INFRASTRUCTURE VICTORIA

Improving evaluation for social housing: methods and data

May 2018

TABLE OF CONTENTS

SUMMARY

One of the three roles of Infrastructure Victoria is to perform research to improve the development and appraisal of infrastructure projects amongst other issues. As part of this research program, in late 2016, Infrastructure Victoria released the paper, Moving From Evaluation to Valuation, which discussed improving project appraisals by monetising more economic and social impacts, focussing on the social infrastructure sectors of health, housing and criminal justice. In our latest research we have identified several important issues that, if addressed, could improve the methodology adopted in the social and affordable housing sector around cost benefit analysis and social return on investment analysis. These issues include:

- Identifying causal impacts of a policy or program, rather than just outcomes for those affected by it. Observed outcomes are not necessarily good measures of the impact of a program, due to outcomes being influenced by multiple concurrent factors. The methodology called the 'econometrics of program evaluation' can utilise simple statistical techniques to estimate causal impacts. There is now an extensive academic literature on developing and applying this approach. However it is yet to be broadly applied to the social and affordable housing sector.
- When identifying causal impacts it is important to separately measure any cohort-specific effects of programs. For example the educational outcomes for social housing residents are likely to be greater for children and youth, and employment outcomes are likely to be greater for those identified as in the labour force.

- To support identifying causal impacts for specific cohorts, collection of larger quantities of longitudinal data on individuals in cohorts of interest is required, drawing, where possible, on linked data.

Furthermore, we also take some first steps to providing parameters for use in the cost benefit analysis of investments in social housing, by applying the `econometrics of program evaluation' to estimate its impacts on outcomes for a specific cohort of residents that were tracked in the Journeys Home (JH) Survey. Impacts are estimated by comparing the outcomes for social housing residents with a statistically constructed control group. We also tested our findings from the JH respondents against a separate group of individuals on which data was collected through the HILDA survey. Both data sources are Australia-wide and the results should be interpreted as Australia-wide, rather than Victoria specific.

Our analysis on the JH respondents, indicates that placing an individual that is vulnerable to becoming homeless in social housing means that they are less likely, to be homeless at the end of observation period, than an individual not in social housing. The housing status of individuals in the control group could range from renting in the private sector to sleeping rough. Hence social housing is providing an important safety net. Our analysis also indicated that in the short run placing a JH or HILDA respondent in social housing has similar outcomes in terms of employment, education, physical and mental health to a respondent not in social housing. A JH respondent in social housing was also shown to have similar outcomes in terms of incarceration to a respondent not in social housing. These findings are subject to several caveats, outlined in section 4.

1. INTRODUCTION

One of the three roles of Infrastructure Victoria is to perform research to improve the development and appraisal of infrastructure projects amongst other issues. As part of this research program, in late 2016, Infrastructure Victoria released the paper, Moving From Evaluation to Valuation, which discussed improving project appraisals by monetising more economic and social impacts, focussing on the social infrastructure sectors of health, housing and criminal justice. The release of this paper was followed by research in each of the sectors to further improve the evidence that can be drawn on for cost-benefit analysis of infrastructure projects.

In this paper we further develop the discussion of how to improve the methodology around evidence for cost benefit analysis. First, we explain how measuring the outcome of a program may not yield an estimate of the causal impact of a program required for cost benefit analysis. We then outline an approach, referred to as the econometrics of program evaluation, of how to estimate causal impacts. Though there is now an extensive academic literature developing and applying this approach it is yet to substantially influence business cases and policy discussions around social housing. We highlight that the approach can be implemented using relatively simple statistical techniques like comparing averages though more sophisticated techniques can be used in more challenging situations. We follow this with some implications for the cost benefit analysis of investments in social housing and make some suggestions around generating better evidence for this. We conclude by reporting on some recent research at Infrastructure Victoria applying the econometrics of program evaluation approach to estimating the causal impacts of social housing. We find that social housing has a substantial impact in reducing homelessness for Australians identified as vulnerable to homelessness but does not result in statistically significant differences in a range of other outcomes. We interpret this outcome as reflecting a range of factors including the highly targeted selection process into social housing and the averaging across cohort-specific effects. Some caveats associated with limitations of the data for analysing social housing are also noted.

The main methodological messages of this paper are threefold. First, cost benefit analysis is more likely to be accurate and relevant when based on estimates of causal impacts generated using the econometrics of program evaluation. Secondly, that it is important to consider and measure cohort-specific effects of programs where these are significant. Thirdly, to support this, agencies need to consider collecting larger quantities of longitudinal data on individuals in cohorts of interest drawing, where possible, on linked administrative data.

2. IMPROVING COST BENEFIT ANALYSIS USING THE ECONOMETRICS OF PROGRAM EVALUATION

Main Messages

- Observed outcomes aren't necessarily good measures of the impact of a program due to outcomes being influenced by multiple factors.
- The causal impact of a program can be estimated by comparing outcomes of program participants with those of a carefully selected control group applying techniques chosen from those in the econometrics of program evaluation.
- When analysing the impact of a program it is important to consider if there are different effects on different cohorts.

When calculating the benefits from a program for a cost benefit analysis using a set of parameters, the parameters should only quantify the size and value of the impact associated with the program. We wouldn't want the value of any parameter to also reflect other factors when making the estimate. For example, the estimate of the effectiveness of a job training program should not be confused with the effects of the economy being in a recession or a boom.

Observing the average outcome for a set of individuals by comparing their status before and after participation in a program will only provide an estimate of the program's impact if there are no other systematic influences on the outcome. These influences could be external factors, like the state of the economy, or actions taken by program participants independent of the program. For example, the participants in a training program could also find jobs through their social networks independent of whether they participated in the training program.

To further illustrate this point using a real example related to social housing, consider the results from the pilot of the Journey to Social Inclusion program – an example of a Housing First type social housing program.[1] In the pilot, a group of homeless people were randomly assigned to two groups. One group (referred to in the study as group J) participated in the pilot program. The other group (referred to as group E) did not receive

[1] A "Housing First" program is an approach to homelessness whereas the homeless person is placed in a home before accessing other services to deal with the problems that contributed to being homeless. This approach contrasts with those that attempt to deal with the problems before placing individuals in housing. For more on Housing First programs see the 2017 paper by Stefan Kertesz and Guy Johnson, *Housing First: Lessons from the United States and Challenges for Australia*.

any additional assistance.[2] Group E effectively was a control group for group J. The program ran for 36 months but data was also collected following the end of the pilot. Data was collected on several outcomes. But for this example we focus on the proportion of people housed for each group by period, following the commencement of the trial, as illustrated in Figure 1.

If we only look at the outcomes for group J, it appears the program is very successful as the proportion housed increases from 20 per cent to over 80 per cent. However, if we compare these outcomes with those for group E, the size of the impact on homelessness is not as clear. Over the same period, the proportion housed of Group E rose from 0 to 60 per cent. This suggests most people in Group E were able to find housing for themselves, without assistance from the program, though it did take longer.

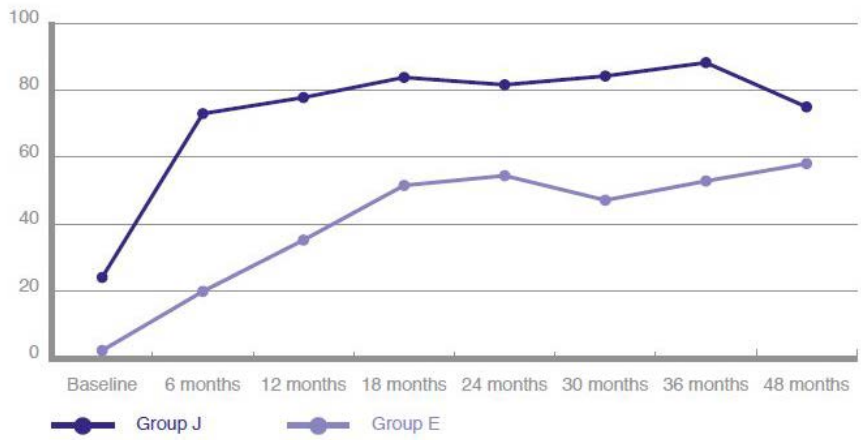

Figure 1: Proportion housed for the program participants (Group J) and the control group (Group E) following the commencement of the trial.

[2] For more about the pilot see the 2014 paper by Guy Johnson, Daniel Kuehnle, Sharon Parkinson, Sandra Sesa and Yi-Ping Tseng, *Sustaining exits from long-term homelessness: A randomised controlled trial examining the 48 month social outcomes from the Journey to Social Inclusion pilot program* published by the Sacred Heart Mission, St. Kilda.

If we didn't also observe outcomes for Group E, the impact of the program on homelessness could have been over-estimated. Not evaluating impacts properly potentially sets up programs to disappoint, and may result in wasted resources.[3]

The analysis done of the pilot of the Journey to Social Inclusion program is a good example of the main features of the econometrics of program evaluation. Indeed, this approach is now used widely in statistics and econometrics as it enables stronger claims about the causal effects of policies rather than just noting correlations between policies and outcomes.[4]

The main features of the econometrics of program evaluation approach are summarised in Figure 2. This approach attempts to solve the problem of estimating a causal impact by adopting a practice similar to that in scientific experiments. First, a treatment is specified and individuals are selected and assigned to either a treatment group or a control group. The control group is observed so to provide an estimate of the outcome for the treatment group if they hadn't have received the treatment. Second, data on the variables measuring the outcomes for each group are collected. Finally, the differences in the outcomes are calculated. The difference in outcomes is referred to technically as the treatment effect or more generally as the effect or impact of the treatment. Typically, it can be statistically tested if the treatment effect is statistically significantly different from zero and therefore

[3] Note the pilot was successful in other ways and the Journey to Social Inclusion program has been expanded (media release, Martin Foley, "A new approach to tackle homelessness", 21, December, 2017.)

[4] for more on this approach see two papers from a Productivity Commission conference; *Evidence based policy: summon the randomistas?* by Andrew Leigh and *Putting the evidence in evidence-based policy* by Jeffrey Smith and Arthur Streetman (2010), the 2014 book by Stephen Morgan and Christopher Winship, *Counterfactuals and causal inference* and the 2018 book by Andrew Leigh, *Randomistas*.

less likely to result from chance. If everything has been done properly, the estimated treatment effect can be interpreted as a causal impact resulting from the treatment.

When implementing the econometrics of program evaluation, the gold standard is to use a randomised controls trial (RCT). In a RCT individuals are randomly assigned to either the treated or control groups. If the population from which the two groups are drawn are sufficiently similar, the treatment effect can be estimated as the difference in the average outcomes for the two groups. This illustrates that if the right data is collected, under certain conditions it is not necessary to do any econometrics to implement the econometrics of program evaluation. Simple descriptive statistics will do. The best data features not only before and after observations for the treatment group but also concurrent observations for the control group.[5]

When analysing outcomes for residents in social housing, we cannot act as if the data was generated from a RCT and compare averages with a randomly selected control group. Doing so is likely to result in biased estimates of the impact of social housing. This is for two sets of reasons which we now discuss in more detail. Implementing the econometrics of program evaluation for social housing typically requires more sophisticated techniques than comparing averages which are typically easily implementable with standard software.[6] It may also require obtaining more data and better data.

[5] Such data is referred to as panel or longitudinal data.
[6] The 2010 paper by Jeffrey Smith and Arthur Streetman provides an introduction to these tools and Susan Athey and Guido Imbens 2017 paper, *The state of applied econometrics: Causality and policy evaluation*, provide a more detailed (though more technical) review of this approach.

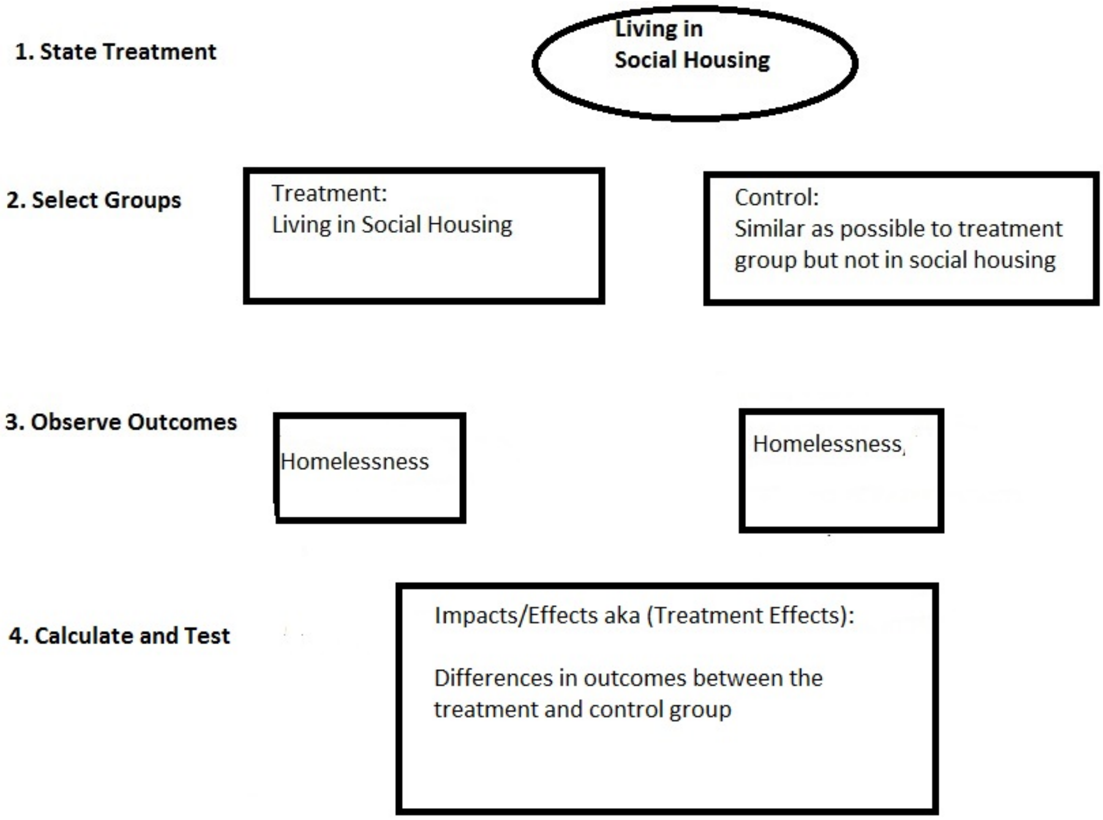

Figure 2: Applying the Program Evaluation Econometrics Approach to Social Housing

The first set of reasons for why more complex techniques are required for analysing outcomes in social housing is that people are not randomly assigned to social housing but undergo an extensive selection process. In particular, people must register to be considered for social housing. Then, even once registered how quickly or even whether at all people are selected to become residents depends on certain criteria.[7] Furthermore, the selection process has changed over time from being focussed on low income workers to giving priority to those in with the greatest needs.[8]

[7] See http://www.housing.vic.gov.au/social-housing for more detail of this process in Victoria.
[8] This is reflected in the analysis by Lucy Groenhart in her 2015 paper, *Employment of Public Housing Residents in Australian Cities*, which highlights the declining labour force participation of tenants in public housing.

The primary implication of the selection process into social housing is that residents in social housing will differ systematically from non-residents. This means differences in outcomes could result not from being in social housing but from the characteristics that meant they were selected into social housing. We cannot estimate the causal effect of being in social housing from comparing the average outcomes in this case.

One way that has been developed to overcome this problem is to use the data on the characteristics of the residents and individuals in the potential control group to construct sample treatment and control groups that are very similar in terms of these characteristics e.g. age, gender, education, employment history, health history etc. Statisticians and econometricians have developed some sophisticated tools to do this, which Infrastructure Victoria used in the research reported in the next section. However, when assessing whether a control group is similar to a treatment group it is important to consider whether the two groups may differ in ways that are not measured in the data available to the researcher. This information could be missing due either to limitations of the data collection process or inherent measurement difficulties. An example is the extent of problems with addiction, mental or physical health. Individuals in the two groups could look identical to a researcher, based on what is measured, but actually differ in the extent to which these problems are present, in ways that bias the results. Differences in outcomes could result not from the program but differences in unobservable characteristics that result from the selection process (where the individual making the selection decision does observe this information). Dealing with this issue may require more sophisticated techniques or better data or both.

The second set of reasons why more complex techniques or better data may be required for analysing social housing is that the treatment effects of going into social housing could differ across cohorts within the residents. For example, the treatment effects on education of going into social housing may be greater for children and young adults. This is because young people, given the chance, are likely to invest in education and training similar to other young people in more fortunate circumstances. By pooling individuals from cohorts with different treatment effects we end up estimating a treatment effect that is an average of the true effects. So, where applicable and possible, different treatment effects should be estimated for different cohorts. This requires sufficient data to estimate cohort-specific effects.

3. IMPLICATIONS FOR COST BENEFIT ANALYSIS AND ADDITIONAL DATA

Main Messages

- When selecting parameters for cost-benefit analysis, better estimates are likely to be obtained by using those estimated as causal effects using the econometrics of program evaluation or similar techniques.
- When selecting benefits and costs to consider in cost benefit analysis of social housing better estimates will be obtained by using those tailored to the relevant cohorts.
- More longitudinal data on cohorts in social housing should be collected via surveys or data linkages so to enable the estimation of cohort specific impacts and to strengthen program evaluation across the social housing portfolio.
- Linked administrative data holds the greatest promise of improving the data available for analysis at least cost in the least invasive way for the people under consideration.

There are two sets of changes that would be desirable for constructing and compiling evidence for evaluations done for business cases.

The first of these is when compiling parameters for cost-benefit analysis of investments in social housing it is important, where possible, to use parameters estimated from studies using the econometrics of program evaluation. Such parameters are more likely to capture causal impacts of programs rather than outcomes due to other causes. This will result in better estimates of the benefits of a program. The limitations of other approaches should be considered and noted when drawing on evidence developed not using this approach.

Similarly, when compiling benefits for cost benefit analysis of investments in social housing, expectations about the measurable economic benefits should be tailored for each cohort being considered. For example, if priority in admission to social housing is given to single women approaching retirement, less weight should be placed on employment outcomes as this is likely to be less important for them than other benefits.

The second set of changes is around generating better data for evaluation. To more richly characterise the broader set of impacts that may occur from social housing, it is necessary to have more detailed data on specific cohorts observed over a long period. For example, the cohort that is most likely to experience improved outcomes in terms of employment and education is young people. And as these benefits are most likely to be realised in the long run, it is important to able have data over the long run. For Australia, the only evidence of the effect of social housing on children is qualitative. Datasets such as Journeys Home and HILDA only target adults. The survey conducted by Phibbs and Young (2005) reports that about half responded that on going into social housing, the education of children had improved and, at most 10 per cent felt it had become worse. The recent US quantitative study of Andersson et al (2016) also find some positive impacts for children going into social housing.

Progress in Victoria, without expensive surveys, is probably best achieved through the analysis of matched administrative data. For example, linking Centrelink data with state government data on education records of children, such as NAPLAN, could allow analysis of potential effects of social housing with the control group being other similar Centrelink recipients not in public housing. The collection of data for treatment and control groups associated with the expansion of the Journey to Social Inclusion Program is a very encouraging development in this direction.

4. APPLICATION TO SOCIAL HOUSING

Main Messages

- Infrastructure Victoria applied the econometrics of program evaluation approach to estimating the impacts of social housing on employment, education, health, incarceration and experiencing homelessness, for Australia.
- The research found large significant reductions in homelessness but little significant impacts in the other areas.
- The lack of significant impacts in other areas is most likely due to a combination of estimating average impacts across all adults rather than impacts on specific cohorts and the selection process for entering social housing.
- The results are subject to caveats around the representativeness of the sample of all social housing residents, the ability to separate recent from historical impacts and from being able to observe only short run impacts for entrants to social housing.

Recently Infrastructure Victoria sought to estimate the impacts of living in social housing on employment, education, health, incarceration and homelessness.[9] The analysis was undertaken using two Australian population based data sources – the Journeys Home

[9] More details on the set of outcomes considered are in the Infrastructure Victoria 2016 paper, *Moving From Evaluation to Evaluation.* The results of this work are described in full in the 2018 paper by David Prentice and Rosanna Scutella, *What are the impacts of living in social housing?*

dataset and the HILDA dataset.[10] The Journeys Home (JH) respondents were adult Centrelink recipients that were either homeless or identified as being at high risk of becoming homeless. The HILDA dataset provides a less vulnerable cohort than the JH data, as it only considers low income renters (social and private rental) drawn from the general population and doesn't include individuals that are homeless. Both datasets exclude children and youth under the age of 18 years old. All of our results should be interpreted as Australia-wide rather than Victorian specific as we used data from respondents all across Australia – not just from within Victoria.

Identifying causal impacts is difficult in the analysis of people vulnerable to becoming homeless, as it is a complex cohort that cannot easily form a control group, as is done in medical and scientific analysis. So in order to account for the selection process into social housing a method called statistical matching was used to construct treatment and control groups within each dataset that are as similar as possible. Furthermore, the main analysis was done using the Journeys Home dataset because the unobservable characteristics of the control group were more likely to be similar to the treatment group. As well as being drawn from the general population rather than from those identified as vulnerable to homelessness, the HILDA dataset features fewer variables relevant to homelessness that can be used to match those in the treatment group with potential control group members. This potentially makes the control group in HILDA less comparable than the control group for JH data set. To tackle the change in the selection process over time, the analysis is

[10] Participants in the Journeys Home dataset were observed, for three years, with about five thousand observations. Participants in the HILDA were observed for a much longer period, with about ten thousand observations.

repeated comparing the changes in individual outcomes for the more narrowly defined treatment group of those who were admitted into social housing during the sample period, and therefore under similar admission criteria.

The research on the Journeys Home dataset reports that placing an individual vulnerable to homelessness in social housing means they are much less likely, compared with other individuals also at risk of homelessness not in social housing, to be homeless in the next observation period. This is an important outcome and is one demonstration of social housing's role as a 'safety net' for vulnerable Australians. However, individuals vulnerable to homelessness in social housing are found, during the survey period, to have similar outcomes in terms of employment, education, physical and mental health, and incarceration to similar individuals in not in social housing. The lack of other statistically significant impacts of social housing is most likely due to two sets of reasons. First, these results are averaged across a sample of adults vulnerable to being homeless in social housing. But different cohorts, for example, young people, may benefit in different ways which aren't strongly reflected in the average effects. Second, this analysis compares individuals selected through a highly targeted approach for a relatively limited supply of social housing and compares them to a similarly needy cohort, albeit not in social housing. The time individuals in the treated and control groups are likely to have spent homeless or at best in insecure housing in chronic disadvantage rather than in education or in building and maintaining human capital in the workforce will contribute to challenges with respect to quickly entering employment or education.

It is also important to note three sets of caveats on the findings due to some limitations of the data for analysing outcomes from social housing:

1. The analysis has considered a specific social housing cohort, being those in social housing being at risk of becoming homeless. This cohort was not constructed to be representative of all social housing residents. In particular it does not include any children or youth under the age of 18 years old. It may be possible however, to identify impacts for the outcomes of employment, education, physical and mental health and incarceration through the same analysis method for other social housing cohorts if the data was available.

2. The outcome measures observed for part of our analysis relates to an individual's status at a point in time. Their status however relates to both their life experience or characteristics prior to being in social housing, and the impact of social housing. The statistical matching process should in theory differentiate the characteristics from the impacts, but it is complex to disentangle the two factors, particularly due to some limitations of the data.

3. When considering the impact of moving in to social housing between those entering social housing and those remaining out of social housing, the individual could move into social housing at any time during the observation period, ranging from the first few days of the period, through to the last day of the period. The timing of entry is not reported in the data. As the observation period is also only six or twelve months long the impacts of social housing captured will only be short run.

Finally, this research also provides a real life example of the importance of thinking about unobservable characteristics when applying the econometrics of program evaluation. The result that residents in social housing, surveyed in HILDA, have statistically significantly worse physical and mental health outcomes than the individuals in the HILDA control group is interpreted as reflecting unobservable differences between the residents in social housing in HILDA and the control group because we were unable to match as well rather than as an impact of social housing.

5. THE NEXT STEPS

Infrastructure Victoria has concluded, for now, its research on improving the cost benefit analysis of infrastructure projects. The next steps, in terms of better practice, collecting better data and improving the analysis should be undertaken within government and the social/community housing sector. DTF could encourage the use in of evidence developed using the econometrics of program evaluation in business cases through its guidelines. And the line agencies and the social/community housing sector, when presenting business cases or more generally making arguments supporting greater investment in social housing could consider more heavily drawing on evidence which has established it measures causal effects. In addition, it would also be useful to generate richer longitudinal data for the estimation of cohort specific effects – possibly by drawing on linked administrative datasets already in existence.

REFERENCES

Andersson, Fredrik, John C. Haltiwanger, Mark J. Kutzbach, Giordano Palloni, Henry O. Pollakowski and Daniel H. Weinberg (2016) "Childhood housing and adult earnings: A between-siblings analysis of housing vouchers and public housing", Centre for Economic Studies Paper, CES 13-48R.

Athey, Susan and Guido W. Imbens (2017), `The state of applied econometrics: Causality and policy evaluation', Journal of Economic Perspectives, 31(2), 3-32.

Groenhart, Lucy (2015), 'Employment of Public Housing Residents in Australian Cities', *Urban Policy and Research*, 33(3), 291-305.

Infrastructure Victoria (2016) *From evaluation to valuation*.

Johnson, Guy, Daniel Kuehnle, Sharon Parkinson, Sandra Sesa and Yi-Ping Tseng (2014) *Sustaining exits from long-term homelessness: A randomised controlled trial examining the 48 month social outcomes from the Journey to Social Inclusion pilot program*. Sacred Heart Mission, St. Kilda.

Kertesz, Stefan G. and Guy Johnson (2017), "Housing First: Lessons from the United States and Challenges for Australia.", Australian Economic Review, 50(2), 220-8.

Leigh, Andrew (2010) "Evidence based policy: summon the randomistas?", Chapter ten of Productivity Commission (2010), *Strengthening Evidence Based Policy in the Australian Federation, Volume I: Proceedings*, Roundtable Proceedings, Productivity Commission, Canberra.

Leigh, Andrew (2018), *Randomistas,* Black Inc.

Morgan, Stephen L. and Christopher Winship (2014), *Counterfactuals and causal inference*, Cambridge University Press, Second Edition.

Phibbs, Peter and Peter Young (2005), *Housing assistance and non-shelter outcomes*, AHURI Final Report No. 74, Australian Housing and Urban Research Institute, Melbourne.

Prentice, David and Rosanna Scutella (2017) 'What are the impacts of living in social housing?', Infrastructure Victoria Technical Paper.

Smith, Jeffrey and Arthur Sweetman (2010) “Putting the evidence in evidence-based policy”, Chapter four of Productivity Commission (2010), *Strengthening Evidence Based Policy in the Australian Federation, Volume I: Proceedings*, Roundtable Proceedings, Productivity Commission, Canberra.

APPENDIX: ALIGNMENT WITH INFRASTRUCTURE VICTORIA’S 30-YEAR STRATEGY

Finding	30-year strategy analysis
Placing an individual that is vulnerable to becoming homeless in social housing means that they are less likely, to become homeless in observation period, than an individual not in social housing	This finding is consistent with the 30-year strategy’s emphasis on improved access to affordable housing for the most vulnerable Victorians.
In the short run placing a JH or HILDA respondent in social housing has similar outcomes in terms of employment, education, physical and mental health to a respondent not in social housing.	This finding is broadly consistent with the 30-year strategy’s contention that many factors contribute to housing stress and homelessness and that these issues can’t be solved with a housing response alone, although it is a critical part of the picture. In addition, the 30-year strategy highlighted that there is no overarching strategy in place to define what interventions are most suited for meeting the varied needs of vulnerable Victorians. As noted in section 4 Caveats, our findings are not necessarily representative of the impacts on all social housing residents and

Finding	30-year strategy analysis
	in particular children and youth under the age of 18 years old. Should data be available for other cohorts to be analysed, the findings may vary. The results may also reflect the importance of housing being provided with other support services, identified in research by others (add reference), particularly for the vulnerable cohort under consideration. Our analysis has not considered the variance in additional support services provided to individuals. Finally, our analysis has not considered the alternative housing status of those not in social housing. If the majority of the group are tenants in the private rental market, the findings support our recommendation that subsidised rental housing is a suitable alternative to social housing. If the majority of the group are sleeping rough, the results are of concern but could be explained by the other points noted in this table.
A JH respondent in social housing was also shown through our analysis to have	The 30-year strategy proposed to provide a broad range of housing from affordable

Finding	30-year strategy analysis
similar outcomes in terms of incarceration to a respondent not in social housing.	rental housing through to supported housing for varied cohorts and crisis housing. The results may reflect the importance of housing being provided with other support services, identified in research by others (add reference). Our analysis has not considered the variance in additional support services provided to individuals.