

Household Economic Inequality in Australia*

GREG KAPLAN

*University of Chicago, Chicago, IL, USA;
University of New South Wales, Sydney, NSW,
Australia; The National Bureau of Economic
Research, Cambridge, MA, USA*

GIANNI LA CAVA AND
TAHLEE STONE

Reserve Bank of Australia, Sydney, NSW, Australia

We use data from the Household Expenditure Survey and Household, Income and Labour Dynamics in Australia Survey to document facts about consumption and income inequality among households in Australia, emphasising the role of the rents imputed to home owners for conclusions about inequality. Consistent with other developed economies, consumption inequality in Australia is lower on average than income inequality. Both have increased since the early 1990s, with income inequality increasing by more. We decompose the trend in income inequality into four components: (i) changes in observed household characteristics; (ii) changes in the returns to unobserved skills; (iii) changes in the size of persistent income shocks; and (iv) changes in the size of transitory income shocks. We find that changes in the size of persistent and transitory income shocks, rather than changes in observed household characteristics, explain most of this trend. Since the middle of the 2000s, the source of income inequality has shifted from transitory to persistent factors, which is consistent with the rise in consumption inequality over the corresponding period. We find that accounting for imputed rents lowers estimates of the level of inequality in Australia, but has a negligible effect on the trends.

*We would like to thank Laura Berger-Thomson, Michael Plumb, John Simon, Peter Tulip, Roger Wilkins, two anonymous referees and the editor for useful comments and suggestions. We would also like to thank Rosetta Dollman for her contribution to an earlier version of the paper and Goustavo Moreira de Souza for outstanding research assistance. 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. 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.

JEL classifications: D12, D31, E21, H31

Correspondence: Greg Kaplan, University of Chicago, Chicago, IL, USA. Email: gkaplan@uchicago.edu

1 Introduction

We examine how the distribution of living standards has evolved in Australia over recent decades by analysing trends in household income, consumption and wealth inequality – which we collectively refer to as ‘household economic inequality’.

Reflecting overseas trends, there is a growing body of Australian research on inequality. For example, Fletcher and Guttman (2013), Greenville *et al.* (2013) and Wilkins (2015) document trends in income inequality in Australia using household survey data. They find that there has been a slight increase in income inequality over recent years. Chatterjee *et al.* (2016) examine the rise in labour income inequality over the past decade. They also document an increase in inequality in labour earnings and find that it is due to residual factors reflecting idiosyncratic risk and unexpected labour market outcomes.

Most of the empirical research to date, particularly for Australia, has focused on inequality in current income or labour earnings. But current income is not necessarily a good guide to welfare. According to the permanent income hypothesis, consumption depends not only on current income, but also on wealth (which is determined by past income) and expectations of future income. We capture each of these dimensions by comparing inequality estimates for current wealth, income and consumption. Consumption is a better indicator of living standards than either current income or wealth because it is more closely connected with households' lifetime budget constraints. Households can smooth temporary fluctuations in income by borrowing and saving. Therefore, annual income is typically more variable than annual consumption, so inequality in annual income typically overstates the level of inequality in household welfare. Moreover, most individuals experience a period of growing income during their early working years, and a period of lower income as they transition to retirement. Overall living standards thus depend more on lifetime income than on current income.

To gauge inequality in living standards, it is therefore 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 but complementary 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 well-being than current income or wealth (see, for example, Slesnick, 1998). 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.

Second, we estimate persistent income inequality by exploiting the panel dimension of the

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 measure the evolution of each type of inequality.

To estimate permanent and transitory residual labour income shocks, Chatterjee *et al.* (2016) use an unobserved components model in the spirit of Gottschalk and Moffitt (1994). More specifically, they assume residual labour income is the sum of two independent processes – a permanent (unit root) component and a transitory (independent and identically distributed (IID)) component. We estimate a more general specification of the income process than Chatterjee *et al.* (2016). First, we analyse total income rather than just labour income. As Wilkins (2015) and others show, much of the rise in inequality in recent years in Australia is due to an increase in capital income at the top of the distribution. Moreover, given our interest in capturing changes in the welfare-relevant distribution, it would seem appropriate to focus on a broader measure of income that captures both capital and labour income. Second, we estimate a more flexible model in which the persistent component is an AR(1) process rather than a random walk, and in which the transitory component is MA(1) rather than IID. This specification allows for some mean reversion in the persistent component of income (which is important since persistent factors, such as promotions, have less than permanent effects) and for some limited serial correlation in transitory income (which is important since temporary factors, such as layoffs, have effects that linger for more than a year). Third, we allow for the possibility that transitory and permanent shocks may be correlated, for example if an employee is simultaneously rewarded with a bonus and a promotion.

We construct estimates of household economic inequality using several sources of data. Our consumption inequality estimates come primarily 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.

A key contribution of our paper is to highlight the importance of adjusting estimates of inequality to include the housing services accruing to

(and consumed by) owner-occupier households. We refer to these services as ‘net imputed rent’ and compute them as the difference between gross imputed rent and the housing expenses paid by owner-occupiers.

We are not the first to point out that the distribution between household income and consumption is significantly affected by including sources of non-monetary income, such as imputed rent for owner-occupiers. However, to the best of our knowledge, the distributional impact of imputed rent has not been explored since the 1990s in Australia (Saunders & Siminski, 2005). Australia experienced a strong nationwide housing boom in the early to mid-2000s. It is therefore worth considering how the distribution of imputed rent evolved over the 2000s and how this, in turn, affected the distribution of consumption and income.

Consistent with this, we also highlight the role that housing prices and debt can play in affecting estimates of household economic inequality. Changes in the distribution of housing prices and debt have an obvious and direct impact on living standards by affecting the wealth distribution. But they also potentially have indirect effects on the distributions of income and consumption to the extent that they affect the measurement of net imputed rent.

We are also not the first to examine consumption inequality in Australia. However, most Australian studies have examined trends in the distribution of non-durables expenditure (Barrett *et al.*, 1999, 2016) or total expenditure (e.g. Harding & Greenwell, 2002; Bray, 2014). In contrast, we extend these estimates of household expenditure to cover household consumption, which includes both spending on durable goods (e.g. motor vehicles) and the imputed service flow associated with owner-occupier housing. As we show, the inclusion of these spending categories can have a significant effect on estimates of the cross-sectional distribution of household consumption. We also cover a longer time period than previous studies.

In summary, our key findings include the following:

- Household consumption inequality is consistently lower than income inequality based on estimates from repeated cross-sectional surveys. Consumption inequality has risen over recent decades, but income inequality has risen by more.
- The rise in income inequality is due to an increase in the variance of unobserved shocks, particularly since the mid-2000s. At least some of the increase in

inequality has been persistent, implying higher inequality in household welfare.

- The inclusion of net imputed rents on owner-occupier dwellings lowers estimates of household economic inequality, but has little effect on the long-run trends in Australia, which is surprising given the strong increase in expenditure devoted to housing.

II Definitions of Household Consumption and Income

(i) Data

The analysis in this paper is primarily based on unit record data from the HES for six different surveys: 1984, 1988/1989, 1993/1994, 1998/1999, 2003/2004 and 2009/2010. The HES is the most comprehensive source of cross-sectional information on household expenditure in Australia.¹ 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 (Appendix S1), we provide alternative estimates of inequality using tax records.

We focus on survey data rather than tax records in measuring inequality because estimates of non-cash income, such as imputed rent to owner-occupiers, are not available in tax data. As we show, these sources of income are important in accurately measuring inequality.

A disadvantage of using household surveys to measure inequality is that they typically under represent very rich individuals. However, the estimates of income inequality presented in the online appendices (Appendix S1) look very similar for Australia based on surveys and tax records. This suggests that our results are not significantly affected by such sampling bias.

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 2003/2004, the ABS incorporated information on salary-sacrificed income into its household

¹ The HILDA Survey has also collected estimates of expenditure on an annual basis since 2006. However, the expenditure definitions have changed over time and are not as comprehensive as the HES.

income estimates for the first time. This is likely to have boosted measured inequality relative 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.²

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 ABS's preferred method 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.

(ii) *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 earning rental income (Saunders & 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.

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 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 S1A, 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 & 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 1993/1994 onwards, we concentrate on the most recent four surveys: 1993/1994, 1998/1999, 2003/2004 and 2009/2010.

In Appendix S1B, 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

² 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.

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 scale to control for household size and composition. The equivalence scale assigns a value of 1 to each household member that is an adult and a value of 0.5 to each child under the age of 16.

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 well-being. 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). The 2009/2010 HES provides estimates of social transfers in kind. The inclusion of such transfers reduces disposable income inequality by about one-quarter, on average. 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.³ Third, we also do not

³ The 2003/2004 and 2009/2010 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.

examine trends in the distribution of leisure time, which is another indicator of household well-being (Attanasio *et al.*, 2015).

III Facts about Household Economic Inequality

(i) 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 0 to 1, where 0 represents perfect equality and 1 represents complete inequality.

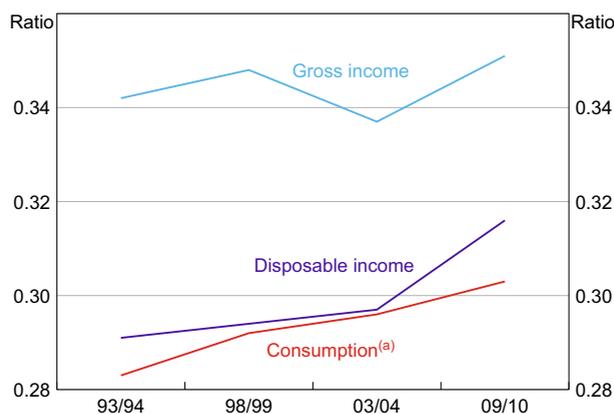
Based on the 2009/2010 HES, the Lorenz curve for gross income indicates that the top 20 per cent of households earned approximately 42 per cent of total household income. In contrast, the bottom 20 per cent of households earned about 7 per cent of income. However, income inequality is reduced to some extent by the redistribution of income from rich households to poor households through government taxes and transfers. As a result, in 2009/2010 the Gini coefficient for disposable income (0.32) was lower than that for gross income (0.35).

In addition, the Gini coefficient for consumption (0.30) was lower than that for 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 1).

In contrast, the HES estimates of disposable income inequality have risen consistently since

FIGURE 1
Consumption and Income Inequality: Gini Coefficient [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings. (a) Also excludes other interest payments.

Source: ABS; authors' calculations.

the early 1990s. This finding is consistent with other studies and largely reflects an increase in capital income inequality, as labour income inequality has not changed much over recent decades (Greenville *et al.*, 2013). According to Greenville *et al.* (2013), labour income inequality has hardly 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 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. 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 indicator of inequality. To complement the analysis, we examine how much of aggregate household income is earned by high-income households (as a proxy for income inequality) and, similarly, how much of aggregate household consumption is accounted for by high-spending households (for consumption inequality).

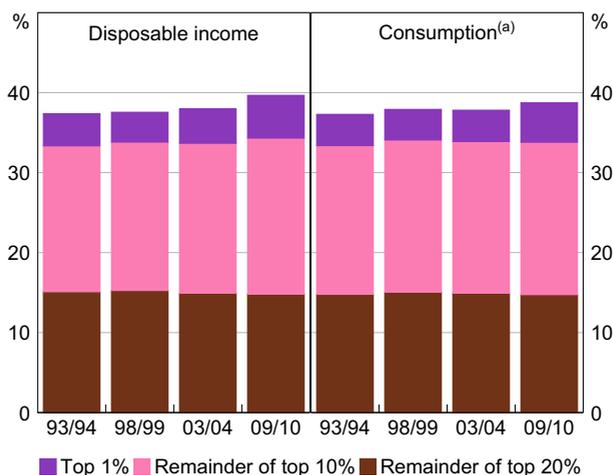
Based on the disposable income estimates, the top 10 per cent of income earners accounted for 22.3 per cent of aggregate household income in 1993/1994 and 24.8 per cent in 2009/2010 (Figure 2). Much of the increase in the share of income held by the top 10 per cent of earners has been due to the very highest earners within the top 1 per cent – their share of total disposable income rose from 4 per cent in 1993/1994 to 5.4 per cent in 2009/2010.⁴

Based on the consumption estimates, the top 10 per cent of spenders accounted for 22.5 per cent of aggregate household consumption in 1993/1994 and 24 per cent in 2009/2010. Again, a

⁴ In Appendix S1C, we provide similar estimates of income inequality using tax record data. The consistency between data sources suggests that HES survey data adequately capture top income earners.

FIGURE 2

Top Income Earners/Consumers: Share of Total Household Income/Consumption [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings. (a) Also excludes other interest payments.

Source: ABS; authors' calculations.

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 1993/1994 to 5 per cent in 2009/2010.⁵

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 IV.

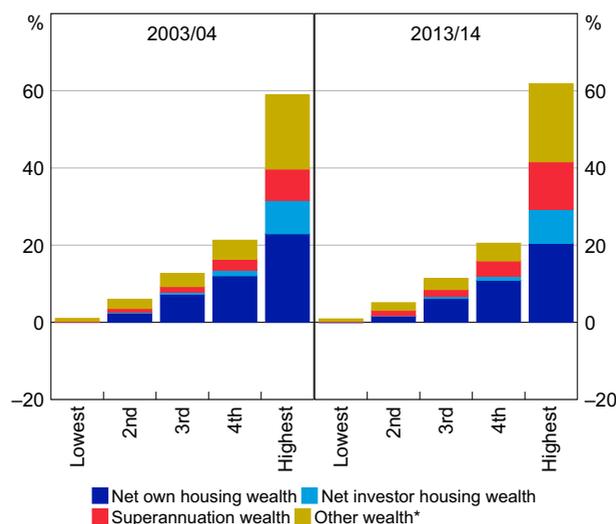
(ii) Wealth Inequality

We mainly focus on consumption and income inequality because we have long-run estimates for these measures of household well-being. But we also briefly consider wealth inequality for two reasons. First, wealth is a potentially important indicator of well-being 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 2013/2014 the Gini coefficient for net wealth was 0.60, which was well above the

⁵ To put it another way, 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 1993/1994 and 2009/2010. 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. More specifically, the top 1 per cent of spenders experienced real growth in consumption of a little 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 3
Net Wealth by Quintile: Share of Total Household Net Wealth [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: *Other wealth includes bank deposits, equities, trusts, businesses, home contents, motor vehicles and other assets not elsewhere classified less personal loans, student loans, other investment loans and credit cards.

Source: ABS.

level of inequality in disposable income. Other indicators of wealth inequality tell the same story. The share of total net wealth held by the top 20 per cent of wealthy households was 62 per cent in 2013/2014. This is 1.5 times the corresponding share of aggregate income for the top 20 per cent of income-earning households.⁶ 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.

Wealth inequality in Australia has risen over the past decade by about the same amount as income inequality. The share of aggregate wealth owned by the wealthiest households (in the top quintile) rose from 59 per cent in 2003/2004 to 62 per cent in 2013/2014 (Figure 3).⁷

⁶ The corresponding share of wealth held by the top 1 per cent was about 13 per cent.

⁷ The share of wealth held by the top 1 per cent of households follows a similar pattern, rising from 12 per cent in 2003/2004 to 13.4 per cent in 2013/2014.

As in most advanced economies, housing wealth is the largest component of aggregate household wealth in Australia. But changes in housing prices have not been the main determinant of changes in wealth inequality over the past decade. Instead, the increase is fully explained by a rise in the value of superannuation wealth of the households in the top wealth quintile; the share of aggregate wealth accounted for by the superannuation of these wealthy households rose by 4.1 percentage points (Figure 3).

(iii) 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 little 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 4). (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.

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

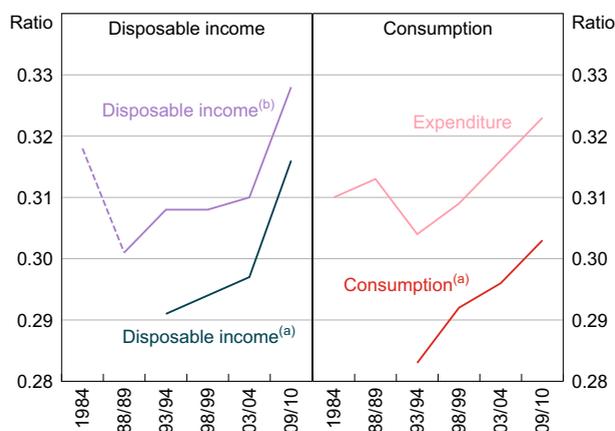
improves their welfare position. On the other hand, many high-income (and high-spending) households are comprised of young renters, meaning that the top of the distribution will not benefit to the same extent by the inclusion of imputed rent in measured income and expenditure.

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. In contrast, housing equity comprises a relatively low share of total net wealth for low-wealth households. Similarly, the average rate of home ownership is higher among low-income households than among low-wealth households. About one-third of households in the lowest wealth quintile are young renter households, while nearly half of those in the bottom income quintile are older home owners.

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

FIGURE 4

The Effect of Imputed Rent on Measured Inequality: Gini Coefficient [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: All measures are population weighted and equalised. (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.

Source: ABS; authors’ calculations.

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.

The adjustment for imputed rent increases the relative income of the very poorest households while reducing the relative income of the very richest households. The ABS estimates for 2013/2014 indicate that the adjustment for imputed rent increases the share of income going to the poorest households (bottom income quintile) by 0.4 percentage points while reducing the share of aggregate income going to the richest households (top income quintile) by 0.3 percentage points (Figure 5).

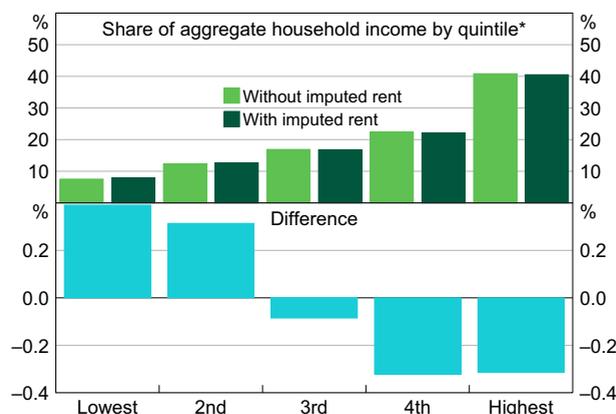
IV Transitory and Persistent Income Inequality

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 (as 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

FIGURE 5
The Effect of Imputed Rent on the Income Distribution, 2013/2014 [Colour figure can be viewed at wileyonlinelibrary.com]



Note: (a) ABS estimates of equivalised household disposable income.
Source: ABS.

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 (ECM) to fully specify the process that generates income over time and decompose income into a highly persistent component and a 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 2015. The main income measure used is real annual household disposable income. This is adjusted to incorporate estimates of net imputed rent. The gross imputed rent is estimated as 5 per cent of the total value of the home of each owner-occupier in the HILDA Survey. This is equivalent to the treatment in the HES data. To construct estimates of net imputed rent we then deduct the usual repayments on the mortgage of each owner-occupier (as reported in the HILDA Survey).

This adjusted measure of disposable income 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 118,396 household-year observations. The number of households in the sample varies by year, ranging from 6,945 to 9,500, with an average of 8,046.

Based on the HILDA Survey, there has been a trend increase in income inequality. Between 2001 and 2015, the share of aggregate disposable income held by the top 1 per cent of earners 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 the natural logarithm of household disposable income. The variance of log income increased by about 8 log points between 2001 and 2015.

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:

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

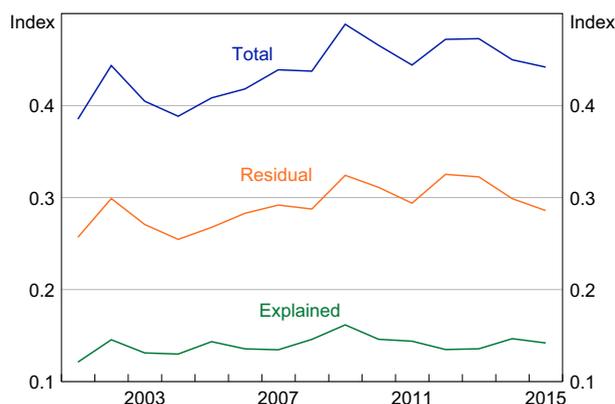
where the dependent variable is the log of equivalised household disposable income ($\ln(Y_{it})$) and the set of explanatory variables (X_{it}) includes the characteristics of the household head, such as level of education, gender, age, employment status, migrant status, indigenous status, and marital status.⁸ The specification also includes state, occupation and industry fixed effects, as well as interaction terms for occupation with industry and state variables.

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 6. 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}).

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.

⁸ The head of each surveyed household is determined by applying the following criteria, in order, until a unique person is selected. These criteria are: in a registered or de facto marriage (and still living together); a lone parent; the person with the highest income; the eldest person.

FIGURE 6
 Decomposition of Variance of Log Household Disposable Income [Colour figure can be viewed at
 wileyonlinelibrary.com]



Notes: All measures are population weighted, equivalised and include net imputed rent for owner-occupied dwellings.
 Sources: Authors' calculations; HILDA Survey Release 15.0.

(i) Error Components Model

We next exploit the panel structure of the HILDA Survey to estimate a flexible ECM. This model is used to decompose residual income inequality and 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:

$$\begin{aligned}\mu_{it} &= \lambda_t \alpha_i + z_{it} + v_{it}, \\ z_{it} &= \phi z_{i,t-1} + \eta_{it}, \\ v_{it} &= \epsilon_{it} + \theta \epsilon_{i,t-1},\end{aligned}\quad (2)$$

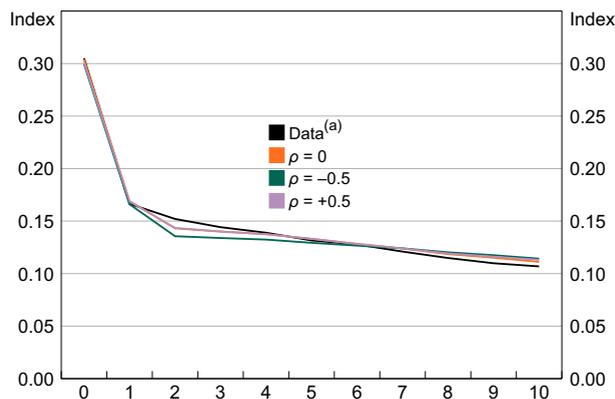
where the 'persistent' component of inequality is a combination of a household fixed effect (α_i) that has a time-varying coefficient (λ_t), with total variance $\lambda_t^2 \sigma_\alpha^2$, and a persistent term (z_{it}) that follows an autoregressive AR(1) process and has variance σ_z^2 . The household 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.

Our choice of ECM is motivated by empirical observation of the autocovariance function of household disposable income in the HILDA Survey data. Figure 7 shows the autocovariance function (averaged across all ages and years) in the data and our estimated models (described below). The large drop between lag 0 and lag 1 suggests the presence of a transitory component; the smaller discrete drop between lag 1 and lag 2 suggests that this transitory component is better captured by an MA(1) than an IID component; the exponential decline after lag 1 suggests the presence of an AR(1) component (and the near-linearity foreshadows our finding that this component is close to a random walk); the fact that the autocovariance is positive even for long lags suggests the presence of a fixed effect.

The η_{it} and ϵ_{it} terms are the respective persistent and transitory shocks to income. We assume that these shocks are mean zero with time-varying variances, $\sigma_{\eta t}^2$ and $\sigma_{\epsilon t}^2$, and a correlation $\rho_{\eta\epsilon}$. We estimate three versions of the model: one in which we assume that the two shocks are

FIGURE 7
Autocovariance Function [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: All measures are population weighted, equalised and include net imputed rent for owner-occupied dwellings. (a) Based on estimates of residual income shown in Figure 6.

Sources: Authors' calculations; HILDA Survey Release 15.0.

independent of each other, $\rho_{\eta\varepsilon}=0$; a second in which we assume that the correlation $\rho_{\eta\varepsilon}$ is exogenously fixed at $+0.5$; and a third in which we assume that the correlation $\rho_{\eta\varepsilon,t}$ is exogenously fixed at -0.5 .⁹ In all cases, we assume that the vector of shocks $(\eta_{it}, \varepsilon_{it})$ is distributed independently of time. Under this error scheme, changes in residual inequality can be driven by changes in four factors: the variance of persistent shocks; the variance of temporary shocks; the covariance between persistent and transitory shocks; and changes in the market price of a household's fixed human capital.

Ideally we would estimate the correlation between the two shocks $\rho_{\eta\varepsilon,t}$ alongside the other model parameters, rather than fixing it exogenously. When the AR(1) parameter $\phi \neq 1$, one can prove that the correlation $\rho_{\eta\varepsilon,t}$ is uniquely identified. However when $\phi=1$, $\rho_{\eta\varepsilon,t}$ is not separately identified from the variance of the fixed effect σ_x^2 .

⁹ The assumptions on $\rho_{\eta\varepsilon}$ have economic content. Constraining $\rho_{\eta\varepsilon}$ to be zero is equivalent to assuming that the structural events that are most likely captured by persistent shocks do not systematically occur at the same times as the structural events that are likely to be captured by transitory shocks. One example that is inconsistent with $\rho_{\eta\varepsilon}=0$ would be if a high-performing employee is rewarded by his or her employer with a simultaneous bonus (likely captured as a transitory shock) and raise (likely captured as a persistent shock).

It will turn out that our estimate for ϕ is sufficiently close to 1 so that we cannot reliably identify $\rho_{\eta\varepsilon,t}$. For this reason we explore the robustness of our results by presenting estimates with a strong positive correlation ($\rho_{\eta\varepsilon,t}=0.5$) and a strong negative correlation ($\rho_{\eta\varepsilon,t}=-0.5$).

We estimate the variance-covariance matrix of the model using the generalised method of moments.¹⁰ 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 9.

The estimated parameters, along with bootstrapped 95 per cent confidence intervals, are reported in Table 1 for the three versions of the model. The estimated autoregressive parameter ϕ is very close to unity, which is consistent with a

¹⁰ When estimating the model we allow for non-parametric age effects in the variance of transitory shocks, in addition to the non-parametric year effects in both shocks and the fixed effects. This allows the model to better capture the age profile of the variance of income, which would otherwise be difficult to match with only the AR(1) component to provide age variation in the cross-sectional variance.

TABLE 1
Error Component Model Parameter Estimates

Parameter	Model		
	$\rho=0$	$\rho=-0.5$	$\rho=+0.5$
ϕ	1.007 (1.002, 1.011)	0.992 (0.983, 0.998)	1.006 (0.997, 1.010)
θ	0.159 (0.133, 0.175)	0.190 (0.163, 0.245)	0.148 (0.124, 0.172)
σ_ϵ^2	0.155 (0.111, 0.202)	0.163 (0.117, 0.212)	0.137 (0.090, 0.188)
σ_η^2	0.005 (0.001, 0.012)	0.001 (0.000, 0.006)	0.005 (0.001, 0.013)
σ_α^2	0.039 (0.021, 0.050)	0.032 (0.021, 0.043)	0.036 (0.010, 0.051)

Notes: The variances of the persistent shocks and the fixed effects are unweighted averages across years. The variance of the transitory shocks is an unweighted average across years and ages. See figures for variance by age and year.
Sources: Authors' calculations; HILDA Survey Release 15.0.

large literature from other countries as well as the roughly linear decline in the autocovariance function shown in Figure 7. The estimated MA parameter θ is 0.15–0.19 depending on the correlation between the transitory and persistent shocks, which is also in the range of existing estimates. The variances of the two shocks and the fixed effects are also reported. Allowing for correlation between the two shocks has only a small effect on the estimates for the remaining parameters, and in the case with a positive correlation (which is most likely the economically relevant case), the point estimates are statistically indistinguishable from the case with zero correlation.

The estimated models fit the data well in several dimensions. First, as seen in Figure 7, the models capture the general shape of the average autocovariance function. Second, Figure 8 shows that the models capture well the age profile of the variance. Third, Figure 9, based on the model with a zero correlation, shows that the model fits the time-series path for the variance of income and thus is an appropriate model through which we can decompose the data.

Figure 9 indicates that on average one-half of the level of residual inequality is due to the variance of transitory income shocks, but this varies from around two-fifths to around three-fifths.¹¹ 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 30–40 per cent of the total cross-sectional variation in household income. The remaining variation in income across households is mainly due to variation in persistent shocks (one-third to one-half), although some inequality is also explained by variation in unobserved fixed human capital, particularly in the middle years of the sample.

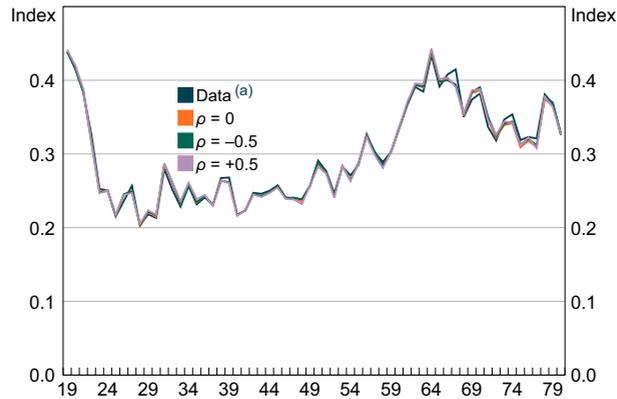
The decomposition in Figure 9 also indicates that the first half of the sample period (2001–2008) was characterised by a slight decline in transitory income inequality and a small (and largely offsetting) increase in persistent income inequality (due to an increase in the variance of the fixed effect). This appears to reflect developments in the Australian labour market over that period. In particular, the unemployment rate fell noticeably between 2001 and 2008, 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

¹¹ The ECM parameter estimates and associated standard errors are shown in detail in Appendix S1D.

FIGURE 8

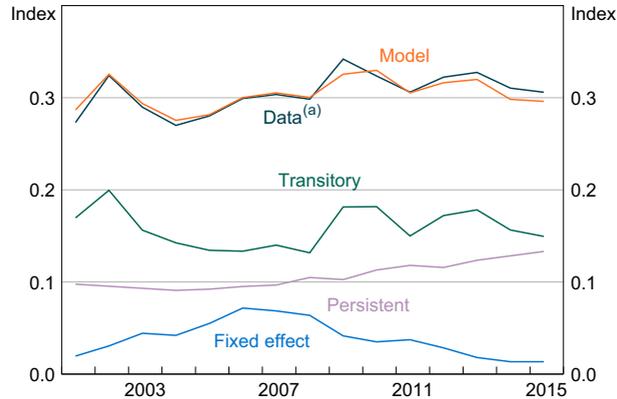
Variance of Disposable Income by Age of Household Head [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: All measures are population weighted, equalised and include net imputed rent for owner-occupied dwellings. (a) Based on estimates of residual income shown in Figure 6.
 Sources: Authors' calculations; HILDA Survey Release 15.0.

FIGURE 9

Decomposition of Variance of the Logarithm of Household Disposable Income [Colour figure can be viewed at wileyonlinelibrary.com]



Notes: All measures are population weighted, equalised and include net imputed rent for owner-occupied dwellings. (a) Based on estimates of residual income inequality shown in Figure 6.
 Sources: Authors' calculations; HILDA Survey Release 15.0.

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 mid-2000s is consistent with the small increase in

TABLE 2
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		71	29		58	42
2nd	19	53	28	23	34	42
3rd	23	50	26	30	30	40
4th	26	54	19	39	31	29
Highest	26	74		43	57	

Sources: Authors' calculations; HILDA Survey Release 15.0.

consumption inequality reported in the HES over a similar period.

(ii) 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 (which have been adjusted for net imputed rent). 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.

Over 70 per cent of households that are either in the top or bottom income quintile remain in that same quintile from 1 year to the next (Table 2). More notably, nearly 60 per cent of these households are in the same quintile a decade later.¹² For households in the middle income quintiles, about half change quintiles each year and about a third are in the same quintile a decade later. This evidence suggests that both persistent and

¹² The results reported in Table 2 are robust to using a balanced panel and a sample of working-age households.

transitory movements are occurring within the income distribution over time, though there is some variation across income quintiles.¹³

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 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, $w_t = Y_t/Y$, 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

¹³ 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.

inequality over time. For a strictly convex inequality measure $I[Y]$, the index takes 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. 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.¹⁴

To assess whether the share of persistent inequality (or income immobility) has changed over time, we divide the HILDA Survey 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 2001–2015, with average R values for total and residual household disposable income of 0.81 and 0.71, respectively (Table 3).¹⁵ However, we do observe a slight rise in the share of persistent inequality or immobility from the mid-2000s onwards, with higher R values reported for the second half of the sample period.¹⁶ This increase in persistent inequality is consistent with the rise in the variance of persistent shocks observed in Figure 9 over the same period, as well as the small increase in consumption inequality reported in the HES between 2003/2004 and 2009/2010.

V Conclusion

We document some new facts about economic inequality among households in Australia over

¹⁴ 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 III(i).

¹⁵ 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 *et al.* (2014) report Shorrocks R values of 0.83 and 0.85 for the United States and Germany over the late 1990s to early 2000s.

¹⁶ 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 & Jenkins, 1998).

TABLE 3
Shorrocks R Values for Household Disposable Income

Year	Theil index		Gini coefficient	
	Total	Residual*	Total	Residual*
2001–2005	0.80	0.71	0.91	0.85
2002–2006	0.79	0.70	0.91	0.85
2003–2007	0.80	0.71	0.91	0.85
2004–2008	0.79	0.70	0.91	0.84
2005–2009	0.78	0.69	0.91	0.84
2006–2010	0.83	0.71	0.93	0.86
2007–2011	0.86	0.74	0.94	0.87
2008–2012	0.84	0.73	0.93	0.87
2009–2013	0.81	0.70	0.92	0.85
2010–2014	0.81	0.70	0.92	0.86
2011–2015	0.78	0.67	0.91	0.85

Note: *Residual income estimated using Equation (1).

Sources: Authors' calculations; HILDA Survey Release 15.0.

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

We find that the treatment of housing income and expenditure is important to estimates of household economic inequality. In particular, incorporating the net imputed rent earned (and spent) by home owners significantly lowers estimates of inequality in both income and consumption. However, this adjustment has little effect on the trends in household economic inequality. This is somewhat surprising, given that a growing share of Australian household budgets have been devoted to housing over recent decades.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Online Appendix.

REFERENCES

- Attanasio, O., Hurst, E. and Pistaferri, L. (2015), 'The Evolution of Income, Consumption, and Leisure Inequality in the US, 1980–2010', in Carroll, C.D., Crossley, T.F., and Sabelhaus, J. (eds), *Improving the Measurement of Consumer Expenditures*. University of Chicago Press, Chicago, IL: 100–40.
- Barrett, G., Crossley, T. and Worswick, C. (1999), 'Consumption and Income Inequality in Australia', CEPR Discussion Paper No. 404, Centre for Economic Policy Research, Research School of Economics, Australian National University, Canberra.
- Barrett, G., Donald, S. and Hsu, Y.-C. (2016), 'Consistent Tests for Poverty Dominance Relations', *Journal of Econometrics*, **191**, 360–73.
- Bayaz-Ozturk, G., Burkhauser, R.V. and Couch, K.A. (2014), 'Consolidating the Evidence on Income Mobility in the Western States of Germany and the United States from 1984 to 2006', *Economic Inquiry*, **52**, 431–43.
- Beech, A., Dollman, R., Finlay, R. and La Cava, G. (2014), 'The Distribution of Household Spending in Australia', RBA Bulletin, pp. 13–22.
- Bray, J.R. (2014), 'Changes in Inequality in Australia and the Redistributive Impacts of Taxes and Government Benefits', in Podger, A. and Trewin, D. (eds), *Measuring and Promoting Wellbeing: How Important is Economic Growth?* ANU Press, Acton, ACT: 423–76.
- Chatterjee, A., Singh, A. and Stone, T. (2016), 'Understanding Wage Inequality in Australia', *Economic Record*, **92**, 348–60.
- DeBacker, J., Heim, B., Panousi, V., Ramnath, S. and Vidangos, I. (2013), 'Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns', *Brookings Papers on Economic Activity*, **46**, 67–142.
- Fletcher, M. and Guttman, B. (2013), 'Income Inequality in Australia', *Economic Roundup*, (2), 35–56.
- Friedman, M. (1957), *A Theory of the Consumption Function*. Princeton University Press, Princeton, NJ.
- Gottschalk, P. and Moffitt, R. (1994), 'The Growth of Earnings Instability in the U.S. Labor Market', *Brookings Papers on Economic Activity*, **25**, 217–72.
- Greenville, J., Pobke, C. and Rogers, N. (2013), *Trends in the Distribution of Income in Australia*. Productivity Commission, Melbourne.
- Harding, A. and Greenwell, H. (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.
- Jarvis, S. and Jenkins, S.P. (1998), 'How Much Income Mobility Is There in Britain?', *Economic Journal*, **108**, 428–43.
- Saunders, P. and Siminski, P. (2005), 'Home Ownership and Inequality: Imputed Rent and Income Distribution In Australia', *Economic Papers*, **24**, 346–67.
- Shorrocks, A. (1978), 'Income Inequality and Income Mobility', *Journal of Economic Theory*, **19**, 376–93.
- Slesnick, D.T. (1998), 'Empirical Approaches to the Measurement of Welfare', *Journal of Economic Literature*, **36**, 2108–65.
- Wilkins, R. (2013), 'Evaluating the Evidence on Income Inequality in Australia in the 2000s', Working Paper 26/13, Melbourne Institute Working Paper Series, Melbourne Institute of Applied Economic and Social Research, The University of Melbourne.
- Wilkins, R. (2015), 'Measuring Income Inequality in Australia', *Australian Economic Review*, **48**, 93–102.
- Yates, J. (1994), 'Imputed Rent and the Income Distribution', *Review of Income and Wealth*, **40**, 43–66.