Australian National University | CENTRE FOR SOCIAL RESEARCH & METHODS

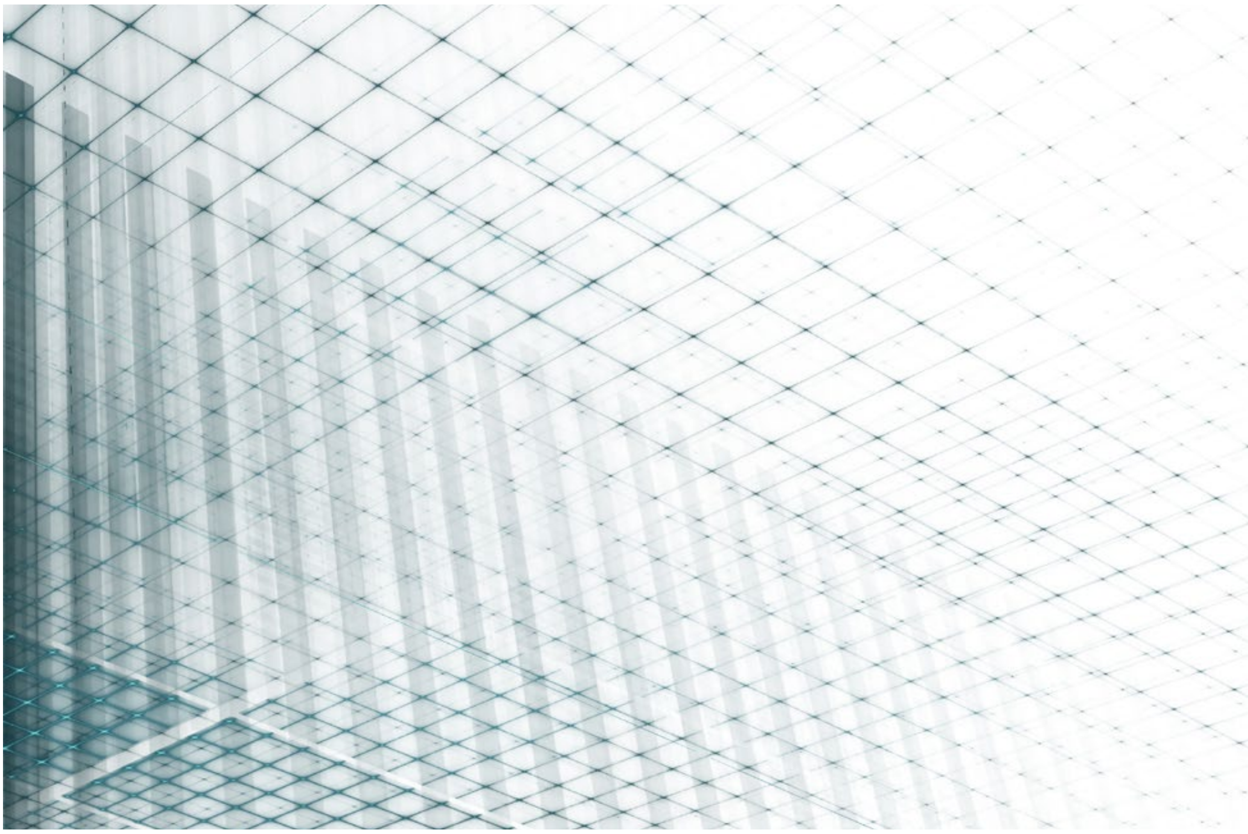

Optimal policy modelling: a microsimulation methodology for setting the Australian tax and transfer system

B Phillips, R Webster and M Gray

CSRM WORKING PAPER

NO. 10/2018

Series note

The ANU Centre for Social Research & Methods (CSRM) was established in 2015 to provide national leadership in the study of Australian society. CSRM has a strategic focus on:

- development of social research methods
- analysis of social issues and policy
- training in social science methods
- providing access to social scientific data.

CSRM publications are produced to enable widespread discussion and comment, and are available for free download from the CSRM website (http://csrm.cass.anu.edu.au/research/publications).

CSRM is located within the Research School of Social Sciences in the College of Arts & Social Sciences at the Australian National University (ANU). The centre is a joint initiative between the Social Research Centre and the ANU. Its expertise includes quantitative, qualitative and experimental research methodologies; public opinion and behaviour measurement; survey design; data collection and analysis; data archiving and management; and professional education in social research methods.

As with all CSRM publications, the views expressed in this Working Paper are those of the authors and do not reflect any official CSRM position.

Professor Matthew Gray
Director, ANU Centre for Social Research & Methods
Research School of Social Sciences
College of Arts & Social Sciences
The Australian National University
December 2018

Working Paper No. 10/2018

ISSN 2209-1858
ISBN 978-1-925715-14-9

An electronic publication downloaded from http://csrm.cass.anu.edu.au/research/publications.

For a complete list of CSRM working papers, see http://csrm.cass.anu.edu.au/research/publications/working-papers.

ANU Centre for Social Research & Methods
Research School of Social Sciences
The Australian National University

Optimal policy modelling: a microsimulation methodology for setting the Australian tax and transfer system

B Phillips, R Webster and M Gray

Ben Phillips is an Associate Professor and Director of the Centre for Economic Policy Research, ANU Centre for Social Research & Methods.

Richard Webster is a Senior Research Officer, ANU Centre for Social Research & Methods.

Matthew Gray is the Director of the ANU Centre for Social Research & Methods.

Abstract

The complexity of the social security system makes it challenging for policy makers to assess what changes should be made to the system to achieve policy objectives, and the implications of changes to the system. This paper describes the results of an initial attempt to develop a new methodology and modelling tool for optimising the social security system to achieve a particular outcome. The illustrative case used is minimising relative income poverty. We do this by using a microsimulation approach in which we alter welfare payments (or other parameters) to minimise household poverty, subject to a range of constraints, such as the overall social security budget or relationships between payment rates. The relationship between payment rate and poverty gap is then estimated using a linear regression model that provides parameter values for an equation that describes how changes in payment rates affect the poverty gap. This equation can be used to determine 'optimal' payment rates, subject to constraints such as a budget constraint or changes from current payment levels.

Acknowledgments

Previous versions of this paper were presented at the 2018 Foundation for International Studies in Social Security Conference, Sigtuna, Sweden; at the Australia–Korea Tax and Welfare Workshop; at the ANU Centre for Social Research & Methods Seminar Series; and to the Grattan Institute. We are grateful to attendees at the presentations for valuable comments. The authors are grateful to David Stanton for valuable discussions about the work. We are also grateful to members of the PolicyMod reference group for helpful comments.

Acronyms

ANU	Australian National University
CSRM	ANU Centre for Social Research & Methods
FTB	Family Tax Benefit
OECD	Organisation for Economic Co-operation and Development

Contents

Tables and figures

1 Introduction

The current Australian social security system provides a social safety net for Australians who require financial assistance to help meet their basic costs of living because age, disability, unemployment, caring responsibilities or other factors limit their ability to be in paid employment. The system also provides targeted assistance to families with dependent children, based on income level. The system helps to alleviate poverty and redistributes income from higher-income to lower-income households.

Over time, the system has evolved into a complex system of payments that vary in eligibility requirements (e.g. disability, age, whether a person is studying, whether a person has dependent children, the age of dependent children), payment rates, thresholds for private income above which the rate of government benefit is reduced, rate of withdrawal of payment as private income increases, indexing of payments to increases in the cost of living, and treatment of the incomes of other people in the income unit.[1]

The complexity of the social security system makes it challenging for policy makers to assess what changes should be made to the system to achieve policy objectives, and the implications of changes to the system. This can be posed as a question: How could the system be optimised to better achieve a policy goal, such as poverty reduction, subject to a budget constraint or some other constraint?

In this paper, we describe the results of an initial attempt to develop a new methodology and modelling tool for optimising the social security system to achieve a particular outcome. The illustrative case used is minimising relative income poverty. We do this by using a microsimulation approach that involves altering welfare payments (or other parameters) to minimise household poverty, subject to a range of constraints, such as the overall social security budget or relationships between payment rates. The simulations are undertaken using the ANU Centre for Social Research & Methods microsimulation model of the Australian tax and transfer system (PolicyMod). We have chosen to use relative income poverty to illustrate our methodology for two reasons. First, relative income poverty is widely used as an outcome measure for assessing how social security systems are operating. Second, relative poverty measures are straightforward to calculate and thus are a simpler starting point for testing this new methodology than some other measures.

In principle, the problem of determining the rates of payment that result in the lowest poverty gap could be solved by running the microsimulation model repeatedly while varying the payment rates. However, this approach is not practicable because the number of times the model would need to be run with different combination of payment rates is enormous, and this would take an infeasible time. To overcome this problem, we have developed a new methodology that drastically reduces the number of simulations required. Our methodology involves first creating a dataset that relates different combinations of the rate of social security payments to the total poverty gap in Australia using a microsimulation model of the Australian tax and transfer system. In the version of the work reported in this paper, 2500 combinations of the rate of social security payments are simulated. The relationship between payment rate and poverty gap is then estimated using a linear regression model that provides parameter values for an equation that describes how changes in payment rates affect the poverty gap. This equation can be used to determine 'optimal' payment rates, subject to constraints such as a budget constraint or changes from current payment levels.

Establishing statistical relationships between payment levels and the policy objective variable (poverty) significantly reduces the size of the problem by allowing use of standard mathematical programming techniques to optimise payment rates to achieve a particular objective. This approach means that it is not necessary to simulate a vast number of combinations of payment rates.

The modelling in this paper optimises outcomes with respect to poverty. The social security system also has important impacts on work incentives (e.g. effective marginal tax rates), income inequality and horizontal equity. The results of our research should be taken with this limitation in mind. The methodology developed in this paper could be extended to optimise other criteria, such as effective marginal tax rates or measures of inequality. We intend to extend the work to a larger range of payments, payment parameters and policy objectives in the future.

We have not been able to identify other examples of this type of approach to modelling of the social security system. There are some examples of the use of microsimulation techniques to optimise a system subject to constraints, although with substantial differences from the approach used in this paper. Ericson and Flood (2012) used microsimulation techniques to model the impacts of six possible broad designs of the Swedish tax system, to identify the design that maximised social welfare. Within each design, a number of tax system parameter values were used, resulting in the modelling of 80 different tax system designs. The authors assessed which of the 80 systems examined was optimal. Aaberge and Colombino (2013) undertook a similar style of analysis for the Norwegian system; they searched policy settings for four marginal tax rates, three income thresholds and a lump-sum transfer to find an 'optimal' income tax.

Our approach differs from this earlier work in several ways. The main difference is that the time-intensive nature of the existing approaches means that they are limited to looking at only a relatively small number of policy options, for a tax policy that has a relatively simple structure (although the specific rules and their application in the tax system are complex). Our approach enables us to deal with a much more complicated and multidimensional social security system, and to consider a very large number of possible policy settings.

Simulation techniques are widely used in other areas of optimisation and operations research – for example, in areas such as traffic flows (Papageorgiou et al. 2009), public transport (Malandraki et al. 2015) and manufacturing (Salim et al. 2017).

The remainder of this paper is structured as follows. Section 2 describes the methodology, key underlying assumptions and variable construction. Section 3 provides an overview of the performance of the model, and Section 4 describes an illustrative example of our approach, including detailed model results and some simple distributional modelling. Section 5 provides some conclusions, with a discussion of the strengths and weaknesses of the approach, and further applications for future research.

2 Methodology

2.1 Description of methodology

The modelling approach involves two steps. The first is to estimate the statistical relationship between social security system payments and the poverty gap. In the second step, nonlinear optimisation methods are used to find the policy parameters or payment rates that minimise the poverty gap, subject to a range of constraints. The payment levels from the optimisation method can then be fed back into PolicyMod to obtain the 'actual' poverty gap and further details on the distributional impact of the optimal policy settings.

In this paper, we model payments for the unemployed, single parents, the disabled, carers, the aged, families and rent assistance. These payments account for around 80% of social security cash payments.[2] The payment rates for the following five payments are allowed to vary: Newstart Allowance, Parenting Payment (single) pension, Rent Assistance, Age Pension and Family Tax Benefit (FTB) payments. The rate of the Disability Support Pension and Carer Payment are set by the rate of the Age Pension; thus, although we do not separately model these payments, they are in effect taken into account. The modelling is undertaken for the 2018–19 financial year.

Poverty is measured using the total poverty gap measure for various measures of household income. A household is defined as being in poverty if its income level (see definitions below) is less than half the value of the median household disposable income across all households (the poverty line). The total poverty gap is then defined as the difference between the poverty line and household income for households below the poverty line. The total dollar gap for all households in poverty is used as the metric rather than an average gap because the average gap can be affected by compositional change as households move into and out of poverty.

The poverty gap for household *i* calculated using PolicyMod (denoted by PM) is given by equation 1:

$$Poverty\ Gap_{PM,i} = (Disposable\ Income_{PM,i} - Poverty\ Line_{PM}) \times Equivalence\ Scale_{PM,i} \quad \textbf{(1)}$$

The total poverty gap is given by equation 2:

$$Policy\ Gap_{PM} = \sum_{i=1}^{N} w_i \cdot min(0, Poverty\ Gap_{PM,i}) \quad \textbf{(2)}$$

where *N* is the total number of households and w_i is the PolicyMod weight for household *i*.

A range of income measures are considered when defining poverty, including equivalised household disposable income and equivalised household disposable income after housing costs. Household disposable income is equivalised using the modified OECD (Organisation for Economic Co-operation and Development) equivalence scale.

2.1.1 Step 1 – estimating the statistical relationship between payment level and poverty gap

The first step involves running a microsimulation model of the tax and transfer system (PolicyMod) with randomly perturbed payment levels. In this paper, we run the model 2500 times, with the level of each of the five payments that are allowed to vary being randomly perturbed. For each of the 2500 simulations, the poverty gap is calculated for each household and summed to calculate the total poverty gap (see equation 1).[3]

Perturbations of payment level are restricted to be within the range of 70% below to 70% above the current payment level. Each simulation provides a unique poverty gap estimate.[4] Most payments have a range of payment parameters. For example, family payments vary by the age of the child, and the Age Pension payment depends on marital status. To simplify the problem, we take into account the complexity of payment parameters by applying the same random perturbation to the payment index for each payment parameter within a given payment type. For example, where our random index for family payments was 0.8, we reduce each payment rate for FTB to a factor of 0.8.

In theory, simulating random perturbations of payment levels could be used to search for the combination of payment levels that minimises the poverty gap. Indeed, for a much simpler model involving only two payments, this is quite feasible. For example, if we ran PolicyMod for 50 increments of both the Age Pension and the Newstart Allowance, 2500 simulations would be required. With the run time of PolicyMod roughly 15 seconds, an optimal solution could be obtained in around 10 hours. However, extending the analysis to five payment types increases the run time exponentially so that it would take years to solve, even with more sophisticated grid search techniques.

To ensure that the problem remains tractable, and that the modelling remains flexible enough to allow experimentation and scalability into the future, an alternative solution is required. The method that we have developed involves using a linear regression model to estimate the relationship between payment levels and the total poverty gap. The regression model is represented in equation 3:

$$Poverty\ Gap = \alpha_0 + \beta_1 X + \beta_2 X^2 + \beta_3 X^3 + \varepsilon \quad \textbf{(3)}$$

where *Poverty Gap* is the total poverty gap; X is a vector of payment rates (operationalised as indexes set randomly between 0.3 and 1.7); X^2 and X^3 are the squared and cubic versions of the payment rates; $\beta_1, \beta_2\ and\ \beta_3$ are vectors of coefficients; and ε is an error term following a standard normal distribution with mean of 0 and standard deviation of 1.

Payment levels are included in the model as a polynomial to take into account the fact that the relationship between payment level and total poverty gap may be nonlinear. As payment levels increase, the number of households being moved out of poverty increases; thus, at some point, the marginal impact of increases in payment rates on the poverty gap diminishes.

2.1.2 Step 2 – solving for payment levels that minimise the poverty gap subject to constraints

Given the aim of estimating the relationship between payment levels and the poverty gap to enable the simplification of a complex system into a simpler mathematical problem that can be solved using constrained optimisation methods, a range of other variables that might explain the poverty gap are not included. Examples of such variables are family size and composition, and employment status.

Since the setting of social security payment rates is subject to constraints (e.g. budget expenditure, relativity of payment rates), the poverty gap needs to be minimised subject to constraints. In principle, a wide range of constraints can be imposed. In this paper, for illustrative purposes, we impose three of what we expect would be the most common constraints in policy applications:

- a budget constraint
- bands within which payment rates can change relative to current payment rates
- a constraint that a specific payment be no more than a certain proportion of another payment.

The budget constraint is operationalised by multiplying an index of change for each payment by each payment's share of the existing budget. For example, if each payment were increased by an index of 1.1, this would imply that the budget for the five payments was also increased by a factor of 1.1, or 10%. This is a modest simplification of the real situation, because each payment has a different taper rate and a different share of recipients on the maximum rate – for example, a 10% increase in a maximum payment rate of the Age Pension may not lead to an exact corresponding 10% increase in the total Age Pension budget.

The constraints on changes in payment rates implemented in this paper are that payment increases must be no more than 1.6 times the current rate, and reductions must be no more than 0.6 times the current rate. The constraint on payment rate relativities is that the Newstart Allowance payment must be no more than 90% of the Age Pension rate. This constraint is considered to be a realistic example of what a policy maker might wish to impose – it is commonly argued (Treasury 2009) that, because the period of receipt of unemployment payments is expected to be much less than for the Age Pension or disability-related payments, the unemployment payment does not need to be as generous. Another reason that unemployment benefit rates may be constrained to be less generous is that policy makers are concerned to ensure a strong financial incentive for the unemployed to find paid employment. There could also be political reasons for this type of constraint. The example used in this paper illustrates how such constraints can be operationalised.

Formally, the optimisation problem is set out in equation 4, where we minimise the poverty gap as estimated in the cubic regression model in equation 3:

$$\min \ Poverty\ Gap = \hat{\alpha}_0 + \sum_{=1}^{V} \sum_{j=1}^{J} \hat{\beta}_{ij} x_i^j$$

subject to:

$x_i^{min} \leq x_i \leq x_i^{max}$

(upper and lower bounds for payment *i*), and

$x_i^{min} \leq x_i \leq x_i^{max}$

(payment *i* constrained as a maximum proportion relative to payment k), and

$B \leq \overline{B}$

(budget constraint) **(4)**

where $\hat{\alpha}_0$ is the constant estimated in equation 3, $\hat{\beta}_{ij}$ are coefficients for the ith payment raised to the power *j* estimated in equation 3, x_i^j is the payment type *i* index raised to the power *j*, *V* is the total number of payment types, *J* is the number of polynomial terms, γ is the maximum proportion of payment *k*, *B* is the index of budget expenditure, and $\overline{B}$ is the maximum value of the index of budget expenditure.

The index of budget expenditure is given by:

$$B = \sum_{i=1}^{v} x_i . \theta_i \qquad \textbf{(5)}$$

where θ_i is the current budget share for payment *i* such that $\sum_{i=1}^{V} \theta_i = 1$.

The objective is to find the payment indexes x_i that minimise the poverty gap. This is a constrained nonlinear optimisation problem. In our case, we have included squared and cubed terms in the objective function. The objective function is the estimated regression equation estimated in equation 3. The minimisation of the objective function is subject to three constraints:

- Payments are constrained to be within a minimum and maximum range above the current level.
- Some of the payments are constrained to be a maximum proportion of another payment's value.
- There is a budget constraint, B, which is set at a specified level. A value of 1 implies a budget-neutral result, whereas a value greater than 1 implies an expansionary budget, and a value below 1 implies a contractionary budget. For example, a value of 1.2 would allow the budget to increase by 20%, whereas a value of 0.8 requires a 20% reduction in the budget for the selected payments.

The optimisation problem (set out in equation 4) is solved using the SAS Operations Research software Proc NLP procedure. A version of the Newton–Raphson solution method is used. The Newton–Raphson method is a standard numerical technique for finding local optimal solutions. Global solutions are not guaranteed.[5]

Proc NLP is able to solve our cubic model in less than 1 second, which allows great flexibility and speed in finding solutions to a variety of different versions of the problem – such as after housing costs, or raw poverty gaps at either the household or income-unit level. We can also solve the problem for a large number of options for the budget constraint, allowing a range of solutions to be mapped against the allowable budget.

2.2 Definition of the poverty gap

In this paper, the poverty line is defined as an equivalised household income of less than half the median household income across households. The poverty gap is the total difference between household income and 50% of median income for each household where equivalised household income is below the poverty line. Equivalised income is calculated by applying the modified OECD scale;[6] total household disposable income is divided by the sum of the modified OECD weightings for people in the household to yield a single adult representation of income. There is no agreement in the literature about which equivalence scale should be used (Gray & Stanton 2010). Empirical studies have found that choice of equivalence scales does affect the relative poverty rates of different household types, and thus the choice of equivalence scales is expected to affect the 'optimal' payment levels. Choice of equivalence scale is an important practical consideration for policy makers. In this paper, we do not consider alternative forms of equivalence scales, focusing rather on developing a methodology. Such considerations would be a worthwhile topic for further research.

The poverty line can either be recalculated for each new set of payment amounts, meaning that the poverty line is a 'moving target', or be held constant at a baseline level. Using the approach of recalculating the poverty line can increase the poverty rate and gap for high levels of some payments, particularly the Age Pension, which has a relatively large number of recipients compared with the other payment types considered. For this reason, the analysis in this paper uses a fixed poverty line based on the current levels of social security payments.

The analysis is this paper excludes households with zero or negative incomes from calculations of both the poverty line and the poverty gap.

Housing costs are an essential item in the family budget and often a significant component. They may vary dramatically by region and particularly by age. As a result, we have also included in our modelling a version of the poverty gap that deducts housing costs from disposable income.

2.2.1 Unit of analysis

Whether poverty should be measured at the household or income-unit level is debated in Greenwell et al. (2001). We do not engage in this debate, but note that, with the optimal policy, social security policy settings are expected to differ depending on the unit of analysis; there are arguments for and against household and income-unit level in the context of social security settings.

On the one hand, eligibility for social security is largely focused on the income unit, with income and asset testing generally undertaken at the income-unit level. The household is sometimes a better representation of resource sharing, but sometimes not. For example, in a share-house household, some resources may be shared (e.g. housing costs), but incomes are usually not shared. For a household with extended family, resources and incomes may both be shared. This paper therefore provides both income-unit and household-level analysis.

While the household is the unit of analysis most commonly used for poverty analysis, in this paper we also include income-unit poverty measures because the social security system is largely defined around the income unit and, in some instances, the income unit is more appropriate for resource sharing than households.

2.3 PolicyMod

PolicyMod is a detailed microsimulation model of the Australian tax and transfer system. The model is based on the 2015–16 Australian Bureau of Statistics (ABS) Survey of Income and Housing. This survey has around 18 000 households, which we use for simulating the tax and transfer system. The survey has detailed information for each person, income unit and household, which enables the model to accurately simulate the complexity of the tax and transfer system. Because the ABS survey data for 2015–16 are unlikely to closely match up with administration numbers for the tax and social security system, and our year of interest is 2018–19, we make a number of adjustments to dollar values for incomes and payment levels. We also benchmark the population to known population estimates from the ABS and official administration data for most of the major social security payments.

2.4 Behavioural assumptions

The approach taken in this paper makes a number of behavioural assumptions. First, we assume that people do not optimise between payments. In reality, faced with different payment levels, recipients may, through choice or changed circumstance, move to an alternative payment. For example, if the Parenting Payment were reduced below the Newstart Allowance payment, which is currently less generous than the Parenting Payment, recipients may switch payments. We do not attempt to model this kind of payment-optimising behaviour. Second, we assume that there are no other behavioural responses to changes in system parameters. In reality, policy is often designed to bring about behaviour changes in areas such as labour supply decisions.

In principle, it is possible to build in behavioural change to our modelling approach. This is left for future research.

3 Performance of the model

The methodology used in this paper to estimate optimal policy is based on a model that summarises, and greatly simplifies, the relationship between the poverty gap and payment levels. We then use an optimisation method to derive optimal payment levels. For this methodology to be successful, we must firstly ensure that the regression model 'fit' between the poverty gap and the payment levels is a strong and reliable representation. The second step in our methodology is finding the optimal solution through our constrained optimisation procedure. This solution needs to be a 'global' maximum that matches up quite closely with PolicyMod. Ensuring a global maximum is not straightforward. Ensuring that the solution is closely replicated in PolicyMod is more easily tested by running the optimal solution back into PolicyMod and re-estimating the poverty gap – ideally finding a poverty gap that closely resembles that of the optimal solution. In this section, we consider how well our methodology performs in optimising payment levels to minimise the poverty gap.

3.1 Poverty gap equation

This section provides an overview of the results of estimation of the poverty gap equation. As outlined above, the poverty gap equation involves estimating the relationship between the total poverty gap and payment rates. The model is estimated using ordinary least squares. The regression models provide a robust statistical relationship between the poverty gap and our five payment levels. When modelled using the payment level, and squared and cubic payment levels, most models have an R-square statistic close to 0.99, indicating that a regression model is a very good estimator of the actual poverty gap as estimated in PolicyMod. Appendix A provides the detailed regression estimation results for poverty gaps at the household and income-unit level before and after housing costs.

Modelling was also undertaken for raw poverty rates (numbers of people) in both raw and after-housing-costs forms. We find that the model fit was not as good for raw poverty numbers, as a result of the binary nature of these poverty estimates. A large number of people can fall in or out of poverty at certain points. Maximum payment rates for, say, the Age Pension can shift above and below the poverty threshold. Because of the large number of people on these payments, the modelling results obtained are affected by discontinuities, sometimes leading to nonconvergent results for our optimal policy modelling algorithm. We prefer the poverty gap approach in this paper and do not present the raw poverty gap results.

3.2 Comparison of optimisation method and microsimulation modelling for calculating poverty gap

This section compares the poverty gap estimated from the equation linking payment levels to the poverty gap with the actual poverty gap for the particular payment levels calculated using PolicyMod. If the equation linking payment levels to the poverty gap works well, there will be a very close relationship between the poverty gap resulting from the regression estimated poverty gap for these payment rates and the poverty gap for these payment rates calculated using PolicyMod. The poverty rates calculated using the standard microsimulation process are the benchmark against which the results of the model can be compared.

Figures 1 and 2 show the relationship between the poverty gap that is calculated using the optimisation approach presented in the paper (termed 'predicted' poverty gap) and the poverty gap for the payment rates calculated using PolicyMod (termed 'actual' poverty gap). Figure 1 presents this information for the household poverty gap, and Figure 2 for the income-unit poverty gap.

If the predicted poverty gap matches exactly the actual poverty gap calculated using PolicyMod, the relationship between the two will be described by a 45° line. Figures 1 and 2 show that the data points are tightly clustered around the 45° line, indicating that the two are highly correlated. The R-square values for the relationship between predicted and actual poverty gaps at the household and income-unit levels are 0.996 and 0.992, respectively. This implies that the relationship between the estimated and actual poverty gaps are highly correlated.

Figure 1 PolicyMod and estimated poverty gap, household level, 2018

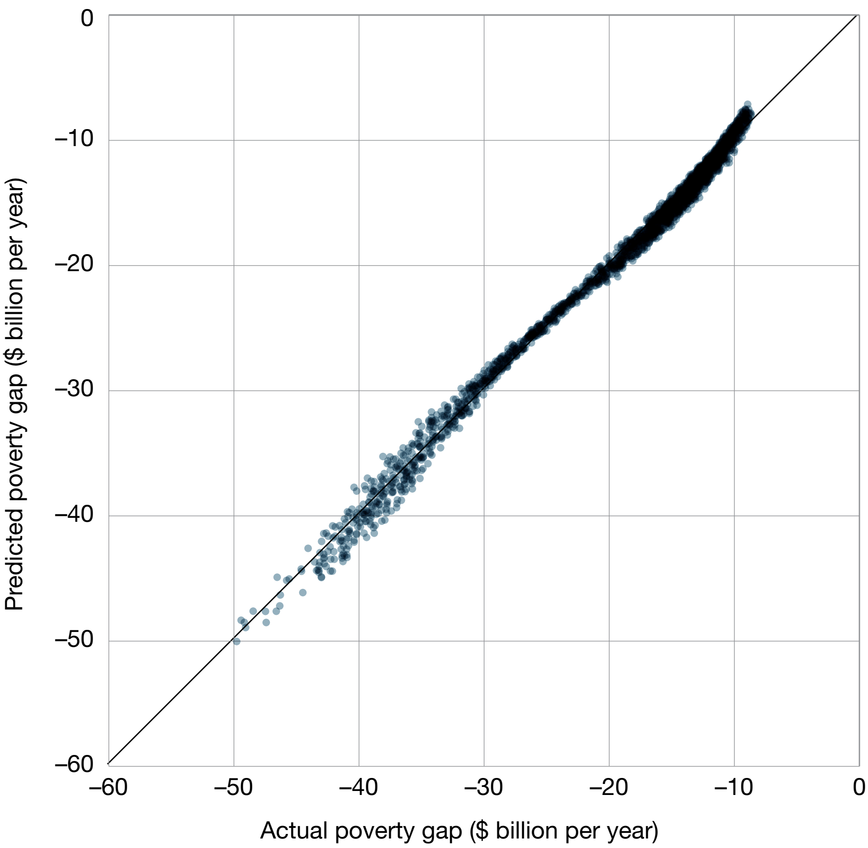

Note: The predicted poverty gap is calculated using the equation summarising the relationship between payment rates and the total poverty gap. The actual poverty gap is estimated using the PolicyMod microsimulation at the household level and then aggregating the poverty gap across households.

Source: PolicyMod

Figure 2 PolicyMod and estimated poverty gap, income-unit level, 2018

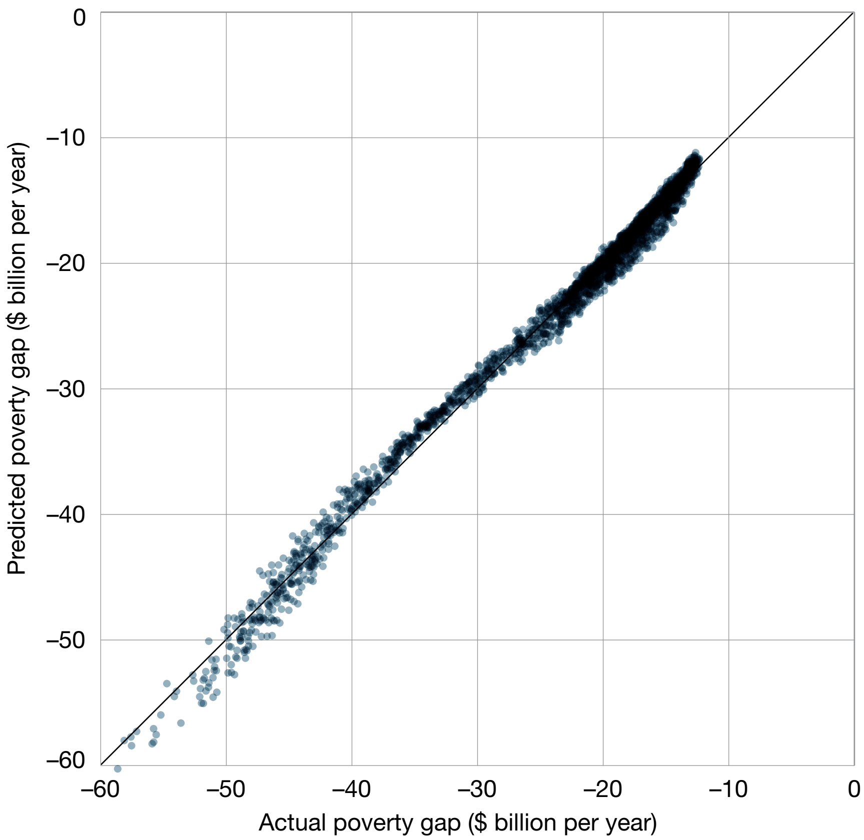

Note: The predicted poverty gap is calculated using the equation summarising the relationship between payment rates and the total poverty gap. The actual poverty gap is estimated using the PolicyMod microsimulation at the income-unit level and then aggregating the poverty gap across income units.

Source: PolicyMod

4 Illustrative example of application of the optimisation approach

4.1 Description of the policy question and constraints

This section describes the results of using the optimisation approach to analyse illustrative examples of policy problems. The examples chosen are designed to demonstrate how our approach works, the impact of various choices (including budget constraints) and how the unit of analysis affects the 'optimal' policy settings. The analysis is intended to illustrate how our methodology works and demonstrate its potential.

We use the model to help inform two hypothetical policy questions:

- What are the full rates of the different social security payments that minimise the extent of poverty experienced by Australian households as the total expenditure on social security is increased or decreased relative to the current level, and what are the implications for the extent of poverty in Australia?
- Without changing total expenditure on social security payments, what should the full rates of the different payments be to minimise the extent of poverty experienced by Australians?

The optimisation of payment rates to minimise poverty in Australia uses the poverty gap as the measure of poverty. We conduct the analysis for four specific measures of poverty:

- household poverty gap
- household after-housing-costs poverty gap
- income-unit poverty gap
- income-unit after-housing-costs poverty gap.

In the terminology of our methodology, the poverty gap is the objective function that is being minimised.

The objective function is minimised subject to three constraints:

- A budget constraint specifies the total size of social security expenditure. Social security payment levels are optimised for a range of social security expenditures, from 80% to 120% of the current overall budget for selected payments, including no change from the current level of expenditure.
- The change in each payment is constrained to a maximum of 160% and a minimum of 60% of its current level.
- The maximum Newstart Allowance payment for singles and couples is constrained to be a maximum of 90% of the maximum Age Pension single rate and couple rate, respectively (this is imposed because of the expectation that the Age Pension should be more generous than the payment to the unemployed).

These constraints are designed to ensure that the optimal solution for each parameter is bound by realistic changes, given tight federal budgets and political realities of changing welfare payments. We have also set the payment movement constraints to be wide enough to allow the model to move payments in a meaningful way, to demonstrate the extent of change that could be feasible. Naturally, we acknowledge that, even with these constraints, implementation of the changes would likely prove enormously politically difficult. As noted earlier, these results only relate to the changes that optimise one policy objective: poverty. The solutions may not necessarily provide payment levels that are sensible or reasonable from the perspective of other objectives. The research presented in this paper is a demonstration of a new methodology rather than a prescription for a new social security system in Australia.

Given these constraints, SAS Proc NLP finds optimal solutions for each objective function and associated constraints. For the analysis presented in the paper, this amounts to 164 solutions – that is, 41 separate budget constraint problems for each of the four objective functions: household poverty gap, income-unit poverty gap, household after-housing-costs poverty gap and income-unit after-housing-costs poverty gap.

4.2 Optimisation of payment levels – household poverty gap

This section presents the results of optimising selected payment levels for annual budget expenditure ranging from $20 billion less than current expenditure to $20 billion more than current expenditure (which is around $100 billion per year) – that is, 80–120% of current levels. The analysis is for the poverty gap measured at the household level. The implications of poverty measure (before or after housing costs) and unit of analysis (household or income unit) are considered in Section 4.4.

Figure 3 shows the household poverty gap for a welfare budget that varies between 80% and 120% of current levels (the poverty gap is derived from applying of the optimised payment levels to equation 1). For no change in the social security budget (around $100 billion per year for selected payments), the poverty gap could be reduced by around 7.7%.[7] An increase in the budget of $10 billion per year (around 10%), according to the modelling, reduces the poverty gap for households by 22.6% by setting payment rates at their poverty-minimising level. A reduction in payments of 10% would lead to an increase in poverty of around 11.5%. A reduction in payments of around 5% could lead to an unchanged poverty gap where optimal payment levels were set.

These results suggests that poverty could be reduced significantly by adjusting existing payment rates without increasing total social

Figure 3 **Household poverty gap with optimised payment rates, by level of social security expenditure, 2018**

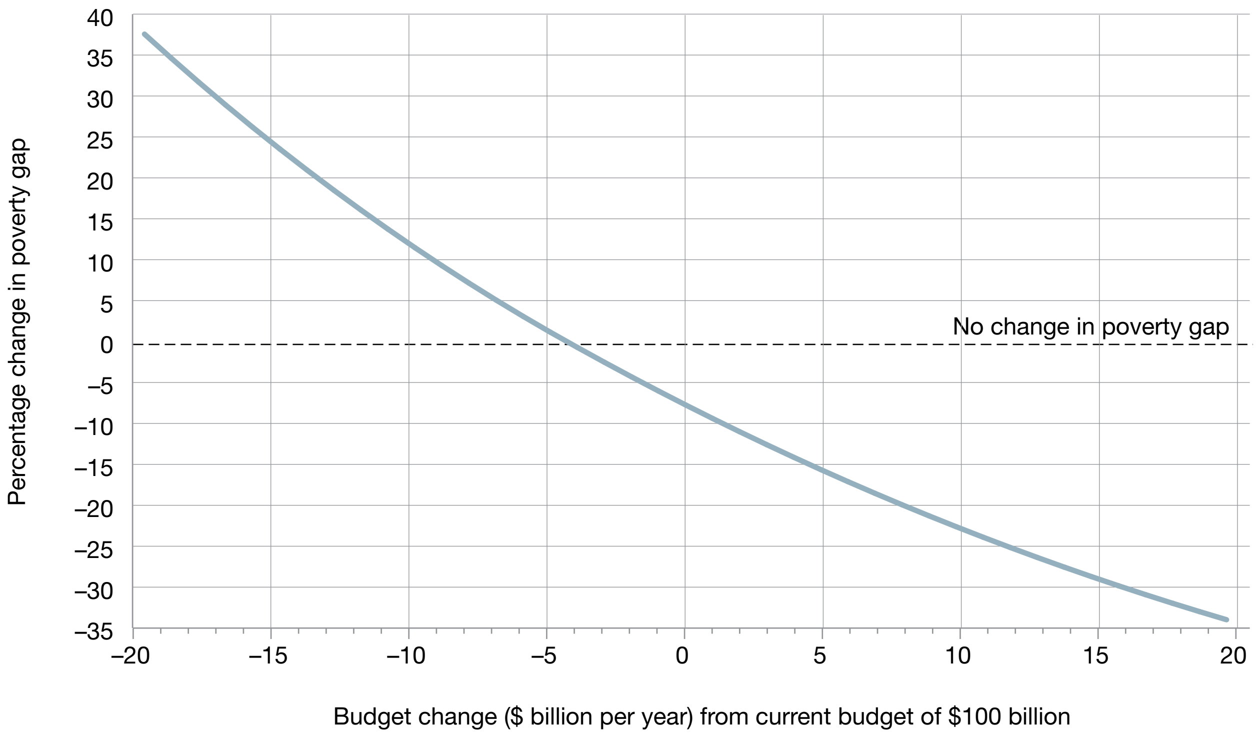

Note: The poverty gap is estimated using the equation summarising the relationship between payment rates and total poverty gap.
Source: PolicyMod

security expenditure. Although significant reductions can be made to the poverty gap, it is worth noting that considerable poverty remains, even with relatively large increases in the welfare budget.

A further complexity, and a potential reason for the significant poverty gap that exists even when social expenditure is increased and optimal payment settings are applied, is the underlying data used in this analysis. The poverty gap is calculated based on weekly income. Some people and households may have variable working hours, so that, for the week when they are surveyed, their hours and income do not accurately reflect their labour market income over a longer period, such as a year. In other cases, business income can be lumpy and highly variable, so that weekly income is not a good reflection of income over a longer period.

Figure 4 shows the levels of payments that minimise the household poverty gap at different levels of social security expenditure, and how these compare with current payments. At the current level of expenditure, the optimisation procedure suggests that, to minimise the poverty gap, the Newstart Allowance should be increased substantially from $551 to $821 per fortnight and the Age Pension single rate from $902 to $915 per fortnight. The modelling suggests that the increases in these payments would be offset by reductions in the Parenting Payment (single) from $770 to $737 per fortnight, FTB Part A for children under 13 years of age from $218 to $154 per fortnight and Rent Assistance from $137 to $131

Figure 4 **Optimal payment levels compared with current payment levels for household poverty gap, by welfare budget level, 2018**

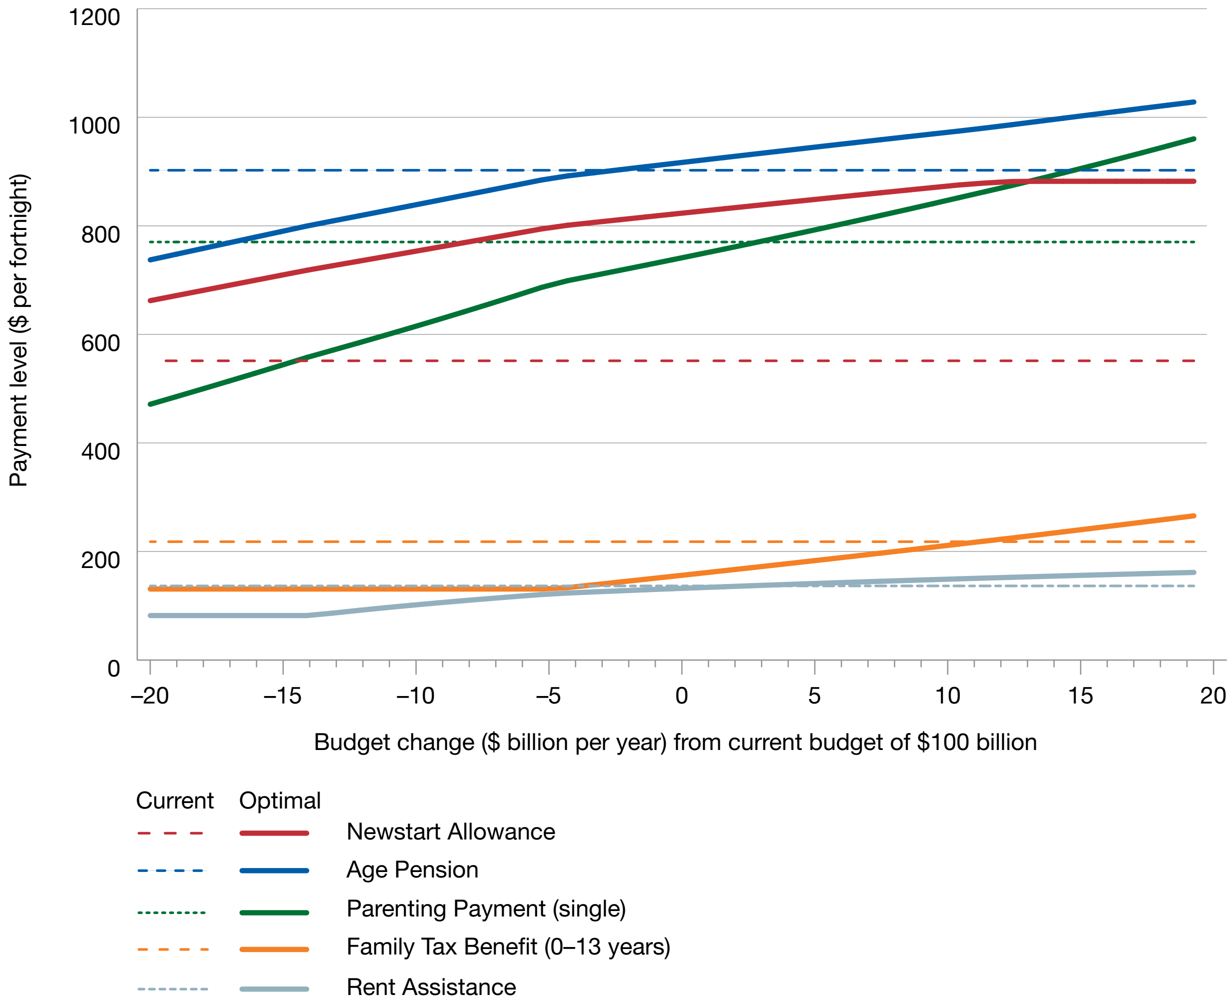

Source: PolicyMod

per fortnight.[8] As noted above, setting payments at these levels reduces the poverty gap by 7.7% (Figure 3), from $14 billion to $12.9 billion. This is a substantial increase in the efficiency of the social security system in reducing poverty.

For an increase in social security payments of $20 billion (roughly 20%), all payments are increased from their current levels. Rent Assistance and FTB are increased by the smallest amount (Figure 4). The Parenting Payment (single) and the Age Pension are both increased substantially above their current levels. Increases to the Newstart Allowance taper off because the solution is constrained to increases for any one payment of no more than 160% of the current level. Without such a constraint, we expect that the Newstart Allowance would be increased more substantially.

For a $20 billion reduction in budget for selected payments, we find that Rent Assistance and family payments are constrained by the binding limit of 60% reduction. The Age Pension is reduced compared with current levels, the Parenting Payment is reduced significantly, and the Newstart Allowance would still be significantly higher than current payment rates.

Tables 1 and 2 show the distributional results that are derived for the budget-neutral optimal

Table 1 **Impact of changing from current to optimal payment level (optimised on household poverty gap) on household disposable income, by household type and income, 2018**

	Change in annual income ($)					
Household type	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	Total
Couple with children	–166	–2477	–1,108	–96	24	–671
Couple only	264	66	209	42	8	350
Lone person	778	136	–2	0	0	617
Other	2074	1274	538	412	84	1233
Single parent	–1025	–2136	–2164	–1933	–288	–1207
Total	133	–24	–162	18	23	0

Note: Income quintiles are for equivalised disposable household income calculated for the whole population.
Source: PolicyMod

Table 2 **Percentage impact of changing from current to optimal payment level (optimised on household poverty gap) on household disposable income, by household type and income, 2018**

	Change in annual income (%)					
Household type	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	Total
Couple with children	–0.4	–3.1	–1.1	–0.1	0.0	–0.3
Couple only	0.5	0.1	0.3	0.0	0.0	0.2
Lone person	3.1	0.4	0.0	0.0	0.0	0.8
Other	4.4	1.8	0.5	0.3	0.0	0.7
Single parent	–2.8	–3.9	–2.9	–2.0	–0.2	–1.4
Total	0.3	0.0	–0.2	0.0	0.0	0.0

Note: Income quintiles are for equivalised disposable household income calculated for the whole population.
Source: PolicyMod

solution. These results are based on applying the optimal payment levels to PolicyMod to derive the 'actual' outcomes for distributional impacts, which the constrained optimisation problem does not provide. The optimal payments do lead to a reduction in household disposable income for families with children ($671 per year for couples with children and $1207 per year for single parents). Lone persons, couple-only families and 'other' household types (including share households) would all be better off. In percentage terms (Table 2), the disposable household income of single parents would be reduced by 2.8% for income quintile 1, 3.9% for income quintile 2, 2.9% for income quintile 3 and 2.0% for income quintile 4. There is little impact on income quintile 5.[9]

The intuition for these changes in payment rates is that people on the Newstart Allowance tend to be in households with high poverty gaps relative to households with people who receive FTB or Rent Assistance. A single-parent family with two young children on the maximum rate of the Parenting Payment and maximum FTB Part A for each child, and also receiving FTB Part B is, by definition, above the poverty line. Their disposable income is around $39 900 per year or $480 per week in equivalised terms. With the household equivalised poverty line at around $448 per week ($409 per week on an income-unit basis), these families can have reductions in their payments and either remain above the poverty line or incur only a modest poverty gap.

A single-person household on the Newstart Allowance receives only $14 000 per year or $275 per week in equivalised terms (the same as actual disposable income for a lone-person household) and therefore requires a very large increase in this payment to move out of poverty. In fact, the maximum increase in social security expenditure of $20 billion in our constrained optimisation problem still leaves this group below the poverty line.

Table 3 shows the poverty gap (total and per capita) for the main source of income where government payments have been split between the pensions, the Parenting Payment, allowances and other social security payments. For households, the main beneficiaries are allowee households whose average per capita poverty gap decreases from $3947 to $1493 per year – a drop of 62%. Partly offsetting this gain would be an increase in the average poverty gap for those on the Parenting Payment (single) from $267 to $891 per year. Table 3 also shows that 56% of the total poverty gap applies to households whose main source of income is not social security payments. Allowee households make up around

Table 3 Household poverty gap per year, by main source of household income, 2018

	Base world		Optimal policy	
Main source of income	**Average ($)**	**Total ($ million)**	**Average ($)**	**Total ($ million)**
Zero and negative income	–1 739	–141	–1 601	–129
Wages and salary	–101	–1 699	–102	–1 711
Business	–257	–273	–276	–293
Pensions	–634	–2 417	–640	–2 403
Parenting Payment	–267	–74	–891	–302
Allowances	–3 947	–2 485	–1 493	–1 402
Other welfare	–1 366	–836	–2 376	–707
Other income	–2 481	–5 164	–2 465	–5 130
Total	–516	–13 089	–476	–12 078

Note: The reduction in the poverty gap estimated in this table uses actual poverty gap estimates from PolicyMod rather than the regression-based constrained optimisation method. These two methods do not produce the same results.

Source: PolicyMod

19% of the gap (modestly higher than pension households), despite their population share being only 2.5%.

Table 3 also shows that a large share of the poverty gap in Australia belongs to households whose main source of income is 'other income'. These households are most heavily reliant on income sources such as share dividends, superannuation income and bank interest. The large share of the poverty gap belonging to households outside the social security system provides a limit to the extent that the social security can lower poverty and the related poverty gap. This finding also raises questions about the poverty gap measure and whether these households are truly in 'poverty'. We find that these households tend to be low-income but high-asset households. They also tend to have much lower rates of financial stress.[10] Further research into these household types, including an income measure that overcomes this issue, would be worthwhile.

The constrained optimisation algorithm finds that taking money from welfare recipients above the poverty line and giving money to those below the poverty line is the most efficient allocation of payments to minimise the poverty gap.

4.3 Optimisation of payment levels – income-unit poverty gap

As discussed in Section 2, the social security system largely operates at the income-unit level rather than the household level. It can be argued that income units are more likely than households to share income – although this is far from clear, and the reverse may also be argued. We do not take a position on this but present results from both levels. This section presents the results of setting payment levels to minimise the income-unit poverty gap.

Figure 5 shows the overall aggregate impact on the poverty gap at the income-unit level of optimising payments to minimise the poverty gap at the income-unit level. For comparison, we also include the household poverty gap. Again, we see

Figure 5 Income-unit poverty gap with optimised payment rates, by level of social security expenditure, 2018

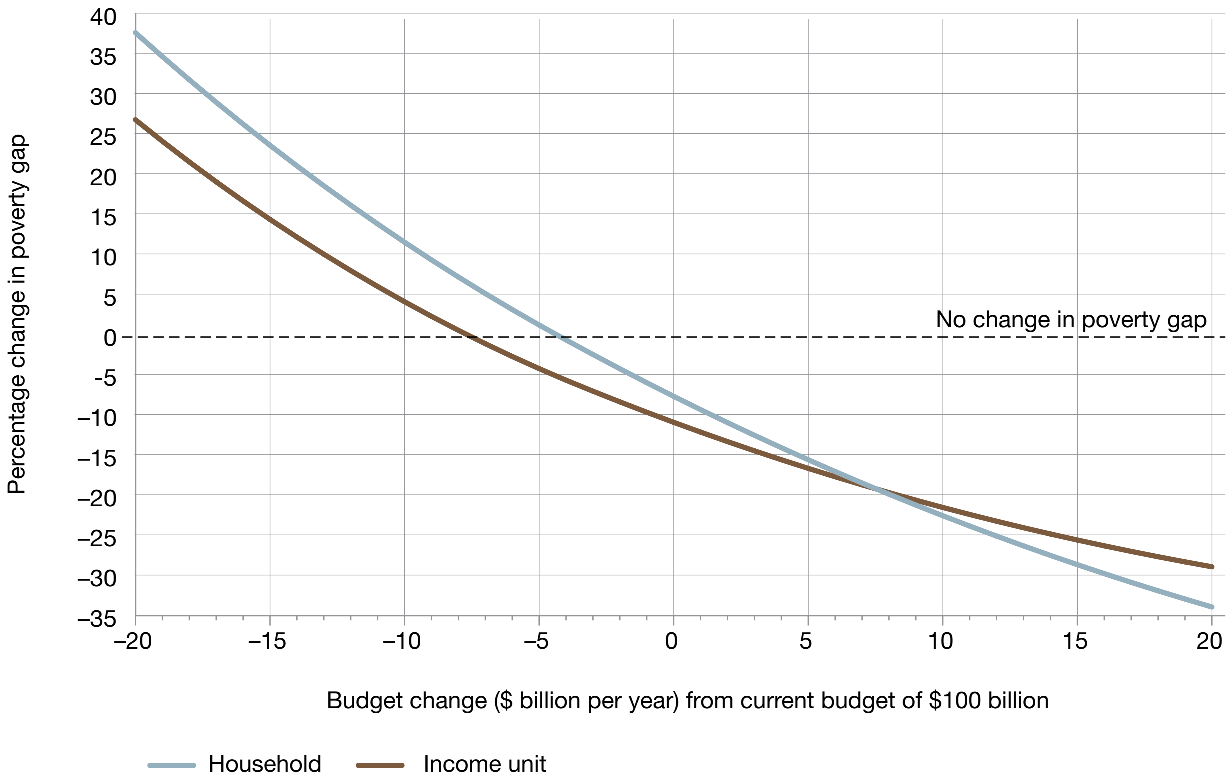

Source: PolicyMod

significant poverty reductions – in fact, larger than at the household level of analysis. For a budget-neutral reallocation of payment rates, the model produces a poverty gap 11% lower than the existing level. A 20% increase in budget would reduce the poverty gap by 29%. The modelling suggests that a 7% reduction in the welfare bill for our five payments would leave the poverty gap unchanged where payments are set at their optimal levels with regard to the income-unit poverty gap.

Figure 5 also shows that, where the social security system is expanded or contracted, the impact at the income-unit level is diminished relative to the household-based poverty gap model.

The payment levels that minimise poverty gaps at the income-unit level for a range of social security budgets are shown in Figure 6. A comparison of the payment levels for the income-unit poverty gap with those for payment optimised to minimise the household poverty gap (Figure 4) shows that, for the budget-neutral scenario, the results for the income-unit and household poverty gaps are similar. These involve significant reductions in family payments, a substantial increase in the Newstart Allowance, a modest increase in the Age Pension and some offsetting reductions in the Parenting Payment (single).

The distributional impact of the optimal policy modelling settings when the objective is minimising the income-unit poverty gap is very

Figure 6 Optimal payment levels compared with current payment levels for income-unit poverty gap, by welfare budget level, 2018

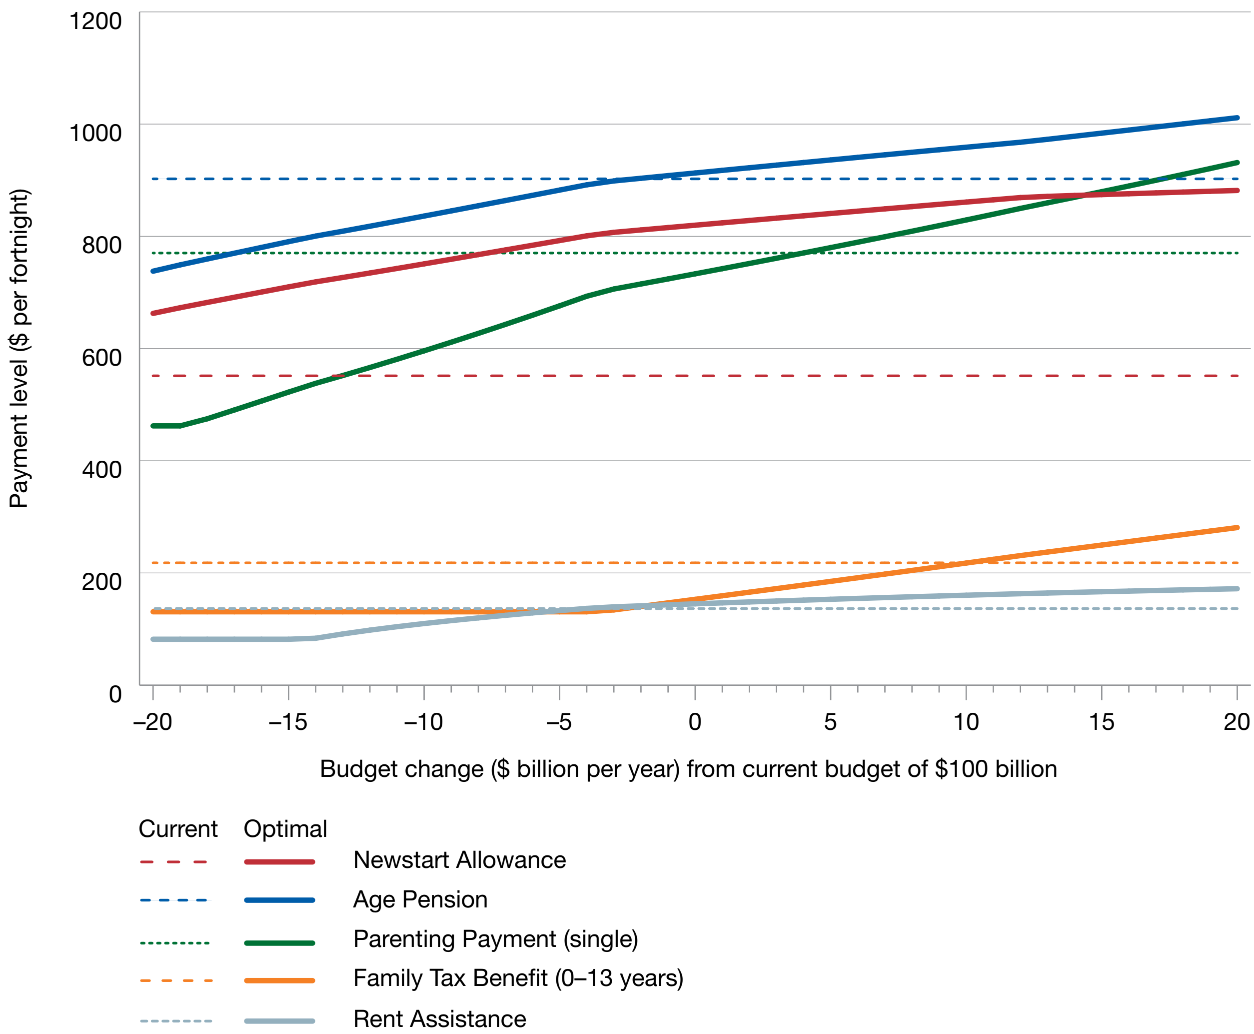

Source: PolicyMod

similar to that for the objective of minimising the household poverty gap (Tables 4 and 5). Again, couples with children and single-parent families would lose the most, particularly in low-income families. Couples and lone people would make modest gains, as would 'other' households.

Table 6 shows the poverty gap results for income units rather than households (as per Table 3). For income-unit poverty, the main beneficiaries from the optimal allocation of payments are allowee households whose average poverty gap (across all persons) decreases from $4358 to $1383 per year – a drop of 68%. Partly offsetting this gain would be an increase in the average poverty gap for those on the Parenting Payment (single) from $97 to $382 per year. Of the total poverty gap, 44% applies to households whose main source of income is not social security payments.[11] Allowee households make up around 38% of the gap (higher than pension households at 11.8%), despite their population share being only 5%.

Table 4 Impact of changing from current to optimal payment level (optimised on income-unit poverty gap) on household disposable income, by household type and income, 2018

	Change in annual income ($)					
Household type	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	Total
Couple with children	–183	–2494	–1071	–96	24	–708
Couple only	270	100	223	44	11	152
Lone person	774	132	–2	0	0	386
Other	2069	1269	514	402	81	825
Single parent	–1066	–2119	–2155	–1930	–283	–1721
Total	134	–24	–156	17	23	0

Note: Income quintiles are for equivalised disposable household income calculated for the whole population.
Source: PolicyMod

Table 5 Percentage impact of changing from current to optimal payment level (optimised on income-unit poverty gap) on household disposable income, by household type and income quintile, 2018

	Change in annual income (%)					
Household type	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	Total
Couple with children	–0.4	–3.2	–1.1	–0.1	0.0	–0.4
Couple only	0.5	0.2	0.3	0.0	0.0	0.1
Lone person	3.0	0.4	0.0	0.0	0.0	0.5
Other	4.4	1.8	0.5	0.3	0.0	0.5
Single parent	–2.9	–3.8	–2.8	–2.0	–0.2	–2.0
Total	0.3	0.0	–0.2	0.0	0.0	0.0

Note: Income quintiles are for equivalised disposable household income calculated for the whole population.
Source: PolicyMod

Table 6 **Income-unit poverty gap per year, by main source of household income, 2018**

Main source of income	Base world		Optimal policy	
	Average ($)	Total ($ million)	Average ($)	Total ($ million)
Zero and negative income	–731	–202	–748	–207
Wages and salary	–101	–1 585	–102	–1 607
Business	–283	–315	–300	–333
Pensions	–318	–1 276	–313	–1 239
Parenting Payment	–97	–41	–382	–191
Allowances	–4 358	–4 121	–1 383	–1 738
Other welfare	–990	–643	–1 487	–479
Other income	–1 188	–2 618	–1 146	–2 526
Total	–426	–10 801	–328	–8 320

Note: The reduction in the poverty gap estimated in this table uses actual poverty gap estimates from PolicyMod rather than the regression-based constrained optimisation method. These two methods do not produce the same results.

Source: PolicyMod

4.4 Sensitivity of results to poverty measure used

This section focuses on assessing how the results differ according to poverty measure used: household versus income unit, and before housing costs versus after housing costs.

Figure 7 shows that different policy objectives lead to considerable differences in the potential improvement to the poverty gap under a budget-neutral constraint. The largest gain is for the income-unit poverty gap. When we include after-housing-costs income as the objective, the income-unit poverty gap reduction is smaller but still quite substantial, at around 7.6%. The household versions of improvement to the poverty gap are not as significant as the income-unit versions. In particular, the after-housing-costs household poverty gap is only reduced by around 6.2% for the budget-neutral constraint.

Why is poverty reduction larger at the income-unit level than at the household level? Although no formal testing has been undertaken, we would expect that, since the social security system is targeted at the income-unit level, it would make sense that improving that targeting would lead to greater benefits than at the household level. Households often have multiple income units, some of which can benefit from welfare increases and some that cannot.

Figure 8 shows the changes to current payment levels that the optimal policy modelling estimates for each policy objective under the budget-neutral constraint. The general finding is that the Newstart Allowance would be increased substantially, modest increases (in some instances very modest reductions) would occur in the Age Pension and the Parenting Payment, and the other payments would generally be lowered (quite substantially for family payments), for the non-housing-adjusted income measures.

The reductions for family payments suggests that these payments are currently paid to many households and income units that are either not in poverty or have low poverty gaps. FTB can go to families, admittedly at a tapered rate, with incomes over $100 000; from a poverty perspective, it is perhaps not as well targeted with respect to income as allowances and pensions. Family benefits have other objectives, such as horizontal and vertical equity.

It is important to remember that the Australian social security system is already tightly targeted towards low-income households (Whiteford 2013). PolicyMod estimates that the current social security system in Australia distributes around 72.8% of payments to the bottom 40% of households. This is largely unchanged under the optimal policy settings.

Appendix B shows the after-housing-costs equivalent distributional tables for optimal policy modelling payment levels for both households and income units.

Figure 7 **Poverty gap reduction by poverty measure, budget-neutral optimisation, 2018**

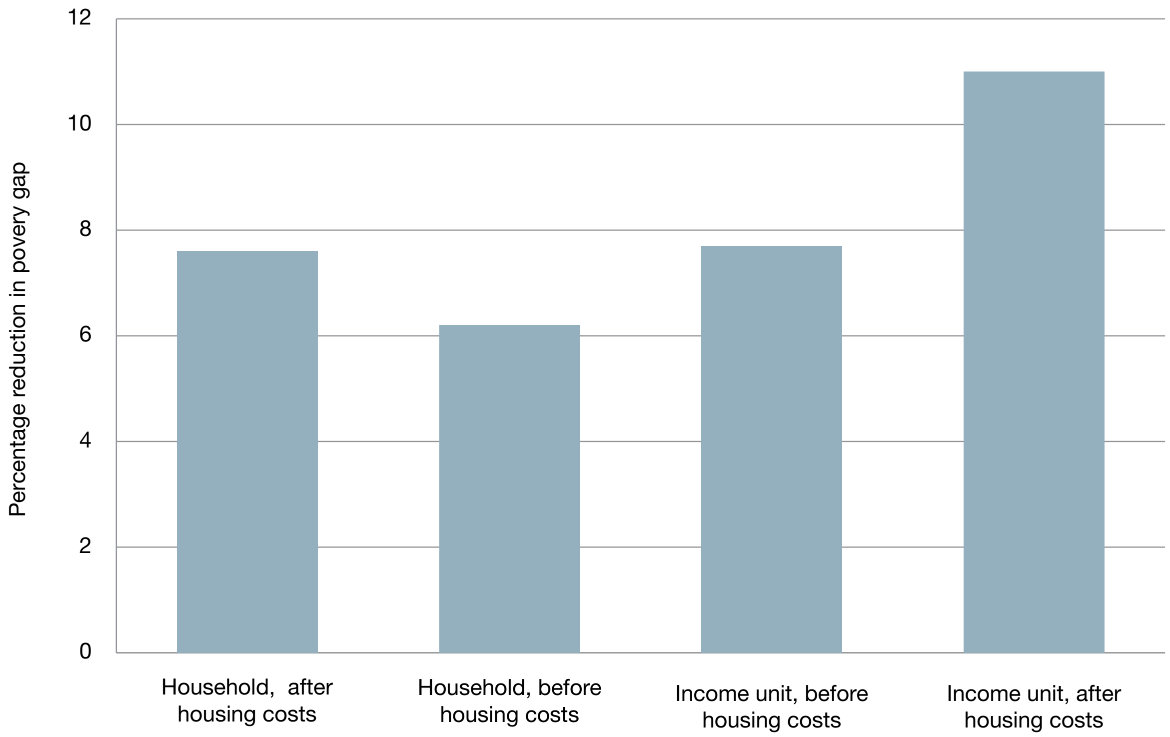

Source: PoliycMod

Figure 8 **Optimal payment level by poverty gap measure, budget-neutral optimisation**

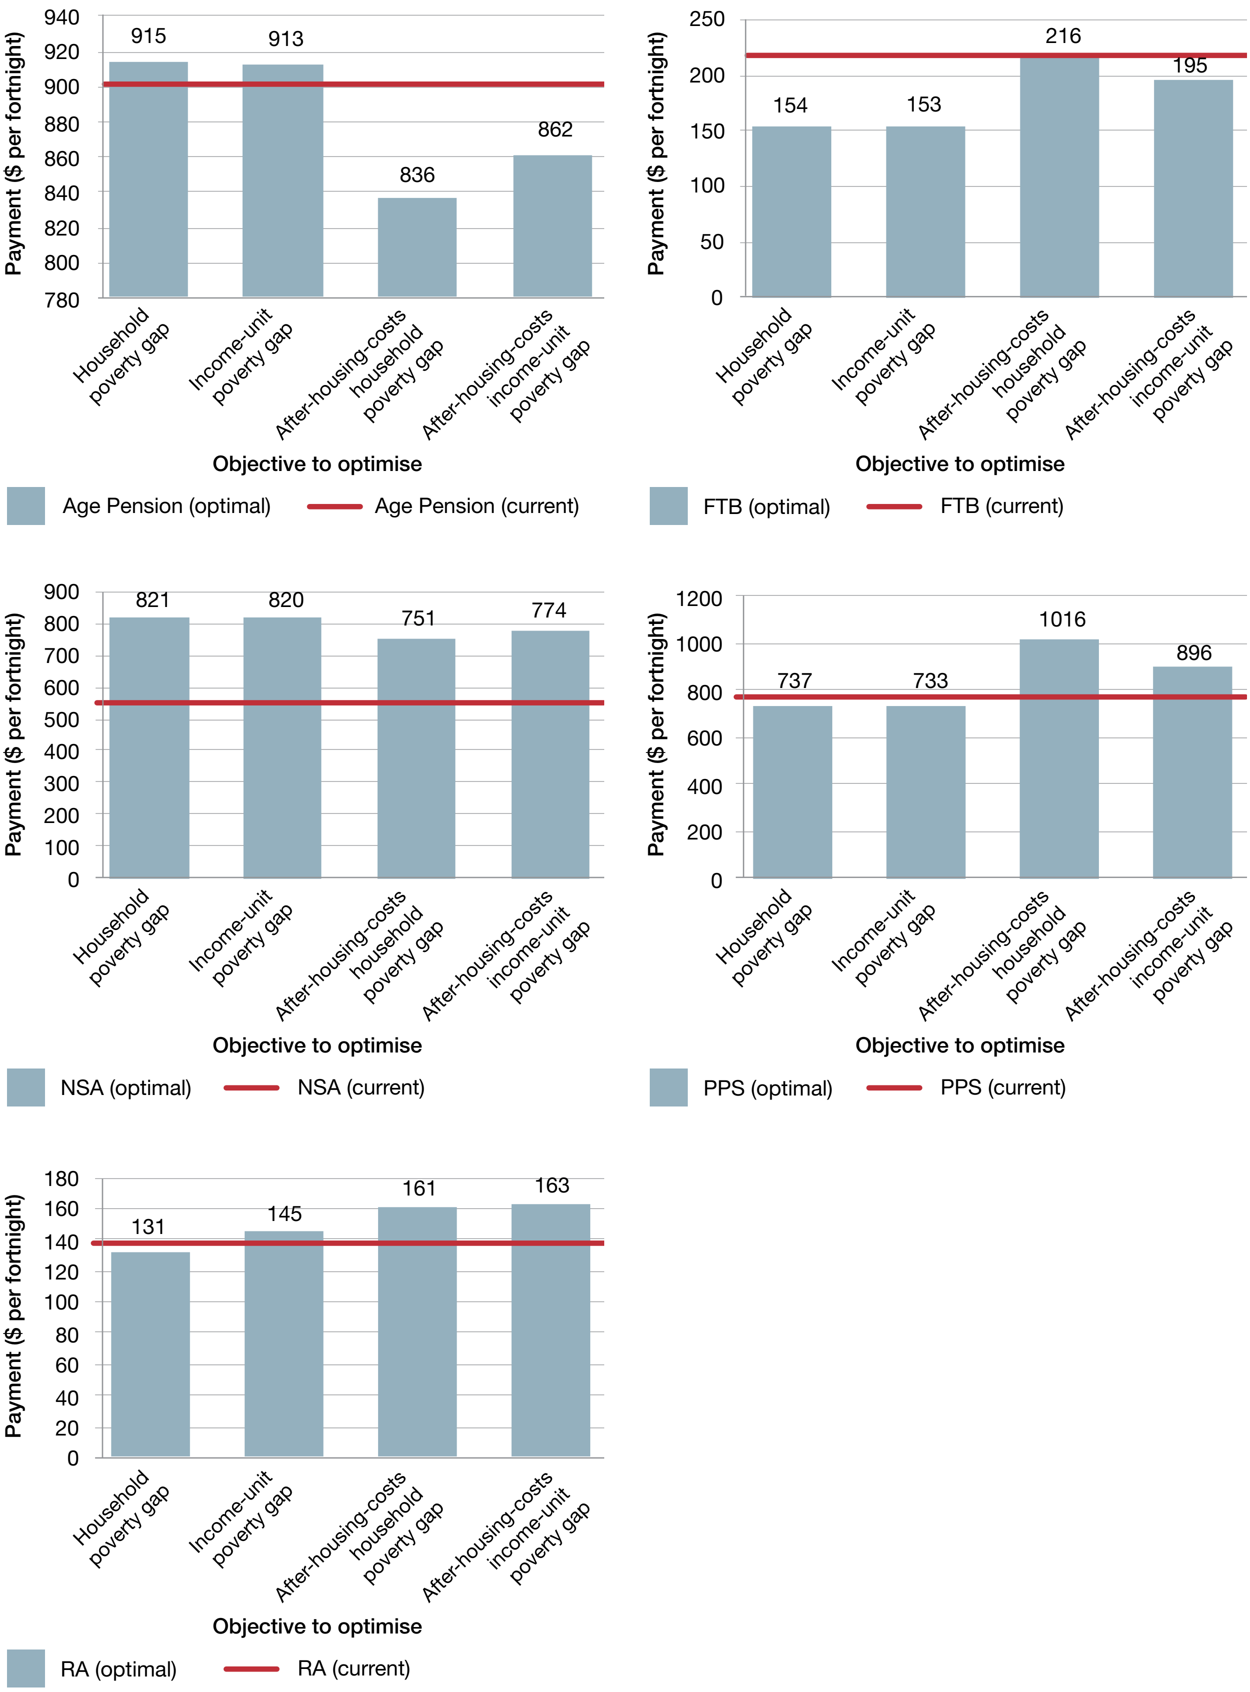

FTB = Family Tax Benefit; NSA = Newstart Allowance; PPS = Parenting Payment (single); RA = Rent Assistance
Source: PolicyMod

5 Conclusions

The modelling results presented in this paper consider the relationship between the total poverty gap and the payment levels of various social security payments. Using this relationship, we estimate 'optimal' payment levels that minimise the poverty gap for households or income units in Australia.

Important changes to the social security system during the past two decades include more generous pensions for single people in 2009 (the Harmer review), tighter income tests for receipt of FTB, and shifting some Parenting Payment recipients onto the less generous Newstart Allowance payment. Despite these changes, the current social security system is broadly based on settings from two or more decades ago. Changes have occurred in how payment levels are indexed – some are indexed to the Consumer Price Index and others to the generally higher changes in wages. This means that the relative rates of some of the key payments change substantially over time. As a result, some recipients are at increasing risk of poverty, and the system is arguably not providing an adequate safety net for some.

Given demographic, labour market and broader economic changes over recent decades, the analysis presented in this paper provides useful perspective on how well our current system performs in minimising poverty. This paper does not attempt to model a completely different social security system; rather, it considers what maximum rates of payment should be set for our main payments to minimise the poverty gap.

The paper demonstrates that it is possible to estimate relationships between payment rates and the poverty gap that closely fit the actual data, and then to use this relationship to determine the optimal payment rates that minimise the poverty gap.

The research finds that altering payments could materially reduce the poverty gap without requiring additional government expenditure. The change that would have the biggest impact on the poverty gap would be to increase the Newstart Allowance payment from its current level of $551 per fortnight to around $800 per fortnight. The exact increase varies depending on whether the household or income-unit poverty gap is being used, and on whether the before-housing-costs or after-housing-costs poverty measure is used. The modelling presented in this paper suggests that the Age Pension should remain roughly at its current level. It suggests that the Parenting Payment (single) should be increased for the after-housing-costs version of the poverty gap, but modestly reduced for the before-housing-costs poverty measure. To offset the increases in these pensions and allowances, reductions were estimated for family payments. When we consider the after-housing-costs version of income, family payments are not materially reduced, and the increase in the Newstart Allowance is funded through modest reductions in the Age Pension.

The optimal policy modelling results in this paper relate only to poverty reduction, which is only one objective of the tax and transfer system. It is important to remember that the results in this paper are only relevant to this single objective and therefore do not necessarily match up well with other objectives of the system.

A further finding in the paper is that, even with significant increases in the welfare budget, a large poverty gap still remains. Are the existing payments well targeted? Are they adequate? Are there many households that do not receive welfare that should? Are there issues with the poverty gap and the associated poverty rate where we include households that are potentially not in financial stress because of their high asset levels or only temporarily have a low income? All these questions are worthy of further investigation.

In summary, this paper provides a new methodology for setting social security payment levels. We use the optimisation approach that we have developed to determine the level of social security payments that minimise a range of poverty gap measures. Further research is planned into the method, including considering other policy objectives such as financial stress, effective marginal tax rates, and alternative social security systems through basic income or universal credit schemes. The methodology could also be expanded to the personal income tax system. Although the methodology works well for minimising the poverty gap, the analysis does raise a number of important questions around the poverty gap measure, and the underlying data and methodology that may refine and improve our approach.

Appendix A Model parameters

Table A.1 Model parameters

Variable	Household poverty gap		Household after-housing-costs poverty gap		Income-unit poverty gap		Income-unit after-housing-costs poverty gap	
	Coefficient	*P* value	Coefficient	*P* value	Coefficient	*P* value	Coefficient	*P* value
Intercept	−85.28	<0.0001	−92.09	<0.0001	−106.03	<0.0001	−107.77	<0.0001
Pension	112.24	<0.0001	105.15	<0.0001	155.58	<0.0001	131.29	<0.0001
Pension squared	−68.70	<0.0001	−68.05	<0.0001	−107.65	<0.0001	−90.23	<0.0001
Pension cubed	13.73	<0.0001	14.86	<0.0001	24.59	<0.0001	20.84	<0.0001
Newstart Allowance	5.22	<0.0001	6.50	<0.0001	6.45	<0.0001	7.82	<0.0001
Newstart Allowance squared	2.00	0.0224	1.01	0.0972	4.59	<0.0001	3.16	<0.0001
Newstart Allowance cubed	−1.32	<0.0001	−0.88	<0.0001	−2.63	<0.0001	−2.02	<0.0001
Family Tax Benefit	7.44	<0.0001	9.82	<0.0001	6.71	<0.0001	9.18	<0.0001
Family Tax Benefit squared	−2.85	0.001	−2.74	<0.0001	−2.46	0.0245	−2.58	<0.0001
Family Tax Benefit cubed	0.42	0.1476	0.24	0.2343	0.28	0.4377	0.15	0.4604
Rent Assistance	1.13	0.1614	2.10	0.0002	1.05	0.3057	2.50	<0.0001
Rent Assistance squared	0.02	0.9839	−0.15	0.808	0.23	0.834	−0.53	0.3922
Rent Assistance cubed	−0.19	0.5018	−0.23	0.2519	−0.28	0.4335	−0.11	0.5944
Parenting Payment	3.85	<0.0001	3.11	<0.0001	4.43	<0.0001	4.81	<0.0001
Parenting Payment squared	−2.13	0.0154	−0.42	0.4952	−2.64	0.0173	−2.00	0.0014
Parenting Payment cubed	0.44	0.136	−0.15	0.4528	0.54	0.1389	0.25	0.2292

Note: Shaded cells are not significant at the 10% level.

Appendix B Distributional impacts for after-housing-costs optimal policy modelling payment levels

Table B.1 After-housing-costs household poverty gap, 2018

Household type	Change in annual income ($)					Total
	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	
Couple with children	1681	421	136	–23	–10	203
Couple only	–1062	–942	105	32	10	–538
Lone person	–343	–960	–109	–6	–5	–373
Other	1069	–329	–69	41	88	110
Single parent	2677	2751	1577	1137	371	2142
Total	–437	189	167	75	19	0

Note: Shaded cells are not significant at the 10% level.
Source: PolicyMod

Table B.2 After-housing-costs household poverty gap, share of disposable income, 2018

Household type	Change in annual income (%)					Total
	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	
Couple with children	3.6	0.5	0.1	0.0	0.0	0.1
Couple only	–2.1	–1.8	0.1	0.0	0.0	–0.4
Lone person	–1.3	–2.9	–0.2	0.0	0.0	–0.5
Other	2.3	–0.5	–0.1	0.0	0.0	0.1
Single parent	7.2	5.0	2.1	1.2	0.2	2.5
Total	–1.1	0.3	0.2	0.1	0.0	0.0

Note: Shaded cells are not significant at the 10% level.
Source: PolicyMod

Table B.3 After-housing-costs income-unit poverty gap, 2018

Household type	Change in annual income ($)					Total
	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	
Couple with children	1095	–573	–301	–57	2	–112
Couple only	–575	–542	157	38	12	–282
Lone person	63	–557	–69	–4	–4	–96
Other	1424	249	134	165	83	364
Single parent	1298	915	227	–44	118	704
Total	–71	34	16	10	19	0

Note: Shaded cells are not significant at the 10% level.

Source: PolicyMod

Table B.4 After-housing-costs income-unit poverty gap, share of disposable income, 2018

Household type	Change in annual income (%)					Total
	Income quintile 1	Income quintile 2	Income quintile 3	Income quintile 4	Income quintile 5	
Couple with children	2.3	–0.7	–0.3	0.0	0.0	–0.1
Couple only	–1.2	–1.0	0.2	0.0	0.0	–0.2
Lone person	0.2	–1.7	–0.1	0.0	0.0	–0.1
Other	3.0	0.4	0.1	0.1	0.0	0.2
Single parent	3.5	1.7	0.3	0.0	0.1	0.8
Total	–0.2	0.1	0.0	0.0	0.0	0.0

Note: Shaded cells are not significant at the 10% level.

Source: PolicyMod

Notes

1. The Australian social security system is described in *A guide to Australian Government payments: 20 September – 31 December 2018* (https://www.humanservices.gov.au/organisations/about-us/publications-and-resources/guide-australian-government-payments) and the *Social security guide* (http://guides.dss.gov.au/guide-social-security-law).

2. The main social security payments not included in the optimisation problem are Carer Allowance, Veterans' Affairs, Paid Parental Leave, Child Care Subsidy and Youth Allowance.

3. This exercise was replicated with 10 separate samples of 2500 simulations. We found that the regression estimates and resulting optimal policy results were very similar for each sample.

4. Similar results were obtained over a narrower range of between 35% above and 35% below the current payment rates.

5. Although a global solution is not guaranteed, we have experimented with different starting values for finding the optimal solution, and obtain broadly similar optimal solutions for different samples and different policy objectives. This in no way guarantees that our solution is 'global', but does add to our confidence that our optimal solution is global.

6. The modified OECD scale assigns a weight of 1 for the reference person (first adult), 0.5 for the second and subsequent persons aged 15 years and over, and 0.3 for each child aged under 15 years.

7. In this instance, our modelled estimate of the poverty gap ($14 billion) is slightly higher than that estimated by PolicyMod ($13.1 billion). However, both the model and PolicyMod estimate reductions in the poverty gap from the 'optimal' solution of 7.7% compared with the current payment level poverty gap.

8. The increase in the Newstart Allowance payment is constrained in the modelling to be a maximum of 90% of the Age Pension single rate. Without this constraint, the optimisation procedure increases the Newstart Allowance above this level. The rent assistance reduction is bounded by the constraint that no payment should be lower than 0.6 of the current level.

9. The sample size for quintile 5 single parents is very small, so caution should be taken when interpreting this result.

10. The ABS 2015–16 Survey of Income and Housing shows that households with 'other income' as the main source of income and in the lowest income decile have negligible financial stress rates. These households have average net assets of around $1.2 million – significantly higher than other low-income households.

11. The overall reduction in the 'actual' poverty gap calculated using PolicyMod is larger than the gap 'estimated' by the optimisation routine. The 'actual' gap is 23% lower than the 'estimated' gap of 11%.

References

Aaberge R & Colombino U (2013). Using a microeconometric model of household labour supply to design optimal income taxes. *Scandinavian Journal of Economics* 115(2):449–475.

Ericson P & Flood L (2012). A microsimulation approach to an optimal Swedish income tax. *International Journal of Microsimulation* 5(2):2–21.

Gray M & Stanton D (2010). Costs of children and equivalence scales: a review of methodological issues and Australian estimates. *Australian Journal of Labour Economics* 13(1):99–115.

Greenwell H, Lloyd R & Harding A (2001). *An introduction to poverty measurement issues*, NATSEM Discussion Paper 55, National Centre for Social and Economic Modelling, University of Canberra, Canberra.

Malandraki M, Papamichail I, Papageorgiou M & Dinopoulou V (2015). Simulation and evaluation of a public transport priority methodology. *Transportation Research Procedia* 6:402–410.

Papageorgiou G, Damieanou P, Pitsillides A & Loannou P (2009). Modelling and simulation of transportation systems: a scenario planning approach. *Automatika* 50(1):39–50.

Salim B, Zied H & Nidhal R (2017). An optimal mathematical modeling for manufacturing/remanufacturing problem under carbon emission constraint. In: *Proceedings of the 7th International Conference on Modeling, Simulation, and Applied Optimization*, Sharjah, United Arab Emirates, 4–6 April 2017.

Treasury (2009). *Australia's future tax system*, Commonwealth of Australia, Canberra.

Whiteford P (2013). Distributional outcomes and the architecture of the Australian tax-transfer system. In: Podger A, Trewin D, Wanna J & Whiteford P (eds), *Towards a stronger, more equitable and efficient tax-social security system*, Academy Papers 1/2013, Academy of the Social Sciences in Australia, Canberra.

CENTRE FOR SOCIAL
RESEARCH & METHODS

+61 2 6125 1279
csrm.comms@anu.edu.au

The Australian National University
Canberra ACT 2601 Australia

www.anu.edu.au

CRICOS PROVIDER NO. 00120C