The effects of neighbourhood and workplace income comparisons on subjective wellbeing

by

Shakked Noy

A thesis
submitted to the Victoria University of Wellington
in fulfilment of the requirements for the degree of
Master of Commerce

Victoria University of Wellington
(2020)

The effects of neighbourhood and workplace income comparisons on subjective wellbeing*

Shakked Noy[†]

February 18, 2020

Abstract

We investigate how the incomes of a person's neighbours and coworkers affect her happiness, using survey data on subjective wellbeing linked to unprecedentedly rich administrative data on the characteristics of survey respondents' peer groups. Linear regressions of subjective wellbeing on peer income variables establish that people care exclusively about their ordinal rank within their peer income distribution, that workplace rank matters much more than neighbourhood rank, and that workplace comparisons are driven primarily by fairness concerns. We confirm that our results reflect a causal effect of peer income by implementing sensitivity analyses, identifying off changes in peer income over time for immobile people, exploiting plausibly exogenous moves between workplaces triggered by mass layoffs, and testing for the effects of unobservable group-level confounders.

Keywords: subjective wellbeing, income comparisons, relative income

JEL Codes: D63, I31, J31

*Thanks to Statistics New Zealand for data access, and to the VUW School of Economics and Finance for funding access to the data. Thanks to Isabelle Sin for invaluable supervision. Thanks to Ilan Noy and seminar participants at Motu Economic & Public Policy Research for helpful comments. Thanks to my parents for everything.

[†]Victoria University of Wellington

0 Disclaimer

The results in this paper are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI), managed by Statistics New Zealand.

The opinions, findings, recommendations, and conclusions expressed in this paper are those of the author(s), not Statistics NZ or Victoria University of Wellington.

Access to the anonymised data used in this study was provided by Statistics NZ under the security and confidentiality provisions of the Statistics Act 1975. Only people authorised by the Statistics Act 1975 are allowed to see data about a particular person, household, business, or organisation, and the results in this paper have been confidentialised to protect these groups from identification and to keep their data safe.

Careful consideration has been given to the privacy, security, and confidentiality issues associated with using administrative and survey data in the IDI. Further detail can be found in the Privacy impact assessment for the Integrated Data Infrastructure available from www.stats.govt.nz.

The results are based in part on tax data supplied by Inland Revenue to Statistics NZ under the Tax Administration Act 1994. This tax data must be used only for statistical purposes, and no individual information may be published or disclosed in any other form, or provided to Inland Revenue for administrative or regulatory purposes.

Any person who has had access to the unit record data has certified that they have been shown, have read, and have understood section 81 of the Tax Administration Act 1994, which relates to secrecy. Any discussion of data limitations or weaknesses is in the context of using the IDI for statistical purposes, and is not related to the data's ability to support Inland Revenue's core operational requirements.

1 Introduction

Is a person’s happiness affected by the incomes of her peers? This question is both important and difficult to answer. It is important because an affirmative answer could provide us with a unified explanation of puzzling empirical phenomena like the Easterlin Paradox (Clark et al. (2008)), wage compression in the workplace (Dube et al. (2019)), or even the existence of unemployment (Akerlof and Yellen (1990)). And it is difficult to credibly answer because of the endogeneity of peer income in observational contexts.

In this paper, we use an unprecedentedly rich dataset to examine the relationship between a person’s subjective wellbeing and the incomes of that person’s neighbours and coworkers. Our dataset consists of survey data on the subjective wellbeing of 32,000 respondents, linked to comprehensive administrative data on the income, residential location, employer, and demographic characteristics of every person in New Zealand over a 15-year period. We use the administrative data to identify the neighbours and coworkers of each survey respondent, which allows us to calculate several measures of peer income for every respondent in our sample. We then run linear regressions of subjective wellbeing on neighbourhood and workplace peer income variables, as well as a variety of controls for individual and group-level covariates.

We establish three baseline results. First, a person’s subjective wellbeing is highly sensitive to her ordinal income rank among her peers: a 10 percentage point increase in a person’s ordinal income rank increases her happiness by 1/10th the happiness premium associated with eliminating moderate physical pain. Conditional on a person’s ordinal rank, her subjective wellbeing is uncorrelated with the other features of her peer income distribution, such as the median income among her peers. Second, workplace income rank matters much more than neighbourhood income rank in predicting a person’s subjective wellbeing. Third, within the workplace, workers are sensitive only to wage differentials that are unexplained by differences in observable characteristics.

These baseline results are purely correlational, and may or may not reflect a causal effect of peer income. We therefore use four supplementary strategies to determine whether our results have a causal interpretation. First, we conduct an informal coefficient stability analysis, and use the techniques developed by Cinelli and Hazlett (forthcoming) to calculate the minimum strength that unobserved confounders would need to have in order to explain away our results. We argue that unobserved confounders of the required strength are unlikely to exist. Second, we restrict our attention to immobile people and identify off changes in their peer income distribution over time. The results from this specification are similar to our main estimates, suggesting that endogenous selection into peer groups cannot be driving our

results. Third, we estimate specifications that exploit plausibly exogenous moves between workplaces triggered by mass layoffs or firm shutdowns. Once again, we obtain results similar to our main estimates, which provides further confirmation that endogenous selection is unlikely to be driving our results. Fourth, we use lagged peer income variables to run placebo tests that check for the influence of time-invariant group-level confounders (like the quality of local amenities in a neighbourhood). We find no evidence that this sort of omitted variable bias is affecting our results.

None of these strategies are by themselves ironclad, but the cumulative evidence they provide strongly suggests that the associations we observe are causal. Any remaining endogeneity driving our results must be due to omitted factors which are uncorrelated with the individual or group-level characteristics we control for, are equally strongly correlated with both the levels and changes of our peer income variables, apply equally to displaced and non-displaced workers, and covary with peer income over time but are not causally affected by it. We believe it is highly unlikely that there exist important omitted variables that satisfy these conditions, and therefore conclude that our estimates reflect a causal effect of peer income on wellbeing.

We make several contributions to the existing literature on income comparisons. Our first contribution consists of our ability to test whether the correlations we observe are causal. Most studies in the income comparisons literature use a simple “control on observables” strategy that is vulnerable to omitted variable bias from two sources: selection into peer groups based on unobservable individual characteristics, and the correlation between peer income variables and unobservable group-level characteristics that affect subjective wellbeing (like social capital or the presence of local amenities). We address these two sources of endogeneity using the supplementary strategies we described above, which allows us to produce estimates that are substantially more credible than the estimates in the existing literature.

We thereby provide causal evidence on income comparisons that is different in nature from the existing causal evidence. The existing evidence on the causal effects of income comparisons comes from several sources. The first is a set of laboratory experiments, which demonstrate that people are willing to pay to “burn” other people’s money in order to increase their relative standing (Zizzo and Oswald (2001)); that people are willing to take risky gambles in order to avoid placing last within an income distribution (Kuziemko et al. (2014)); and that the effort exerted by workers does *not* depend on the wages of their peers (Charness and Kuhn (2007), Goerg et al. (2010)) but their labour supply does (Bracha et al. (2015)). The second source of causal evidence is a pair of field experiments showing that information about peer wages reduces the job satisfaction of low-paid workers (Card et al.

(2012)) and that worker effort is sensitive to coworker wages (Breza et al. (2018)). Finally, two quasi-experimental studies show that retail workers who experienced lower quasi-random pay rises than their peers were subsequently more likely to quit (Dube et al. (2019)) and that making Norwegian tax records publicly available increased the gap in subjective wellbeing between wealthy and poor Norwegians (Perez-Truglia (forthcoming)).

The evidence we provide differs from these existing pieces of evidence in two key ways. First, we provide evidence of the causal effects of income comparisons *on subjective wellbeing*: with the exception of Perez-Truglia (forthcoming), none of the aforementioned studies has subjective wellbeing as its outcome variable. Second, we use observational data that covers a wider variety of more organic contexts than the comparatively narrow lab and field experiments we reviewed above.

Our second contribution to the literature consists of our ability to provide insight into the exact nature of the relationship between peer income and subjective wellbeing. A person’s happiness could conceivably depend on a variety of features of her peer income distribution: for example, people might care about the median income among their peers, the standard deviation of income among their peers, the top decile of income among their peers, or their ordinal rank among their peers. Distinguishing between these possibilities is crucial, because both the policy and academic implications of income comparisons are highly sensitive to the exact nature of those comparisons. For example, the existence of income or wage comparisons is sometimes invoked as a justification for redistributive taxation, or as an explanation of wage compression in the workplace.¹ But if people care primarily about their ordinal income rank, then redistribution or wage compression can do nothing to ameliorate the unhappiness created by income comparisons.

Despite the importance of identifying precisely which features of the peer income distribution people care about, most existing studies of income comparisons control for only one or two features of the peer income distribution, and are therefore unable to distinguish between the various competing possibilities.² Moreover, since the different features of the peer income distribution are strongly correlated with each other, existing studies are vulnerable to a form of omitted variable bias that makes it difficult to interpret their results. For example, a study that finds a negative effect on happiness of median peer income may simply be picking up the omitted effect of the standard deviation of peer income; and if that were the case, it would drastically change the interpretation of the study’s results.

Since our dataset allows us to observe the full distribution of peer income for each survey respondent, we are able to control for any potentially salient feature of the peer income

¹See Boskin and Sheshinski (1978) and Dube et al. (2019).

²No study of which we are aware has controlled for more than two features of the peer income distribution.

distribution. In particular, we choose to control for the median income among a person’s peers, that person’s ordinal rank among her peers, the standard deviation of income among her peers, and the top and bottom deciles of income among her peers. Our results suggest that people exclusively care about their ordinal income rank among their peers: people are happier the higher their ordinal rank is, and conditional on their ordinal rank they are unaffected by changes in the other features of the peer income distribution.

Several existing studies emphasise the importance of ordinal rank concerns, including Brown et al. (2008), Clark et al. (2009a), Powdthavee (2009), and Clark et al. (2010). We extend on these studies by demonstrating that ordinal rank is the *only* peer income variable that matters, by showing that this result remains even when we control for the standard deviation and top and bottom deciles of peer income, and by illustrating the robustness of this result.

Our third contribution arises from our ability to control simultaneously for both neighbourhood and workplace peer income variables, something which no previous study has done using objective measures of peer income.³ Controlling simultaneously for neighbourhood and workplace comparisons is important because failing to do so results in omitted variable bias, since neighbourhood and workplace peer income measures are strongly correlated with each other. Consider a study which finds a negative correlation between a person’s subjective wellbeing and the incomes of that person’s coworkers. This result could reflect the fact that people care about the incomes of their coworkers; alternatively, it could be explained by the fact that people care about the incomes of their *neighbours*, and the incomes of a person’s coworkers are correlated with the incomes of her neighbours. Distinguishing between these possibilities requires controlling for both neighbourhood and workplace peer income variables.

We find that for employed people, workplace comparisons predominate: when we include both neighbourhood and workplace peer income variables in our regressions, only the workplace variables are strongly predictive of subjective wellbeing. This result corroborates the available survey evidence: Clark and Senik (2010) and Goerke and Pannenberg (2015) both find that when survey respondents are directly asked about who they compare their income with, “coworkers” is the most common answer. We therefore show that people’s self-reports about income comparisons accurately describe their actual behaviour.

Finally, we contribute to the literature by showing that workplace comparisons . More specifically, we are interested in whether workers care about wage differentials *simpliciter*, or whether they care only about wage differentials that they perceive to be unfair. In order

³Goerke and Pannenberg (2015) control for both workplace and neighbourhood comparisons, but they rely on a self-reported 1-5 point scale as their measure of relative income.

to investigate this possibility, we run simple earnings regressions that allow us to decompose each worker’s wage into a component that is explained by their observable characteristics and an unexplained component. We use this decomposition to proxy for fairness considerations by assuming that individual i feels that they unfairly earn less than individual j if individual i ’s unexplained wage residual is smaller or more negative than individual j ’s unexplained wage residual.

Our results suggest that workplace comparisons are mostly motivated by fairness concerns. When we include both peer wage and peer residual variables in the same regression, we find that the only robust predictor of a worker’s subjective wellbeing is that worker’s rank among her peers in terms of unexplained wage residuals. The fact that workers are only sensitive to differences in rank created by *unexplained* wage differences suggests that concerns about fairness underlie workplace wage comparisons.

We therefore contribute to a small literature on the role that fairness concerns play in workplace comparisons. Breza et al. (2018) find, in a field experiment, that wage inequality reduces worker effort; but when workers can see that their higher-earning peers are simply being rewarded for their greater productivity, this effect disappears.⁴ We extend this result by providing suggestive evidence that, in real-world contexts where productivity is difficult to observe, workers use the observable characteristics of their coworkers as a proxy for those coworkers’ productivity when deciding whether wage differentials are fair.

In a slightly different vein, Clark and Oswald (1996) find that workers’ job satisfaction is negatively correlated with the gap between their actual wage and the wage that their observable characteristics predict they “ought” to earn. D’Ambrosio et al. (2018) find a similar result for the gap between a worker’s wage and the wage that they report would be “fair” for someone working in their occupation. However, neither of these studies controls for peer wages, so they cannot rule out the possibility that their findings simply reflect the omitted influence of peer wages. Since we are able to control for peer wages, we can more confidently infer that our results reflect fairness-based wage comparisons.

The rest of this paper proceeds as follows. In Section 2, we describe our dataset and display summary statistics for our estimation samples. In Section 3, we report our baseline results: estimates from regressions of subjective wellbeing on a set of peer income variables and controls. Subsequently, in Section 4 we conduct a series of tests to verify whether our Section 3 results reflect causal effects. In Section 5, we check that our results are robust to a variety of different sample restrictions and peer group definitions, including definitions

⁴Similarly, Bracha et al. (2015) find, in a lab experiment, that pay disparities affect workers’ labour supply, but this effect disappears when the pay disparities are attributable to differences in the quality of essays that the workers submitted.

that restrict a person’s peer group to consist only of people with the same sex, ethnicity, education level, or age as that person. Finally, in Section 6 we conclude.

2 Data

2.1 Sample construction

We use a combination of administrative and survey data drawn from Statistics New Zealand’s Integrated Data Infrastructure (IDI). The IDI contains a large collection of microdata sets, all of which are linked by unique person, business, and location identifiers. These datasets range from surveys (like the Census or Household Economic Survey) to comprehensive administrative records on tax, business, education, and criminal justice.

Our estimation sample comes from five biennial waves of the General Social Survey (GSS) between 2008 and 2016. The GSS asks respondents a series of questions about their personal circumstances, including a question about their “overall life satisfaction,” which we use as a measure of subjective wellbeing. Altogether, the five waves of the GSS contain 43,020 observations, all of whom are distinct individuals (the GSS is purely cross-sectional). We will refer to individuals from this sample as “GSS people.”

In order to calculate peer income measures for our GSS people, we use administrative data on tax, employment, and residential location that covers all individuals living in New Zealand and spans from 2003 to 2018. For each GSS person, we use address and employment data to determine where that person lived and worked during the month they were interviewed for the GSS. We then use the same data to identify all other individuals who lived in the same neighbourhood or worked in the same workplace during the relevant month. These individuals form the GSS person’s neighbourhood and workplace peer groups, respectively. Finally, we link all of our observations onto tax records to obtain detailed measures of their pre-tax incomes and wage earnings, and onto demographic characteristics data to obtain their age, sex, ethnicity, and education level. This allows us to observe the full distribution of income within each GSS person’s peer group, as well as the basic demographic characteristics of that peer group. More detail on each of these data sources is available in Section 2.2.

Overall, we can successfully link 76% of the GSS sample (32,643 people) on to income tax records and a neighbourhood peer group. Of these 32,643 people, our subsample of people employed at workplaces with at least 10 employees consists of 13,920 individuals (43%). This is a relatively low percentage of our overall sample, given that the national employment rate tends to hover around 60-65%. The disparity between the national employment rate and the employment rate in our sample is explained by three factors. First, the GSS sample

has an unusually high concentration of retirement-age people: 26% of the GSS sample is over 65, compared with 12% of the general population. Second, as we describe below, we count as “employed” only people who are employed and earning wages specifically during their GSS interview month. Third, we drop from our employed sample GSS people who are employed in small workplaces (workplaces with <10 peers), since peer income measures for those people would be very noisy.

This gives us two samples: a “full sample” of 32,643 people, which we use in our neighbourhood peer income regressions, and an “employed subsample” of 13,920 people, which we use in our workplace peer wage regressions.

2.2 Data sources

Our residential address, income, employment, and demographic characteristics data are drawn from a variety of administrative sources.

We derive our residential address data from an “address notifications” dataset compiled by Statistics New Zealand. Each time an individual provides a government ministry with her address (for example, by filing her taxes or filling out a school enrolment form), this information is recorded by Statistics New Zealand and used to construct a series of snapshots of that individual’s address at a variety of points in time. We fill in the gaps between these snapshots by assuming that, if person i lives in a particular location at time t , then person i continues to live in that location until the next time we observe a record of person i ’s address. This process of imputing an individual’s residential location at each point in time is obviously imperfect, especially for people who interact with the government infrequently, but any measurement error will at worst bias our results towards zero.

Our employment and wages data are drawn from the Linked Employer-Employee Database, which provides us with longitudinal information on the monthly pre-tax earnings and employer details of every worker in New Zealand. These data are constructed from a monthly tax schedule that employers file with the Inland Revenue Department, and a Business Register maintained by Statistics New Zealand. We use the monthly earnings data to calculate our workplace peer wage variables.

To calculate our neighbourhood peer income variables, we use annual tax returns filed with the Inland Revenue Department, which contain information on the total pre-tax annual income of the individual filing the return, as well as a breakdown of that total into different income sources. In particular, these tax returns allow us to determine whether an individual received income from welfare benefits in the past year.

Our age, sex, and ethnicity data comes from a Statistics New Zealand resource that

aggregates data from a variety of administrative sources to calculate the birthdate, sex, and ethnicity of every person in New Zealand.

Obtaining data on education is more challenging. For each observation in our sample, we impute that person’s highest qualification achieved at each point in time using a combination of Ministry of Education data (which records every qualification obtained in New Zealand between 1997 and 2018⁵), Ministry of Social Development data (which collects self-reported education data from welfare and superannuation beneficiaries), and self-reported qualifications data from the 2013 Census. Since some individuals have not obtained a qualification during the period covered by the Ministry of Education data, did not respond to the relevant Census question, and have not interacted with the Ministry of Social Development, we end up missing education data for 19% of our sample. We keep these observations and code their education as “missing.”

Finally, we use two supplementary data sources to construct controls for neighbourhood-level characteristics that are correlated with neighbourhood peer income and could affect the subjective wellbeing of neighbourhood residents. First, we use criminal offence records from the Ministry of Justice to construct measures of the average crime rate, and the average seriousness of the crimes committed, in each area unit. Unfortunately, these records only begin in July of 2014, so for each area unit, we simply take the average crime rate and average crime seriousness over the 2014-2018 period, and use these numbers as a time-invariant measure of the level of crime in that area unit.

Second, we use data on business locations to control for the presence in each neighbourhood of businesses which could plausibly be considered “social amenities.” We classify a variety of business types as being social amenities, including cafes, libraries, hospitals, sports and recreational facilities, and supermarkets.⁶ After calculating the number of amenities of each type in every neighbourhood, we conduct a factor analysis to construct a single variable that proxies for the overall presence of social amenities in a particular neighbourhood.

2.3 Subjective wellbeing data

Our measure of subjective wellbeing comes from a GSS question that reads as follows:

I am going to ask you a very general question about your life as a whole these days. This includes all areas of your life.

⁵Though records of school qualifications are incomplete for the first several years of this period.

⁶The full list of business types we classify as social amenities is: Supermarket and Grocery Stores, Fruit and Vegetable Retailing, Liquor Retailing, Recreational Goods Retailing, Cafes and Restaurants, Pubs, Taverns, and Bars, Libraries and Archives, Public Administration, Public Order and Safety Services, Preschool and School Education, Tertiary Education, Adult, Community, and Other Education, Medical and Other Health Care Services, Social Assistance Services, Arts and Recreation Services, and Religious Services.

Looking at showcard 17, where X is completely dissatisfied and Y is completely satisfied, how do you feel about your life as a whole?⁷

Unfortunately, the 2008-2012 waves of the GSS ask respondents to rate their life satisfaction on a 1-5 point scale, while the 2014 and 2016 waves provide respondents with a 0-10 point scale. In order to maximise the size of our sample, we standardise both of the life satisfaction variables to have a mean of 0 and a standard deviation of 1, and then pool together all five waves of the GSS. We recognize that this will distort the distribution of subjective wellbeing and may thereby affect our results, so in Section 5 we run robustness checks to confirm that this pooling of data does not qualitatively affect our estimates.

Figure 2 plots the distributions of subjective wellbeing in our sample. The 2008-2012 and the 2014-2016 distributions both display patterns that are common in the subjective wellbeing literature:⁸ the distributions are heavily skewed towards higher scores, and they peak at scores of 4 (out of 5) and 8 (out of 10), respectively. As a result, our pooled-sample, standardised subjective wellbeing variable peaks at a value of -0.09, which corresponds to a score of 4 for a 2008-2012 observation and a score of 8 for a 2014-2016 observation.

2.4 Peer group definitions

We define a person’s “neighbourhood” as the *area unit* in which they reside. Area units, which contain 1,860 income earners on average, are roughly the size of a small urban suburb; Figure 1 displays area unit boundaries overlaid on a map of central Wellington, to give a sense of their size. We choose to define neighbourhoods at the area unit level because area units are the natural level of interaction in an urban environment:⁹ people walk or drive around their area unit, visit shops and parks within their area unit, and so on. As a result, the incomes of other residents of the area unit will be highly salient.

A person’s neighbourhood peer group is defined to be everyone else in that person’s area unit for whom we have income data. This is a relatively broad peer group definition; we also replicate our main analyses using narrower definitions which restrict a person’s neighbourhood peer group to consist only of neighbours of the same sex, ethnicity, education level, or age bracket as that person. The results of these analyses are reported in Section 5; we show that narrower definitions produce qualitatively identical results.

We define a “workplace” as being a Permanent Enterprise Number–Permanent Business Number pair, which corresponds to a particular business operating at a particular estab-

⁷This is the wording from the 2014-2016 waves; the corresponding questions in the 2008-2012 waves have very minor differences in wording.

⁸See, for example, the distribution of subjective wellbeing in Clark et al. (2018).

⁹98% of our observations are from urban areas.

lishment location (for example, a specific McDonald’s outlet).¹⁰ A GSS person’s workplace is defined to be the place at which she was working during the month she was interviewed for the GSS. If a GSS person held multiple jobs during her interview month, we define her workplace as the job at which she collected the highest gross earnings in that month, and we use her wage from that job (rather than her total wage earnings) in our regressions.

As with our neighbourhood peer groups, a person’s workplace peer group is defined to be all other employees at that workplace (including those for whom the job is not their highest-paying job). If a GSS person’s workplace peer group contains fewer than 10 people, we drop that person from our workplace comparison dataset. In Section 5, we report the results of supplementary analyses that use narrower definitions of a person’s workplace peer group. Switching to narrower peer group definitions does not substantively affect our results.

2.5 Peer income variables

We control for several different features of a person’s peer income distribution, to allow us to distinguish between the many different ways people’s happiness could be affected by the incomes of their peers.

Most studies of income comparisons assume that people care about the median or mean income among their peers, as a natural point of comparison (see for example Luttmer (2005)). In addition, a number of studies argue that people care about their ordinal rank in the peer income distribution (Brown et al. (2008), Clark et al. (2009a), Powdthavee (2009), Clark et al. (2010)). People might also care about the standard deviation of income among their peers: people may dislike inequality, or enjoy socioeconomic diversity.¹¹ Finally, people might make unidirectional comparisons, in the sense that they primarily compare themselves to people at the top or bottom of their peer income distribution: people might be frustrated by the fact that they will never earn as much as their wealthiest peers, or relieved that they are not as badly off as the poorest of their peers.¹²

Correspondingly, we choose to control for the following peer income variables. For person i in neighbourhood g , we control for the median log income in neighbourhood g , person i ’s

¹⁰Note that for businesses with multiple establishment locations (like restaurant or retail chains), there is some measurement error involved in Statistics New Zealand’s assignment of workers to specific establishment locations. This means that there is a degree of measurement error in the creation of our workplace peer groups.

¹¹There is a large literature on the relationship between income inequality and subjective wellbeing, though the vast majority of the literature is at the cross-country, macro level, and has produced conflicting results; see Schneider (2016) for a review of the evidence. There has been comparatively little work on the effects of more localized, neighbourhood-level inequality, with the exception of Knight et al. (2009), who find a positive effect of local income inequality, and Morawetz et al. (1977), who find a negative effect.

¹²There is some evidence for this possibility: Godechot and Senik (2015) find that workers’ job satisfaction is sensitive to the wages of the top 1% of earners in their workplace.

ordinal income rank in neighbourhood g , the standard deviation of log income in neighbourhood g , the median log income within the top income decile in neighbourhood g , and the median log income within the bottom income decile in neighbourhood g . A person’s ordinal income rank is a number that ranges from 0 (the bottom rank) to 1 (the top rank). “Income” in this context means total pre-tax annual income.

Similarly, for person i in workplace w , we control for the median log wage in workplace w , person i ’s ordinal wage rank in workplace w , the standard deviation of log wages in workplace w , the median log wage within the top wage decile in workplace w , and the median log wage within the bottom decile in workplace w . “Wage” in this context means a worker’s gross earnings from workplace w in the relevant month. We restrict to wage income from workplace w because we expect that workers will compare salaries with each other, rather than comparing overall incomes.

To test whether workplace comparisons are motivated by fairness concerns, we also calculate our workplace comparison variables using wage residuals instead of wages. In order to calculate wage residuals, we take our full set of employment data, which consists of month-level data on all workers in New Zealand between 2003 to 2018. We partition this dataset by industry-year, using detailed industry codes from the Australia and New Zealand Standard Industrial Classification, which divides the economy into about 700 distinct industries.

Within each industry-year, we run a regression of the form

$$\ln w_{itg} = \alpha + \beta \text{indiv}_{it} + \gamma \text{workplace}_{tg} + \delta_{tg} + \varepsilon_{itg} \quad (1)$$

where i indexes individuals, t indexes months, and g indexes workplaces. $\ln w_{itg}$ is a worker’s log monthly wage earnings from workplace g , indiv_{it} is a vector of individual characteristics (age, sex, ethnicity, education level, and workplace-specific tenure), workplace_{tg} is a vector of workplace characteristics (number of employees, employee demographics, turnover), and δ_{tg} are month fixed effects and labour market area fixed effects.¹³

We take the residuals from these regressions and treat them as the “unexplained component” of a worker’s log monthly wage. We then calculate all of our workplace comparison variables using peer residuals as well as peer wages.

Because we use objective measures of peer income (rather than self-reports from respondents about their perceived relative income), we implicitly assume that the people in our sample know roughly how much their neighbours and coworkers earn. This, we believe, is a reasonable assumption. People often know the occupations of their neighbours, and the incomes of those neighbours are further revealed through their patterns of conspicuous con-

¹³Labour market area boundaries are defined using the results of Papps and Newell (2002).

sumption: the size of their house, the make of their car, the frequency with which they go on vacation, and so on. All of these sources of information combine to give most people a reasonable sense of their position within their neighbourhood peer income distribution. Meanwhile, in the workplace, norms supporting pay secrecy do present a barrier to discovering the salaries of one’s coworkers. However, promotions and titles usually serve as a good proxy for a particular coworker’s relative earnings. Once again, this gives most workers a rough sense of how much they are paid relative to all of their coworkers.

2.6 Descriptive statistics

Table 1 reports some basic descriptive statistics for our GSS sample. Note that all counts involved in the calculation of these statistics are randomly rounded to base 3, in order to comply with Statistics New Zealand’s confidentiality requirements.

Column 1 of Table 1 displays the individual and neighbourhood characteristics of our full sample. Our full sample is majority European, middle-aged on average, and earns slightly more than the New Zealand median income. Interestingly, members of our full sample are relatively uneducated: more than 50% have never earned a qualification beyond a high school diploma.

98% of our full-sample observations live in an urban area, and 15% moved into their current area unit some time in the past year. In the average area unit, 3 people report being victims of crime in a given month.

Column 2 displays summary statistics for our employed sample; members of this subsample are younger on average, better educated, and higher-earning than members of our full sample. About 4% of them are in their first month at their current job, and about 2% stay at their current job for 2 months or fewer. The average workplace in our overall employment dataset contains 40.34 employees, but the average worker in our sample has 340.28 coworkers (due to the fact that workers are more likely to be sampled from larger workplaces).

3 Baseline Results

In this section, we report the results of simple linear regressions of subjective wellbeing on our peer income variables and controls. All our regressions in this section have the same basic form: for individual i in group g at time t , the regression equation is

$$SWB_{itg} = \alpha + \beta ref_{itg} + \gamma indiv_{it} + \delta group_{tg} + \theta_t + \varepsilon_{itg} \quad (2)$$

where ref_{itg} is a vector of peer income variables, $indiv_{it}$ is a vector of individual characteristics, $group_{tg}$ is a vector of peer-group-level characteristics, and θ_t are year fixed effects. β is the vector of treatment effect coefficients. Our subjective wellbeing variable, SWB_{itg} , is standardised to have a mean of 0 and a standard deviation of 1.

Table 2 lists all of the right-hand side covariates that we include in our models, including precise variable definitions. We control for a range of individual characteristics, including income, age, sex, ethnicity, education, employment and welfare status, household composition, self-reported physical health and loneliness, and past income. Our neighbourhood-level controls include the crime rate in the neighbourhood, the presence of social amenities in the neighbourhood, the demographic characteristics of the neighbourhood’s residents, and the levels of ethnic and income segregation in the neighbourhood. Finally, our workplace controls include the turnover and average tenure in the workplace, as well as the demographic characteristics of the workplace’s employees.

Each regression is run either on our full sample (consisting of all 32,643 GSS people) or on the subsample of 13,920 employed GSS people. We run linear regressions, rather than the ordered probit regressions that are common in the literature, for ease of interpretation and so that our baseline results are comparable with the results of our sensitivity analysis and causality checks in Section 4.

To confirm that our controls are well-behaved and replicate the results from the wellbeing literature, we run test regressions of subjective wellbeing on our full set of controls without any peer income variables. The results of these regressions are reported in Table 3. The coefficients on the controls are consistent with the existing literature on the correlates of subjective wellbeing: men are less happy than women, wellbeing exhibits a U-shaped relationship with age, welfare recipients are less happy, more educated people are happier, physically healthier people are happier, and lonelier people are less happy. Interestingly, Māori and Pasifika are happier than Europeans, conditional on the other covariates. Residents of neighbourhoods with a higher concentration of social amenities are happier, but there is no correlation between the crime rate in a neighbourhood and the wellbeing of residents of that neighbourhood.

3.1 Specification buildup

Thanks to the richness of our dataset, we are able to control for a variety of measures of peer income and for a large set of group-level characteristics. In order to illustrate how the inclusion of these covariates affects our results, we start with basic regressions that predict subjective wellbeing using median peer income and our full set of individual-level controls.

We then add in additional peer income variables and our group-level controls, and show that our estimated coefficients are highly sensitive to the introduction of these additional covariates. This demonstrates the importance of controlling for a variety of peer income variables, in order to avoid omitted variable bias.

3.1.1 Neighbourhood results

Table 4 reports the results from five regressions. Column 1 displays the results from a regression of subjective wellbeing on median peer income and a full set of individual-level controls, as well as year fixed effects. Column 2 adds a control for ordinal income rank; Column 3 adds a control for the standard deviation of peer income; Column 4 adds controls for the top and bottom deciles of peer income; and Column 5 adds neighbourhood-level controls.

In Column 1, the coefficient on median peer income is small and insignificant, suggesting that a person’s happiness is unaffected by the incomes of her peers. This is a counterintuitive result, and one that contradicts the majority of findings in the income comparisons literature.

Adding a term for ordinal income rank in Column 2 clarifies matters. The coefficient on median peer income in Column 1 is small and insignificant not because of the absence of a relationship between peer income and subjective wellbeing, but because the median peer income variable in Column 1 is capturing two strong effects that go in opposite directions. On the one hand, people are happier the higher their ordinal rank is within the neighbourhood income distribution (the coefficient on ordinal income rank is positive and highly significant). On the other hand, conditional on their ordinal rank, people are happier the higher the median income among their peers (the coefficient on median peer income is also positive and highly significant).¹⁴

This tells an interesting story. People do care about the incomes of their peers, and they are unhappy when they earn less than their neighbours. However, people’s relative income concerns are purely ordinal: a person becomes happier if she moves up the ordinal income ladder, but not if she simply increases the cardinal amount by which her income exceeds the incomes of her peers. In fact, conditional on a person’s ordinal income rank, she is happier the more her peers earn, plausibly because higher-earning neighbours are associated with better neighbourhood amenities and higher levels of social capital.

The Column 2 results also explain why the coefficient on median peer income in Column 1 is small and insignificant. An increase in median peer income, conditional on a person’s own

¹⁴This is very similar to the results in Clark et al. (2009a), who examine the happiness of residents of Danish neighbourhoods using a dataset similar to our own. Clark et al. (2009a) only control for median peer income and ordinal income rank, so Columns 3-5 extend on their results.

income, corresponds to both a decrease in that person’s ordinal income rank (which reduces her subjective wellbeing) and an increase in the median income among her peers (which increases her subjective wellbeing). These effects mostly cancel out, which is why Column 1 gives the impression that there is no relationship between peer income and subjective wellbeing. The results in Columns 1 and 2 therefore emphasise how controlling for only one parameter of the peer income distribution can obscure crucial nuances in the relationship between peer income and subjective wellbeing.

Column 3 adds a control for the standard deviation of income in the neighbourhood. This addition demonstrates that the positive coefficient on median peer income in Column 2 was largely reflecting the omitted influence of the standard deviation of peer income. In Column 3, the coefficient on median peer income has halved in size and is no longer statistically significant, and the subjective wellbeing of people in our sample is instead strongly positively correlated with the standard deviation of income in their neighbourhood. This positive correlation may be due to the beneficial effects of neighbourhood socioeconomic diversity, or it may reflect the fact that neighbourhoods with a high dispersion of income contain residents who possess more opportunities for upward mobility.¹⁵

Column 4, which adds controls for the top and bottom deciles of the neighbourhood income distribution,¹⁶ shows that the positive coefficient on the standard deviation of peer income in Column 3 was in turn reflecting the omitted influence of the top decile of peer income. According to Columns 4 and 5, doubling the top decile of income in a neighbourhood (while holding the other peer income variables constant) increases the wellbeing of the neighbourhood’s residents by 0.07-0.10 of a standard deviation on average. The positive effect of the top decile of peer income likely reflects one of two things. It could be the result of an “aspiration” or “signalling” effect, whereby the incomes of the wealthiest residents serve as a positive signal of the future incomes of other residents. Alternatively, it could reflect the fact that the wealthiest residents of a neighbourhood are disproportionately influential in determining that neighbourhood’s endowment of public goods and other amenities.

Throughout this process, the coefficient on ordinal income rank remains positive and highly significant. Column 5 shows that moving up 10 percentage points in the neighbourhood income distribution (holding the other peer income variables constant) is associated with a 0.02 standard deviation increase in subjective wellbeing. This is a substantial effect:

¹⁵Knight et al. (2009) find a similar positive correlation in rural Chinese counties, and suggest that diversity or greater opportunities are driving the correlation. The results of Chetty et al. (2014) confirm that young residents of more socioeconomically diverse neighbourhoods have more opportunities for upwards mobility.

¹⁶Recall that the “top decile” variable is the median income in the top income decile in the neighbourhood, and the “bottom decile” variable is analogously defined.

it is 10% of the difference in happiness between people who report experiencing physical pain that “moderately” interferes with their daily activities and people who report that physical pain “does not at all” interfere with their daily activities.¹⁷

The coefficient on own income is small and statistically insignificant in Columns 3-5, suggesting that increases in a person’s income affect that person’s subjective wellbeing only through improvements in her ordinal income rank. This is an unusually extreme result; most studies in the literature find that controlling for peer income reduces, but does not eliminate, the positive effect of own income (see Clark et al. (2018)).

Two broad patterns in Table 4 are worth emphasizing. First, the coefficients on the peer income variables are highly sensitive to exactly which of the peer income variables we include in the regression. For example, the coefficient on median peer income more than doubles when we add a control for a person’s ordinal rank in her neighbourhood, then halves and becomes insignificant when we add a control for the standard deviation of neighbourhood income, then switches sign when we add controls for the top and bottom decile of neighbourhood income. This coefficient behaviour confirms our theoretical predictions: studies which just control for median peer income, and omit other parameters of the peer income distribution, will be strongly influenced by omitted variable bias.

Second, the coefficient on ordinal income rank is unusually stable. While the other peer income coefficients fluctuate in magnitude as we add in additional peer income variables or the group-level controls, the ordinal rank coefficient is barely affected by the introduction of additional covariates: it varies by less than 5% of its initial magnitude across all four specifications. This is encouraging, since it suggests that the ordinal rank coefficient is less vulnerable to omitted variable bias. We discuss this more in depth in Section 4.

3.1.2 Workplace results

Table 5 presents the same set of results as Table 4, except we restrict to our employed subsample and use workplace peer wage variables instead of neighbourhood peer income variables.

The patterns we observed in Table 4 are replicated in Table 5. Column 2 shows that a worker’s subjective wellbeing is strongly positively correlated with her ordinal wage rank; conditional on her ordinal rank, she is happier the higher the median wage of her coworkers. The positive effect of median peer wages disappears once we control for additional parameters of the peer wage distribution. Column 3 demonstrates that a worker’s subjective wellbeing is positively correlated with the standard deviation of wages in her workplace; Column 4 shows

¹⁷This presumably includes both people who do not experience any physical pain, and people who experience pain that does not interfere with their daily activities.

that a worker’s subjective wellbeing is also positively correlated with the bottom decile of wages in her workplace. These positive coefficients may reflect the fact that workplaces with a high standard deviation of wages have more opportunities for internal career advancement,¹⁸ and workplaces with a high bottom decile of wages adopt a management style that treats their lowest-paid workers relatively well.

Once again, the coefficient on ordinal wage rank is very stable (it varies by only 18% of its initial magnitude), while the coefficients on the other peer income variables change substantially in magnitude from column to column. As in Table 4, this suggests that the coefficient on ordinal wage rank is comparatively less vulnerable to omitted variable bias.

Next, we test whether workplace comparisons are primarily driven by fairness concerns. In Columns 1-5 of Table 6, we replicate Table 5 with peer residual variables instead of peer wage variables. As in Tables 4 and 5, the main story here is the coefficient on ordinal residual rank, which is large, significant, and remains exceptionally stable even as we add additional covariates (it varies by less than 1% of its initial magnitude).

However, these results could simply be driven by the omitted influence of the peer wage variables, which are correlated with the peer residual variables. In order to distinguish between wage-based and fairness-based workplace comparisons, we include both the peer wage and peer residual variables in the same regression. The results are reported in Column 6 of Table 6: they show that a worker’s subjective wellbeing is positively correlated with her ordinal rank in the workplace in terms of unexplained wage residuals, and positively correlated with the standard deviation of wages in her workplace. Conditional on a worker’s ordinal residual rank, her subjective wellbeing is uncorrelated with her ordinal wage rank.

These results suggest that workplace comparisons arise primarily from concerns about fairness. In particular, the positive coefficient on ordinal residual rank suggests that workers are upset when their coworkers earn larger wage premiums than they do. By contrast, the tiny and insignificant coefficient on ordinal wage rank suggests that workers do not care about wage differences that are explained by differences in observable characteristics.

The coefficient on the standard deviation of peer wages remains positive and significant, suggesting once again that people are happier when employed in workplaces with more opportunities for career advancement. Interestingly, the coefficient on a person’s own wage residual is negative and highly significant: conditional on their relative position within the workplace and on their total wage, people dislike having a larger proportion of their wage be unexplained by their observable characteristics. It is not clear how to interpret this result; one possibility is that people dislike having a large proportion of their wage be unexplained

¹⁸Or that workplaces with a high standard deviation of wages contain workers from a wider diversity of occupations, which may be good for workplace morale.

because they are concerned about arousing jealousy or frustration among their coworkers.

3.2 Preferred results

Having shown how the addition of a more comprehensive set of peer income variables shapes our results, we now present the estimates from our preferred (fully-controlled) specifications. These estimates are displayed in Table 7. Column 1, which displays the effects of neighbourhood peer income for our full sample, replicates Column 5 of Table 4. Column 2, which replicates Column 6 of Table 6, displays the effects of workplace peer wages and residuals for our employed sample. Column 3 contains both neighbourhood and workplace peer income variables, for our employed sample.

Table 7 shows three things. First, a person’s ordinal rank within her peer group is highly predictive of her subjective wellbeing. This is a remarkably robust result that remains stable across multiple samples and specifications. Second, workplace comparisons seem to be primarily driven by fairness concerns: as we noted above, workers care about their ordinal rank in terms of unexplained wage residuals, not total wages. Third, for employed people, only workplace comparisons appear to matter. In Column 3, the only peer income variables that significantly predict subjective wellbeing are workplace variables (ordinal residual rank, top decile of peer residuals, and standard deviation of peer wages). The estimated coefficients on the neighbourhood peer income variables are quite noisy, so we cannot rule out an effect of neighbourhood peer income, but we find no evidence that neighbourhood peer income matters once we condition on workplace peer income.¹⁹

The overall story we draw from these results is as follows. People care about their income rank within their peer group, plausibly because it plays a role in determining their overall status. Rank in the workplace matters more than rank in the neighbourhood, perhaps because relative income is more salient in the workplace or because workplace peer groups have tighter social bonds than neighbourhood peer groups. And rank in the workplace is determined by the unexplained component of a worker’s wage (so workers adjust for observable qualifications when they make workplace comparisons).

What explains the fact that people are primarily sensitive to their ordinal income rank, and are insensitive to cardinal differences conditional on their ordinal rank? One speculative explanation is as follows: people are able to estimate their ordinal position in their peer

¹⁹An alternative explanation of this result is as follows: the workplace peer wage variables matter more than the neighbourhood peer income variables not because people care more about workplace comparisons than about neighbourhood comparisons, but because people care more about peer *wages* than about peer *incomes*. In order to test for this explanation, we rerun the Column 3 regression using neighbourhood peer wage variables rather than neighbourhood peer income. This produces qualitatively identical results, which rules out this alternative explanation.

income distribution with a reasonable degree of accuracy, but lack information about the exact cardinal differences between themselves and their peers. It is, after all, much easier to correctly answer the question “does your neighbour/coworker earn more than you?” than it is to give an accurate answer to “by how much does your neighbour’s/coworker’s income exceed your own?” If people are roughly aware of their ordinal rank but are ignorant of the cardinal differences between them and their peers, it makes sense that their subjective wellbeing will be primarily or exclusively dependent on their ordinal rank.

The stories we just offered assume that the correlations we have documented in this section reflect a causal relationship between peer income and subjective wellbeing. We now test whether this is likely to be the case.

4 Tests for causation

The results in Section 3 are correlational only, and do not necessarily reflect the true causal effect of income comparisons. This is because the simple regression framework we (and most other studies in the literature) adopt is vulnerable to omitted variable bias from two sources.

The first source of omitted variable bias is individual selection into peer groups. The incomes of a person’s peers are plausibly correlated with that person’s unobservable characteristics: for example, more ambitious people may choose to live or work among higher-achieving peers, whereas more complacent people might choose to reside with peers who are similar to them. If these unobservable characteristics also affect a person’s subjective wellbeing, this creates an endogeneity problem. The direction of the bias created by individual selection is theoretically ambiguous, but in this section we will primarily test for individual selection that biases our estimates away from zero.

The second source of omitted variable bias is the correlation between peer income and other group-level characteristics that affect subjective wellbeing. For example, neighbourhoods with higher median incomes may also tend to have superior natural amenities. Similarly, workplaces with higher standard deviations of wages may have more segregated workplace cultures. Since these group-level characteristics are likely to affect the subjective wellbeing of group members, this creates an additional endogeneity problem. Once again, the direction of the resulting bias is theoretically ambiguous, but we will primarily test for omitted variable bias that pushes our estimates away from zero.

Distinguishing between these two sources of omitted variable bias is important because different strategies are required to tackle each source. For example, several studies in the literature (Luttmer (2005), Clark et al. (2009b)) use panel data with individual fixed effects, which helps them mitigate individual selection problems but does nothing to fight the omitted

influence of group-level characteristics.²⁰ In this section, we pursue one general strategy (sensitivity analysis), and three more specific strategies, two of which deal with individual selection and one of which deals with the omitted influence of fixed amenities.

4.1 Sensitivity analysis

We begin by analysing the sensitivity of our estimated coefficients to the potential presence of unobserved confounders. It is (by definition) impossible to directly test for the influence of unobserved confounders, but our observed results can be indirectly informative about the potential extent of omitted variable bias. Altonji et al. (2005) and Oster (2019) show that, under certain assumptions, the responsiveness of our peer income coefficients to the introduction of observable controls can provide us with useful information about the vulnerability of those coefficients to unobserved confounders. And Cinelli and Hazlett (forthcoming) develop several formal measures of the robustness of a set of estimates to the presence of unobserved confounders. We therefore begin by providing visual evidence on the stability of our peer income coefficients, and then calculate Cinelli and Hazlett’s (forthcoming) measures and discuss the results.

4.1.1 Coefficient stability analysis

The intuition behind this strategy is simple. Most unobserved confounders that could affect our results are likely to be closely related to our observable controls. For example, unobservable psychological characteristics that affect wellbeing (like a person’s optimism or vulnerability to mental health problems) are strongly related to socioeconomic status, which we can (albeit imperfectly) control for. Similarly, unobservable neighbourhood-level characteristics that affect wellbeing are likely to be correlated with the demographic characteristics of the neighbourhood’s residents, or with the other neighbourhood-level variables that we include in our regressions. The same is true for unobservable workplace-level variables.

As a result, the degree to which our coefficients shift as we add our observable controls is informative about the degree to which our coefficients could be influenced by omitted variable bias. If our coefficients are highly sensitive to the introduction of additional controls, this indicates that the effects of peer income are difficult to disentangle from the effects of a variety of other correlated variables, meaning that endogeneity is likely to be a problem even once we’ve included all of our observable controls. By contrast, if our coefficients

²⁰The natural parallel to this is to use group fixed effects, which would help mitigate the second source of endogeneity. No one (to our knowledge) has done this because, given the requirement of working with survey data, it is difficult to find datasets with many individuals in the same peer group *and* several different peer groups.

remain relatively stable, this suggests that confounding variables are not driving the observed correlation between our peer income variables and subjective wellbeing, and therefore that endogeneity is unlikely to be a significant concern.

Figures 3 and 4 plot the evolution of our peer income coefficients as we add controls into our main regressions. Figure 3 displays our neighbourhood peer income coefficients, and Figure 4 displays our workplace peer residual coefficients (which are estimated using regressions that also include the peer wage variables). The size of each dot in Figures 3 and 4 is proportional to the partial R^2 value of the corresponding control with the outcome. This is important information to display since, as Oster (2019) points out, the fact that a coefficient remains stable when we add a control variable means nothing if the control variable explains very little of the variation in the outcome variable.

The neighbourhood ordinal rank coefficient in Figure 3 is quite stable relative to the other peer income coefficients. While the standard deviation and bottom decile coefficients decline monotonically towards zero as we add controls, and the top decile coefficient fluctuates just above zero, the coefficient on ordinal income rank stabilizes after the addition of a few basic controls and then remains remarkably flat.

The two controls that have the biggest marginal effect on the ordinal income rank coefficient are the age controls (which cause the coefficient to increase by 50%) and a dummy variable for “receiving welfare benefits” (which causes the coefficient to decrease by 65%). Age is a major determinant of both a person’s ordinal income rank and her subjective wellbeing, so it makes sense that adding it as a control will substantially affect our estimate of the effect of ordinal rank. Meanwhile, the fact that introducing the “receives welfare” dummy has such a large negative effect on the ordinal rank coefficient is very interesting. It suggests that the ordinal rank variable is proxying for a wider sense of relative socioeconomic deprivation that is also partially captured by the “receives welfare” dummy.

Meanwhile, the coefficient on ordinal residual rank in Figure 4 is exceptionally stable and even increases gently in magnitude as we add controls. However, this stability is not unique to the ordinal rank variable: all of the other peer residual variables remain generally stable, though they are all statistically insignificant.

Overall, the results from Figures 3 and 4 are encouraging. The stability of the ordinal rank coefficients suggests that they are unlikely to be vulnerable to omitted variable bias. However, Figures 3 and 4 also give us reasons to be cautious. The neighbourhood rank coefficient is quite sensitive to the introduction of the first few individual-level controls, so it is possible that individual selection on unobservables is driving part of our results. We therefore make it a priority to test for the presence of unobservable selection, which we do in Sections 4.2 and 4.3.

4.1.2 Robustness analysis

We now conduct a more formal analysis of the vulnerability of our results to unobserved confounders, using the setup and procedures developed in Cinelli and Hazlett (forthcoming). Consider an augmented version of Equation (2), with the addition of a vector of unobserved confounding variables $unobs_{itg}$.

$$SWB_{itg} = \alpha_{aug} + \beta_{aug}ref_{itg} + \gamma_{aug}indiv_{it} + \delta_{aug}group_{tg} + \theta_{aug}FE_t + \lambda_{aug}unobs_{itg} + \xi_{itg} \quad (3)$$

The bias caused by the omission of the unobserved confounders is equal to the difference between the augmented coefficient β_{aug} and the β coefficient from Equation (2). The magnitude of this bias depends on two factors: the association of $unobs_{itg}$ with the peer income variables ref_{itg} , and the association of $unobs_{itg}$ with the outcome variable SWB_{itg} .

Cinelli and Hazlett (forthcoming) propose several measures of the vulnerability of our treatment effect estimates to the omission of $unobs_{itg}$. The primary measure is the Robustness Value, which measures the minimal degree of association that would need to hold between both $unobs_{itg}$ and ref_{itg} and between $unobs_{itg}$ and SWB_{itg} in order to explain away our results. The Robustness Value for a particular peer income coefficient T is equal to

$$RV = \frac{1}{2} \left(\sqrt{\left(\frac{R_{Y \sim T|C}^2}{1 - R_{Y \sim T|C}^2} \right)^2 + 4 \frac{R_{Y \sim T|C}^2}{1 - R_{Y \sim T|C}^2} - \frac{R_{Y \sim T|C}^2}{1 - R_{Y \sim T|C}^2}} \right) \quad (4)$$

where $R_{Y \sim T|C}^2$ is the partial R^2 of the peer income variable T with subjective wellbeing, conditional on the other peer income variables and our full set of controls.

The Robustness Value for the coefficient on ordinal income rank in the neighbourhood (from Column 1 of Table 7) is 2.9%, which means that unobserved confounders would have to explain at least 2.9% of the residual variation in subjective wellbeing and ordinal income rank in order to “explain away” our results. The Robustness Value for the coefficient on ordinal residual rank (from Column 2 of Table 7) is 3.2%. The Robustness Value for the coefficient on ordinal residual rank when both the neighbourhood and workplace peer income variables are included (Column 3 of Table 7) is 3.1%.

These numbers are very low, but remember that it is extremely difficult to explain much of the observed variation in subjective wellbeing. A person’s age, sex, education, ethnicity, employment status, and welfare benefit status combined only explain 2.7% of the residual variation in subjective wellbeing. Thus unobserved confounders of the required importance would need to be more important than all of those variables combined in determining the observed distribution of subjective wellbeing.

Is it likely that there exist unobserved confounders that are strong enough to explain away our results? We believe it is highly unlikely. These confounders would have to satisfy three conditions. First, they would need to be very important in explaining the residual variation in subjective wellbeing. Second, they would need to be mostly uncorrelated with the observable characteristics we have controlled for (otherwise, our observable controls would capture a large part of the effects of these unobserved confounders). Third, they would need to be strongly correlated with our ordinal income rank variables.

It is entirely plausible that there exist unobserved confounders that satisfy the first and second conditions. But it is very hard to think of an unobserved confounder that would satisfy all three, since it would need to be uncorrelated with most of our observables *but* strongly correlated with ordinal income rank.

We therefore conclude that omitted variable bias is unlikely to be driving the entirety of our ordinal rank results. Nevertheless, we still implement three additional strategies to confirm that this is the case.

4.2 Immobile specification

Our second strategy deals with concerns about individual selection into peer groups. In particular, we are concerned that a person’s ordinal rank in her peer income distribution is correlated with unobservable individual characteristics that affect her subjective wellbeing. In order to test for this possibility, we restrict our attention to immobile people (people who have remained in the same peer group for ≥ 5 years) and instrument for our peer income variables using the change in those peer income variables over the last 5 years. This (and the fact that we control for a 5-year lag of income) allows us to identify our peer income coefficients purely off changes in the peer income distribution over time. The resulting estimates will be less influenced by selection bias if the correlation between an immobile person’s unobservable characteristics and the changes in her peer income distribution over time is weaker than the correlation between her unobservable characteristics and the current levels of her peer income distribution.²¹

There are good *a priori* reasons to believe that this condition will be satisfied. It is plausible that a person will select into a peer group based in part on the distribution of income in that group. But once that person has settled in to her new peer group, future changes in the group’s distribution of income are beyond that person’s control. And the high costs of moving are likely to induce that person to stay, even if the changes in the

²¹Technically, we also need to assume that for each peer income variable $peer_i$, we have $\text{var}(peer_i) > \text{var}(\Delta peer_i)$. It’s easy to verify that this is the case in our dataset.

group’s income distribution are not to her liking. As a result, changes in the peer income distribution are more likely to be exogenous than the levels of the peer income distribution.

We can also conduct an indirect empirical test of this condition, by checking whether the correlation between an immobile person’s *observable* characteristics and the past changes in her peer income variables is weaker than the correlation between her observable characteristics and the current values of her peer income variables. In order to do this, we take our two peer income variables of interest (neighbourhood income rank and workplace residual rank). We regress both ordinal rank variables on our full set of individual characteristics, and then regress both of our “change in ordinal rank” variables on the same set of individual characteristics. The R^2 values from the first two regressions are 0.705 and 0.657, and the R^2 values from the second two are 0.499 and 0.524. The fact that the latter two values are substantially lower than the first two shows that much less of the variation in our “change in ordinal rank” variables is correlated with individual observable characteristics, which suggests that those variables are generally more exogenous than our normal ordinal rank variables.

On the basis of these theoretical and empirical considerations, we conclude that our “change in peer income” variables are less correlated with individual unobservable characteristics than our normal peer income variables, and therefore that our estimates in this section should be less influenced by selection bias than our estimates from Section 3. Of course, directly comparing our estimates from this section to our Section 3 estimates would be inappropriate; since immobile people are different, on average, from our full sample, we may obtain different estimates than the ones in Section 3 simply due to heterogeneity in the causal effect of peer income. To account for this fact, we take our sample of immobile people and use them to run both our Section 3 regressions *and* IV regressions where we instrument for our normal peer income variables using the 5-year change in those peer income variables. We then compare the coefficients from the two sets of regressions.

The results of these regressions are presented in Table 8. Column 1 replicates the specification from Column 1 of Table 7 for our immobile neighbourhood sample, and Column 3 replicates the specification from Column 2 of Table 7 for our immobile workplace sample. Columns 2 and 4 display the results from the corresponding IV regressions. The IV first stages are very strong: the average (minimum) first-stage F -statistics for the neighbourhood peer income variables are 4,087.2 (359.2), and for the workplace peer wage and residual variables they are 50.5 (10.7).

In both our neighbourhood and workplace immobile samples, the OLS and IV coefficients are not significantly different from each other, and the IV coefficient is larger. The IV coefficient on neighbourhood income rank is statistically insignificant, but this is due to the larger standard errors involved in IV estimation, not to any difference in the magnitude of

the coefficient.

If the assumption underlying this section is accurate, this shows that individual selection into peer groups cannot be driving our results.

4.3 Displacement specification

Our strategy in this section is an alternative attempt to detect whether individual selection is influencing our results. Rather than looking at immobile people, we look at movers, and restrict our attention to moves between workplaces that result from mass layoffs or firm shutdowns. More specifically, we restrict our attention to GSS people who left their previous job as the result of a mass layoff or firm shutdown, and run our standard regressions on this subsample. This strategy will produce results that are less influenced by selection bias if the correlation between a worker’s unobservable characteristics and the income distribution in her workplace is weaker for workers who lost their previous job as part of a mass layoff.

Once again, there are good *a priori* reasons to believe that this will be the case. Workers who leave their job as part of a mass layoff or firm shutdown are more likely to have lost their job involuntarily; moreover, workers who lose their job involuntarily often face an urgent need to find new employment, and are therefore less likely to be picky when attempting to find a new job. As a result, there will be less of a correlation between their unobservable characteristics and the characteristics of the new workplace they end up in.

The main difficulty confronting this strategy is the problem of correctly identifying instances where a mass layoff or firm shutdown occurs. For example, in firms with high turnover, a large voluntary exodus of employees may look like a mass layoff. And as Dixon and Stillman (2009) point out, Statistics New Zealand is prone to changing a firm’s identifying number when an ownership change, merger, or partial outsourcing of its functions occurs. This can create the appearance of a firm shutdown in the data, when in reality nothing much has happened.

In order to deal with these difficulties, we construct a series of increasingly strict definitions of “mass layoff.” The stricter the definition, the more likely it is that the associated moves are truly the result of an exogenous mass layoff.

We identify mass layoffs as follows. For each firm in our data, we calculate the percentage of employees who leave the firm in every month that the firm is present in our dataset. A particular month is classified as involving a mass layoff if:

First, the number of employees who leave in that month is at least 20 or at least $X\%$ of the firm’s workforce, where $X \in \{20, 40, 60, 80, 100\}$. The higher the percentage threshold, the stricter the definition.

Second, the percentage of employees who leave in that month is at least one standard deviation higher than the median percentage of employees who left that firm each month in the prior year. This requirement ensures that natural churn in workplaces with high turnover is not classified as a mass layoff.

Third, the percentage of employees who leave in that month is at least half a standard deviation higher than the percentage of employees who left the firm in the same month last year. This ensures that regularly scheduled events (like the departure of a group of summer interns or apprentices) are not classified as mass layoffs.

Fourth, of the employees who leave the firm, it's not the case that at least 50% of them move to the same new workplace. This requirement ensures that ownership changes, mergers, or partial outsourcings are not classified as mass layoffs.

These requirements give us 5 definitions of “mass layoff,” ordered by the strictness of the percentage threshold. We restrict to employed GSS people whose last job ended with a mass layoff, and who found their current job within six months of losing their previous one. This gives us 5 samples of employed GSS people, one for each mass layoff definition. The most inclusive sample (corresponding to the weakest definition) contains 1,926 workers, and the most restrictive sample (corresponding to the strongest definition) contains 1,578 workers. For each of these samples, we use one-to-one propensity score matching to construct a comparison sample of GSS workers who were *not* displaced at their previous job. The outcome variable in this matching procedure is a dummy for satisfying the relevant mass layoff definition, and we match on our full set of individual-level controls.

We then run our workplace peer income regressions on each of the layoff samples and each of the comparison samples. The intuition behind this strategy is borrowed from Section V.C. of Chetty and Hendren (2018). If our results are upwards-biased due to unobservable selection, and if mass layoffs do weaken the degree of unobservable selection, then we should expect to see systematic differences between the estimated coefficients in two ways. First, we should observe that the comparison coefficients are systematically larger (in absolute magnitude) than the layoff-based coefficients, since the comparison coefficients are more heavily influenced by selection bias. Second, we should observe the layoff-based coefficients diminishing in absolute magnitude as the definition of “mass layoff” becomes more strict (and therefore the probability that the layoff is truly exogenous increases).

The layoff-based and comparison coefficients are plotted in Figure 5 (we only plot the coefficients for variables that were statistically significant in our Section 3 regressions). There is no evidence that the comparison coefficients are systematically larger than the layoff coefficients. In fact, the opposite is true: each of the layoff coefficients is larger than the corresponding comparison coefficient, though the differences are not statistically significant.

If anything, this suggests that individual selection is biasing our estimates *towards* zero. Meanwhile, there is no evidence that the layoff coefficients decrease in magnitude due to the increasing strictness of the definition of “mass layoff:” in fact, there is an upwards trend in magnitude for the layoff-based residual rank coefficients, but there is a similar upwards trend for the comparison coefficients, so this appears to be due to some idiosyncratic change in the sample composition.

We therefore find no evidence that individual selection is biasing our workplace results away from zero.

4.4 Lagged peer income placebo test

We now turn our attention to the second source of endogeneity: omitted variable bias arising from group-level characteristics. In this section, we test for the influence of fixed group-level characteristics that are correlated with peer income, like the convenience of a neighbourhood’s location or the level of hazard at a workplace. Ideally, we would simply include group fixed effects in our regressions, but unfortunately our dataset contains very few instances where multiple GSS people are in the same neighbourhood or workplace.

Instead, we pursue a strategy analogous to the one in Section V.D. of Chetty and Hendren (2018). We begin by restricting our attention to GSS people who moved into their current neighbourhood or workplace sometime in the past 5 years.²² We then calculate 5-year lags of all of our peer income variables “as if” the people in our restricted sample had been in their current location 5 years ago. That is, for individual i in location l at time t , we calculate her peer income variables at time $t - 5$ “as if” individual i had been in location l at time $t - 5$.

Finally, we include these lagged peer income variables in our standard regressions, and check whether they attract significant coefficients. The intuition behind this strategy is simple. Any fixed local amenities that are correlated with peer income should be just as correlated with the lagged values of peer income as they are with the current values. So if our results are driven by the omitted influence of fixed amenities, then the coefficients on the lagged peer income variables should attract significant coefficients. Contrapositively, if the coefficients on the lagged variables are close to zero, this type of endogeneity cannot be driving our results.

The results of this test are displayed in Tables 9 and 10. In both the neighbourhood and the workplace regressions, the lagged variables are all very close to zero and statistically insignificant, and adding them into the regression barely affects the coefficients on the current

²²We make this sample restriction to rule out the possibility of lagged effects of the peer income variables, for example due to reference-dependent preferences.

peer income variables. This suggests that endogeneity due to fixed group-level confounders is unlikely to be driving our results.

This strategy cannot detect the influence of time-varying group-level confounders that are correlated with changes in peer income over time. However, this is not very concerning. Consider an unobservable group-level characteristic that covaries with peer income over time, like the level of social capital in a neighbourhood. This characteristic is likely to covary with peer income because *it is causally affected by it*. For example, the level of social capital in a neighbourhood is plausibly causally affected by the reactions of residents to the distribution of income in the neighbourhood. As a result, we would not want to control for social capital even if we were able to do so, since we would thereby ignore one of the mechanisms through which income comparisons affect subjective wellbeing. Social capital would be a “bad control” in the sense of Angrist and Pischke (2009). More broadly, any unobservable group-level characteristic that covaries with peer income over time is likely to be a bad control, so the fact that we are unable to detect or control for them is not a significant problem.

5 Robustness Tests

In this section, we conduct a series of tests to ensure that our results are robust to a variety of different specifications and sample definitions. The results of these tests are displayed in Figure 6, where we plot coefficient magnitudes for our two ordinal rank variables. The first column displays the coefficient magnitudes from our main specifications (Columns 1 and 2 of Table 7), and the rest of the columns display coefficients estimated from the robustness tests we describe below.

We begin by checking whether splitting up our estimation sample affects our results. As we described in Section 2.3, the 2008-2012 waves of the GSS use a 1-5 point subjective wellbeing scale, while the 2014-2016 waves use a 0-10 point scale. Consequently, pooling together all five waves of the GSS creates a distorted distribution of subjective wellbeing, and this might affect our results. We therefore reestimate our main regressions separately on our 2008-2012 sample and on our 2014-2016 sample. The results are displayed in the second and third columns of Figure 6: as is evident, the estimated coefficients from each subsample are almost exactly the same as the estimated coefficients from our main regressions.

Next, we check whether excluding retirement-aged individuals from our estimation sample changes our results. Retirement-aged people are much less likely to be employed, and live mostly off of pensions or superannuation income, so it seems possible that they have left the “rat race” of income comparisons. As a result, including them in our regressions may dilute

(or otherwise affect) our estimates of the effects of peer income. To test for this possibility, we run our main regressions on the subsample of GSS people who are under 65. The results are displayed in the fourth column of Figure 6: the resulting coefficient on neighbourhood income rank is about 50% larger than our main estimate, which suggests that the inclusion of retirement-aged people dilutes our main estimate. By contrast, the resulting coefficient for workplace residual rank is very similar to our main estimate, which makes sense given that by restricting to employed people we already exclude most retirement-aged people (though about 7% of our employed sample is of retirement age).

Finally, we test whether narrower peer group definitions affect our estimates. A wide body of evidence in psychology argues that people primarily compare themselves to their observably similar peers (see the “similarity hypothesis” in Festinger (1954)). As a result, we may want to restrict the definition of a peer group so that a person’s peer group only contains individuals that are observably similar to her.

To test whether doing so affects our results, we recalculate our peer income variables using four narrower definitions of a person’s peer group. The first definition restricts person i ’s peer group to consist of everyone in the same neighbourhood/workplace with the same sex as person i ; the second restricts to people in the same neighbourhood/workplace of the same ethnicity as person i ; the third restricts to people in the same neighbourhood/workplace with the same level of education as person i ; and the fourth restricts to people in the same neighbourhood/workplace within the same 10-year age bracket as person i .²³

The results from these narrower definitions are displayed in the final four columns of Figure 6. The coefficient on neighbourhood income rank is mostly unaffected by the use of narrower peer group definitions. By contrast, the estimated coefficient on workplace residual rank is smaller than our main estimate for three out of the four narrower definitions. It is unclear why this is the case.

Overall, Figure 6 shows that our main estimates are highly robust. There is relatively little variation in the magnitude of the coefficients, and no variation in the qualitative results: all of the coefficients from our robustness checks are positive and statistically significant.

6 Conclusion

We investigate the relationship between a person’s subjective wellbeing and the incomes of her neighbours and coworkers. Using a rich dataset that allows us to study neighbourhood and workplace comparisons simultaneously, we establish four broad results.

²³The age brackets are 0-19, 20-29, 30-39, 40-49, 50-59, 60-69, and 70+.

First, peer income has a causal effect on subjective wellbeing. Naive linear regressions demonstrate the existence of a strong correlation between peer income and subjective wellbeing, and a set of followup tests confirm that this correlation is unlikely to be driven by omitted variable bias.

Second, people care specifically about their ordinal rank within their peer income distributions. This has a number of important implications. It implies that compression of the peer income distribution (through redistributive taxation, or wage compression within the workplace) does nothing to ameliorate the effects of income comparisons. It provides evidence for the hypothesis that people care about income comparisons because income helps to determine their rank within a status hierarchy. And it is accompanied by a set of negative results: we find no evidence that people care about the variance of the income within their peer group, or about the cardinal distance between their own income and the incomes of their peers.

Third, workplace comparisons matter more than neighbourhood comparisons. On the one hand, this is an intuitive result for a variety of reasons: most people interact more with their colleagues than their neighbours, salary concerns are highly salient in the workplace, and so on. On the other hand, it is difficult to reconcile this result with the observation that most conspicuous consumption takes place in the neighbourhood (large houses, fancy cars, and so on).

Fourth, workplace comparisons are motivated by fairness concerns. Of all our results, this one is the most informative about the psychological mechanisms underlying income comparisons. The fact that people seem to adjust for observable characteristics when making wage comparisons in the workplace is very interesting. One possible interpretation of this result is that income comparisons help determine a person's self-esteem, and people naturally adjust for differences in background when comparing how well they have done with how well their peers have done.

Overall, we have provided some novel insights into the nature of the relationship between subjective wellbeing and peer income, and increased the credibility of existing estimates by showing that omitted variable bias cannot explain away the observed correlation between peer income and subjective wellbeing. We hope that future research will further probe the mechanisms underlying income comparisons, and produce even more rigorous tests of whether the association between peer income and subjective wellbeing is causal.

References

- Akerlof, George A. and Janet L. Yellen, “The Fair Wage-Effort Hypothesis and Unemployment,” *The Quarterly Journal of Economics*, 1990, 105 (2), 255–283.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools,” *Journal of Political Economy*, 2005, 113 (1), 151–184.
- Angrist, Joshua and Jörn-Steffen Pischke, *Mostly Harmless Econometrics*, Princeton University Press, 2009.
- Boskin, Michael J. and Eytan Sheshinski, “Optimal Redistributive Taxation When Individual Welfare Depends on Relative Income,” *The Quarterly Journal of Economics*, 1978, 92 (4), 589–601.
- Bracha, Anat, Uri Gneezy, and George Loewenstein, “Relative Pay and Labour Supply,” *Journal of Labor Economics*, 2015, 33 (2), 297–315.
- Breza, Emily, Supreet Kaur, and Yogita Shamdasani, “The Morale Effects of Pay Inequality,” *The Quarterly Journal of Economics*, 2018, 133 (2), 611–663.
- Brown, Gordon D. A., Jonathan Gardner, Andrew J. Oswald, and Jing Qian, “Does Wage Rank Affect Employees’ Well-being?,” *Industrial Relations*, 2008, 47 (3), 355–389.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, “Inequality at Work: the Effect of Peer Salaries on Job Satisfaction,” *American Economic Review*, 2012, 102 (6), 2981–3003.
- Charness, Gary and Peter Kuhn, “Does Pay Inequality Affect Worker Effort? Experimental Evidence,” *Journal of Labor Economics*, 2007, 25 (4), 693–723.
- Chetty, Raj and Nathaniel Hendren, “The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects,” *The Quarterly Journal of Economics*, 2018, 133 (3), 1107–1162.
- , —, Patrick Kline, and Emmanuel Saez, “Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1553–1623.
- Cinelli, Carlos and Chad Hazlett, “Making Sense of Sensitivity: Extending Omitted Variable Bias,” *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, forthcoming.
- Clark, Andrew E. and Andrew J. Oswald, “Satisfaction and Comparison Income,” *Journal of Public Economics*, 1996, 61 (3), 359–381.
- and Claudia Senik, “Who Compares to Whom? The Anatomy of Income Comparisons in Europe,” *The Economic Journal*, 2010, 120 (544), 573–594.
- , David Masclet, and Marie Claire Villeval, “Effort and Comparison Income: Experimental and Survey Evidence,” *Industrial and Labor Relations Review*, 2010, 63 (3), 407–426.
- , Nicolai Kristensen, and Niels Westergård-Nielsen, “Economic Satisfaction and Income Rank in Small Neighbourhoods,” *Journal of the European Economic Association*, 2009, 7 (2-3), 519–527.
- , —, and —, “Job Satisfaction and Co-worker Wages: Status or Signal?,” *The Economic Journal*, 2009, 119 (536), 430–447.

- , **Paul Frijters**, and **Michael A. Shields**, “Relative Income, Happiness, and Utility: an Explanation for the Easterlin Paradox and Other Puzzles,” *Journal of Economic Literature*, 2008, 46 (1), 95–144.
- , **Sarah Flèche**, **Richard Layard**, **Nattavudh Powdthavee**, and **George Ward**, *The Origins of Happiness: the Science of Well-Being over the Life Course*, Princeton University Press, 2018.
- D’Ambrosio**, **Conchita**, **Andrew E. Clark**, and **Marta Barazzetta**, “Unfairness at Work: Well-Being and Quits,” *IZA DP No. 11318*, 2018.
- Dixon**, **Sylvia** and **Steven Stillman**, “The Impact of Firm Closures on Workers’ Future Labour Market Outcomes,” *Statistics New Zealand Working Paper*, 2009.
- Dube**, **Arindrajit**, **Laura Giuliano**, and **Jonathan Leonard**, “Fairness and Frictions: the Impact of Unequal Raises on Quit Behavior,” *American Economic Review*, 2019, 109 (2), 620–663.
- Festinger**, **Leon**, “A Theory of Social Comparison Processes,” *Human Relations*, 1954, 7 (2), 117–140.
- Godechot**, **Olivier** and **Claudia Senik**, “Wage Comparisons in and out of the Firm. Evidence from a Matched Employer-Employee French Database,” *Journal of Economic Behavior & Organization*, 2015, 117, 395–410.
- Goerg**, **Sebastian J.**, **Sebastian Kube**, and **Ro’i Zultan**, “Treating Equals Unequally: Incentives in Teams, Workers’ Motivation, and Production Technology,” *Journal of Labor Economics*, 2010, 28 (4), 747–772.
- Goerke**, **Laszlo** and **Markus Pannenberg**, “Direct Evidence for Income Comparisons and Subjective Well-being across Reference Groups,” *Economics Letters*, 2015, 137, 95–101.
- Knight**, **John**, **Lina Song**, and **Ramani Gunatilaka**, “Subjective Well-being and its Determinants in Rural China,” *China Economic Review*, 2009, 20, 635–649.
- Kuziemko**, **Ilyana**, **Ryan W. Buell**, **Taly Reich**, and **Michael I. Norton**, “‘Last-Place Aversion’: Evidence and Redistributive Implications,” *Quarterly Journal of Economics*, 2014, 129, 105–149.
- Luttmer**, **Erzo F. P.**, “Neighbors as Negatives: Relative Earnings and Well-being,” *The Quarterly Journal of Economics*, 2005, 120 (3), 963–1002.
- Morawetz**, **David**, **Ety Atia**, **Gabi Bin-Nun**, **Lazaros Felous**, **Yuda Gariplerden**, **Ella Harris**, **Sami Soustiel**, **George Tombros**, and **Yossi Zarfaty**, “Income Distribution and Self-rated Happiness: Some Empirical Evidence,” *The Economic Journal*, 1977, 87 (347), 511–522.
- Oster**, **Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business and Economic Statistics*, 2019, 37 (2), 187–204.
- Papps**, **Kerry L.** and **James O. Newell**, “Identifying Functional Labour Market Areas in New Zealand: A Reconnaissance Study Using Travel-to-Work Data,” *IZA DP No. 443*, 2002.
- Perez-Truglia**, **Ricardo**, “The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment,” *American Economic Review*, forthcoming.
- Powdthavee**, **Nattavudh**, “How Important is Rank to Individual Perception of Economic Standing? A Within-Community Analysis,” *Journal of Economic Inequality*, 2009, 7, 225–248.

- Reardon, Sean F.**, “Measures of Income Segregation,” *CEPA Working Papers*, 2011.
- Schneider, Simone M.**, “Income Inequality and Subjective Wellbeing: Trends, Challenges, and Research Directions,” *Journal of Happiness Studies*, 2016, *17*, 1719–1739.
- Zizzo, Daniel John and Andrew J. Oswald**, “Are People Willing to Pay to Reduce Others’ Incomes?,” *Annales d’Économie et de Statistique*, 2001, *63/64*, 39–65.

Figures

Figure 1: Area unit boundaries in central Wellington

Notes: this is a screenshot from the Statistics New Zealand datafinder, accessible here <https://datafinder.stats.govt.nz/layer/27771-area-unit-2017-generalised-version/>.

Figure 2: Distributions of subjective wellbeing

Figure 3: Neighbourhood peer income coefficients

Notes: this figure displays coefficients from linear regressions of subjective wellbeing on our neighbourhood peer income variables and controls. The base model is a regression of subjective wellbeing on the peer income variables and year fixed effects, and each column adds an additional control into the regression. The size of each dot is proportional to the partial R^2 of the column's control with the outcome. Error bars represent 95% confidence intervals calculated using robust standard errors.

Figure 4: Workplace peer residual coefficients

Notes: this figure displays peer residual coefficients from linear regressions of subjective wellbeing on our workplace peer wage & residual variables and controls. The base model is a regression of subjective wellbeing on the peer wage & residual variables and year fixed effects, and each column adds an additional control into the regression. The size of each dot is proportional to the partial R^2 of the column's control with the outcome. Error bars represent 95% confidence intervals calculated using robust standard errors. "LMA"="Labour market area."

Figure 5: Displacement specification coefficients

Notes: this figure displays coefficients from linear regressions of subjective wellbeing on workplace peer wage and residual variables. Blue/circular dots represent estimates where the sample is restricted to people who lost their previous job as part of a mass layoff. The horizontal axis represents five increasingly strict definitions of “mass layoff,” as described in Section 4.3. For each laid-off sample, we construct a comparison population using one-to-one propensity score matching; the red/triangular dots represent estimates from regressions run on the relevant comparison population.

Figure 6: Robustness check coefficients

Notes: this figure displays the income rank coefficient from regressions of subjective wellbeing on neighbourhood peer income variables and a full set of controls, and the residual rank coefficient from regressions of subjective wellbeing on workplace peer wage and residual variables and a full set of controls. Each coefficient represents a different robustness check.

Tables

Table 1: Descriptive statistics

Variable	Full sample		Employed sample	
	Mean	S.D.	Mean	S.D.
<i>Demographics</i>				
Male dummy	0.46	–	0.47	–
Age	50.71	(18.51)	44.20	(13.85)
Total income (2019 NZD)	44,566	(44,463)	61,780	(48,290)
Employed dummy	0.43	–	–	–
Receives welfare benefits	0.14	–	0.07	–
<i>Highest educational qualification</i>				
Missing qualifications info	0.19	–	0.12	–
No qualifications	0.21	–	0.17	–
School qualification	0.32	–	0.34	–
Postschool qualification	0.15	–	0.17	–
Undergraduate qualification	0.09	–	0.13	–
Postgraduate qualification	0.05	–	0.07	–
<i>Ethnicity</i>				
Sole European	0.76	–	0.73	–
Sole Māori	0.06	–	0.06	–
Māori and European	0.05	–	0.06	–
Pacific or Pacific-European	0.03	–	0.04	–
Asian or Asian-European	0.06	–	0.08	–
MELAA or European-MELAA	0.01	–	0.01	–
Other	0.01	–	0.01	–
<i>Other individual characteristics</i>				
Lives in urban area	0.98	–	0.95	–
Is “never lonely”	0.64	–	0.64	–
Moved into nhd this year	0.15	–	0.17	–
Started at workplace this month	–	–	0.04	–
≤ 2 months spent at this workplace	–	–	0.02	–
Monthly wage (2019 NZD)	–	–	5,086	(5,272)
Residual of log wage	–	–	0.05	(0.70)
<i>Neighbourhood characteristics</i>				
Average monthly # of crimes	3.07	(4.87)	3.08	(4.87)
<i>Workplace characteristics</i>				
Number of coworkers	–	–	341.28	(838.47)
Monthly turnover (% of employees)	–	–	0.10	–
# of observations	32,643		13,920	

Table 2: Variable descriptions

Variable	Description
Sex	Male dummy.
Age	Age (in years), and age squared/100.
Employment status	Dummy for being employed in the GSS interview month.
Education	Highest qualification achieved: categories are (1) no qualifications, (2) high school qualification, (3) post-school qualification, (4) undergraduate qualification, (5) postgraduate qualification.
Ethnicity	Categories are: (1) Sole European, (2) Sole Maori, (3) Maori and European, (4) Pacific or Pacific-European, (5) Asian or Asian-European, (6) Middle Eastern/Latin American/African or European-MELAA, (7) Other.
Urban	Dummy for living in an urban area.
Household composition	Dummies for household composition (couple with children/single parent with children/one-person household etc).
Household income	Self-reported household income, from the GSS.
Past income	Annual income from 5 years ago.
Receives welfare	Dummy for receiving nonzero income from welfare benefits (excluding parental leave and superannuation).
Physical health	Factor variable constructed from two GSS questions about being limited by physical health problems and feeling physical pain.
Loneliness	GSS question about self-rated loneliness on a 1-5 point scale.
Housing quality	Factor variable constructed from three GSS questions about overall housing quality/presence of mold in house/problems with cold during winter.
Tenure in nhd	Number of months spent continuously living in current area unit.
Tenure in workplace	Number of continuous years in which the person spent at least one month working at this workplace.
Crime rate in nhd	Average number of crimes committed in this area unit over the 2014-2018 period, and average intensity of the crimes committed (intensity is measured using Ministry of Justice seriousness scores).
Social amenities in nhd	Factor variable constructed from the numbers of businesses of various types located in the area unit, including restaurants, supermarkets, hospitals, and sports facilities. A full list of the business types is available in Footnote 6.
Segregation in nhd	Measures of income and ethnic segregation between the meshblocks within each area unit, and measures of the ethnic diversity in each area unit, using the measures from Reardon (2011). Meshblocks are smaller geographic units containing
Nhd demographics	Number of people in area unit, average values of all education/ethnicity dummies in area unit, percent of area unit income that is wage income, percent of people in area unit receiving welfare benefits, employment rate in nhd.
Workplace demographics	Number of people in workplace, % of employees who left/entered workplace during this month, average values of all education/ethnicity dummies in workplace.

Table 3: Test regressions

Variable	Neighbourhood		Workplace	
	Coef	S.E.	Coef	S.E.
Income (log)	0.017***	(0.005)	0.032***	(0.012)
Household income (log)	0.039***	(0.006)	0.032***	(0.010)
Income 5 years ago (log)	0.008	(0.005)	0.008	(0.011)
Male dummy	-0.111***	(0.011)	-0.063***	(0.016)
Age	-0.025***	(0.002)	-0.035***	(0.004)
Age squared/100	0.029***	(0.002)	0.040***	(0.004)
Employed dummy	0.002	(0.012)	—	—
Receives welfare benefits	-0.288***	(0.021)	-0.184***	(0.038)
Missing qualifications info	-0.001	(0.018)	0.040	(0.031)
School qualification	0.010	(0.015)	0.009	(0.024)
Postschool qualification	-0.000	(0.018)	0.007	(0.028)
Undergraduate qualification	0.040*	(0.021)	0.013	(0.031)
Postgraduate qualification	0.074***	(0.025)	0.032	(0.036)
Sole Māori	0.074***	(0.025)	0.050	(0.036)
Māori and European	-0.014	(0.024)	0.004	(0.034)
Pacific or Pacific-European	0.116***	(0.033)	-0.013	(0.044)
Asian or Asian-European	0.004	(0.023)	0.009	(0.031)
MELAA or MELAA-European	-0.026	(0.050)	-0.050	(0.069)
Other ethnicity	0.116**	(0.045)	0.103	(0.064)
Lives in urban area	0.001	(0.033)	-0.050	(0.040)
Physical health factor variable (higher = better)	0.501***	(0.026)	0.348***	(0.044)
Lonely all of the time	-0.590***	(0.059)	-0.360***	(0.092)
Lonely most of the time	-0.889***	(0.041)	-0.934***	(0.068)
Lonely some of the time	-0.550***	(0.018)	-0.568***	(0.027)
Lonely a little of the time	-0.257***	(0.013)	-0.285***	(0.019)
Average monthly # of crimes committed in nhd	-0.001	(0.001)	—	—
Social amenities (higher = more amenities)	0.003**	(0.001)	—	—
Moved into nhd this year	-0.045**	(0.021)	—	—
Started at workplace this month	—	—	0.112**	(0.046)
≤ 2 months spent at this workplace	—	—	-0.215**	(0.088)
Turnover at this workplace	—	—	-0.032	(0.069)
Year fixed effects	Yes		Yes	
Household composition fixed effects	Yes		Yes	
Neighbourhood demographic characteristics	Yes		No	
Neighbourhood ethnic & income segregation	Yes		No	
Workplace demographic characteristics	No		Yes	
R^2	0.171		0.139	
# of observations	32,643		13,920	

Notes: this table displays coefficients from linear regressions of subjective wellbeing on our full set of controls, with no peer income variables. Standard errors are robust. The number of observations has been randomly rounded to base 3. Asterisks denote significance at *p<0.10, **p<0.05, ***p<0.01.

Table 4: Neighbourhood peer income regressions

Variable	(1)	(2)	(3)	(4)	(5)
Income (log)	0.016*** (0.005)	-0.013* (0.007)	-0.011 (0.007)	-0.012 (0.008)	-0.012 (0.008)
Median peer income	0.030 (0.022)	0.070*** (0.023)	0.034 (0.024)	-0.044 (0.035)	-0.062 (0.069)
Ordinal income rank		0.206*** (0.038)	0.196*** (0.038)	0.204*** (0.038)	0.202*** (0.038)
Std dev of peer income			0.113*** (0.020)	-0.073 (0.149)	-0.238 (0.178)
Top decile of peer income				0.154*** (0.041)	0.120** (0.059)
Bottom decile of peer income				-0.049 (0.048)	-0.076 (0.055)
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes	Yes
Group-level controls	No	No	No	No	Yes
R^2	0.169	0.169	0.170	0.171	0.172
# of observations	32,643	32,643	32,643	32,643	32,643

Notes: this table displays coefficients from linear regressions of subjective wellbeing on neighbourhood peer income variables and other covariates, for our full sample. The median peer income variable is the median log income in the nhd, the ordinal rank variable ranges from 0 (bottom rank) to 1 (top rank), the standard deviation variable is the standard deviation of log income in the nhd, and the top/bottom decile variables are the median log income in the top/bottom income decile in the nhd. All regressions include the full set of individual-level characteristics described in Table 2. Column 5 adds the full set of neighbourhood-level characteristics described in Table 2. Standard errors are robust. The number of observations has been randomly rounded to base 3. Asterisks denote significance at *p<0.10, **p<0.05, ***p<0.01.

Table 5: Workplace peer wage regressions

Variable	(1)	(2)	(3)	(4)	(5)
Wage (log)	0.019*	-0.026	-0.029*	-0.032*	-0.032*
	(0.011)	(0.017)	(0.017)	(0.017)	(0.017)
Median peer wage	0.028*	0.076***	0.098***	0.041	0.000
	(0.015)	(0.020)	(0.023)	(0.033)	(0.037)
Ordinal wage rank		0.157***	0.165***	0.167***	0.185***
		(0.046)	(0.046)	(0.046)	(0.047)
Std dev of peer wage			0.064**	0.245***	0.176**
			(0.032)	(0.076)	(0.079)
Top decile of peer wage				-0.008	-0.009
				(0.011)	(0.011)
Bottom decile of peer wage				0.078***	0.073**
				(0.029)	(0.030)
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes	Yes
Group-level controls	No	No	No	No	Yes
R^2	0.134	0.135	0.135	0.135	0.148
# of observations	13,920	13,920	13,920	13,920	13,920

Notes: this table displays coefficients from linear regressions of subjective wellbeing on workplace peer wage variables and other covariates, for our employed sample. The median peer wage variable is the median log wage in the workplace, the ordinal rank variable ranges from 0 (bottom rank) to 1 (top rank), the standard deviation variable is the standard deviation of log wages in the workplace, and the top/bottom decile variables are the median log wages in the top/bottom wage decile in the workplace. All regressions include the full set of individual-level characteristics described in Table 2. Column 5 adds the full set of workplace-level characteristics described in Table 2. Standard errors are robust. The number of observations has been randomly rounded to base 3. Asterisks denote significance at * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Workplace peer residual regressions

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Residual of log wage	0.013 (0.011)	-0.077*** (0.019)	-0.076*** (0.019)	-0.078*** (0.019)	-0.070*** (0.019)	-0.111** (0.043)
Median peer residual	0.015 (0.032)	0.100*** (0.035)	0.095** (0.037)	0.010 (0.054)	0.011 (0.055)	0.123 (0.086)
Ordinal residual rank		0.254*** (0.045)	0.253*** (0.045)	0.253*** (0.045)	0.256*** (0.045)	0.254*** (0.064)
Std dev of peer resid			-0.014 (0.032)	-0.131 (0.131)	-0.027 (0.136)	-0.304 (0.192)
Top decile of peer resid				0.095** (0.042)	0.063 (0.044)	0.084* (0.045)
Bottom decile of peer resid				-0.009 (0.044)	0.021 (0.045)	-0.084 (0.076)
Wage (log)						0.039 (0.040)
Median peer wage						-0.101 (0.064)
Ordinal wage rank						0.020 (0.065)
Std dev of peer wage						0.267** (0.128)
Top decile of peer wage						-0.013 (0.011)
Bottom decile of peer wage						0.100* (0.054)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes	Yes	Yes
Group-level controls	No	No	No	No	Yes	Yes
R^2	0.133	0.135	0.135	0.135	0.149	0.150
# of observations	13,920	13,920	13,920	13,920	13,920	13,920

Notes: this table displays coefficients from linear regressions of subjective wellbeing on workplace peer residual variables and other covariates, for our employed sample. Covariates are the same as in Table 5, except with peer residuals instead of peer wages. Column 6 adds peer wage variables. Standard errors are robust. The number of observations has been randomly rounded to base 3. Asterisks denote significance at *p<0.10, **p<0.05, ***p<0.01.

Table 7: Main results

Variable	(1)	(2)	(3)
Income (log)	-0.012 (0.008)		0.002 (0.021)
Median peer income	-0.062 (0.069)		-0.110 (0.107)
Ordinal income rank	0.202*** (0.038)		0.041 (0.087)
Std dev of peer income	-0.238 (0.178)		0.293 (0.272)
Top decile of peer income	0.120** (0.059)		-0.051 (0.093)
Bottom decile of peer income	-0.076 (0.055)		0.080 (0.085)
Residual of log wage		-0.105** (0.044)	-0.106** (0.044)
Median peer residual		0.118 (0.087)	0.119 (0.087)
Ordinal residual rank		0.247*** (0.065)	0.241*** (0.065)
Std dev of peer resid		-0.295 (0.194)	-0.328* (0.193)
Top decile of peer resid		0.075* (0.045)	0.090** (0.045)
Bottom decile of peer resid		-0.081 (0.077)	-0.090 (0.076)
Wage (log)		0.033 (0.040)	0.033 (0.042)
Median peer wage		-0.097 (0.065)	-0.104 (0.065)
Ordinal wage rank		0.031 (0.066)	0.007 (0.067)
Std dev of peer wage		0.308** (0.129)	0.265** (0.128)
Top decile of peer wage		-0.015 (0.011)	-0.015 (0.011)
Bottom decile of peer wage		0.112** (0.055)	0.099* (0.054)
Full set of controls	Yes	Yes	Yes
R^2	0.172	0.142	0.153
# of observations	32,643	13,920	13,920

Notes: coefficients from regressions of subjective wellbeing on neighbourhood and/or workplace peer income variables, with a full set of controls. Columns 1 and 2 replicate Columns 5 and 6 of Tables 4 and 6, respectively. Column 3 includes the full set of neighbourhood and workplace peer income variables, and full sets of neighbourhood-level and workplace-level controls. Standard errors are robust. Asterisks denote significance at * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Immobile specifications

Variable	Neighbourhood		Workplace	
	OLS	IV	OLS	IV
Income/residual	-0.010 (0.010)	-0.016 (0.018)	-0.087 (0.104)	-0.276 (0.258)
Median peer income/residual	-0.015 (0.092)	0.659 (0.696)	0.080 (0.186)	0.191 (1.573)
Ordinal income/residual rank	0.162*** (0.052)	0.203 (0.129)	0.395*** (0.132)	0.787*** (0.249)
Std dev of peer income/resid	-0.125 (0.235)	-1.446 (2.988)	-0.095 (0.415)	-0.644 (1.026)
Top decile of peer income/resid	0.042 (0.079)	0.050 (0.604)	0.058 (0.096)	0.298 (0.389)
Bottom decile of peer income/resid	-0.030 (0.073)	-0.433 (0.915)	-0.100 (0.155)	-0.240 (0.716)
Year fixed effects	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes
Group-level controls	Yes	Yes	Yes	Yes
R^2	0.168	0.165	0.172	0.162
# of observations	18,126	18,126	3,768	3,768

Notes: coefficients from OLS and 2SLS regressions of subjective wellbeing on neighbourhood peer income variables and workplace peer wage & residual variables. The OLS regressions are the same specification as Columns 1 and 2 of Table 7, except restricted to the sample of immobile people (people who remained in the same neighbourhood/workplace for at least 5 years prior to their interview month). The IV regressions instrument for the peer income/wage/residual variables with the change in those variables over the last 5 years. Average (minimum) first-stage F -statistics are 4087.2 (359.2) for the neighbourhood IV regression and 50.5 (10.7) for the workplace IV regression. Standard errors are robust. Asterisks denote significance at * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Placebo test for neighbourhood variables

Variable	No placebo	Placebo
Income (log)	-0.029** (0.014)	-0.029** (0.014)
Median peer income	-0.113 (0.126)	-0.135 (0.162)
Ordinal income rank	0.293*** (0.072)	0.293*** (0.074)
Std dev of peer income	-0.137 (0.326)	-0.199 (0.337)
Top decile of peer income	0.122 (0.110)	0.235 (0.186)
Bottom decile of peer income	-0.065 (0.101)	-0.062 (0.103)
Median peer income (5-year lag)		0.028 (0.138)
Ordinal income rank (5-year lag)		0.000 (0.042)
Std dev of peer income (5-year lag)		0.106 (0.084)
Top decile of peer income (5-year lag)		-0.129 (0.181)
Bottom decile of peer income (5-year lag)		0.014 (0.069)
Year fixed effects	Yes	Yes
Individual-level controls	Yes	Yes
Group-level controls	Yes	Yes
R^2	0.188	0.188
# of observations	10,437	10,437

Notes: Column 1 displays coefficients from an OLS regression of subjective wellbeing on our neighbourhood peer income variables and a full set of controls (the same specification as Column 1 of Table 7, except restricted to people who moved into their current neighbourhood sometime in the past 5 years). Column 2 displays coefficients from the same regression, with the addition of 5-year lags of the peer income variables that are calculated “as if” each person lived in their current neighbourhood 5 years ago. Note that this sample is not simply the complement of the “immobile” sample from Table 8, since it excludes people who are not present in the data 5 years ago. Standard errors are robust. Asterisks denote significance at *p<0.10, **p<0.05, ***p<0.01.

Table 10: Placebo test for workplace variables

Variable	No placebo	Placebo
Residual of log wage	-0.206* (0.114)	-0.189 (0.118)
Median peer residual	-0.136 (0.136)	-0.110 (0.139)
Ordinal residual rank	-0.131 (0.116)	-0.148 (0.117)
Std dev of peer resid	-0.124 (0.196)	-0.106 (0.207)
Top decile of peer resid	0.127* (0.068)	0.125* (0.069)
Bottom decile of peer resid	0.770*** (0.294)	0.739** (0.306)
Median peer residual (5-year lag)		-0.042 (0.045)
Ordinal residual rank (5-year lag)		0.075 (0.091)
Std dev of peer resid (5-year lag)		0.021 (0.165)
Top decile of peer resid (5-year lag)		-0.027 (0.051)
Bottom decile of peer resid (5-year lag)		0.153 (0.207)
Year fixed effects	Yes	Yes
Individual-level controls	Yes	Yes
Group-level controls	Yes	Yes
Peer wage variables	Yes	Yes
R^2	0.181	0.182
# of observations	4,485	4,485

Notes: Column 1 displays coefficients from an OLS regression of subjective wellbeing on our workplace peer wage and residual variables and a full set of controls (the same specification as Column 2 of Table 7, except restricted to people who moved into their current workplace sometime in the past 5 years). Column 2 displays coefficients from the same regression, with the addition of 5-year lags of the peer wage/residual variables that are calculated “as if” each person worked in their current workplace 5 years ago. Note that this sample is not simply the complement of the “immobile” sample from Table 8, since it excludes people who are not present in the data 5 years ago and people whose workplace did not exist 5 years ago.

Standard errors are robust. Asterisks denote significance at * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.