy-Topel standard errors are in parentheses.

Table 3.11: Exposure effect estimates: by population subgroup

	Baseline	Gender		Parent		Destination	
		Male	Female	Poorer	Richer	Worse	Better
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Early	0.011	0.011	0.011	0.026	0.003	0.009	0.019
	(0.007)	(0.009)	(0.010)	(0.011)	(0.009)	(0.016)	(0.014)
Late	0.042	0.049	0.032	0.035	0.045	0.039	0.037
	(0.003)	(0.005)	(0.005)	(0.006)	(0.004)	(0.008)	(0.008)
Post-outcome	0.008	0.031	-0.015	0.038	-0.006	0.001	-0.018
	(0.013)	(0.018)	(0.019)	(0.022)	(0.016)	(0.028)	(0.030)
N	264,500	135,100	129,400	124,000	140,500	132,300	132,100

Notes: Estimates of the exposure effects $\gamma_{\bar{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure for the full sample and various subpopulations. These represent the expected boost to an individual's household income rank associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. The subpopulations considered are, in order: males and females; those with parental income rank ≤ 50 (or not); and those with $\Delta_{odps} < 0$ (or not). Murphy-Topel standard errors are in parentheses.

Table 3.12: Moves within non-mining states: exposure effects with local labor market controls

	Baseline	With child location controls			
	(1)	(2)	(3)	(4)	(5)
Early	0.014	0.017	0.014	0.000	0.021
	(0.010)	(0.009)	(0.009)	(0.028)	(0.017)
Late	0.041	0.027	0.023	0.029	0.025
	(0.005)	(0.005)	(0.005)	(0.011)	(0.016)
Endogenous controls					
child location ...		X	X		
... X child cohort ...			X		
Exogenous controls					
Ender predictions				X	
Ender predictions, kink at 15					X
N	160,200	160,200	160,200	126,500	126,500

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure with various controls for child location. These represent the expected boost to an individual's household income rank associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. Child location controls cover: none (1); child location fixed effects (2); child location-cohort fixed effects (3); and exogenous controls (4) and (5). Child location is simply the corresponding SA4 for those filing a tax return with a valid postcode. Standard errors are in parentheses (Murphy-Topel for (1)-(3), naive otherwise).

Table 3.13: Moves within cities: exposure effects with local labor market controls

	Baseline	With child location controls			
	(1)	(2)	(3)	(4)	(5)
Early	0.015	0.015	0.016	0.031	0.037
	(0.014)	(0.013)	(0.013)	(0.036)	(0.021)
Late	0.041	0.027	0.024	0.024	0.014
	(0.007)	(0.007)	(0.007)	(0.014)	(0.021)
Endogenous controls					
child location ...		X	X		
... X child cohort ...			X		
Exogenous controls					
Ender predictions				X	
Ender predictions, kink at 15					X
N	83,200	83,200	83,200	68,500	68,500

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure with various controls for child location. These represent the expected boost to an individual's household income rank associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. Child location controls cover: none (1); child location fixed effects (2); child location-cohort fixed effects (3); and exogenous controls (4) and (5). Child location is simply the corresponding SA4 for those filing a tax return with a valid postcode. Standard errors are in parentheses (Murphy-Topel for (1)-(3), naive otherwise).

Table 3.14: Parent and peer influences on household income rank at age 30

	(1)	(2)	(3)	(4)	(5)	(6)
Parent rank	0.170	0.170	0.168			
	(0.002)	(0.002)	(0.002)			
Peers	0.022	0.040	0.038	0.007	0.024	0.029
	(0.011)	(0.011)	(0.013)	(0.035)	(0.039)	(0.049)
Specification						
Window width	3	5	7	3	5	7
Family fixed effects				X	X	X
N	716,500	606,300	499,700	716,500	606,300	499,700

Notes: coefficients from equation 3.5 — the regression of a child's household income rank at age 24 on: their parent household income rank; and their peers mean parent rank expressed as a deviation from the 3-, 5- or 7-year moving average. Peers are defined by postcode and financial year of birth and exclude the individual in question. A peer's primary parent must have been a permanent resident of the postcode — not filing from outside it — from 1991 to the year in which the child turned 20. Robust standard errors, clustered by postcode, are in parentheses.

3.C Validation exercises

This Appendix replicates validation exercises conducted by Chetty and Hendren (2018a), with largely comforting results. The first set of tests considers the robustness of the estimates to more general specifications and later ages of observation, the remainder examine in more detail the key identifying assumption — that selection effects do not vary with the age at move of the child.

3.C.1 Specification and age at observation

In Figure 3.14 I show that the patterns of exposure effects observed in Figure 3.1 emerge even if using the more general specification in equation 3.2. This more general specification replaces parametric controls for origin and disruption effects with fixed effects for each combination of parent income decile, cohort, origin and age at move. Age-invariant selection effects, positive exposure effects and the pronounced sensitivity of the teenage years all remain apparent.

In Table 3.16 I switch attention to the models in which place effects are explicitly modeled as a function of exposure to place. Once again moving from the baseline model (column (1)) to one where parametric controls for origin and disruption effects are replaced by fixed effects (column (2)) has little effect on the estimates — if anything the sensitivity of the teenage years is even more pronounced. Lifting the age at which income is measured from 24 to 26, 28 or 30 also leaves the general conclusions unchanged.

3.C.2 Family fixed effects

The key identifying assumption behind the methodology here, and in Chetty and Hendren (2018a), is that selection effects do not vary with the age at move of the child. This seems unlikely to be true in a strict sense — certainly observables appear to differ slightly by age at move (Appendix Figure 3.8) — but it remains unclear whether the extent of any variation is sufficient to meaningfully bias the results.

An obvious place to begin testing this assumption is through the addition of family fixed effects to control for any fixed differences between families moving

Figure 3.14: Place exposure effect estimates for child income rank in adulthood.

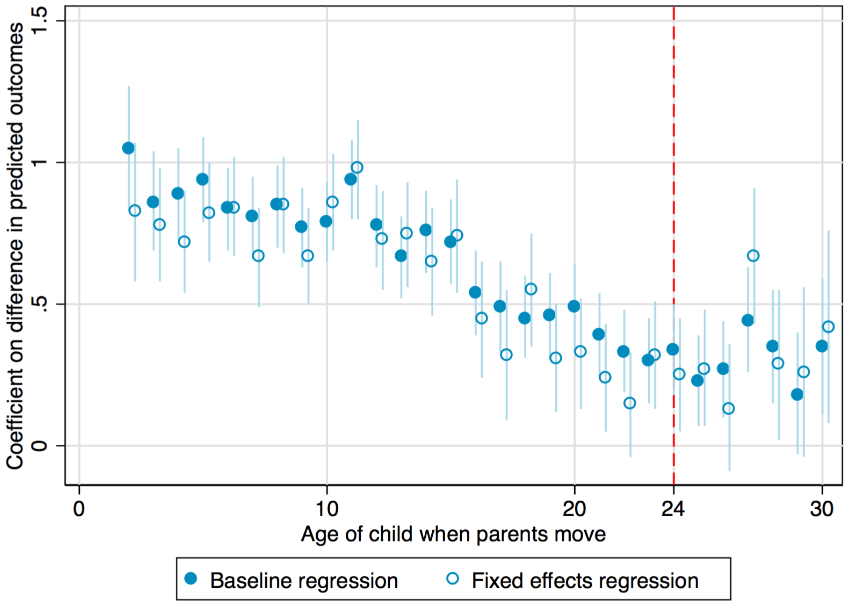

Notes: Estimated coefficients b_m from equations 3.2 and 3.3. The b_m capture the expected boost to an individual's household income rank at age 24 from moving at age m to a place with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their age at parent move m with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls vary across the specifications. Equation 3.2 includes indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin capture, alongside fixed effects for each combination of parent income decile, origin, cohort and age at move. Equation 3.3 discards the fixed effects and includes instead: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); and disruption effects (via indicators for age at move and their interaction with parental rank). This replicates Figure IV from Chetty and Hendren (2018a).

with children at different ages. I also consider family-sex fixed effects given the evidence in Table 3.11 of heterogenous exposure effects by child sex. In these fixed effect tests, identification comes from comparing siblings in different cohorts who thus differ in both their length of time exposed to the destination (the e_m) and in the predicted outcomes of their destination relative to the origin Δ_{odps} (since these are allowed to vary by birth cohort s). This requires a greater degree of precision in the measurement of the predicted outcomes to avoid attenuation bias, so more stringent sample restrictions on the estimated precision in Δ_{odps} are also considered.

The results are comforting. With family fixed effects, the estimated exposure

Table 3.15: Exposure effect estimates: more general specification and later ages of observation

	Baseline	General	Later age of observation		
	(1)	(2)	(3)	(4)	(5)
Early	0.011	-0.008	0.001	0.001	-0.013
	(0.007)	(0.009)	(0.012)	(0.023)	(0.052)
Late	0.042	0.052	0.044	0.045	0.044
	(0.003)	(0.005)	(0.005)	(0.006)	(0.007)
Post-outcome	0.008	0.013	-0.010	-0.004	0.033
	(0.013)	(0.019)	(0.015)	(0.016)	(0.021)
Age of observation	24	24	26	28	30
N	264,500	264,500	221,000	181,900	142,200

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-24 ($m \in \{25, ..., 34\}$) exposure, with either a more general set of controls (2) or for a later age of observation (3)-(5). These represent the expected boost to an individual's household income rank at the given age associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. In (1) and (3)-(5) controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. In (2) all but the last control is replaced by a much larger set of fixed effects for each combination of parent decile, origin, cohort and age at move. Murphy-Topel standard errors are in parentheses.

effect falls modestly from 0.042 to around 0.03. With family-sex fixed effects, the fall is even less pronounced, with the estimates remaining at around 0.04. This suggests heterogeneity by child sex is important in the Australian setting. I also examine the selection effect — the expected boost to an individual's household income rank from having their parent move to a destination with 1 percentile rank higher expected outcomes *after* the child turns 24. With family-sex fixed effects this selection effect is halved and no longer statistically significant. It falls further towards zero as the sample is restricted to moves where the difference in origin and destination predicted outcomes is more precisely estimated.

Table 3.16: Exposure effect estimates: more general specification and family fixed effects

	Baseline	Family fixed effects			Family-sex fixed effects		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Early	0.011	-0.025	-0.012	-0.005	-0.021	-0.002	0.001
	(0.007)	(0.010)	(0.013)	(0.013)	(0.015)	(0.013)	(0.013)
Late	0.042	0.028	0.032	0.030	0.039	0.035	0.040
	(0.003)	(0.006)	(0.008)	(0.008)	(0.043)	(0.008)	(0.009)
Post-outcome	0.008	0.040	0.011	0.025	0.018	-0.036	-0.028
	(0.013)	(0.017)	(0.043)	(0.022)	(0.028)	(0.022)	(0.026)
Selection	0.292	0.365	0.293	0.287	0.140	0.097	0.047
	(0.068)	(0.104)	(0.145)	(0.133)	(0.361)	(0.123)	(0.140)
Sample							
s.e. on Δ_{odps}	< 2	< 2	< 1.75	< 1.5	< 2	< 1.75	< 1.5
N	264,500	263,100	228,300	175,400	263,100	228,300	175,400

Notes: Estimates of the exposure effects $\gamma_{\bar{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure, with either family or family-sex fixed effects. These represent the expected boost to an individual's household income rank at age 24 associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank and cohort in the destination versus the origin. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. Attention is restricted to families with five or fewer children. Murphy-Topel standard errors are in parentheses.

3.C.3 Exogenous moves

A remaining concern is that there may be time-varying differences between families moving with children at different ages. Relationship breakdown, job loss or promotion could all give rise to moves, and themselves matter for outcomes in proportion to the time a child is exposed to them. The next test considers subsamples of moves that are more plausibly exogenous — moves out of locations in years with unusually large outflows for that location — and then re-estimates the exposure effects.

Let k_{pt} be the number of families leaving postcode p in financial year t as a proportion of the average number of families leaving the same postcode from

1991 to 2014. As in Chetty and Hendren (2018a), many of those postcode-years with the highest relative outflows k_{pt} are associated with external shocks (such as mine closures in the Australian setting).[46] As noted by Chetty and Hendren (2018a), while moves in subsamples with high values of k_{pt} may be more often for exogenous reasons, the destinations may still reflect endogenous choices. I follow them in instrumenting for Δ_{odps} and y_{ops} by $E[\Delta_{odps}|p,q]$ and $E[y_{ops}|p,q]$ — the mean Δ_{odps} and y_{ops} for all movers in the sample from postcode p and in parental income decile q. I also present OLS estimates that do not account for endogenous choice of destination.

Figure 3.15 shows the estimated late childhood exposure effect and its 95% confidence interval for subsamples drawn from moves that were part of progressively larger relative outflows from a postcode. I consider moves where k_{pt} was above its median value, 55^{th} percentile and so on to the 95^{th} percentile. The results are mixed. Below the 80^{th} percentile of relative postcode outflows the OLS exposure effect estimates are relatively close to the baseline estimate of 0.042. Beyond that point the estimates fall substantially, with negative point estimates and large standard errors for moves in the top decile of relative outflows. The IV estimates are more stable, but less precisely estimated. The average IV exposure effect estimate is 0.027, an attenuation of 30% relative to the baseline, with a less pronounced fall in point estimates in the top decile of relative outflows. In their (IV) estimates, Chetty and Hendren (2018a) see a similar attenuation of around 20% on average, but if anything less attenuation of point estimates for the top decile.

Figure 3.15 provides some comfort that the results are not driven by other factors correlated with *moderately* large relative postcode outflows, but the same cannot be said for the largest outflows. This validation exercise is thus less conclusive in the Australian setting than it appeared in the United States. One explanation is the failure of the identifying assumption — perhaps selection effects do vary with age. That said, the other validation exercises make this explanation more challenging to uphold. A more benign explanation may that the largest relative postcode outflows in Australia tend to be coupled with other factors that

[46]Postcode-years with less than ten families leaving are dropped to avoid have high relative outflows that are driven by small underlying populations. I use the same threshold as in Chetty and Hendren (2018a), purely to remain as close as reasonable to their specification.

Figure 3.15: Place exposure effect estimates for progressively larger displacement shocks.

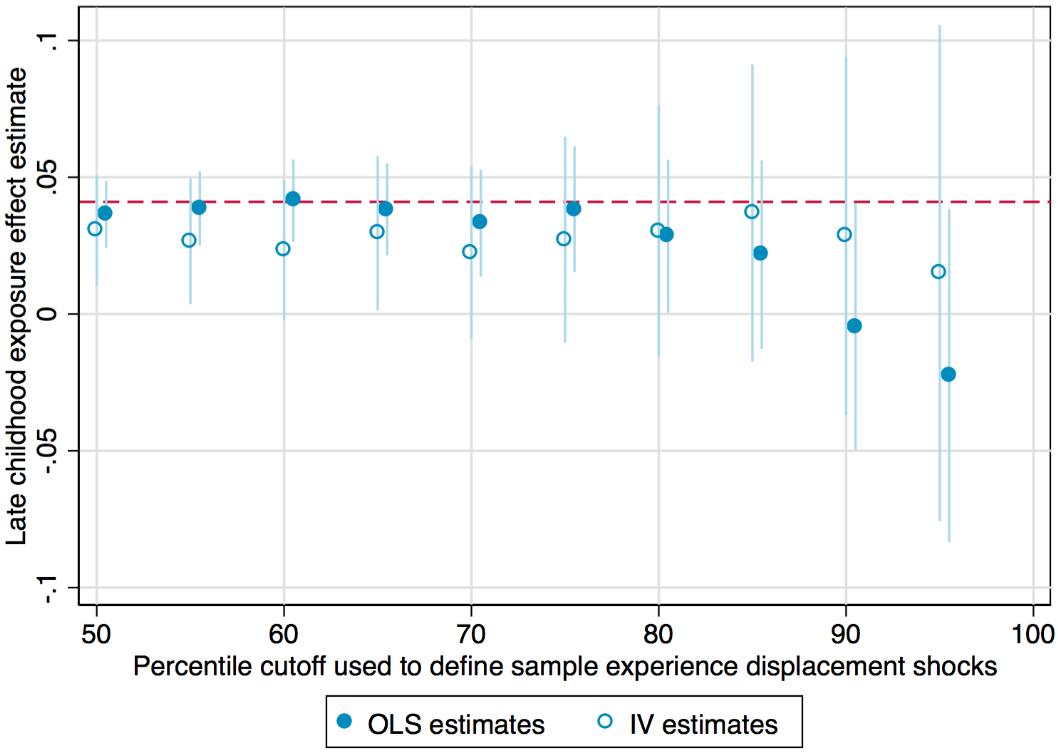

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ and 95% confidence intervals from equation 3.4 for late childhood exposure, for subsamples of those moving out of postcodes in years with progressively higher relative outflows. These are identified by first calculating, for each postcode p and financial year t, the number of families leaving the postcode divided by the average annual number of families leaving the postcode from 1991 to 2014 (call it k_{pt}). Each individual in the 1-time mover sample is thus associated with a value of k_{pt} that indicates whether they were part of a relatively small $k_{pt} << 1$ or large outflow $k_{pt} >> 1$. The chart estimates the exposure effects for those with values of k_{pt} above its median value, its 55^{th} percentile and so on. OLS estimates are presented, alongside IV estimates where the origin and destination outcomes are instrumented for as described in the text. The IV estimates replicate Figure VI from Chetty and Hendren (2018a).

mitigate the effects of exposure to the destination. Indeed, fundamental differences in the treatment effects experienced by those choosing to move versus those forced to move are apparent in Chyn (2018).[47] This would be a threat to the

[47]While Chyn (2018) finds larger treatment effects for those forced to move, this need not contradict the attenuation apparent in Figure 3.15 if, as seems plausible, the appropriate speci-

external rather than the internal validity of the baseline estimates. For example, I argue in Section 3.6.1 that the observed pattern of exposure effects can partly be explained by the fact that increased time exposed to the destination increases the probability an individual ends up working there, in the associated local labor market. It is quite plausible that being forced out of an origin due to job loss or natural disaster breaks this relationship between childhood exposure to a place and adult location. This would result in the attenuation or even disappearance of the exposure effects in adult outcomes, depending on the extent to which local labor markets are responsible for such patterns.

3.C.4 Placebo test

A final test shows the outcomes of movers converges to those of permanent residents in a manner that picks up more than just the *persistent* differences in outcomes between the destination and origin. Rather, movers converge to the *cohort-specific* outcomes of permanent residents.

This greatly limits the potential for unobserved factors to explain away the exposure effects. For example, it seems unlikely that unobserved shocks when parents move — such as to income, wealth or family status — are correlated with the as-yet-unobserved *cohort-specific* predicted outcomes for a child. Such shocks seem far more likely to be correlated, if at all, with the persistent features of a place. A fuller and more formal discussion of these issues can be found in Chetty and Hendren (2018a).

First, I show the best predictor of a mover's outcome is based on the experience of movers in their cohort, rather than those of surrounding cohorts. Following Chetty and Hendren (2018a) I run two sets of regressions. In the first thirteen regressions I re-estimate the baseline specification in equation 3.4 as if an individual's financial year of birth was $s+l$ rather than s, where $l \in \{-6, ..., 6\}$. The resulting late childhood exposure effect estimates γ are in the solid dots in Figure 3.16. Reflecting high serial correlation in a location's predicted outcomes, the exposure effects are all around the baseline estimate of 0.04. In the second single regression I re-estimate the baseline specification but include the lags and leads for the origin

fication of the treatment effect changes alongside its magnitude for exogenous shocks.

and difference terms. Where these lags or leads fall outside the sample window, the predicted outcomes are set to zero and an indicator I_l for the absence of that lag or lead is set to one. This gives rise to the specification below:

$$\begin{aligned} y_i = & \sum_{s=1978}^{1991} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}) + \sum_{m=1}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) \\ & + \sum_{s=1978}^{1991} I(s_i = s)(\kappa_s \Delta_{odps}) \\ & + \sum_{l \in \{-6,-5,\ldots,5,6\}} \left(\sum_{\tilde{m} \in M} \delta_{\tilde{m}} + \gamma_{\tilde{m}} e_{\tilde{m}} \right) \Delta_{odp,s+l} \\ & + \sum_{l \in \{-6,-5,\ldots,5,6\}} \alpha_l \bar{y}_{po,s+l} + \omega_l I_l + \varepsilon_i \end{aligned} \tag{3.7}$$

The results are in the hollow dots in Figure 3.16, and support a causal interpretation of the exposure effect estimates. The exposure effect estimate for the true cohort is only slightly attenuated. Further, while this estimate is statistically different from zero (with a p-value of 0.0057), the lags and leads are jointly insignificant (with a p-value of 0.20 on the joint test). It follows that any selection process giving rise to the observed exposure effects must do so in a way that is correlated not just with the persistent features of a place, but its cohort-specific features — a far more onerous requirement.

In Chetty and Hendren (2018a) they go further, and show movers outcomes converge to those of permanent residents in distribution and in a gender-specific manner as well. I attempted but failed to replicate the first of these. Whereas Chetty and Hendren (2018a) find that the best predictor of a mover being in the top or bottom decile of the income distribution is the predicted probability of the same outcome, rather than the predicted mean outcome, I found the opposite.[48] However, I feel the most likely explanation of this lies in the lower precision of the Australian predictions for permanent residents due to smaller geographic units and reduced geographic variation. The predicted probabilities of making the top or bottom decile are particularly imprecise and thus, if they capture more noise than signal, it is quite plausible that the predicted mean ranks may give a better indicator of the likely distributional outcomes of movers. As was noted by

[48] Results available on request. Specifications were as detailed in Chetty and Hendren (2018a).

Figure 3.16: Place exposure effect estimates: event study.

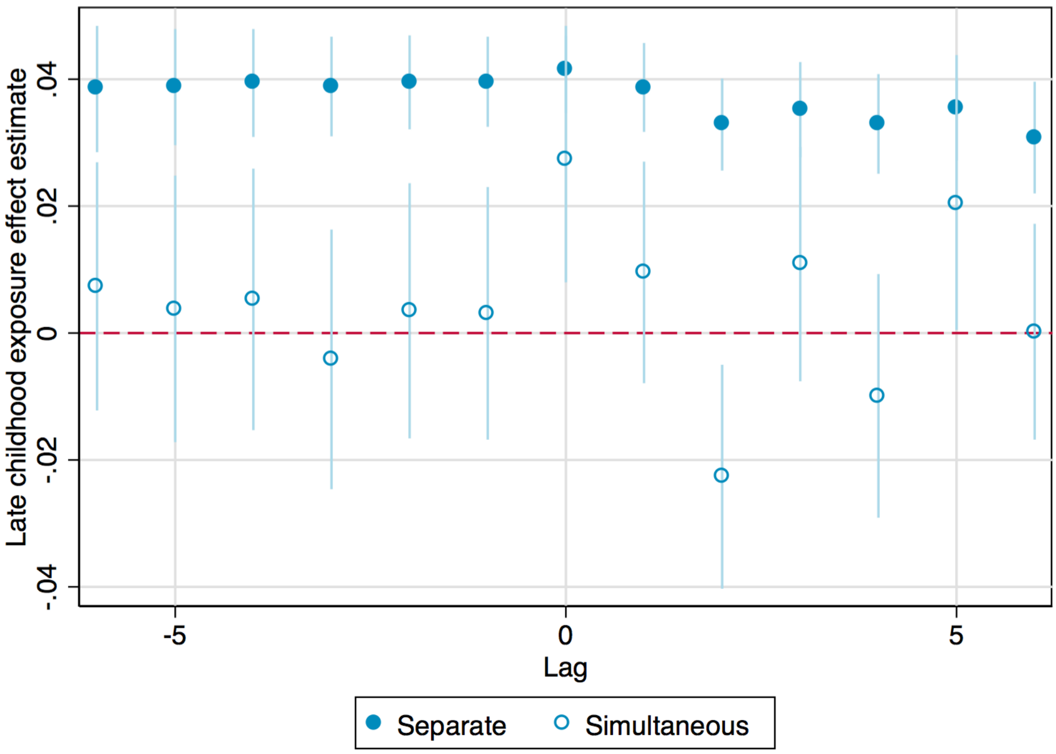

Notes: Estimates of the exposure effects $\gamma_{\tilde{m}}$ and 95% confidence intervals from equations 4 and 3.7 for late childhood exposure, when predicted outcomes are derived from a birth cohort that is not necessarily your own (solid dots) or when predicted outcomes for your birth cohort are included alongside those for neighboring cohorts (hollow dots). Thus the solid dots represent coefficients from thirteen separate regressions, using the predicted outcomes for those in financial year of birth cohort $s + l$ rather than an individual's actual birth cohort s, where $l \in \{-6, ..., 6\}$. The hollow dots run a single regression that includes the origin and difference in predicted outcome terms for all neighboring cohorts as in equation 3.7. Both these specifications allow for cohort effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. This replicates Figure VII from Chetty and Hendren (2018a).

Chetty and Hendren (2018a), an advantage of their data was the ability to conduct validation exercises requiring both large samples and significant variation.

3.C.5 Summary

The results outlined in this section provide comfort as to the internal validity of the research design introduced in Chetty and Hendren (2018a), both generally and in the Australian setting. Any unobserved factor explaining the observed exposure effects would need to operate within the family in proportion to time exposed and be able to replicate the cohort-specific outcomes of permanent residents. The examination of exogenous moves left a question mark over external validity, but could be consistent with one of the explanations for exposure effects put forward in this chapter — that time spent growing up in a location matters largely because it influences where you end up working.

3.D Generated regressors, precision and valid inference

The equations estimated in this chapter (and in Chetty and Hendren (2018a)) fall into the more general class of two-step estimation, where regressors in the model of interest are generated from an auxiliary model. In particular, in the first step, the expected outcomes y for children born into a particular location l, cohort s, and parental household income rank p are predicted based on the sample of permanent residents of that location:

$$y_i = \alpha_{ls} + \beta_{ls} p_i + \varepsilon_i \tag{3.8}$$

This model provides predicted values for the movers — denoted $\bar{y}_{ops}$ and $\bar{y}_{dps}$ — where we take their location l to be either their origin o or destination d respectively. Let $\hat{\beta}_1$ be the vector of estimated coefficients and X_{1o} and X_{1d} the matrices of observations indicating a mover's origin or destination respectively, along with their cohort and parent rank.

In the second step, these predicted values are used to generate regressors — $\bar{y}_{ops}$ and $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — for inclusion in a model for the outcomes of the movers:

$$\begin{aligned} y_i &= g(x_{2i}, \beta_2, \bar{y}_{dsp}, \bar{y}_{osp}) + \varepsilon_i \\ &= g(x_{2i}, \beta_2, x_{1d}\hat{\beta}_1, x_{1o}\hat{\beta}_1) + \varepsilon_i \end{aligned} \tag{3.9}$$

This is a classic example of the use of generated regressors. As noted in Pagan (1984), generated regressors pose a number of potential econometric issues. Perhaps most notably, while coefficients estimated from Equation 3.9 are generally consistent, the standard errors will not be, as they fail to account for uncertainty in the generated regressors. Perhaps reasonably, given they restrict attention to commuting zones with populations over 250,000, where the generated regressors are fairly precisely estimated, Chetty and Hendren (2018a) do not consider this issue. However, given the Australian data is marked by smaller geographies and less geographic variation, this issue seems worth considering in more detail here.

3.D.1 Valid inference

Murphy and Topel (1985) provide a procedure for calculating asymptotically correct standard errors in the fairly general circumstances. From the presentation in Greene (2003) the Murphy-Topel estimated covariance matrix for the model, given the two steps are estimated on different samples, is:

$$M = \hat{V}_2 + \hat{V}_2 \hat{C} \hat{V}_1 \hat{C}^T \hat{V}_2 \tag{3.10}$$

where $\hat{V}_1$ and $\hat{V}_2$ are the estimated covariance matrices for models 1 and 2 respectively and:

$$\hat{C} = \sum_{i=1}^{n} \left(\frac{\partial \ln f_{i2}}{\partial \hat{\beta}_2} \right) \left(\frac{\partial \ln f_{i2}}{\partial \hat{\beta}_1^T} \right) \tag{3.11}$$

where f_{i1} and f_{i2} are the contributions of observation i to the likelihood functions of models 1 and 2 respectively. Now, we can follow the presentation in Hole (2006) and apply the chain rule to observe that:

$$\begin{aligned} \frac{\partial \ln f_{i2}}{\partial \hat{\beta}_2} &= \frac{\partial \ln f_{i2}}{\partial (x_{i2}\hat{\beta}_2)} \frac{\partial (x_{i2}\hat{\beta}_2)}{\partial \hat{\beta}_2} \\ &= \frac{\partial \ln f_{i2}}{\partial (x_{i2}\hat{\beta}_2)} x_{i2} \\ &= \frac{\partial \ln f_{i2}}{\partial \hat{y}_{mover}} x_{i2} \end{aligned}$$

and:

$$\begin{aligned} \frac{\partial \ln f_{i2}}{\partial \hat{\beta}_1} &= \frac{\partial \ln f_{i2}}{\partial (x_{i1o}\hat{\beta}_1)} \frac{\partial (x_{i1o}\hat{\beta}_1)}{\partial \hat{\beta}_1} + \frac{\partial \ln f_{i2}}{\partial (x_{i1d}\hat{\beta}_1)} \frac{\partial (x_{i1d}\hat{\beta}_1)}{\partial \hat{\beta}_1} \\ &= \frac{\partial \ln f_{i2}}{\partial (x_{i2}\hat{\beta}_2)} \left(\frac{\partial (x_{i2}\hat{\beta}_2)}{\partial (x_{i1o}\hat{\beta}_1)} x_{i1o} + \frac{\partial (x_{i2}\hat{\beta}_2)}{\partial (x_{i1d}\hat{\beta}_1)} x_{i1d} \right) \\ &= \frac{\partial \ln f_{i2}}{\partial \hat{y}_{mover}} \left(\frac{\partial \hat{y}_{mover}}{\partial \bar{y}_{ops}} x_{i1o} + \frac{\partial \hat{y}_{mover}}{\partial \bar{y}_{dps}} x_{i1d} \right) \end{aligned}$$

In both equations the first term is simply the score vector for model 2 — for simplicity denote its elements s_{i2}. The second equation includes derivatives in the brackets that simply pick up the estimated coefficients on the predicted values. The resulting estimate of $\hat{C}$ is as follows:

$$\hat{C} = X_2^T \mathrm{Diag} \left\{ s_{i2}^2 \frac{\partial \hat{y}_{mover}}{\partial \bar{y}_{ops}} \right\} X_{1o} + X_2^T \mathrm{Diag} \left\{ s_{i2}^2 \frac{\partial \hat{y}_{mover}}{\partial \bar{y}_{dps}} \right\} X_{1d} \tag{3.12}$$

The above easily extends to the case where predicted values for neighboring cohorts are also included in the regression. The implementation of these standard errors in STATA has been outlined in Hardin (2002) and simplified in Hole (2006).

3.D.2 Precision-based sample restrictions

Finally, throughout this chapter, analysis is restricted in to those for whom the difference in predicted outcomes Δ_{odps} is more precisely estimated. The distribution of the standard error in Δ_{odps} for the 1-time movers sample is shown in Figure 3.17. For most of the analysis, I require $\Delta_{odps} < 2$, thus restricting attention to around the 80% of the sample for whom Δ_{odps} is most precisely estimated.

Figure 3.17: Distribution of standard error in difference in predicted outcomes for permanent residents of the destination and the origin

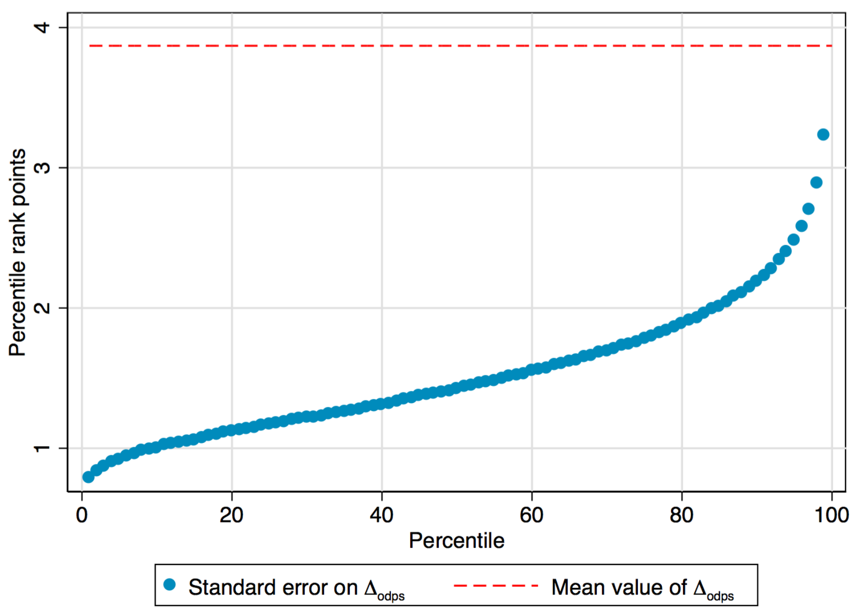

Notes: For the 1-time mover sample, shows the distribution of the estimated standard errors on the key generated regressor: Δ_{odps}. Also shows the mean value of this regressor.

Key findings are robust to this precision-based sample restriction. In Table 3.17, exposure effect estimates are shown for the baseline case, and for increasing levels of precision in Δ_{odps}. The results are not particularly sensitive to the choice of the precision-based sample restriction, with the late childhood exposure effect estimates all close to the baseline estimate of 0.042 and always larger than the early childhood exposure effect estimate.

Table 3.17: Exposure effect estimates: varying levels of precision in Δ_{odps}

	Baseline	Increasing levels of precision				
	(1)	(2)	(3)	(4)	(5)	(6)
Early	0.011	0.006	0.008	0.011	0.030	0.010
	(0.007)	(0.006)	(0.006)	(0.007)	(0.009)	(0.023)
Late	0.042	0.039	0.040	0.042	0.040	0.046
	(0.003)	(0.003)	(0.003)	(0.003)	(0.005)	(0.013)
Post-outcome	0.008	-0.001	0.001	0.008	0.015	0.052
	(0.013)	(0.021)	(0.012)	(0.013)	(0.018)	(0.056)
Sample restrictions						
s.e. on Δ_{odps}	< 2	none	< 2.5	< 2	< 1.5	< 1
N	264,500	312,900	297,800	264,500	176,300	30,200

Notes: Estimates of the exposure effects $\gamma_{\bar{m}}$ from equation 3.4 for early ($m \in \{2, ..., 11\}$), late ($m \in \{12, ..., 24\}$) or post-outcome ($m \in \{25, ..., 34\}$) exposure for the full baseline sample and larger or smaller samples based on varying restrictions on the standard error on Δ_{odps}. These represent the expected boost to an individual's household income rank at age 24 associated with an additional year at this stage of life in a destination with 1 percentile rank higher expected outcomes for permanent residents. They are estimated by regressing the adult ranks y_i of those whose parents move once in their childhood on the interaction of their time exposed to the destination at each life stage with $\Delta_{odps} = \bar{y}_{dps} - \bar{y}_{ops}$ — the difference between the expected outcomes for permanent residents of the same parent percentile rank and cohort in the destination versus the origin. the difference between the expected outcomes for permanent residents of the same parent percentile rank p and cohort s in the destination d versus the origin o. Controls capture: cohort and origin effects (via indicators for cohort and their interactions with predicted outcomes for permanent residents of the origin); disruption effects (via indicators for age at move and their interaction with parental rank); and indicators for cohort interacted with Δ_{odps} to capture potential mis-measurement of the origin. Murphy-Topel standard errors are in parentheses.

3.E Intergenerational data construction

This Appendix describes the creation of the Australian Taxation Office's (ATO) de-identified intergenerational dataset. It is based on information provided by the ATO and those involved in the construction of the dataset.

3.E.1 Overview

The dataset begins with the universe of federal tax returns from the 1991 to 2015 financial years, linked across individuals. This provides comprehensive information on individual incomes — the key challenge is linking parents and children.

Australia does not have two sources of parent-child links commonly used internationally. Birth register information is held by state and territories, and there is no national register as there is for Nordic countries. Further, parents are generally not required to provide identifying information for their children on tax returns, as family benefits are administered separately as cash transfers.[49] This rules out the methodology underlying Chetty et al. (2014), which uses the fact that parents' tax returns in the United States report their children's social security numbers.

Instead, parent-child links were formed by matching individuals to parents based on their reported residential addresses. Individuals report a residential address when they register for a tax file number — a unique personal identifier that is the closest Australian analogue to a social security number. The vast majority of individuals do this before they turn 17. These individuals are then linked to their likely parents based on residential addresses reported in tax returns. These links are disciplined by a set of more direct links available for a subset of individuals.

Address matching is behind the Statistics Canada dataset used in numerous widely-cited studies of intergenerational mobility (e.g. Corak and Heisz (1999); Oreopoulos (2003); Corak and Piraino (2011)). Yet the Australian institutional background, described below, means the ATO intergenerational dataset delivers a much higher match rate. Corak and Heisz (1999) report that they have parent links for around 49% of their selected Canadian cohorts. For the Australian cohorts

[49]Linking tax returns to this separate administrative database would have failed to provide complete parent-child links, as cash transfers have been and remain highly targeted, rather than universal.

studied in this chapter, the link rate is around 92%, in line with that achieved by Chetty et al. (2014).

3.E.2 Institutional background

Address matching delivers high quality parent-child links in Australian tax data because most individuals register for a tax file number (TFN) with the ATO while still young and living in the family home. This reflects strong incentives to do so. Since its introduction in 1989, or shortly afterwards, a TFN has been needed to:

- avoid paying higher withholding tax rates on labor and capital income;
- apply for unemployment, disability or family benefits; and
- apply for concessional loans for higher education.

As a result most transitions from childhood to independence — be it work, welfare or higher education — reward or require registering for a TFN. For example, of those born in Australia in the 1980 financial year and with a TFN by the time they were 30, over 90% had registered by age 17, and over 99% by age 20 (see Figure 3.18). Importantly, a residential address is captured for most of these children at the point of registration, and is typically of sufficient quality to match to a geocoded address.

3.E.3 Family linking procedure

The ATO dataset focuses on those born between the 1970 and 2000 income years (inclusive). Those born earlier are difficult to link to parents as many will have left the family home before the tax return panel begins in 1991. Similarly, many of those born later were yet to register for a TFN at the time the dataset was constructed.

Family links were generated for all individuals in the relevant birth cohorts, whether or not they were born in Australia. However, for the file used in this research, attention was restricted to those born in Australia. Country of birth is not directly observed in the tax data, but a good proxy for those born in Australia was derived based on other administrative information. From the 1978 birth cohort

Figure 3.18: Proportion registered for a TFN by age (Australia-born, 1980 birth cohort)

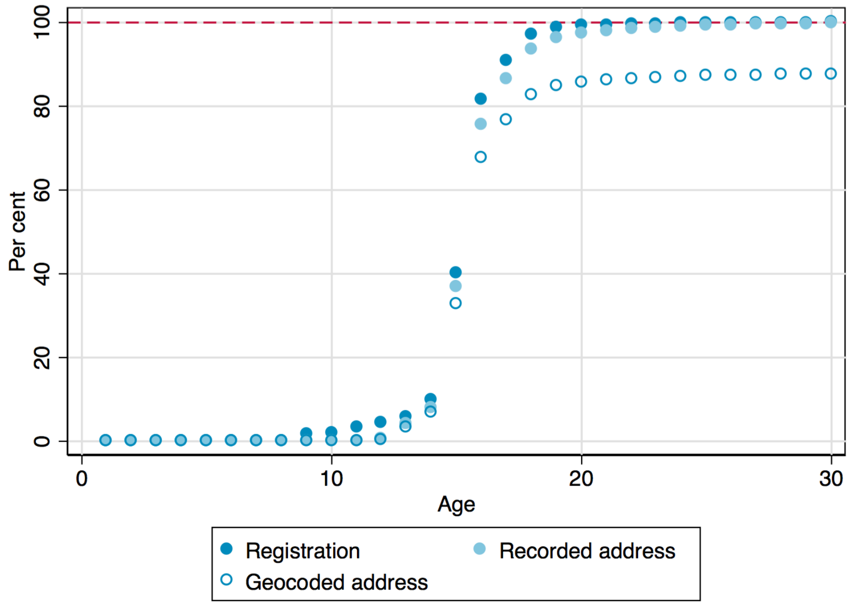

Notes: Darker blue dots show the proportion of registered clients born in Australia in the 1980 financial year who had registered by the given age. The lighter blue and hollow dots show the respective proportions with a recorded address and an address matched to a geocoded address by that age.

onwards this proxy performs particularly well, with the resulting Australian-born annual birth cohorts deviating by at most 1.5% from the population benchmark for the 1978-1991 cohorts.

3.E.4 Family Tax Assistance links

Between the 1997 and 2000 income years, Family Tax Assistance (FTA) allowed low-to-middle income families to claim a higher effective tax free threshold. Low income families could claim the entire benefit through the payments system. However, middle income families had to provide the given names and dates of birth of their children on their tax returns. This provides a relatively direct source of family links for a subset of the child population.[50] These direct links then informed

[50]FTA claims do not necessarily imply a biological parent-child relationship, though in most cases the claimant will be a biological parent or primary carer.

the algorithm for generating family links from the more widely available address links.

Initially, the details of all children a parent claimed between 1997 and 2000 were collected. This included a child's first name, date of birth and potential last names — while a child's actual last name is not listed, potential last names as inferred from those of their claiming parent and that parent's spouse. Duplicate claims were dropped and the remaining claims formed a base population of FTA children. FTA children were then linked to their adult selves among individuals registered for a TFN. A sequence of matches was performed, with only unmatched children passed to the next stage:

- *Perfect matches*: the first name, last name and date of birth match a unique individual;
- *First name error*: the last name and date of birth match a unique individual, where the two first names to have a levenshtein string edit distance of at most two;
- *DOB error*: the first name and last name match a unique individual, where the two years of birth are the same;
- *First name and DOB error*: the last name matches a unique individual, where the two first names have a levenshtein string edit distance of at most two and the two years of birth are the same;
- *Last name error*: the first name and date of birth match a unique individual.

Well over 70% of claimed children in each year were perfectly matched to an adult client. Fuzzy but unique matches allowed over 85% of claimed children in each year to be matched. These matches are spread across the birth cohorts of interest, with large numbers of those born from the 1980s onwards having FTA links.

3.E.5 Address links

As FTA could only provide family links for a selective subset of the child population, the primary source of family links was based on shared residential addresses.

As a first step, children were linked to all individuals who had ever lived at an address the child had lived at.[51] This forms the set of potential siblings and parents.

Siblings

First, individuals were linked as siblings if:

- they had been at the same address within five years of one another;
- they both lived at that address before they turned 20;
- they had less than a 13 year age gap; and
- they had the same earliest last name.

These links were 'filled' out to ensure transitivity.[52] At the end of the parent linking process individuals were also linked as siblings if they shared the same parents.

Parents

Individuals were then linked to parents. First, potential parents who were particularly young at the birth of the child (under 15 years of age) or old (45 years of age for women, 55 years of age for men) are dropped. Then the subsample of children who were *perfectly* matched as FTA children and have parent links as a result was isolated. A logistic regression was run on this subsample on the outcome that a potential parent is an FTA parent. The independent variables used in this regression were:

- potential parent sex interacted with an indicator for whether the potential parent and child share a last name;
- potential parent sex interacted with a quartic in parental age at birth either side of the median age at birth for that sex (29 for men, 27 for women)[53];

[51]The address was not required to be concurrent, as address histories in the tax data have gaps, and non-concurrent shared addresses in tax data may have been concurrent in reality.

[52]That is, if Alice is Bob's sibling and Bob is Charlie's sibling then Alice is Charlie's sibling. As a result children with more than a 13 year age gap may be identified as siblings if, for example, they share a sibling in common.

[53]Over the period 1975-1990, as calculated from Australian Bureau of Statistics (2017b).

- an indicator for whether the potential parent and child address histories imply they were at the same address at the same time;
- the length of overlap (in years) of a concurrent address episode;
- the distance (in years) separating a non-concurrent address episode; and
- the age of the child when first at the address (categorical between 13 and 25), interacted with whether the address episode was concurrent or not.

In the final step children were linked to their most probable parent, based on the logistic model's out-of-sample predictions for the probability a potential parent was an FTA parent. Each child was linked to the potential parent with the highest predicted probability of being an FTA parent, conditional on that probability being greater than 0.5. At the chosen threshold, less than 4% of the address-derived parents for the FTA subsample failed to match the FTA parent. Given the FTA links are not infallible, this seems reasonable. FTA link parents were then used for those children with no parent.

Postcode links

Address matching is limited by the absence of complete address histories. Further back in the panel tax filers are less likely to have a recorded residential address. However, residential postcodes are reliably recorded — they are captured for the vast majority of tax filers in each of the years between 1991 and 2015. To exploit this, a set of supplementary links is based on residential postcode histories.

For all children, the postcode of their first address was extracted, typically their address when registering for a TFN. Children were also assigned their earliest recorded last name. Children were then linked to all individuals in the same postcode in the same year and with the same last name. This formed the set of potential parents. As in the address matching, potential parents who are particularly young at the birth of the child (under 15 years of age) or old (45 years of age for women, 55 years of age for men) were dropped. Given the large number of potential postcodes and last names, most children end up with a relatively small set of potential parents.

In the next step the subsample of children who perfectly matched FTA children and have parent links as a result is again isolated. Once again, a logistic regression was run on the binary outcome that a potential parent is an FTA parent. The independent variables used in this regression were:

- potential parent sex interacted with quartics in parental age at birth either side of the median age at birth for that sex (29 for men, 27 for women)[54];
- an indicator for if the potential parent's spouse is among the alternatives, also interacted with a categorical variable for the number of potential parents for the child (top coded at ten);
- the age of the child when first at the postcode (categorical between 13 and 25).

In the final step children were linked to their most probable parent in the same manner as for the address matching. The exception is that here a slightly more conservative threshold was set — children were only linked to parents if the estimated probability of the potential parent being an FTA parent was greater than 0.75. The accuracy of this algorithm was only a little worse than the address matching. At the chosen threshold, only 8% of the supplementary parents for the FTA subsample fail to match the FTA parent.

Sibling links

In this step, individuals with siblings were linked to the most probable parent for their family. The steps were as follows:

- look across groups of siblings (families);
- identify the most probable parent for each family — that is, the potential parent with the highest estimated probability of being a true parent[55]; and
- match individuals to the resulting most probable parent for their family.

[54] Over the period 1975-1990, as calculated from Australian Bureau of Statistics (2017b).

[55] FTA parents were assigned a probability of 1 and supplementary parents were all assigned to 0.75.

Reassuringly, this process showed a great deal of consistency in the parent-child links. For children already linked to a parent, that parent is not replaced, or replaced by their spouse, in 90% of cases. Once children are matched to their most probable parent, those parents are matched to their earliest reported spouse over the period 1991 to 2015.

Resulting parent-child links

The parent-child links resulting from this process are shown in Figure 3.19. From 1978 to 1991 onwards the sample closely matches the size of the Australian-born population, deviating from the population benchmark by at most 1.5%. The proportion of the population linked to parents averages around 92% over this period as well. For the birth cohorts examined in this chapter, 88% of links are derived from shared residential address, 4% from FTA, 2% from postcodes and 6% from siblings.

Figure 3.19: Sample coverage rates relative to the population of interest (%)

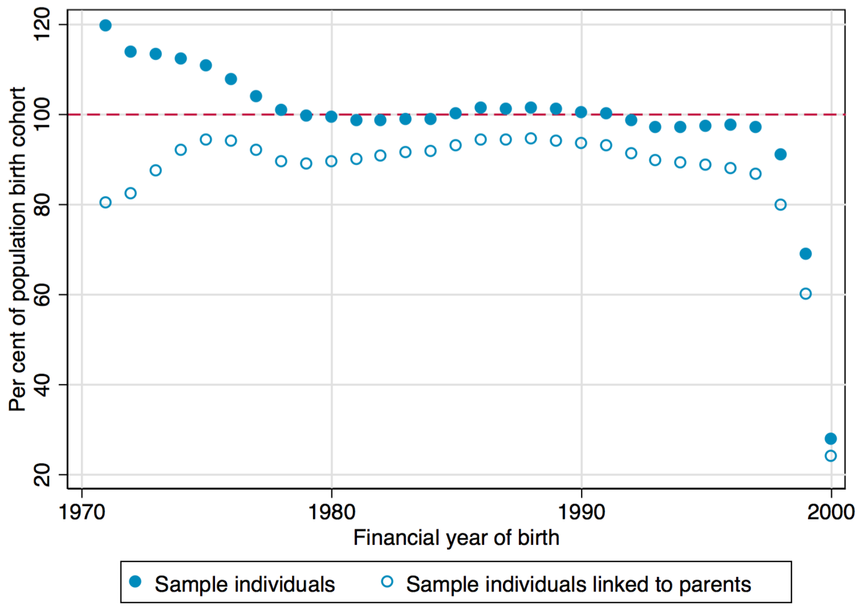

Notes: Shows the number of individuals in the sample, and the number linked to parents, as a percentage of the relevant population of interest. The population of interest is taken as the number of births in Australia (Australian Bureau of Statistics (2017b)), or for financial years prior to 1976 (where this data is not available) the estimated resident population aged zero on the last day (30 June) of the relevant financial year (Australian Bureau of Statistics (2017a)). Where both series are available the deviate by at most 2%.

Chapter 4

What drives second generation success? The role of education, culture and social context

Abstract

I explore the role of education, culture and social context in the intergenerational income mobility of second generation migrant communities in Australia. I present a new decomposition of intergenerational income mobility, and find a central role for differences in education mobility in driving differences in income mobility between migrant communities. Further, differences in the cultural values migrant communities bring with them, and the context of their migration, are associated with large differences in second generation educational achievement and attainment. Second generation migrants from countries that outperform on tests of student achievement, or face higher income penalties in the first generation, tend to have better educational outcomes. I use a rich array of survey and test score data to show the outperformance of migrants from poorer backgrounds emerges late in adolescence, and is reflected in attainment, aspirations and the perceived returns to education, but not in school test scores.

The author has been supported by a Sir Roland Wilson scholarship, funded by the Australian Treasury and the Australian National University. The helpful suggestions and advice of Robert

4.1 Introduction

Intergenerational relationships in outcomes within a country often mask considerable heterogeneity — be it across geographic, ethnic or family lines. What then drives some children to succeed while others struggle? What role do differences in abilities, in economic and institutional factors, or in beliefs and preferences that may be specific to a community's culture or social context play? These questions are particularly salient in considering the intergenerational mobility of migrants. Migrants are both socially and economically important in many advanced economies, making the outcomes of the second generation an issue of much interest. They are also a group where variation in some of these factors may be more pronounced or more plausibly isolated than in the broader community.[1]

I make two main contributions to the literature on the intergenerational mobility of migrants (see Sweetman and van Ours (2015) for a recent survey). First, I present a decomposition of intergenerational income mobility, and use it to shed light on the potential drivers of second generation outcomes across a large set of source countries. Second, I examine the drivers of educational success, using an array of individual-level data tracking test scores, aspirations and attainment through adolescence and into adulthood. The Australian setting is ideal due to its large, diverse and longstanding migrant population providing significant variation — both across source countries and over time — and ability to precisely identify

Breunig, Andrew Leigh, Bhash Mazumder, Peter Siminski and Zachary Ward have been greatly appreciated. I thank Alan Jenner of the Australian Bureau of Statistics for his expertise, advice and work in extracting the historical Census data. This paper uses data from the Youth in Focus project. The Youth in Focus project was jointly funded by the Australian Government Department of Education, Employment and Workplace Relations, the Australian Government Department of Families, Housing, Community Services and Indigenous Affairs and the Australian Research Council, and carried out by The Australian National University. This paper uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Project was initiated and is funded by the Australian Government Department of Social Services (DSS) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this paper, however, are those of the author and should not be attributed to either DSS or the Melbourne Institute, or other data custodians, funding organisations or employers.

[1] There is a long tradition in economics of using the experiences of migrants to shine a light on the role of factors that may be more difficult to observe in the broader community. For example, migrants have been used to study the role of networks in job search (Munshi (2003)), and culture in female labour force participation and fertility decisions (Fernandez and Fogli (2009)).

and link migrant generations in a wide variety of census, test score and survey data.

I begin with group-level data linking generations of migrants across Australian Censuses. Similar approaches have been used to examine the intergenerational mobility of migrants in the United States and Canada (Borjas (1993); Card et al. (2000); Aydemir et al. (2009)). I develop a simple decomposition of the residual from a standard intergenerational income mobility regression into residuals from underlying regressions representing education mobility, and the returns to education in both generations. In this framework, exceptional income mobility will reflect some combination of: the unwinding of particularly low or high returns to education for the first generation; exceptional education mobility; or particularly low or high returns to education in the second generation. All three channels are plausible, and have some support in the existing literature. The advantage of this decomposition is that it allows all three to be considered alongside one another. Further, the decomposition is allowed to differ across migrant communities, rather than averaging across potentially diverse experiences.

I find a central role for differences in education mobility in driving differences in income mobility between migrant communities. In particular, the upward income mobility of Asian Australians is mostly accounted for by their upward educational mobility. However, the decomposition varies greatly by source country and region. Lower second generation returns to education play an important role in the mobility of Middle Eastern and South American Australians. And the unwinding of particularly high and low first generation returns for migrants from the United Kingdom and Mainland South East Asia respectively influence the observed mobility of the second generation. Finally, the mobility of migrant communities is relatively persistent by source country over a thirty-year period, despite significant changes in policy over this period.

The importance of education and persistence of source country effects motivates the remainder of the paper. What persistent factors do migrant communities share that influence outcomes — does culture play a role? To examine the potential role of culture, I adopt the 'epidemiological approach' from the broader literature on culture and economics (Fernández (2011)).[2] The approach hinges on the idea

[2]While culture is a difficult concept to define, in this context it is best thought of as systematic

that migrants may more readily bring their culture with them than their economic and institutional environment. Central to the approach is the use of a quantitative proxy for the outcome in question. I use a source country's performance on the Programme for International Student Assessment (PISA), adjusted for their income level. In the source country, PISA performance reflects the decisions and investments of individual families and children, and broader economic and institutional settings, such as the quality of the education system. Performance will thus capture an element of cultural attitudes towards education.[3] If this proxy has explanatory power over second generation Australians, born and educated in Australia, it likely reflects the influence of cultural traits that migrants bring with them, rather than the environment they leave behind.[4]

Of course, migrant communities share more than just their culture — they also share the experience of migration and integration itself. In particular, a large body of research has examined the origins and evolution of the wage gap between migrants and similarly skilled natives.[5] Motivated by this, I use the income gap between first generation migrants and similarly educated natives as an additional control variable; a proxy for the context of migration. This may capture many relevant features of the migration experience, such as the factors motivating migration, and the adjustment costs faced on arrival.

I find sizeable and statistically significant associations between these cultural and contextual proxies and second generation education. Being born to a father from a source country that outperforms on PISA by a full standard deviation in student test scores — roughly the gap between Vietnam and New Zealand, or New Zealand and Lebanon — is associated with an additional 0.5 years of education. Being born to a father from a community facing an income penalty that is 10 percentile rank points higher — roughly the gap between Vietnam and Italy, or

variation in beliefs and preferences across source countries. This definition is borrowed from Fernández (2011).

[3]Indeed, cross-country studies of the influence of education policies on test scores work hard to exclude cultural factors as a potential omitted variable driving their conclusions (Hanushek and Woessmann (2011)).

[4]Naturally, migrants may choose to settle in locations and schools that resemble the economic and institutional environment in their source country. I thus test the robustness of my results to the inclusion of controls that attempt to account for the fact that migrants may take some of their environment with them.

[5]For example, Chiswick (1978); Borjas (1985); Borjas et al. (1987); Lubotsky (2007).

Italy and England — is associated with an additional 0.3 years of education.

The fact that a higher income penalty is *positively* associated with second generation success is potentially surprising. Ordinarily we might expect the opposite — higher incomes may relax financial constraints to investing in education, or may indicate higher unobserved ability or other advantages that may be shared with the second generation, such as geography or a lack of discrimination. But there are also multiple credible mechanisms running in the other direction. For example, it may be first generation migrants who face higher income penalties are more likely to have migrated due to altruistic concerns, and thus also invest more in their children. Alternatively, they may come from countries where educational opportunities were limited, and hence parents may possess and pass on to their children higher unobserved ability.

I shed light on these mechanisms using rich test and survey data. In particular, I show the association between a higher income penalty and success is likely driven by aspirations rather than ability or parental investments that pay early dividends — the association is relatively weak when looking at test scores at age 15, but strong when looking at educational attainment. Further, those from communities with higher income penalties plan on more education, and see it as more important to 'getting ahead'. The results are consistent with a world in which individuals face the same constraints as econometricians (at least) in determining the true return to education (Manski (1993)) and form their expectations based on the world they see around them. In this sense, a higher income penalty in the first generation could influence expectations, aspirations and subsequent attainment. That said, it may also be that parental altruism or other pre-migration differences, as proxied for by the income penalty, only reveal themselves late in the second generation.

Omitted variables are the key constraint on any causal interpretation of the results. I address these concerns in a variety of ways that make the case for a causal interpretation more compelling. First, the cultural and contextual proxies retain their explanatory power when controlling for other explanations for 'ethnic capital', such as more traditional human capital externalities (as in Borjas (1992)) or community size (as in Gang and Zimmermann (2000)). Both proxies remain important in the survey and test score data where precise controls for an individual's family background and institutional setting can be included. PISA outperformance appears most strongly associated with similar measures of student achievement, while the income penalty is most strongly associated with eventual educational attainment. Finally, both these associations remain when estimation is restricted to within source regions, rather than drawing on variation between them. This latter test exploits a key advantage of the Australian setting and data — a large, diverse and longstanding migrant population that allows variation within source regions and over time to be exploited. While the selectiveness and composition of migrants differs dramatically across nations, the patterns of source country effects is often similar (Sweetman and van Ours (2015)), and the results presented here thus offer insights that may apply beyond the Australian setting.

4.2 Related literature

Migrants and their recent descendants are a large and growing share of the population of the developed economies (Sweetman and van Ours (2015)). Their outcomes have an important bearing on destination countries, and of course the wellbeing of the individuals concerned. Unsurprisingly then there is a large body of work on the intergenerational mobility of migrants.

Much of the literature on the intergenerational mobility of migrants considers the general degree of persistence in the economic status of migrant communities, relative to that observed in the population as a whole. In pioneering work, Borjas found higher levels of persistence among migrant communities, with the outcomes for an individual reflecting not just those of their parents, but those of their parents' ethnic group (Borjas (1992, 1993)). Borjas noted that multiple mechanisms could underlie this 'ethnic capital'. There have been many subsequent studies of the mobility of migrants and on the outcomes of second generation migrants in general (Sweetman and van Ours (2015)).[6] A common observation is that the migrant source country matters — and this is a focus of this paper. Why do some migrant communities do particularly well, or not?

A number of existing studies have examined the outcomes of second genera-

[6]For example, migrant mobility papers have reexamined the United States (Card et al. (2000)), including during the age of mass migration (Abramitzky et al. (2014)), but also investigated Canada (Aydemir et al. (2009, 2013)) and Sweden (Hammarstedt and Palme (2012)).

tion Australians. Khoo et al. (2002) compares outcomes for second generation Australians by source country — however, data limitations prevent them from exploring what lies behind the diversity of second generation outcomes. More recent studies have looked at specific aspects of second generation success. For example, Messinis (2009) examines the role of language and job mismatch in second generation earnings. Cobb-Clark and Nguyen (2012) find that young Australians from non-English-speaking background immigrant families have an educational advantage over English-speaking background and Australian-born peers, which largely stems from differences in outcomes between those with similar family backgrounds. Lastly, the performance of the second generation on PISA and other educational outcomes, including the success of Asian Australians and explanatory power of behaviours and cultural traits, has also been examined (Marks (2010); Mendez (2015)).

Finally, this paper can be seen in the context of a rapidly growing literature on intergenerational mobility more generally. Recent contributions have highlighted significant heterogeneity within nations across regional and racial groupings.[7] Such findings complement a more longstanding observation from sibling correlation studies of heterogeneity across families — much of the effect of family on a child's eventual income is transmitted through mechanisms uncorrelated with family income itself (Black and Devereux (2011)).[8] This paper helps shine a light on some such mechanisms.

[7]For example, geographic variation in intergenerational mobility in the United States (Chetty et al. (2014)) and Italy (Güell et al. (2018)). Differences by race in America are also well established (e.g. Hertz (2008); Bhattacharya and Mazumder (2011); Chetty, Hendren, Jones and Porter (2018)).

[8]A few studies of sibling correlations have shone a light on transmission mechanisms for family fixed effects by examining the extent to which correlations can be explained by child or parent characteristics. For example, Mazumder (2008) finds up to half the brother correlation in the United States can be accounted for by their educational attainment and test scores, while Björklund et al. (2010) find a similar level of explanatory power in indicators of parenting practices and attitudes. This paper complements those studies, finding education is pivotal to migrant mobility and pointing to a role for culture and social context.

4.3 Data

I use a mix of full population census data, and detailed test and survey data, drawing on the relative strengths of each in the analysis. The census data allows the intergenerational mobility of migrants to analysed at a group-level for the largest possible set of communities, as even small migrant communities are captured.[9] The survey data misses some of the smaller communities, but allows a much richer analysis of outcomes and set of controls.

4.3.1 Full population Census data

I use customised tabulations from the full unit record files of eight Australian Censuses of Population and Housing. The Australian Census is well placed to examine questions relating to the intergenerational mobility of migrants. Australia has a large and diverse migrant population, and five-yearly censuses have routinely collected information on the income, educational attainment, and own and parent country of birth of the full population.

Second generation sample and outcomes

The 2016 Census is my primary source of information on the second generation, as it is the first since 1996 to ask for parent country of birth.[10] I begin by grouping all Australian-born individuals by intercensal birth cohorts (c) and male parent country of birth (o). I focus on six birth cohorts (those born prior to the 1966, 1971, 1976, 1981, 1986 and 1991 Censuses) and 76 country groupings (75 countries, including Australia, and an other grouping).[11] This results in 456 possible

[9]Previous studies have typically relied on samples, rather than full population counts, of census data. For example, Borjas (1993) uses 1-2% sample files from the United States Census while Aydemir et al. (2009) use 20% sample files from the Canadian Census. Full counts allow even small migrant communities to be studied, increasing variation in source country characteristics, and allowing communities to be studied from their emergence. It also eliminates measurement error due to sampling variation that will result in a mismatch between those observed in the first generation versus those observed in the second.

[10]Data from the full count 2016 Census was extracted using the Australian Bureau of Statistics' Census TableBuilder product.

[11]The countries were chosen based on the size of the second generation population in the relevant cohorts and ability to link from the 2016 Census back to 1976 Census. Some countries no longer exist as a single entity — for example, Yugoslavia and Czechoslovakia.

second generation 'migrant communities' (including six that are actually born to Australia-born fathers).

I then calculate the mean income percentile rank and years of educational attainment for each migrant community. The use of income ranks has a number of advantages. A well known problem in the intergenerational mobility literature is the life-cycle bias that can arise if measuring child incomes too early in their working life, and using it as a proxy for lifetime income (Haider and Solon (2006)). In short, future high income earners experience faster income growth early in their working life, and measuring mobility at this point may understate the degree of intergenerational persistence. The second generation Australians I consider were aged anywhere from 25-54 at the time of the 2016 Census, making this a potential concern. However, income ranks (as opposed to levels) tend to settle down earlier in the lifecycle — for example, using Swedish administrative data Nybom and Stuhler (2017) find rank-rank mobility measures are reasonably stable from age 30 onwards. Similarly, attenuation bias arising from measuring paternal income over too short a window is also less of a concern when working with rank-based measures (Mazumder (2016)).

I calculate mean income percentile ranks as follows. First, individual income is reported in one of 15 brackets. Using all Australian-born individuals in the birth cohort as a benchmark, I convert these to fractile brackets. For example, if 10 per cent of the birth cohort are in the top income bracket (total individual income of AUD156,000 or more), the fractile bracket spans from 0.90-1.00 with midpoint 0.95. For each cohort-origin cell, I then calculate the mean income rank by assigning all those reporting income in a bracket to the corresponding fractile bracket's midpoint, taking the mean and multiplying by 100.

I extract educational attainment in one of the 41 categories, ranging from no education through to a doctorate.[12] I map these categories to years of education, creating a variable ranging from 0 through to 19 years, and once again calculate the mean for the migrant community. This mean is a near-continuous measure of educational attainment, unlike the years of education for any individual, which is necessarily discrete.

[12]Specifically, I use the 3-digit level of the 'highest educational attainment' variable that the Australian Bureau of Statistics derives from individual responses.

First generation sample, outcome and intergenerational links

I link cohort-origin cells of second-generation Australians to their first generation fathers in the 1976 through to the 2006 Censuses (the seven censuses where income is reported and the birth cohorts of interest are still likely to be living at home). I do this by grouping men according to whether they have Australia-born children in one of the six birth cohorts (c) and their country of birth (o).

I then calculate the mean income percentile rank and years of educational attainment for the first generation fathers in each migrant community. To calculate income percentile ranks I again consider individuals in the same birth cohort. The tabulations thus group men into five-year age intervals, as well as by their country of birth and the birth cohorts of any Australia-born children. I focus on men aged 35-39 and 40-44 years to measure income mid-career, but also at a point when children are still likely to be living at home.[13] I also consider men aged 45-49 and 50-54 years in the 1976 Census to capture earlier birth cohorts, as the Censuses prior to this one did not collect income information. Appendix Table 4.7 shows in detail the child and father birth cohorts captured, and their corresponding ages at observation in the censuses. Income is again reported in brackets through the seven censuses, with 12-16 brackets. Using all fathers in the same birth cohort (and with Australia-born children who also share a birth cohort) as a benchmark, I again convert the income brackets to fractile brackets. As before, I calculate the mean income rank for each migrant community, though this time I have different means for each father birth cohort and census year. I average these to give a single number for each cell. I first take simple averages for fathers observed in multiple censuses — for example, those aged 35-39 in one census with a child aged [a,a+5], and then aged 40-44 in the following census with a child aged [a+5,a+10]. I then take averages weighted by the number of children, to produce the mean income percentile rank of the migrant community for the first generation fathers.

Educational attainment is simply measured in the first census in which the

[13]This limits the analysis to children whose fathers were aged 16-44 years at the time of their birth, given I only capture children who are aged 0-19 years. This is not a problem for the analysis, as it includes the vast majority of children — my analysis of Australian Bureau of Statistics data suggests over 95% of children will be captured (based on paternity ages from 1975-1991 — data is not available prior to 1975) (Australian Bureau of Statistics (2017b)).

child's birth cohort is captured. By placing no restrictions on the father's age, I assume that educational attainment at paternity is a reasonable proxy for lifetime educational attainment. I extract years of education as a categorical variable, again ranging from no education to doctorate.[14] I map these categories to years of education, creating a variable ranging from 0 through to 19 years, and once again calculate the mean for the migrant community.

4.3.2 Summary statistics and sample selection

As noted earlier, an advantage of the full count Census data is the ability to examine even small migrant communities. However, the confidentialisation process applied by Australian Bureau of Statistics when extracting the data will make particularly small cells unreliable. As such I restrict attention to cohort-origin cells with more than 200 individuals in the second generation. This appears to strike a reasonable balance between maintaining the ability to look at emergent communities, while also minimising measurement error.

A potential concern in a pseudo-panel such as this is mismatch between the populations used to calculate the first and second generations. Selective outmigration, family dissolution or measurement errors could all lead to the first generation and second generation outcomes being observed for different populations for each migrant community, and a resulting bias in the estimates. In Appendix 4.C I consider this issue in detail. The populations of second generation individuals captured as children in the historical censuses and as adults in the 2016 Census typically line up well, and I exclude those migrant communities where there is a change of more than 30% in magnitude. In Appendix 4.C I show the key associations persist under alternative sample specifications.

Table 4.1 presents summary statistics for the first and second generations of the migrant communities in the sample.[15] Panel A presents the variation across all migrant communities by cohort and origin, while Panel B presents the variation within migrant communities from the same origin, by first subtracting the mean

[14] The classification was created by the Australian Bureau of Statistics to reflect the 'highest educational attainment' variable available in the 2016 Census.

[15] I include only the true migrant communities, dropping the six observations for those born to Australia-born fathers.

value of the variable in question for the origin. There is significant variation in both cases. Migrant communities are slightly poorer than average in the first generation, with a mean income rank of 48.6, and slightly richer than average in the second generation, with a mean income rank of 51.8. The size of migrant communities varies dramatically, with a 10^{th} percentile population of 300 and a 90^{th} percentile population of 12,000. Importantly, Panels A and B show there is substantial variation in the outcome variables and community size both between and within source countries, allowing me to exploit the panel dimension of the data.

4.3.3 Individual-level intergenerational data

I provide further evidence of a role for culture and social context and explore specific transmission mechanisms using two rich sources of microdata.

First, the Longitudinal Surveys of Australian Youth (LSAY). The LSAY program follows large, nationally representative samples of students from their mid-teens through to their mid-20s. There are six cohorts to date, beginning in 1995, 1998, 2003, 2006, 2009 and 2015. This roughly captures those born in 1980, 1983, 1988, 1991, 1994 and 2000 and hence lines up with the last three cohorts explored in the census data (and beyond). The first waves includes test scores in reading and mathematics and (often) educational aspirations, while later waves capture university entrance scores and educational attainment.[16] Own and parent country birth are also available. Finally, detailed demographic information allow family background to be more precisely controlled for, and school fixed effects to control for institutional setting.

Lastly, I use the Youth in Focus dataset. This dataset includes both administrative data and survey responses and is based on a random sample of all those born between 1 October 1987 and 31 March 1988, and appearing in the records of the Australian Government welfare agency (Centrelink) between 1991 and July 2006. This dataset has own and parent country of birth, and parent educational attainment and income. It also asks parents and children about the importance

[16] Since 2003, the initial survey wave has been integrated with the OECD Programme for International Student Assessment (PISA).

Table 4.1: Variation in outcomes across and within migrant communities

	Mean	SD	p10	p50	p90
Panel A: Between migrant communities — raw data					
First generation outcomes					
Mean income rank	48.6	10.3	33.6	49.8	62.0
Mean education years	11.9	1.3	9.8	12.0	13.4
Number of children ('000)	4.8	10.1	0.3	1.1	12.0
Second generation outcomes					
Mean income rank	51.8	3.9	47.8	51.3	57.2
Mean education years	14.0	0.6	13.3	13.9	14.8
Number of adults ('000)	4.5	8.7	0.3	1.1	12.1
Covariates					
PISA outperformance	0.0	0.3	-0.3	-0.0	0.3
Income penalty	2.0	7.1	-5.9	1.1	12.7
Panel B: Within migrant communities — net of mean for country of origin					
First generation outcomes					
Mean income rank	0	3.5	-3.5	-0.2	4.6
Mean education years	0	0.4	-0.5	-0.0	0.5
Number of children ('000)	0	3.1	-1.3	0.0	1.5
Second generation outcomes					
Mean income rank	0	1.5	-1.7	-0.0	1.7
Mean education years	0	0.3	-0.4	0.0	0.3
Number of adults ('000)	0	2.7	-1.1	-0.0	1.5
Covariates					
PISA outperformance	0	0.0	0.0	0.0	0.0
Income penalty	0	3.5	-4.2	0.1	3.6
N (distinct origin-cohort)			277		
N (distinct origin)			63		

Notes: Presents summary statistics for key outcome variables between and within the sample of migrant communities. Migrant communities are defined by shared paternal country of birth and birth cohort. As described in the text, I drop migrant communities with populations of less than 200 or with a discrepancy of more than 30% in the populations observed in the historical censuses versus the 2016 Census. Panel A presents the statistics for the raw data, whereas Panel B presents first subtracts the mean for the country of origin, to illustrate the variation over time in the sample.

of various mechanisms for upward mobility — or 'how people get ahead in life'.

4.3.4 Cultural and social context proxies

Finally, to investigate source country effects I use quantitative proxies linked to the characteristics of interest. This takes the analysis beyond the 'black box' of simply examining source country fixed effects, which facilitates thinking about both potential transmission mechanisms and omitted variables. It is also much more parsimonious, and in some cases allows variation over time to be exploited.

As a proxy for the cultural value placed on education, I use country-level outperformance on PISA tests in the country of origin. There is a strong relationship between PISA test scores and national income (see Appendix Figure 4.5). To avoid picking up these income effects, I first regress combined average PISA test scores across science, reading and maths on log of gross national income per capita and take the residual as my proxy — effectively the outperformance of the country on these tests relative to other countries with a similar level of income.[17] I do this for each year of available data from 2006 to 2015, assigning to each country its average PISA outperformance. This approach of looking at PISA performance conditional on national income mirrors the exceptional educational mobility I am seeking to explain among second generation migrants — their performance relative to others with similar endowments.[18]

I use the income rank gap between the first generation and similarly educated natives as a measure of the social context. To do this I use the 1% Confidentialised Unit Record Files (CURFs) for the 1981-2006 Censuses.[19] I first calculate the mean income rank for the native fathers at each level of education. I then calculate the predicted first generation income rank for each migrant community had they earned the same income ranks as the native fathers at each level of education. I

[17] I do not further control for national education spending or other factors that frequently enter cross-country education production functions, due to both the proliferation of choices entailed in doing that, and the fact that these factors will in many cases reflect the cultural value placed on education, which I am trying to capture.

[18] An alternative approach would be to use source country measures of the cultural values thought to matter. For example, Figlio et al. (2016) find a measure of long-term orientation is associated with better third grade reading and math scores, and larger test score gains over time, while Mendez and Zamarro (2018) report associations between some non-cognitive skills and education, occupation and wage outcomes. However, the approach in this paper remains agnostic on the question of which of many cultural values may matter.

[19] CURFs were only produced for the 1981 Census onwards.

subtract this from the observed income rank to obtain the income penalty. The summary statistics are in Table 4.1. Across all migrant communities the mean income penalty is 2 percentile rank points, that is, an income rank 2 percentile points below a group of natives with the same education distribution. There is a positive skew — higher income penalties are more common than negative penalties, as few migrant communities earn higher returns than similar educated natives.

4.4 Model

I now describe the model I use to decompose the intergenerational income rank mobility of migrant communities. The aim of this exercise is to better understand what drives the differences in second generation outcomes across migrant communities. I begin with a group-level version of the intergenerational income mobility regression:

$$y_{2,oc} = \alpha_c + \beta_c y_{1,oc} + \varepsilon_{oc} \tag{4.1}$$

where $y_{2,oc}$ and $y_{1,oc}$ are mean own and father income ranks for those born to fathers from origin o and into birth cohort c. Both the intercept and slope are allowed to vary by birth cohort. The residual ε_{oc} captures what I refer to as the 'exceptional' income mobility of each particular migrant community — the extent to which second generation outcomes exceed those expected based on the full suite of migrant communities. An analogous education mobility regression is:

$$e_{2,oc} = a_c + b_c e_{1,oc} + \zeta_{oc} \tag{4.2}$$

where $e_{2,oc}$ and $e_{1,oc}$ are own and father years of education. In this case, ζ_{oc} captures the 'exceptional' education mobility of each particular migrant community. Income and education in both generations can be related to one another via regressions of the form:

$$y_{g,oc} = \rho^0_{g,c} + \rho^1_{g,c} e_{g,oc} + \eta_{g,oc} \tag{4.3}$$

where $g \in \{1, 2\}$ indicates the generation. The resulting coefficients give the returns, in income percentile rank points, to varying levels of education. The residuals $\eta_{g,oc}$ capture the excess returns experienced by individual migrant communities.

To understand what drives the exceptional income mobility of particular migrant communities, it is helpful to express the residual from equation (4.1) in terms of underlying education mobility and the returns to education captured in equations (4.2) and (4.3):

$$\begin{aligned}\varepsilon_{oc} =& \eta_{2,oc} + \rho^1_{2c}\zeta_{oc} - \frac{b_c\rho^1_{2c}}{\rho^1_{1c}}\eta_{1,oc} \\ &+ \left[\left(\rho^0_{2,c} + \rho^1_{2,c}\left(a_c + b_c\frac{y_{1,oc} - \rho^0_{1,c}}{\rho^1_{1,c}}\right)\right) - (\alpha_c + \beta_c y_{1,oc})\right]\end{aligned} \tag{4.4}$$

This decomposition has an intuitive interpretation. A migrant community's exceptional income mobility is simply the sum of their excess return earned in the second generation ($\eta_{2,oc}$), the return to their exceptional education mobility ($\rho^1_{2c}\zeta_{oc}$), the unwinding of their excess return in the first generation ($-\frac{b_c\rho^1_{2c}}{\rho^1_{1c}}\eta_{1,oc}$) and a final bracketed term that is equal for all communities with the same first generation income rank $y_{1,oc}$. The final term simply reflects the difference in predicted second generation income rank generated by either the income mobility relationship or the underlying education mobility and return relationships. The first three terms capture distinct and commonly hypothesised explanations for second generation success or otherwise — namely that the second generation may earn particularly high or low returns in the labour market, get more or less education, or simply appear to do better or worse because of the low or high returns earned by their parents.

In the formulation above, exceptional intergenerational mobility in income (ε_{oc}) and education (ζ_{oc}), and excess returns to education ($\eta_{g,oc}$), are measured relative to expectations based on group-level regressions estimating the underlying relationships across the full sample of migrant communities. An obvious extension is to measure exceptional mobility and returns relative to expectations based on the native population — those born to Australian fathers. This allows for exceptional mobility or returns to accrue to migrant communities as a whole. This is possible if the relationships in equations (4.1)-(4.3) are first estimated or assumed for the native population. I do this drawing the returns to education from Census microdata and estimates of intergenerational income and education mobility from the existing literature and the Household and Income Labour Dynamics in Australia (HILDA) survey. Further details are in Appendix 4.D.

4.5 Results

4.5.1 Decomposing residual income rank mobility

I now turn to the results. In this section I discuss the intergenerational income rank mobility of Australian migrant communities. I then relate this to the underlying intergenerational mobility in education, and the returns to education in the first and second generations.

Figure 4.1: Intergenerational income mobility for Australian migrant communities (1987-91 cohort)

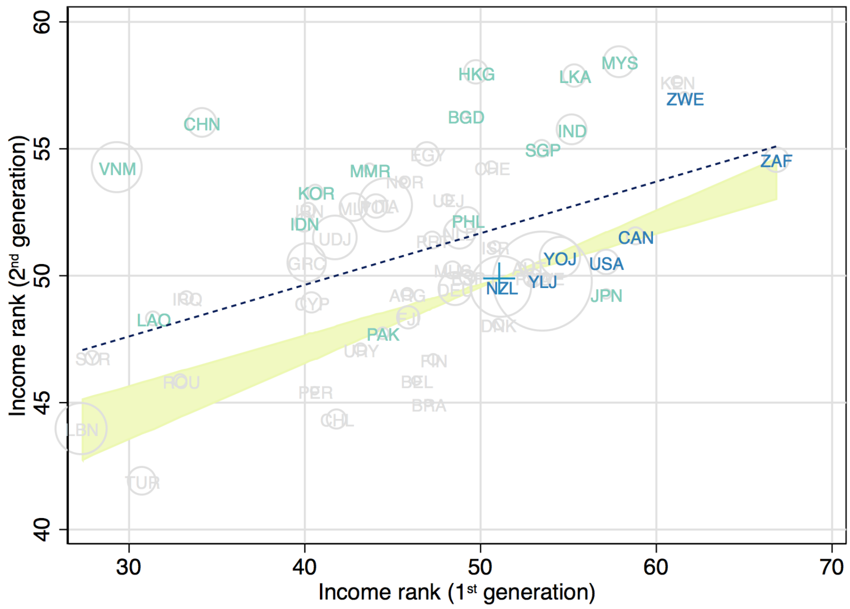

Notes: Plots the average individual income rank of second generation Australians against that of their fathers, by father source country. Circle sizes increase with cell population. Circles are labelled with ISO country codes, with different colours for Asian countries and English-speaking countries. A short-dash line of-best fit is shown, based on an unweighted regression. The cross marks the location of those born to Australian fathers, from which the shaded region extrapolates, assuming a rank-rank correlation between 0.20-0.30.

Figure 4.1 shows the mean income ranks in the first and second generations for all migrant communities in the most recent birth cohort.[20] There is clear intergenerational persistence in the outcomes of Australian migrant communities — higher income ranks in the first generation are associated with higher income ranks in the second generation. Regressing second generation on first generation income rank yields the dashed line-of-best fit, with a rank-rank slope of 0.24. This is similar to recent estimates of rank-rank mobility in the Australian population at large (Murray et al. (forthcoming); Deutscher (2018)).

There are also clear patterns in the outliers — some groups of source countries do better than others. Asian source countries tend to outperform relative to other migrant communities. This includes countries where the first generation is above or well below the mean income rank and of a range of different sizes. Conversely, Latin American source countries often underperform. The English speaking countries also modestly underperform other migrant communities but are, perhaps unsurprisingly, in line with expectations for the native population.

Why then do we see these patterns of second generation success? The framework introduced earlier suggests three potential explanations for a country performing particularly well. It could be that the community in question outperforms in education mobility. Alternatively, it could be that they earned a lower return on first generation education or a higher return on second generation education. The same three mechanisms, working in the opposite direction, can explain a particularly poor performance.

Figure 4.2 shows the underlying relationship between second and first generation education, and the returns to education in each generation. While all graphs show a clear relationship, there are equally clear outlying groups. For example, the Asian source countries are striking in their exceptional educational mobility. The differences in the returns to education are, not surprisingly, most pronounced in the first generation. Some source countries receive lower returns in only the first generation (such as many Asian countries) while for others lower returns persist (Middle Eastern countries). In Appendix Table 4.8 I show the source country effects apparent in these charts tend to persist, as most of the residual variation

[20] That is, for second generation Australians born between 1987 and 1991. The most recent birth cohort is chosen to maximise the number of source countries plotted.

in the intergenerational income mobility, education mobility, or return regressions can be explained with country fixed effects.[21]

[21]This can also be confirmed visually by replicating Figures 4.1 and 4.2 for the other five birth cohorts. Results available on request.

Figure 4.2: Intergenerational education mobility and returns to education for Australian migrant communities (1987-91 cohort)

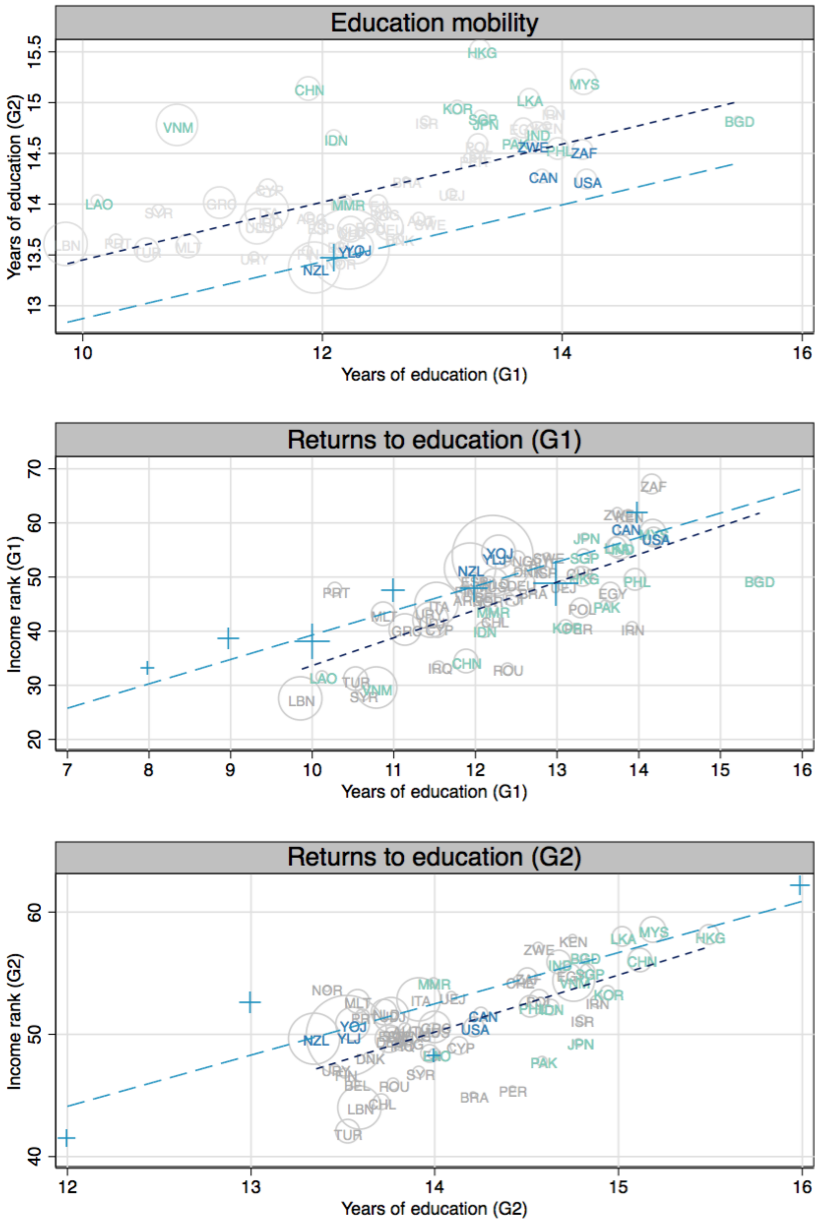

Notes: Plots: the average years of education of second generation Australians against that of their fathers, by father source country (Education mobility); the average father income rank against average father years of education (Returns to education (G1)); or the average individual income rank of second generation Australians against their average years of education (Returns to education (G2)). Circles are labelled with ISO country codes, with different colours for Asian countries and English-speaking countries. Short-dash lines-of-best fit are shown, based on unweighted regressions. The cross marks the location of those born to Australian fathers in the same group-level dataset (Education mobility) or in associated individual level data (Returns to education). The long-dash lines present an indicative benchmark for those born to Australian fathers based on the data presented or, as discussed in the text, the Household and Income Labour Dynamics in Australia (HILDA) survey (Education mobility).

Finally, in Figure 4.3 I explicitly decompose the residual intergenerational income mobility of migrant communities using the approach outlined in Section 4.4. To simplify the presentation, I collapse the decompositions to regional groupings, weighting each individual migrant community's decomposition by population size.[22] Those born to Australian fathers are in the row shaded grey and serve as a useful comparison group. There are a number of clear observations from the decomposition. The exceptional mobility of the Asian source countries is predominantly due to their exceptional education mobility. Relatively poor education is also the leading contributor to eight of the ten communities with the poorest income mobility — the Middle Eastern and South American countries the exceptions, where relatively poor second generation returns are more important.

[22]I use the second level of the Australian Bureau of Statistics' Standard Australian Classification of Countries (SACC) Australian Bureau of Statistics (2016). The SACC is essentially based on geographic proximity, grouping neighbouring countries into progressively larger areas on the basis of similarity in terms of social, cultural, economic and geopolitical characteristics. Some regional groupings (e.g. Central America) are not represented in the dataset.

Figure 4.3: Decomposition of exceptional income mobility

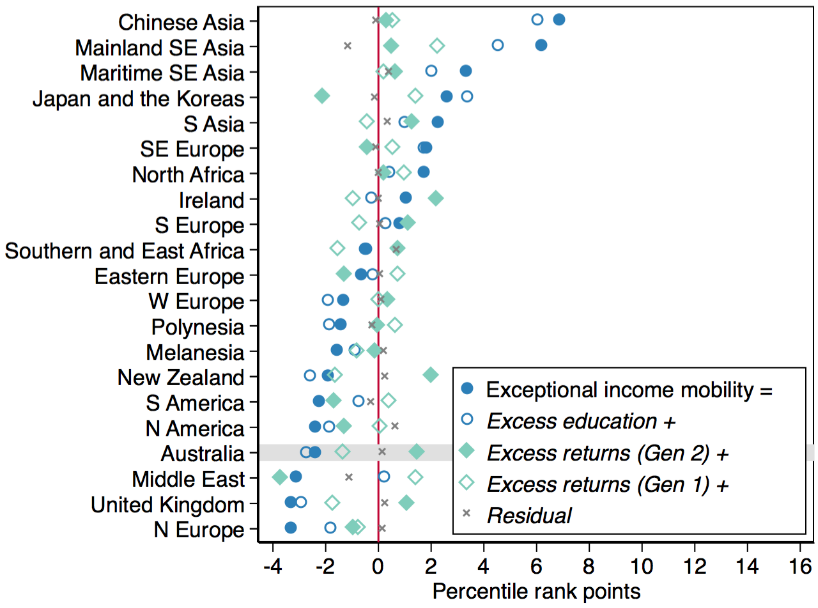

Notes: Decomposes exceptional intergenerational income mobility — the residual from a regression of average second generation individual income rank on average father income rank — into four parts as in equation 4.4. The decomposition consists of components attributable to: excess educational mobility (defined as for exceptional income mobility); excess returns in first or the second generations, and a residual component. The decomposition is performed for individual migrant communities, with the resulting decomposition then aggregated into regions (weighting according to population size) to aid visualisation. The decompositions within the regional groupings are typically quite similar.

There are notable differences and similarities in the decomposition when exceptional mobility and returns are measured relative to the native population. In Appendix Figure 4.6 I present the same decomposition based on reasonable lower and upper bounds on the income rank-rank correlation for the underlying population. Exceptional education mobility remains the dominant contributor to exceptional upward income mobility, but across all countries excess returns now play a more prominent role as well. This reflects the fact that the deviations of migrant communities from underlying native experiences tend to be less idiosyncratic and more of a shared migrant experience when it comes to their labour market returns (compared to their intergenerational mobility). For example, for Mainland SE Asia, the unwinding of below average returns received in the first generation contributes around 3.4 percentile rank points to the exceptional mobility of the second generation. For the United Kingdom, the unwinding of above average returns detracts around 1.4 percentile rank points. These highlight two extremes of the migrant experience. The Mainland SE Asian countries — Cambodia, Laos, Myanmar, Thailand, Vietnam — have been sources of large refugee outflows. It is not surprising that the first generation earned lower returns on their education than the native population. In contrast, migration from the United Kingdom is more likely to be in response to specific economic opportunities. Past research has hypothesised and found evidence that migrants between similar countries may display "negative assimilation", with their relative earnings declining as the economic rents that may have prompted their migration dissipate (Chiswick and Miller (2011, 2012)). The results in this paper suggest an intergenerational variant of this negative assimilation, whereby economic rents earned only by the first generation lead to reduced intergenerational income mobility in the second generation.

4.5.2 The sources of residual educational mobility

Educational mobility plays a significant role in the income mobility of migrant communities. In this section I explore a wide variety of possible explanations. Why is it that migrant communities typically get more education than expected for natives? And why do migrant communities differ so much among one another? To explore these questions, I add additional variables to an otherwise standard intergenerational educational mobility regression:

$$e_{2,oc} = a + be_{1,oc} + \gamma X_{oc} + \delta M_o + \zeta_{oc} \tag{4.5}$$

where $e_{2,oc}$ and $e_{1,oc}$ are mean own and father years of education for those born to fathers from origin o into birth cohort c. All specifications include an indicator variable M_o flagging the recent migrant communities, those whose fathers were born in a country other than Australia. A range of other covariates, corresponding to potential explanations for the educational mobility of migrants, are in the matrix X_{oc} — these are discussed as they are introduced. I estimate robust standard errors, clustered by country of origin.[23]

Table 4.2 presents the results. In column (1) I run the straight intergenerational educational regression. The coefficient on first generation years of education of 0.32 is broadly similar to the 0.28 seen in the native population.[24] The R^2 for the regression is 0.47 — much of the variation in second generation education between migrant communities is explained by first generation education levels. On average, migrants end up with half a year more education than would be expected of those born to similarly educated Australian fathers. In column (2) I add country of origin fixed effects, which lifts the R^2 to 0.87. The fact that so little variation remains indicates a high degree of persistence in exceptional education mobility by country of origin.

What is it about particular migrant communities that explains these outcomes? In column (3) I add the proxies for culture and the migration context to the initial specification in column (1). The R^2 now jumps up to 0.60, explaining around 24% of the residual variation from column (1). The coefficients on both the culture and context proxies are both sizeable and statistically significant. Being born to a father from a source country that outperforms on PISA by a full standard deviation — roughly the gap between Vietnam and New Zealand, or New Zealand and Lebanon — is associated with an additional 0.5 years of education. Being born to a father from a community facing an income penalty that is 10 percentile rank points higher — roughly the gap between Vietnam and Italy, or Italy and England — is associated with an additional 0.3 years of education.

In column (4) I explore explanations based on ethnic networks. For example, it may be beneficial to be from a larger and more established migrant community (as suggested in Gang and Zimmermann (2000)). I thus add the natural logarithm of migrant community size to the regression. There may also be human capital

[23] In those specifications that include the PISA outperformance measure an additional concern is inconsistency in the standard errors due to a failure to account for the presence of a generated regressor (Pagan (1984); Murphy and Topel (1985)). However, bootstrapping standard errors to account for this did not change the conclusions in this or subsequent sections.

[24] This need not have been the case, but may reflect a degree of balance between forces that would be expected to lead to higher and lower persistence among migrant communities. It may also be that the factors that lead migrants to do better or worse in the second generation than natives with similarly educated fathers are relatively uncorrelated with their father's years of education.

Table 4.2: Explanations of exceptional education mobility

	Baseline	Country FE	Culture and context controls	Ethnic network controls	Regional and Country FE			Cohort varying slope and intercept
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Years of education (G1)	0.324***	0.352***	0.372***	0.348***	0.279***	0.370***	0.440***	0.292***
	(0.045)	(0.053)	(0.056)	(0.051)	(0.035)	(0.087)	(0.081)	(0.077)
PISA outperformance (‘culture’)			0.466**	0.488**	0.086	0.082		0.492**
			(0.222)	(0.206)	(0.114)	(0.157)		(0.220)
Income penalty (‘context’)			0.031***	0.025***	0.017***	0.014*	0.025**	0.017**
			(0.007)	(0.007)	(0.006)	(0.007)	(0.010)	(0.008)
Years of education (G1) * NESDC				0.029**	0.010	-0.061	-0.118	0.034***
				(0.012)	(0.007)	(0.103)	(0.098)	(0.012)
Log community size				-0.007	-0.011	-0.020	-0.030	-0.001
				(0.045)	(0.023)	(0.030)	(0.075)	(0.046)
Migrant	0.508***		0.411***	0.111	0.166			0.123
	(0.066)		(0.072)	(0.197)	(0.109)			(0.204)
Specification:								
No. of regional FE	0	45	0	0	7	16	45	0
Country fixed effects		X					X	
R^2	0.470	0.868	0.595	0.631	0.803	0.811	0.884	0.696
$adj.R^2$	0.465	0.831	0.587	0.620	0.790	0.791	0.849	0.670
N	209	209	209	209	209	209	209	209

Notes: Results from group-level regressions of second generation mean years of education on first generation mean years of education, with an indicator variable for those born to overseas born fathers, column (1). Additional covariates are added in columns (2)-(8) as follows. Column (2) adds country of origin fixed effects. Column (3) instead controls for culture and context through the PISA outperformance and income penalty variables described in the text. Column (4) adds further controls for ethnic network effects, including the community size and allowing the return to first generation education to differ for those from non-English speaking backgrounds. Columns (5)-(7) add increasingly fine region of origin fixed effects to the specification in column (4) taken from the Standard Australian Classification of Countries (Australian Bureau of Statistics (2016)) — the 1 digit level (column (5)), 2 digit level (column (6)) and country of origin fixed effects (column (7)). Column (8) instead modifies the specification in column (4) by allowing cohort specific intercepts and coefficients on first generation education. Robust standard errors are clustered at a country of origin level. Significance levels are indicated as follows: *** $p < 0.01$; ** $p < 0.05$; and * $p < 0.1$.

externalities, whereby individual education is influenced not just by their paternal education, but the average paternal education of their community. This is the formulation of ethnic capital originally introduced by Borjas (1992). For example, suppose the true *individual-level* education mobility relationship is:

$$e_{2,i,oc} = \tilde{a} + \tilde{b}e_{1,i,oc} + \tilde{\gamma}e_{1,oc} \tag{4.6}$$

where $\tilde{\gamma}$ is a human capital externality. Averaging this over individuals i within migrant communities oc would yield a group-level relationship:

$$e_{2,oc} = \tilde{a} + (\tilde{b} + \tilde{\gamma})e_{1,oc} \tag{4.7}$$

Estimating $\tilde{\gamma}$ is straightforward in individual-level data, but ordinarily impossible in grouped data. One potential test for the presence of such effects can be conducted if we assume that human capital externalities are negligible ($\tilde{\gamma} \approx 0$) for some communities. For example, it seems far less likely that the average paternal education of the first generation matters for second generation migrants from English-speaking developed countries. Thus, I include the interaction of $NESDC_o$ — an indicator variable for whether the origin is not an English-speaking developed country — with average paternal education to see if it matters more for countries where human capital externalities seem more plausible. There are, of course, a wide range of interpretations to a positive coefficient. Peer, role model or network effects are all possible, but so to are mechanisms arising from statistical discrimination or simply omitted variables. The results are at least consistent with modest human capital externalities, with an increase of one year in the mean years of education of the fathers associated with a 0.029 (s.e. 0.012) increase in child years of education. In contrast, the size of an individual's migrant community is not associated with second generation outcomes.[25]

In columns (5)-(7) I take the baseline specification from column (4) and add fixed effects at increasingly fine geographic levels, from region through to country of origin. The association between PISA outperformance and second generation educational attainment disappears. In contrast the association between the income

[25] This is also true if migrant community size is included without taking logs.

penalty and education remains significantly positive. Finally, in column (8) I allow both the intercept and slope of the intergenerational education relationship to vary by cohort. The associations remain, and are only partially attenuated in the case of the income penalty.

What might drive the association between our proxies for culture and social context and second generation education? Further analysis is needed if we are to rule out some of the more obvious omitted variable explanations, and read underlying causal mechanisms into these associations.

For the cultural proxy, it may be that migrants bring more than their culture with them to Australia. For example, migrants may bring something of their origin country institutions with them if they actively seek out schools and communities that reflect the environment they left behind. This is a potential concern with the epidemiological approach — institutional and environmental factors may be more portable than they appear. More generally, PISA outperformance may be correlated with unobserved human capital that is not reflected in first generation education or income ranks. More detailed family background and institutional controls are possible in the individual-level data.

For the social context proxy, a number of credible mechanisms may give rise to a positive association between income penalty and second generation educational attainment. It may be that first generation migrants who face larger income penalties are more likely to be motivated by altruistic concerns, and thus invest more in their children — a selection mechanism. Alternatively, such migrants may come from countries where educational opportunities were limited, and hence the parents posses and pass on to their children higher unobserved ability — a simple omitted variable. Explicitly controlling for parental altruism or latent ability is not possible in the data. In both these cases, however, we may expect to see this association emerging early in an individual's schooling, something I am able to test in the following section.

4.5.3 Individual-level data, validity and transmission mechanisms

Motivated by the above discussion, I now use test score and survey data to better control for omitted variables and selection, explore transmission mechanisms. I begin with the Longitudinal Surveys of Australian Youth (LSAY). I estimate individual-level regressions of the form:

$$y_{i,ocs} = \alpha_{cs} + \beta X_i + \gamma Z_{os} + \delta M_o + \eta W_o + \varepsilon_{i,ocs} \tag{4.8}$$

where $y_{i,ocs}$ is an education outcome for individual i from school s, survey cohort c and born in Australia to a father born in an origin country o. In X_i I include an individual's: sex; age and its square; and years of education and occupational scores for both parents. In Z_{os} I include the mean years of education for parents for all individuals in the same school s and born to fathers born in origin country o. The controls in Z_{os} may pick up peer or role model effects, but may also reflect own-parent ability in the likely event that parent education and occupation are reported with some error. I also allow for survey cohort and school fixed effects α_{cs}. The variable M_o simply indicates those with overseas born fathers, picking up a generic second generation effect. Finally, the key independent variables, proxying for culture and context, are in W_o. Since the culture proxy varies only by origin, I cluster standard errors by origin o.

I begin with a visual overview of the results. Figure 4.4 shows the association between PISA outperformance, the income penalty, and normalised educational outcomes at various ages. Two features are of note. First, the associations between PISA outperformance ('culture') and outcomes falls with age, while the associations between the income penalty ('context') and outcomes rise. Second, while the associations between PISA outperformance and outcomes falls when controlling for ability, the associations between the income penalty and outcomes do not.

Detailed results are in Table 4.3. In columns (1) and (2) of Table 4.3 I report the results where the independent variables are standardised maths and reading scores in Year 9 respectively. The coefficient on PISA outperformance is positive

Figure 4.4: Associations between PISA outperformance and income penalty and normalised educational outcomes at various ages

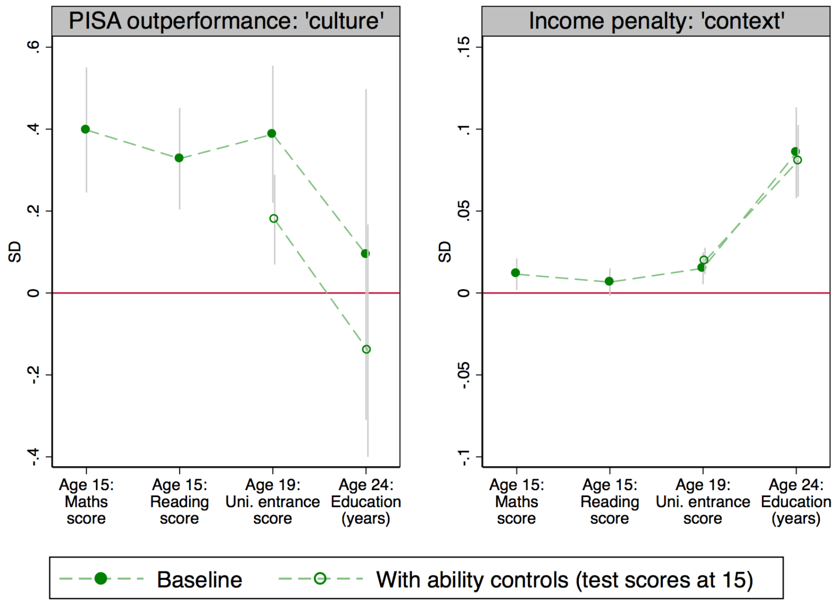

Notes: Presents the associations from Table 4.3 between PISA outperformance or the income penalty and education outcomes, normalised here to have standard deviation one. The associations are simply the coefficients from individual-level regressions of Year 9 math scores, reading scores, university entrance scores and years of education on both variables and other covariates. These other covariates include an individual's: sex; age and its square; years of education and occupational scores for both parents; the mean years of education for parents for all individuals in the same school born to fathers born in the same origin; an indicator for an overseas born father; and cohort and school fixed effects. For the hollow dots, Year 9 math and reading scores are also included as covariates. Confidence intervals (95%) are based on robust standard errors, accounting for clustering by paternal country of birth.

and strongly significant. The PISA outperformance variable is calculated in such a way that a one unit increase is roughly equivalent to an increase of one standard deviation on the PISA tests themselves. Thus individuals with fathers from countries that outperform on PISA by one standard deviation, themselves outperform by around a third of a standard deviation in maths and reading test scores at age 15 in Australia. In contrast the associations between the income penalty and test scores are smaller and only statistically significant for maths scores. In columns (3) and (4) I consider university entrance scores (a percentile rank) as the dependent variable. Once again, PISA outperformance is strongly associated with higher scores. However, a higher income penalty is now strongly positively associated with student achievement. And whereas the coefficient on PISA outperformance more than halves once conditioning on Year 9 test scores, the coefficient on the income penalty rises. Finally, I consider educational attainment. In this case, PISA outperformance loses its explanatory power, while the income penalty is strongly associated with educational attainment, and again very little of the association is explained by test scores.[26]

In Table 4.3 I use the survey weights to account for the heavy attrition between waves in which the successive outcomes are measured. Nonetheless, it is plausible that the differences across outcomes are driven by inadequacies in this method of controlling for selective attrition. Table 4.9 largely puts these concerns to rest. In columns (1) and (2) I replicate the specification from the same columns in Table 4.3, but using the much smaller sample of individuals surveyed in Wave 10 and with non-missing educational attainment. The coefficients on the key independent variables are broadly similar in size and significance. In columns (3)-(6) I examine the years of education based on either the child's reporting in Wave 1 of their own plans or their parents' plans for their education. Foreshadowing the associations with eventual educational attainment, a higher income penalty is associated with higher aspirations while PISA outperformance has no significant positive bearing on aspirations, once controlling for test scores.

The joint analysis of school test scores, educational aspirations and attainment points to the culture and context proxies operating through two different channels. The culture proxy — PISA outperformance — is most strongly associated with achievement, and particularly conditional on this, has relatively little bearing on educational attainment. It may be that the influence of source country fades with time in the destination country, or that the cultural proxy used here isolates beliefs and practices that matter more for educational achievement than attainment. In contrast, the context proxy operates primarily through university entrance scores and educational attainment. This is an important restriction. If the earlier association between income penalty and educational attainment in the

[26]The fact that second generation migrants from many source countries improve their relative performance through their educational career has been noted previously (for example, by Marks (2010) in the Australian context). The findings here tie that more closely to the context of migration.

Table 4.3: Influence on culture and context on individual youth achievements

	Age 15		Age 19		Age 24	
	Maths score	Reading score	Uni. score		Years of education	
	(1)	(2)	(3)	(4)	(5)	(6)
PISA outperformance (‘culture’)	0.39***	0.32***	6.52***	3.01***	0.09	-0.14
	(0.08)	(0.06)	(1.44)	(0.94)	(0.21)	(0.16)
Income penalty (‘context’)	0.01**	0.01	0.25***	0.33***	0.09***	0.08***
	(0.00)	(0.00)	(0.08)	(0.07)	(0.01)	(0.01)
Migrant	0.05**	0.04**	1.16**	0.81*	0.05	0.05
	(0.02)	(0.02)	(0.51)	(0.41)	(0.10)	(0.08)
Female	-0.14***	0.32***	2.65***	3.71***	0.73***	0.66***
	(0.00)	(0.01)	(0.23)	(0.24)	(0.02)	(0.03)
Father’s:						
Education (years)	0.03***	0.03***	0.41***	0.25***	0.09***	0.07***
	(0.00)	(0.00)	(0.05)	(0.03)	(0.00)	(0.01)
Occupation score	0.00***	0.00***	0.08***	0.05***	0.01***	0.00***
	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)
Mother’s:						
Education (years)	0.02***	0.02***	0.33***	0.14***	0.05***	0.04***
	(0.00)	(0.00)	(0.07)	(0.04)	(0.00)	(0.00)
Occupation score	0.00***	0.00***	0.05***	0.04***	0.01***	0.00***
	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)
Mean peer parent education (years):						
Fathers	0.08***	0.07***	1.11***	0.61*	0.09	0.04
	(0.02)	(0.02)	(0.39)	(0.34)	(0.06)	(0.04)
Mothers	0.06**	0.05**	1.02**	0.65*	0.14***	0.08**
	(0.02)	(0.02)	(0.48)	(0.37)	(0.05)	(0.04)
Ability controls						
Y9 math score				5.96***		0.44***
				(0.14)		(0.02)
Y9 reading score				4.90***		0.38***
				(0.23)		(0.04)
R^2	0.18	0.19	0.21	0.41	0.20	0.29
$adj.R^2$	0.17	0.18	0.18	0.38	0.16	0.26
N	44,556	44,561	9,945	9,902	9,333	9,276

Notes: Results from individual-level regressions of Year 9 math scores (column (1)), reading scores (column (2)), university entrance scores (columns (3) and (4)) and years of education (columns (5) and (6)). A quadratic in age, and school and cohort fixed effects, are included in all specifications. To capture the later life outcomes, the regressions in columns (3) and (4), and (5) and (6), are restricted to those observed in Waves 5 and 10 respectively and weighted using survey-supplied weights to adjust for attrition. Robust standard errors are reported in parentheses, accounting for clustering by paternal country of birth. Significance levels are indicated as follows: *** $p < 0.01$; ** $p < 0.05$; and * $p < 0.1$.

census data were driven by greater unobserved parental ability it would likely show up in an association with test scores at much younger ages. Similarly, if it were driven by greater parental altruism and consequently higher parental investments in children, that too might be expected to show up by age 15. Instead, it seems the association is mostly driven by factors that influence late educational achievement and attainment. The educational aspirations reported in Table 4.9 are one such factor.

As a final test of the robustness of the associations between the cultural proxy and scholastic achievement, and the income penalty and educational attainment — I examine whether they persist when considering variation within major source regions. In Table 4.4 I repeat the individual-level regressions as specified in equation (4.8) with either the combined maths and reading score, or years of education as the dependent variable. This time I also include dummy variables for migrant paternal region of birth (Europe, Asia or Other) and the interactions of these with the culture or context proxy. Comfortingly, the coefficients on the culture and context proxies remain jointly significant and similar in magnitude when they are allowed to vary by region of origin.

4.5.4 Adolescent econometricans?

Why do some families aspire to more education than others? One answer may be that they anticipate, rightly or wrongly, greater returns. As stressed by Manski (1993), adolescents attempting to estimate the returns to education, and plan accordingly, face much the same problem as econometricians. Past empirical work has indeed found expected returns to be influential in the schooling choices of adolescents, including in the Australian context (Wilson et al. (2005); Roussel (2004)). The same is true for parents considering the education of their children. Faced with the decisions and outcomes of past generations, and potentially unable to observe ability, both parents and children try to infer the return to education. It should be unsurprising that those confronted with different data may reach different conclusions, just as econometricians do. This is a potential explanation for the results. Migrant communities with low education levels and receiving higher income penalties in the first generation may well infer higher returns to education

Table 4.4: Influence on culture and context on outcomes, within paternal regions of origin

	Achievement: Combined test score		Attainment: Years of education	
	(1)	(2)	(3)	(4)
PISA outperformance ('culture')				-0.21
				(0.45)
X Europe	0.24*	0.21*		
	(0.12)	(0.12)		
X Asia	0.25***	0.24***		
	(0.07)	(0.05)		
X Other	0.24***	0.26***		
	(0.09)	(0.09)		
Income penalty ('context')		0.00		
		(0.00)		
X Europe			0.09***	0.09***
			(0.01)	(0.01)
X Asia			0.06**	0.09
			(0.02)	(0.06)
X Other			0.07***	0.08***
			(0.02)	(0.03)
p-value on test that culture (1) and (2) or context (3) and (4) proxies are:				
jointly zero	0.00	0.00	0.00	0.00
jointly equal	1.00	0.93	0.32	0.71
R^2	0.19	0.19	0.20	0.20
$adj.R^2$	0.17	0.17	0.17	0.17
N	44,857	44,473	9,571	9,333

Notes: All specifications include: a quadratic in age; father and mother years of education and occupational scores; mean peer father and mother years of education; school, cohort and region of origin fixed effects. The regressions in columns (3) and (4) are restricted to those observed in Wave 10 and weighted using survey-supplied weights to adjust for attrition. Robust standard errors are reported in parentheses, accounting for clustering by paternal country of birth. Significance levels are indicated as follows: *** $p < 0.01$; ** $p < 0.05$; and * $p < 0.1$.

than the population at large, and adjust their aspirations accordingly.

I present tentative evidence for this channel using the Youth in Focus dataset. Beyond rich demographic information, this dataset includes youth and parent assessments of the important of various mechanisms for 'getting ahead in life'. In almost all cases the surveyed parent was the youth's natural mother, and I drop from the analysis the rare exceptions to this. The precise wording of the questions to parents was as follows:

> *Now we have some questions about how people get ahead in life generally. For each question, I would like you to tell me whether it is 'extremely important', 'fairly important', 'not too important', 'does not matter at all' or 'undesirable, a bad thing'.*
>
> - *In order to get ahead in life, how important is it to come from a wealthy family?*
> - *To get ahead in life, how important is it to have well-educated parents?*
> - *How important is it for a person to have a good education?*
> - *How important is a person's own ambition?*
> - *How important is it for a person to have a job?*

The same questions were asked of youth, with minor modifications. The 'extremely important' option was replaced by 'very important' and the 'undesirable, a bad thing' option was dropped, and youth were not asked about the importance of coming from a wealthy family. Interesting raw differences in responses to these questions have been noted in past research — Cobb-Clark and Nguyen (2012) note that immigrant mothers from non-English speaking backgrounds rate both own and paternal education as more important than immigrant mothers from English-speaking backgrounds and Australian-born mothers. I take this analysis further by controlling for other factors that may influence responses and examining how responses vary by more finely grained source country characteristics (the cultural and contextual proxies). To analyse responses, I estimate individual-level ordered logistic regressions of the form:

$$y_{i,o} = \alpha + \beta X_i + \eta W_o + \varepsilon_{i,o} \tag{4.9}$$

where y_i is the response given by an individual i born in Australia to a mother born in an origin country o. Alternatively, it may be the response given by their mother.

In X_i I include an individual's: sex; years of education for both parents; maternal age and its square; and household earnings. The variable M_o simply indicates those with overseas born mothers, picking up a generic second generation effect. Finally, the key independent variables — the culture and context proxies — are in W_o.[27]

Table 4.5 presents the results for parent (Panel A) and child (Panel B) responses. For the parents, there is only one association between the culture and context proxies and the perceived importance of various factors to 'getting ahead' that is significant at the 5 per cent level. A higher income penalty is associated with parents placing more importance on education. This finding is tempered a little when looking at children, for whom a higher income penalty does not carry any equally strong associations. Interestingly, however, both young women and their parents tend to place more importance on an individual's education. This could also be consistent with an income penalty, in this case for a gender rather than a migrant community, influencing the perceived return to education and flowing through to aspirations and attainment. In both Tables 4.3 and 4.9, young women aspire to and attain more years of education.

4.6 Conclusion

The size and salience of migrant communities in advanced economies makes their intergenerational mobility an important issue — what drives second generation success or otherwise? This paper adds to the large literature on second generation outcomes (Sweetman and van Ours (2015)). I exploit the large, diverse and long-standing migrant population available in the Australian setting, and a suite of data sources. I explore why some migrant communities do better or worse in second generation outcomes than might be expected based on first generation outcomes.

I begin with a simple decomposition of exceptional intergenerational income mobility. This decomposes the residual from a standard intergenerational income mobility regression into components that include the contributions of: excess re-

[27] I no longer cluster standard errors by origin o, as this results in smaller standard errors, perhaps reflecting the relatively small number of clusters, including many with very small populations.

Table 4.5: Perceived importance of various factors to 'getting ahead' regressed on explanatory variables

	(1)	(2)	(3)	(4)	(5)
	Importance of parent...		Importance of individual...		
	...wealth	...education	...education	...ambition	...job
Panel A: Parent responses					
PISA outperformance ('culture')	0.64	0.42	-0.50	-0.41	-0.72
	(0.48)	(0.49)	(0.56)	(0.47)	(0.52)
Income penalty ('context')	-0.01	0.03	0.07**	-0.02	-0.00
	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Migrant	0.14	0.25	-0.07	-0.30	-0.27
	(0.17)	(0.17)	(0.18)	(0.21)	(0.22)
Female	-0.03	0.07	0.30**	0.10	-0.11
	(0.11)	(0.11)	(0.12)	(0.14)	(0.15)
N	1,166	1,171	1,170	1,173	1,173
Panel B: Child responses					
PISA outperformance ('culture')		0.05	-0.10	-0.26	-0.23
		(0.41)	(0.43)	(0.50)	(0.42)
Income penalty ('context')		-0.02	0.00	-0.05*	-0.02
		(0.02)	(0.02)	(0.03)	(0.02)
Migrant		0.12	0.19	0.64**	-0.20
		(0.17)	(0.18)	(0.25)	(0.18)
Female		-0.17	0.32***	0.29**	-0.30**
		(0.11)	(0.12)	(0.14)	(0.12)
N		1,175	1,172	1,175	1,174

Notes: Results from individual-level ordered logistic regressions of the perceived importance of various factors to getting ahead on culture and context proxies, and migrant and female dummies. Controls for maternal and paternal years of education and log household earnings are also included in the regressions, coefficients available on request. Naive standard errors are reported in parentheses. Significance levels are indicated as follows: *** $p < 0.01$; ** $p < 0.05$; and * $p < 0.1$.

turns in the second generation; the unwinding of excess returns in the first generation; and exceptional educational mobility.[28] There is significant heterogeneity across migrant communities in both their exceptional income mobility and the relative contributions of these factors — very different pictures of migrant mobility emerge across regions of origin. Nonetheless, exceptional education mobility emerges as driving differences between migrant communities, and explaining the striking upward mobility of Asian Australians.

What then drives exceptional education mobility? I find some evidence for a role for culture using an 'epidemiological approach' (Fernández (2011)). Sec-

[28]This decomposition, or variants, could readily be applied to individuals, families or other groups.

ond generation Australians from countries that outperform on PISA end up with more education themselves. An even stronger relationship emerges when looking at individual-level data on test scores in Australia that allow for much more detailed controls for parental background, and the peer and school environment. This weighs against a potential concern with the methodological approach — that migrants may bring not just their 'culture' with them when migrating to Australia, but also a tendency to choose institutions and environments that mimic those in their origin. Lasting effects of culture on educational attainment are, however, not apparent in the individual-level data.

I also find evidence that the context of migration matters, and indeed that it may matter more for educational attainment than culture. Second generation Australians from migrant communities that experience a higher income penalty in the first generation (relative to similarly educated natives) end up with more education, on average. This effect is apparent at an individual level in educational aspirations and attainment — but not in test scores. In survey data, parents from these relatively poorer communities view a person's education as more important to 'getting ahead' than those from richer communities. Young women — who also experience a much studied wage gap — and their parents also place more importance on education. These findings would be consistent with an immigrant family's social context influencing the value they perceive in education, and hence their educational aspirations and the attainment of their children.

This paper thus provides fresh evidence on the empirical importance of education to upward mobility, and the potential roles of culture and social context in influencing educational outcomes. Variations in culture and social context are not purely between migrant communities, and further exploration of the way in which values placed on education, the value perceived in it and their influence on educational achievement, aspirations and attainment would be well worthwhile.

4.A Additional charts

Figure 4.5: Relationship between PISA test scores and national income

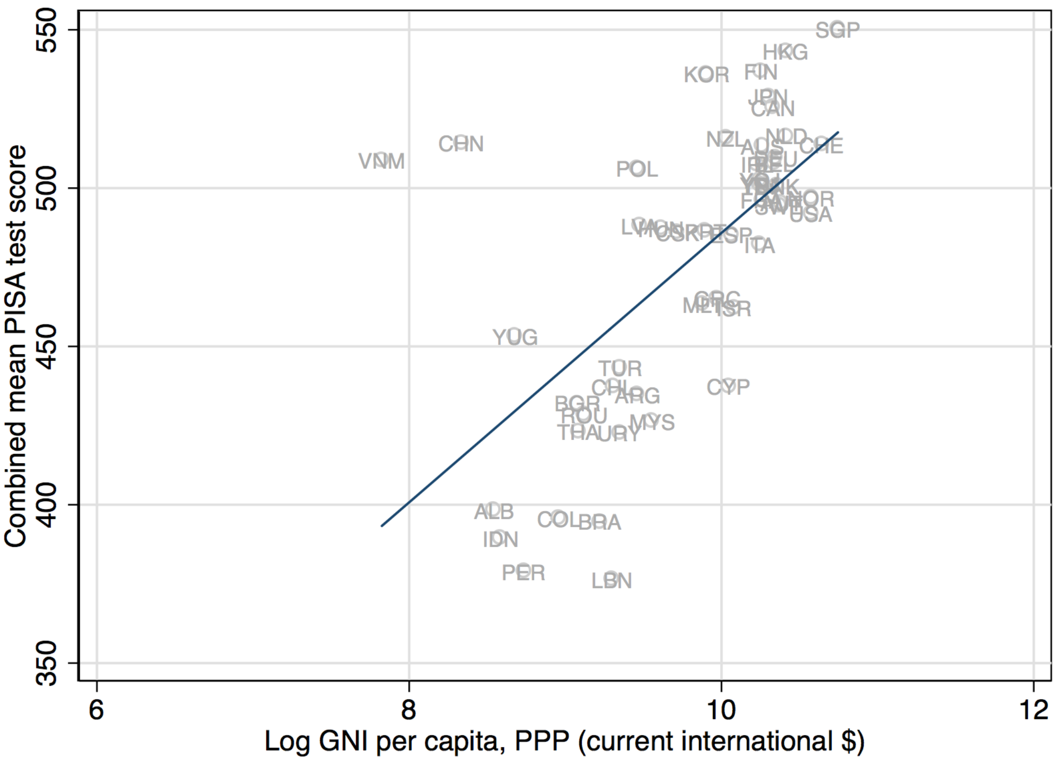

Notes: Illustrates the relationship between the national mean PISA test scores, averaged across mathematics, reading and science, and log of gross national income per capita, in purchasing power parity terms (current international dollars). For countries observed more than once over 2006, 2009, 2012 and 2015, the average values are plotted.

Figure 4.6: Decomposition of exceptional income mobility

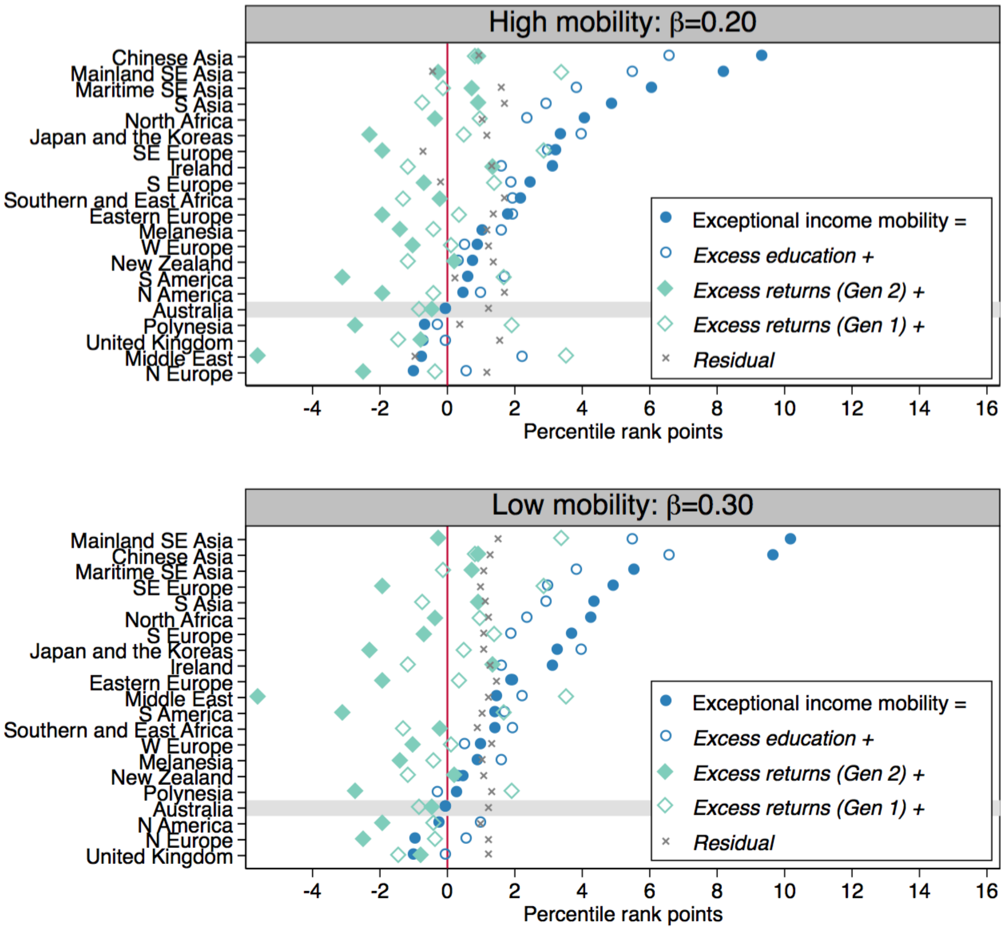

Notes: Decomposes exceptional intergenerational income mobility — the residual from a regression of average second generation individual income rank on average father income rank — into four parts as in equation 4.4. The decomposition consists of components attributable to: excess educational mobility (defined as for exceptional income mobility); excess returns in first or the second generations and a residual component. The decomposition is performed for individual migrant communities defined by country of origin and birth cohort, with the resulting components then aggregated into regions (weighting according to population size) to aid visualisation.

4.B Additional tables

Table 4.6: List of countries, second generation population and first generation income (1987-91 birth cohort)

Country	ISO Alpha-3	Region	Subregion	Population	First generation income rank
Australia	AUS	Oceania And Antarctica	Australia	747,419	51.1
New Zealand	NZL		New Zealand	17,751	51.2
Papua New Guinea	PNG		Melanesia	1,471	53
Cook Islands	COK		Polynesia	283	35.5
Fiji	FJI			2,233	45.8
Tonga	TON			1,269	28.8
Samoa	WSM			966	29.4
England	YLJ	North-West Europe	United Kingdom	46,654	53.5
Northern Ireland	YNJ			765	55
Scotland	YOJ			7,660	54.5
Wales	YPJ			1,892	57.4
Ireland	IRL		Ireland	3,415	54.8
Austria	AUT		Western Europe	858	52.7
Belgium	BEL			268	46.9
Switzerland	CHE			628	50.8
Germany	DEU			5,044	48.6
France	FRA			1,040	45.9
Netherlands	NLD			4,574	48.8
Denmark	DNK		Northern Europe	542	50.7
Finland	FIN			547	47.3
Norway	NOR			118	46.6
Sweden	SWE			276	53.4
Spain	ESP	Southern And Eastern Europe	Southern Europe	1,049	49.2
Italy	ITA			13,500	44.6
Malta	MLT			3,635	42.8
Portugal	PRT			1,311	47.3
Albania	ALB		South Eastern Europe	72	24.4
Bulgaria	BGR			56	46.2
Cyprus	CYP			1,863	40.4
Greece	GRC			6,778	40.1
Romania	ROU			776	32.9
Yugoslavia	UDJ			8,882	41.6
Hungary	HUN		Eastern Europe	1,056	45.4
Latvia	LVA			57	93
Poland	POL			2,536	44.1
Czechoslovakia	UEJ			486	48.3
Ukraine	UKR			151	52.8
Egypt	EGY	North Africa And The Middle East	North Africa	2,572	47
Iran	IRN		Middle East	911	40.2
Iraq	IRQ			750	33.5
Israel	ISR			461	51
Lebanon	LBN			12,576	27.4
Syria	SYR			858	28
Turkey	TUR			3,667	30.8
Cambodia	KHM	South-East Asia	Mainland South-East Asia	1,938	26
Laos	LAO			883	31.4
Myanmar	MMR			510	43.7
Thailand	THA			346	44.2
Vietnam	VNM			11,947	29.4
Indonesia	IDN		Maritime South-East Asia	1,551	39.9
Malaysia	MYS			4,468	57.9
Philippines	PHL			3,118	49.3
Singapore	SGP			1,322	53.5
China	CHN	North-East Asia	Chinese Asia	3,787	34.2
Hong Kong	HKG			2,598	49.8
Taiwan	TWN			303	36
Japan	JPN		Japan and the Koreas	417	57.3
South Korea	KOR			1,050	40.7
Bangladesh	BGD	Southern And Central Asia	Southern Asia	173	49.7
India	IND			4,002	55.2
Sri Lanka	LKA			2,400	55.4
Pakistan	PAK			561	44.7
Canada	CAN	Americas	Northern America	1,335	58.8
United States of America	USA			2,703	57.2
Argentina	ARG		South America	632	45.8
Brazil	BRA			176	46.1
Chile	CHL			1,739	41.7
Colombia	COL			161	39.4
Ecuador	ECU			75	42.1
Peru	PER			205	40.5
Uruguay	URY			597	43.2
Kenya	KEN	Sub-Saharan Africa	Southern and East Africa	443	61.1
Mauritius	MUS			1,305	48.4
South Africa	ZAF			2,695	66.9
Zimbabwe	ZWE			526	61.9

Table 4.7: Data structure — Census years and the ages of children and their fathers when father education and income is observed

Child cohort	Census year and ages of children and fathers						
Father cohort (age at birth)	1976	1981	1986	1991	1996	2001	2006
1962-66	10-14	15-19					
All cohorts	All						
1922-26 (40)	50-54						
1927-31 (35)	45-49						
1931-36 (30)	40-44						
1937-41 (25)	35-39 →	40-44					
1942-46 (20)		35-39					
1962-66	5-9	10-14	15-19				
All cohorts	All						
1927-31 (40)	45-49						
1932-36 (35)	40-44						
1937-41 (30)	35-39 →	40-44					
1942-46 (25)		35-39 →	40-44				
1947-51 (20)			35-39				
1971-76	0-4	5-9	10-14	15-19			
All cohorts	All						
1932-36 (40)	40-44						
1947-41 (35)	35-39 →	40-44					
1942-46 (30)		35-39 →	40-44				
1947-51 (25)			35-39 →	40-44			
1952-56 (20)				35-39			
1977-81		0-4	5-9	10-14	15-19		
All cohorts		All					
1937-41 (40)		40-44					
1942-46 (35)		35-39 →	40-44				
1946-51 (30)			35-39 →	40-44			
1952-56 (25)				35-39 →	40-44		
1957-61 (20)					35-39		
1982-86			0-4	5-9	10-14	15-19	
All cohorts			All				
1942-46 (40)			40-44				
1947-51 (35)			35-39 →	40-44			
1952-56 (30)				35-39 →	40-44		
1956-61 (25)					35-39 →	40-44	
1962-66 (20)						35-39	
1987-91				0-4	5-9	10-14	15-19
All cohorts				All			
1947-51 (40)				40-44			
1952-56 (35)				35-39 →	40-44		
1957-61 (30)					35-39 →	40-44	
1962-66 (25)						35-39 →	40-44
1967-71 (20)							35-39

Notes: Illustrates the census years and the ages of children and their fathers when father education and income is observed. For each migrant community, the first generation education is taken from a single observation when the second generation children are first recorded in the historical censuses. In contrast, first generation income is observed when the father is 35-39 years old and/or 40-44 years, to minimise lifecycle bias. Observations at older ages are taken to avoid missing earlier cohorts born to fathers otherwise too old to be captured.

Table 4.8: Persistence of source country effects

	Mobility regressions				Return regressions			
	Income		Education		Generation 1		Generation 2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
G1 income rank	0.23	0.16						
	(0.04)	(0.05)						
G1 education (years)			0.31	0.19	6.52	5.47		
			(0.04)	(0.06)	(0.56)	(0.79)		
G2 education (years)							5.72	5.55
							(0.42)	(0.78)
Specification:								
Country fixed effects		X		X		X		X
R^2	0.39	0.87	0.57	0.95	0.74	0.96	0.74	0.91
$adj.R^2$	0.36	0.83	0.55	0.93	0.73	0.94	0.72	0.88
N					283			

Notes: Results from regressions of: second generation income rank on first generation income rank (intergenerational income mobility, columns (1) and (2)); second generation educational attainment on first generation educational attainment (intergenerational education mobility, columns (3) and (4)); first generation income rank on first generation educational attainment (returns to education, columns (5) and (6)); and second generation income rank on second generation educational attainment (returns to education, columns (7) and (8)). Slopes and intercepts are allowed to vary by birth cohort across all specifications, with the marginal effect of the relevant regressor at means shown. Even number columns include country of origin fixed effects. Standard errors in parentheses.

Table 4.9: Influence on culture and context on individual youth achievements and aspirations at age 15

	Math score	Reading score	Aspired years of education for child			
			Child's aspirations		Parent's aspirations	
	(1)	(2)	(3)	(4)	(5)	(6)
PISA outperformance ('culture')	0.33***	0.21**	0.19***	-0.03	0.03	-0.13***
	(0.07)	(0.07)	(0.05)	(0.05)	(0.04)	(0.03)
Income penalty ('context')	0.01	0.00	0.05***	0.04***	0.04***	0.03***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Migrant	0.01	0.02	0.13***	0.11***	0.14***	0.11***
	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Female	-0.12***	0.34***	0.58***	0.51***	0.54***	0.43***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Father's:						
Education (years)	0.03***	0.03***	0.06***	0.04***	0.04***	0.02***
	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)
Occupation score	0.00***	0.00***	0.01***	0.00***	0.01***	0.00***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Mother's:						
Education (years)	0.02**	0.03***	0.05***	0.03***	0.04***	0.02***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Occupation score	0.00**	0.00	0.00***	0.00***	0.00***	0.00***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Mean peer parent education (years):						
Fathers	0.05***	0.06***	0.08***	0.04***	0.06***	0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Mothers	0.08***	0.06***	0.05***	0.01	0.01	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Ability controls						
Y9 math score				0.28***		0.16***
				(0.02)		(0.02)
Y9 reading score				0.37***		0.36***
				(0.02)		(0.02)
R^2	0.21	0.24	0.13	0.23	0.19	0.31
$adj.R^2$	0.18	0.21	0.12	0.21	0.16	0.28
N	9,300	9,303	37,403	37,210	17,944	17,944

Notes: Results from individual-level regressions of Year 9 math scores (column (1)), reading scores (column (2)), child aspired years of education (columns (3) and (4)) and parent aspired years of education (for the child) (columns (5) and (6)). A quadratic in age, and school and cohort fixed effects, are included in all specifications. The regressions in columns (1) and (2) are restricted to those observed in Wave 10 with non-missing educational attainment and weighted using survey-supplied weights to adjust for attrition. Educational aspirations are from Wave 1. Robust standard errors are reported in parentheses, accounting for clustering by paternal country of birth. Significance levels are indicated as follows: *** $p < 0.01$; ** $p < 0.05$; and * $p < 0.1$.

4.C Attrition and related concerns

This paper uses data from multiple censuses to create a pseudo-panel following migrant communities across generations. In a true panel, selective attrition would be a concern. In a pseudo-panel such as this, we might instead worry about selective mismatch between the populations used to measure first and second generation outcomes. A mismatch between the first and second generation populations could arise in a number of ways. For example, family dissolution may result in some fathers being missed in the first generation, while outmigration may result in some children being missed in the second generation.[29] Reporting or processing errors could also result in individuals being misclassified as first or second generation, which may be more common for those with low English proficiency.

In Table 4.10 I examine this mismatch. I show the percentage discrepancy between the second generation populations in the historical censuses in which either income (Panel A) or education (Panel B) is observed and the 2016 Census. I show the mismatch for both those born to Australia-born fathers, and its distribution for those born to overseas-born fathers. There is a moderate mismatch between first generation income measurement populations and the 2016 Census for both those born to Australia-born fathers and the median migrant community (Panel A), with higher populations in the 2016 Census. This is more pronounced for the most recent birth cohorts. These facts are consistent with family dissolution resulting in some fathers being missed in the first generation. There is much less of a mismatch — less than 6% in absolute magnitude across all cohorts — between first generation education measurement and the 2016 Census for those born to Australia-born fathers and the median migrant community. This likely reflects

[29] Outmigration is a common concern in the migrant assimilation literature (Dustmann and Görlach (2015)), however theory and past analysis suggests it may be less of a concern here. Borjas and Bratsberg (1996) develop a theoretical model of return migration that suggests it should vary negatively with migration costs, such as those associated with distance, and positively with mean income in the source country. These predictions are then shown to hold true in the United States data — outmigration is more common for migrants from nearby wealthy countries. Australia's distance from the rest of the world, particularly wealthy countries, suggests return migration should be uncommon. And indeed it is — for example, Cobb-Clark and Stillman (2013) report outmigration rates between the 1996 and 2001 Censuses of less than 10% across 88 countries of origin, with a modal rate of just over 1%. Further, outmigration only a problem if it results in us observing one generation but not the other — for example, fathers returning to their source country without their children, or children returning without their fathers.

the fact that while income is measured when fathers are aged 35 years or older (to minimise lifecycle bias), education is measured at the earliest point following the birth of their children. Thus families are more likely intact at the point at which education is measured.

Table 4.10: Percentage growth in 2^{nd} generation population between historical censuses and 2016 Census

Birth cohort	Australia-born fathers	Overseas-born fathers p10	p50	p90	N
Panel A: Historical censuses used for 1^{st} generation income					
1966	8.6	-23.2	6.6	52.3	64.0
1971	8.4	-12.7	3.4	56.3	71.0
1976	8.1	-16.0	0.1	30.9	74.0
1981	10.7	-16.8	-1.7	20.5	74.0
1986	15.9	-10.5	6.2	26.1	74.0
1991	22.8	1.5	16.7	46.0	74.0
Pooled	12.4	-15.4	5.1	38.0	431.0
Panel B: Historical censuses used for 1^{st} generation education					
1966	0.0	-42.5	-3.2	61.9	66.0
1971	-2.6	-22.2	-4.8	30.4	69.0
1976	-4.3	-23.2	-5.8	28.6	73.0
1981	-1.0	-13.6	2.6	24.8	74.0
1986	2.1	-7.7	6.0	23.6	74.0
1991	3.6	-7.1	3.9	31.7	74.0
Pooled	-0.4	-19.6	1.3	28.6	430.0

Notes: Shows the percentage growth in the 2^{nd} generation population between the historical censuses and 2016 Census. The historical census populations are based on the number of children in the father's household in the censuses in which either income (Panel A) or education (Panel B) is measured. The growth is shown for those born to Australia-born fathers, and its distribution is shown for those born to overseas-born fathers.

In Figures 4.7 and 4.8 I explore mismatch in a little more detail to motivate the baseline sample selection. In Figure 4.7 I present a simple scatterplot of the second generation population as observed in historical Censuses versus the 2016 Census. There is, unsurprisingly, a close relationship between the two. However, this relationship is much weaker for the very small populations in the the historical Censuses. This could be due to some first generation migrants being incorrectly classified in the 2016 Census — if the second generation population is very small in a birth cohort relative to the first generation then a small misclassification rate can still result in a large percentage increase in the reported second generation. I take a conservative approach and restrict attention to migrant communities with populations of more than 200 in both the historical and 2016 Censuses. In Figure 4.8 I present histograms showing the distribution of the growth between the historical and 2016 Censuses, for both income and education measurement. Few communities have large negative growth, which would have been consistent with large second generation outmigration. In the baseline analysis I allow for a mismatch of up to 30% in magnitude, excluding only the extreme tails of the distribution.

Figure 4.7: 2^{nd} generation population in the 2016 Census versus historical Censuses

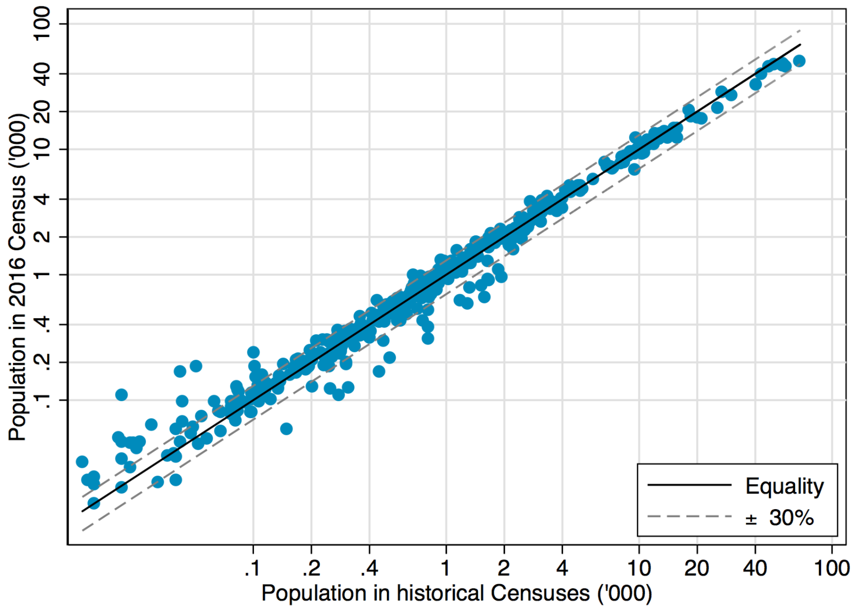

Notes: For all second generation migrant communities defined by country of origin and intercensal birth cohort, plots the population in the 2016 Census against that in the historical census in which the education of the first generation fathers is observed.

The core associations examined in this paper persist under alternative sample restrictions. In Table 4.11 I replicate the specification in column (4) of Table 4.2. Restricting the sample to communities with less mismatch between the generations *increases* the magnitude of the coefficients on PISA outperformance and the income penalty. The difference in the coefficients may reflect an attenuation bias

Figure 4.8: Histograms of % growth in 2^{nd} generation population from historical Censuses to the 2016 Census

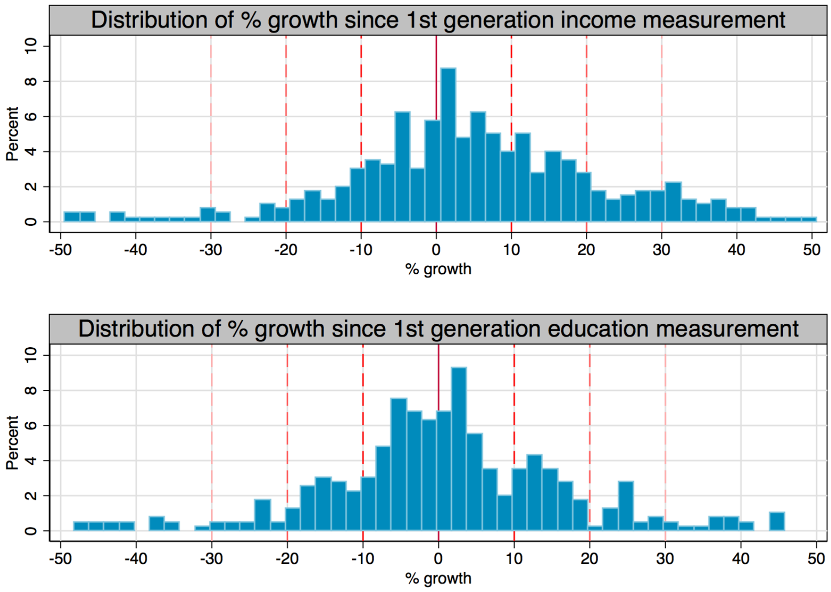

Notes: For all second generation migrant communities defined by country of origin and intercensal birth cohort, plots the distribution of the percentage growth between the population in the historical census in which either income or education is observed and that in the 2016 Census.

introduced by the mismatch, but could also be driven by idiosyncratic differences between the small samples.

Finally, it is also possible to examine the extent to which mismatches over a shorter time horizon are selective given many birth cohorts are observed with their fathers in multiple historical Censuses. In Table 4.12 I examine changes from one historical census to the next in the population of second generation children living with their fathers, and the outcomes of those fathers. I present this by the age bracket of the child in the earlier of the two censuses. As expected, over the five year interval between censuses, the population of children tends to fall — by an average of 5-10% for children initially observed at 0-4 years, increasing to 20-25% for children initially observed at 10-14 years. Again, this is consistent with family dissolution. The mean income rank and years of education of the fathers also tends to fall, which is consistent with less advantaged fathers being less

Table 4.11: Explanations of exceptional educational mobility — varying sample restrictions

Minimum population	100			200		
Maximum mismatch	±30%	±20%	±10%	±30%	±20%	±10%
Years of education (G1)	0.342***	0.360***	0.395***	0.348***	0.366***	0.399***
	(0.047)	(0.053)	(0.065)	(0.051)	(0.057)	(0.071)
PISA outperfomance	0.506**	0.541**	0.685*	0.488**	0.504**	0.682*
(‘culture’)	(0.200)	(0.213)	(0.343)	(0.206)	(0.213)	(0.343)
Income penalty	0.025***	0.030***	0.041***	0.025***	0.030***	0.041***
(‘context’)	(0.007)	(0.007)	(0.014)	(0.007)	(0.007)	(0.014)
Years of education (G1)	0.029**	0.032**	0.032**	0.029**	0.033**	0.032**
* NESDC	(0.012)	(0.012)	(0.014)	(0.012)	(0.012)	(0.014)
Log community size	-0.007	-0.000	-0.004	-0.007	-0.008	0.000
	(0.041)	(0.046)	(0.064)	(0.045)	(0.049)	(0.070)
Migrant	0.119	0.100	0.114	0.111	0.054	0.128
	(0.181)	(0.203)	(0.307)	(0.197)	(0.215)	(0.322)
R^2	0.622	0.640	0.677	0.631	0.648	0.670
$adj.R^2$	0.612	0.628	0.654	0.620	0.635	0.647
N	224	182	93	209	173	92

Notes: Results from group-level regressions of second generation mean years of education on first generation mean years of education, with an indicator variable for those born to overseas born fathers, and controls for: culture and context through the PISA outperformance and income penalty variables described in the text; ethnic network effects, including the community size and allowing the return to first generation education to differ for those from non-English speaking backgrounds. Robust standard errors are clustered at a country of origin level. Significance levels are indicated as follows: *** $p < 0.01$; ** $p < 0.05$; and * $p < 0.1$.

likely to live with their children with time. But the fall in outcomes is very modest, averaging less than 3%. This provides more comfort that mismatch between the populations underlying the observed first and second generation outcomes is unlikely to significantly bias the results.

Table 4.12: Ratio of 2^{nd} generation populations and 1^{st} generation outcomes in consecutive censuses

	Ratio of populations		Ratio of outcomes	
Cohort age in first census (years)	Mean	SD	Mean	SD
Panel A: Consecutive 1^{st} generation income observations				
0-4	0.945	0.094	1.014	0.072
5-9	0.933	0.090	0.997	0.071
10-14	0.836	0.097	0.988	0.074
Panel B: Consecutive 1^{st} generation education observations				
0-4	0.891	0.127	1.009	0.031
5-9	0.853	0.109	0.972	0.029
10-14	0.739	0.095	0.973	0.038

Notes: Presents the slippage in 2^{nd} generation population totals and 1^{st} generation outcomes when observing a migrant community in two consecutive historical censuses, by the age of the 2^{nd} generation in the first census.

4.D Australian-born benchmarks for mobility and returns

This Appendix describes how I decompose the income mobility of migrants relative to natives, rather than simply relative to other migrant communities. All that is required is to estimate, or assume based on prior literature, the coefficients for equations (4.1)-(4.3), for the native population. The exceptional mobility or excess returns for each migrant community is then the difference between observed and predicted values, and the decomposition in equation (4.4) can be calculated.

Benchmarking to natives is relatively straightforward when it comes to the returns to education. I am able to estimate individual level variants (or equivalents) of equation (4.3) for the first and second generation of natives using data from the same Censuses. For the second generation I use tabulations of income and education for the natives in each birth cohort in the 2016 Census, weighting regressions by the number of natives with each level of education. For the first generation I use the 1% Confidentialised Unit Record Files (CURFs) for the 1981-2006 Censuses.[30]

In contrast, it is not possible to estimate intergenerational relationships for natives in the Census data, as the links are possible only at group level.[31] Thus, for each birth cohort I only know the mean income and education of both generations of natives — a single data point in each relationship. I supplement this with estimates for the relevant slopes.

For educational mobility, I use the Household Income and Labour Dynamics in Australia (HILDA) survey. The HILDA survey is a nationally representative longitudinal survey of Australian non-remote private dwellings, with 16 waves of data currently available. Crucially, it asks respondents for both their own educational attainment and (since Wave 5) that of their parents. I begin with the full unbalanced panel and then restrict attention to individuals aged 25 or more and born in Australia to Australian-born fathers. I take the most recent

[30]CURFs were only produced for the 1981 Census onwards.

[31]The Australian Bureau of Statistics now has a program of linking individuals across Censuses, but the resulting longitudinal file only captures those from the 2006 Census onwards. Further, the Australian Census has not asked for state of birth in recent history, which rules out linking natives using the methodology in Aaronson and Mazumder (2008).

observation for each individual.[32] The selected sample does not fit those underlying any of the weights provided with the survey, so I instead weight all individuals in a given sex, cohort and educational attainment cell so as to reproduce the distribution of individuals across the same cells observed in the 2016 Census. I then estimate the individual level equivalent of equation (4.2), again allowing cohort specific slopes and intercepts. The results are in column (3) of Appendix Table 4.13 (other columns present alternative specifications). The slope coefficients are relatively similar across all cohorts, with a p-value of 0.35 on the test that they are jointly equal to their average value of 0.28. I thus use 0.28 as the underlying slope coefficient relating father and child educational attainment in the native population.

For the rank-rank relationship, I use a lower bound of 0.2 and an upper bound of 0.3. These choices are motivated by recent estimates using both survey and comprehensive tax data. Using HILDA survey data, Murray et al. (forthcoming) find a rank-rank correlation of 0.27 for those born in 1984-86, and aged 28-31 years at the time of measurement. Using tax data, Deutscher (2018) estimates a rank-rank correlation of 0.17 for Australian born children in 1978-85, and aged 30 years at the time of measurement. Finally, note that while the rank-rank correlation is less precisely estimated than its educational equivalent, it is also less central to the decomposition in equation (4.4) — it appears only in the second line, in a term that is equal for migrant communities with the same first generation income. Hence the decomposition is less sensitive to this parameter, and in particular the first three terms will remain unchanged.

[32]This minimises the extent to which we miss educational qualifications accrued latter in life. Parental education is only asked for once, so this response is used for each individual.

Table 4.13: Education mobility for those born to Australia-born fathers

	Coefficient on paternal years of education					
Birth cohort	(1)	(2)	(3)	(4)	(5)	(6)
Pooled	0.279			0.286		
	(0.011)			(0.011)		
1962-66		0.260	0.269		0.290	0.290
		(0.013)	(0.025)		(0.014)	(0.026)
1967-71		0.281	0.245		0.298	0.254
		(0.013)	(0.026)		(0.013)	(0.028)
1972-76		0.285	0.298		0.303	0.315
		(0.012)	(0.028)		(0.013)	(0.029)
1977-81		0.292	0.330		0.305	0.352
		(0.012)	(0.028)		(0.012)	(0.028)
1982-86		0.277	0.256		0.286	0.260
		(0.012)	(0.028)		(0.012)	(0.027)
1987-91		0.268	0.267		0.272	0.289
		(0.011)	(0.029)		(0.012)	(0.028)
Specification:						
Weighted	X	X	X			
Cohort varying...						
...slopes		X	X		X	X
...intercepts			X			X
p-value on test of:						
Slope(s) equal 0.28	0.905	0.005	0.353	0.613	0.000	0.133
Slope(s) equal		0.002	0.254		0.000	0.127
R^2	0.093	0.095	0.096	0.093	0.096	0.097
N	6,643	6,643	6,643	6,643	6,643	6,643

Notes: Results from a regression of child years of education on father years of education for children and fathers both born in Australia in the Household Income and Labour Dynamics in Australia (HILDA) survey. The most recent observation for all surveyed individuals is taken, provided they are 25 years or older at the time of the survey and have non-missing responses to the relevant questions. Standard errors in parentheses.

Chapter 5

Baby Bonuses: natural experiments in cash transfers, birth timing and child outcomes

Abstract

We use the 1 July 2004 introduction of the Australian Baby Bonus to identify the effect of family income on child test scores at grade three. Using a difference-in-differences design, we find no evidence the Baby Bonus improved child outcomes in aggregate, but some evidence of a modest effect for children from disadvantaged backgrounds. We examine whether birth shifting associated with the Baby Bonus and two other Australian maternity payments had negative long-term effects on children. Despite widespread concerns about this unintended treatment, regression discontinuity estimates provide no clear evidence of lasting health or educational consequences.

This chapter was co-authored with Robert Breunig. The authors would like to thank the Australian Curriculum, Assessment and Reporting Authority (ACARA), Social Research Centre (SRC), Department of Education and Training (DET), Australian Institute of Health and Welfare (AIHW) and Australian Taxation Office (ATO) for the data used in this paper. Discussions with Steve Croft (ACARA), Megan O'Connell (SRC), Phil Aungles (DET) and Georgina Chambers (AIHW) were particularly helpful in refining data requests. The research plan was approved by the Australian National University Human Research Ethics Committee, protocol number 2014/783. Nathan Deutscher acknowledges the support of the Sir Roland Wilson Foundation in

5.1 Introduction

In this paper we examine the long-term effects on educational outcomes of the 2004 Australian Baby Bonus – a cash transfer of AUD3,000 to parents of babies born on or after 1 July 2004.[1] We use a difference-in-difference methodology, comparing those born in the months either side of the introduction of Baby Bonus in 2004 with those born in the same windows in 2003 and 2005. This strategy is invalid at the cut-off date due to the well-documented birth shifting that occurred around the introduction of the policy. However, this birth shifting occurs within a very short and relatively precise period around the policy introduction and by excluding those who are most likely to have been affected by birth shifting, we can identify the long-term impacts of the policy.

We are also interested in the unintended birth shifting associated with the sharp introduction of family payments paid at the birth of a child. (Gans and Leigh (2009) documented unprecedented birth shifting accompanying the Baby Bonus). We examine the long-term consequences of three such Australian programs: the 1996 Maternity Allowance; the 2004 Baby Bonus; and the 2006 expansion of the Baby Bonus. All three policies induced birth shifting.

The paper thus makes two contributions. The first is to exploit a natural experiment in cash transfers at birth to identify their effect on educational outcomes at grade three. In exploiting a date of birth cut off in this way, this paper is most closely related to the relatively recent and growing literature evaluating expansions in maternity leave provisions. The paper's second contribution is to present the first tests for long-term harms associated with birth shifting events – in doing so we shed light on the complex set of influences determining who shifts births.

We find no evidence that the Baby Bonus improved child test scores in aggregate. Indeed, we can reject at a 5% level the hypothesis that the Bonus shifted average test scores by more than 1.8% of a standard deviation lower or 1.5% of a standard deviation higher. This cautions against an overly optimistic (or pessimistic) take on the effect of cash transfers on child outcomes based on the existing literature. There is some evidence in support of a modest effect for more

his PhD studies. All findings and views in this paper are those of the authors and should not be attributed to others.

[1] US$2,100 in 2004 dollars.

disadvantaged families.

Given the size of the payment (AUD3,000) it may seem unsurprising that there is little effect. However, governments frequently sell modest interventions of this nature by touting large medium- and long-term benefits. The literature (see section 5.2.2) also suggests that modest interventions can have significant effects if they are well-aimed and well-timed.

A relatively complex picture emerges about the effect of large birth shifting events, with no 'smoking gun' suggesting they lead to obviously worse outcomes. Any lasting disparities around the policy changes appear just as likely to be due to differences in family background as to the actual impact of birth-shifting.

We begin by describing the rationale for family payments, the related literature and the Australian policy environment, which has been relatively atypical among OECD countries. The methodology is outlined in Section 5.3. In Section 5.4, the data are described along with the construction of outcome variables and the analysis sample. The results are presented and alternative specifications tested in Section 5.5. Section 5.6 places the results in the context of existing literature and concludes.

5.2 Background

5.2.1 Family Benefit Programs

In 2011, all OECD nations provided some form of cash benefits to families, at an average cost of 1.3% of GDP.[2] Despite their ubiquity, there is a great deal of variety in how and why countries boost family incomes. Parental leave policies look to encourage parents, often mothers, to care for their children during infancy. Broader cash transfers to families can be seen as a form of redistribution, in recognition of the costs of raising children, or an investment in later child outcomes. This latter view, of children as an investment rather than a cost, has gained increasing traction in public and academic discourse.

Despite significant public expenditure and interest in family benefit programs, the literature on the *causal* effect of family income on child outcomes remains

[2] https://data.oecd.org/socialexp/family-benefits-public-spending.htm

relatively modest. This reflects the difficulty in finding genuinely exogenous sources of variation in family income. As noted in Currie and Almond (2011), there are few policy examples that increase incomes without potentially affecting child outcomes through other channels.

5.2.2 Related Literature

A number of previous papers have attempted to identify the causal effects of cash transfers using instrumental variable techniques exploiting variations in cash transfers over time, household type and location. In Dahl and Lochner (2012), variations in the amount of Earned Income Tax Credit that households are eligible for are exploited to identify relatively modest effects of household income on children's test scores. They find an additional USD1,000 in current income raises combined math and reading test scores by 6 per cent of a standard deviation, with larger effects for children from disadvantaged families. Exploiting variations in Canadian child tax benefits, Milligan and Stabile (2011) find larger effects on test scores, and also identify positive effects on physical health for boys and mental health for girls.

These studies focus on the contemporaneous effects of family income on child outcomes. Yet there is a rich and well known literature on the potential for large effects from early life experiences on later childhood and adult outcomes; see Currie and Almond (2011). In Morris et al. (2005), the authors draw on seven random assignment welfare and poverty evaluations to examine the age-specific pattern of effects of welfare policies on child achievement. Earnings supplement programs are found to be most effective between ages 2 and 5 (although the treatment effect for those aged 0 to 1 is not statistically different).

The Australian Baby Bonus offers a natural experiment to identify causal effects on child outcomes of family income *at birth*. Based on the existing literature, it is difficult to know a priori what kind of effect to expect from the Baby Bonus. With the full universe of test score data we have the precision to detect effects even smaller than those discussed above. The median net benefit from the Baby Bonus was around AUD1,700 (in 2004 dollars), making it larger than the aforementioned increases lifting test scores by several percentage points of a standard deviation. However, the Baby Bonus was a one-off payment, whereas the increases in tax benefits were likely longer lasting (though precisely how much so is unclear). The timing of the Baby Bonus at birth could either reduce its effect – if the effect of income fades over time – or increase its effect – if early childhood is a more sensitive period with long lasting consequences (as in Morris et al. (2005)). This latter perspective is also consistent with economic models of child development with potentially large returns to early childhood investments (such as in the influential work of Cunha and Heckman (2007)). The more pessimistic view of early childhood investments would leave us expecting no discernible effect from the Baby Bonus, even with the full universe of test scores. However, the more optimistic strands of the literature leave it as a real possibility, making the Baby Bonus an interesting test case.

In theory, cash transfers may influence child outcomes through a variety of mechanisms. They may directly lift investment in a child's health and education. They may indirectly affect child development, such as through reduced parental financial stress. Harmful effects are also plausible. For example, parents may spend more on goods with spillover costs to the child (such as tobacco or alcohol). Parents may also have more children in response to the Baby Bonus, resulting in a potential quantity-quality trade-off.[3] Our estimates are of the combined effect of the cash transfer, implicitly including the individual effects of all such mechanisms.

There is already a modest literature pointing to short- and long-run consequences of the Baby Bonus. The unintended birth shifting and associated potential for harm was noted by Gans and Leigh (2009). In Drago et al. (2011), the Baby Bonus is found to have (very modestly) lifted fertility rates following its introduction. Our paper is the first looking at the net effect on child outcomes of the payment *and* its unintended consequences. A concurrent working paper, Gaitz and Schurer (2017), exploits variation in sibling birthdates to isolate an exogenous increase in family income from the Baby Bonus.[4] They find no clear evidence of an effect of the Baby Bonus on child outcomes, or on potential transmission mechanisms such as parental wellbeing, behaviour and labour supply. Their study

[3]Family size is not available in our data, so we cannot separately identify such effects.

[4]The use of sibling birthdate variation is by necessity — the cohorts in their dataset were all born either well before or after the introduction of the Baby Bonus so it is only the siblings that allow examination of birthdates either side of the cut off.

is, however, constrained by the small sample sizes of only a few hundred children available in their survey data.

Examining whether the unintended birth shifting that accompanied the Baby Bonus had any effect on child outcomes is also of broad interest. There are now numerous examples of birth shifting in response to announcement effects, and the scheduling of births for a wider range of reasons is increasingly common. For example, birth delays and increases in birthweight have been identified in response to German parental leave reforms; see Tamm (2013). In the opposite direction, Schulkind and Shapiro (2014) find relatively small birth bring-forwards in response to the United States' tax system, which are nonetheless associated with lower birthweight and lower Agpar scores. More broadly, birth shifting has been identified by parents seeking to avoid inauspicious days and care providers wishing to avoid weekends, public holidays and conferences; see Gans and Leigh (2012); Gans et al. (2007).

5.2.3 Policy Background

Australia has a long history of relatively unconditional cash transfers for new parents. The first 'Baby Bonus' was a near universal £5 maternity allowance for births from midnight 10 October 1912; see Kewley (1973).[5] The allowance was intended to help pay for the attendance of a health professional at childbirth, with then Prime Minister Fisher suggesting it would:

> *"be the means of helping poor parents to tide over an anxious period, and ensuring that their offspring's health, and perhaps life, shall not be jeopardised in the dawn of existence ... I venture to say that it will not only protect many mothers, but to a great extent will preserve child life, and send it into the world equipped."*

While popular, maternity allowance was found to have had little effect on infant and maternal health and was allowed to erode in value before being abolished in 1978 (Figure 5.1).

[5] Near universal in that the allowance was not paid where the mother 'was an Asiatic or an Aboriginal of Australia, Papua or the Pacific Islands'; see Daniels (2009).

In recent times, Baby Bonuses have filled a hole left by the absence of mandatory paid maternity leave, introduced in Australia in 2011. In 1996, Australia and the United States were the only two OECD nations without mandatory paid maternity leave. Political pressure to lift support to new parents without disadvantaging stay-at-home mothers led to the reintroduction of, and increases to, cash transfers to new parents. It is this atypical experience that provides a series of natural experiments in the provision of cash transfers at birth.[6]

Figure 5.1: Cash transfers on birth of a child.

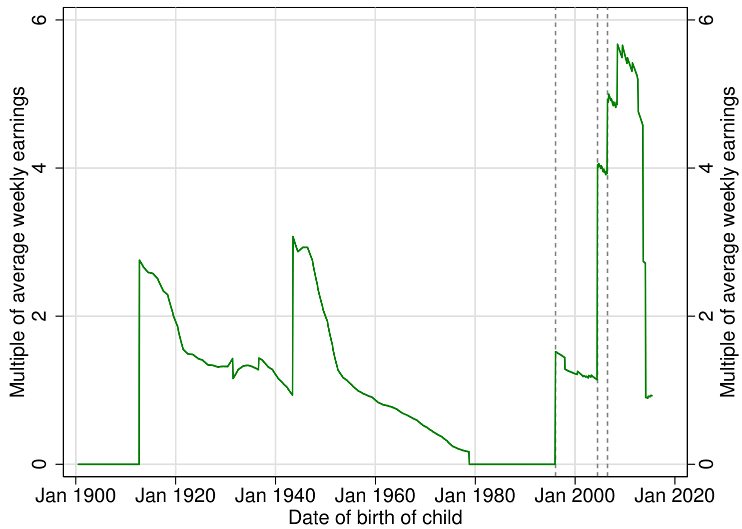

Notes: Dashed lines mark the policy changes examined in this paper: the 1 February 1996 reintroduction of Maternity Allowance; the 1 July 2004 introduction of the Baby Bonus; and the 1 July 2006 increase to the Baby Bonus. The above ignores income testing in place from 1 July 1931 to 1 July 1943 and from 1 February 1996 onwards and uses the minimum payment for a child where payments have varied according to the number of existing children in the family. Does not include the First Child Tax Refund, discussed in more detail below, or Paid Parental Leave, introduced from 1 January 2011. Historical payment rates are taken from Daniels (2009); *Family Assistance Guide* (2016) and average weekly earnings from Hutchinson (2016).

[6] In contrast to the original Baby Bonus, and perhaps more realistically, the stated aspirations for recent cash transfers to new parents have been more modest, and framed around assisting with the costs of parenthood. There may also have been political motivations, with the major policy changes — the introduction of Maternity Allowance in 1996 and the Baby Bonus in 2004 — coinciding with federal election years.

5.2.4 The Baby Bonus

The primary focus of this paper is the 1 July 2004 introduction of the Baby Bonus, an unconditional cash grant of AUD3,000 on the birth of a child.[7] The Baby Bonus was the largest increase in unconditional maternity payments in Australia since WWII (Figure 5.1). The policy was announced only a few months prior to its introduction with the delivery of the Commonwealth Budget on 11 May 2004. As such, any children conceived in anticipation of the Bonus would have had a due date well into the following year.

Interpreting the effect of the Baby Bonus is complicated by the fact it was not a uniform, unconditional increase in family incomes in *net* terms. The Baby Bonus replaced Maternity Allowance (a smaller means-tested cash transfer) and the variable First Child Tax Refund.[8] Nevertheless, the Baby Bonus was an unambiguous and large increase in entitlements for most families. Table 5.9 provides information about the distribution of the First Child Tax Rebate and the net benefits of the Baby Bonus. Around 75% of births in June 2004 would have been better off under the new policy, while 44% would have benefited by the full \$2,157 increase in cash transfers (as they were ineligible for the refund). Importantly, the change was more beneficial, on average, to parents from disadvantaged backgrounds. For the mean (median) June 2004 birth the Baby Bonus would have been an increase of around \$900 (\$1,700) at birth — around one (two) week's average earnings.[9]

In examining the long-run consequences of birth shifting, we also draw on the next two largest policy changes since WWII – the reintroduction of Maternity Allowance on 1 February 1996 and an increase to the Baby Bonus on 1 July 2006.

[7] This was originally named Maternity Payment, but has always been popularly known as the Baby Bonus, and was eventually renamed as such.

[8] The First Child Tax Refund is described in detail in Appendix 5.C.

[9] Any financial assistance provided by the First Child Tax Refund is likely to have been more heavily discounted than suggested in Table 5.9. First, the refund was paid out over the five years following the birth of the child, lowering its net present value. Second, the refund depended on the (uncertain) future taxable income of the primary carer, which would lower the certainty equivalent cash transfer at birth for the risk averse. Finally, the refund was far from salient — claims were still being submitted up to five years after they could have been paid out. These features could explain the lack of any sizeable bring-forward in births by those who would have been better off under the refund compared to the Baby Bonus. Our examination of Australian Taxation Office data, not presented here, shows no significant increase in number or size of claims around the cut-off.

In Table 5.1 we present key details of these policies, including their gross value and estimates of the share of births delayed in the fortnight prior to the policy. The Baby Bonus estimates are taken from Gans and Leigh (2009), the Maternity Allowance estimate uses the same method and data as those authors to ensure comparability.[10] The introduction of the Baby Bonus induced by far the largest shift in births. Assuming birth shifting in line with the other two policy changes implies an average net value for the Baby Bonus change of $1,200-2,700, broadly consistent with the figures in Table 5.9.[11]

Table 5.1: Births delayed in response to changes in cash transfers at birth.

Policy	Date of effect	Conception effect?	Gross value (2004 dollars)	Share of births moved
Maternity Allowance	1 Feb. 1996	No	AUD1,030	4%
Baby Bonus	1 July 2004	No	AUD3,000	10%
Baby Bonus	1 July 2006	Possible	AUD780	6%

Notes: Conception effect indicates whether the policy was announced more than 9 months prior to its date of effect. Estimates of the share of births moved in response to the introduction and increase to the Baby Bonus are from Gans and Leigh (2009), Column 2 of Tables 1 and 8 respectively. They take the logarithm of the daily birth count for the fortnights either side of the relevant threshold and regress on week, day of year and public holiday fixed effects, with an indicator for the fortnight after the relevant policy change. They use data from 1975-2004 for the 1 July 2004 Baby Bonus introduction and from 1975-2006 (excluding 2004) for the 1 July 2006 Baby Bonus increase. We use the same data and method and all years from 1975-2006 to estimate birth shifting for the 1 February 1996 Maternity Allowance reintroduction.

[10] Maternity Allowance is not examined in Gans and Leigh (2009) but its introduction induced birth shifting that is apparent in the data underlying that paper.

[11] Comparing the re-introduction of Maternity Allowance on 1 February 1996 and the increase to the Baby Bonus on 1 July 2006 reveals a clear increase in the propensity of parents to shift births in response to policy incentives. This may in part reflect the greater prevalence of caesarean sections, which rose from around 20% to 30% of births over this period.

5.3 Methodology

5.3.1 Effect of intended treatment

We use a difference-in-differences design to estimate the treatment effect of the additional family income provided by the Baby Bonus. This design compares the outcomes of children born shortly before and after the introduction of the Baby Bonus (in 2004) with those born in the same windows the years either side (2003 and 2005). Similar approaches have been used to evaluate expansions in maternity leave coverage in numerous countries (as an example, see Dustmann and Schönberg (2012)). In particular, we estimate:

$$y_i = \beta_0 + \beta_1 \mathbb{I}_{2004,i} + \beta_2 \mathbb{I}_{afterJuly,i} + \beta_3 \mathbb{I}_{2004,i}.\mathbb{I}_{afterJuly,i} + \beta_4 X_i + \varepsilon_i \tag{5.1}$$

where, for child i, y_i is their test score, $\mathbb{I}_{2004}$ indicates whether they were born in 2004 and $\mathbb{I}_{afterJuly,i}$ indicates whether they were born on or after 1 July. The vector of covariates X_i includes: sex; location (dummy variables for each state or territory and for whether the child lives in a metropolitan area); language background other than English; state specific cubic terms in the child's age-at-test; and parental demographics.[12]

As previously noted the introduction of the Baby Bonus induced a large shift in births, with around 1,000 births delayed until after 1 July 2004 (Gans and Leigh (2009)). This poses a threat to the validity of the identification strategy by potentially altering the composition and outcomes of those born either side of 1 July in the treatment year. To address this we exclude those born in the fortnight either side of 1 July. This is likely to be sufficient as the birth shifting was highly localised — nearly 20% of the 'shifted' births occurred on the first day of the policy and nearly 80% within the first fortnight (Gans and Leigh (2009)).[13] In addition,

[12]A student is defined as being of a language background other than English if either the student, the student's mother or the student's father speaks a language other than English at home. Parental demographics are captured by categorical variables covering the mother's and father's education — unknown, Year 9, Year 10, Year 11, Year 12, certificate, diploma, and bachelor degree or higher — and occupation — unknown, not in paid work, blue collar, white collar, management, and senior management.

[13]Note birth shifting outside this window need not imply that any births were moved more by than a fortnight to cross the 1 July threshold. Rather it may reflect flow-on birth reschedulings

we show (see Table 5.7 below) that the results also hold when excluding all those born a month either side of the policy introduction date.

5.3.2 Analysis of an unintended 'treatment': birth shifting events

Birth shifting also provides an opportunity to examine an unintended consequence of the policy that was highly criticised at the time. As Gans and Leigh noted "with more data, this event provides an opportunity for health researchers and economists to study the impact of a large disruption in a well-developed, modern medical system". This is the focus of the second part of this paper. This is a descriptive exercise. How do those born either side of the policy changes differ in their background characteristics and outcomes? How do these facts fit with potential stories about who shifts births, and whether birth shifting events have harmful consequences?

Birth shifting is associated with four key differences between those born immediately before and after the Baby Bonus (or other major policy changes).

- **Hospital crowding:** the primary concern raised by Gans and Leigh (2009) was the additional pressure placed on the hospital system. This was also the focus of media reporting at the time. In all, 1,005 children were born on 1 July 2004, over 200 more than what might have otherwise been expected (Gans and Leigh (2009)).

- **Increased gestation length and birthweight:** a further concern was based on overly long gestations or high birthweights being associated with poorer infant health; see Thorngren-Jerneck and Herbst (2001). Thus, if planned births are postponed for too long, it is possible that child outcomes will be compromised.

- **Differences in family backgrounds:** it may be that some families are better able to shift births around the introduction date than others. These families may also differ in ways (observable or not) that matter for their

by providers to smooth out their workflows — for example moving births from early to late June or early to late July.

children's later life success. This could raise equity concerns around the introduction of the policy.

- **Differences in policy environment:** finally, birth shifting is motivated by the difference in policy across the threshold — in our case the parents of those born after the changes typically received more financial support.

It is not possible to disentangle the first three of these effects – the causal effect of the last is estimated (for the Baby Bonus) in the first part of the paper. Nonetheless, it is possible to imagine a case where a relatively compelling picture of long-run harms emerges. If the policy itself has negligible effect on outcomes (as we find), and no or 'positive' selection into birth shifting sits alongside consistently worse outcomes after the threshold, then that would seem to strengthen the argument that the birth shifting event itself harmed long-run outcomes.

The existing work by Gans and Leigh (2009) provides a good sense of the magnitude of birth shifting that accompanied the introduction of the Baby Bonus. The birth shifting and potential hospital crowding is thus relatively well understood. We use a regression discontinuity approach to estimate other discontinuities in characteristics and outcomes around 1 July 2004.[14] As recommended in Lee and Lemieux (2010) we use local linear regressions either side of the threshold, allowing both intercept and slope to vary. Let h_i be the health or family background variable of interest. We now estimate the regression below in a window around the introduction of the Baby Bonus, where t is number of days after 1 July 2004 and $\mathbb{I}_{afterJuly,i}$ remains an indicator for whether the individual was born on or after 1

[14]The introduction of the Baby Bonus is a classic case where a regression discontinuity design would not be a valid approach to estimating the treatment effect of the policy alone. Nonetheless, regression discontinuity can be used in this case to estimate the differences associated with (rather than strictly caused by) the policy.

July 2004.[15]

$$h_i = \alpha_0 + \alpha_1 \mathbb{I}_{afterJuly,i} + \alpha_2 t + \alpha_3 \mathbb{I}_{afterJuly,i} t + \alpha_4 X + \varepsilon_i \tag{5.2}$$

The estimated discontinuity in the expected value of h at 1 July 2004 is now α_1, and standard hypothesis testing can be used. For some of the outcomes observed either side of the Baby Bonus we can modify this slightly to a regression discontinuity – difference-in-differences approach, estimating:

$$\begin{aligned} h_i = \alpha_0 + \alpha_1 \mathbb{I}_{2004,i} + \alpha_2 \mathbb{I}_{afterJuly,i} + \alpha_3 \mathbb{I}_{2004,i}.\mathbb{I}_{afterJuly,i} + \alpha_4 t + \alpha_5 \mathbb{I}_{2004,i} t + \\ \alpha_6 \mathbb{I}_{afterJuly,i} t + \alpha_7 \mathbb{I}_{2004,i}.\mathbb{I}_{afterJuly,i} t + \alpha_8 X + \varepsilon_i \end{aligned} \tag{5.3}$$

where α_3 is now the parameter of interest.[16] This approach has the advantage of increasing precision and controlling for any pre-existing financial year discontinuities.[17]

A vector of covariates X is included in the regressions to control for day of week effects in health and family background variables. In testing for a discontinuity in demographic variables and test scores we use the same vector of covariates as

[15] A key question in the regression discontinuity literature is the choice of bandwidth. Two broad approaches may be used: estimating an optimal bandwidth given the joint distribution of all variables, beginning with estimating a rule-of-thumb bandwidth, and cross-validation procedures (Lee and Lemieux (2010)). In our case, for local linear regression with a rectangular kernel, the rule of thumb bandwidth is given by:

$$bw_{ROT} = 2.702 \left[\frac{\tilde{\sigma}^2 R}{\sum_{i=1}^{N} \{\tilde{m}''(x_i)\}^2} \right]^{\frac{1}{5}}$$

where $\tilde{m}''(.)$ is the second derivative (curvature) of a regression of h on the independent variables in Equation 5.2, $\tilde{\sigma}$ is the estimated standard error of the regression and R is the range of the assignment variable, here date of birth, over which the regression is estimated. For simplicity and given the nature of the results, we simply present the regression discontinuity estimates over a range of bandwidths, including for the rule of thumb bandwidth when estimating any discontinuity in outcomes. Following Lee and Lemieux (2010) we use a quartic in the assignment variable when calculating the rule of thumb bandwidth.

[16] We thank an anonymous referee for suggestion a regression discontinuity – difference-in-differences approach. An earlier version of this paper used regression discontinuity exclusively, though with similar results Deutscher and Breunig (2016).

[17] For example, the interstate migration of children who started schooling in those states with 1 July cut off dates could generate such a discontinuity in our sample.

in the earlier difference-in-differences design.[18] To complement this analysis, we run similar regressions testing for lasting consequences from the two other policy changes in Table 5.1.

5.4 Data and sample selection

5.4.1 Child outcomes data

Early child outcomes are drawn from Australia's national system of school testing: the National Assessment Program — Literacy and Numeracy (NAPLAN).[19] Since 2009, students in Years 3, 5, 7 and 9 have sat the NAPLAN tests in mid May of each year. The tests cover public and private schools and jurisdictions have committed to ensuring maximum participation.[20] Students are tested across five domains — reading, numeracy, spelling, grammar and written work. Raw marks in each of these domains are mapped onto an achievement scale spanning all four year levels and ranging from approximately 0 to 1000.[21] These measures of student achievement form our initial outcome variables. Throughout the paper we normalise test scores to have mean zero and standard deviation one.

In Australia, primary and secondary education policy is primarily the responsibility of the six state and two territory authorities. An important difference between states and territories is in their school commencement policies: children can only start school if they will be six years old by the relevant cut off date. The degree to which children eligible to start school can be held back ('redshirted') also differs across states. In the years examined, three jurisdictions — Queens-

[18]We only test for a discontinuity in expected test scores adjusted for observable family background traits.

[19]The NAPLAN data used in this publication are sourced from the Australian Curriculum, Assessment and Reporting Authority (ACARA) and are available from ACARA in accordance with its Data Access Protocols.

[20]Students can be exempted from one or more NAPLAN tests if they have significant or complex disability, or if they are from a non-English-speaking background and arrived in Australia less than one year before the tests. However, exemption is not automatic and parents may choose for their child to participate. Exempted students will still appear in our data, just with missing test scores. The issues posed by missing test scores are considered below.

[21]Details of the process of mapping raw marks onto these scales are in Australian Curriculum, Assessment and Reporting Authority (2015); many of the procedures are similar if not identical to those used in the Programme for International Student Assessment (PISA).

land, Western Australia and the Northern Territory — have 1 July cut off dates. They also have virtually no redshirting. As a result those born either side of 1 July 2004 in these jurisdictions are in different test years and differ by nearly a year in the age at which they sit the test. This difference in age at test induces a large discontinuity in test scores (Appendix 5.A, Figure 5.6) which would render the regression discontinuity approach invalid. These jurisdictions are dropped from the sample. The remaining jurisdictions account for around two thirds of the population of interest and have slightly higher test scores than the national average.[22]

We focus on those born from 2003 to 2005. These children can end up in different test years based on state and territory policies and whether they were held back (Appendix 5.A, Figure 5.7). We restrict attention to those sitting the test in the year in which they turn 8, 9 or 10. This is necessary as we only have NAPLAN data for Year 3 students sitting the tests from 2011 to 2015.[23] It is also a fairly benign restriction, as it captures the vast majority of each cohort: only 0.05% of those born in 2004 in our data took the test outside this window.

For the difference-in-differences design the key identifying assumption is that of common trends across treatment and control. Here this requires that the relationship between date of birth and test score remains the same across birth years — the relative penalties or bonuses from being born in a certain part of the year do not change. This could fail in a number of ways. First, NAPLAN tests change between years, and different tests may, unintentionally, have different age-at-test effects. A large amount of effort goes into the process of mapping raw marks to the NAPLAN achievement scale to provide a measure that is comparable across the years (Australian Curriculum, Assessment and Reporting Authority (2015)), but this may nonetheless be insufficient for our purposes. Indeed, student gain scores are only published for reading, numeracy and writing. Trends in redshirting may also influence the relationship between when a child is born during the year and

[22]The cut off dates for the remaining states vary: 1 January in South Australia and Tasmania; 1 May in Victoria and the Australian Capital Territory; and 1 August in New South Wales.

[23]Specifically, we have NAPLAN data from 2013, 2014 and 2015. This includes students' past results, which allows us to infer the Year 3 test results for students who sat the tests in 2011 and 2012 and were eligible to sit the Year 5 tests in 2013 and 2014. This gives us Year 3 test results from 2011-2015 inclusive.

their eventual test scores. Some parts of the year are more prone to the influence of redshirting than others: Figure 5.7 shows the distribution of children across test years changes most rapidly prior to 1 May. Importantly, test year is not something that can sensibly be controlled for in our regressions as it is in part endogenous to ability — students are probably more likely to be held back, all else equal, if it is felt they are not yet ready for school.

To inform the choice of tests and windows we estimate the treatment effects for a placebo policy introduced on 1 July 2005, using 2005 as the treatment year and 2003 as the control. This tests the common trends assumption for the two control years. Of course, it is possible there are variations in the age-at-test effect in 2004 for the chosen tests and windows that are not apparent across 2003 and 2005, so we later test the sensitivity of our results in other ways. The outcome variables are (normalised) test scores for reading, numeracy, spelling, grammar and written work. Windows are progressively widened to look one, two or six months either side of 1 July (less the fortnight either side of July 1). A hybrid window from May to December is also included to make the most of the relative stability in test year distribution in the second half of the year. Results are in Table 5.2. Highly statistically significant differences are found for the spelling and written tests across most windows and for the full year window across all but one test. These deviations from the common trends assumption are also apparent from a visual inspection of test scores by date of birth, as shown in Appendix 5.A, Figure 5.8. Based on these results we restrict attention to reading and numeracy test scores and the first three windows.[24,25] Reading and numeracy test scores are

[24]Grammar scores are excluded as they arise from the same test as spelling scores. While spelling and grammar have different questions on this test, this could nonetheless generate an interdependency.

[25]This data driven rejection of spelling, grammar and written work test scores as outcome variables also has a sound intuitive basis — there are reasons why these scores may be less discriminating and thus harder to equate across years. The reading and numeracy scores are based on separate tests which, for Year 3 students in 2014, had 38 and 35 questions respectively (marked either as right or wrong). In contrast, spelling and grammar scores were based on only 25 and 26 questions respectively, and were bundled into the same test. While other tests are year-level specific, the written work score is based on a persuasive response to a writing prompt that is identical across year-levels. This is marked on ten criteria, and while the total mark is out of 48 there was evidence of multi-modal distribution in 2014 that made equating scores more difficult, see Australian Curriculum, Assessment and Reporting Authority (2015).

Table 5.2: Estimated treatment effects of placebo policy on component tests.

Window	Jun-Jul	May-Aug	May-Dec	Jan-Dec
	(1)	(2)	(3)	(4)
Reading	-0.007	-0.018	-0.008	-0.008
	(0.021)	(0.012)	(0.010)	(0.007)
Numeracy	-0.022	-0.006	-0.000	0.036***
	(0.021)	(0.013)	(0.010)	(0.007)
Spelling	-0.048**	-0.027**	-0.021**	-0.032***
	(0.021)	(0.013)	(0.010)	(0.007)
Grammar	-0.023	-0.019	-0.011	-0.017**
	(0.021)	(0.012)	(0.010)	(0.007)
Written	-0.025	-0.025**	-0.028***	-0.080***
	(0.020)	(0.012)	(0.010)	(0.007)

Notes: Table presents DiD estimated treatment effects of a placebo policy introduced on 1 July 2005. That is, we estimate equation 5.1 except using the 2005 birth cohort as the treatment and the 2003 birth cohort as the control. All windows exclude a fortnight either side of 1 July. Covariates include: sex; location (dummy variables for each state or territory and for whether the child lives in a metropolitan area); language background other than English; state specific cubic terms in the child's age-at-test; and categorical variables for parental education and occupation. Robust standard errors in parentheses. Significance levels indicated as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

summed to generated a combined test score.[26]

Reading and numeracy test scores are available for 90% of students. Missing test scores are potentially problematic for two reasons. First, missing test scores are not random. Students with missing scores are more likely to have parents with no further education beyond high school and who are unemployed; children from such backgrounds who are not missing test scores have lower test scores on average. Second, the pattern of missing test scores across the calendar years does not follow the common trends assumption, as is clear from Appendix A, Figure

[26] This is primarily done to (potentially) lift precision given reading and numeracy scores are not perfectly correlated. We also report key results separately for reading and numeracy scores.

5.9. However, if we assume those with missing test scores tend to perform worse, then the trends observed will, if anything, bias our results downwards.[27]

In Table 5.3 we present summary statistics from the NAPLAN data for the three birth cohorts — 2003, 2004 and 2005. Mean test scores and demographics are very similar across all three test years, though with a trend towards slightly more educated parents in higher status jobs over time. Of course, what matters for our identification strategy is not that the levels in child characteristics are similar across the treatment and control years, but that the trends within birth years are. In Table 5.4 we do this by estimating treatment effects for our observed child characteristics. The results are encouraging, with the common trends assumption not rejected (i.e. no statistically significant treatment effect on the covariate) for the three shortest windows that are the focus of our analysis.

5.4.2 Births and birth outcomes data

As a prelude to looking for lasting consequences of birth shifting we extend the work on intermediate infant health outcomes by Gans and Leigh (2009). To do this we use Australian Institute of Health and Welfare (AIHW) data on daily birth outcomes in 2004. These are drawn from the National Perinatal Data Collection, which includes information from midwives and other health professionals attending births. In particular, we have information on the distribution of gestational age at birth and birthweight by date of birth cohort. In Gans and Leigh (2009), the authors had birthweight data for only three jurisdictions (covering 36% of births) and did not have gestational age data. Further details of the AIHW data can be found in Australian Institute of Health and Welfare (2014).[28]

[27] There is a lower prevalence of missing scores in the second half of 2004 relative to what one would expect based on trends in 2003 and 2005 and the prevalence of missing scores in the first half of 2004. If those missing scores do worse, this will lead scores in the second half of 2004 to be lower than expected.

[28] Data on daily birth counts, used in Table 5.1 and Figure 5.4 is from the Australian Bureau of Statistics (ABS). These data are precisely as used and described in Gans and Leigh (2009) and have been provided by those authors. They are drawn from state and territory birth registers and includes all registered births in Australia. Further details can be found in Australian Bureau of Statistics (2014).

Table 5.3: Summary statistics

Birth year	2003 (control)	2004 (treatment)	2005 (control)
Mean reading score	0.068	0.065	0.049
	(0.994)	(0.948)	(0.970)
Mean numeracy score	0.069	0.062	0.097
	(0.987)	(0.950)	(1.000)
Age at test	8.665	8.658	8.651
	(0.381)	(0.374)	(0.369)
Female (%)	48.5	48.5	48.8
	(50.0)	(50.0)	(50.0)
LBOTE (%)	26.0	26.1	26.3
	(43.9)	(43.9)	(44.0)
Living in a city (%)	75.5	75.5	75.8
	(43.0)	(43.0)	(42.9)
Living in New South Wales (%)	47.7	47.6	48
	(49.9)	(49.9)	(50.0)
Highest parental...			
...education=dropout (%)	11.1	10.5	10
	(31.4)	(30.6)	(30.1)
...education=high school (%)	19.3	18.7	18.1
	(39.5)	(39)	(38.5)
...education=degree (%)	35.2	36.9	38
	(47.8)	(48.2)	(48.6)
...job=unemployed (%)	9.2	9.1	9.2
	(28.9)	(28.8)	(28.9)
...job=senior manager (%)	24.3	24.9	25.4
	(42.9)	(43.2)	(43.5)
N	186,777	184,949	194,302

Notes: presents summary statistics on the sample used for estimation, excluding jurisdictions with a 1 July cut off dates for school starting.

Table 5.4: Tests of common trends: estimated treatment effects of policy on child characteristics

Window	Jun-Jul (1)	May-Aug (2)	May-Dec (3)	Jan-Dec (4)
Female	-0.314	-0.412	-0.085	-0.388
	(0.954)	(0.559)	(0.450)	(0.301)
LBOTE	0.476	-0.360	0.031	0.244
	(0.841)	(0.485)	(0.403)	(0.294)
Living in a city	0.265	0.011	0.191	0.003
	(0.828)	(0.482)	(0.394)	(0.268)
Living in New South Wales	-1.320	-0.905	-0.199	0.544*
	(0.962)	(0.561)	(0.461)	(0.318)
Highest parental...				
...education=dropout	0.051	0.006	0.099	0.009
	(0.583)	(0.346)	(0.278)	(0.185)
...education=high school	0.194	0.505	0.202	0.104
	(0.751)	(0.437)	(0.356)	(0.238)
...education=degree	-0.613	-0.162	0.202	0.185
	(0.927)	(0.542)	(0.437)	(0.307)
...job=unemployed	0.713	0.098	-0.024	-0.456**
	(0.556)	(0.338)	(0.274)	(0.178)
...job=senior manager	-0.852	-0.457	0.284	-0.097
	(0.830)	(0.485)	(0.391)	(0.276)

Notes: Table presents DiD estimated treatment effects of the Baby Bonus on key child characteristics. That is, we estimate equation 5.1 except using child characteristics as the dependent variable and excluding covariates. LBOTE indicates language background other than English. All windows exclude a fortnight either side of 1 July. Robust standard errors in parentheses. Significance levels indicated as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

5.4.3 Additional child outcomes data

We use the NAPLAN test score data to look at the combined treatment effect of the 1 July 2004 Baby Bonus introduction. However, the time periods spanned by the NAPLAN data prevent us from looking at the re-introduction of Maternity Allowance on 1 February 1996 or the increase in the Baby Bonus on 1 July 2006. To look at the effects of these increases in cash transfers at birth we use two alternative outcome measures.

For Maternity Allowance, we use aggregate data on university offers by precise date of birth. This covers all offers made to Australian residents by Australian universities in either 2013 or 2014 — the years in which the majority of those born around 1 February 1996 will have been completing high school and applying for university. More information on the Higher Education Statistics is available at Department of Education and Training (2016). For the Baby Bonus increase, we use unit record data including precise date of birth from the Australian Early Development Census (AEDC).[29] The AEDC is a triennial census of children in their first year of full-time school. Data are collected from schools and teachers using the Australian version of the Early Development Instrument (AvEDI), adapted from the Canadian Early Development Instrument. The data include information on child development across five domains: physical health and wellbeing; social competence; emotional maturity; language and cognitive skills; and communication skills and general knowledge. Children have scores from 0-10 in each domain — a higher score indicates a higher level of development. We sum these scores to give a single outcome metric.[30] More information on the AEDC is available in Australian Government (2013).

[29]The AEDC is funded by the Australian Government Department of Education and Training. The findings and views reported are those of the authors and should not be attributed to the Department or the Australian Government.

[30]This combined score is coded as missing if any one of a child's five domain scores is missing or flagged as invalid. Domain scores are flagged as invalid for children who have been in the class for less than a month, are less than four years old, have special needs or where teachers complete less than 75 per cent of the items in any given domain. This affects less than 7 per cent of children. We also conducted the analysis keeping 'invalid' scores — this had no meaningful impact on the results.

5.5 Results

We begin with the difference-in-differences estimates of the treatment effect of the additional family income from the Baby Bonus, including estimates for population subgroups and alternative specifications. We drop the data for one fortnight on either side of the threshold to eliminate the effect of birth shifting and to identify the effect of the additional family income.

We then extend the work of Gans and Leigh (2009) looking at birth shifting in response to the introduction of the Baby Bonus. This provides a better understanding of the potential treatment applying immediately after 1 July 2004 — additional family income, hospital crowding, and differences in infant health and family background. We use a regression discontinuity approach to assess whether this was associated with any effect on child outcomes. We also test for discontinuities in outcomes associated with the 1 February 1996 Maternity Allowance reintroduction and 1 July 2006 Baby Bonus increase.

5.5.1 Child outcomes: difference-in-differences

In Figure 5.2 we plot average test scores by date of birth for the treatment and control groups in fortnightly bins either side of 1 July. This provides an immediate feel for the results that follow and allows a visual inspection of the assumptions underlying the difference-in-differences design. Note the inverted U-shape apparent for both treatment and control groups simply reflects the well known age-at-test effect — older children tend to get better results. In our sample, the interaction of school entry policies and redshirting behaviours result in average age-at-test following an inverted U-shape, see Figure 5.10 in Appendix 5.A. Similarly, the apparent discontinuity between April and May reflects a discontinuity in age induced by the school starting age cut off in the state of Victoria and in the Australian Capital Territory.

Figure 5.2 provides some further comfort with respect to the common trends assumption required for the difference-in-differences design. Here this requires the test score by date of birth profiles to be parallel, aside from any difference emerging with the introduction of the policy. This assumption appears reasonable in the

Figure 5.2: Mean test score by date of birth

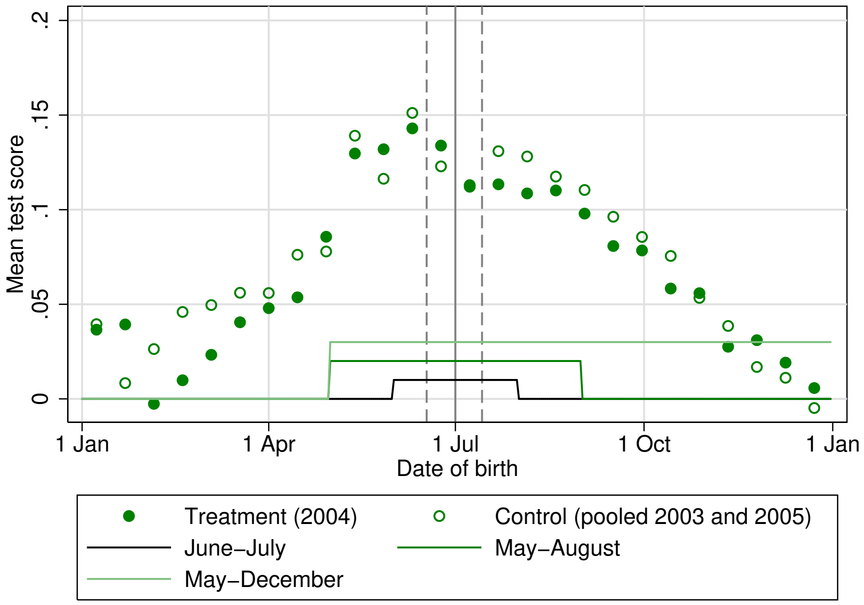

Notes: means are calculated for fortnightly bins either side of 1 July and dashed lines indicate the exclusion window — a fortnight either side of 1 July.

May to December window — the test scores profiles have the same basic shape. Perhaps the clearest violations are in the beginning of the year (and we have already rejected using a full year window of analysis) and around the introduction of the Baby Bonus.

A visual inspection of Figure 5.2 leaves mixed impressions as to any impact of the Baby Bonus. There is a hint of a *negative* discontinuity around the 1 July threshold, which is investigated further in Section 5.5.2. Yet when we consider wider windows, scores in the treatment year (on average a little lower than those in the control) do not appear to have been shifted by the introduction of the Baby Bonus.

Table 5.5 provides difference-in-differences estimates of the effect of the additional family income provided by the Baby Bonus over progressively wider windows of analysis. All regressions exclude those born a fortnight either side of introduction and report estimates with and without covariates. Heteroskedastic-robust standard errors, clustered within schools, are reported alongside the results. The

Table 5.5: Estimated treatment effects of the Baby Bonus

Window	Jun-Jul		May-Aug		May-Dec	
	(1)	(2)	(3)	(4)	(5)	(6)
Covariates	No	Yes	No	Yes	No	Yes
β_3	-0.009	-0.014	-0.017	-0.014	-0.005	-0.002
	(0.020)	(0.018)	(0.012)	(0.010)	(0.009)	(0.008)
N	46,349	45,971	132,531	131,477	306,178	303,630
R-squared	0.000	0.201	0.000	0.199	0.001	0.193

Notes: Table presents DiD estimated treatment effects of the Baby Bonus. We estimate equation 5.1 with varying window widths and with and without controls. The controls include: sex; location (dummy variables for each state or territory and for whether the child lives in a metropolitan area); language background other than English; state specific cubic terms in the child's age-at-test; and categorical variables for parental education and occupation. Robust standard errors in parentheses. Significance levels indicated as follows: *** p<0.01, ** p<0.05, * p<0.1.

results confirm the visual impression left by Figure 5.2. All point estimates are small and none are statistically different from zero. In an encouraging sign for the validity of the design, none of the estimates are statistically different from each other — violations of the common trends assumption could have seen the window of analysis and/or the inclusions of covariates leading to significantly different estimates.

The precision of our estimates allows us to reject even small negative or positive effects of the Baby Bonus. Using the treatment effect in column (6) as our benchmark, a 95% confidence interval for the treatment effect of the Baby Bonus is $[-0.018, 0.015]$. Thus we can reject at the 5% level the hypothesis that the Baby Bonus shifted average test scores by more than 1.8% of a standard deviation lower or 1.5% of a standard deviation higher.

Subpopulation estimates

It may not be surprising that the small, temporary boost to family income provided by the Baby Bonus had no discernible effect across the population as a whole. Improved outcomes are more plausible in disadvantaged households with lower household incomes and more significant credit constraints. To test for effects here, we repeat regressions (2), (4) and (6) from Table 5.5 for different population subgroups: those where the highest level of parental education is high school dropout, those where the highest level of parental education is high school (including dropouts), those with a university degree, and those where the highest status parental occupation is not employed, and those where the highest status parental occupation is senior management.[31] The results are in Table 5.6. Covariates are included in all regressions to increase precision.

The results are mixed, though they are consistent with the introduction of the Baby Bonus improving outcomes for children from disadvantaged backgrounds. While only one of the point estimates is statistically significant, the treatment effects for those from disadvantaged backgrounds are typically positive and larger than those from more advantaged backgrounds. For the children of those with no more than a high school education, the estimated treatment effect on test scores is 4% of a standard deviation, and statistically significant at the 5% level. These children make up 19% of the sample and have an average percentile rank of 36.8. That said, the effects are still relatively modest in size.[32] A caveat on the statistical significance of these results is that the uncertainty in transforming raw grades to NAPLAN scores is not accounted for: true standard errors will be slightly larger than presented here.

Sensitivity analysis

We test the sensitivity of the above results in a number of ways. In Table 5.7, we repeat regression (3) from Table 5.6 under alternative specifications. In (1), all those born in June and July are dropped to allow for birth shifting over a wider window. In (2) and (3), the control year of birth is restricted to 2003 or 2005 respectively. Finally, in (4) and (5), results are estimated separately for New

[31] Differences by migrant status or family structure may also be of interest, but are not as clearly identified in the data.

[32] While the relatively more disadvantaged group of children with high school dropout parents has a smaller and less significant treatment effect, it is not statistically significantly different from that for the larger group that includes high school graduates. In addition, children from severely disadvantaged backgrounds may face challenges that are not as easily ameliorated by additional family income as those a little less worse off.

Table 5.6: Estimated treatment effects of Baby Bonus policy: subpopulations

Window	Jun-Jul (1)	May-Aug (2)	May-Dec (3)
Panel A: Full sample			
β_3	-0.014	-0.014	-0.002
	(0.018)	(0.010)	(0.008)
N	45,971	131,477	303,630
R-squared	0.201	0.199	0.193
Panel B: Highest parental education = high school dropout			
β_3	0.031	0.029	0.029
	(0.053)	(0.031)	(0.025)
N	4,617	13,246	29,920
R-squared	0.081	0.074	0.072
Panel C: Highest parental education = high school or less			
β_3	0.004	0.022	0.040**
	(0.039)	(0.023)	(0.018)
N	8,383	24,002	54,566
R-squared	0.119	0.114	0.112
Panel D: Highest parental education = university degree			
β_3	-0.009	-0.019	-0.002
	(0.030)	(0.018)	(0.014)
N	17,223	49,221	115,725
R-squared	0.093	0.086	0.079
Panel E: Highest parental occupation = not employed			
β_3	-0.041	0.040	0.033
	(0.059)	(0.035)	(0.028)
N	3,891	11,236	25,324
R-squared	0.143	0.136	0.128
Panel F: Highest parental occupation = senior manager			
β_3	0.018	0.002	0.009
	(0.036)	(0.021)	(0.017)
N	11,812	33,455	78,481
R-squared	0.119	0.118	0.108

Notes: Table presents DiD estimated treatment effects of the Baby Bonus. We estimate equation 5.1 for different population subgroups and window widths. Covariates are included in all regressions and include: sex; location (dummy variables for each state or territory and for whether the child lives in a metropolitan area); language background other than English; state specific cubic terms in the child's age-at-test; and categorical variables for parental education and occupation. Robust standard errors in parentheses. Significance levels indicated as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table 5.7: Key estimated treatment effects under alternative specifications.

	Robustness test				
	(1)	(2)	(3)	(4)	(5)
Exclude	± month	± fortnight	± fortnight	± fortnight	± fortnight
Control	pooled	2003	2005	pooled	pooled
States	all	all	all	NSW	VIC
Panel A: Full sample					
β_3	-0.001	-0.004	0.000	-0.009	0.002
	(0.010)	(0.010)	(0.010)	(0.012)	(0.014)
N	257,659	196,975	205,968	147,821	109,495
R-squared	0.192	0.192	0.198	0.217	0.163
Panel B: Highest parental education = high school dropout					
β_3	0.014	0.018	0.043	0.025	0.002
	(0.031)	(0.030)	(0.029)	(0.036)	(0.042)
N	25,303	19,934	19,690	14,814	10,044
R-squared	0.072	0.077	0.070	0.069	0.058
Panel C: Highest parental education = high school or less					
β_3	0.039*	0.035	0.046**	0.051*	0.007
	(0.022)	(0.022)	(0.021)	(0.027)	(0.030)
N	46,183	36,100	36,315	25,692	19,543
R-squared	0.112	0.111	0.119	0.116	0.091
Panel D: Highest parental education = university degree					
β_3	-0.005	-0.001	-0.002	-0.014	0.008
	(0.017)	(0.017)	(0.016)	(0.021)	(0.022)
N	98,502	73,512	80,331	55,886	44,373
R-squared	0.077	0.081	0.080	0.089	0.069
Panel E: Highest parental occupation = not employed					
β_3	0.048	0.017	0.048	0.020	-0.006
	(0.034)	(0.034)	(0.032)	(0.044)	(0.041)
N	21,433	16,219	17,323	11,120	11,240
R-squared	0.127	0.131	0.137	0.123	0.124
Panel F: Highest parental occupation = senior manager					
β_3	0.009	0.020	-0.002	0.006	0.003
	(0.020)	(0.019)	(0.019)	(0.024)	(0.027)
N	66,669	50,240	53,871	39,488	27,463
R-squared	0.107	0.112	0.110	0.130	0.083

Notes: Table presents DiD estimated treatment effects of the Baby Bonus. We estimate equation 5.1 for different population subgroups and choices of: the exclusion window; control birth cohort; and the sample (all states, New South Wales or Victoria). Covariates are included in all regressions and include: sex; location (dummy variables for each state or territory and for whether the child lives in a metropolitan area); language background other than English; state specific cubic terms in the child's age-at-test; and categorical variables for parental education and occupation. Robust standard errors in parentheses. Significance levels indicated as follows: *** p<0.01, ** p<0.05, * p<0.1.

South Wales and Victoria — the two most populous states in Australia, accounting for 85% of our sample. This latter test is particularly useful as the two states have different education policies (including school entry cut off dates).

The results are mixed. First, the estimates change little when both June and July are excluded. For the remaining specifications the main finding of no statistically significant effect (and an ability to reject even modest effects) is confirmed. The suggestive evidence of a positive effect for those from disadvantaged backgrounds is more tenuous: the case for such an effect is weaker when the control year of birth is 2003 and is absent in the Victorian subsample. The earlier effects could thus be driven by unobserved changes in education practices in New South Wales, for example, rather than the introduction of Baby Bonus. An alternative explanation is that there are other differences in state policies that make the Baby Bonus more beneficial in New South Wales compared to Victoria, for example differences in social housing policy or other support services.

As a final test, we extend the earlier exercise estimating treatment effects for a placebo policy introduced on 1 July 2005 in Appendix 5.B, Table 5.10. This time we report results for reading, numeracy and combined test scores by population subgroup for a May to December window of analysis. It is possible the initial failure to find significant effects for the placebo policy only held for the separate tests and/or the full population, yet none of the estimates are statistically significant, which provides some support for the validity of our design.

5.5.2 Descriptive analysis of birth shifting events

Differences in family background and infant health

Policy-induced birth shifting can change a child's start in life through hospital crowding and infant health. Birth shifting may also result in differences in family background either side of the policy threshold. These differences and any long run repercussions shine a light on potential unintended harms from birth shifting and, perhaps more benignly, any inequities resulting from differing abilities or propensities to schedule births.

Large birth shifting in response to the Baby Bonus has already been documented by Gans and Leigh (2009) — there were more births on 1 July 2004 than any prior day in Australian history. As reported in Table 5.1, it is estimated that some 10% and 6% of births were shifted across the fortnights either side of the introduction and increase to the Baby Bonus respectively. For the reintroduction of Maternity Allowance, we find less birth shifting with 4% of births moved. These are all sizeable increases in demand on the health system.

Birth shifting also has direct, mechanical consequences for infant health: delays, all else being equal, lead to longer gestations and higher birthweights. Figure 5.3 shows mean gestation length and birthweight for all births in 2004. There is a clear spike in birthweight in the fortnight following the Baby Bonus' introduction. Interestingly, gestation length appears to rise either side of the introduction of the policy. This could be consistent with unsuccessful attempts to shift births into July, or providers delaying procedures to fill scheduling gaps left by those successfully shifting.

Figure 5.3: Mean birthweight and gestation length by date of birth (2004)

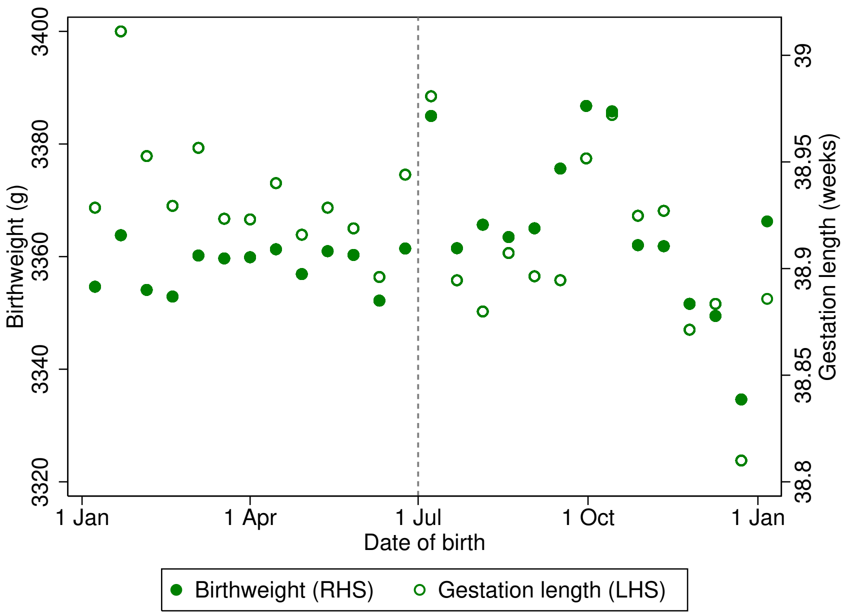

Notes: means are calculated for fortnightly bins either side of 1 July 2004. As control cohorts are not needed for the regression discontinuity design, they are not plotted here.

Table 5.8 presents estimates of the discontinuity in expected parent and child characteristics at 1 July 2004. Panels A and B use the demographic variables in the NAPLAN data. Panel C uses infant health variables aggregated across child date of birth cohorts. In both cases heteroskedastic-robust standard errors are reported. Standard errors are clustered by date of birth in the first two panels, while in the second the regression is weighted by the number of births in each date of birth cohort.

The results are consistent with and extend earlier findings. Gans and Leigh (2009) find only modestly significant differences in parental age and none in mean income of parental postcode. Our results suggest that, if anything, parents with more education or higher status occupations were marginally less likely to shift births. This could have gone either way. Well-off parents may be more informed about policy and better placed to negotiate with physicians. On the other hand, they may be less likely to take risks in delaying a birth for the sake of a relatively small amount of money.[33]

There are larger and more statistically significant differences in other parent and child characteristics. Those born on 1 July are less likely to live in a metro area and have a language background other than English, and more likely to be female. Perhaps those outside cities were better able to shift births due to greater slack in the hospital system, while those without an English language background were less informed or less able to negotiate with physicians. The rise in the proportion of females may be due to sex differences in fetal health making physicians more willing to delay female births.[34] In any case, a complex set of factors likely determines whether parents have the incentives, information, willingness and ability to shift their child's birth, so it is not surprising to see differences in some areas but not others.

The primary interest in birth shifting is the extent to which it influences, and potentially harms child health. Panel C shows clear effects of birth shifting on infant health. Those born immediately after the Baby Bonus had higher birthweights and were more likely to be of high birthweight, with an increase in expected birthweight of 30-60 grams and an increase in the expected probability of

[33]Given the introduction of the Baby Bonus was, as discussed earlier, not a uniform increase in cash transfers at birth in *net* terms, well-off parents also had less incentive in absolute terms to shift births.

[34]Selective survival is another explanation, although Gans and Leigh (2009) fail to find any effect of birth shifting from the Baby Bonus on infant mortality.

high birthweight of 1-2 percentage points.[35] While there is no statistically significant increase in mean gestation length, this may reflect the fact that gestation lengths were rising prior to the introduction of the Baby Bonus. The estimates are nonetheless of the expected signs, consistent with longer gestations.

These differences in infant health are modest. They also apply to those whose birth dates were a conscious choice — not those born very prematurely, for example. Thus, despite the concerns raised by Gans and Leigh (2009) on this front, it would be surprising if they had any effects on later life outcomes. To our minds, the more plausible mechanism for harms arising from birth shifting is that hospital congestion may lead to poorer care, for example, due to rationing and physician fatigue. This only requires that physicians fail to account the negative externalities when scheduling a birth (rather than also ignoring effects on the child).

Differences in later life outcomes

There are clearly differences in the health settings, health outcomes and families of those born before and after the introduction of the Baby Bonus. We now test for differences in later life outcomes. We do this by testing for a discontinuity in test scores either side of 1 July 2004. To complement this analysis, we investigate two other major episodes of birth shifting — the reintroduction of Maternity Allowance on 1 February 1996 and the increase to the Baby Bonus on 1 July 2006.

The results are foreshadowed in Figure 5.4 showing the size and later life outcomes of birth cohorts around each policy change. Although it resulted in the smallest share of births moved, the introduction of Maternity Allowance is associated with a relatively large discontinuity in the university offer rate — those born immediately after appear significantly less likely to be offered a place at university.[36] The impression left by the average NAPLAN test scores around the introduction of the Baby Bonus is less clear, and if anything average AEDC development scores are higher for births following the increase to the Baby Bonus.

[35] Some of this difference may be explained by the more advantaged families having higher birthweight children, on average, as well as the mechanical increase in birthweight expected from longer gestations.

[36] The university offer rate is, more precisely, the number of university offers for those born on a given day as a proportion of the original size of that birth cohort. We are unable to adjust for migration.

Table 5.8: Parent and child characteristics

Window	± 7 days	± 14 days	± 21 days	± 28 days
	(1)	(2)	(3)	(4)
Panel A: Parent highest education or occupation				
Dropout	0.002	0.002	0.009	0.007
	(0.012)	(0.010)	(0.009)	(0.007)
Degree	0.004	-0.021*	-0.013	-0.012
	(0.014)	(0.011)	(0.011)	(0.010)
Not employed	-0.002	-0.008	-0.005	-0.002
	(0.006)	(0.007)	(0.006)	(0.006)
Senior manager	0.009	-0.010	-0.001	-0.001
	(0.017)	(0.013)	(0.011)	(0.009)
Panel B: Other parent and child characteristics				
Metropolitan	-0.032**	-0.030***	-0.024**	-0.015*
	(0.015)	(0.011)	(0.009)	(0.009)
LBOTE	-0.049***	-0.044***	-0.030***	-0.025***
	(0.010)	(0.007)	(0.007)	(0.007)
Female	0.023**	0.023***	0.017**	0.016***
	(0.009)	(0.007)	(0.007)	(0.006)
Panel C: Child cohort characteristics				
Mean birthweight (g)	62.5*	54.0**	30.4*	40.4**
	(26.0)	(20.1)	(16.8)	(15.1)
Low birthweight share	-0.016	-0.003	0.004	-0.002
	(0.010)	(0.009)	(0.007)	(0.006)
Normal birthweight share	0.005	-0.017	-0.016*	-0.008
	(0.021)	(0.012)	(0.009)	(0.007)
High birthweight share	0.011	0.019*	0.012*	0.010
	(0.020)	(0.010)	(0.007)	(0.007)
Mean gestation length (weeks)	0.025	0.031	0.044	0.044
	(0.065)	(0.056)	(0.042)	(0.041)
Pre-term share	-0.016	-0.009	-0.007	-0.006
	(0.018)	(0.012)	(0.008)	(0.007)
Term share	0.011	0.009	0.003	0.001
	(0.021)	(0.013)	(0.009)	(0.008)
Post-term share	0.005	0.002	0.003	0.004
	(0.004)	(0.003)	(0.002)	(0.002)

Notes: Table presents RD-DiD (Panels A and B) and RD (Panel C) estimated differences in expected parent and child characteristics either side of 1 July 2004. We estimate equations 5.3 and 5.2 for the relevant characteristic, including covariates as discussed in Section 5.3. Robust standard errors in parentheses. Significance levels indicated as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Figure 5.4: Birth cohort sizes (lines) and later life outcomes (dots) around three policy changes.

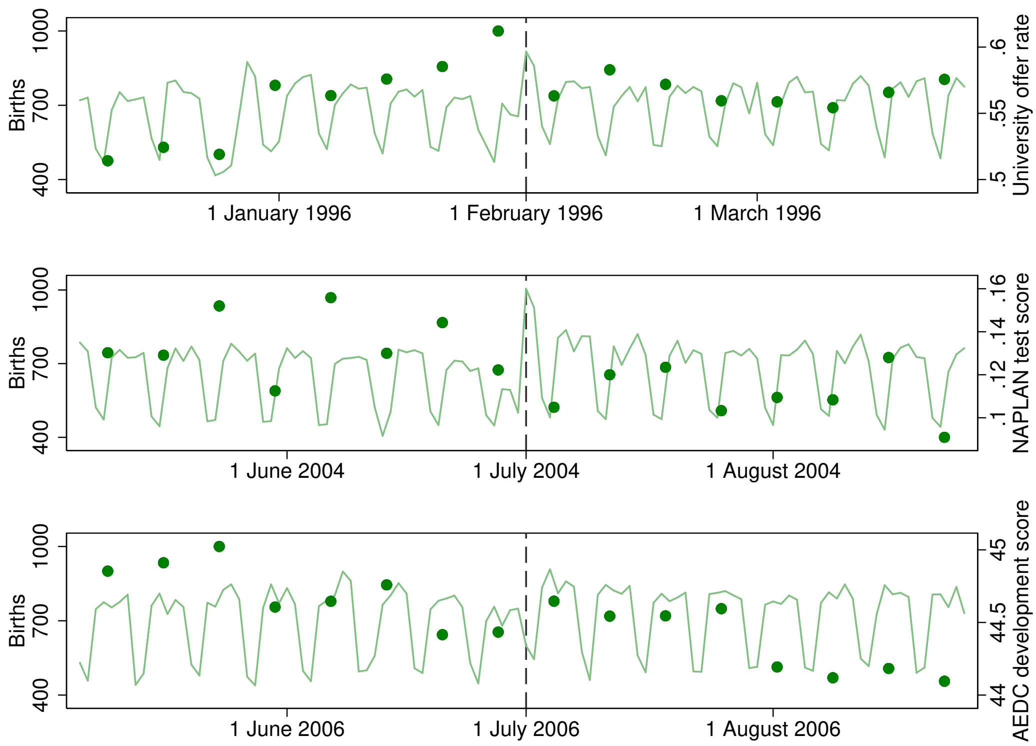

Note: means shown for the outcome variables in weekly bins either side of the policy change. In the first panel, 1 January birth dates are excluded from calculation of the mean due to their significant overrepresentation relative to registered births — this could be due to migrants with known birth years but unknown birth dates assuming 1 January of their relevant date of birth.

The intuitive feel for the results from Figure 5.4 is mostly born out by the regression discontinuity estimates in Figure 5.5. The introduction of Maternity Allowance is associated with a fall in the probability of receiving an offer of a university place by over 5 percentage points.[37] The effect is significant at the 5 per cent level. There is fall in expected NAPLAN test scores of 5% of a standard deviation associated with the introduction of the Baby Bonus, which is just significant at the same level. The effect on average AEDC development scores is not

[37]The discontinuity in university offers is entirely due to a discontinuity in university applications, rather than in the rate of offers per application, as apparent in Appendix 5.A, Figure 5.11. This suggests no change in ability relative to aspirations at the time of applying to university.

significant.[38]

These mixed results are not easy to interpret, particularly as they relate to different policy changes. Unfortunately, we do not have a common outcome measure that we can employ for all three policy changes. Also, it is not possible to disentangle the effects of differences in family background and birth shifting. For example, if well-off parents are less likely to shift births then the discontinuities in both test scores and in university offers for their children may simply reflect the intergenerational transmission of ability and aspirations.

[38]We also tested for a discontinuity in AEDC development scores separately across the five developmental domains, and again failed to find clear evidence of a discontinuity.

Figure 5.5: Discontinuities in outcomes around three policy changes

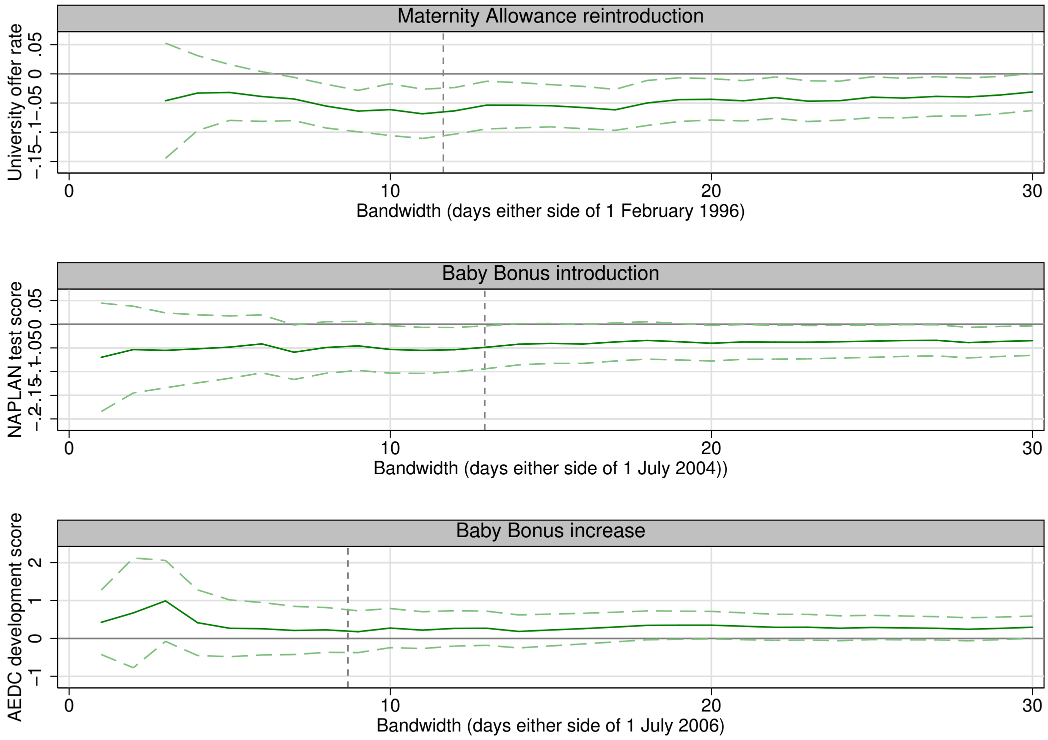

Notes: presents RD/RD-DiD estimated differences in expected child outcomes either side of three major birth shifting events (on 1 February 1996, 1 July 2004 and 1 July 2006). We estimate equation 5.2 or 5.3 for the relevant outcome, including covariates as discussed in Section 5.3 over a range of bandwidths. Vertical line indicates rule of thumb bandwidth.

5.6 Conclusion

Cash transfers to families are a central part of many governments' responses to poverty and an important component of social welfare expenditure. These payments have a variety of motives, including ameliorating current economic hardship and improving child outcomes.[39] However, it is hard to assess the efficacy of cash

[39]There are of course other policy arguments for cash transfers to families with children, such as supporting higher fertility rates or simple redistribution to those bearing the costs of raising children. A fuller examination of these, while necessary for a full welfare analysis of such programs, is beyond the scope of this paper.

transfers in improving child outcomes in the absence of large-scale experiments. This makes natural experiments a crucial source of evidence.

In this paper, we use a difference-in-difference design to assess the effect of the additional family income provided by the Baby Bonus on child outcomes, here measured by a combined reading and numeracy test score in grade 3. We find no evidence that the AUD3,000 of additional family income improved outcomes in aggregate. There is some evidence for a modest effect on children from disadvantaged backgrounds — around 4% of a standard deviation — but this effect is not apparent across all specifications. This research complements the existing literature that finds relatively modest aggregate effects of family income on child outcomes. These effects are largest among children from disadvantaged backgrounds, as in Dahl and Lochner (2012), and when provided early in life, as in Morris et al. (2005). However, our results caution against an overly optimistic interpretation of this literature — such as the potential for higher payoffs for payments provided earlier in life — as the precision of our estimates allows us to reject even small aggregate effects.

There are a variety of explanations for these findings. The first is simply that the boost to family income was small — a one-off increase of around 1 to 2 weeks average earnings. Despite the large body of evidence on the importance of early life experiences to later outcomes (see Currie and Almond (2011)), in a wealthy country with high quality public services, such as Australia, the *average* returns to *additional* investments in early childhood may be relatively low. Alternatively, even if such investment opportunities exist, parents may be unaware, unwilling or unable to use the additional funds to make them.

Our results suggest that in Australia, any benefits of payments at birth to a child's later cognitive skills are extremely small in aggregate. They also suggest that if governments want to have substantial impact on medium- and long-term outcomes for children that this may require a substantial commitment of resources. Of course, as noted earlier there are other potential motivations for family benefits, including as a form of redistribution and or a recognition of the costs of parenting.

The unprecedented birth shifting that accompanied the Baby Bonus triggered significant academic and media concern — most notably that hospital congestion would lead to poorer care and infant health outcomes. Our results on this point are inconclusive. Most compelling are the complex set of differences in parent and child characteristics either side of policy changes. Looking for discontinuities around three large policy changes, we find marginally better early development indicators, but significantly lower test scores and rates of university offers. Since we also find evidence of negative selection into birth shifting – less well-off parents are more likely to shift births – these latter results could simply reflect unobservable differences in family background.

This mix of findings suggests sharp date-of-birth cut offs are best argued against on a precautionary basis, if at all – there is no ‘smoking gun’ indicating long run harms for policy changes of the sizes examined here. It certainly appears possible that parents and health professionals navigated the largest recorded birth shifting events in Australian history without any harmful consequences.

5.A Additional charts

Figure 5.6: Binned mean test score by month of birth: four largest states

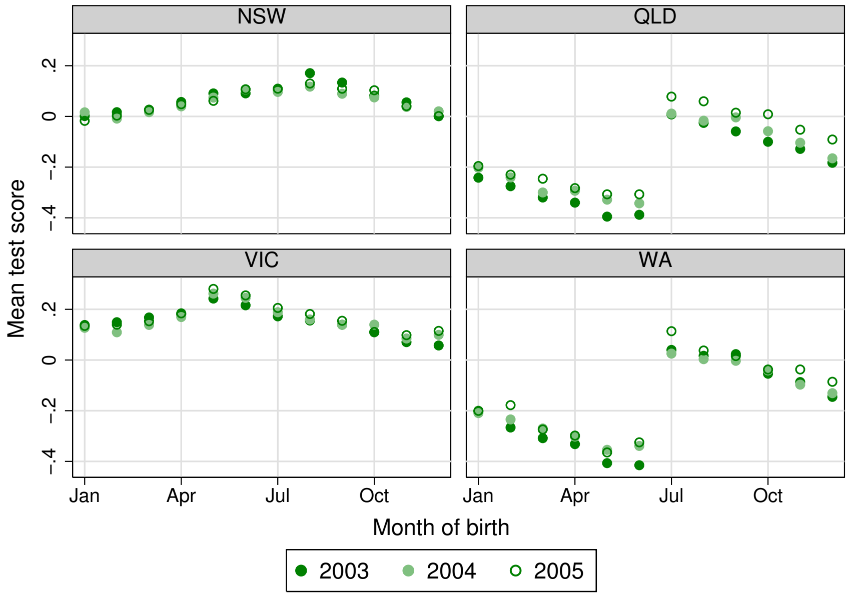

Note: Outcome variable is the summed reading and numeracy score, normalised across all years and states and territories to have mean zero and standard deviation one.

Figure 5.7: Distribution of individuals across test years by date of birth

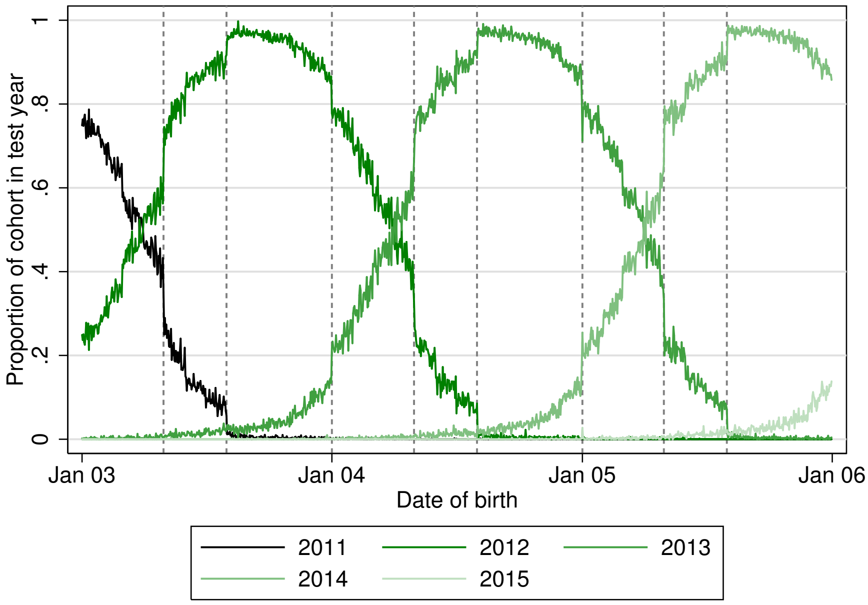

Notes: Excludes jurisdictions with 1 July cut off dates. Vertical lines indicate remaining cut off dates for school entry (30 April, 31 July and 1 January) and correspond to discontinuous jumps in the probability of sitting the test in a later year.

Figure 5.8: Binned mean test score by month of birth: test components

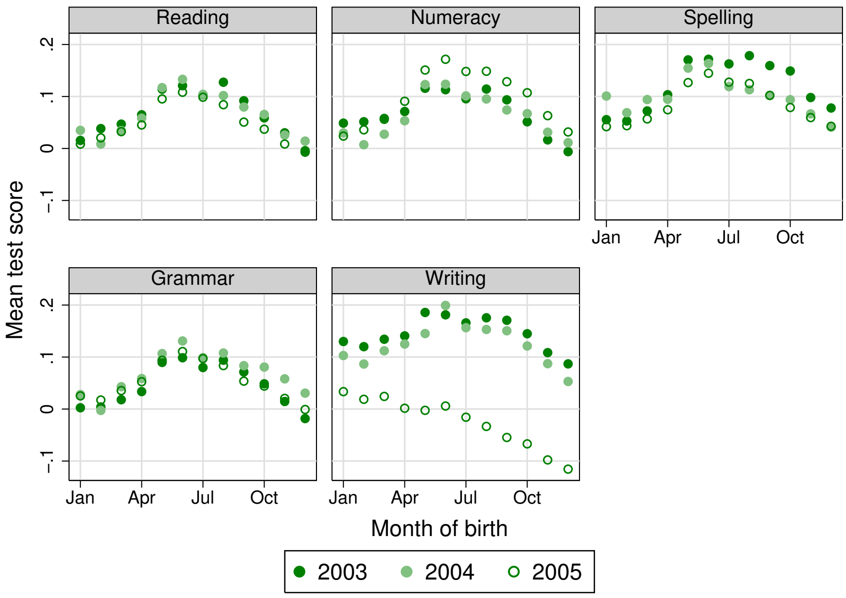

Note: Excludes those states and territories with 1 July cut off dates.

Figure 5.9: Missing values by year of birth

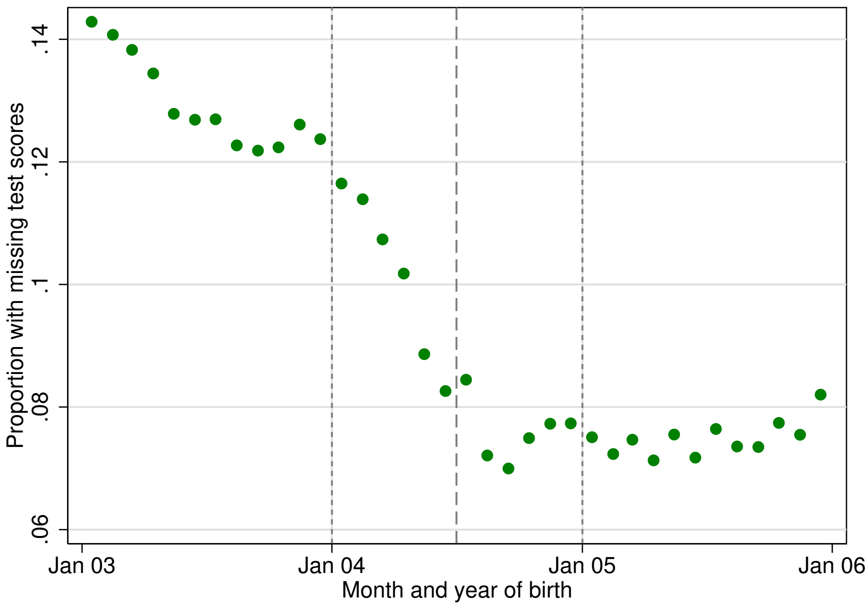

Note: excludes those states and territories with 1 July cut off dates. The higher rate of missing values for earlier months likely reflects the increasing share in earlier birth cohorts taking the test in 2011 or 2012, where we infer the Year 3 test scores from the past test scores of Year 5 students two years later. For example, for 2013 we have both direct data on Year 3 test scores and indirect data from the past test scores of Year 5 students in 2015. The rate of missing values in the direct data is 7.6% compared to 13.6% in the indirect data. The higher rate of missing values in the indirect data could stem from: a failure to link an individual to their past test scores (so past scores are recorded missing, despite actually existing); and migration resulting in the Year 5 population being larger than the actual eligible population of Year 3 students two years prior. In either case, the trends in missing values will, if anything, bias estimates downwards. As a further check, we estimate regression (3) from Table 5.7 restricting attention to those sitting the tests in the year in which they turned 9 or 10. This ensures only direct test data is used. The estimates did not change in a substantive way (results available on request).

Figure 5.10: Mean age-at-test by date of birth

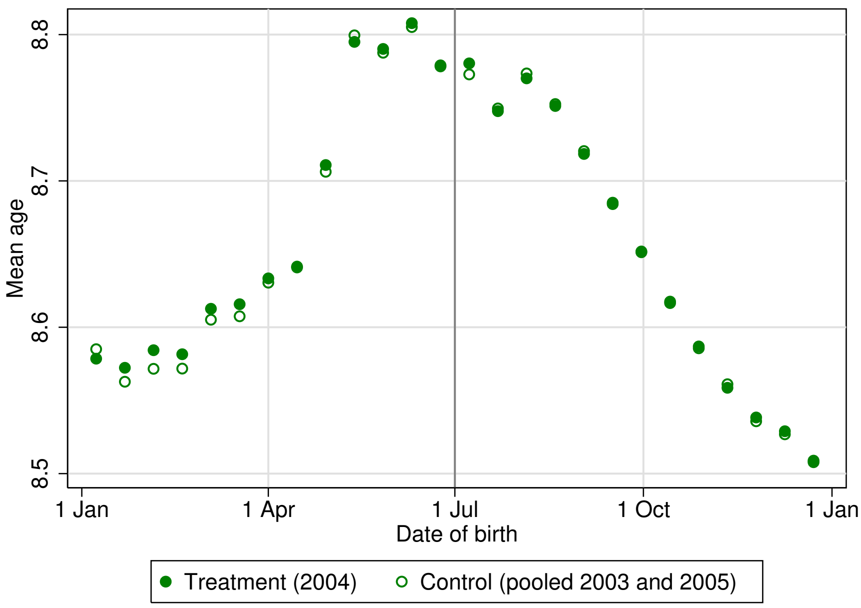

Note: Excludes those states and territories with 1 July cut off dates.

Figure 5.11: Decomposition of effect of Maternity Allowance reintroduction

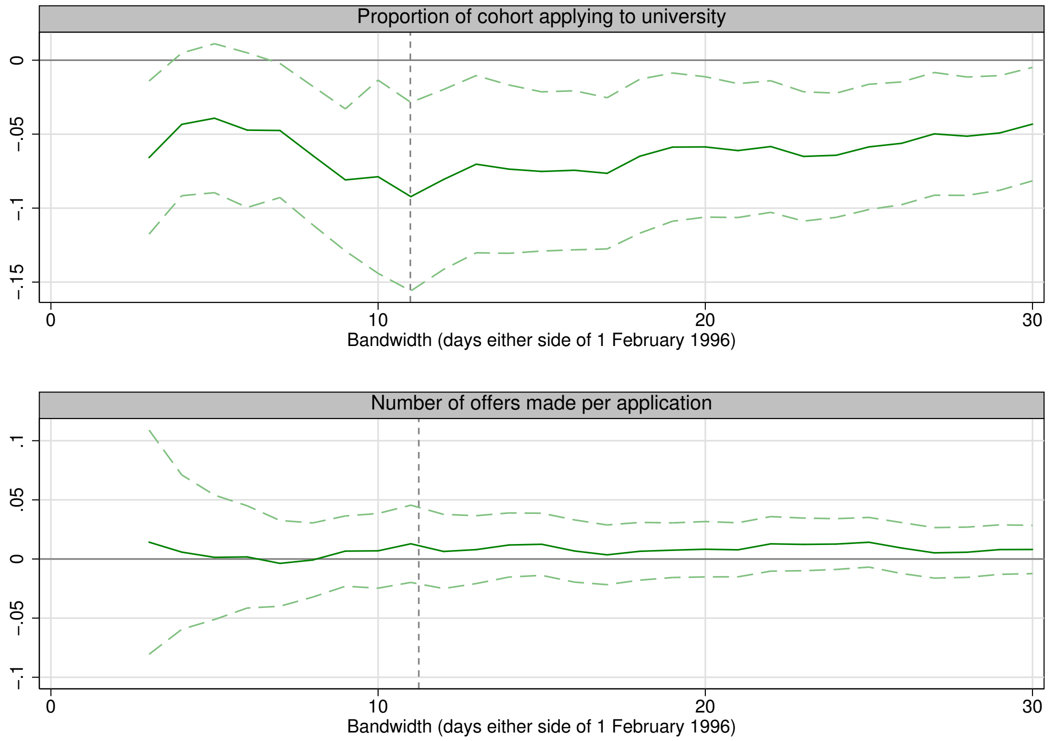

Notes: presents RD estimated differences in expected child outcomes either side of the Maternity Allowance reintroduction on 1 February 1996. We estimate equation 5.2 for the relevant outcome, including covariates as discussed in Section 5.3 over a range of bandwidths. Vertical line indicates rule of thumb bandwidth.

5.B Additional tables

Table 5.9: Distribution of First Child Tax Refund claims and (hypothetical) net gain from Baby Bonus introduction (June 2004 births).

	Mean	Percentile				
		10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Panel A: All carers						
Refund	\$1,280	0	0	\$485	\$2,266	\$3,348
Net gain	\$877	\$2,157	\$2,157	\$1,672	-\$109	-\$1,191
Panel B: Carers in bottom quintile by base income						
Refund	\$1,093	0	0	\$957	\$2,269	\$2,282
Net gain	\$1,064	\$2,157	\$2,157	\$1,200	-\$112	-\$125
Panel C: Carers in top quintile by base income						
Refund	\$2,179	0	0	\$331	\$3,907	\$6,958
Net gain	-\$22	\$2,157	\$2,157	\$1,826	-\$1,750	-\$4,801

Notes: All figures are expressed in September quarter 2004 dollars and calculated from Australian Taxation Office data. Base income is the primary carer's income in the year prior to the birth of the child. Net gain is the \$2,157 increase in cash transfers (the \$3,000 Baby Bonus minus the \$843 Maternity Allowance) minus the refund claim. The net gain will be an underestimate for some families who were not have eligible for Maternity Allowance. We can not determine Maternity Allowance eligibility from the aggregated tax data.

Table 5.10: Estimated treatment effects of placebo policy by population subgroups.

	All	Dropout	High school or less	Degree	Not employed	Senior manager
	(1)	(2)	(3)	(4)	(5)	(6)
Reading	-0.008	-0.039	-0.014	-0.002	-0.030	0.022
	(0.010)	(0.030)	(0.022)	(0.016)	(0.034)	(0.019)
Numeracy	-0.000	-0.008	-0.006	0.004	-0.038	0.018
	(0.010)	(0.030)	(0.023)	(0.016)	(0.034)	(0.020)
Combined	-0.005	-0.026	-0.012	0.001	-0.033	0.022
	(0.010)	(0.030)	(0.022)	(0.016)	(0.034)	(0.019)

Notes: DiD estimated treatment effects of a placebo policy introduced on 1 July 2005. We estimate equation 5.1 using the 2005 birth cohort as the treatment and the 2003 birth cohort as the control. The May-December window width, excluding a fortnight either side of 1 July is used. Covariates include: sex; location (dummy variables for state/territory and metropolitan area); language background other than English; state specific cubic terms in the child's age-at-test; and categorical variables for parental education and occupation. Robust standard errors in parentheses. Significance levels indicated as follows: *** p<0.01, ** p<0.05, * p<0.1

5.C First Child Tax Rebate

The First Child Tax Refund was a refundable tax offset that was designed to pay back to mothers the tax paid on their income (up to a limit) in the year prior to having their first child. (In Australia, the unit of taxation is the individual.) The primary carer of a first child born between 1 July 2001 and 30 June 2004 could claim the refund at the end of each of the first five financial years following the birth. The base annual refund was equal to a fifth of the income tax paid in the financial year prior to the birth of the child, bounded by \$500 below and \$2,500 above. This was then scaled down in proportion to the time the child was in their care and the degree to which the primary carer's taxable income I_t remained below that in the year before the child was born I_0. Thus, the refund in year t was:

$$\text{(base annual refund)} \cdot (1 - I_t/I_0).\text{(days in claimant's care/365)} \tag{5.4}$$

As a result, families could receive anywhere between nothing and \$12,500 through the refund. In practice, there were 736,000 births registered in Australia in the eligibility window for the refund and an average refund paid of \$1,793 for each birth. As the refund only applied to the first birth after 1 July 2001, average refunds fell over time. In the month prior to the introduction of the Baby Bonus the mean real refund paid over the following five years was \$1,280 while the median was only \$485 (Table 5.9).

As refunds differed substantially according to individual circumstances, it might be expected that the average refund would increase prior to its abolition as individuals sorted across the threshold date to qualify for the Baby Bonus or the First Child Tax Refund according to whichever was most financially attractive. In fact, there is only a very modest increase in the average refund in the days immediately prior to 1 July 2004. This could be due to the refund being: less salient (and harder to calculate) than the bonus; and heavy discounting due to its delayed payment and uncertainty regarding the primary carer's return to work.

Chapter 6

Conclusion

This thesis contributes to the vast literature on intergenerational mobility and early childhood human capital formation. I provide the most detailed and precise picture of intergenerational income mobility in Australia to date. I also provide fresh insights into some of the causal mechanisms at play, in a variety of different settings, thereby contributing to the broader international literature.

I introduce a new Australian intergenerational dataset in Chapter 2, co-authored with Bhashkar Mazumder. We use this to provide a brief overview of intergenerational mobility in Australia, paying special attention to a variety of measurement issues. Australia emerges as a relatively mobile country, though with meaningful geographic variation.

In Chapter 3, I explore whether this geographic variation reflects a causal effect of place or simply differences between families. I follow Chetty and Hendren (2018a) in exploiting variation in the age at which individuals move to identify the causal effect of exposure to place. In doing so, I present the first replication of their highly influential study, but also extend it to examine how exposure effects vary over childhood, and underlying mechanisms, in more detail. I show that place matters most in the teenage years, and present evidence that place effects may be partly driven by local labour market entry and peer effects.

In Chapter 4, I follow a different tradition in economics of using migrants as a way of identifying the potential influences of culture (Fernández (2011)). I show most of the differences in intergenerational income mobility across migrant communities stem from differences in intergenerational education mobility, rather than differential returns to education — though with some notable exceptions. I find evidence of underlying roles for culture and social context influencing educational achievement (test scores) and attainment (years of education) respectively. Perhaps most notably, migrant communities whose fathers were poorer, conditional on their education, tend to see more value in, aspire to, and attain more education.

In Chapter 5, co-authored with Robert Breunig, I examine the causal effect of family income itself on early childhood human capital. We use the introduction of the Australian Baby Bonus — a cash transfer of AUD3,000 to parents of babies born on or after 1 July 2004. The Bonus had the intended effect of boosting family income, but also led to sizeable birth shifting and discontinuities in infant health outcomes across the threshold (Gans and Leigh (2009)). We find no evidence the intended or unintended treatments affected child outcomes in aggregate, and only modest evidence of an effect for those from more disadvantaged backgrounds.

This thesis shows Australia enjoys a high degree of intergenerational mobility. But it also shines a light on the complex mechanisms potentially at play in driving persistence in economic outcomes from one generation to the next. I examine the roles of place, local labour markets, culture and social context, and family income itself using a range of empirical strategies. This thesis is far from definitive; planned improvements in the scope and accessibility of Australian administrative data will allow still deeper explorations of intergenerational mobility. Yet it is nonetheless an important step on that path, greatly improving our knowledge about intergenerational mobility in Australia, and using our experiences to contribute to the international literature.

Bibliography

Aaronson, D. (1998). Using sibling data to estimate the impact of neighborhoods on children's educational outcomes, *Journal of Human Resources* **33**(4): 915–946.

Aaronson, D. and Mazumder, B. (2008). Intergenerational economic mobility in the United States, 1940 to 2000, *Journal of Human Resources* **43**(1): 139–172.

Abramitzky, R., Boustan, L. P. and Eriksson, K. (2014). A nation of immigrants: Assimilation and economic outcomes in the age of mass migration, *Journal of Political Economy* **122**(3): 467–506.

Angrist, J. D. (2014). The perils of peer effects, *Labour Economics* **30**: 98–108.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*, Princeton University Press, Princeton, N.J.

Australian Bureau of Statistics (1991). Census of Population and Housing, 1991, *Cat. no. 2101.0*, ABS, Canberra.

Australian Bureau of Statistics (2010). Family Characteristics Survey, Australia, 2009-10, *Cat. no. 4442.0*, ABS, Canberra.

Australian Bureau of Statistics (2011). Australian Statistical Geography Standard (ASGS): Volume 1 – Main Structure and Greater Capital City Statistical Areas, July 2011, *Cat. no. 1270.0.55.001*, ABS, Canberra.

Australian Bureau of Statistics (2014). *Births, Australia, 2013*, Australian Bureau of Statistics, Canberra. Cat. no. 3301.0.

Australian Bureau of Statistics (2016). Standard Australian Classification of Countries (SACC), 2016, *Cat. no. 1269.0*, Australian Bureau of Statistics, Canberra.

Australian Bureau of Statistics (2017a). Australian Demographic Statistics, June 2017, *Cat. no. 3101.0*, ABS, Canberra.

Australian Bureau of Statistics (2017b). Births, Australia, 2016, *Cat. no. 3301.0*, ABS, Canberra.

Australian Bureau of Statistics (2017c). Consumer Price Index, Australia, March 2017, *Cat. no. 6401.0*, ABS, Canberra.

Australian Bureau of Statistics (2017d). Labour Force, Australia, Detailed, November 2017, *Cat. no. 6291.0.55.001*, ABS, Canberra.

Australian Curriculum, Assessment and Reporting Authority (2015). *National Assessment Program — Literacy and Numeracy 2014: Technical Report*, Sydney.

Australian Government (2013). *A Snapshot of Early Childhood Development in Australia 2012 - AEDI National Report*, Australian Government, Canberra.

Australian Institute of Health and Welfare (2014). Australia's mothers and babies 2012, *Perinatal statistical series no. 30. cat. no. PER 69*, Australian Institute of Health and Welfare, Canberra.

Aydemir, A., Chen, W.-H. and Corak, M. (2009). Intergenerational earnings mobility among the children of Canadian immigrants, *The Review of Economics and Statistics* **91**(2): 377–397.

Aydemir, A., Chen, W.-H. and Corak, M. (2013). Intergenerational education mobility among the children of Canadian immigrants, *Canadian Public Policy* **39**(Supplement 1): S107–S122.

Basu, S. (2018). Age-of-arrival effects on the education of immigrant children: A sibling study, *Journal of Family and Economic Issues (forthcoming)* .

Bell, A. M., Chetty, R., Jaravel, X., Petkova, N. and Van Reenen, J. (2017). Who Becomes an Inventor in America? The Importance of Exposure to Innovation. NBER Working Paper No. 24,062.

Bhattacharya, D. and Mazumder, B. (2011). A nonparametric analysis of black–white differences in intergenerational income mobility in the United States, *Quantitative Economics* **2**(3): 335–379.

Bingley, P., Cappellari, L. and Tatsiramos, K. (2016). Family, community and long-term earnings inequality. IZA Discussion Paper No. 10,089.

Björklund, A., Lindahl, L. and Lindquist, M. J. (2010). What more than parental income, education and occupation? An exploration of what Swedish siblings get from their parents, *The BE Journal of Economic Analysis & Policy* **10**(1).

Björklund, A., Roine, J. and Waldenström, D. (2012). Intergenerational top income mobility in Sweden: Capitalist dynasties in the land of equal opportunity?, *Journal of Public Economics* **96**(5-6): 474–484.

Black, S. E. and Devereux, P. J. (2011). Recent developments in intergenerational mobility, *Handbook of Labor Economics*, Vol. 4, Elsevier, chapter 16, pp. 1487–1541.

Black, S. E., Devereux, P. J. and Salvanes, K. G. (2013). Under pressure? The effect of peers on outcomes of young adults, *Journal of Labor Economics* **31**(1): 119–153.

Bleakley, H. and Chin, A. (2004). Language skills and earnings: Evidence from childhood immigrants, *Review of Economics and Statistics* **86**(2): 481–496.

Borjas, G. J. (1985). Assimilation, changes in cohort quality, and the earnings of immigrants, *Journal of Labor Economics* **3**(4): 463–489.

Borjas, G. J. (1992). Ethnic capital and intergenerational mobility, *The Quarterly Journal of Economics* **107**(1): 123–150.

Borjas, G. J. (1993). The intergenerational mobility of immigrants, *Journal of Labor Economics* **11**(1, Part 1): 113–135.

Borjas, G. J. and Bratsberg, B. (1996). Who Leaves? The Outmigration of the Foreign-Born, *The Review of Economics and Statistics* **78**(1): 165–176.

Borjas, G. J. et al. (1987). Self-Selection and the Earnings of Immigrants, *American Economic Review* **77**(4): 531–553.

Boserup, S. H., Kopczuk, W. and Kreiner, C. T. (2013). Intergenerational wealth mobility: Evidence from Danish wealth records of three generations, *Univ. of Copenhagen mimeo* .

Bratberg, E., Davis, J., Mazumder, B., Nybom, M., Schnitzlein, D. D. and Vaage, K. (2017). A comparison of intergenerational mobility curves in Germany, Norway, Sweden, and the US, *The Scandinavian Journal of Economics* **119**(1): 72–101.

Breunig, R. V. and Carter, A. (2018). Do earned income tax credits for older workers prolong labor market participation and boost earned income? Evidence from Australia's Mature Age Worker Tax Offset. Tax and Transfer Policy Institute Working Paper 15/2018.

Brown, B. B. and Larson, J. (2009). Peer relationships in adolescence, *Handbook of Adolescent Psychology*, Vol. 2, John Wiley & Sons, chapter 3.

Bütikofer, A., Dalla-Zuanna, A. and Salvanes, K. G. (2018). Breaking the links: Natural resource booms and intergenerational mobility. Discussion Paper Series in Economics 19/2018, Norwegian School of Economics, Department of Economics.

Card, D., DiNardo, J. and Estes, E. (2000). The More Things Change: Immigrants and the Children of Immigrants in the 1940s, the 1970s, and the 1990s, *in* G. Borjas (ed.), *Issues in the Economics of Immigration*, National Bureau of Economic Research and University of Chicago Press, Chicago.

Case, A., Fertig, A. and Paxson, C. (2005). The lasting impact of childhood health and circumstance, *Journal of Health Economics* **24**(2): 365–389.

Chadwick, L. and Solon, G. (2002). Intergenerational income mobility among daughters, *American Economic Review* **92**(1): 335–344.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R. and Porter, S. R. (2018). The opportunity atlas: Mapping the childhood roots of social mobility. National Bureau of Economic Research Working Paper No. 25,147.

Chetty, R. and Hendren, N. (2018a). The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects, *The Quarterly Journal of Economics* **133**(3): 1107–1162.

Chetty, R. and Hendren, N. (2018b). The impacts of neighborhoods on intergenerational mobility II: County-level estimates, *The Quarterly Journal of Economics* **133**(3): 1163–1228.

Chetty, R., Hendren, N., Jones, M. R. and Porter, S. R. (2018). Race and Economic Opportunity in the United States: An Intergenerational Perspective. National Bureau of Economic Research Working Paper No. 24,441.

Chetty, R., Hendren, N. and Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment, *American Economic Review* **106**(4): 855–902.

Chetty, R., Hendren, N., Kline, P. and Saez, E. (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States, *The Quarterly Journal of Economics* **129**(4): 1553–1623.

Chiswick, B. R. (1978). The effect of Americanization on the earnings of foreign-born men, *Journal of Political Economy* **86**(5): 897–921.

Chiswick, B. R. and Miller, P. W. (2012). Negative and positive assimilation, skill transferability, and linguistic distance, *Journal of Human Capital* **6**(1): 35–55.

Chiswick, B. R. and Miller, W. (2011). The "negative" assimilation of immigrants: a special case, *ILR Review* **64**(3): 502–525.

Chyn, E. (2018). Moved to opportunity: The long-run effects of public housing demolition on children, *American Economic Review (forthcoming)* .

Cobb-Clark, D. A., Dahmann, S., Salamanca, N. and Zhu, A. (2017). Intergenerational disadvantage: Learning about equal opportunity from social assistance receipt. Life Course Centre Working Paper No. 2017-17.

Cobb-Clark, D. A. and Nguyen, T.-H. (2012). Educational attainment across generations: The role of immigration background, *Economic Record* **88**(283): 554–575.

Cobb-Clark, D. A. and Stillman, S. (2013). Return migration and the age profile of retirement among immigrants, *IZA Journal of Migration* **2**(1): 20.

Corak, M. and Heisz, A. (1999). The intergenerational earnings and income mobility of Canadian men: Evidence from longitudinal income tax data, *Journal of Human Resources* **34**(3): 504–533.

Corak, M. and Piraino, P. (2011). The intergenerational transmission of employers, *Journal of Labor Economics* **29**(1): 37–68.

Cunha, F. and Heckman, J. (2007). The technology of skill formation, *The American Economic Review* **97**(2): 31–47.

Currie, J. and Almond, D. (2011). Human capital development before age five, *Handbook of Labor Economics* **4**: 1315–1486.

Dahl, G. B. and Lochner, L. (2012). The impact of family income on child achievement: Evidence from the earned income tax credit, *The American Economic Review* **102**(5): 1927–1956.

Damm, A. P. and Dustmann, C. (2014). Does growing up in a high crime neighborhood affect youth criminal behavior?, *American Economic Review* **104**(6): 1806–1832.

Daniels, D. (2009). Social security payments for people caring for children, 1912 to 2008: a chronology, *Parliament of Australia, Department of Parliamentary Services, Parliamentary Library, Background Note* .

Davis, J. and Mazumder, B. (2018). Racial and ethnic differences in the geography of intergenerational mobility. SSRN Working Paper.

Department of Education and Training (2016). *Higher Education Statistics.*
URL: *https://www.education.gov.au/higher-education-statistics*

Deutscher, N. (2018). Place, jobs, peers and the teenage years: exposure effects and intergenerational mobility. Tax and Transfer Policy Institute Working Paper No. 10/2018.

Deutscher, N. and Breunig, R. (2016). Baby bonuses: Natural experiments in cash transfers, birth timing and child outcomes. Life Course Centre Working Paper Series Number 2016-25.

Drago, R., Sawyer, K., Shreffler, K. M., Warren, D. and Wooden, M. (2011). Did Australia's Baby Bonus increase fertility intentions and births?, *Population Research and Policy Review* **30**(3): 381–397.

Dustmann, C. and Görlach, J.-S. (2015). Selective out-migration and the estimation of immigrants' earnings profiles, *Handbook of the Economics of International Migration*, Vol. 1, Elsevier, pp. 489–533.

Dustmann, C. and Schönberg, U. (2012). Expansions in maternity leave coverage and children's long-term outcomes, *American Economic Journal: Applied Economics* **4**(3): 190–224.

Family Assistance Guide (2016).
URL: *http://guides.dss.gov.au/family-assistance-guide/3/6/4*

Fernández, R. (2011). Does culture matter?, *Handbook of Social Economics*, Vol. 1, Elsevier, pp. 481–510.

Fernandez, R. and Fogli, A. (2009). Culture: An empirical investigation of beliefs, work, and fertility, *American Economic Journal: Macroeconomics* **1**(1): 146–77.

Figlio, D., Giuliano, P., Özek, U. and Sapienza, P. (2016). Long-term orientation and educational performance. NBER Working Paper No. 22541.

Gaitz, J. and Schurer, S. (2017). Bonus skills: Examining the effect of an Australian unconditional cash transfer on child development. Life Course Centre Working Paper Series Number 2017-04.

Gang, I. N. and Zimmermann, K. F. (2000). Is child like parent? Educational attainment and ethnic origin, *Journal of Human Resources* pp. 550–569.

Gans, J. S. and Leigh, A. (2009). Born on the first of July: An (un) natural experiment in birth timing, *Journal of Public Economics* **93**(1): 246–263.

Gans, J. S. and Leigh, A. (2012). Bargaining over labour: Do patients have any power?, *Economic Record* **88**(281): 182–194.

Gans, J. S., Leigh, A. and Varganova, E. (2007). Minding the shop: The case of obstetrics conferences, *Social Science & Medicine* **65**(7): 1458–1465.

Gaviria, A. and Raphael, S. (2001). School-based peer effects and juvenile behavior, *Review of Economics and Statistics* **83**(2): 257–268.

Greene, W. H. (2003). *Econometric analysis*, 5th edition edn, Prentice Hall, Upper Saddle River, N.J.

Güell, M., Pellizzari, M., Pica, G. and Rodríguez Mora, J. V. (2018). Correlating social mobility and economic outcomes, *The Economic Journal* **128**(612): F353–F403.

Haider, S. and Solon, G. (2006). Life-cycle variation in the association between current and lifetime earnings, *American Economic Review* **96**(4): 1308–1320.

Hammarstedt, M. and Palme, M. (2012). Human capital transmission and the earnings of second-generation immigrants in Sweden, *IZA Journal of Migration* **1**(1): 4.

Hanushek, E. A. and Woessmann, L. (2011). The economics of international differences in educational achievement, *Handbook of the Economics of Education*, Vol. 3, Elsevier, pp. 89–200.

Hardin, J. W. (2002). The robust variance estimator for two-stage models, *Stata Journal* **2**(3): 253–266.

Hertz, T. (2008). A group-specific measure of intergenerational persistence, *Economics Letters* **100**(3): 415–417.

Hole, A. R. (2006). Calculating Murphy-Topel variance estimates in Stata: A simplified procedure, *Stata Journal* **6**(4): 521–529.

Hutchinson, D. (2016). *Weekly wages, average compensation and minimum wage for Australia from 1861-Present.*
URL: *https://www.measuringworth.com/auswages/*

Katz, L. F., Kling, J. R. and Liebman, J. B. (2001). Moving to opportunity in Boston: Early results of a randomized mobility experiment, *The Quarterly Journal of Economics* **116**(2): 607–654.

Kewley, T. H. (1973). *Social security in Australia, 1900-72*, Sydney University Press.

Khoo, S.-E., McDonald, P., Giorgas, D. and Birrell, B. (2002). Second generation Australians: report for the Department of Immigration and Multicultural and Indigenous Affairs, *Technical report*, Department of Immigration and Multicultural and Indigenous Affairs, Canberra.

Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics, *Journal of Economic Literature* **48**: 281–355.

Leigh, A. (2007). Intergenerational mobility in Australia, *The BE Journal of Economic Analysis & Policy* **7**(2).

Lubotsky, D. (2007). Chutes or ladders? A longitudinal analysis of immigrant earnings, *Journal of Political Economy* **115**(5): 820–867.

Ludwig, J., Liebman, J. B., Kling, J. R., Duncan, G. J., Katz, L. F., Kessler, R. C. and Sanbonmatsu, L. (2008). What can we learn about neighborhood effects from the Moving to Opportunity experiment?, *American Journal of Sociology* **114**(1): 144–188.

Manski, C. F. (1993). Adolescent econometricians: How do youth infer the returns to schooling?, *Studies of supply and demand in higher education*, University of Chicago Press, pp. 43–60.

Manski, C. F. (2000). Economic analysis of social interactions, *Journal of Economic Perspectives* **14**(3): 115–136.

Marks, G. (2010). Improvements over the educational career of immigrant students, *Australian Journal of Education* **54**(2): 133–154.

Mazumder, B. (2005). Fortunate sons: New estimates of intergenerational mobility in the United States using social security earnings data, *Review of Economics and Statistics* **87**(2): 235–255.

Mazumder, B. (2008). Sibling similarities and economic inequality in the US, *Journal of Population Economics* **21**(3): 685–701.

Mazumder, B. (2016). Estimating the intergenerational elasticity and rank association in the United States: Overcoming the current limitations of tax data, *Research in Labor Economics* **43**: 83–129.

Mendez, I. (2015). The effect of the intergenerational transmission of noncognitive skills on student performance, *Economics of Education Review* **46**: 78–97.

Mendez, I. and Zamarro, G. (2018). The intergenerational transmission of noncognitive skills and their effect on education and employment outcomes, *Journal of Population Economics* **31**(2): 521–560.

Mendolia, S. and Siminski, P. (2016). New estimates of intergenerational mobility in Australia, *Economic Record* **92**(298): 361–373.

Messinis, G. (2009). Earnings and languages in the family: Second-generation Australians, *Economic Record* **85**(s1).

Milligan, K. and Stabile, M. (2011). Do child tax benefits affect the well-being of children? Evidence from Canadian child benefit expansions, *American Economic Journal: Economic Policy* pp. 175–205.

Mitnik, P., Bryant, V., Weber, M. and Grusky, D. B. (2015). New estimates of intergenerational income mobility using administrative data, *Statistics of Income, Internal Revenue Service. mimeo (in preparation)* .

Morris, P., Duncan, G. J. and Clark-Kauffman, E. (2005). Child well-being in an era of welfare reform: the sensitivity of transitions in development to policy change., *Developmental Psychology* **41**(6): 919.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the US labor market, *The Quarterly Journal of Economics* **118**(2): 549–599.

Murphy, K. M. and Topel, R. H. (1985). Estimation and inference in two-step econometric models, *Journal of Business & Economic Statistics* **3**(4): 370–79.

Murray, C., Clark, R., Mendolia, S. and Siminski, P. (forthcoming). Direct measures of intergenerational income mobility for Australia, *Economic Record* .

Nybom, M. and Stuhler, J. (2017). Biases in standard measures of intergenerational income dependence, *Journal of Human Resources* **52**(3): 800–825.

OECD (2018). *A Broken Social Elevator? How to Promote Social Mobility.*
URL: *https://www.oecd-ilibrary.org/content/publication/9789264301085-en*

Oreopoulos, P. (2003). The long-run consequences of living in a poor neighborhood, *The Quarterly Journal of Economics* **118**(4): 1533–1575.

Overman, H. G. (2002). Neighbourhood effects in large and small neighbourhoods, *Urban Studies* **39**(1): 117–130.

Pagan, A. (1984). Econometric issues in the analysis of regressions with generated regressors, *International Economic Review* **25**(1): 221–247.

Page, M. E. and Solon, G. (2003a). Correlations between brothers and neighboring boys in their adult earnings: The importance of being urban, *Journal of Labor Economics* **21**(4): 831–855.

Page, M. E. and Solon, G. (2003b). Correlations between sisters and neighbouring girls in their subsequent income as adults, *Journal of Applied Econometrics* **18**(5): 545–562.

Raaum, O., Salvanes, K. G. and Sørensen, E. Ø. (2006). The neighbourhood is not what it used to be, *The Economic Journal* **116**(508): 200–222.

Roussel, S. (2004). The influence of expected earnings on the educational choices of Australian youth, *Unpublished Manuscript* .

Sacerdote, B. (2011). Peer effects in education: How might they work, how big are they and how much do we know thus far?, *Handbook of the Economics of Education*, Vol. 3, Elsevier, chapter 4, pp. 249–277.

Schulkind, L. and Shapiro, T. M. (2014). What a difference a day makes: quantifying the effects of birth timing manipulation on infant health, *Journal of Health Economics* **33**: 139–158.

Sharkey, P. and Faber, J. W. (2014). Where, when, why, and for whom do residential contexts matter? Moving away from the dichotomous understanding of neighborhood effects, *Annual Review of Sociology* **40**: 559–579.

Solon, G. (1992). Intergenerational income mobility in the United States, *The American Economic Review* pp. 393–408.

Solon, G. (1999). Intergenerational mobility in the labor market, *Handbook of labor economics*, Vol. 3, Elsevier, pp. 1761–1800.

Sweetman, A. and van Ours, J. C. (2015). Immigration: what about the children and grandchildren?, *Handbook of the Economics of International Migration*, Vol. 1, Elsevier, pp. 1141–1193.

Tamm, M. (2013). The impact of a large parental leave benefit reform on the timing of birth around the day of implementation, *Oxford Bulletin of Economics and Statistics* **75**(4): 585–601.

Thorngren-Jerneck, K. and Herbst, A. (2001). Low 5-minute Apgar score: A population-based register study of 1 million term births, *Obstetrics & Gynecology* **98**(1): 65–70.

Van den Berg, G. J., Lundborg, P., Nystedt, P. and Rooth, D.-O. (2014). Critical periods during childhood and adolescence, *Journal of the European Economic Association* **12**(6): 1521–1557.

Wilson, K., Wolfe, B. and Haveman, R. (2005). The role of expectations in adolescent schooling choices: Do youths respond to economic incentives?, *Economic Inquiry* **43**(3): 467–492.

Wodtke, G. T. (2013). Duration and timing of exposure to neighborhood poverty and the risk of adolescent parenthood, *Demography* **50**(5): 1765–1788.