



RESERVE BANK
OF AUSTRALIA

Research Discussion Paper

Household Economic Inequality in Australia

Rosetta Dollman, Greg Kaplan, Gianni La Cava
and Tahlee Stone

RDP 2015-15

The Discussion Paper series is intended to make the results of the current economic research within the Reserve Bank available to other economists. Its aim is to present preliminary results of research so as to encourage discussion and comment. Views expressed in this paper are those of the authors and not necessarily those of the Reserve Bank. Use of any results from this paper should clearly attribute the work to the authors and not to the Reserve Bank of Australia.

Enquiries:

Phone: +61 2 9551 9830

Facsimile: +61 2 9551 8033

Email: rbainfo@rba.gov.au

Website: <http://www.rba.gov.au>

Figures in this publication were generated using Mathematica.

The contents of this publication shall not be reproduced, sold or distributed without the prior consent of the Reserve Bank of Australia and, where applicable, the prior consent of the external source concerned. Requests for consent should be sent to the Head of Information Department at the email address shown above.

ISSN 1448-5109 (Online)

Household Economic Inequality in Australia

Rosetta Dollman, Greg Kaplan, Gianni La Cava and Tahlee Stone

Research Discussion Paper
2015-15

December 2015

Economic Group
Reserve Bank of Australia

We would like to thank Laura Berger-Thomson, Michael Plumb, John Simon, Peter Tulip and Roger Wilkins for useful comments and suggestions. The views expressed in this paper are our own and do not necessarily reflect those of the Reserve Bank of Australia. We are solely responsible for any errors.

Authors: dollmanr, lacavag and stonet at domain rba.gov.au

Media Office: rbainfo@rba.gov.au

Abstract

We document some new stylised facts about consumption and income inequality (or ‘economic inequality’) among households in Australia. Based on household-level information from the Household Expenditure Survey we find that consumption inequality is lower on average than income inequality, but that income and consumption inequality have both increased a little since the early 1990s, with income inequality increasing by more. These findings are broadly similar to the changes in income and consumption inequality documented in other developed economies.

We provide insight into the welfare implications of these changes using panel data from the Household, Income and Labour Dynamics in Australia Survey. We decompose the broad trends in income inequality into four statistical components: (i) changes in observed household characteristics; (ii) changes in the returns to unobserved skills; (iii) changes in the size of persistent income shocks (reflecting events such as promotions and long-term unemployment); and (iv) changes in the size of transitory income shocks (reflecting events such as bonuses, short-term unemployment and short-term illness).

The reported trends in income inequality do not appear to be due to changes in observed household characteristics, but rather to changes in the size of persistent and transitory income shocks. Since the middle of the 2000s, at least some of the increase in income inequality has been due to persistent factors, a conclusion that is consistent with the rise in consumption inequality over the corresponding period.

JEL Classification Numbers: D6, D12, D31, E21, H31

Keywords: inequality, income, consumption, imputed rent

Table of Contents

1.	Introduction	1
2.	Definitions of Household Consumption and Income	4
2.1	Data	4
2.2	Imputing Housing Expenditure and Income	5
3.	Stylised Facts About Household Economic Inequality	8
3.1	Long-run Trends in Consumption and Income Inequality	8
3.2	Wealth Inequality	13
3.3	Housing Prices, Imputed Rent and Inequality Estimates	16
3.4	Consumption Inequality by Type of Household	19
4.	Cross-country Comparisons of Income and Consumption Inequality	20
5.	Transitory and Persistent Income Inequality	24
5.1	Error Components Model	27
5.2	Income Mobility	29
6.	Conclusion	32
	Appendix A: Imputed Rent and the Distributions of Household Expenditure and Income	34
	Appendix B: Alternative Survey Measures of Inequality	36
	Appendix C: Income Inequality Estimates Using Tax Records	40
	References	43
	Copyright and Disclaimer Notices	47

Household Economic Inequality in Australia

Rosetta Dollman, Greg Kaplan, Gianni La Cava and Tahlee Stone

1. Introduction

Since the early 1990s, real per capita consumption and disposable income in Australia have both risen by an average of close to 2 per cent annually. However, aggregate trends can mask important changes over time in the distribution of income and spending across households.

We explore how the distribution of living standards has evolved over recent decades by examining trends in household income and consumption inequality (which we will refer to as ‘household economic inequality’). We also explore some short-run trends in wealth inequality.

Most of the empirical research to date, particularly for Australia, has focused on inequality in current income. But current income is not necessarily a good guide to welfare. Since most individuals experience a period of growing income during their early working years, and a period of lower income as they transition to retirement, and since individuals can borrow and save to smooth out temporary fluctuations in income, overall living standards depend more on lifetime income than on current income. To gauge inequality in living standards, it is better to focus on that part of household income which is due to factors that are likely to persist through time, since this persistent component of income (reflecting things like promotions and long-term unemployment) is likely to be more strongly correlated with lifetime income than the transitory component of income (reflecting things like bonuses, short-term illness and temporary lay-offs).

These underlying factors are not easily observed in available datasets. We thus take two indirect approaches to estimate the degree of inequality in the persistent (and hence welfare-relevant) component of income for Australia. First, we follow other studies that suggest that consumption is a more appropriate measure of household

wellbeing than current income or wealth (see, for example, Slesnick (1998)).¹ Under this approach, we use repeated cross-sections of the Australian Bureau of Statistics (ABS) Household Expenditure Survey (HES) to examine how consumption inequality has evolved, relative to inequality in current income and wealth, over recent decades. We also explore some of the drivers of these changes over time.

Our second approach to estimating persistent income inequality is to exploit the panel dimension of the Melbourne Institute's Household, Income and Labour Dynamics in Australia (HILDA) Survey. By tracking the same households across time we are able to estimate a statistical model of household income dynamics that allows the distribution of temporary and persistent income to evolve separately over time. Through the lens of the estimated model, we can then measure the evolution of each type of inequality.

Our paper is motivated by the recent shift in focus in macroeconomics from the time series dynamics of household consumption and income to an exploration of the cross-sectional distributions of these variables across households and how the distributions change over time. Research into the distributions of household income and spending is an important input into identifying emerging risks to financial stability. It can also broaden our understanding of how monetary and fiscal policies affect the economy. The sensitivity (or resilience) of the household sector to shocks can be affected by which households are saving and which are borrowing at a given time. For example, aggregate spending will be particularly sensitive to changes in interest rates if there is a relatively large share of households that are constrained from borrowing.

Australian research on inequality has increased of late. For instance, Fletcher and Guttman (2013), Greenville, Pobke and Rogers (2013) and Wilkins (2015) document trends in current income inequality in Australia using household survey data. They find that there has been a slight increase in income inequality over recent years which has largely been driven by an increase in capital income at the top of the distribution. Some Australian studies have also examined trends

¹ As consumption is a choice variable, it is more closely connected with the lifetime wealth constraint faced by households than is current income. If some households smooth temporary fluctuations in income by borrowing and saving, then income will tend to be more variable than consumption at a point in time and hence income will overstate the level of inequality in household welfare.

in expenditure inequality (e.g. Harding and Greenwell 2002; Bray 2014) and non-durable consumption inequality (Barrett, Crossley and Worswick 1999). Bray (2014) finds that expenditure inequality has been lower and more stable than income inequality over the past three decades, although it appears to have been increasing since the late 1990s.

We construct estimates of household economic inequality using several sources of data. Our consumption inequality estimates primarily come from the HES. Nevertheless, we explore measures of inequality using other data sources, including the HILDA Survey, the ABS Survey of Income and Housing (SIH) and data based on individual income tax records provided by the Australian Taxation Office (ATO).

We will consider a wide range of inequality measures, including indicators that are specifically designed to measure inequality, such as the Gini coefficient, along with simpler measures that are advocated in Piketty (2014), such as the share of total household income held by the highest-earning households and the variance of the logarithm of income. Our paper makes the following contributions:

1. To the best of our knowledge, we are the first to quantify the extent to which the trends in income inequality in Australia are due to changes in observable characteristics and to changes in the distribution of unobserved persistent and temporary income shocks.²
2. We construct estimates of consumption inequality that are broader than previous Australian studies.
3. We examine in detail how housing prices can affect estimates of household economic inequality. In particular, we show how it can have different effects on estimates of wealth inequality as compared with income and consumption inequality.

2 We focus on trends in inequality for *household* income in this paper in order to allow for direct comparisons with the consumption trends, which are only available at a household level. We have also examined the trends in *individual* income inequality, but the key results are basically the same.

Our key findings are that:

1. Consumption inequality is lower on average than income inequality and has risen by less since at least the early 1990s.
2. The rise in consumption inequality has reflected a relatively large increase in spending by the highest-spending households (within the top 1 per cent of the distribution).
3. The increase in income inequality over the past decade has not been due to observable factors, such as an ageing population or rising educational attainment. Instead, it has reflected an increase in the variance of unobserved shocks, particularly since the middle of the 2000s. At least some of the increase in income inequality has been persistent, implying higher inequality in household welfare.

2. Definitions of Household Consumption and Income

2.1 Data

The analysis in this paper is primarily based on unit record data from the HES for six different surveys: 1984, 1988/89, 1993/94, 1998/99, 2003/04 and 2009/10. The HES is the most comprehensive source of cross-sectional information on household expenditure in Australia.³ For comparability with the spending estimates, we focus on the HES estimates of income. To examine the drivers of inequality we also examine measures based on the HILDA Survey. In the Appendices, we provide alternative estimates of inequality using tax records.

It is not straightforward to use the HES to derive a long time series of either expenditure or income. A key obstacle to making time series comparisons of income inequality is that the ABS has developed more sophisticated ways to measure income over time. For example, in 2003/04, the ABS incorporated information on salary sacrificed income into their household income estimates for the first time. This is likely to have boosted measured inequality relative

³ The HILDA Survey also collects annual estimates of expenditure. However, the expenditure definitions have changed over time and are not as complete as the HES.

to earlier surveys as high-income earners are more likely to engage in salary sacrificing. Despite this, in each HES, the ABS provides estimates of income based on definitions from earlier surveys. This helps us match the income measures over time to generate a reasonably comparable time series. Moreover, we have found that the income definitions can affect the estimated level of inequality in any given survey, but the broad trends in measured inequality are similar regardless of the definition of income.⁴

In this paper we examine inequality in both gross and disposable household income. This allows us to examine the role of government taxes and transfers in affecting inequality. We follow the ABS in defining disposable income as gross income after deducting personal income tax and the Medicare levy. In addition to changes in the definitions of income, the ABS also changed the way it collects household-level tax data. Prior to the late 1980s, the tax data are calculated using a combination of actual reported taxes and imputations, but the tax data for the later surveys are entirely imputed, which is now the preferred method of the ABS for estimating taxes in household surveys. This complicates comparisons of inequality in disposable income before and after the early 1990s (Barrett *et al* 1999). Partly for this reason, we mainly focus our analysis on the period since the early 1990s.

2.2 Imputing Housing Expenditure and Income

To construct our preferred estimates of household consumption and income we adjust the raw data. Most importantly, we add a service-flow equivalent of housing expenditure for owner-occupiers (or ‘net imputed rent’) to both the consumption and income estimates. Imputed rent is the value of housing services that owner-occupiers receive from living in a rent-free dwelling and it constitutes a significant component of non-cash household income and consumption.

Most guidelines for the compilation of income distribution statistics recommend the inclusion of imputed rent in both consumption and income. Conceptually, the inclusion of imputed rent as part of income treats owner-occupiers as if they were renting the home from themselves, so they are simultaneously paying rent and

4 Despite these caveats, the ABS publishes its own time series of income inequality estimates based on the SIH. The trends in the SIH estimates broadly align with those identified in this paper. Wilkins (2013) provides a very detailed discussion of the relative merits of the inequality estimates obtained from the various data sources.

earning rental income (Saunders and Siminski 2005). The imputed rent adjustment essentially makes estimates of consumption and income for renters comparable to those of owner-occupiers. Doing otherwise can lead to unintuitive results.

To see this, consider the following example. Persons A and B live next door to each other in identical homes. They are the same in all respects; they pay the same amount of rent, spend the same amount on other goods and services, and they have the same income and wealth. Suppose that person A decides to buy the home they currently rent by running down their savings in a bank deposit. In contrast, person B continues to rent. Without any imputed rent adjustment, person A's measured expenditure falls relative to person B because they no longer pay rent. And their measured income also falls relative to person B because they lose the interest earnings on their deposit account. In a sense, without adjusting for imputed rent, person A would appear 'worse off' than person B simply because they became a home owner.

In contrast, with an imputed rent adjustment, person A's consumption is unchanged as, under reasonable assumptions, the rent that is imputed is the same as the existing market rent. Their income is also unchanged to the extent that the imputed rent is the same as the interest they previously earned on their savings account.⁵ In other words, with the adjustment for imputed rent, persons A and B are still basically in the same welfare position as before, despite person A becoming a home owner.⁶

Net imputed rent is equal to the estimated market rent of a dwelling ('gross imputed rent') less housing costs normally paid by a landlord such as mortgage interest, rates, insurance and repairs. Total household 'consumption' is then equal

5 This requires some 'hand waving', as standard theory would suggest the returns on the home and the bank deposit need not be the same every period but only over the lives of the two assets.

6 Similar logic applies if person A borrows the full amount to buy the home rather than selling an existing asset. For simplicity, assume it is an interest-only mortgage loan. Without the imputed rent adjustment, person A's measured expenditure falls relative to person B because the home owner pays interest rather than rent, and interest payments are not part of consumption. Person A's measured income also falls as interest payments are deducted from income (otherwise interest earnings and interest payments are treated differently). With the imputed rent adjustment, person A's consumption is unchanged by the home ownership decision (as before). And their income is unchanged to the extent that the imputed rent is equal to the interest paid on the loan. If the imputed rent is larger (smaller) than the interest paid then person A's income rises (falls) relative to person B.

to total household ‘expenditure’ on goods and services plus net imputed rent. Similarly, ‘adjusted’ income is equal to reported income plus net imputed rent. In Appendix A, we provide a more detailed description of the differences between household consumption and expenditure.

The gross imputed rent estimates are based on the self-reported value of each owner-occupier’s dwelling; weekly gross imputed rent is defined to be equal to 5 per cent of the self-reported value of the owner-occupier’s dwelling (divided by 52 weeks). The choice of 5 per cent for the ‘imputed rental yield’ is based on previous Australian research (Yates 1994; Saunders and Siminski 2005). The benefit of this approach to estimating imputed rent is that it is straightforward to implement and it fully utilises the available self-reported data on dwelling values. As the ABS has only made information on the reported dwelling value publicly available from 1993/94 onwards, we concentrate on the most recent four surveys: 1993/94, 1998/99, 2003/04 and 2009/10.

In Appendix B, we provide estimates of inequality using an alternative measure of imputed rent based on a hedonic modelling framework. This modelling approach estimates the market value of the rental equivalent for owner-occupied dwellings using information on comparable rented dwellings. This alternative approach allows the implied rental yield to vary over time. A comparison of the two approaches highlights the fact that measures of inequality are somewhat sensitive to the treatment of housing income and expenditure. Nevertheless, the general trends in household economic inequality are fairly similar under this alternative approach. We find that consumption and income inequality have increased since the early 1990s using either the baseline or alternative approach to estimating imputed rent. For a more detailed discussion of the inequality estimates using this alternative approach, see Beech *et al* (2014).

Our estimates of consumption deduct both mortgage interest payments and interest payments on other forms of debt (e.g. personal loans and credit cards) from total expenditure. Interest payments do not represent a flow of services to the household. All income and consumption estimates are population weighted and divided by an equivalence factor to control for household size and composition.⁷

⁷ The estimated trends in inequality presented in this paper are largely unaffected by the use of an equivalence factor.

There are some caveats to our consumption estimates. First, consumption is a better guide to living standards than current income, but it is still not a complete measure of household wellbeing. Most notably, our estimates do not include measures of consumption of public goods (e.g. recreational facilities), social transfers in kind (e.g. government-funded goods and services such as public health care and education), or goods that are produced within the home. Data limitations prevent us from constructing these broader estimates of consumption. By excluding items such as social transfers in kind, we will tend to overstate the level of economic inequality (Barrett *et al* 1999).⁸ But it is less clear whether the exclusion of these items affects the estimated trends in inequality. Second, we do not convert all durable goods expenditure to a service-flow equivalent because we do not have long-run household-level data on the stock of such durable goods. However, we have found that excluding spending on particular durable goods, such as motor vehicles, has little discernible effect on our inequality estimates.⁹ Third, we also do not examine trends in the distribution of leisure time, which is another indicator of household wellbeing (Attanasio, Hurst and Pistaferri 2014).

3. Stylised Facts About Household Economic Inequality

3.1 Long-run Trends in Consumption and Income Inequality

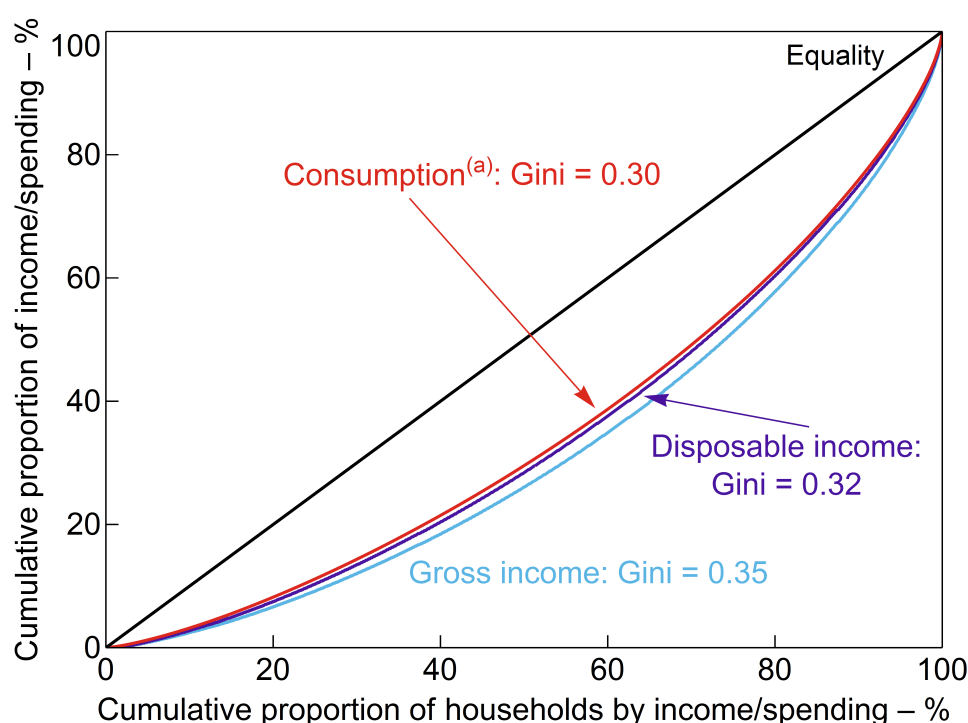
There are several different indicators of inequality that are typically used in the literature. The most common measure of inequality is the Gini coefficient, which is derived from the Lorenz curve. The Lorenz curve shows the share of spending (or income) by households ranked by spending (or income). The further the curve is below the 45 degree line, the less equal the distribution. Correspondingly, the Gini coefficient is calculated as the area between the Lorenz curve and the 45 degree line divided by the total area under the 45 degree line. The Gini coefficient ranges from zero to one, where zero represents perfect equality and one represents complete inequality.

8 The most recent HES for 2009/10 provides estimates of social transfers in kind. The inclusion of such transfers reduces disposable income inequality by about one-quarter, on average.

9 The 2003/04 and 2009/10 HES provide information on the total (net) value of the stock of vehicles owned by households. Replacing the reported expenditure on motor vehicles with the service-flow equivalent (measured as 10 per cent of the net value of the stock of vehicles) slightly reduces the estimates of expenditure inequality for the two survey years, and leads to a larger increase in measured inequality between the two periods.

Based on the 2009/10 HES, the Lorenz curve for gross income indicates that the top 20 per cent of households earned approximately 42 per cent of total household income (Figure 1). In contrast, the bottom 20 per cent of households earned about 7 per cent of income. However, income inequality is reduced to some extent by the redistribution of income from rich households to poor households through government taxes and transfers. As a result, in 2009/10 the Gini coefficient for disposable income (0.32) was lower than for gross income (0.35).

Figure 1: Lorenz Curves
2009/10



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

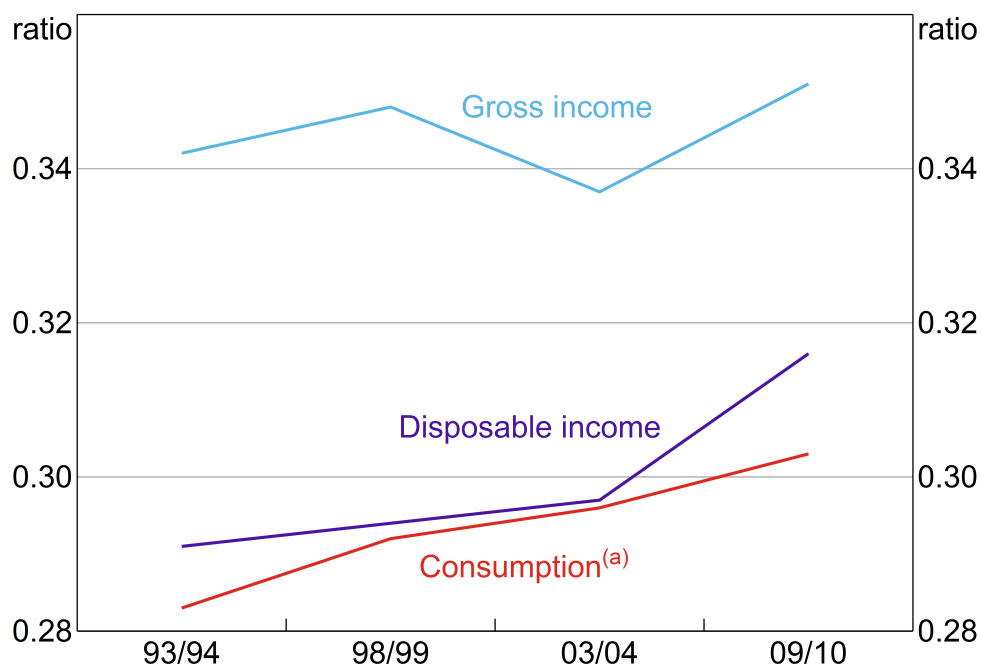
(a) Also excludes other interest payments

Sources: ABS; Authors' calculations

In addition, the Gini coefficient for consumption (0.30) was lower than that of disposable income (0.32), suggesting that economic inequality is further reduced by the ability of households to borrow and save to offset temporary changes in income. The Lorenz curve for consumption indicates that the highest-spending households (in the top 20 per cent) accounted for approximately 39 per cent of total spending in the economy. The lowest-spending households (in the bottom 20 per cent) accounted for about 8 per cent of total spending.

Next, we examine how inequality in both consumption and income has evolved over recent decades. Based on the Gini coefficient, gross income inequality is a bit higher than in the early 1990s, although it moderated in the early 2000s (Figure 2).

Figure 2: Consumption and Income Inequality
Gini coefficient



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

(a) Also excludes other interest payments

Sources: ABS; Authors' calculations

In contrast, the HES estimates of disposable income inequality have risen consistently since the early 1990s. This finding is consistent with other studies and largely reflects an increase in capital income inequality, as labour income inequality has been little changed over recent decades (Greenville *et al* 2013). According to Greenville *et al*, labour income inequality has been little changed over the past two decades due to two offsetting effects. On the one hand, high-income households have benefited relatively more from rising hourly wages for full-time employees and an increase in the share of part-time employment (which have tended to increase inequality since higher-income households are more likely to be double-income households). On the other hand, low-income households have benefited relatively more from the reduction in the share of jobless households

(which has tended to reduce inequality), which is consistent with the substantial trend decline in the unemployment rate since the early 1990s.

Consumption inequality has been consistently lower than both gross and disposable income inequality. Furthermore, the increase in consumption inequality has also been less pronounced than the increase in disposable income inequality since the early 1990s. In Section 4 we show that these trends have also been observed in other advanced economies. We explore drivers of these changes in the next section of the paper.

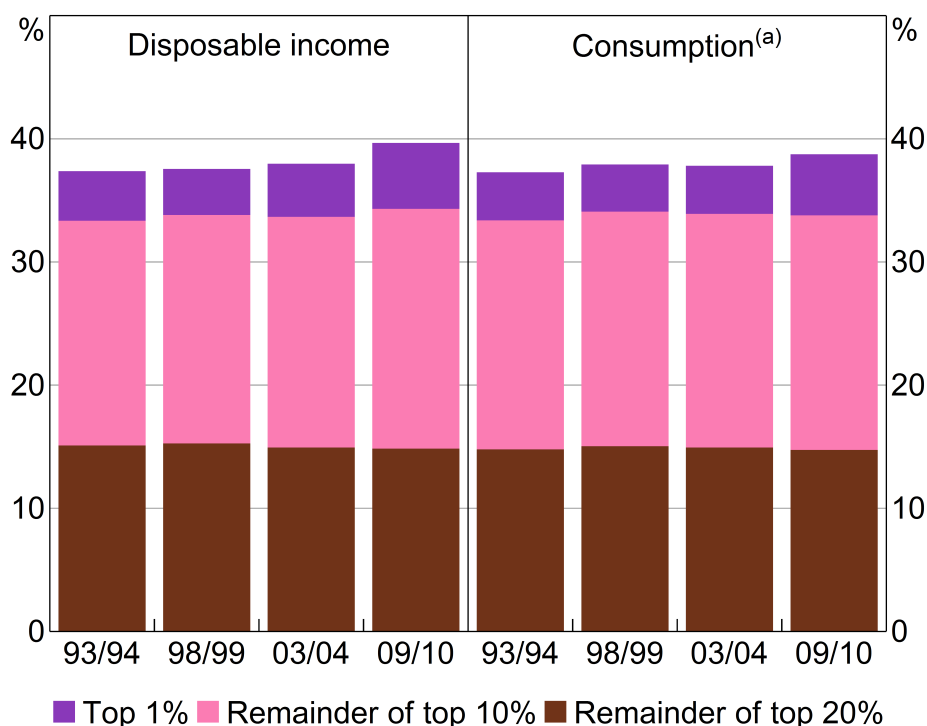
The Gini coefficient is a useful indicator for summarising distributions. However, it does not identify which parts of the distribution are responsible for any changes over time. It is also not a particularly intuitive measure of inequality. To complement the analysis, we examine how much of aggregate household income is earned by the high-income households (as a proxy for income inequality) and, similarly, how much of aggregate household consumption is accounted for by the high-spending households (for consumption inequality).

Based on the disposable income estimates, the top 10 per cent of income earners accounted for 22.3 per cent of aggregate household income in 1993/94 and 24.8 per cent in 2009/10 (Figure 3). Much of the increase in the share of income held by the top 10 per cent of earners has been due to the very highest earners within the top 1 per cent – their share of total disposable income rose from 4 per cent in 1993/94 to 5.4 per cent in 2009/10.¹⁰

Based on the consumption estimates, the top 10 per cent of spenders accounted for 22.5 per cent of aggregate household consumption in 1993/94 and 24 per cent in 2009/10. Again, a substantial part of the increase in the share of consumption for the top 10 per cent is due to the very biggest spenders within the top 1 per cent – their share of aggregate consumption rose from 3.9 per cent in 1993/94 to 5 per cent in 2009/10.

¹⁰ In Appendix C we provide similar estimates of income inequality using tax record data. The consistency between data sources suggests that HES survey data adequately captures top income earners.

Figure 3: Top Income Earners/Consumers
Share of total household income/consumption



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

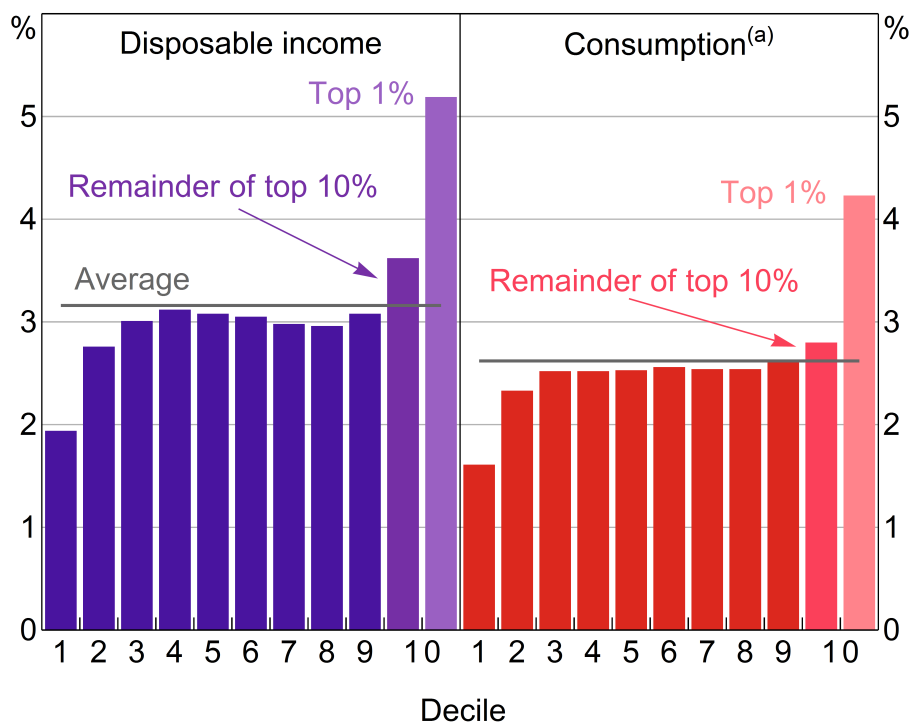
(a) Also excludes other interest payments

Sources: ABS; Authors' calculations

To further examine different parts of the income and consumption distributions, we break down each distribution into deciles and separate out the top 1 per cent. We then examine the relative growth in income and consumption for each decile and the top 1 per cent. Based on this, the top 1 per cent of earners experienced a relatively large increase in real disposable income of just over 5 per cent per annum between 1993/94 and 2009/10 (Figure 4). In contrast, the bottom 99 per cent of households experienced real annualised income growth of about 3 per cent, with the bottom decile experiencing much lower growth than the rest of the distribution. The top 1 per cent of spenders have also experienced faster growth in real consumption than other households over recent decades, though the difference in growth is less pronounced than in the case of income (Figure 4). More specifically, the top 1 per cent of spenders experienced real growth in consumption of a bit over 4 per cent. In contrast, households in the bottom 99 per cent of

the consumption distribution experienced real average annualised growth of about 2.5 per cent, with the bottom decile experiencing growth of just 1.6 per cent.

Figure 4: Real Consumption and Income Growth
By decile, average annualised, 1993/94–2009/10



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

(a) Also excludes other interest payments

Sources: ABS; Authors' calculations

An interpretation of the differing trends in income and consumption inequality is that some of the increase in income inequality has been due to an increase in the variance of transitory income shocks, which households have been able to smooth through borrowing and saving. This is consistent with the permanent income hypothesis, which postulates that consumers spend in line with their permanent income and borrow and save to offset temporary fluctuations in income. We explore the separate trends in permanent and transitory income in Section 5.

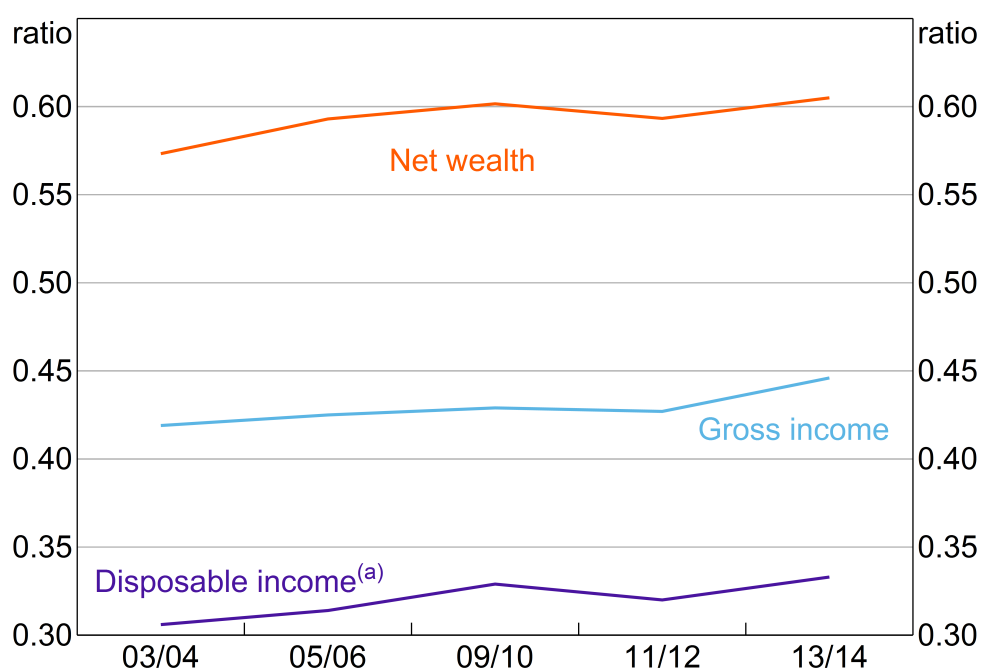
3.2 Wealth Inequality

We mainly focus on consumption and income inequality because we have long-run estimates for these measures of household wellbeing. But we also briefly consider

wealth inequality for two reasons. First, wealth is a potentially important indicator of wellbeing in its own right. Second, it highlights the important role of housing prices in affecting measured inequality. To develop the most complete picture we consider estimates of wealth inequality from both the HES and SIH.

Wealth is much more concentrated within the household population than either income or consumption. In 2013/14 the Gini coefficient for net wealth was 0.60, which was well above the level of inequality in disposable income (Figure 5). Other indicators of wealth inequality tell essentially the same story. For instance, the share of total net wealth held by the top 20 per cent of wealthy households was 62 per cent in 2013/14. This is 1.5 times the corresponding share of aggregate income for the top 20 per cent of income-earning households.¹¹ It is common in both Australia and other advanced economies to find that wealth is much more skewed towards the top of the distribution than either income or consumption.

Figure 5: Income and Wealth Inequality
Gini coefficient

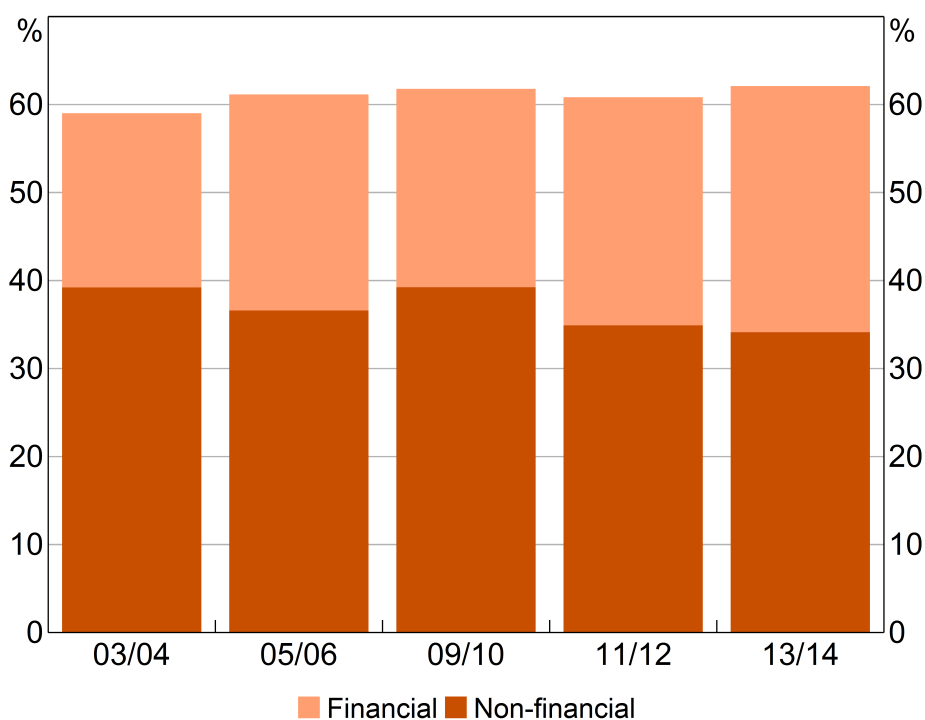


Note: (a) Equivalised and excludes net imputed rent for owner-occupied dwellings
Source: ABS

¹¹ The corresponding share of wealth held by the top 1 per cent was about 13 per cent.

Wealth inequality has risen modestly in Australia over the past decade or so by about the same amount as income inequality. The rise in wealth inequality has been due to changes at both the top and bottom of the distribution. Between 2003/04 and 2013/14, the net wealth of the top wealth quintile rose at an annualised rate of 3.3 per cent, while the net wealth of the bottom wealth quintile increased by an annualised rate of 1.1 per cent. Consistent with this, the share of total net wealth accounted for by the top wealth quintile rose from 59 per cent in 2003/04 to 62 per cent in 2013/14 (Figure 6).¹²

Figure 6: Wealth of Top 20% of Wealthy Households
Share of total household net wealth



Source: ABS

¹² The share of wealth held by the top 1 per cent of households follows a similar pattern, rising from 12 per cent in 2003/04 to 13.4 per cent in 2013/14.

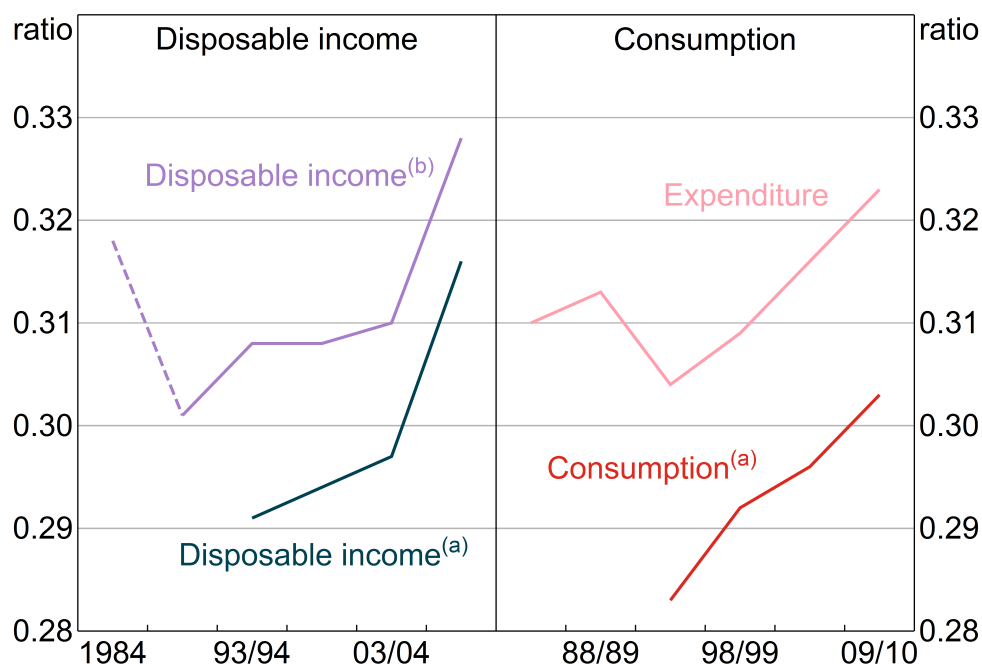
As in most advanced economies, housing wealth is the largest component of aggregate household wealth in Australia. Somewhat surprisingly though, changes in housing prices have not been the main determinants of changes in wealth inequality over the past decade. Instead, the rise in wealth inequality has been due to a rise in inequality in financial wealth and, in particular, an increase in the value of holdings of private superannuation and debt instruments (such as debentures and bonds) for the wealthiest households.

3.3 Housing Prices, Imputed Rent and Inequality Estimates

The estimates of household consumption and income inequality presented in the previous section include a service-flow equivalent of housing expenditure for owner-occupiers (or ‘net imputed rent’). It is worth discussing the adjustment for net imputed rent in detail as it has a significant effect on estimates of both the level and cross-sectional distribution of consumption and income in the economy.

The household surveys show that the inclusion of net imputed rent significantly reduces the level of inequality in both income and spending. Based on the Gini coefficient, the addition of net imputed rent reduces measured inequality in spending by a bit over 6 per cent, on average. This is shown by the fact that total consumption is more equally distributed across households than goods and services expenditure, on average (Figure 7). (Recall that household ‘consumption’ is the sum of household ‘expenditure’ and net imputed rent.) Similarly, the addition of net imputed rent reduces inequality in disposable income by just under 5 per cent on average.

Figure 7: Household Economic Inequality
Gini coefficient



Notes: All measures are population weighted and equivalised

(a) Includes net imputed rent for owner-occupied dwellings; consumption also excludes other interest payments

(b) Dashed line indicates that tax estimates are not imputed

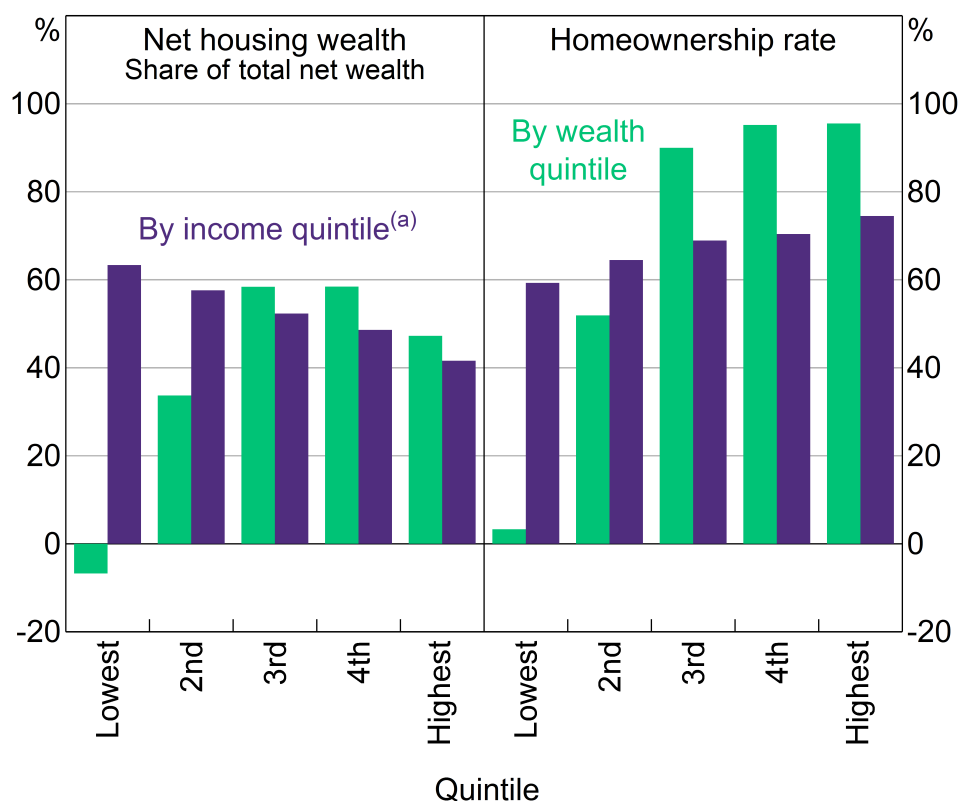
Sources: ABS; Authors' calculations

The inclusion of net imputed rent has an equalising effect as it disproportionately benefits low-income (and low-spending) households. This is because the home is typically the largest asset for these households, and as a result, the net imputed rent paid (and earned) on that asset is a relatively large fraction of the household's budget. For example, older (retired) households are likely to have a relatively low level of income (and spending), but a large proportion of these households own their own home outright and, therefore, adding net imputed rent to their measured income (and expenditure) significantly improves their welfare position. On the other hand, many high-income (and high-spending) households are comprised of young renters, meaning that the top of the distribution will not benefit to the same extent by the inclusion of imputed rent in measured income and expenditure.

Essentially, the equalising effect of imputed rent on income is due to the fact that low-income households are not the same as low-wealth households. Net housing

wealth (or housing equity) is a relatively large share of total net wealth for low-income households (Figure 8). In contrast, housing equity comprises a relatively low share of total net wealth for low-wealth households (Figure 8). Furthermore, the average rate of home ownership is higher amongst low-income households than amongst low-wealth households (Figure 8). About one-third of households in the lowest wealth quintile are young renter households that are likely to be credit constrained. But nearly half of the households in the bottom income quintile are older home owners.

Figure 8: Distribution of Housing Wealth
2013/14



Note: (a) Equivalised household disposable income
Sources: ABS; Authors' calculations

These differences between the income and wealth distributions also imply that changes in housing prices can have different effects on the trends in estimated income and wealth inequality. For instance, if we adjust for net imputed rent, a rise in housing prices can cause estimates of income (and consumption) inequality to fall, all other things being equal, as lower-income owner-occupier households benefit disproportionately from the higher housing prices. At the same time, higher

housing prices cause estimated wealth inequality to rise as wealthier households benefit disproportionately.

3.4 Consumption Inequality by Type of Household

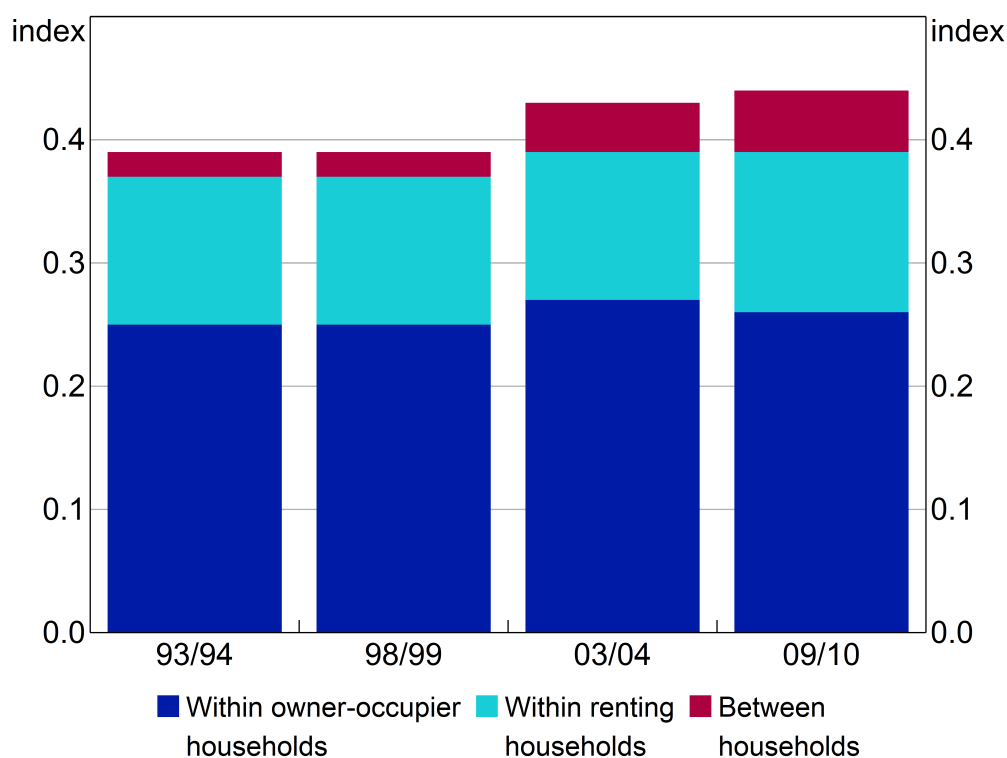
Given the important role of home ownership in affecting measured inequality, it is useful to decompose overall consumption inequality into inequality *within* both renting households and owner-occupier households as well as inequality *between* these two groups of households. We do this using analysis of variance (ANOVA); for each survey year, the observed variance in consumption is essentially partitioned into three components: 1) the variance in consumption for owner-occupiers ('within-owners'); 2) the variance in consumption for renters ('within-renters'); and 3) the difference between the average level of consumption for owner-occupiers and renters ('between owners and renters').¹³

Based on this, we find that much of the trend increase in consumption inequality over the past decade has been due to a rise in inequality between the two groups (Figure 9). In particular, owner-occupier households consume more goods and services than renter households, on average, and this has become increasingly true over time.

13 More specifically, we decompose the overall variance of log consumption at a point in time as:

$$V(\ln(C)) = \frac{\sum_i \sum_j (\ln(C_{ij}) - \overline{\ln(C)})^2}{N} = \frac{\sum_i \sum_j (\ln(C_{ij}) - \overline{\ln(C_j)})^2}{N_j} \left(\frac{N_j}{N}\right) + \frac{\sum_i \sum_j (\overline{\ln(C_j)} - \overline{\ln(C)})^2}{N}$$

where i denotes households and j denotes the number of household groups (in this case, there are two groups – renters and owner-occupiers). Also, $\overline{\ln(C_j)} = \sum_i \ln(C_{ij})/N_j$ denotes the average level of log consumption within each household group, $\overline{\ln(C)} = \sum_i \sum_j \ln(C_{ij})/N$ denotes the average level of log consumption across all households, N_j denotes the total number of households within each group while N denotes the total number of households. The first component on the right-hand side of the equation is the 'within-groups' component. The second component is the 'between-groups' component.

Figure 9: Decomposition of Consumption Inequality

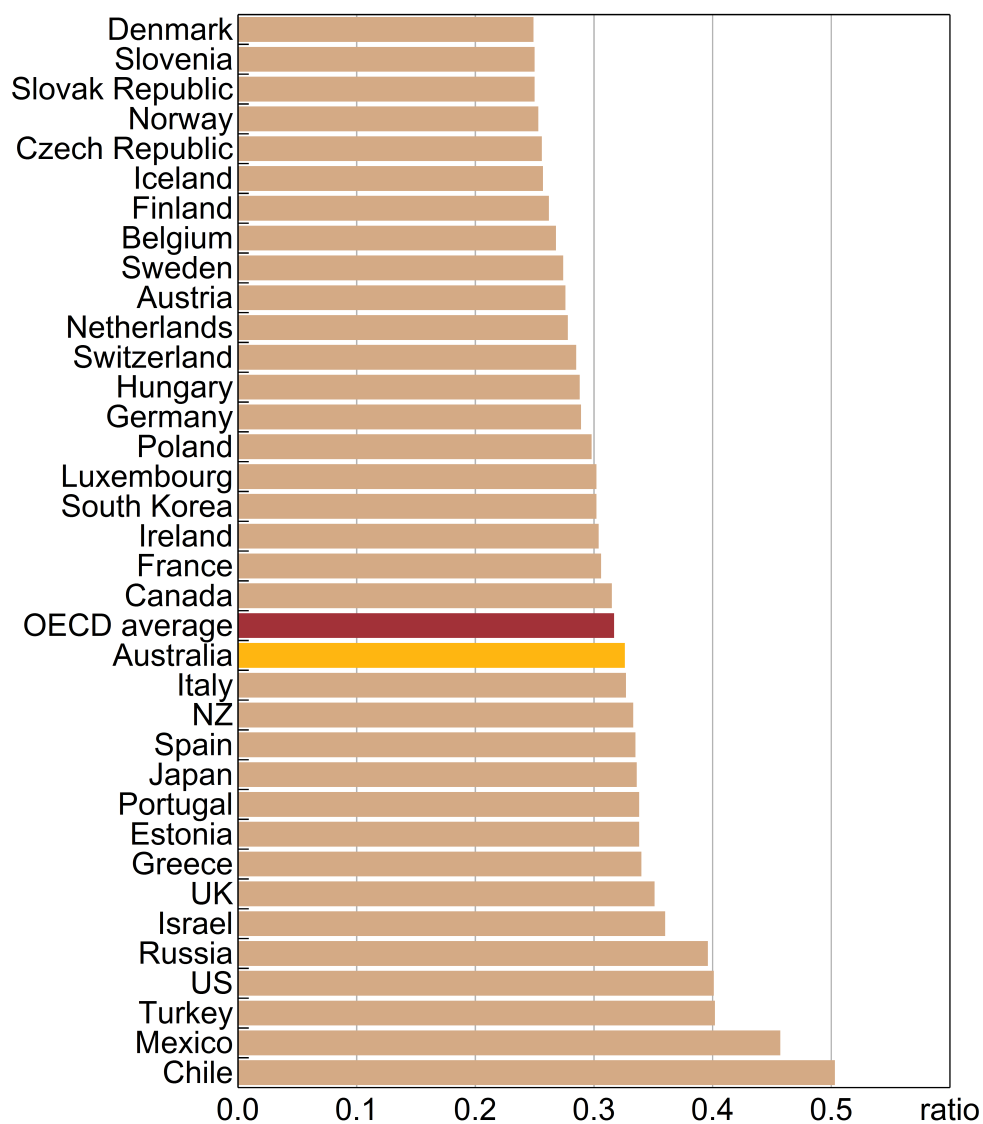
Notes: Decomposition of the variance of the logarithm of consumption; includes net imputed rent for owner-occupied dwellings and excludes other interest payments

Source: ABS; Authors' calculations

4. Cross-country Comparisons of Income and Consumption Inequality

We can compare the inequality estimates for Australia with corresponding estimates in other countries. The Organisation for Economic Co-operation and Development (OECD) has produced estimates of income inequality based on the Gini coefficient that allow for comparisons across countries. According to the latest estimates, the level of inequality in Australia is only slightly higher than the OECD average (Figure 10).

Figure 10: Gini Coefficients
Household disposable income

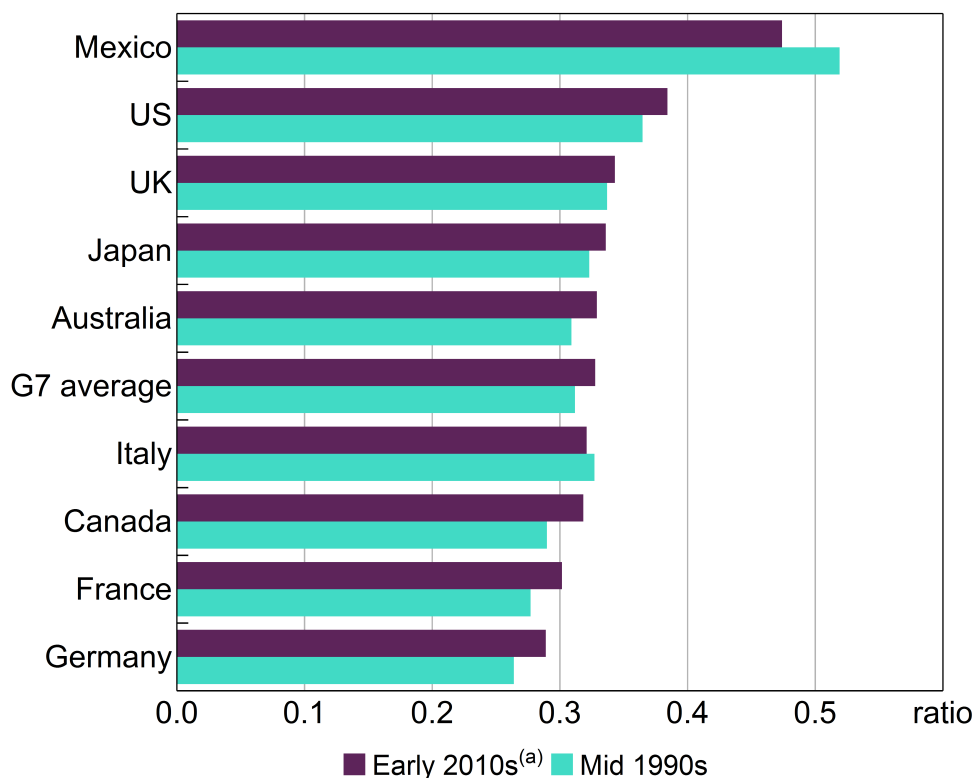


Note: Survey year underpinning estimates varies from 2009 to 2014

Sources: Authors' calculations; OECD

Compared with the average of the G7 countries, Australia appears to have experienced a similar increase in disposable income inequality between the mid 1990s and late 2000s. (Figure 11).

Figure 11: Gini Coefficients
Household disposable income, averages



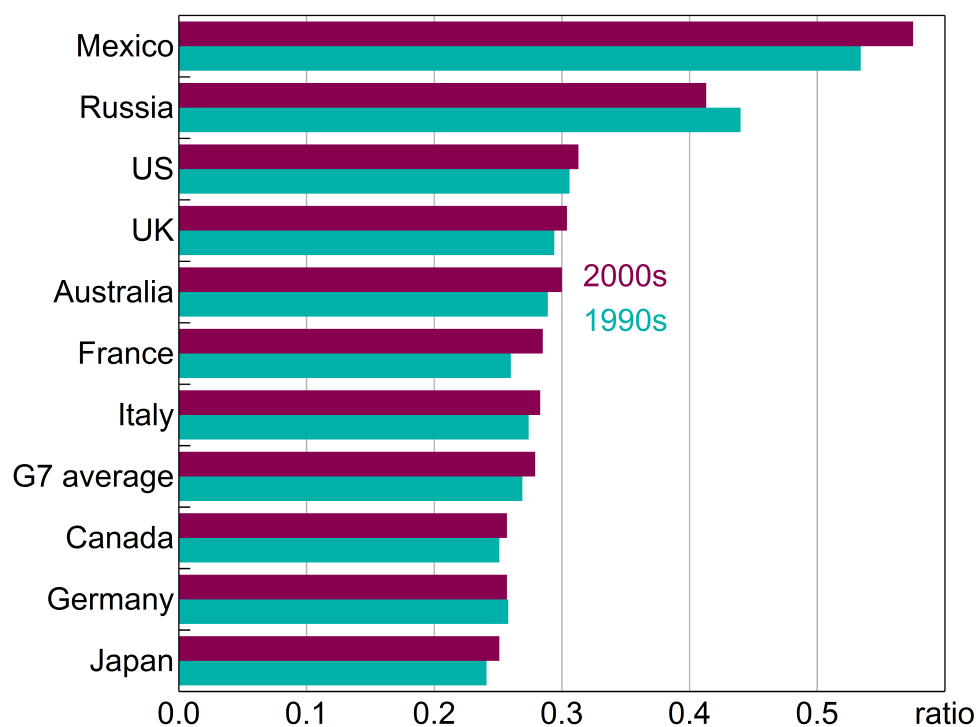
Note: (a) Estimate for Japan is based on 2009 data

Sources: Authors' calculations; OECD

It is more difficult to find comparable estimates of consumption inequality for other developed economies. However, Krueger *et al* (2010) provide a summary of some country-specific studies on consumption inequality. It is possible to use the results of these studies to construct average estimates of inequality for each of the past couple of decades. These estimates are shown, alongside Australia, in Figure 12. The results should be treated with caution as most of the studies

exclude durable goods from the consumption estimates, which leads to a more equal distribution and may affect the measured trends in inequality.¹⁴

Figure 12: Gini Coefficients
Household consumption, decade-level averages



Sources: Authors' calculations; Krueger *et al* (2010)

Notwithstanding this caveat, the results of these country-specific studies suggest that, as in the case of Australia, consumption inequality is typically lower and more stable over time than income inequality. These studies also imply that consumption inequality in Australia is higher than in countries such as Canada and Japan but lower than in the United States. Comparing the level of inequality

¹⁴ For more details on consumption inequality in developed countries, see the *Review of Economic Dynamics* special issue on 'Cross-Sectional Facts for Macroeconomists', which includes studies for the United States (Heathcote, Perri and Violante 2010), Canada (Brzozowski *et al* 2010), the United Kingdom (Blundell and Etheridge 2010), Germany (Fuchs-Schündeln, Krueger and Sommer 2010), Italy (Jappelli and Pistaferri 2010), Spain (Pijoan-Mas and Sánchez-Marcos 2010) and Sweden (Domeij and Flodén 2010). The evidence for US consumption inequality is particularly mixed, with some studies indicating that consumption inequality has been broadly unchanged (see, for example, Krueger and Perri (2005) and Meyer and Sullivan (2012)), while other studies suggest that it has risen in line with disposable income (see, for example, Aguiar and Bils (2011) and Attanasio *et al* (2014)).

in the 2000s to that in the 1990s, it appears that Australia experienced a slight increase in inequality very similar to that of the G7 countries.

5. Transitory and Persistent Income Inequality

According to Friedman's (1957) permanent income hypothesis, the distribution of household consumption should closely resemble the distribution of permanent income. So an alternative way to examine how household economic inequality has evolved, and to understand its welfare implications, is to explore whether changes in income inequality have been driven by persistent or transitory shocks to income.

The distinction between persistent and transitory income can be important for a couple of reasons (outlined in DeBacker *et al* (2013)). First, the distinction may help to understand the determinants of higher annual cross-sectional inequality. For example, if higher inequality is due to more persistent shocks to income, then potential explanations could include structural changes in the labour market and institutional changes that affect employers' remuneration policies. If, instead, higher inequality is due to temporary income fluctuations, then this could reflect changes in factors such as job mobility, workplace flexibility or the development of a bonus culture. Second, the distinction helps to inform welfare evaluations of changes in inequality. A change in income inequality that persists over time will have a larger welfare effect than a change in income inequality that is only temporary, especially if there are no constraints on households that prevent them from smoothing their consumption.

To separately identify the persistent and transitory income shocks driving inequality, we need to be able to track individual households over time. The HES surveys a different cross-section of households every time, so it is not useful for this. Instead, to explore the dynamics of household income, we turn to longitudinal household-level data from the HILDA Survey.

Our analysis takes two separate approaches to investigate household income dynamics. We first adopt an error components model to fully specify the process that generates income over time and decompose income into a highly persistent component and another transitory component that allows for some (limited) serial correlation. We then present estimates of income mobility as another way to consider dynamic changes in income using a household panel.

The HILDA Survey data cover the period from 2001 to 2013. The main income measure used is real annual household disposable income. This measure is population weighted and divided by an equivalence factor to control for household size and composition to make the estimates as consistent as possible with those obtained using the HES data. Households must be present for at least three consecutive years of the survey, and those with non-positive income and missing demographic information are excluded from the sample. The final sample consists of about 19 000 households and almost 100 000 observations.

Based on the HILDA Survey, there has been a trend increase in income inequality. Between 2001 and 2013, the share of aggregate disposable income held by the top 1 per cent of highest-earning households rose from 4.6 per cent to nearly 6 per cent. The top 10 per cent of earners saw their share of aggregate income rise from 22.7 per cent to 25 per cent over the same period. This is consistent with the trend increase in income inequality observed in the HES data over the 2000s. For most of the analysis in this section of the paper it is more useful to work with a measure of inequality based on the variance of household disposable income. The variance of income increased by about 8 log points between 2001 and 2013.

To quantify the extent to which the rise in income inequality is due to persistent and transitory factors, we first estimate the portion of inequality explained by *observed* differences across households using the following least squares regression:

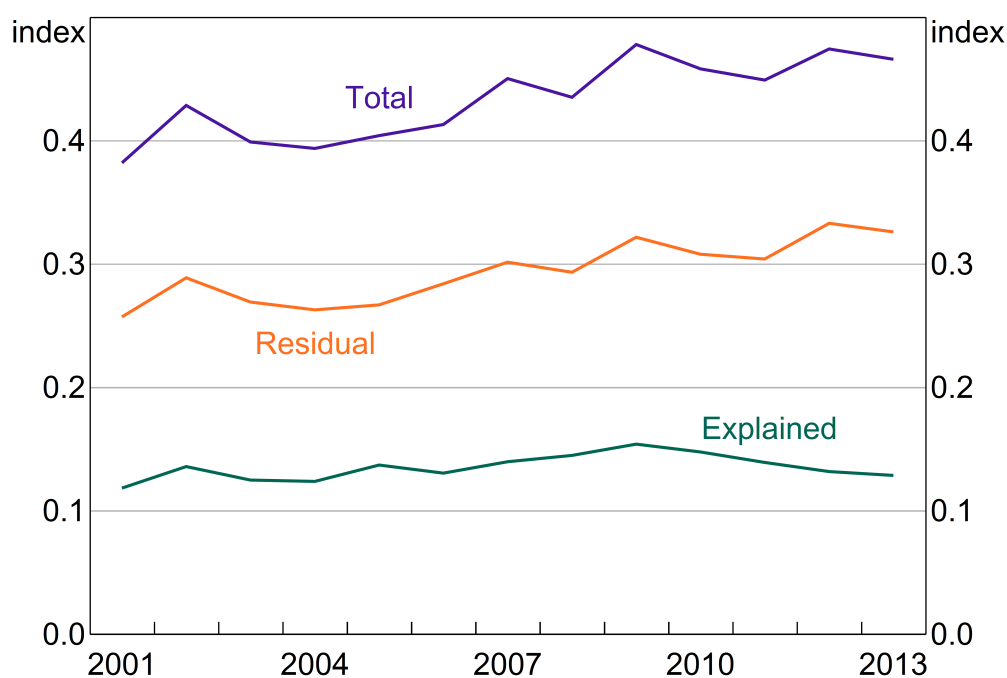
$$\ln(Y_{it}) = X_{it}'\beta_t + \mu_{it} \quad (1)$$

where the dependent variable is the log of equivalised household disposable income ($\ln(Y_{it})$) and the set of explanatory variables (X_{it}) includes the characteristics of the household head, such as level of education, gender, age, employment status, migrant status, indigenous status, and marital status. The specification also includes state, occupation and industry fixed effects, as well as interaction terms for occupation with industry and state variables. The regressions are estimated separately for each year in the sample.

We then take the estimated residuals ($\hat{\mu}_{it}$) from Equation (1) for each household i in year t and calculate the variance of these residuals each year. We plot this variance, which we label ‘residual’ income inequality, together with total and explained income inequality in Figure 13. We define ‘explained’ inequality as the variation in

income over time explained by the observable household characteristics included in the set of explanatory variables (X_{it}).

Figure 13: Household Income Inequality
Variance of the logarithm of household disposable income



Note: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

Sources: Authors' calculations; HILDA Release 13.0

Looking at changes over time, we find some evidence that the rise in income inequality between 2001 and 2009 was due to changes in observable characteristics. But most of the rise in inequality cannot be explained by observed inequality. Instead, on average, residual inequality accounts for about 70 per cent of total income inequality. Residual inequality bears a close resemblance to total inequality, with the two estimates displaying similar upward trends and short-term fluctuations. This suggests that observable factors, such as an ageing population and rising educational attainment, have played limited roles in explaining changes in inequality over the past decade and the unobserved, dynamic component of income has been the main determinant of rising income inequality.

5.1 Error Components Model

We next use an error components model (ECM) to decompose residual income inequality. The ECM has been standard in the inequality literature since Gottschalk and Moffitt (1994). Many international studies have used this model to estimate the dynamics of inequality over time, particularly for the United States using the Panel Study of Income Dynamics Survey.¹⁵ To date, this approach has not been possible for Australia due to the absence of a sufficiently long panel dataset.

To address this, we use the panel structure of the HILDA Survey and a flexible specification of the ECM to examine the dynamics of household income inequality for Australia.

As before, the residual of log equivalised disposable income for household i in year t is estimated from the regression described by Equation (1). The dynamics of the residual are then modelled by the following process:

$$\begin{aligned}\mu_{it} &= \lambda_t \alpha_i + z_{it} + v_{it} \\ z_{it} &= \rho z_{it-1} + \eta_{it} \\ v_{it} &= \varepsilon_{it} + \theta \varepsilon_{it-1}\end{aligned}\tag{2}$$

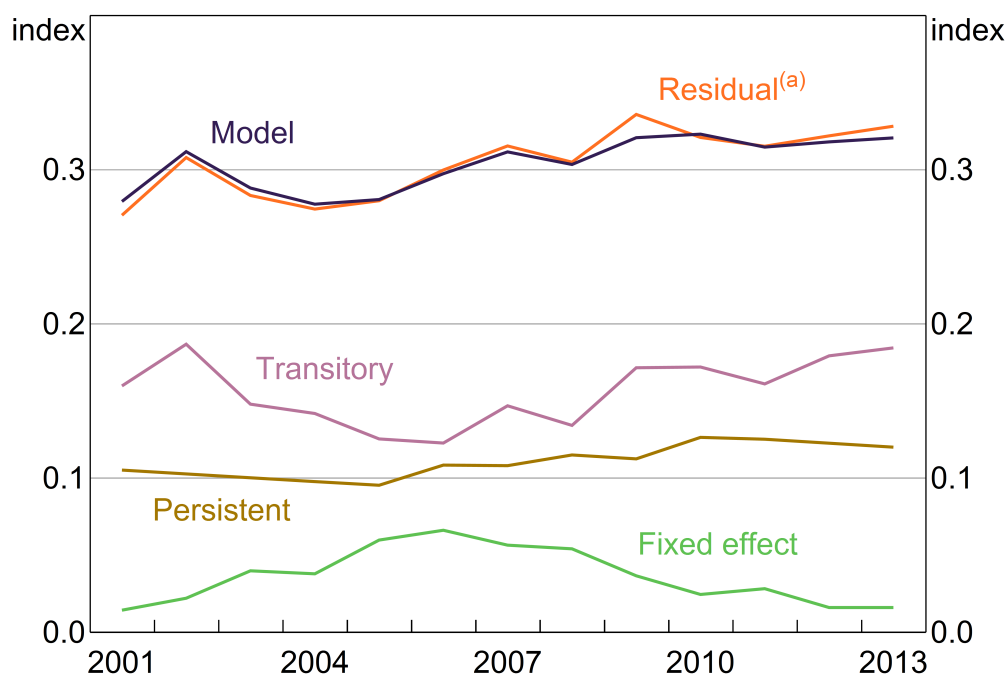
where the ‘persistent’ component of inequality is a combination of a household fixed effect (α_i) that has a time-varying coefficient (λ_t), with total variance of $\lambda_t^2 \sigma_\alpha^2$, and a highly persistent term (z_{it}) that follows an autoregressive AR(1) process and has variance of σ_z^2 . The household fixed effect captures unobserved time-invariant factors such as skill or ability (i.e. human capital). The time-varying coefficient captures the ‘market price’ for human capital. The AR(1) term (z_{it}) captures other shocks to income that persist over time, such as promotions that affect the level of wage income or possibly a long-term health condition. The temporary component (v_{it}) is specified as a moving average MA(1) process. This specification allows temporary income factors, such as lay-offs and bonuses, to have effects that persist for more than a year. This is motivated by empirical observation of the autocovariances of household income in the HILDA data.

¹⁵ Prominent examples include MaCurdy (1982), Abowd and Card (1989), Blundell, Pistaferri and Preston (2008) and Heathcote *et al* (2010).

The η_{it} and ε_{it} terms are the respective persistent and transitory shocks to income with mean zero and time-varying variances, $\sigma_{\eta t}^2$ and $\sigma_{\varepsilon t}^2$, respectively. These shocks are assumed to be uncorrelated with each other and independently and identically distributed across households over time. Under this error scheme, changes in residual inequality can be driven by changes in three different factors: 1) the variance of persistent shocks; 2) the variance of temporary shocks; and/or 3) changes in the market price of a household's fixed human capital.

We estimate the variance-covariance matrix of the model using generalised method of moments. This procedure essentially estimates a parameter vector of interest by minimising the weighted sum of squared distances between the population moments implied by the model and their empirical counterparts. The parameter vector is then used to construct estimates of the variances of transitory and persistent shocks to income, and the variance of the return to fixed human capital, shown in Figure 14.

Figure 14: Household Income Inequality
Variance of the logarithm of household disposable income



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

(a) Equivalent to residual income inequality shown in Figure 13

Sources: Authors' calculations; HILDA Release 13.0

Our preferred estimates indicate that about two-thirds of the level of residual inequality is due to the variance in transitory income shocks. Given that about 70 per cent of the variation in total income across households is due to unobserved characteristics, this suggests that temporary shocks explain close to one-half of the total cross-sectional variation in household income. The remaining variation in income across households is mainly due to variation in persistent shocks, though some inequality is also explained by variation in unobserved fixed human capital.

The model also indicates that the first half of the sample period (2001 to 2006) was characterised by a slight decline in transitory income inequality and a small (and largely offsetting) increase in persistent income inequality (due to an increase in the variance of the fixed effect). This appears to reflect developments in the Australian labour market over that period. In particular, the unemployment rate fell noticeably between 2001 and 2006 and this is likely to have disproportionately benefited lower-income workers who may be more exposed to temporary income shocks. The increase in the variance of the fixed effect in the early to mid 2000s reflects a rise in the ‘price’ that the market was willing to pay for unobserved ability, which may also be due to the relatively strong labour market at the time.

The trend in overall income inequality in the latter half of the sample period appears to reflect an increase in the variance of both transitory and persistent income shocks. There also appears to be a slight jump in transitory income inequality around the time of the global financial crisis. This suggests that the crisis had a limited effect on the distribution of consumption across households. In general, households are more able to insulate their consumption from transitory rather than persistent shocks to income. The slight rise in the variance of persistent income shocks since the middle of the 2000s is consistent with the small increase in consumption inequality reported in the HES over a similar period.

5.2 Income Mobility

To further quantify the extent to which the trends in inequality are persistent we next estimate the degree of mobility in the income distribution. Income mobility has a direct bearing on the degree of persistence in inequality. For example, if household income is relatively immobile and the same households are ranked as high-income from one year to the next, then this suggests that the inequality is persistent. In contrast, if household income is fairly mobile on average then

high-income households may move down the income rankings the following year, suggesting that the inequality is temporary.

To examine mobility, we divide the sample into quintiles based on the residual income estimates. We then estimate the share of households that move up, down or stay in the same quintile over time. We do this for both 1-year and 10-year windows to measure short-term and long-term mobility.

Table 1: Transition Matrix for Household Disposable Income
Per cent

Income quintile	From t to $t+1$			From t to $t+10$		
	Down	Same	Up	Down	Same	Up
Lowest		58	42		46	54
2nd	23	38	39	24	27	49
3rd	33	34	32	38	23	39
4th	40	37	23	50	26	25
Highest	41	59		55	45	

Sources: Authors' calculations; HILDA Release 13.0

Just over half of households that are either in the top or bottom income quintile remain in that same quintile from one year to the next (Table 1). More notably, a bit over 40 per cent of these households are in the same quintile a decade later.¹⁶ For households in the middle income quintiles, about a third change quintiles each year and about a quarter are in the same quintile a decade later. This evidence suggests that both persistent and transitory movements are occurring within the income distribution over time, though there is some variation across income quintiles.¹⁷

To directly quantify the extent to which the observed mobility reflects permanent or temporary transitions we use a key indicator of income mobility – the Shorrocks R index (Shorrocks 1978). This index provides a direct link between mobility and the relative contribution of persistent and transitory inequality by defining

¹⁶ The results reported in Table 1 are robust to using a balanced panel and a sample of working-age households.

¹⁷ The skew in the income distribution could account for some of the differences in estimated mobility across the income quintiles. The positive skew implies that there is greater dispersion in income within the top income quintile than in any of the lower quintiles. This implies, for example, that a negative income shock of a particular size will be more likely to cause a given household to fall out of the second highest quintile than out of the highest quintile, all other things being equal.

immobility as the ratio of persistent inequality to average annual or total inequality over the same period. The Shorrocks R index is defined as:

$$R = \frac{I[Y]}{\sum_{t=1}^T w_t I[Y_t]} \quad (3)$$

where the numerator, $I[Y]$, is a multi-year inequality value estimated from household incomes aggregated over T years; the denominator is the weighted average of single-year inequality values, $I[Y_t]$, over the T -year period. The weight assigned to each year, $w_t = \frac{Y_t}{Y}$, is the ratio of average household income in year t (Y_t) to average total household income (Y) earned over the entire period.

The Shorrocks R index reflects the relative contribution of persistent to total income inequality over time. The index can take values between 0 and 1. The higher the value of the index, the higher the share of persistent or long-term inequality, and the less mobility there is in the income distribution.

To calculate the index, it is necessary to use an inequality measure, $I[Y]$, that is a strictly convex function (Shorrocks 1978). We use two common inequality indices that meet this condition to estimate mobility for both total and residual household disposable income – the Theil index and the Gini coefficient.¹⁸

To assess whether the share of persistent inequality (or income immobility) has changed over time, we divide the HILDA panel into moving 5-year windows and estimate the Shorrocks R value for each sample window.

Based on the Theil index, the Shorrocks R value remains relatively stable over the nine 5-year windows spanning 2001 to 2013, with average R values for total and residual household disposable income of 0.81 and 0.70, respectively (Table 2).¹⁹

18 The Theil index is defined as $H = \frac{1}{N} \sum_{i=1}^N \frac{x_i}{\bar{x}} \log\left(\frac{x_i}{\bar{x}}\right)$ where x_i is the income of household i and \bar{x} is mean household income. This index measures the distance the population is away from perfect equality. If all households have the same income, then the index is equal to 0, signalling perfect equality. If one household has all the income, then the index is equal to 1, implying perfect inequality. The Gini coefficient is defined in Section 3.1.

19 The estimated degree of persistent inequality appears to be lower for Australia than for some other advanced economies based on the Theil index. For instance, Bayaz-Ozturk, Burkhauser and Couch (2014) report Shorrocks R values of 0.83 and 0.85 for the United States and Germany over the late 1990s to early 2000s.

Table 2: Shorrocks R Values for Household Disposable Income

Year	Theil index		Gini coefficient	
	Total	Residual ^(a)	Total	Residual ^(a)
2001–2005	0.80	0.70	0.91	0.85
2002–2006	0.79	0.70	0.91	0.85
2003–2007	0.79	0.69	0.91	0.85
2004–2008	0.78	0.69	0.91	0.84
2005–2009	0.79	0.70	0.91	0.85
2006–2010	0.84	0.71	0.93	0.86
2007–2011	0.87	0.74	0.94	0.87
2008–2012	0.83	0.72	0.93	0.87
2009–2013	0.79	0.68	0.91	0.85

Note: (a) Residual income estimated using Equation (1)

Sources: Authors' calculations; HILDA Release 13.0

However, we do observe a slight rise in the share of persistent inequality or immobility from the middle of the 2000s onwards, with higher R values reported for the second half of the sample period.²⁰

This increase in persistent inequality is consistent with the rise in the variance of persistent shocks observed in Figure 14 over the same period, as well as the small increase in consumption inequality reported in the HES between 2003/04 and 2009/10.

6. Conclusion

We document some new facts about economic inequality among households in Australia over recent decades. We find that consumption inequality is lower on average than income inequality due to the ability of households to smooth consumption by borrowing and saving. Income and consumption inequality have both increased a little since the early 1990s, but income inequality has risen by slightly more. These findings are in line with the changes in income and consumption inequality documented in other developed economies.

We also provide new estimates of household-level income dynamics for Australia. The broad trends in consumption and income inequality do not appear to be due

²⁰ The same pattern is observed for the Gini coefficient, although the Gini coefficient typically reports higher values for R than other inequality measures that place weight on the extremes of the income distribution (Jarvis and Jenkins 1998).

to changes in observed household characteristics, but rather to changes in the distribution of unobserved shocks. The increase in income inequality over the past decade has reflected similar-sized increases in the variance of transitory and persistent income shocks. The rise in persistent income inequality since the middle of the 2000s is consistent with the rise in consumption inequality over the same period.

Our results also suggest that monetary policy can affect inequality to the extent that changes in interest rates influence asset prices, and such interest-sensitive assets are not distributed equally across the household population. But our results also indicate that the direction of these effects is unclear *a priori*. For instance, lower interest rates may lead to higher housing prices which, in turn, boost wealth inequality given that wealthy households are more likely to own their homes. But if income and spending are adjusted to account for imputed rent, our results also imply that lower interest rates could boost imputed rent (relative to market rents) and disproportionately benefit relatively low-income home owners, reducing measured inequality in income and consumption, at least in the short term.

Appendix A: Imputed Rent and the Distributions of Household Expenditure and Income

To understand how the addition of imputed rent affects the measured distributions of consumption and income it helps to consider some simple accounting exercises.

Suppose a renter has the following single period budget constraint:

$$E + R = (Y - T) - S$$

Where total expenditure consists of non-housing expenditure (E) and expenditure on housing services which, for the renter, is simply the value of market rent paid (R). The total resources available each period consist of non-housing income (Y), such as wages and salaries, less net government taxes (T) and saving (S).

A home owner has a corresponding budget constraint:

$$E + M = (Y - T) - S$$

This expression is essentially the same as that of the renter, except that the housing expenditure of the home owner is given by mortgage interest payments (and other costs of maintaining a home) (M). Now suppose we add gross imputed rent (IR) to both sides of the home owner's budget constraint (and just move mortgage interest payments to the right of the constraint):

$$E + IR = (Y - T) + (IR - M) - S$$

Now the home owner spends an amount equal to gross imputed rent on housing services and earns 'housing income' measured by net imputed rent (which is just gross imputed rent less mortgage interest payments) ($NIR = IR - M$). The addition of gross imputed rent to both sides of the household budget constraint does not affect *aggregate* household saving. But it affects the distribution of both spending and income if home owners are not identical to renters.²¹

In essence, we add *net* imputed rent to non-housing expenditure when we redefine household 'expenditure' as 'consumption'. In other words, we replace mortgage

²¹ We could also write down the budget constraint for a housing investor but it does not affect the overall story.

interest payments (and associated housing costs, such as maintenance) with gross imputed rent. This causes measured inequality to fall because it affects low-spending households by more than high-spending households.

To demonstrate how consumption and expenditure inequality can follow different trends over time, consider the following simple decomposition.

Suppose that the level of imputed rent (IR) is simply given by the imputed rental yield (r) multiplied by the housing price (HP):

$$IR = r * HP$$

Correspondingly, the level of mortgage interest payments (M) is equal to the mortgage interest rate (i) multiplied by the outstanding stock of mortgage debt (D):

$$M = i * D$$

It follows that we can write net imputed rent (as a share of income) as:

$$NIR/(Y + NIR) = (r - i * D)/HP * (HP/(Y + NIR))$$

This decomposition shows that the distribution of net imputed rent (relative to total income) is a function of the distribution of three factors:

1. The ‘spread’ between the imputed rental yield and the mortgage interest rate ($r - i$)
2. The housing debt-to-price (or ‘leverage’) ratio (D/HP)
3. The house price-to-income ratio ($HP/(Y + NIR)$).

Changes in any of these factors can drive differential changes in measures of inequality based on consumption and expenditure over time. Furthermore, they drive changes in measures of income, depending on whether the estimates adjust for net imputed rent or not.

Appendix B: Alternative Survey Measures of Inequality

In this Appendix we adopt an alternative definition of imputed rent to cross-check the results for consumption and income inequality in the paper.

B.1 Hedonic Estimates of Imputed Rent

The alternative measure of imputed rent is calculated using the ABS hedonic modelling methodology, which estimates the market value of the rental equivalent for owner-occupied dwellings using the dwelling characteristics available in the HES.

The dependent variable is the logarithm of gross imputed rent for the sample of owner-occupier households.²² There are two types of explanatory variables in the model – Type I and Type II variables.

The Type I variables describe the characteristics of the dwelling, including:²³

- State
- Section of state (i.e. capital city or balance of state)
- Type of dwelling structure
- Number of bedrooms.

The Type II variables describe the households renting the dwellings, including:²⁴

- Household income
- Type of landlord (i.e. public or private landlord).

²² The explanatory variables used to estimate the market rent are not as extensive as those used by the ABS as we only have access to the basic confidentialised unit record files.

²³ Prior to the 2003/04 survey, the ‘section of state’ variable is not included. The ‘state’ variable is not included in the 1988/89 survey.

²⁴ Prior to the 2003/04 survey, landlord type is not included so we use tenure type to capture similar information.

The model for market renters is:

$$\ln(R_i) = \alpha_0 + \sum_{j=1}^J \beta_j X_{ij} + \sum_{k=1}^K \delta_k Z_{ik} + \phi M_i + \varepsilon_i$$

where $\ln(R_i)$ is the natural logarithm of the weekly rent of household i , X_{ij} is the set of Type I variables and Z_{ik} is the set of Type II variables for each household, M_i is the estimated Inverse Mills ratio and ε_i is a normally distributed error term with a mean of zero and a standard deviation of σ .

The Inverse Mills ratio is calculated using the Heckman procedure. This adjusts for possible bias that could result from non-random selection. The Inverse Mills ratio is the ratio of the probability distribution function and the cumulative distribution function of the standard Z-score estimated in the probit model:

$$P(i \in \text{'renter'}) = \Phi(X\beta)$$

where X is a set of dwelling and household characteristics (e.g. age and educational attainment of household head) that determine whether a household rents or not.²⁵

Next, we make an adjustment to the intercept to control for the Type II variables. Specifically, the intercept is adjusted to the mean for renters:

$$\hat{\alpha}_0^{adj} = \hat{\alpha}_0 + \sum_{k=1}^K \hat{\delta}_k \bar{Z}_k$$

where $\hat{\alpha}_0^{adj}$ is the adjusted intercept for imputation, $\hat{\alpha}_0$ is the intercept estimate of the basic model, $\hat{\delta}_k$ is the estimated coefficient for each Type II variable, and \bar{Z}_k is the mean of each Type II variable.

A scaling factor is then applied to preserve the relationship between the observed and model rent estimates for private market renters. The imputed rent distribution is re-positioned to the original median rent observed in the private market renters' data.

$$ScalingFactor = median(R) - median(\hat{R}^{renter})$$

²⁵ Prior to 1998/99, the education variable is not included.

and so:

$$\hat{R}_i^{OODadj} = \hat{R}_i^{OOD} + ScalingFactor$$

where \hat{R}_i^{OOD} is the estimated gross imputed rent for owner-occupier household i .

Since the model does not fully explain the variation in rent for high-value dwellings, an extrapolation method is used to adjust gross imputed rent for high-value dwellings. A ceiling rent is determined by a visual inspection of the modelled results. The ceiling rent is divided by the estimated annual average rate of return for all owners. A regression model is fitted to owners below the house value cut-off:

$$r_i = \theta_0 + \theta_1 \frac{1}{h_i} + \theta_2 \left(\frac{1}{h_i}\right)^2 + \varepsilon_i$$

where $r_i = \frac{\hat{R}_i^{OODadj}}{h_i}$ is the rental rate of return for owner-occupier household i and h_i is the value of that household's dwelling.

The gross imputed rent for owners with house values above the cut-off were recalculated using the statistically significant coefficients in the above equation. The extrapolated imputed rent is given by:²⁶

$$\bar{R}_j^{OODhighvalue} = p_j \hat{r}_j$$

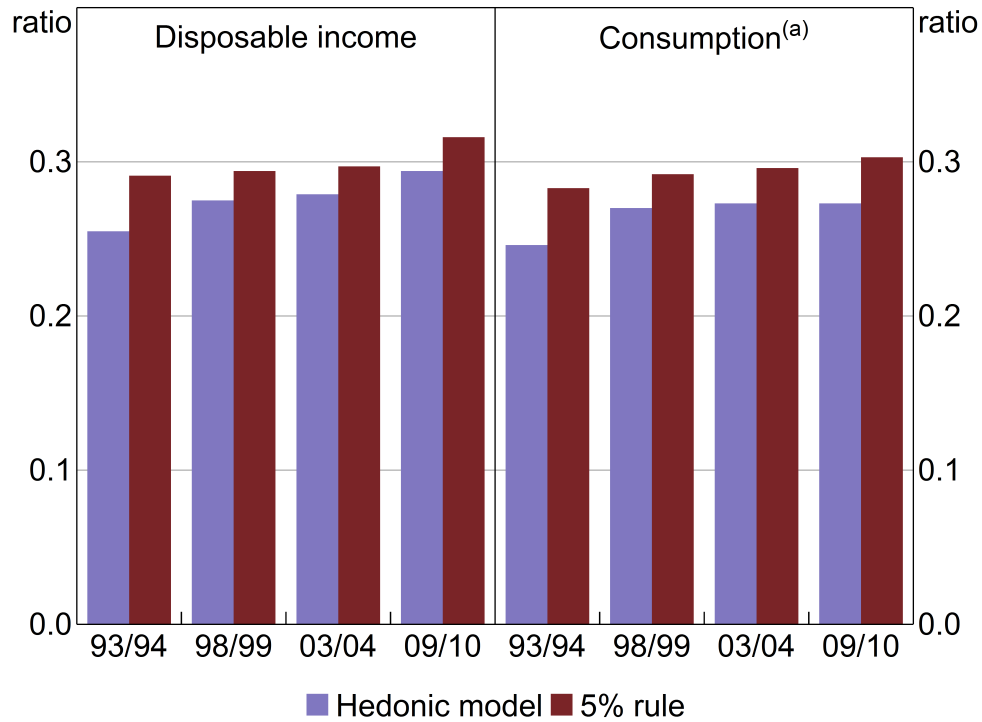
where $\bar{R}_j^{OODhighvalue}$ is the adjusted gross imputed rent for the j th owner-occupied dwelling, p_j is the j th home value and \hat{r}_j is the estimated rental rate of return. For further details, see Australian Bureau of Statistics (2008).

B.2 Alternative Estimates of Inequality

Using the alternative measure of imputed rent results in a lower level of inequality for both income and consumption (as measured by the Gini coefficient) compared with the simple approach in the main paper (Figure B1). The upward trend for disposable income inequality is apparent under both approaches. However, the rise in consumption inequality is less apparent under the hedonic modelling approach.

²⁶ We are not able to use this extrapolation method for the 1975/76, 1984 and 1988/89 surveys because the estimated housing value of the dwelling is not included.

Figure B1: Household Economic Inequality
Gini coefficients



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings

(a) Also excludes other interest payments

Sources: ABS; Authors' calculations

Appendix C: Income Inequality Estimates Using Tax Records

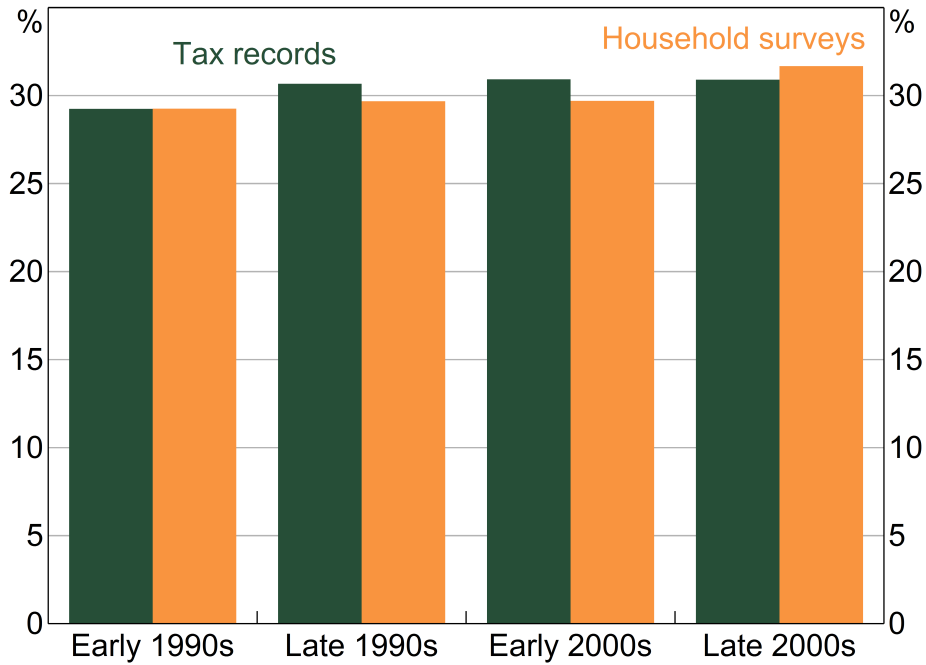
Income tax records provide an alternative way of measuring income inequality. These estimates are useful because tax records typically cover a broader population of individuals than household surveys (and hence better capture top income earners). There is also a long and consistent history of tax records, providing us with a sense of long-run trends. The main limitations of tax data relative to household surveys are that we generally cannot examine the underlying characteristics that are driving any changes over time and we can only measure inequality in *individual* incomes.

Estimates of individual income inequality based on tax records are very similar to those based on the SIH data (Figure C1). Since the early 1990s, relative to the tax records, the survey data has underestimated the share of income going to the top 10 per cent of earners by about 0.4 percentage points, on average. This suggests that the analysis based on the survey data is reasonably representative of the broader population of individuals covered by the tax data. Moreover, the consistency over time between the two series suggests that the survey data captures genuine trends in income inequality.²⁷

Using the Australian tax data, Leigh (2013) shows that the top shares in gross income follow a U-shape, decreasing from the early 1940s to the 1980s and then increasing after this period (Figure C2). To put this in an international context, cross-country comparisons can be made using the tax data collected by Atkinson and Morelli (2014). It is clear from Figure C3 that the U-shape pattern is common to several developed Anglo-Saxon countries, such as the United States, the United Kingdom, New Zealand and Canada (Atkinson and Morelli 2014). In contrast, the share of aggregate income accruing to the highest income earners has remained fairly steady since the 1950s in countries such as France, Germany and Japan (Piketty 2014).

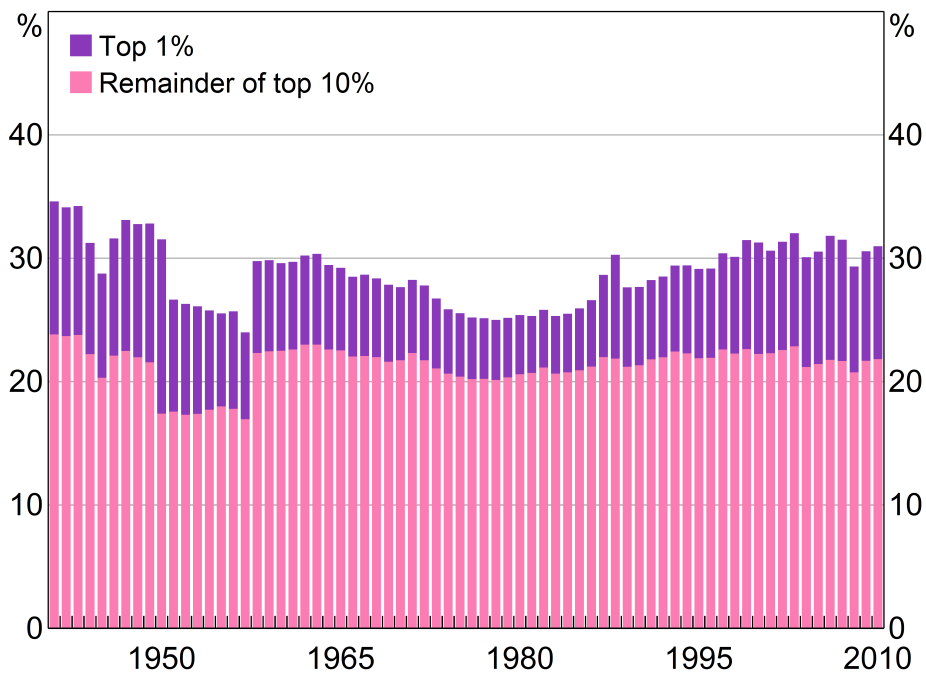
²⁷ For a more complete discussion of the uses and limitations of tax and survey data, see Wilkins (2013).

Figure C1: Income of Top 10% of Income Earners
Share of total gross income



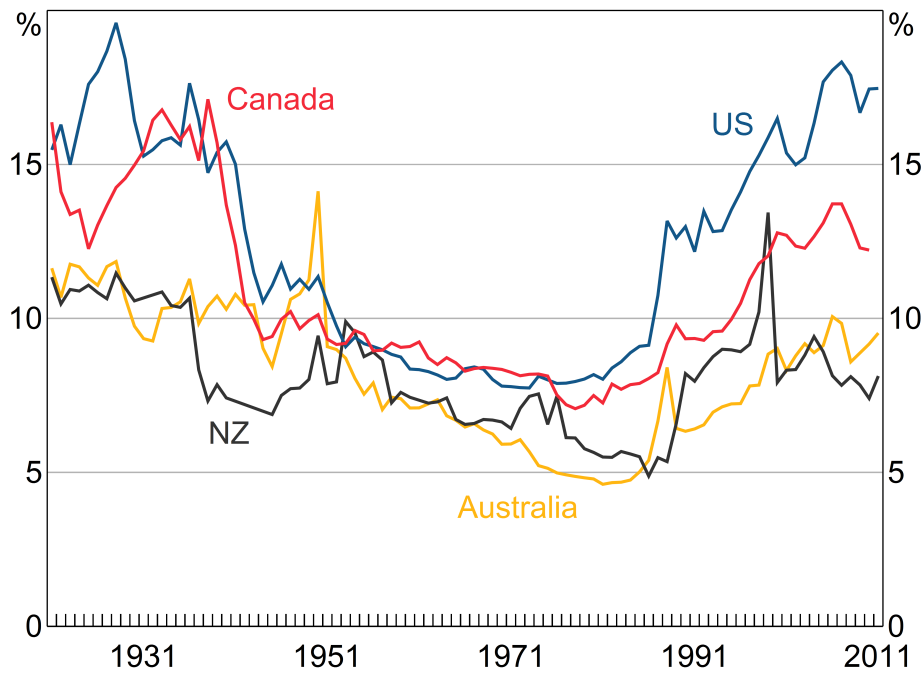
Sources: ABS; ATO; Authors' calculations

Figure C2: Income of Highest Earners
Share of total gross income, financial year



Sources: ATO; Authors' calculations; www.andrewleigh.org

Figure C3: Income of Top 1% of Income Earners
Share of total gross income



Note: Estimates based on income taxation statistics

Source: Atkinson and Morelli (2014)

References

Abowd JM and D Card (1989), ‘On the Covariance Structure of Earnings and Hours Changes’, *Econometrica*, 57(2), pp 411–445.

Aguiar MA and M Bils (2011), ‘Has Consumption Inequality Mirrored Income Inequality?’, NBER Working Paper No 16807.

Atkinson AB and S Morelli (2014), ‘Chartbook of Economic Inequality’, Society for the Study of Economic Inequality Working Paper ECINEQ WP 2014 - 324.

Attanasio O, E Hurst and L Pistaferri (2014), ‘The Evolution of Income, Consumption, and Leisure Inequality in the United States, 1980–2010’, in CD Carroll, TF Crossley and J Sabelhaus (eds), *Improving the Measurement of Consumer Expenditures*, NBER Studies in Income and Wealth, Vol 74, University of Chicago Press, Chicago, pp 100–140.

Australian Bureau of Statistics (2008), ‘Experimental Estimates of Imputed Rent, Australia, 2003-04 and 2005-06’, ABS Cat No 6525.0.

Barrett G, TF Crossley and C Worswick (1999), ‘Consumption and Income Inequality in Australia’, Australian National University Centre for Economic Policy Research Discussion Paper No 404.

Bayaz-Ozturk G, RV Burkhauser and KA Couch (2014), ‘Consolidating the Evidence on Income Mobility in the Western States of Germany and the United States from 1984 to 2006’, *Economic Inquiry*, 52(1), pp 431–443.

Beech A, R Dollman, R Finlay and G La Cava (2014), ‘The Distribution of Household Spending in Australia’, *RBA Bulletin*, March, pp 13–22.

Blundell R and B Etheridge (2010), ‘Consumption, Income and Earnings Inequality in Britain’, *Review of Economic Dynamics*, 13(1), pp 76–102.

Blundell R, L Pistaferri and I Preston (2008), ‘Consumption Inequality and Partial Insurance’, *The American Economic Review*, 98(5), pp 1887–1921.

Bray JR (2014), ‘Changes in Inequality in Australia and the Redistributive Impacts of Taxes and Government Benefits’, in A Podger and D Trewin (eds), *Measuring and Promoting Wellbeing: How Important is Economic Growth?: Essays in Honour of Ian Castles AO and a Selection of Castle’s Papers*, Australia and New Zealand School of Government Monograph, ANU Press, Canberra, pp 423–475.

Brzozowski M, M Gervais, P Klein and M Suzuki (2010), ‘Consumption, Income, and Wealth Inequality in Canada’, *Review of Economic Dynamics*, 13(1), pp 52–75.

DeBacker J, B Heim, V Panousi, S Ramnath and I Vidangos (2013), ‘Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns’, *Brookings Papers on Economic Activity*, Spring, pp 67–122.

Domeij D and M Flodén (2010), ‘Inequality Trends in Sweden 1978–2004’, *Review of Economic Dynamics*, 13(1), pp 179–208.

Fletcher M and B Guttman (2013), ‘Income Inequality in Australia’, *Economic Roundup*, 2, pp 35–54.

Friedman M (1957), *A Theory of the Consumption Function*, NBER General Series, No 63, Princeton University Press, Princeton.

Fuchs-Schündeln N, D Krueger and M Sommer (2010), ‘Inequality Trends for Germany in the Last Two Decades: A Tale of Two Countries’, *Review of Economic Dynamics*, 13(1), pp 103–132.

Gottschalk P and R Moffitt (1994), ‘The Growth of Earnings Instability in the U.S. Labor Market’, *Brookings Papers on Economic Activity*, 1994(2), pp 217–254, 270–272.

Greenville J, C Pobke and N Rogers (2013), *Trends in the Distribution of Income in Australia*, Productivity Commission Staff Working Paper, Productivity Commission, Canberra.

Harding A and H Greenwell (2002), ‘Trends in Income and Consumption Inequality in Australia’, Paper presented at the 27th General Conference of the International Association for Research in Income and Wealth, Stockholm, 18–24 August.

Heathcote J, F Perri and GL Violante (2010), ‘Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006’, *Review of Economic Dynamics*, 13(1), pp 15–51.

Jappelli T and L Pistaferri (2010), ‘Does Consumption Inequality Track Income Inequality in Italy?’, *Review of Economic Dynamics*, 13(1), pp 133–153.

Jarvis S and SP Jenkins (1998), ‘How Much Income Mobility is There in Britain?’, *The Economic Journal*, 108(447), pp 428–443.

Krueger D and F Perri (2005), ‘Does Income Inequality Lead to Consumption Equality? Evidence and Theory’, Federal Reserve Bank of Minneapolis Research Department Staff Report No 363.

Krueger D, F Perri, L Pistaferri and GL Violante (2010), ‘Cross-Sectional Facts for Macroeconomists’, *Review of Economic Dynamics*, 13(1), pp 1–14.

Leigh A (2013), *Battlers & Billionaires: The Story of Inequality in Australia*, Redback, Collingwood.

MaCurdy TE (1982), ‘The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis’, *Journal of Econometrics*, 18(1), pp 83–114.

Meyer BD and JX Sullivan (2012), ‘Winning the War: Poverty from the Great Society to the Great Recession’, *Brookings Papers on Economic Activity*, Fall, pp 133–183.

Pijoan-Mas J and V Sánchez-Marcos (2010), ‘Spain is Different: Falling Trends of Inequality’, *Review of Economic Dynamics*, 13(1), pp 154–178.

Piketty T (2014), *Capital in the Twenty-First Century*, trans A Goldhammer, The Belknap Press of Harvard University Press, Cambridge.

Saunders P and P Siminski (2005), ‘Home Ownership And Inequality: Imputed Rent And Income Distribution In Australia’, *Economic Papers*, 24(4), pp 346–367.

Shorrocks A (1978), ‘Income Inequality and Income Mobility’, *Journal of Economic Theory*, 19(2), pp 376–393.

Slesnick DT (1998), ‘Empirical Approaches to the Measurement of Welfare’, *Journal of Economic Literature*, 36(4), pp 2108–2165.

Wilkins R (2013), ‘Evaluating the Evidence on Income Inequality in Australia in the 2000s’, Melbourne Institute Working Paper Series No 26/13.

Wilkins R (2015), ‘Measuring Income Inequality In Australia’, *The Australian Economic Review*, 48(1), pp 93–102.

Yates J (1994), ‘Imputed Rent and Income Distribution’, *The Review of Income and Wealth*, 40(1), pp 43–66.

Copyright and Disclaimer Notices

HILDA

The following Disclaimer applies to data obtained from the HILDA Survey and reported in this RDP.

Disclaimer

The Household, Income and Labour Dynamics in Australia (HILDA) Survey was initiated and is funded by the Australian Government Department of Social Services (DSS), and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views based on these data should not be attributed to either DSS or the Melbourne Institute.



RESERVE BANK OF AUSTRALIA