150 to 180 public housing residents in Brisbane. They found mixed responses regarding employment outcomes after going into public housing.

There is relatively little evidence on other outcomes such as health or education. Phibbs and Young (2005) find survey respondents reported improvements in health and the education of their children. Wood et al (2016) examined the effects of entering public housing in WA on health service usage of those formerly homeless. They find significant positive outcomes following entering social housing for health service usage (in that usage of health services declines once in social housing) for those formerly homeless. Although this work has an important advantage of measuring a health related outcome, it has three limitations. The most important of these is that there is no counterfactual in the analysis so we don't know what the health service usage was for those that didn't manage to enter social housing. In addition, health service use may not always be a good proxy for health. In some cases, people moving out of homelessness may be better off with more extensive treatment for certain health conditions like mental illness. Finally, social housing is not solely for the formerly homeless which limits the extent to which the findings can be generalised to other social housing residents.

There has been one study that looks at the effect of public housing on homelessness. Johnson, Scutella, Tseng and Wood (2017), who apply a different approach and focus on modelling the transition in and out of homelessness, confirm the protective effect of social housing. Those in public housing, in particular, are substantially less likely to enter homelessness than similarly vulnerable people in the private rental market.

We will also briefly cover the international literature. This is not only useful for comparisons but also, potentially, to suggest how impacts might be different if social housing in Australia was provided on a larger scale as suggested by, for example, Infrastructure Victoria (2016a). Fitzpatrick and Stephens (2007) and Scanlon et al (2015) distinguish between roughly two groups of countries in terms of their provision of social housing. One group, whom Scanlon et al (2015) term universalist, provide social housing on a substantial scale (between 20 to 30 per cent of the housing stock) for a wide variety of households. The other group, whom Scanlon et al (2015) refer to as dualist, provide social housing, on a much smaller scale, specifically for low-income households. This latter group includes Australia and the United States.

Unfortunately, there are no broad studies of the impact of social housing for universalist systems. Most studies that have been done are on the United States. US housing policies, including social housing, have most recently been reviewed in Olsen and Zabel (2015). This review concludes that housing assistance, including social housing, tends to reduce employment participation. There have been relatively few studies on the impacts on health which have mainly found no significant or substantial effects. There doesn't appear to have been a study of the effect of social housing specifically on homelessness.

Another set of useful evidence has come from a set of programs that provide housing (not necessarily social housing), accompanied by more intensive support services, for the chronically homeless – referred to as "Housing First" (Kertesz and Johnson, 2017). Kertesz and Johnson's review reports, subject to some concerns about how the programs were carried out and evaluated, that this approach appears to reduce the extent of homelessness but that it has been hard to find statistically significant positive effects on physical or mental health. They suggest the latter may be because either that benefits may take longer to arrive than the typical evaluation period (1-2yrs) or that the sample includes significant numbers of people whose condition is unlikely to improve, even with housing.

2. Econometric framework and data

In this section, we introduce the econometric framework used to analyse these questions. This is followed by an introduction of the two datasets we use to estimate the impacts and analyse their robustness. We conclude with a discussion of the threats to identification when estimating the impacts and how we attempt to deal with them.

3.1 Econometric framework

We adopt a quasi-experimental approach to estimating the impact of social housing. The first step is to characterise living in social housing as a treatment. The next step is to construct a control group for those observed in social housing. Outcomes are then compared so to estimate the average treatment effect on the treated (ATET) as follows:

$$ATET_{SH}=E[(Y_{SHi}-Y_{0i})|T_i=SH]= E[(Y_{SHi})|T_i=SH] - E[(Y_{0i})|T_i=SH]$$
$$= E[(Y_{SHi})|X=x, T_i=SH] - E[(Y_{0i})| X=x, T_i=0]$$

where: Y_{Shi} refers to the outcome of each i social housing resident while in social housing, and Y_{0i} provides the counterfactual, i.e. the outcome of each social housing resident *i* if they hadn't been subject to the treatment. Since the counterfactual outcome for each individual is not observed, it is not possible to calculate $E[(Y_{0i})|T_i=SH]$. This is replaced by the average outcome of the matched non-treated (control) group $E[(Y_{0i})| X=x, T_i=0]$. Then we can estimate the ATET.

In addition to comparing residents and non-residents of social housing, we also compare the ATET for entrants into social housing with matched individuals who do not enter social housing. As discussed in more detail below, this is to address the threat to identification by including in the treatment group individuals selected into social housing under different selection processes (or with a different time to be affected by living in social housing) such that they have different unobservable characteristics.

Note though that the interpretation of the results from these two approaches is subtly different. When comparing residents with non-residents, the differences will be in the probability of achieving each outcome relative to that in the control group. When comparing entrants with non-residents, the way to interpret these values is that that the estimates represent a net gain in the outcome in question for entrants to social housing relative to the matched control group. A negative treatment effect therefore represents a net loss in the outcome relative to the control group; that is that either the improvement in the outcome was smaller than the average improvement for the control group or that there was a larger deterioration in the outcome in question for entrants than there was for the control group on average.

3.2 Data Sources

There is no survey that is designed to focus on social housing residents accompanied with a matching control group. Instead we draw on two surveys which include, to differing degrees, individuals who are in social housing and those who are potentially eligible for living in social housing.

The primary data source for this paper is Journeys Home. Journeys Home (hereafter JH) is a nationally representative longitudinal survey of a sample of particularly disadvantaged adults. The sample is drawn from a population of Centrelink customers that have been identified as vulnerable to being homeless. There are six waves between 2011 and 2014. For further detail of the JH survey see Wooden et al (2012) and Scutella, Tseng and Wooden (2017).

The individuals represented in the JH survey are likely to be more similar to individuals who recently entered social housing, due to their particular vulnerability.

However they may not necessarily be representative of most social housing residents per se. In particular, it does not include any children or youth under the age of 18 years old.

Therefore, to check the robustness of our results, we also use the HILDA dataset. HILDA is a nationally representative household longitudinal survey of individuals who are at least fifteen years old. The survey features fifteen waves between 2001 and 2015. Although the survey includes a wide range of questions relevant to our research questions, not all questions are asked in all waves. For further detail on the HILDA survey see Watson and Wooden (2012). The individuals in the HILDA survey are drawn from the general public so, they are less likely to be similar in terms of their unobservable characteristics in terms of their vulnerability to those entering and living in social housing. HILDA also does not include any children or youth under the age of 18.

3.2.1 Outcomes, treatment and control groups

The treatment analysed in this paper is residing in social housing, as defined in Table 2. If a person lives in either public or community housing, then they are recorded as living in social housing.

The outcomes we consider are also defined in Table 2. The outcomes fall into four groups consistent with the approach taken in Infrastructure Victoria (2016b). These are employment, education, physical and mental health. In addition, we take advantage of the focus of JH on people at risk of homelessness to consider two additional outcomes: homelessness and incarceration.

The outcomes, as measured, differ to the extent that they can quickly respond to treatment. Employment and some of the health measures can potentially change relatively quickly as a response to entering social housing. The level of education will be affected by decisions made in the past, though it can improve.

For the main Journeys Home analysis, the control group is selected from all surveyed individuals not in social housing at the time. No further restrictions are undertaken because, to be surveyed in Journeys Home, the individual must have been determined to be at least vulnerable to becoming homeless. As being homeless is one of the factors (but not the only factor) that can result in an individual obtaining priority access to housing, this makes it more likely that individuals in Journeys Home have the required characteristics that would make them, at least in principle, eligible for social housing.

The specific control group is constructed using a Nearest Neighbour method. Specifically, each individual in the treatment group is matched to its nearest neighbour by minimising a weighted function of the differences between selected covariates using the Mahalanobis Distance Method(MDM).[5] These covariates are listed in Table 3.

The first step of our robustness analysis is to construct an alternative control group using Propensity Score Matching (PSM).[6] There are three stages in propensity score matching:

1. Estimate the probability of being observed in the treatment group by using a probit. This is referred to as the propensity score. The set of variables we use in the probit are the same covariates that were used previously as listed in Table 3.
2. Use a rule to construct matching observations for the treatment observations, subject to the treatment and controls having a common support. This is done with replacement. Because JH (and HILDA) are longitudinal surveys we also need to confine matching within waves to avoid potential issues with unobservable

[5] To do this we use the Stata teeffects nnmatch command.
[6] To do this we use the user written Psmatch2 command with kernel estimator, and bootstrap the standard errors.

unobservable characteristics upon entering. This may be in part due to accessing services that they have not been able to do so while, for example, homeless. Though in Australia it is generally not the case in public housing, and to some extent in community housing, that there is strict and exclusive tying of housing and services. This means individuals in the control group with similar needs to those in social housing may also be accessing similar services. But, again, without observing how long an individual has been in social housing, we cannot directly control for this.

We adapt our methodology in two ways to deal with these threats to identification. First, to allow for the time it takes for benefits from social housing to occur, we consider the value of the dependent variable, i.e. the outcome, at t+1 whereas the treatment and any controls are at t.[7] So we allow for a six month to year-long lagged effect. Limitations on the survey length for Journeys Home and the size of the treatment group prevent a more flexible analysis of the timing of effects.

Second, to reduce potential problems of different unobservable characteristics associated with different cohorts and different treatment effects from being in social housing for different periods of time we also analyse differences in changes in outcomes. Specifically, we consider changes in outcomes associated with changes in treatment status i.e. the changes in outcomes following moving into social housing. Note we are no longer able to compare outcomes for homelessness as the treated individuals, by definition, cannot be homeless. There is a further limitation on these results in that we can only measure short run impacts. This is due to the combination of the observation periods being only six to twelve months and our not being able to observe how long the individual has been in social housing after entry but before being observed.

[7] This is also done to break the connection between the covariates used to identify a control group and the outcomes of interest.

period following social housing the probability of being homelessness is around 0.13 lower for social housing residents relative to similar individuals not in social housing, who feature a homelessness rate of about 0.2. Thus social housing is providing people with more housing stability which is an important outcome. And the robustness of this result is supported by Johnson et al. (2017) finding a similar result despite applying a substantially different approach to the same data. However, as we now demonstrate, this stability doesn't seem to be translating to changing other outcomes.

We find no statistically significant robust impacts on employment, education, self-assessed mental or physical health using either matching method. We obtain negative treatment effects on incarceration and having a long-term health condition which are statistically significant at 10 per cent and 5 per cent. However, these treatment effects are not found to be statistically significant when using propensity score matching so we do not emphasise them.

Results for employment parallel the results achieved by the Productivity Commission (2015) using the Centrelink dataset. The Productivity Commission also reports lower employment rates among social housing residents when comparing averages across treatment and control samples (as we do in Table 5). But, once matching has taken place, which is broadly analogous to controlling for observable characteristics the differences disappear. Our results differ from those found in Productivity Commission (2015) and Dockery et al (2008) using the administrative data for South Australia and Western Australia. Our results are potentially stronger as we compare contemporaneous outcomes for the treated and controls rather than use data on controls reported upon application which, given the length of time individuals remain on waiting lists, may not be up to date. This doesn't rule out the possibility of positive impacts on employment for certain cohorts but this would require more extensive data to analyse. Unlike the more qualitative analysis of Phibbs and Young (2005) we don't find an improvement in self-assessed health outcomes. It would be interesting to explore the changes associated with other health related outcomes, like in Wood et al (2016) but this data is not available in Journeys Home.

In Table 8, which reports the results from comparing changes in outcomes for those who have entered social housing to the changes in outcomes for those who did not, we see qualitatively similar results to those of Table 7. The differences in the changes in outcomes were statistically insignificantly different from zero for most outcomes examined. The only exception is the probability of improving educational attainment if the kernel based propensity score matching method is used. In this case, the probability of improving educational attainment for those who moved into social housing is 0.023 lower than that for those who didn't move into social housing. However, the effect is not robust as using the nearest neighbour method it is not statistically significant.

Therefore, in general, the results for the narrower sample of individuals, those entering social housing, are similar to those comparing those in social housing with those not in – there are no consistently statistically significant treatment effects for education, employment, incarceration or health.

The most likely explanation for these results is that access to social housing has, for some time, being highly targeted to the most vulnerable members of society. Their situation may be such that while they are better off being in social housing, this does not translate systematically into quick differential changes in employment or education due to age, family commitments, disabilities or other issues. Similarly, the physical or mental health conditions associated with their vulnerability to homelessness may remain even after entering social housing, even though their housing situation has become more stable. Another possibility is that there may be cohort-specific effects that are being averaged out but may emerge in a more detailed analysis of a larger dataset.

4.3 Results using HILDA

The second stage of our robustness analysis is to repeat the analysis in the preceding section using individuals from a subset of the HILDA dataset. The variables we use are summarised in Tables A.4 to A.6 – the limitations of which have already been discussed. The next step is to discuss the outcomes of the matching process, which we report in Tables A.7 and A.8.

Table A.7 in the appendix demonstrates the matching procedure works quite well for constructing a control group for social housing residents in the HILDA dataset. Table A.8 shows though that the matching procedure does not work as well when constructing a control group for those HILDA participants entering social housing. Many of the standardised differences after matching are actually further from zero than they were prior to the matching procedure. This suggests we need to be more cautious about assigning the differences in outcomes to entering social housing for this case.

The results reported in Tables 9 and 10 show that using the HILDA data does not result in finding significant improvements in outcomes arising from social housing. Indeed, Table 9 reports that for physical and mental health, social housing residents have statistically significantly worse outcomes. And this finding is largely the same regardless of the matching methods used. Table 10 though finds similar results to Table 8. All of the treatment effects are significant except for one case. Using the nearest neighbour method, we find a significantly positive effect on education. As this result is not robust to matching method we do not emphasise it.

The much more negative results for social housing using the HILDA dataset, compared with the Journeys Home dataset are suggestive of two potential limitations of using HILDA to analyse outcomes associated with social housing. First, HILDA featured fewer covariates to undertake the matching. Hence, instead of picking up the effect of social housing we are instead picking up differences in the determinants of health status that are unobservable in the HILDA dataset but have been controlled for with participants in the Journeys Home dataset. In addition, there may also be other important unobservable characteristics of individuals that differ between social housing residents and other low-income renters in the HILDA dataset which do not differ across individuals in the Journeys Home dataset.

4. Conclusion

In this study we have taken some first steps in improving the cost benefit analysis of investment in social housing by estimating the impacts of being in social housing on individuals, identified as being vulnerable to homelessness, in Australia. We apply the latest and most comprehensive datasets, Journeys Home and HILDA, for analysing social housing in Australia to consider, simultaneously, the impacts of living in social housing on employment, education, physical and mental health, incarceration and homelessness. Using these datasets enables comparing the outcomes for existing and new social housing residents with similar individuals in the private rental market.

In general, we find placing a vulnerable individual in social housing means they are less likely, compared with other individuals also at risk of homelessness not in social housing, to become homeless. This demonstrates social housing's role as a 'safety net' for vulnerable Australians. In addition, in the short run, individuals in social housing are found to have similar outcomes in terms of employment, education, physical and mental health, and incarceration to similar individuals not in social housing. This is most likely due to the highly targeted approach to selecting residents into a relatively limited supply of social housing and the averaging across of cohort specific effects.

These results appear to be robust. The result that social housing reduces the likelihood of homelessness is robust to using different matching methods to constructing treatment and control groups. The result that social housing has no robust positive effect on other outcomes also does not vary greatly by matching methods or whether we consider new or current residents.

These results all have parallels in previous and the contemporary literature. Johnson, Scutella, Wood and Tseng (2017) utilise Journeys Home to look at factors contributing to risks of entering homelessness. Social housing here, and public housing in particular, substantially lowered the probability that the vulnerable Journeys Home sample had of entering homelessness. The Productivity Commission (2015), as well as many authors, analysing Australian and international data, also fail to find strong positive treatment effects on employment. This could be because of the trend occurring, both in Australia and internationally, of social housing increasingly being allocated to those with the greater needs that simply providing housing, while of benefit, is not enough in itself to translate into improvements in other outcomes.

This work, in the context of previous research, is suggestive that there are several ways that research in social housing proceed. Before discussing them, it is important to note that all of them require new data that is currently not available to researchers. Probably the most important limitation of our work and the way in which future research could improve on is how the effects could vary across different groups living in social housing. For example, it may be the case that younger people may benefit in a different way by having stable