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

Shakked Noy†‡ Isabelle Sin†

November 3, 2020

Abstract

We investigate how a person's happiness is affected by the incomes of her neighbours and coworkers. Using an unprecedentedly rich combination of administrative and survey data, we establish two central results. First, a person's happiness is primarily sensitive to her ordinal rank within her peer income distribution: people are happier the higher their income rank. Second, workplace rank matters much more than neighbourhood rank. 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 access to the data. Funding: the Victoria University of Wellington School of Economics and Finance funded access to the data, but had no involvement in the research process. Declarations of interest: none.

†Motu Economic and Public Policy Research

‡Corresponding author. Email: shakked.noy@gmail.com, postal address: 97 Cuba Street, Wellington, New Zealand.

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 authors, not Statistics NZ, Motu Economic and Public Policy Research, 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 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 exceptionally 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 the 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 two baseline results. First, a person's subjective wellbeing is primarily sensitive to her ordinal income rank among her peers: a 10 percentage point increase in a person's ordinal income rank is associated with an increase in happiness equal to 1/10th the happiness premium associated with eliminating moderate physical pain. Conditional on a person's ordinal rank, there is suggestive evidence that her subjective wellbeing is positively correlated with other features of her peer income distribution, such as the median income or standard deviation of income among her peers, but these results are less robust. Second, workplace income rank matters more than neighbourhood income rank in predicting a person's subjective wellbeing.

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 coefficient stability analysis, and use the techniques developed by Cinelli and Hazlett (2020) 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 is unlikely to 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 definitive, but the cumulative evidence they provide suggests that the associations we observe are causal. Any remaining endogeneity driving our results must be due to omitted factors that 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 unlikely that there exist important omitted variables that satisfy these conditions, and therefore conclude that our estimates likely reflect a causal effect of peer income on wellbeing.

We make several contributions to the existing literature on income comparisons and subjective wellbeing. 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 more credible than most estimates in the existing literature.

That being said, our analyses still face important limitations (mainly the fact that we are unable to use individual or group fixed effects), as we discuss in Section 4.5.

A small number of existing studies provide causal evidence of the effects of income comparisons on subjective wellbeing. Perez-Truglia (2020) shows, using a quasi-experimental design, that the publication of Norwegian tax records increased the income-happiness gradient in Norway. Similarly, Card et al. (2012) show in a field experiment that providing University of California employees with information about the salaries of their peers reduced the job satisfaction of relatively low-paid employees.

In addition, a number of existing studies use field experiments, lab experiments, or quasi-experimental designs to estimate the causal effects of income comparisons on non-wellbeing outcome variables (Zizzo and Oswald (2001), Kuziemko et al. (2014), Charness and Kuhn

(2007), Goerg et al. (2010), Bracha et al. (2015), Breza et al. (2018), Cullen and Perez-Truglia (2018), Dube et al. (2019)).

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 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, most 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 distribution. In particular, we choose to control for the median income among a person's peers, that person's ordinal rank among her peers, and the standard deviation of income among her peers. In supplementary specifications, we replace the "median peer income" control with controls for the top and bottom deciles of peer income. Our results suggest that the most stable predictor of a person's subjective wellbeing is their ordinal income rank among their peers: people are happier the higher their ordinal rank is, a remarkably robust result barely varies across all our specifications. We also find some suggestive evidence that, conditional on their ordinal rank, people's wellbeing is positively associated with the median

[1] See Boskin and Sheshinski (1978), Oswald (1983), and Dube et al. (2019).

[2] A small number of existing studies do consider more than two peer income variables. For example, Godechot and Senik (2015) consider (a) the median wage among a worker's coworkers, (b) the top 1% of wages among a worker's coworkers, and (c) the median wage among similar workers in the region.

income, standard deviation of income, or top/bottom decile of income among their peers, but these results are much less robust to changes in specification.

The importance of ordinal rank concerns may reflect either a fact about people's information or a fact about their utility functions. One hypothesis is that people are primarily sensitive to their ordinal rank because they lack information about the exact cardinal differences between themselves and their peers.[3] An alternative hypothesis is that people are aware of both the ordinal and cardinal differences between themselves and their peers, but care primarily about the ordinal differences, perhaps because they ultimately care about their position in an ordinal status hierarchy that is determined by income.

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 most important peer income variable, 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.[4] We are thereby able to test whether people care more about comparisons to their neighbours or comparisons to their coworkers. Moreover, since neighbourhood and workplace peer income measures are strongly correlated with each other, we avoid the omitted variable bias that complicates the interpretation of studies that only control for one of the two sets of 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 a person's ordinal wage rank in her workplace is strongly predictive of her 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. Once again, there are two possible interpretations of this result. One interpretation is that

[3]It is 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?"

[4]Goerke 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, while we use objective measures of peer income. Godechot and Senik (2015) consider two comparison groups for the workers in their dataset (the worker's coworkers, and similar workers in the region), so they consider both a workplace and a geographic reference group, but they do not directly consider everyone in the worker's neighbourhood as a reference group.

people are primarily affected by workplace comparisons because they have more information about the incomes of their coworkers; the other is that people intrinsically care more about workplace comparisons, perhaps because status concerns are more salient in the workplace.

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 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 is a cross-sectional survey that covers New Zealand residents aged 15 and over. Each wave of 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 GSS waves in our dataset contain 43,020 observations (each of which is a distinct individual). 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 13,917 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 17% 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 13,917 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 and from superannuation 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 tertiary qualification obtained in New Zealand from 1994 onwards, and every secondary school qualification obtained from 2007 onwards), 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 neighbourhood. Unfortunately, these records only begin in July of 2014, so for each neighbourhood, we take the average number of crimes per month and the average seriousness of the crimes committed over the 2014-2018 period, and use these numbers as a time-invariant measure of the level of crime in that neighbourhood. In order to measure the seriousness of different crimes, we use seriousness scores from the Ministry of Justice.

Second, we use data on business locations to control for the presence in each area unit 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.[5] After calculating the number of amenities of each type in every area unit, we conduct a factor analysis and use the regression method to construct a single variable that proxies for the overall presence of social amenities in a particular area unit.

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.
>
> Looking at showcard 17, where 0 is completely dissatisfied and 10 is completely satisfied, how do you feel about your life as a whole?[6]

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. In Section 5 we run robustness checks to confirm that this pooling of data does not qualitatively affect our results.

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:[7] 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

[5]The full list of business types we classify as social amenities is: supermarkets, fruit and vegetable retailers, cafes and restaurants, government administrative centres, primary and secondary schools, tertiary educational institutions, medical institutions, zoos, botanic gardens, nature reserves, performing arts institutions, gyms, sports, recreation clubs, and amusement parks.

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

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

sense of their size. We define neighbourhoods at the area unit level because area units are the natural level of interaction in an urban environment:[8] 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. Our GSS individuals belong to 1,557 distinct area units.

A person's neighbourhood peer group is defined to be everyone else in that person's area unit for whom we have income data.[9] 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 establishment 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. Our employed GSS individuals belong to 9,153 distinct workplaces.

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

[8] 98% of our observations are from urban areas.

[9] Our income data covers all individuals who receive taxable income or any kind of government benefit, and should therefore cover a very high percentage of the New Zealand adult population.

[10] Note 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. However, this is unlikely to be a significant problem, since the wage distributions in the different outlets of a single business are likely to be quite similar to each other, so even if an employee is accidentally matched to the wrong outlet location, the measures of peer income that we calculate for them should still be fairly accurate.

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), Barrington-Leigh and Helliwell (2008)). 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 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 ordinal income rank in neighbourhood g, and the standard deviation of log income in neighbourhood g. In supplementary specifications, we replace the "median peer income" variable with controls for the mean log income within the top income decile in neighbourhood g and the mean log income within the bottom income decile in neighbourhood g. We refrain from including the median peer income and top and bottom decile variables in the same specification because they are very strongly correlated with each other, resulting in a massive inflation of standard errors. 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, and the standard deviation of log wages in workplace w. In supplementary specifications, we replace "median peer wages" with the mean log wage within the top wage decile in workplace w, and the mean 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 to capture the idea that workers compare salaries with each other, rather than comparing overall incomes.

[11] There 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.

[12] There 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.

More formally, consider a peer group g containing N individuals, indexed by $i \in \{1, \dots, N\}$. Let x_i denote either an individual's log income (in the neighbourhood case) or her log wage (in the workplace case). Letting M denote the median function and letting $1_{()}$ denote an indicator function, we define the following peer income variables for individual j in group g.

- $median_g = M(\{x_i\}_{i=1}^N)$ [The median income in group g.]
- $rank_{jg} = \frac{1}{N-1} \sum_{i \neq j} 1_{(x_j \geq x_i)}$ [This is a variable that varies continuously between 0 and 1, and equals 0 for the person with the lowest or lowest-equal income in the peer group and 1 for the person with the highest or highest-equal income in the peer group. Essentially, it is equal to individual j's rank in the distribution (which is a number between 0 and $N-1$, with $N-1$ being the highest rank), divided by $N-1$.]
- $sd_g = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2}$ [The standard deviation of income in group g.]
- $top_g = \frac{1}{\sum_{i=1}^N 1_i^T} \sum_{i=1}^N 1_i^T x_i$ where 1^T is an indicator for being in the top income decile in group g. [The average income among individuals in the top income decile in group g.]
- $bottom_g = \frac{1}{\sum_{i=1}^N 1_i^B} \sum_{i=1}^N 1_i^B x_i$ where 1^B is an indicator for being in the bottom income decile in group g. [The average income among individuals in the bottom income decile in group g.]

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. Members of our full sample are relatively uneducated: more than 50% have never earned a qualification beyond a high school diploma. This may be due to the high concentration of older people (who tend to have fewer formal qualifications) in our sample.

Ninety-eight percent of our full-sample observations live in an urban area, and 15% moved into their current area unit some time in the past year. A majority of our full-sample observations report that they are "never lonely" (64%), and a slim majority report that they are "not at all" hindered by physical pain (55%).

Column 2 displays summary statistics for our employed subsample. 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 employees, but the average worker in our sample has 341 coworkers (due to more workers being sampled from larger workplaces).

Table 2 displays the correlation matrices of our neighbourhood and workplace peer income variables in the full GSS sample. The correlations in these matrices are intuitive: the median and standard deviation of income in a peer group are positively correlated, the median and top decile are positively correlated, and so on.

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

$$SWB_{itg} = \alpha + \beta ref_{itg} + \gamma indiv_{it} + \delta group_{tg} + \theta_t + \varepsilon_{itg} \tag{1}$$

where ref_{itg} is own income and 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 3 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, the total income of other members of the household, 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 13,917 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 4. 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. Māori and Pasifika are happier than Europeans, conditional on the other covariates. There is no correlation between a person's happiness and the presence of social amenities in their neighbourhood or the crime rate in their neighbourhood.

We omit the subjective loneliness and physical health variables from subsequent regressions, since these variables are likely to be affected by any within-individual heterogeneity in self-reporting that affects our subjective wellbeing variable. For example, an individual who is feeling unusually pessimistic on the day they are surveyed may give themselves a low wellbeing score and a high loneliness score. This will create the (spurious) impression of a relationship between wellbeing and loneliness, even if none exists. However, in Section 5 we run a robustness check that shows the omission of these controls does not qualitatively affect our results.

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

3.1.1 Neighbourhood results

Table 5 reports the results from five regressions. Column 1 displays the results from a regression of subjective wellbeing on median peer income, a full set of individual-level controls, and 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 swaps out the median peer income variable for controls for the top and bottom deciles of peer income; and Column 5 switches back to the median peer income variable and adds neighbourhood-level controls.

In Column 1, the coefficient on median peer income is small, positive, and significant, suggesting counterintuitively that a person's happiness is *increasing* in the income of her peers. This result contradicts the majority of findings (though not all findings) in the income comparisons literature.[13]

Adding a term for ordinal income rank in Column 2 clarifies matters. The coefficient on median peer income in Column 1 is positive not because people are unconditionally happier when they live near wealthier peers, 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, or because living in a wealthier neighbourhood enhances a person's status relative to the rest of society.

The Column 2 results also explain why the coefficient on median peer income in Column 1 is small and positive. An increase in median peer income, conditional on a person's own 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). The latter effect is stronger in Column 1, which is why Column 1 gives the impression that there is a small positive 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 partially reflects the omitted influence of the standard deviation of peer income. In Column 3, the coefficient on median peer income has nearly halved in size, and the subjective wellbeing

[13]See, for example, Dittmann and Goebel (2010) and Knies (2012).

[14]This 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) control only for median peer income and ordinal income rank, so Columns 3-5 extend on their results.

of people in our sample is 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 replaces the median peer income variable with controls for the top and bottom deciles of the neighbourhood income distribution,[16] shows a strong positive correlation between the top decile of income in a neighbourhood and the wellbeing of its residents. 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 the neighbourhood's endowment of public goods and other amenities.

Column 5 shows that the addition of neighbourhood-level controls causes the positive coefficients on the median and standard deviation of peer income from Column 3 to disappear. This suggests that the positive correlation between a person's subjective wellbeing and the median/standard deviation of their peer income distribution is driven by the endogenous influence of other group-level characteristics.

Throughout this process, the coefficient on ordinal income rank remains stable, 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 0.03 standard deviation increase in subjective wellbeing. This effect is equal to 13% 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 (see Table 4).[17]

The coefficient on own income is positive, significant, and declining in magnitude across Columns 1-5. This result is consistent with the existing literature, which finds that controlling for peer income reduces, but does not eliminate, the positive effect of own income on subjective wellbeing (see Clark et al. (2018)).

[15]Knight 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.

[16]Recall that the "top decile" variable is the mean income in the top income decile in the neighbourhood, and the "bottom decile" variable is analogously defined.

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

Two broad patterns in Table 5 are worth emphasizing. First, the coefficients on the peer income variables are sensitive to which peer income variables we include in the regression, due to the strong correlations between the different peer income variables (see Table 2). 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 nearly halves when we add a control for the standard deviation of neighbourhood income. This coefficient behaviour confirms our prediction that studies that control only for median peer income, and omit other parameters of the peer income distribution, will be strongly influenced by omitted variable bias, and will therefore convey a misleading picture of the relationship between the peer income distribution and subjective wellbeing.

Second, the coefficient on ordinal income rank is unusually stable. While the other peer income coefficients fluctuate in magnitude and sign 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 suggests that ordinal income rank is a much more robust predictor of subjective wellbeing than the cardinal peer income variables.

A relevant question is whether we find this unstable relationship between subjective wellbeing and the cardinal peer income variables because the cardinal peer income variables are all strongly correlated, making their coefficients insignificant due to large standard errors when they are all included in the same regression. However, as we show in Appendix Table A1, this is not the case: even if we include each cardinal peer income variable individually alongside ordinal income rank, their coefficients are all still small and insignificant as long as we include group-level controls. Thus the unique robustness of the ordinal rank variable is not simply attributable to the fact that it is less strongly correlated with the other peer income variables.

3.1.2 Workplace results

Table 6 replicates for the employed subsample the analysis presented in Table 5, using workplace peer wage variables instead of neighbourhood peer income variables. It shows similar patterns to Table 5. Column 1 displays a weak positive correlation between a worker's subjective wellbeing and the median wage in her workplace. Column 2 shows that this weak positive correlation obscures a strong positive relationship between a worker's subjective wellbeing and her ordinal wage rank, and conditional on her wage rank a worker's subjective wellbeing is positively associated with her coworkers' median wages. Columns 3 and 4 show that a worker's subjective wellbeing is positively correlated with the standard deviation and bottom decile of wages in her workplace. This may reflect the fact that workplaces with a high dispersion of wages contain more opportunities for career mobility, and workplaces with a high bottom decile of wages adopt a management style that treats their lowest-paid workers relatively well. Finally, unlike Table 5, when we add workplace-level controls in Column 5 the coefficient on median peer wages remains positive and significant, though it decreases in magnitude. The coefficient on ordinal wage rank actually increases in magnitude with the addition of group-level controls.

Once again, in Table 6 the coefficient on ordinal wage rank is very stable (it varies by only 13% of its initial magnitude), while the coefficients on the other peer wage variables change substantially in magnitude across specifications.[18]

3.2 Preferred specifications

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 5. Column 2, which replicates Column 5 of Table 6, displays the effects of workplace peer wages for our employed sample. Column 3 contains both neighbourhood and workplace peer income variables for our employed sample.

Table 7 shows two main 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, for employed people, only workplace comparisons appear to matter. In Column 3, the only peer income variable that significantly predicts subjective wellbeing is ordinal wage rank in the workplace. 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 primarily about

[18]As before, we can question whether this because ordinal wage rank is less strongly correlated with the other peer wage variables. Analogously to Appendix Table A1, we demonstrate in Appendix Table A2 that the fact that the coefficients on the cardinal peer wage variables in Column 5 are mostly small and statistically insignificant cannot be explained by the fact that the cardinal variables are strongly correlated with each other. The median peer wage and standard deviation of peer wage variables now have marginally significant effects on subjective wellbeing, but this is not a very robust result across our specifications.

[19]An 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 (results not presented). This produces qualitatively identical results, which rules out this alternative explanation.

their income rank within their peer group, either because they are better informed about their income rank or because they care more about ordinal comparisons than cardinal ones. Similarly, rank in the workplace matters more than rank in the neighbourhood, either because people are better informed about their coworkers' incomes or because they care more about workplace comparisons.

There is suggestive evidence that people's subjective wellbeing is positively correlated with other features of their peer income distributions, but this evidence is not robust enough to produce a definitive conclusion. Moreover, the most plausible explanations of why the cardinal peer income variables might be positively associated with subjective wellbeing implicitly invoke omitted variable bias. For example, explaining the positive correlation between median peer income and subjective wellbeing by saying that neighbourhoods with higher median incomes have more social capital or better amenities concedes that the positive correlation is driven by endogeneity.

In the Appendix, we present four brief extensions on these baseline results. We show that the effects of median peer income do not appear to be different for people below the median income in their peer group, compared to people above the median; that the effects of ordinal income rank do not appear to vary with the size of a person's peer group; that exchanging the top decile of income for the top 1% of income in our regressions does not qualitatively affect our results; and that using non-logged measures of peer income does not affect our results either.

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 (Brodeur and Flèche (2019)). 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 address the influence of omitted group-level characteristics.[20] In this section, we pursue one general strategy (sensitivity analysis), and three specific strategies, two of which deal with individual selection and one of which deals with the omitted influence of fixed amenities.

One final source of endogeneity is possible: the effect of coworkers' earnings on effort and therefore wages. There is some evidence (Breza et al. (2018)) that workers reduce their effort when they observe their coworkers unfairly earning more than them. If workers earn incentivized wages, then reductions in a worker's effort will decrease that worker's wage. If, moreover, wage decreases directly affect workers' subjective wellbeing, then reductions in a worker's effort will decrease her subjective wellbeing by decreasing her wages. It follows that if all these assumptions hold, then increases in coworker wages will reduce a worker's subjective wellbeing in part *through* decreases in her wage. Controlling for a worker's own wage in our regressions therefore partially confounds the effect of coworker wages. At worst, however, controlling for a worker's own wage will bias our estimate of the effects of coworker wages towards zero.

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 vul-

[20]The 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.

nerability of those coefficients to unobserved confounders. Furthermore, Cinelli and Hazlett (2020) develop several formal measures of the robustness of a set of estimates to the presence of unobserved confounders. We therefore begin by providing informal visual evidence on the stability of our peer income coefficients, and then calculate Cinelli and Hazlett's (2020) 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, for which we can (albeit imperfectly) control. 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 the 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 the coefficients 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 main 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 wage coefficients. 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 coefficients on the median and standard deviation of peer income decline almost monotonically towards zero as we add controls, 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 60%). Age is a major determinant of both a person's ordinal income rank and her subjective well-being, so it makes sense that adding it as a control will substantially affect our estimate of the effect of ordinal rank. The fact that introducing the "receives welfare" dummy has a large negative effect on the ordinal rank coefficient is interesting, and perhaps 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 wage rank in Figure 4 is exceptionally stable and does not fluctuate at all. By contrast, the median and standard deviation of peer wage coefficients decline monotonically towards zero. The contrast between the stability of the ordinal rank coefficient and the shrinking of the cardinal coefficients is even more striking in Figure 4 than in Figure 3.

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 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. Moreover, this visual evidence is informal; as Gelbach (2016) notes, it depends on the order in which we add covariates to the model.

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 (2020). Consider an augmented version of Equation (1), with the addition of a vector of unobserved confounding variables $unobs_{itg}$.

$$SWB_{itg} = \alpha_{\text{aug}} + \beta_{\text{aug}} ref_{itg} + \gamma_{\text{aug}} indiv_{it} + \delta_{\text{aug}} group_{tg} + \theta_t^{\text{aug}} + \lambda_{\text{aug}} unobs_{itg} + \xi_{itg} \quad (2)$$

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 (1). 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} (conditional on the other covariates in the model).

Cinelli and Hazlett (2020) 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_{ig} and between $unobs_{itg}$ and SWB_{itg} in order to explain away our results. The Robustness Value for the coefficient on a peer income variable T is equal to

$$RV = \frac{1}{2}\left(\sqrt{\left(\frac{R^2_{Y\sim T|C}}{1-R^2_{Y\sim T|C}}\right)^2 + 4\frac{R^2_{Y\sim T|C}}{1-R^2_{Y\sim T|C}}} - \frac{R^2_{Y\sim T|C}}{1-R^2_{Y\sim T|C}}\right) \tag{3}$$

where $R^2_{Y\sim T|C}$ 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 4.2%, which means that unobserved confounders would have to explain at least 4.2% 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 wage rank (from Column 2 of Table 7) is 4.3%. When neighbourhood peer income variables are added (Column 3 of Table 7) the Robustness Value for ordinal wage rank falls to 3.2%.

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 1.4% of the residual variation in subjective wellbeing. Thus unobserved confounders of the required importance would need to be two or three times as important as all of those variables combined in determining the observed distribution of subjective wellbeing.

We believe it is unlikely that there exist unobserved confounders that are strong enough to explain away our results. These confounders would have to satisfy three conditions. First, they would need to be very important in explaining the residual variation in subjective well-being. 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 plausible that there exist unobserved confounders that satisfy the first and second conditions. However, it is 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. As additional evidence for this conclusion, we implement three additional strategies.

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 (Bottan and Perez-Truglia (2017)). 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 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.

If these arguments are correct, our estimates in this section should be less influenced by selection bias than our estimates from Section 3. To test whether this affects our results, 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. If the coefficients are similar, then individual selection is unlikely to be heavily influencing our results.

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.

[21] Technically, 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.

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 259 (143), and for the workplace peer wage variables they are 130 (21).

In both our neighbourhood and workplace immobile samples, the OLS and IV coefficients are similar to the Table 7 estimates and are not significantly different. In the workplace case, the IV coefficient is larger. The IV coefficient on neighbourhood income rank is somewhat smaller than the OLS coefficient and is statistically insignificant, which suggests that our neighbourhood results are perhaps partially driven by endogeneity. However, the IV coefficient's statistical insignificance is also attributable to the larger standard errors involved in IV estimation.

If the assumption underlying this section is correct, this shows that individual selection into peer groups is unlikely to be driving our main results.

It is possible that our assumption fails because individuals select into peer groups partially in anticipation of future changes in the distribution of income in those groups. However, as long as our strategy significantly *reduces* the degree of individual selection (even if it does not fully *eliminate* that selection), our findings in this section support a causal interpretation of our baseline 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 empirical challenge 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 occasionally 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 the five following conditions all hold.

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 wage 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 ordinal rank 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 plot the coefficients only for our main variable of interest, ordinal wage rank). There is no evidence that the comparison coefficients are systematically larger than the layoff coefficients: the comparison and layoff coefficients are very similar in magnitude and there are no statistically significant differences. Moreover, 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 a weak upwards trend in magnitude for the layoff-based wage rank coefficients.

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. We then calculate 5-year placebo lags of all of our peer income variables "as if" each individual had been in their current location 5 years ago. More precisely, consider individual i in location ℓ at time t. We observe individual i's income at time $t-5$ and the distribution of income in location ℓ at time $t-5$, and then calculate all of individual i's peer income variables "as if" individual i

had been in location ℓ at time $t-5$. In reality, individual i was *not* in location ℓ at time $t-5$, since we restrict to individuals who moved location within the past 5 years. We make this restriction in order to avoid our lagged peer income variables picking up a causal effect of lagged peer income, for example due to reference-dependent preferences or adaptation.

Having calculated the lagged peer income variables, we include these lagged variables in our standard regressions, and check whether they attract significant coefficients. The intuition behind this strategy is that 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 is unlikely to be driving our results.

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

4.5 Limitations

Despite the promising results of Sections 4.1–4.4, our empirical analysis has a few key limitations.

First, unlike some previous studies in the income comparisons literature, we are unable to include individual fixed effects in our regressions (due to the cross-sectional nature of the

GSS). This means that we are unable to directly control for the influence of unobserved time-invariant individual confounders. Our results in Sections 4.1–4.3 provide suggestive evidence that our results are unlikely to be driven by these sorts of confounders, but we cannot entirely rule this out.

Similarly, we are unable to use peer group fixed effects, since the individuals in the GSS are widely dispersed over a large number of peer groups. Consequently, our results may be affected by unobserved time-invariant group-level confounders. The results of Sections 4.1 and 4.4 suggest that this type of confounder is unlikely to be driving our results, but again we cannot definitively rule out this possibility.

While the results of Sections 4.2–4.4 provide solid evidence that our main estimates have a causal interpretation, this evidence is not conclusive. Our strategies in these sections are based on untestable assumptions, and our estimates in these sections are less statistically precise than our main estimates (due to the smaller sample sizes required by our strategies).

As a result, while we believe we have made a substantive contribution to the credibility of the existing body of evidence on income comparisons, our findings should not be interpreted as definitive evidence that income comparisons have a causal effect on subjective wellbeing, or that they have an effect in precisely the ways suggested by our results.

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 of Figure 6 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. We find that 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.[22] The results are displayed in the fourth column of Figure 6: the resulting coefficient on neighbourhood income rank is 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 wage 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).

We then 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)).[23] As a result, we may want to restrict the definition of a peer group so that a person's peer group contains only individuals who 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.[24]

The results from these narrower definitions are displayed in the next four columns of Figure 6. The coefficient on neighbourhood income rank is declines very slightly when narrower definitions are used; the estimated coefficient on workplace wage rank is similarly affected. On any of the narrower definitions, both coefficients are still positive and statistically significant.

Next, we test whether the inclusion of controls for self-reported physical health and loneliness affect our results. The ordinal rank coefficients when these controls are included are displayed in the final column of Figure 6. Neither ordinal rank coefficient experiences a large drop in magnitude, which is highly encouraging; since potential unobserved confounders are likely to be similar in nature to physical health and loneliness, the fact that the inclusion of these controls does not substantively affect our estimates is a positive sign of the robustness

[22] 65 is the unofficial retirement age in New Zealand, and the age at which government superannuation plans begin paying out.

[23] See also Kingdon and Knight (2007).

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

of our estimates.

Overall, Figure 6 shows that our main estimates are 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 virtually all are 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 three broad results.

First, peer income appears to have 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 strongly but not definitively suggest that this correlation is likely to capture a causal relationship.

Second, people are affected specifically by 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) is unlikely to effectively ameliorate the impacts of income comparisons. It provides suggestive 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 little 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 a lot of conspicuous consumption takes place in the neighbourhood (large houses, fancy cars, and so on).

These insights make a substantial contribution to the existing understanding of the relationship between subjective wellbeing and peer income, and increase the credibility of existing estimates by showing that omitted variable bias cannot explain the observed correlation between peer income and subjective wellbeing.

References

Akerlof, G.A, Yellen, J.L, 1990. The fair wage-effort hypothesis and unemployment. The Quarterly Journal of Economics 105, 255–283.

Altonji, J.G, Elder, T.E, Taber, C.R, 2005. Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. Journal of Political Economy 13, 151–184.

Angrist, J, Pischke, J.S, 2009. Mostly Harmless Econometrics: Princeton University Press.

Barrington-Leigh, C.P, Helliwell, J.F, 2008. Empathy and emulation: Life satisfaction and the urban geography of comparison groups. NBER WP No. 14593 , 1–39.

Boskin, M.J, Sheshinski, E, 1978. Optimal redistributive taxation when individual welfare depends on relative income. The Quarterly Journal of Economics 92, 589–601.

Bottan, N.L, Perez-Truglia, R, 2017. Choosing your pond: Location choices and relative income. NBER WP No. 23615 , 1–47.

Bracha, A, Gneezy, U, Loewenstein, G, 2015. Relative pay and labour supply. Journal of Labor Economics 33, 297–315.

Breza, E, Kaur, S, Shamdasani, Y, 2018. The morale effects of pay inequality. The Quarterly Journal of Economics 133, 611–663.

Brodeur, A, Flèche, S, 2019. Neighbors' income, public goods, and well-being. Review of Income and Wealth 65, 1–39.

Brown, G.D.A, Gardner, J, Oswald, A.J, Qian, J, 2008. Does wage rank affect employees' well-being? Industrial Relations 47, 355–389.

Card, D, Mas, A, Moretti, E, Saez, E, 2012. Inequality at work: the effect of peer salaries on job satisfaction. American Economic Review 102, 2981–3003.

Charness, G, Kuhn, P, 2007. Does pay inequality affect worker effort? experimental evidence. Journal of Labor Economics 25, 693–723.

Chetty, R, Hendren, N, 2018. The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. The Quarterly Journal of Economics 133, 1107–1162.

Chetty, R, Hendren, N, Kline, P, Saez, E, 2014. Where is the land of opportunity? the geography of intergenerational mobility in the united states. The Quarterly Journal of Economics 129, 1553–1623.

Cinelli, C, Hazlett, C, 2020. Making sense of sensitivity: Extending omitted variable bias. Journal of the Royal Statistical Society, Series B (Statistical Methodology) 82, 39–67.

Clark, A.E, Flèche, S, Layard, R, Powdthavee, N, Ward, G, 2018. The Origins of Happiness: the Science of Well-Being over the Life Course: Princeton University Press.

Clark, A.E, Frijters, P, Shields, M.A, 2008. Relative income, happiness, and utility: an explanation for the easterlin paradox and other puzzles. Journal of Economic Literature 46, 95–144.

Clark, A.E, Kristensen, N, Westergård-Nielsen, N, 2009a. Economic satisfaction and income rank in small neighbourhoods. Journal of the European Economic Association 7, 519–527.

Clark, A.E, Kristensen, N, Westergård-Nielsen, N, 2009b. Job satisfaction and co-worker wages: Status or signal? The Economic Journal 119, 430–447.

Clark, A.E, Masclet, D, Villeval, M.C, 2010. Effort and comparison income: Experimental and survey evidence. Industrial and Labor Relations Review 63, 407–426.

Clark, A.E, Senik, C, 2010. Who compares to whom? the anatomy of income comparisons in europe. The Economic Journal 120, 573–594.

Cullen, Z, Perez-Truglia, R, 2018. How much does your boss make? the effects of salary comparisons. NBER WP No. 24841 , 1–52.

Dittmann, J, Goebel, J, 2010. Your house, your car, your education: the socioeconomic situation of the neighborhood and its impact on life satisfaction in germany. Social Indicators Research 96, 497–513.

Dixon, S, Stillman, S, 2009. The impact of firm closures on workers' future labour market outcomes. Statistics New Zealand Working Paper .

Dube, A, Giuliano, L, Leonard, J, 2019. Fairness and frictions: the impact of unequal raises on quit behavior. American Economic Review 109, 620–663.

Ferrer-i-Carbonell, A, 2005. Income and well-being: an empirical analysis of the comparison income effect. Journal of Public Economics 89, 997–1019.

Festinger, L, 1954. A theory of social comparison processes. Human Relations 7, 117–140.

Gelbach, J.B, 2016. When do covariates matter? and which ones, and how much? Journal of Labor Economics 34, 509–543.

Godechot, O, Senik, C, 2015. Wage comparisons in and out of the firm. evidence from a matched employer-employee french database. Journal of Economic Behavior & Organization 117, 395–410.

Goerg, S.J, Kube, S, Zultan, R, 2010. Treating equals unequally: Incentives in teams, workers' motivation, and production technology. Journal of Labor Economics 28, 747–772.

Goerke, L, Pannenberg, M, 2015. Direct evidence for income comparisons and subjective well-being across reference groups. Economics Letters 137, 95–101.

Kingdon, G.G, Knight, J, 2007. Community, comparisons, and subjective well-being in a divided society. Journal of Economic Behavior & Organization 64, 69–90.

Knies, G, 2012. Income comparisons among neighbours and satisfaction in east and west germany. Social Indicators Research 106, 471–489.

Knight, J, Song, L, Gunatilaka, R, 2009. Subjective well-being and its determinants in rural china. China Economic Review 20, 635–649.

Kuziemko, I, Buell, R.W, Reich, T, Norton, M.I, 2014. 'last-place aversion': Evidence and redistributive implications. Quarterly Journal of Economics 129, 105–149.

Luttmer, E.F.P, 2005. Neighbors as negatives: Relative earnings and well-being. The Quarterly Journal of Economics 120, 963–1002.

Morawetz, D, Atia, E, Bin-Nun, G, Felous, L, Gariplerden, Y, Harris, E, Soustiel, S, Tombros, G, Zarfaty, Y, 1977. Income distribution and self-rated happiness: Some empirical evidence. The Economic Journal 87, 511–522.

Oster, E, 2019. Unobservable selection and coefficient stability: Theory and evidence. Journal of Business and Economic Statistics 37, 187–204.

Oswald, A.J, 1983. Altruism, jealousy, and the theory of optimal non-linear taxation. Journal of Public Economics 20, 77–87.

Perez-Truglia, R, 2020. The effects of income transparency on well-being: Evidence from a natural experiment. American Economic Review 110, 1019–1054.

Powdthavee, N, 2009. How important is rank to individual perception of economic standing? a within-community analysis. Journal of Economic Inequality 7, 225–248.

Powdthavee, N, Burkhauser, R.V, Neve, J.E.D, 2017. Top incomes and human well-being: Evidence from the gallup world poll. Journal of Economic Psychology 62, 246–257.

Reardon, S.F, 2011. Measures of income segregation. CEPA Working Papers .

Schneider, S.M, 2016. Income inequality and subjective wellbeing: Trends, challenges, and research directions. Journal of Happiness Studies 17, 1719–1739.

Zizzo, D.J, Oswald, A.J, 2001. Are people willing to pay to reduce others' incomes? Annales d'Économie et de Statistique 63/64, 39–65.

Figures

Figure 1: Area unit boundaries in central Wellington

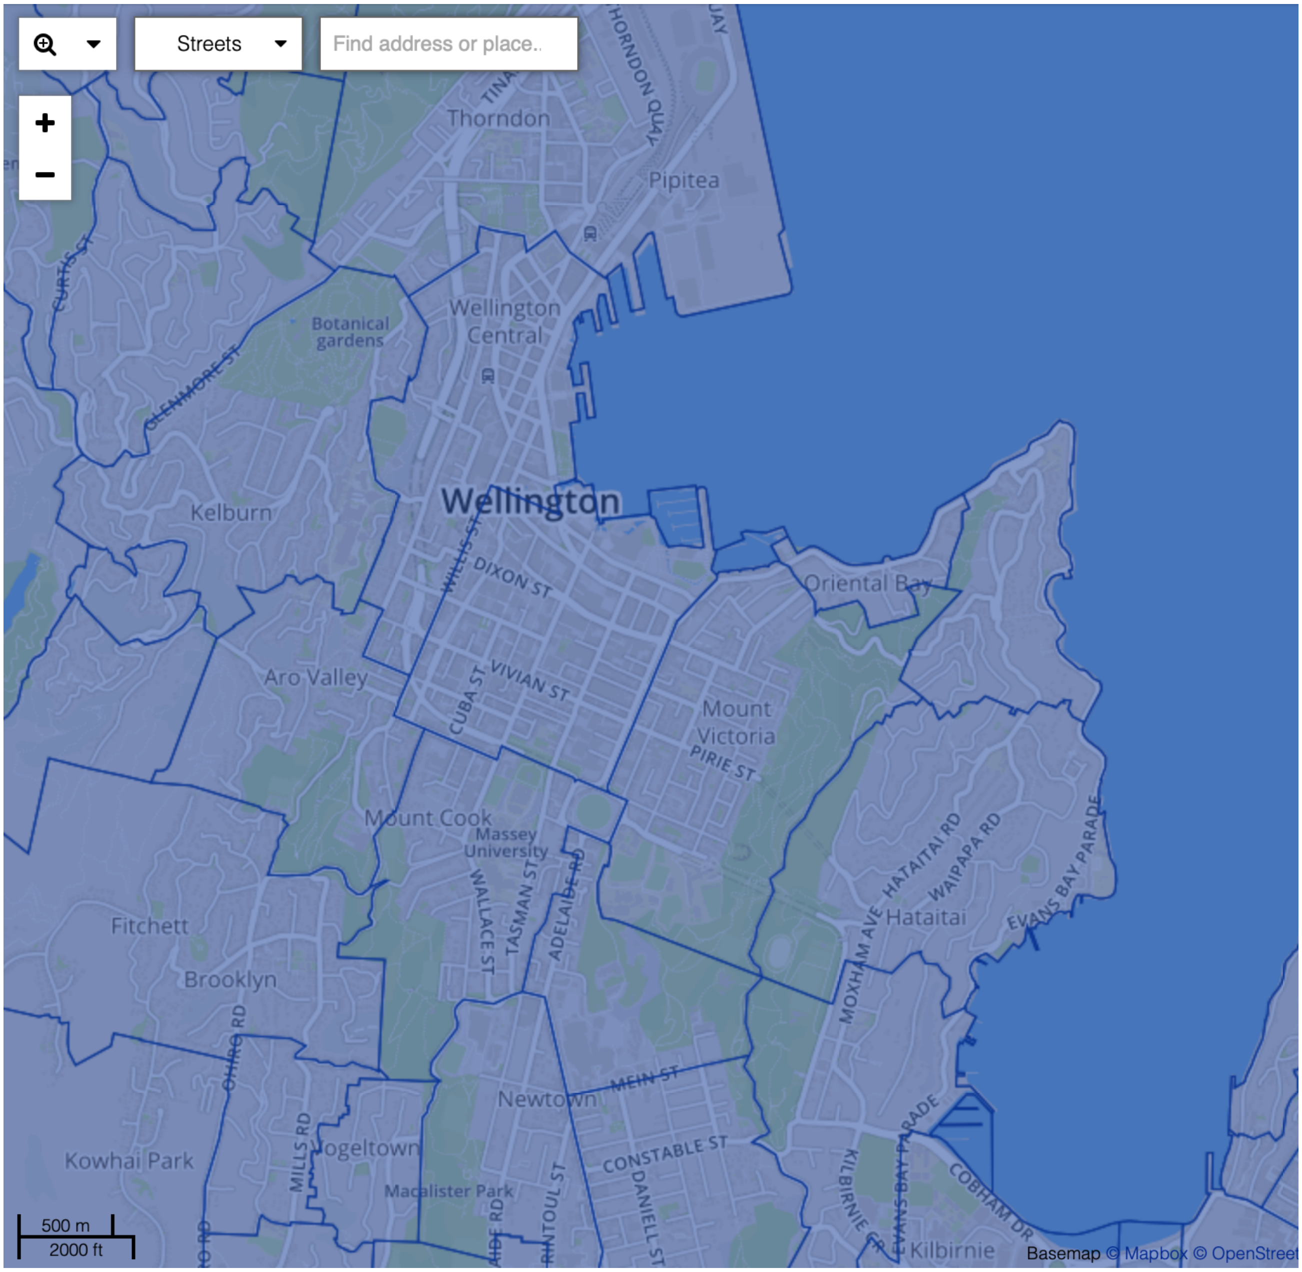

Notes: this is a screenshot from the Statistics New Zealand datafinder, accessible here `https://datafinder.stats.govt.nz/layer/27771-area-unit-2017-generalised-version/`.

Figure 2: Distributions of subjective wellbeing

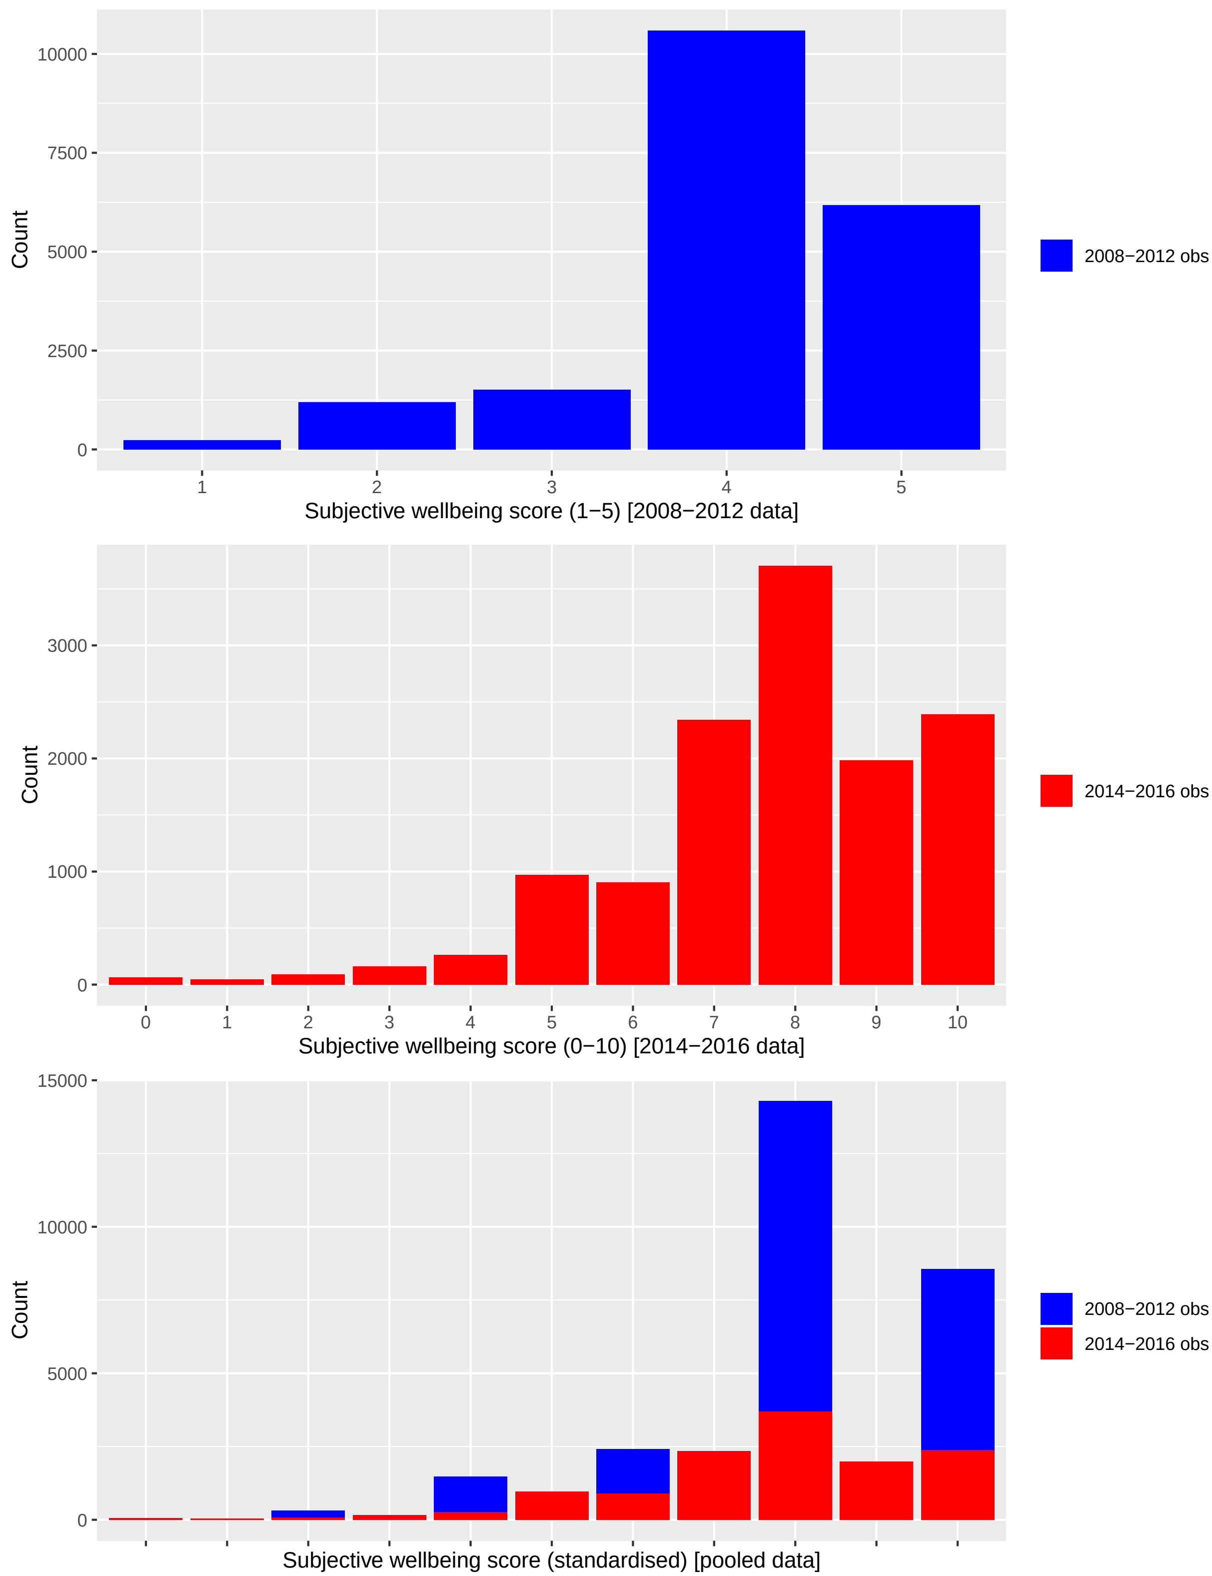

Figure 3: Neighbourhood peer income coefficients

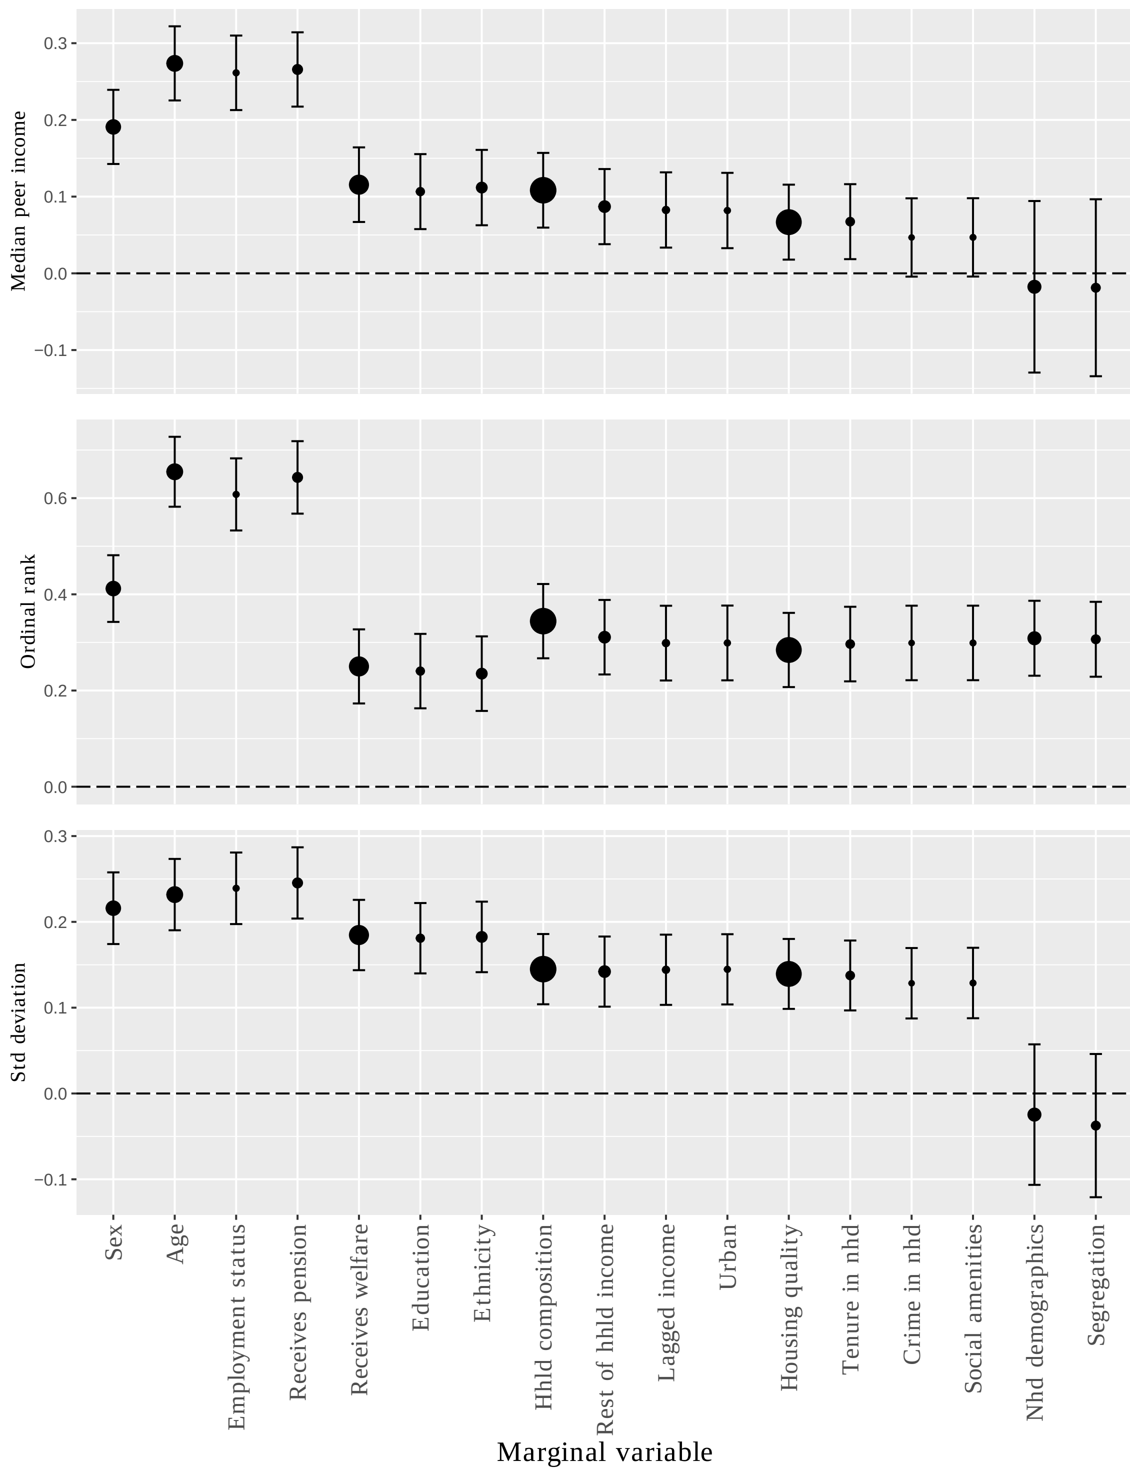

Notes: this figure displays coefficients from linear regressions of subjective wellbeing on our neighbourhood peer income variables and controls. The base model is a regression of subjective wellbeing on the peer income variables and year fixed effects, and each column adds an additional control into 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 wage coefficients

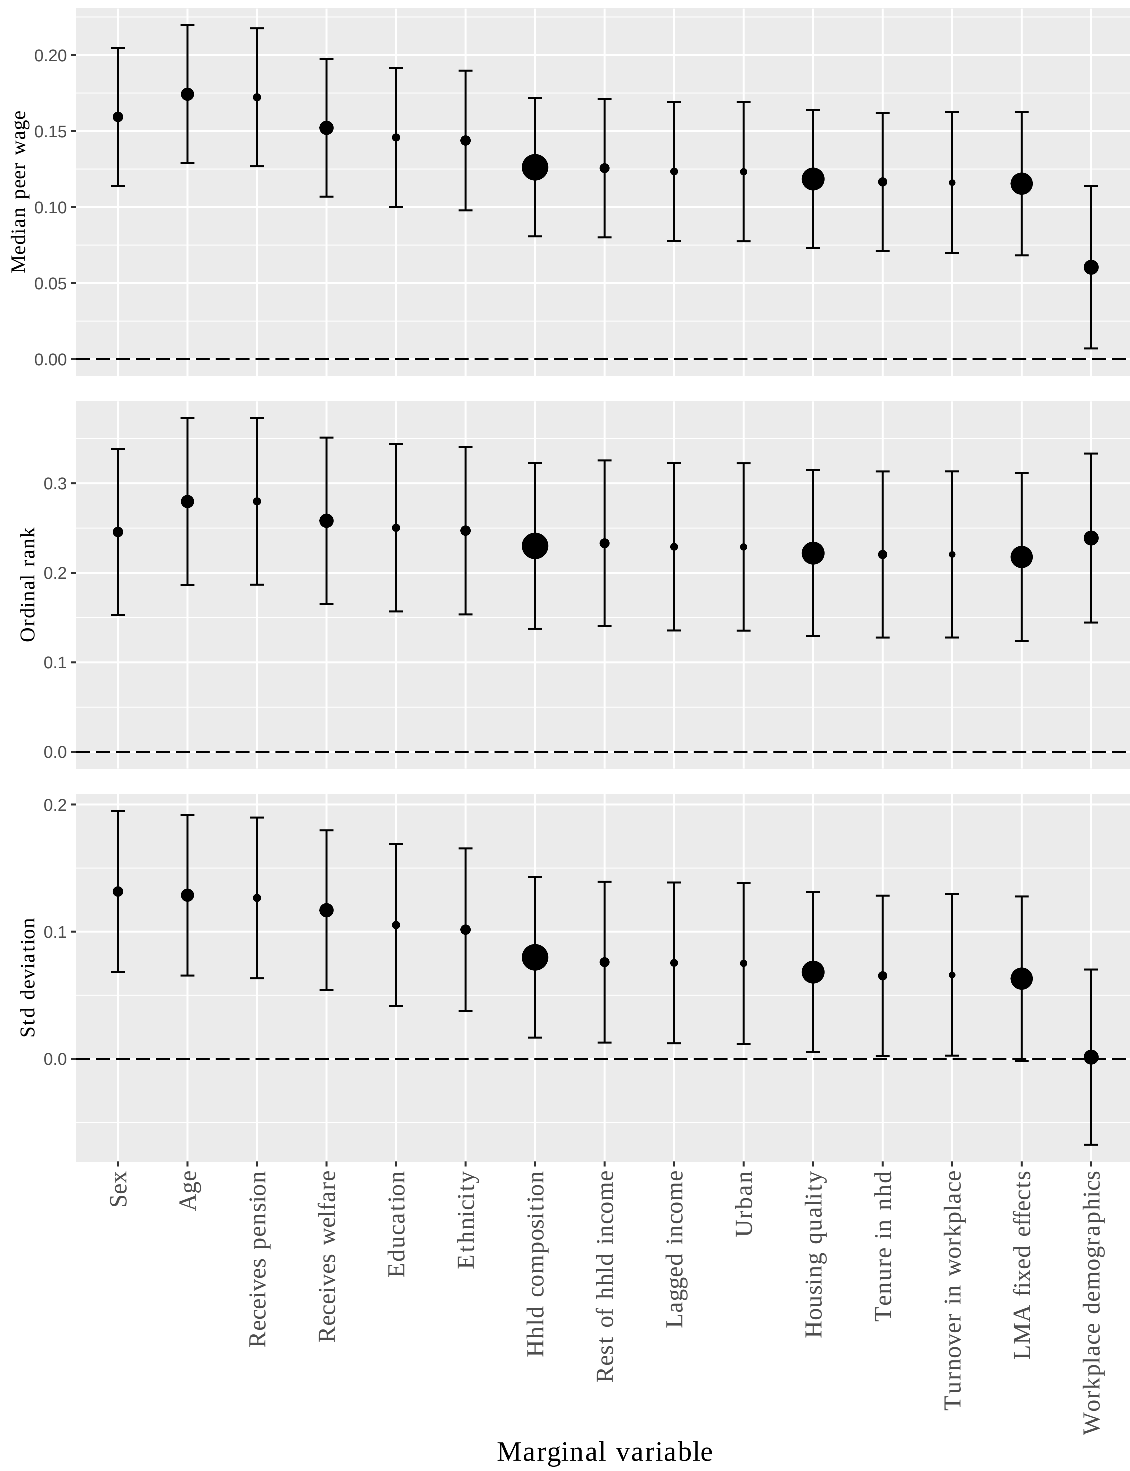

Notes: this figure displays peer wage coefficients from linear regressions of subjective wellbeing on our workplace peer wage variables and controls. The base model is a regression of subjective wellbeing on the peer wage variables and year fixed effects, and each column adds an additional control into 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

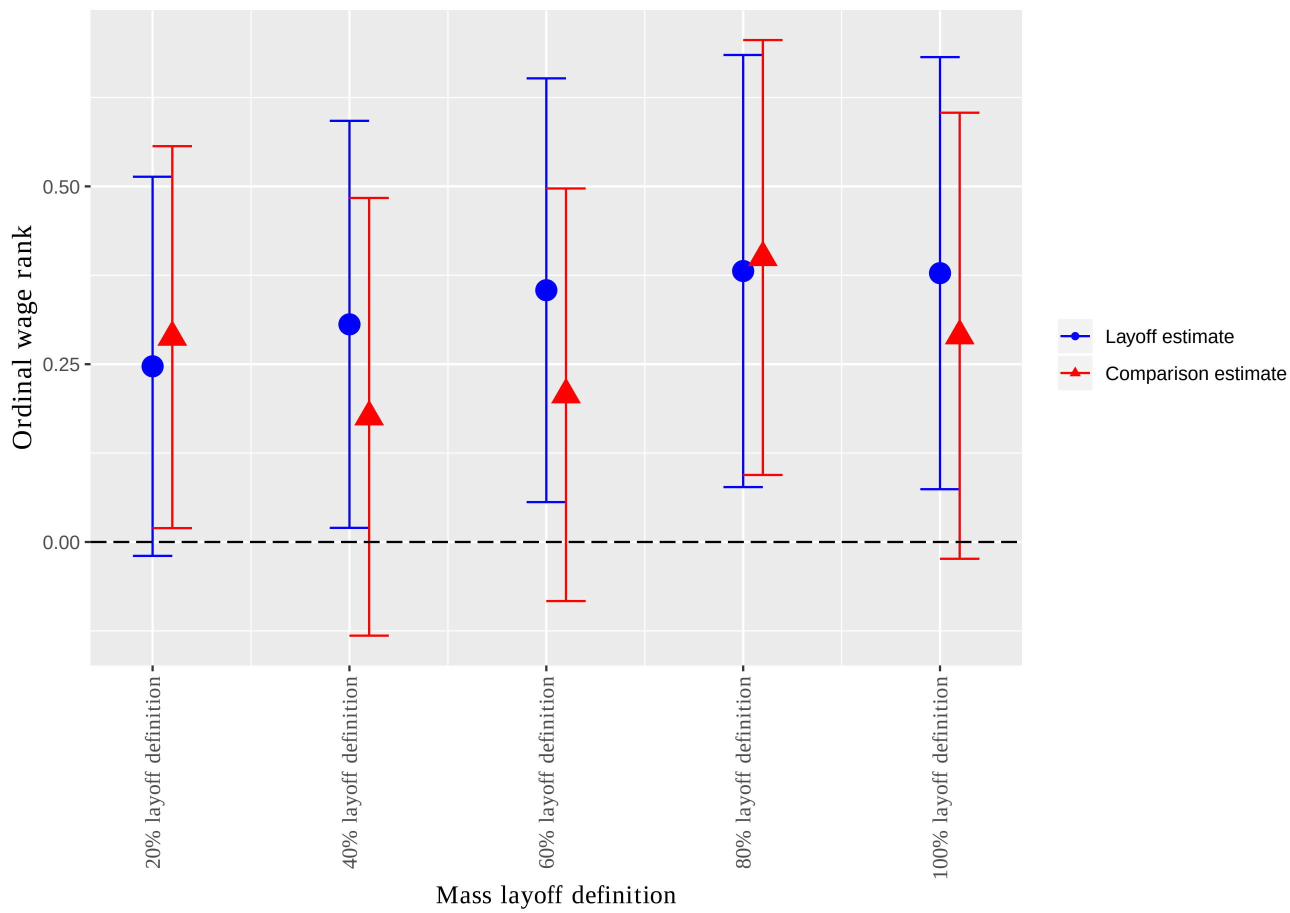

Notes: this figure displays the coefficient on ordinal wage rank from linear regressions of subjective wellbeing on workplace peer wage variables. Blue/circular dots represent estimates where the sample is restricted to people who lost their previous job as part of a mass layoff. The horizontal axis 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

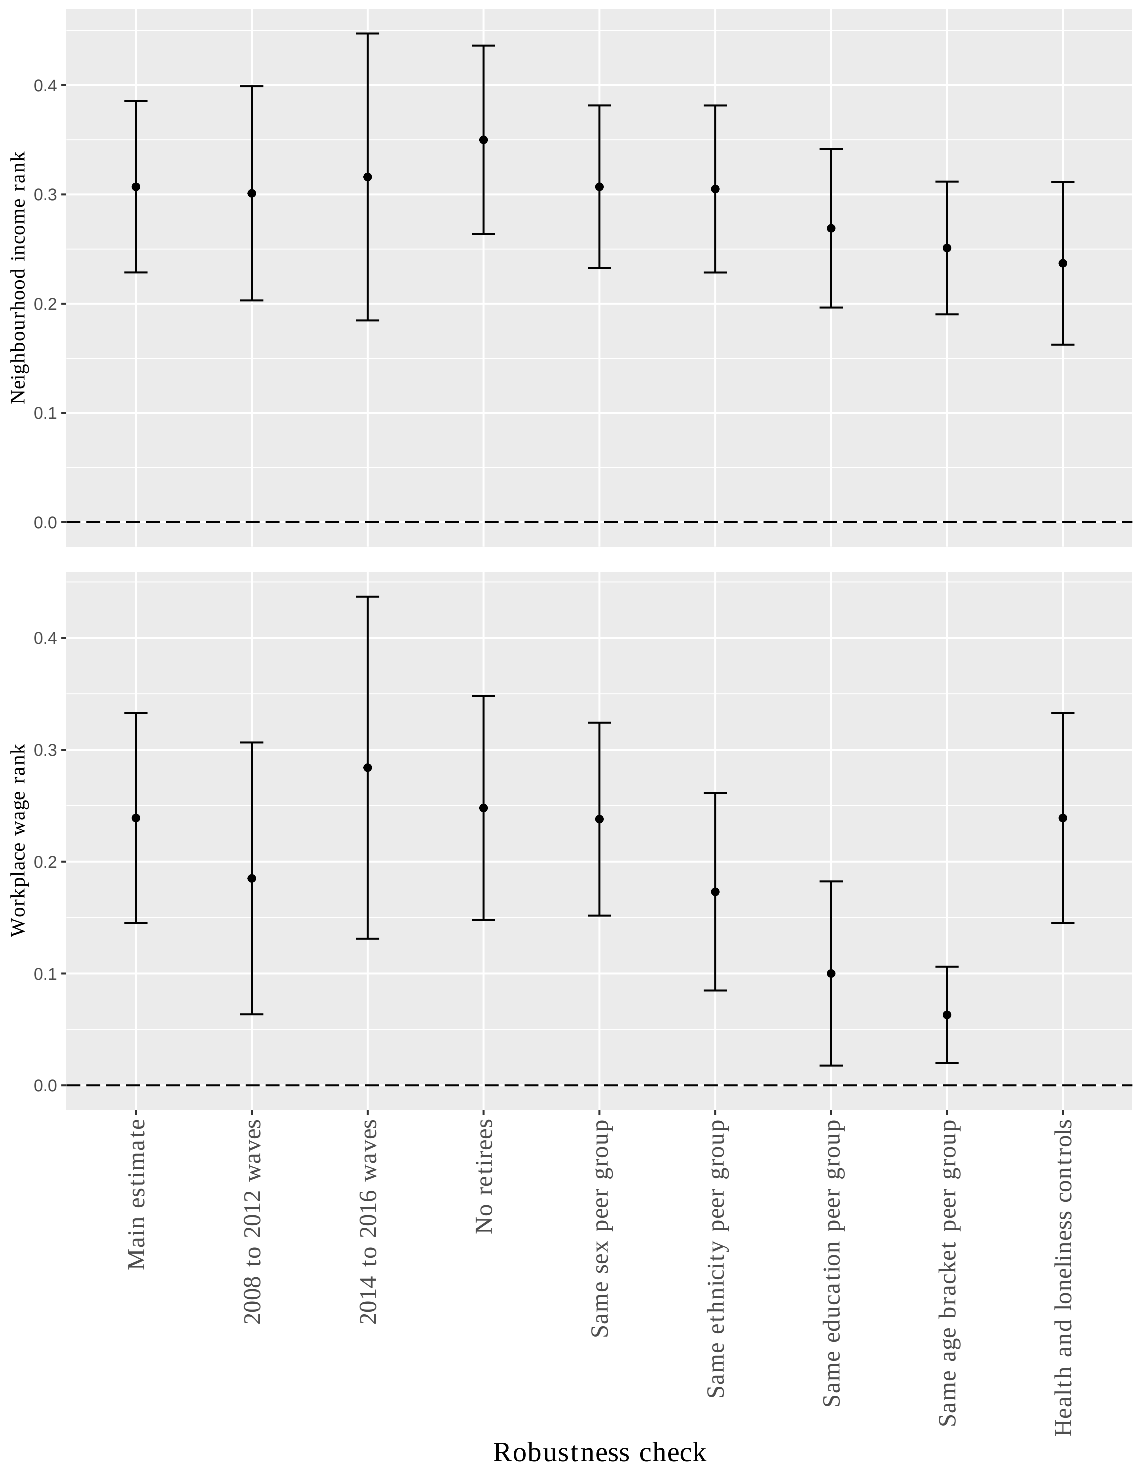

Notes: this figure displays the income rank coefficient from regressions of subjective wellbeing on neighbourhood peer income variables and a full set of controls, and the wage rank coefficient from regressions of subjective wellbeing on workplace peer wage variables and a full set of controls. Each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

Tables

Table 1: Descriptive statistics

Variable	Full sample		Employed sample	
	Mean	S.D.	Mean	S.D.
Demographics				
Male dummy	0.46	–	0.47	–
Age	50.71	(18.51)	44.20	(13.85)
Total income (2019 NZD)	44,566	(44,463)	61,799	(48,290)
Employed dummy	0.43	–	–	–
Receives welfare benefits	0.14	–	0.07	–
Highest educational qualification				
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	–
Ethnicity				
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	–
Other individual characteristics				
Lives in urban area	0.98	–	0.95	–
Is "never lonely"	0.64	–	0.64	–
Never hindered by physical pain	0.55	–	0.62	–
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)	–	–	5,085	(5,293)
Neighbourhood characteristics				
Average monthly # of crimes	105.07	(111.07)	106.17	(116.11)
Workplace characteristics				
Number of coworkers	–	–	341.28	(838.47)
Monthly turnover (% of employees)	–	–	0.10	–
# of observations	32,643		13,917	

Table 2: Correlation matrices of peer income variables

Neighbourhood peer income variables					
	Median	Ordinal rank	Std dev	Top decile	Bottom decile
Median	1				
Ordinal rank	-0.011	1			
Std dev	0.247	-0.005	1		
Top decile	0.730	0.003	0.422	1	
Bottom decile	-0.012	0.002	-0.951	-0.157	1
Workplace peer wage variables					
	Median	Ordinal rank	Std dev	Top decile	Bottom decile
Median	1				
Ordinal rank	-0.069	1			
Std dev	-0.595	0.049	1		
Top decile	0.480	-0.017	-0.085	1	
Bottom decile	0.844	-0.046	-0.892	0.307	1

Notes: this table displays the unconditional correlation matrices of our two sets of peer income variables.

Table 3: 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, and (6) missing qualifications info.
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 (MELAA) 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.).
Income of rest of hhld	Self-reported total household income (from the GSS), minus own income.
Past income	Annual income from 5 years ago.
Receives pension	Dummy for receiving nonzero pension or superannuation income.
Receives welfare	Dummy for receiving nonzero income from welfare benefits (excluding parental leave and superannuation).
Physical health	Self-reported 1-5 point scale about whether the person is limited in her activities by physical pain (from the GSS).
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 5.
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 about 100 inhabitants on average.
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, percent of people in area unit receiving pension income, 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 4: Test regressions

Variable	Neighbourhood Coef	Neighbourhood S.E.	Workplace Coef	Workplace S.E.
Income (log)	0.061***	(0.008)	0.066***	(0.015)
Rest of household's income (log)	0.047***	(0.006)	0.035***	(0.010)
Income 5 years ago (log)	0.010*	(0.005)	0.010	(0.011)
Male dummy	-0.107***	(0.011)	-0.062***	(0.016)
Age	-0.019***	(0.002)	-0.032***	(0.004)
Age squared/100	0.020***	(0.002)	0.037***	(0.005)
Employed dummy	0.022*	(0.012)	-	-
Receives welfare benefits	-0.305***	(0.021)	-0.200***	(0.038)
Receives pension	0.170***	(0.023)	0.053	(0.043)
Education (omitted = no qualifications)				
Missing qualifications info	-0.016	(0.018)	0.033	(0.031)
School qualification	0.016	(0.016)	0.014	(0.024)
Postschool qualification	0.012	(0.018)	0.017	(0.028)
Undergraduate qualification	0.049**	(0.021)	0.020	(0.031)
Postgraduate qualification	0.078***	(0.025)	0.038	(0.036)
Ethnicity (omitted = sole European)				
Sole Māori	0.064**	(0.025)	0.040	(0.036)
Māori and European	-0.017	(0.025)	-0.004	(0.033)
Pacific or Pacific-European	0.092***	(0.033)	-0.036	(0.043)
Asian or Asian-European	-0.006	(0.023)	-0.010	(0.031)
MELAA or MELAA-European	-0.044	(0.050)	-0.055	(0.070)
Other ethnicity	0.098**	(0.045)	0.082	(0.065)
Lives in urban area	0.016	(0.034)	-0.037	(0.040)
Is "never lonely"	0.416***	(0.012)	0.411***	(0.017)
Physical pain (omitted = "not at all")				
Pain interfered "a little" with normal work	-0.115***	(0.012)	-0.127***	(0.018)
Pain "moderately" interfered with normal work	-0.227***	(0.020)	-0.253***	(0.033)
Pain interfered "quite a bit" with normal work	-0.354***	(0.022)	-0.327***	(0.041)
Pain interfered "extremely" with normal work	-0.345***	(0.034)	-0.099**	(0.047)
Average monthly # of crimes committed in nhd	-0.000	(0.000)	-	-
Social amenities (higher = more amenities)	-0.000	(0.000)	-	-
Moved into nhd this year	-0.050**	(0.021)	-	-
Housing quality factor variable	0.177***	(0.012)	-	-
Started at workplace this month	-	-	0.116**	(0.046)
$\leq$ 2 months spent at this workplace	-	-	-0.215**	(0.089)
Turnover at this workplace	-	-	-0.024	(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	0.157		0.126	
# of observations	32,643		13,917	

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 5: Neighbourhood peer income regressions

Variable	(1)	(2)	(3)	(4)	(5)
Income (log)	0.068***	0.020**	0.022**	0.017*	0.019*
	(0.008)	(0.010)	(0.010)	(0.010)	(0.010)
Median peer income	0.051**	0.110***	0.067***		-0.019
	(0.023)	(0.024)	(0.025)		(0.059)
Ordinal income rank		0.308***	0.297***	0.315***	0.307***
		(0.040)	(0.040)	(0.039)	(0.040)
Std dev of peer income			0.138***	0.042	-0.037
			(0.021)	(0.143)	(0.043)
Top decile of peer income				0.136***	
				(0.042)	
Bottom decile of peer income				-0.018	
				(0.045)	
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	0.098	0.100	0.101	0.102	0.104
# 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 3. Column 5 adds the full set of neighbourhood-level characteristics described in Table 3. 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 wage regressions

Variable	(1)	(2)	(3)	(4)	(5)
Wage (log)	0.047***	-0.015	-0.018	-0.006	-0.019
	(0.012)	(0.018)	(0.018)	(0.016)	(0.018)
Median peer wage	0.027*	0.093***	0.114***		0.060**
	(0.015)	(0.021)	(0.023)		(0.027)
Ordinal wage rank		0.213***	0.220***	0.185***	0.239***
		(0.048)	(0.048)	(0.043)	(0.048)
Std dev of peer wage			0.062*	0.302***	0.001
			(0.032)	(0.069)	(0.035)
Top decile of peer wage				-0.008	
				(0.01)	
Bottom decile of peer wage				0.112***	
				(0.022)	
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	0.065	0.067	0.067	0.067	0.082
# of observations	13,917	13,917	13,917	13,917	13,917

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 3. Column 5 adds the full set of workplace-level characteristics described in Table 3. 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)	0.019*		0.025
	(0.010)		(0.024)
Median peer income	-0.019		-0.100
	(0.059)		(0.094)
Ordinal income rank	0.307***		0.100
	(0.040)		(0.089)
Std dev of peer income	-0.037		0.003
	(0.043)		(0.068)
Wage (log)		-0.019	-0.034*
		(0.018)	(0.019)
Median peer wage		0.060**	0.044
		(0.027)	(0.028)
Ordinal wage rank		0.239***	0.190***
		(0.048)	(0.052)
Std dev of peer wage		0.001	-0.007
		(0.035)	(0.035)
Full set of controls	Yes	Yes	Yes
R^2	0.104	0.081	0.086
# of observations	32,643	13,917	13,917

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 Column 5 of Tables 5 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 OLS	Neighbourhood IV	Workplace OLS	Workplace IV
Income/wage	0.026*	0.042**	0.034	-0.067
	(0.013)	(0.021)	(0.046)	(0.129)
Median peer income/wage	0.022	0.291	0.057	0.239
	(0.080)	(0.424)	(0.060)	(0.386)
Ordinal income/wage rank	0.260***	0.131	0.156	0.384*
	(0.054)	(0.125)	(0.103)	(0.216)
Std dev of peer income/wage	-0.085	-0.163	0.123*	0.328*
	(0.058)	(0.203)	(0.070)	(0.190)
Year fixed effects	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes
Group-level controls	Yes	Yes	Yes	Yes
R^2	0.102	0.101	0.099	0.096
# of observations	18,126	18,126	3,774	3,774

Notes: coefficients from OLS and 2SLS regressions of subjective wellbeing on neighbourhood peer income variables and workplace peer wage 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 variables with the change in those variables over the last 5 years. Average (minimum) first-stage F-statistics are 259 (143) for the neighbourhood IV regression and 130 (21) 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 tests

Variable	Neighbourhood		Workplace	
	No placebo	Placebo	No Placebo	Placebo
Income/wage	0.006	0.006	-0.047	-0.048
	(0.019)	(0.019)	(0.034)	(0.034)
Median peer inc/wage	-0.058	-0.077	0.059	0.016
	(0.107)	(0.152)	(0.053)	(0.064)
Ordinal inc/wage rank	0.419***	0.402***	0.201**	0.198**
	(0.074)	(0.077)	(0.086)	(0.090)
Std dev of peer inc/wage	0.045	-0.058	-0.015	-0.093
	(0.076)	(0.098)	(0.069)	(0.074)
Median peer inc/wage (5-year lag)		0.036		0.083
		(0.121)		(0.052)
Ordinal inc/wage rank (5-year lag)		-0.035		-0.026
		(0.044)		(0.060)
Std dev of peer inc/wage (5-year lag)		0.141*		0.163**
		(0.079)		(0.068)
Year fixed effects	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes
Group-level controls	Yes	Yes	Yes	Yes
R^2	0.116	0.116	0.099	0.101
# of observations	10,437	10,437	4,476	4,476

Notes: Columns 1 and 3 display coefficients from an OLS regression of subjective wellbeing on our peer income/wage variables and a full set of controls (the same specifications as Columns 1 and 2 of Table 7, except restricted to people who moved into their current neighbourhood sometime in the past 5 years). Columns 2 and 4 displays coefficients from the same regression, with the addition of 5-year lags of the peer income/wage 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 administrative data 5 years ago. Standard errors are robust. Asterisks denote significance at *p<0.10, **p<0.05, ***p<0.01.

Appendix

In this Appendix, we present four brief extensions on our baseline results.

First, we test whether the effects of median peer income vary according to a person's position in her peer income distribution. A few existing studies (Ferrer-i-Carbonell (2005), Card et al. (2012)) find that peer income has a larger effect on the subjective wellbeing of people who are below the median within their peer income distribution (that is, income comparisons matter more to people who are relatively deprived). In particular, Card et al. (2012) find that their treatment has a small and statistically insignificant effect on average, but when they interact their treatment dummy with a "below the median" dummy, they find strong and statistically significant effects for individuals below the median.

In order to test for this possibility, we rerun our Table 7 regressions with the addition of a dummy for being below the median peer income and an interaction between this dummy and our median peer income variable. The results are reported in the top panel of Appendix Table A3. They show that there is no statistically significant difference in the effects of median peer income for people above or below the median in their peer group.

Next, we test whether the effects of peer income vary with the size of a person's peer group. It is, for example, possible that income comparisons matter more in smaller peer groups, where status concerns are more salient or individuals have better knowledge of their peers' income. To test whether this is the case, we run our Table 3 regressions with the addition of an interaction between the ordinal rank variable and the size of a person's peer group (a main effect for "size of peer group" is already included in our standard set of controls). The results are reported in Appendix Table A3; we find no evidence that the effects of income comparisons change with the size of a person's peer group. The coefficients on the interaction terms, in addition to being statistically insignificant, are very small – they suggest that in order to reach an average treatment effect of zero, our peer groups would have to be 40-50 times larger than their current average sizes.

We next test for the effects on subjective wellbeing of the very top incomes in a neighbourhood or workplace. Godechot and Senik (2015) and Powdthavee et al. (2017) find that people's subjective wellbeing is sensitive to the incomes of the top 1% in their peer groups.

In order to check for this possibility, we rerun our main regressions using the average income in the top 1% of a peer group rather than the average income in the top 10%. The results, which are reported in Appendix Table A3, are qualitatively identical to our main results: the top 1% of peer income still has a small and statistically insignificant relationship with subjective wellbeing.

Finally, we rerun our Table 5 and 6 regressions using non-logged measures of median peer income/wages, standard deviation of peer income/wages, and top and bottom decile of peer income/wages. Our preferred specifications use logged measures of peer income for ease of interpretation and because we believe that percentage differences in income (captured by the log specification) are likely to be more relevant than dollar differences in income (captured by linear specifications). Appendix Table A3, which uses non-logged measures of our cardinal peer income variables, shows that this choice of functional form has no qualitative effect on our results.

Appendix Tables

Table A1: Neighbourhood peer income regressions

Variable	(1)	(2)	(3)	(4)
Income (log)	0.020**	0.019*	0.019*	0.019*
	(0.010)	(0.010)	(0.010)	(0.010)
Median peer income	-0.006			
	(0.057)			
Ordinal income rank	0.305***	0.308***	0.308***	0.308***
	(0.040)	(0.040)	(0.040)	(0.040)
Std dev of peer income		-0.034		
		(0.041)		
Top decile of peer income			0.040	
			(0.043)	
Bottom decile of peer income				0.009
				(0.012)
Year fixed effects	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes
Group-level controls	Yes	Yes	Yes	Yes
R^2	0.104	0.104	0.104	0.104
# of observations	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 a full set of covariates, for our full sample. Unlike Table 7, here the cardinal peer income variables are included one-by-one rather than all together. 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 A2: Workplace peer wage regressions

Variable	(1)	(2)	(3)	(4)
Wage (log)	-0.019	0.005	0.010	-0.007
	(0.018)	(0.015)	(0.015)	(0.016)
Median peer wage	0.060**			
	(0.025)			
Ordinal wage rank	0.239***	0.180***	0.168***	0.209***
	(0.048)	(0.041)	(0.041)	(0.043)
Std dev of peer wage		-0.032		
		(0.032)		
Top decile of peer wage			-0.003	
			(0.010)	
Bottom decile of peer wage				0.023**
				(0.010)
Year fixed effects	Yes	Yes	Yes	Yes
Individual-level controls	Yes	Yes	Yes	Yes
Group-level controls	Yes	Yes	Yes	Yes
R^2	0.081	0.081	0.081	0.081
# of observations	13,917	13,917	13,917	13,917

Notes: this table displays coefficients from linear regressions of subjective wellbeing on workplace peer wage variables and a full set of covariates, for our employed sample. Unlike Table 7, here the cardinal peer wage variables are included one-by-one rather than all together. 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 A3: Extensions, part 1

Specification	Neighbourhood Coef	Neighbourhood S.E.	Workplace Coef	Workplace S.E.
Interaction with "below the median"				
Median peer income/wage	-0.031	(0.073)	0.061**	(0.028)
Median peer income/wage * below the median	0.027	(0.044)	-0.004	(0.028)
Below the median dummy	-0.270	(0.454)	0.047	(0.231)
Full set of controls and peer income variables	Yes		Yes	
R^2	0.104		0.081	
# of observations	32,643		13,917	
Interaction with "size of peer group"				
Ordinal income/wage rank	0.306***	(0.040)	0.245***	(0.049)
Ordinal income rank/wage * size of peer group	-0.006	(0.011)	-0.020	(0.031)
Size of peer group/1000	0.056	(0.108)	0.015	(0.020)
Full set of controls and peer income variables	Yes		Yes	
R^2	0.104		0.081	
# of observations	32,643		13,917	
Top 1% instead of top decile				
Top 1% of peer income/wages	0.017	(0.012)	0.002	(0.002)
Full set of controls and peer income variables	Yes		Yes	
R^2	0.104		0.082	
# of observations	32,643		13,917	
Non-logged cardinal peer variables				
Median peer income/wages (non-logged/10,000)	-0.000	(0.002)	0.011	(0.007)
Std dev of peer income/wages (non-logged/10,000)	-0.000	(0.000)	-0.003	(0.005)
Full set of controls and peer income variables	Yes		Yes	
R^2	0.104		0.082	
# of observations	32,643		13,917	

Notes: this table displays coefficients from four alternative specifications of the regressions from Columns 1 and 2 of Table 7. "Full set of peer income variables" means "all neighbourhood peer income variables" in the neighbourhood case and "all peer wage variables" in the workplace specifications. The first specification adds a dummy for being "below the median income" in a peer group, and an interaction between that dummy and median peer income. The second specification adds an interaction between the "size of peer group" control and our ordinal rank variable. The third specification exchanges our "top decile of peer income" variable for a "top 1% of peer income" variable. The fourth specification exchanges our logged measures of the cardinal peer income variables with non-logged measures. None of the non-displayed coefficients change qualitatively under these alternative specifications. Standard errors are robust. Asterisks denote significance at *p<0.10, **p<0.05, ***p<0.01.