the review of income and wealth

Review of Income and Wealth
Series 0, Number 0, Month 2022
DOI: 10.1111/roiw.12628

RISING TOP-INCOME PERSISTENCE IN AUSTRALIA: EVIDENCE FROM INCOME TAX DATA

BY NICOLAS HÉRAULT

Bordeaux School of Economics, University of Bordeaux, Pessac, France
Melbourne Institute of Applied Economic and Social Research, The University of Melbourne, Parkville, Victoria, Australia

DEAN HYSLOP

Motu Research Wellington, New Zealand

STEPHEN P. JENKINS

Department of Social Policy, London School of Economics and Political Science, London, UK
IZA Institute of Labor Economics, Bonn, Germany

AND

ROGER WILKINS*

Melbourne Institute of Applied Economic and Social Research, The University of Melbourne, Parkville, Victoria, Australia
IZA Institute of Labor Economics, Bonn, Germany

Using income tax administrative data for Australia, we examine levels and trends in the persistence in top-income group membership, focusing on the top 1 percent. Top-income persistence increased markedly between 1991 and 2018, with most of the increase occurring in the mid-2000s and early 2010s. In the mid- to late-2010s, Australian top-income persistence rates were near the top of the range of tax-data estimates for other countries. We decompose the increase into factors associated with (i) changes in the composition of the top-income group and (ii) increases in persistence rates for specific population subgroups. We find that the rise in top-income persistence is accounted for by changes in subgroup persistence rates, notably for individuals aged 35–64, and especially those aged 55–64. We suggest that these effects are partially related to increases in the effective retirement age over the relevant period.

JEL Codes: D31, I31, C81

Keywords: top incomes, income mobility, top-income persistence

1. INTRODUCTION

There is continuing interest in top incomes, especially in information about levels and trends of the share of total income received by the top 1 percent or other

Note: We thank the handling editor Conchita D'Ambrosio, two anonymous referees, Andrew Leigh, and participants at the Tax and Transfer Policy Institute's ALife Conference (March 2021), the TREE seminar in Pau (October 2021), and the World Inequality Conference in Paris, (December 2021), for their helpful comments on earlier versions of this paper. Open access publishing facilitated by The University of Melbourne, as part of the Wiley - The University of Melbourne agreement via the Council of Australian University Librarians.

*Correspondence to: Roger Wilkins, University of Melbourne, Victoria, Australia and IZA, Germany (r.wilkins@unimelb.edu.au).

top-income groups such as the top 0.1 percent. Long time-series of estimates of top income shares now exist for many countries: see the WID. World portal (https://wid.world/) for an extensive collection from around the world. However, much less is known about intertemporal persistence in top-income group membership. This is a notable gap because how we judge estimates of yearly top income shares depends on how much turnover there is at the top. We are more likely to tolerate a top 1 percent income share of 10 percent if different individuals form the top 1 percent each year than if the same people are always at the top. The greater mobility in the former case means greater equality in the opportunity to reach the top of the income distribution than the latter case. The latter case signals an entrenched elite and greater inequality of permanent income. In this paper, we present new evidence about top-income persistence levels and trends using high-quality personal income tax data for Australia.

The advantages of using income tax data for studying top-income mobility are the same as the advantages for studying income inequality in yearly cross-sections. By comparison with household survey data, income tax data suffer much less from top-income under-coverage (whether arising from unit or item nonresponse) and provide significantly larger sample sizes, meaning that top-income group membership can be characterized more accurately. (We elaborate this point below: see Footnote 8.) The longitudinal data we use in this paper, from the Australian Longitudinal Information files ("ALife") produced by the Australian Tax Office (ATO), have these desirable qualities. ALife also contains information about the characteristics of individuals, and so we can examine differences in top-income persistence between different population subgroups.

We make four contributions. First, compared to most top-income mobility studies based on tax data, our analysis is more closely linked to the mainstream literature about top-income shares. That literature refers to income shares held by different groups among the population of all adults, whereas many prominent studies of top-income mobility have instead used data about tax filers—a subset of all adults (mostly individuals who are liable for income taxation). Table 1 demonstrates our point.

Table 1 summarizes the key features of eight earlier studies of top-income persistence and contrasts our work with them. We restrict attention to studies based on administrative record data (typically income tax records) rather than survey data for the reasons given earlier.[1] Only three earlier studies—Aaberge *et al.* (2013) for Norway, Boschini *et al.* (2020) for Sweden, and Saez and Veall (2005) for Canada—have defined top income groups with reference to the total adult population, albeit with a range of lower age-cut offs to define that population.

Our second contribution is to provide a more detailed description of the top-income mobility process than earlier tax-data-based studies. Previous work has taken two main approaches to summarizing persistence, as the far-right-hand column of Table 1 indicates. The first and most common approach, exemplified by, e.g. Auten *et al.* (2013) for the U.S., is to take the individuals who belong to the

[1]Table 1 is not comprehensive. We do not consider studies of mobility in labor earnings (see, e.g. Kopczuk *et al.* 2010 or Martínez 2017). Some other earlier research about top-income mobility is cited by the studies listed in Table 1.

Table 1
Studies of Top-Income Persistence Based on Personal Income Tax Administrative Record Data

Author(s)	Country	Income unit	Population	Period covered	Principal approach to summarizing top-income persistence
Aaberge *et al.* (2013)	Norway	Individual	All aged 16+	1967–2011	T-year average income compared to single-year top-income prevalence
Auten and Gee (2009)	USA	Family	Tax filers (aged 25+)	1996–2005	Top-income transition matrices, and associated income changes
Auten *et al.* (2013)	USA	Family	Tax filers (aged 25–60)	1991–2010	Fraction staying in top-income group
Boschini *et al.* (2020)	Sweden	Individual	All aged 20+	1971–2012	Fraction staying in top-income group; T-year average income compared to single-year top-income prevalence, separately for men and women (not for all adults)
Jenderny (2016)	Germany	Family	Tax filers	2001–2006	T-year average income compared to single-year top-income prevalence
Joyce *et al.* (2019)	UK	Individual	Tax filers	2001–2016	Fraction staying in top-income group, pooled data for 2001–2016
Landais (2008)	France	Individual	Tax filers	1998–2004	Fraction staying in top-income group
Saez and Veall (2005)	Canada	Individual	All aged 20+	1982–2000	Fraction staying in top-income group; T-year average income compared to single-year top-income prevalence
This paper	Australia	Individual	All aged 15+	1991–2018	Fraction staying in top-income group; T-year average income compared to single-year top-income prevalence; income changes associated with leaving and joining the top 1%; income origins and destinations of leavers and joiners

Notes: All studies use a "pre-tax gross income" (excluding capital gains) cash income definition. The tax unit is the individual in all countries for the periods covered, except for the U.S.A. and Germany where it is the family (single adult or married couple). The U.S. family income measure is equalized by the square root of family size. The German family income measure is not adjusted for differences in family size. Population refers to the population used for the income persistence analysis in the paper cited. The principal approaches to summarizing top-income persistence are the first and second approaches explained in the main text.

top-income group of interest in a specific year and calculate the fraction of these individuals who remain in the same top-income group in subsequent years. That is, the focus is on top-income survival rates and their trends over time. This approach ignores the fact that turnover in top-income group membership also arises because non-top-income individuals join the top-income group. Distinctively, we also document top-income re-entry rates for previous top-income group members. (With ALife data, we cannot identify whether a top-income spell is the first such spell.) In addition, following Auten and Gee (2009) and Auten *et al.* (2013), we provide information about the income group destinations of top-income leavers, and the origins of top-income joiners, using graphical summaries of transition matrices, supplementing these pictures with discussion of the income changes that accompany the top-income mobility. In each case, we document how patterns have changed over time.

The second approach in previous literature to summarizing top-income persistence is to take a window T years long and to compare the top-income shares of incomes longitudinally-averaged over the T years ("permanent" shares) with (averaged) yearly top-income shares ("transitory" shares). The greater the difference between the two shares, the more top-income persistence there is. This is an application of Shorrocks' (1978) approach to income immobility in which a top-income share is used as the inequality index rather than the Gini coefficient or other indices based on all incomes. Aaberge *et al.*'s (2013) top-income mobility index encapsulates this idea; we convert their index to an immobility index. Three earlier studies based on tax data have taken the permanent/transitory approach (see Table 1). We apply this approach to our Australian data, using moving windows to describe levels and trends in top-income persistence.

Our third contribution arises from the nature of the specific application—to Australia and covering the 25-year period from 1991 through 2018. The long time span means that we can study trends in persistence as well as levels, and through to a year well after the Global Financial Crisis (GFC). The data used by four of the studies cited in Table 1 cover only up to the mid-2000s or earlier. Joyce *et al.* (2019) have U.K. data covering 2001–2016 but their analysis of top-income persistence is relatively brief (using only one approach) and they examine tax filers rather than all adults. The studies by Aaberge *et al.* (2013) for Norway, and Boschini *et al.* (2020) for Sweden, are more comparable to ours because they use data for all adults as we do, and the data span around 40 years in each case though end in the early 2010s. Our analysis uses data covering a long period too but, unlike the two Nordic studies, we employ multiple approaches to the measurement of top-income persistence in a single study. At the same time, by looking at Australia, we provide a new Anglo country comparison to these two Nordic countries. In all three nations, the share of total income held by the top 1 percent was declining prior to the mid- to late-1970s but increased thereafter, albeit at different rates. In around 1980, the 1 percent share was around 5 percent in Australia, Norway, and Sweden but, by the mid-2000s, was around 7 percent to 8 percent in Australia and Sweden and around 11 percent to 12 percent in Norway. By comparison, the U.S. top 1 percent share was around 17 percent at that time. (See Atkinson *et al.* 2011, Figures 8 and 10) Although cross-national comparisons of top-income persistence are difficult to make because the various studies use different measures and cover different time

periods, we are able to provide some new (but broad-brush) information about how trends in Australia compare with trends for a range of other countries.

Our fourth contribution is analysis of the factors associated with the rise in top-income persistence, focusing on the top 1 percent. ALife provides information about sex, age, principal income source, self-employment status, and federal state of residence, and we use this to examine whether rising persistence in the top 1 percent is associated with changes in persistence for specific population subgroups or changes in the subgroup composition of the top 1 percent. We supplement univariate breakdowns with regression-based decomposition analysis: we apply the methods of Fairlie (2005) to decompose changes over time in (i) the probability of remaining in the top 1 percent from 1 year to the next, and (ii) the probability of entering the top 1 percent.

We show that top-income persistence in Australia increased between 1991 and 2018, with most of the increase occurring in the early- and late-2000s. This picture arises whichever approach to assessing persistence we employ. Subject to caveats about comparability of measures and differences in time periods, top-income persistence levels in Australia in the mid- to late-2010s are towards the top of the range of estimates provided by the studies for other countries cited in Table 1. Comparability issues also bedevil cross-national comparisons of trends in persistence rates. However, despite differences in trends in yearly top-income shares, Australia and Norway both experienced marked increases in top-income persistence in the mid-2000s whereas Sweden—with similar top-income share trends to Australia—did not.

Seeking explanations for the rise in top-income persistence in Australia, we decompose the increase in top-income persistence rates into factors associated with (i) changes in the composition of the top-income group (share of men and women, main income source is labor or non-labor income, self-employed status, and shares of different age groups), vs. (ii) increases in persistence rates for specific groups in the population. We find that changes in the subgroup composition of the top 1 percent do not account for the rise in top-income persistence; rather, it is accounted for by changes in subgroup persistence rates, notably those for individuals aged 35–64, especially 55–64 years. We suggest that these effects are partially related to increases in the effective retirement age over the relevant period.

Our paper proceeds as follows. The ALife dataset is described in Section 2. Section 3 documents levels and trends in top-income persistence using a range of perspectives. We focus on the top 1 percent as the top income group, but report robustness of estimates to various other "top income" definitions ranging from top 10 percent to top 0.1 percent. We examine turnover in the top 1 percent in more detail in Section 4, noting that changes in top 1 percent membership arise because people not in the top 1 percent experience an income rise sufficient to move their incomes above the 99th percentile ($p99$) threshold or people in the top 1 percent experience an income fall sufficient to move them below the $p99$ threshold. However, the effect on top income group membership also depends on the location of the top-income threshold relative to the incomes defining other income groups. How much top income mobility there is depends on how far apart the "rungs of the ladder" are (the rungs are the incomes defining income group boundaries). For someone outside the top 1 percent, located at $p95$ (say), an increase in real income of $10,000 is more likely to lead to top income group membership if the gap between $p95$ and $p99$ is small rather than large. We provide information about the extent to

which the top income rungs have been moving closer together or further apart, and about the income changes per se for those within the top income group and those on its fringes.

Sections 5 and 6 contain our analysis relating the rise in top-1 percent persistence to changes in the top 1 percent persistence within specific population subgroups and changes in the subgroup composition of the top 1 percent. Section 7 contains a summary and conclusions. We provide supplementary estimates in Appendix S1 A–F.

2. Data

We use the ALife longitudinal unit record dataset produced by the Australian Taxation Office (ATO), made available to researchers through a secure remote access facility. ALife is based on a 10 percent random sample of all tax filers observed by the ATO for tax years 1991 through 2018 and contains all the income tax records for these individuals over this period. We use the 2018 release of ALife, compiled in October 2020, by which time tax returns had been finalized for almost all people required to file a tax return for the 2018 tax year. (The Australian tax year runs from July 1 to June 30. We refer to tax years according to the calendar year in which the tax year ended. For example, 1991 refers to the tax year running from July 1, 1990 to June 30, 1991.) Filing is compulsory for those with taxable income above the tax-free threshold. Many who earn less than the tax-free threshold also lodge a return to claim back tax withheld by their employer. The tax-free threshold was $5,249 in 1991, $5,400 from 1992 to 2000, $6,000 from 2001 to 2012 and $18,200 from 2013 onwards. The ALife sample also includes people who never file a tax return that the ATO becomes aware of because of other contact with the government, for example through receipt of government benefits.

In addition to detailed information on income components, deductions, rebates, offsets, and tax liabilities, ALife contains information on year of birth, sex, residential location and, for employed persons, occupation. The tax unit in Australia is the individual, and very little information is available about a tax filer's spouse (if present). For years in which an ALife-sampled individual did not file a tax return, ALife records the individual's information for those years as missing. Full details about ALife, including the variables available, are available at https://alife-research.app/.

Income totals and income components are not top coded, with one exception: in each year, the 24 largest "employment termination" (redundancy) payments in the entire tax filer population are set equal to the level of the 25th-largest payment value. Between 1991 and 2017, this represents an adjustment of between $8 million and $57 million in total and affected between 0 and 7 individuals in ALife each year (see Appendix S1). All these individuals continue to belong to the top 0.5 percent of income recipients after this top coding, which means that the top coding does not affect estimates of transitions into and out of the top 1 percent.

Our approach to examining top incomes using tax return data follows earlier work for Australia, notably Burkhauser *et al.* (2018), who in turn build on the work of Atkinson and Leigh (2007), and is consistent with approaches taken for many other countries (see, for example, Morelli *et al.* 2015). The main income measure

is annual gross taxable income exclusive of realized capital gains. Gross taxable income refers to taxable income before deductions, such as for work expenses and concessional superannuation contributions, and before addition of tax credits such as dividend imputation credits.[2] We exclude taxable realized capital gains from the main analysis for reasons that are explained in detail by Burkhauser *et al.* (2015). Key among these reasons is that realized capital gains on assets held more than 1 year (excluding the family home) only became taxable from July 1, 1986, and only on assets acquired after September 19, 1985. This resulted in a steady rise in the share of realized capital gains entering the tax base from 1986, which would lead to spurious measured increases in top incomes if realized taxable capital gains were included.[3]

Australia has a system of individual-level taxation rather than the family-level taxation that exists in some countries (e.g. the U.S. and Germany). Analyses using Australian income tax data are therefore of the distribution of individual gross taxable income among individuals. Furthermore, since almost all tax filers are aged 15 or over, we examine distributions among individuals aged 15 and over. Consequently, we define top-income groups with reference to the total resident population aged 15 and over.[4] Following Burkhauser *et al.* (2018), we derive top-income shares using income control totals from the household income account of the National Accounts rather than from ALife.[5]

Most of our analysis focuses on the top 1 percent of individuals aged 15+, corresponding to the top 13,499 tax filers in ALife in 1991 and the top 19,962 tax filers in ALife in 2018. We also examine the top 0.1 percent and top 10 percent, with most estimates for these groups reported in Appendix S1 E and F respectively.

We begin by providing information about yearly top-income shares to provide a reference point for our analysis of top-income persistence. Panel (a) of Figure 1 shows our ALife-based estimates of the income shares of the top 0.1 percent, top 1 percent, and top 10 percent income groups by year over the period 1991–2018. It shows that the top 1 percent income share increased from 6.4 percent in 1991 to 9.5 percent in 2018. The increase largely occurred over the 1990s, with the income share of the top 1 percent in 2018 approximately the same as in 2001.[6]

[2]Since July 1, 1987, Australia has had a system of dividend imputation allowing dividend recipients to claim tax credits for the imputed company tax paid on those dividends.

[3]Inclusion of taxable realized capital gains is problematic even absent changes over time in the share of capital gains that are taxed. Realized capital gains typically relate to a period longer than the annual time-frame over which other income sources are measured. In principle, it is all capital gains accrued over the year that should be included, not the taxable capital gains that happened to be realized in that year. Moreover, capital gains are not measured in the National Accounts, which are used to estimate the income "control total".

[4]We derive population control totals from ABS Catalogue No. 3101.0 (Australian Demographic Statistics), table 59.

[5]We derive income control totals from the December 2020 release of the National Accounts. A given year's control total is equal to: Gross mixed income + Compensation of employees + Interest + Dividends + Workers' compensation + Social assistance benefits – Interest payable by unincorporated enterprises – Consumption of fixed capital – Employers' social contributions.

[6]Our estimates are consistent with estimates for the 2004–2014 period reported by Burkhauser *et al.* (2018): see Appendix S1 B.

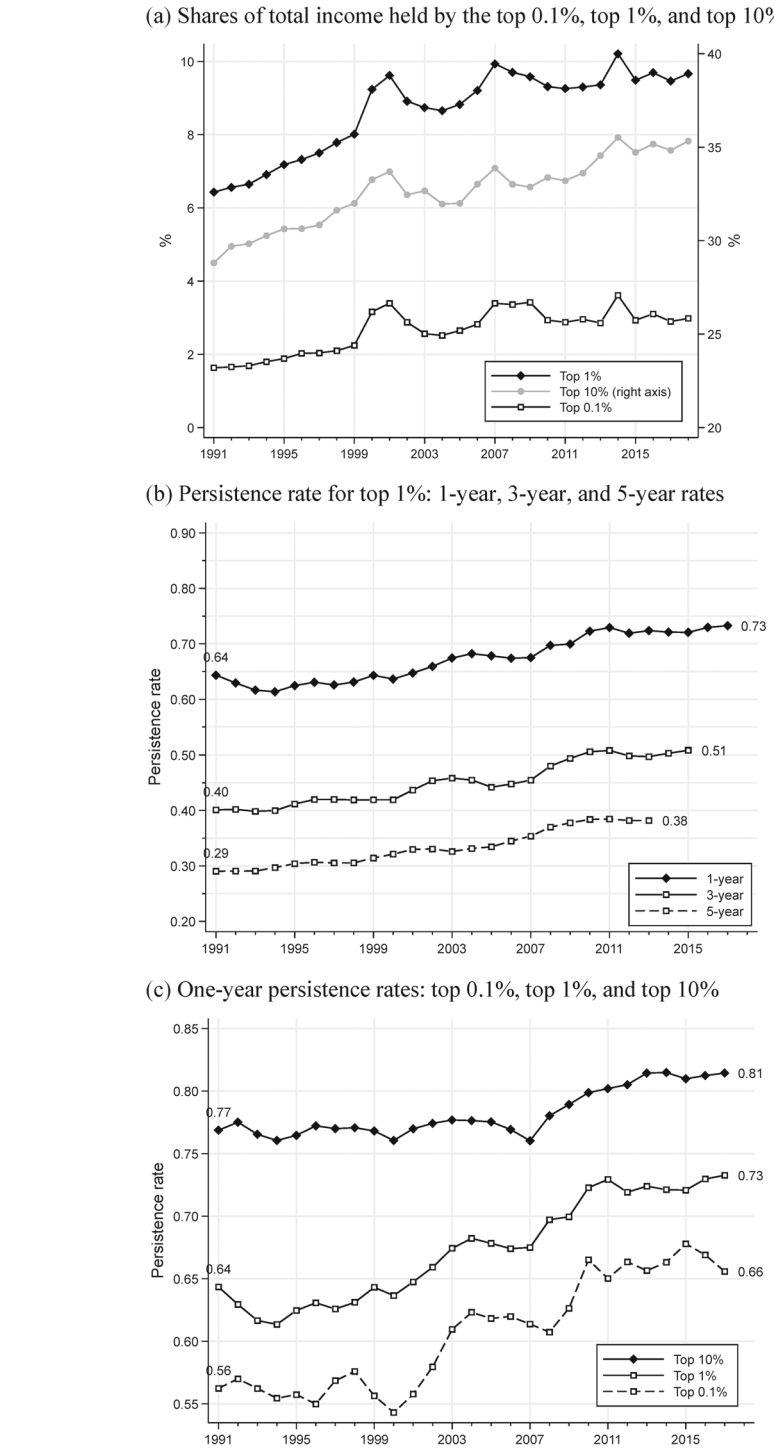

Figure 1. Top-income shares and top-income persistence rates, 1991–2018.

The trends are broadly similar for the top 0.1 percent and top 10 percent. The top 0.1 percent share rose from approximately 1.8 percent to 3 percent over the period but there was little change after 2001. The top 10 percent share rose from approximately 28 percent to 35 percent but, in contrast to the top 1 percent and top 0.1 percent share, the top 10 percent share rose after the mid-2000s.[7]

3. Top-Income Persistence—Multiple Perspectives

3.1. *Estimates of top-income persistence*

We first consider top income persistence using the approach of Auten *et al.* (2013), showing the proportion of those in the top 1 percent in base year t who remain in the top 1 percent in every subsequent year over three time frames: 1 year, 3 years, and 5 years. See panel (b) of Figure 1. As expected, rates of survival in the top 1 percent decrease the longer the time frame considered: the 5-year estimates lie below the 3-year estimates and, in turn, the 3-year estimates lie below the 1-year estimates.[8]

There is a clear rise in persistence according to all three measures over the period as a whole. The one-year survival rate fell from 64 percent in 1991 to 61 percent in 1994 but then rose to reach a peak of 73 percent in 2011 and remained at approximately that level thereafter. Similarly, the three-year persistence rate rose from 40 percent in 1991 to 51 percent in 2011 and the five-year persistence rate increased from 29 percent in 1991 to 38 percent to reach peaks in 2011, with both measures remaining at these levels thereafter.

Panel (c) of Figure 1, focusing on one-year persistence rates, shows that the rise in top-income persistence also holds when we use narrower and broader definitions of the top-income group.[9] However, it is also evident that the rise in persistence is greater the more narrowly defined is the top-income group. The one-year persistence rate rose from 0.77 to 0.81 between 1991 and 2018 for the top 10 percent (a 5 percent increase), from 0.64 to 0.73 for the top 1 percent (14 percent), and from 0.56 to 0.66 for the top 0.1 percent (18 percent).[10]

[7]Appendix S1 C shows the composition of the top 1% income group in terms of main source of income, sex, self-employment status, sex, age, and state of residence. Between 1991 and 2018, the female share of the top 1% has increased, the proportion aged 55–64 has increased, the proportion aged under 35 has decreased, and the proportion residing in Western Australia has increased but, otherwise, no clear trends are evident.

[8]Corresponding estimates from the Australian HILDA panel survey underline the value of using income tax data. We find top 1% persistence estimates (1-year, 3-year, and 5-year) are distinctly lower in HILDA than in ALife and exhibit more volatility. (The relatively large HILDA 95% confidence intervals do not include ALife point estimates for most years.) A counterpart to Figure 1 that includes HILDA estimates is available from the authors on request.

[9]Appendix S1 Figures F2 and G2 present persistence rates for the top 0.1% and top 10% for all three time-frames.

[10]Appendix S1 Figures D1 and D2 show that the rise in top-income persistence is robust to analysis of alternative income definitions, namely, gross taxable income inclusive of taxable realized capital gains, and disposable (after-tax) income. When we change the income definition, we also change the composition of the top-income group.

Figure 1 shows that trends in top-income persistence track the trend in top-income shares, albeit with a lag of 2–3 years. For example, the top 1 percent persistence rate series flattened out around 2010 whereas the growth rate in the yearly top 1 percent share fell from around 2007. There is also some correlation with the business cycle, with persistence rates rising when the economy was growing. Conversely, Australia experienced a short contraction in real Gross Domestic Product growth around 2008/2009 due to the GFC, and there was a sharp recession around 1991, and we see a drop or flattening out in the top-income persistence rate series over the next 3 to 4 years in each case.

3.2. *Cross-national differences in top-income persistence?*

How do these persistence rates for Australia compare with those in other countries? Cross-national comparability of estimates of levels and trends is limited by differences in income tax rules and thus in the definition of taxable income and tax units, and because of differences across studies in methods and time periods covered. Subject to these caveats, our estimates for Australia appear to be near the top of the range of tax-data estimates for other countries.

For Australia, we estimate the fraction of the top 1 percent of all adults aged 15+ remaining there 1 year later to range from below 70 percent in the 1990 s to more than 70 percent in the mid- to late-2010s (panel (b) of Figure 1). By comparison, for the U.S., Auten *et al.* (2013, Table 3) report 1-year persistence rates between 60 percent and 70 percent between 1991 and 2009 for tax filers aged 25–60.[11] The three-year persistence rates for the U.S. are between 29 percent and 37 percent (all referring to the pre-2010 period), whereas these rates have been hovering around 50 percent in Australia since 2010.

For the UK, and using pooled data for tax years 2000/2001 to 2015/2016, Joyce *et al.* (2019, Figure 14) report that 75 percent of the top 1 percent of tax filers remained in the top 1 percent after 1 year, 60 percent remained after 2 years, and 50 percent after 3 years, respectively. These are higher estimates than our corresponding ones for Australia but refer to tax filers and rates are averages for the whole period so trends cannot be seen.

For Sweden, Boschini *et al.* (2020, Figure C4) report one-year persistence rates for the top 0.1 percent and "top 1 percent excluding the top 0.1 percent", with separate estimates for men and women aged 20+ (but not for "all individuals"). In the years around 2010, the one-year persistence rate for men in the "top 1 percent excluding the top 0.1 percent" is around 70 percent and a few percentage points lower for women in the same group. In other words (and with reminders about comparability issues), the Swedish one-year persistence rates appear to be of roughly the same magnitude as our estimates for Australia over the same period. However, the Swedish 5-year persistence rates appear to be lower than our Australian ones for

[11] Auten *et al.* (2013) exclude capital gains from income as we do, but note that there are differences in definitions of income and income unit, and they focus on samples of tax filers rather than all adults (Table 1). Auten et al. restrict their analysis to tax filers aged 25–60, whereas we consider all adults aged 15+. Appendix S1 Figure F1 shows that applying age restrictions to our Australian samples leads to even higher persistence rates.

the years around 2010. The 5-year persistence rate for Swedish men in the "top 1 percent excluding the top 0.1 percent" was around 30 percent and a few percentage points smaller for women in the "top 1 percent excluding the top 0.1 percent" (Boschini *et al.* 2020, Figure 6).

For Germany, Drechsel-Grau *et al.* (2022, Figure 17) show that the 1-year persistence rates in the top 1 percent hovered between 66 percent and 68 percent between 2000 and 2015 for workers but increased from 70 percent to 76 percent for entrepreneurs. Although they do not provide results for the top 1 percent as a whole, which impedes comparability with our own estimates, their estimates suggest Germany's persistence rates are similar to Australia's.

There are persistence rate estimates for the three other countries focusing on the top 0.1 percent. For Australia, Appendix S1 E shows one-year persistence rates for this group of between 55 percent and 60 percent for the 1991–2003 period, above 60 percent for 2003–2010, and above 65 percent for 2010–2018. For France, Landais (2008, Figure 7) reports one-year rates for the top 0.1 percent of tax filers hovering around 65 percent between 1998 and 2004, above our estimates for all adults aged 15+. For Germany, Jenderny (2016, Figure 1) reports rates for tax filers of between 45 percent and 48 percent for the 2001–2005 period, i.e. distinctly lower than our Australian estimates. However, Drechsel-Grau *et al.* (2022, Figure 17) report much higher persistence rates, rising from 55 percent to 73 percent for workers between 2000 and 2015 and from 59 percent to 70 percent for entrepreneurs, thus much more similar to our Australian estimates. The estimates for Canada derived by Saez and Veall (2005) are more comparable with ours because they are based on all adults (aged 20+).

Saez and Veall (2005, Figure 7) report estimates ranging between 50 percent and 60 percent over the period 1982–1998, and hence of roughly the same magnitude as our Australian estimates.

3.3. *Additional perspectives on top-income persistence*

Another approach to summarizing top-income persistence is shown by Table 2. This displays the distribution of the number of years in the top 1 percent over a 7-year period among those ever observed in the top 1 percent over that 7-year period. (This is the top-income analogue to statistics summarizing the distribution of number of years in poverty over a fixed time window.) By contrast with the series shown in Figure 1, this summary shows the extent to which membership of the top 1 percent is intermittent and highlights that turnover in top-income group membership arises from entries to as well as exits from the group.

The "number of years in top 1 percent" measure shows growth in top-income persistence since 1991 as well. The fraction of individuals in the top 1 percent only 1 year out of the seven fell substantially, from one-half for the 1991–1997 period to 38 percent for the 2012–2018 period. In contrast, the fraction spending two to 6 years out of five in the top 1 percent increased from around 41 percent to 48 percent. The fraction spending all 7 years in the top 1 percent increased more dramatically: 8.7 percent of individuals belonged to the top 1 percent in every year during the 1991–1997 period but 14 percent during 2012–2018.

A feature of the T-year persistence rate measures reported in Figure 1 is that they do not take account of how long an individual has already been in the

Table 2
Distribution of the Number of Years in the Top 1% over a Seven-Year Period among Individuals in the Top 1% at Least 1 Year of the 7-Year Period

	Distribution of the number of years in top 1% (in %)							
Period	1	2	3	4	5	6	7	Total
1991–1997	50.0	15.1	9.5	6.6	5.3	4.7	8.7	100
1998–1904	47.9	14.9	9.5	7.4	5.7	4.8	9.8	100
2005–1911	40.2	16.3	11.4	8.4	6.4	5.6	11.7	100
2012–1918	37.6	16.1	11.3	8.1	6.8	6.1	14.1	100

top-income group: the measures are based on samples of the individuals who happen to belong to the top-income group in the base year. In that year, some individuals will have just joined the top 1 percent; some will have been in the top 1 percent for several years already. Because the chances of leaving the top-income group decline with the number of years since entry—there is negative duration dependence in the exit hazard rate (not shown)—Figure 1 provides over-estimates of how long someone starting a spell in the top 1 percent will remain in the top 1 percent.[12] A further feature of this figure is that it provides a one-sided perspective on turnover at the top. By conditioning on top-income group membership, the persistence measures provide no information about entry rates and their trends.

Addressing these issues enriches the description of what "top-income persistence" entails. To do this, in Figure 2 we present information about probability distributions of survival in the top 1 percent for cohorts entering the top 1 percent in different years (panel (a)) and survival outside the top 1 percent (i.e. in the poorest 99 percent) for cohorts leaving the top 1 percent in different years (panel (b)).[13] We group years during 1991–2017 into five sub-periods when defining cohorts. Corresponding estimates for the top 0.1 percent and top 10 percent are shown in Appendix S1 E and F and are similar to those for the top 1 percent discussed here.

Figure 2 shows clearly that the chances of remaining in the top 1 percent for entrants to the top 1 percent between 2007–2011 and 2012–2017 are distinctly larger than the chances of remaining for entry cohorts in the 1990 s (1992–1996, 1997–2001). For example, the probability of remaining at least 3 years in the top 1 percent was around 25 percent for the most recent cohorts but only around 18 percent for the two earliest cohorts. The differential survival chances exist regardless of how long it is since the individuals entered the top 1 percent. For example, the probability of remaining in the top 1 percent for at least 10 years since entry is around 10 percent for the 2007–2011 entry cohort but only around 6 percent for the earliest two cohorts. The picture of top-income persistence increasing over time

[12]This is a standard result from survival analysis about sampling from the "stock" rather than a cohort of entrants.

[13]We cannot estimate how long it takes to first enter the top-income group (i.e. not conditioning on being in top-income group at least once between 1991 and 2017) because of left censoring. ALife has no data about incomes before 1991 and so we cannot consistently estimate when individuals became at risk of entering the top 1% for the first time.

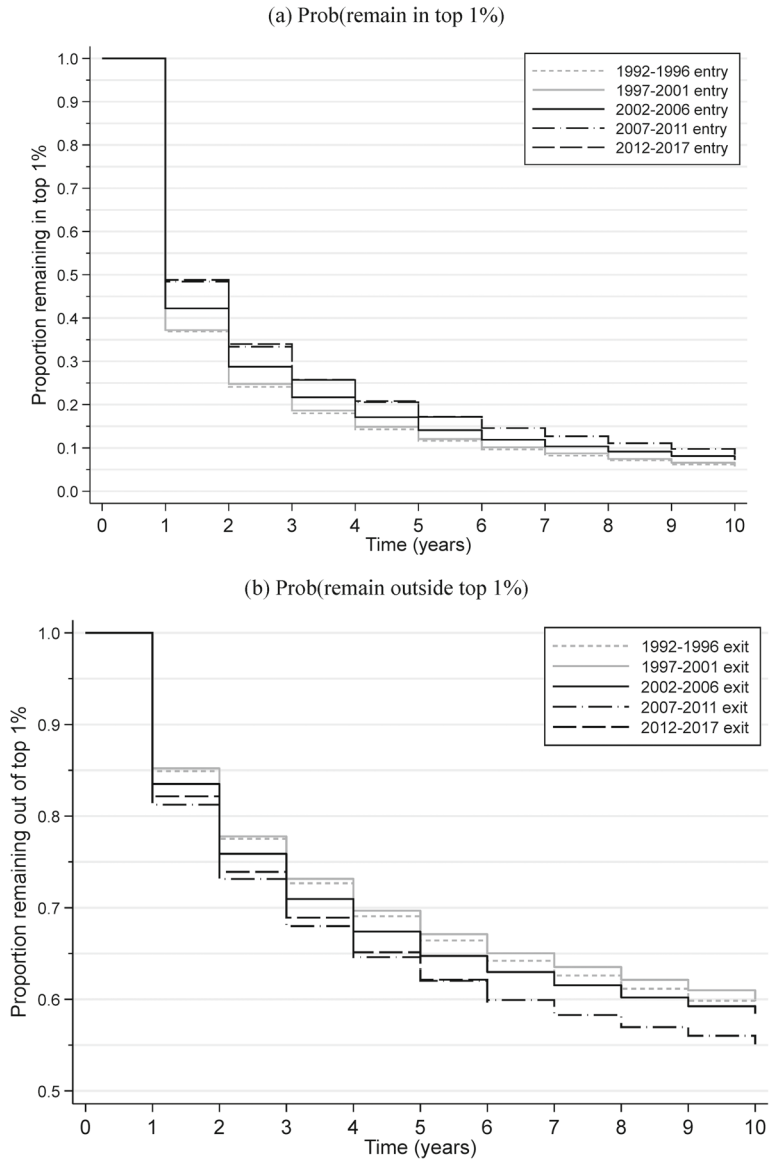

Figure 2. Probabilities of remaining in and outside the top 1%, by duration and top 1% cohort.

provided by this figure is consistent with that provided by Figure 1, but what about the additional role played by re-entry?

Panel (b) of Figure 2 shows that changing patterns of re-entry to the top 1 percent have reinforced top-income persistence. The probability of remaining outside the top 1 percent group for 3 years since having left the group is around 73 percent for the two earliest exit cohorts but around 68 percent for the two most recent cohorts. These cross-cohort differences in probabilities are apparent at all

durations as well. For example, the probability of remaining outside the top 1 percent for 10 years after exiting the group is around 60 percent for two earliest cohorts but only 55 percent for the 2007–2011 cohort.

3.4. *Comparisons of permanent and transitory top-income shares, and an immobility index*

Another approach to summarizing top-income persistence is to compare top-income shares calculated using income defined using a single-year measure with top-income shares calculated using a multi-year, longitudinally-averaged, measure. The greater the gap between the multi-year average of top-income shares of one-year income and the top-income share of multi-year income, the smaller is top-income persistence (the more top-income mobility there is).

The difference between the multi-year averaged one-year top income share and the multi-year top income share is an index of top-income immobility (persistence). This is the complement of Aaberge *et al.*'s (2013) index of top-income mobility, $T(u)$, for the top $100(1-u)$ percent.[14] (We examine—$T(u)$ because of our paper's focus on persistence, and take $u = 0.99$.) Our index is zero if there are no transitions into or out of the top 1 percent in the period. The more transitions there are, and the larger the income movements accompanying those transitions, the more negative (larger in absolute magnitude) is the immobility index. At one extreme, if the same individuals comprise the top 1 percent income group year after year, the immobility index equals zero. At the other extreme, if belonging to the top 1 percent in a specific year were entirely random and the time frame examined long enough, the top-income share for multi-year income would be 1 percent. This means that the upper bound for the mobility index is one minus the average of the yearly top 1 percent shares over the period, i.e. between approximately −5.5 and −8.5 over the period we examine.

Figure 3 presents estimates of top 1 percent shares for incomes longitudinally-averaged over multiple years, as well as estimates of the average of the yearly top 1 percent shares over the same years. Panel (a) presents estimates for a three-year time-frame, while panel (b) presents estimates over a 5-year time-frame. In addition, the figure shows estimates of the top-income immobility index, $-T(0.99)$, calculated using three- and five-year time frames as appropriate.

Figure 3 shows a clear upward trend in the three-year and five-year top 1 percent income shares. (The moving average of the cross-sectional yearly top 1 percent shares is a smoothed version of the estimates shown in Figure 1.) There were two short-lived reductions in both multi-year top 1 percent share series: between 2001 and 2003 and between 2008 and 2011 for the three-year top 1 percent share; and between 2002 and 2004 and between 2008 and 2012 for the five-year top 1 percent share.

The immobility index summarizes how the two income share series translate into top-income persistence. We see a decrease in top-income immobility between 1991 and 2001, followed by a clear upward trend up to 2012 and a slight reversal thereafter. Put differently, there was a decline in top-income persistence in the first

[14]See their equation (2.2).

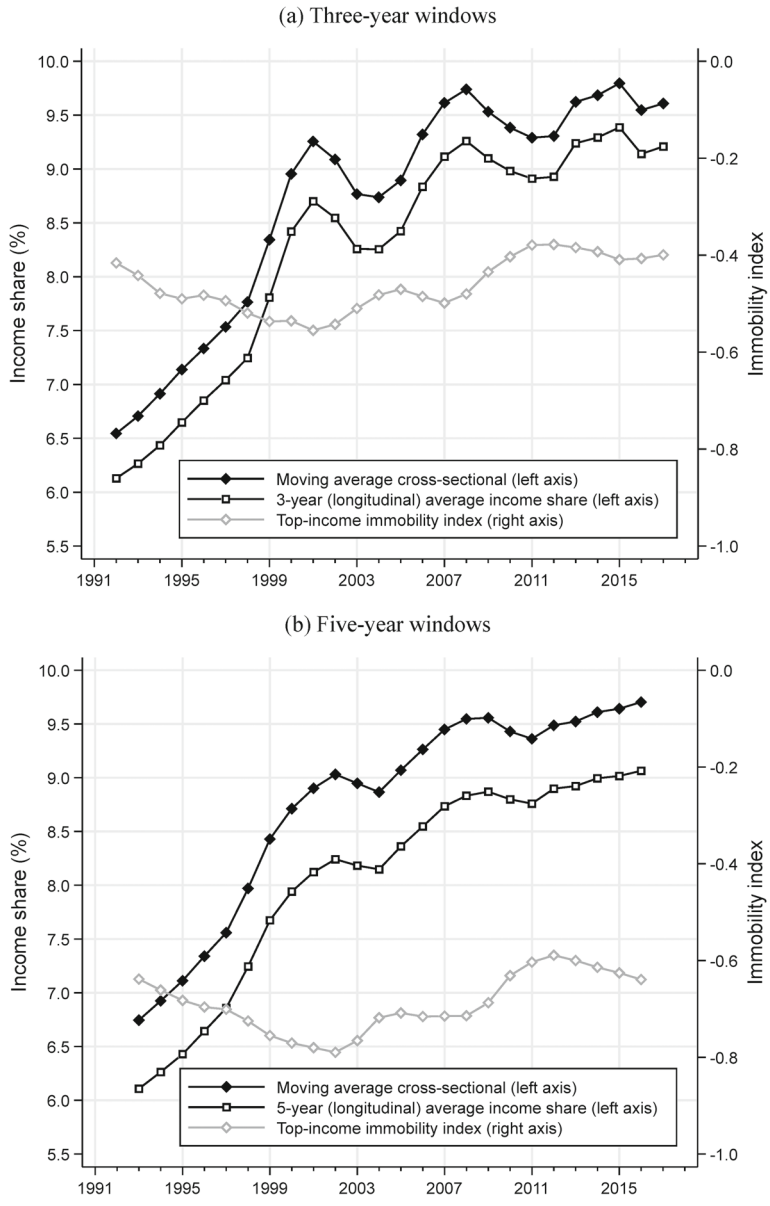

Figure 3. Permanent versus transitory top 1% shares, and immobility index: three- and five-year windows. The *x*-axis labels refer to the middle year of the period. The immobility index is the difference between the two top-income share lines, i.e. $-T$ (0.99), where T (u) is Aaberge *et al.*'s (2013) mobility index for the top 100 ($1-u$) percent.

decade followed by a rise in top-income persistence thereafter. This is a slightly different picture than provided by the top 1 percent persistence rate trends shown in Figure 1 because, although both perspectives show a rise in persistence from the 2000s on, the decrease shown by immobility index $-T(0.99)$ at the start of the period is only apparent for the top 1 percent persistence rate when a 1-year window is used.

There are no estimates for other countries to compare our immobility index with apart for those of Aaberge *et al.* (2013) for Norway. Our three-year estimates for Australia are similar to the values reported by Aaberge *et al.* (2013), Figures 4 and 5) in the 1970s and 1980s but much smaller in magnitude than they found for the 1990s and 2000s. Put differently, the secular trend in top-income immobility in Australia was upwards from around 2000 onwards, whereas top-income immobility in Norway fell steadily between 1990 and the late-2000s. We conjecture that the cross-national differences in persistence trends are related to differences in top-income share trends. As noted in Section 1, top income shares increased significantly more over the period in Norway than they did in Australia.

Returning to the Australian estimates in Figure 3, we observe that although the turning points in the top 1 percent share series and the immobility index series broadly coincide, their trends do not. For example, there are periods when top 1 percent shares increased but these coincide with both periods when the immobility index decreased and periods when it increased. This draws attention to a potential limitation of immobility index $-T(u)$. Periods of growth in the top 1 percent share are also periods in which there tend to be large income gains by the top 1 percent, and these gains can act to increase averaged 1-year top-income shares relative to multi-year top-income shares. That is, it is helpful to distinguish between (im)mobility defined in terms of absolute real income changes and (im)mobility that changes individuals' ranks (and hence also their top-income group membership status potentially). We return to this issue in Section 4.

Complementing the estimates for the top 1 percent in Figure 3, Appendix S1 Figures E5 and E6 present estimates for the top 0.1 percent and Appendix S1 Figures F5 and F6 show estimates for the top 10 percent. The immobility index for the top 0.1 percent, $-T(0.999)$, decreased in the 1990s, followed an inverted V-shape in the early 2000s, and then increased in the late 2000s before plateauing in the 2010s. The immobility index for the top 10 percent, $-T(0.90)$, shows some fluctuations over the period but no clear trend. In this case, the increase in immobility due to increased persistence rates may have been offset by the shift to the right in the income densities, meaning that the income loss required to drop out of the top 10 percent increased, leading to larger income movements but among fewer individuals. We discuss this point further in Section 4.

3.5. *Income group origins and destinations of top 1 percent group members*

We provide further information about levels and trends in top-income persistence by examining the previous-year income origins of entrants to the top 1 percent and the next-year income destinations of those who leave the top 1 percent. There is an analogy with poverty analysis: if a poor person with an income just below the poverty line last year has an income $500 above the poverty line this year, there is less low-income persistence of concern than if this year's income were only $1 above

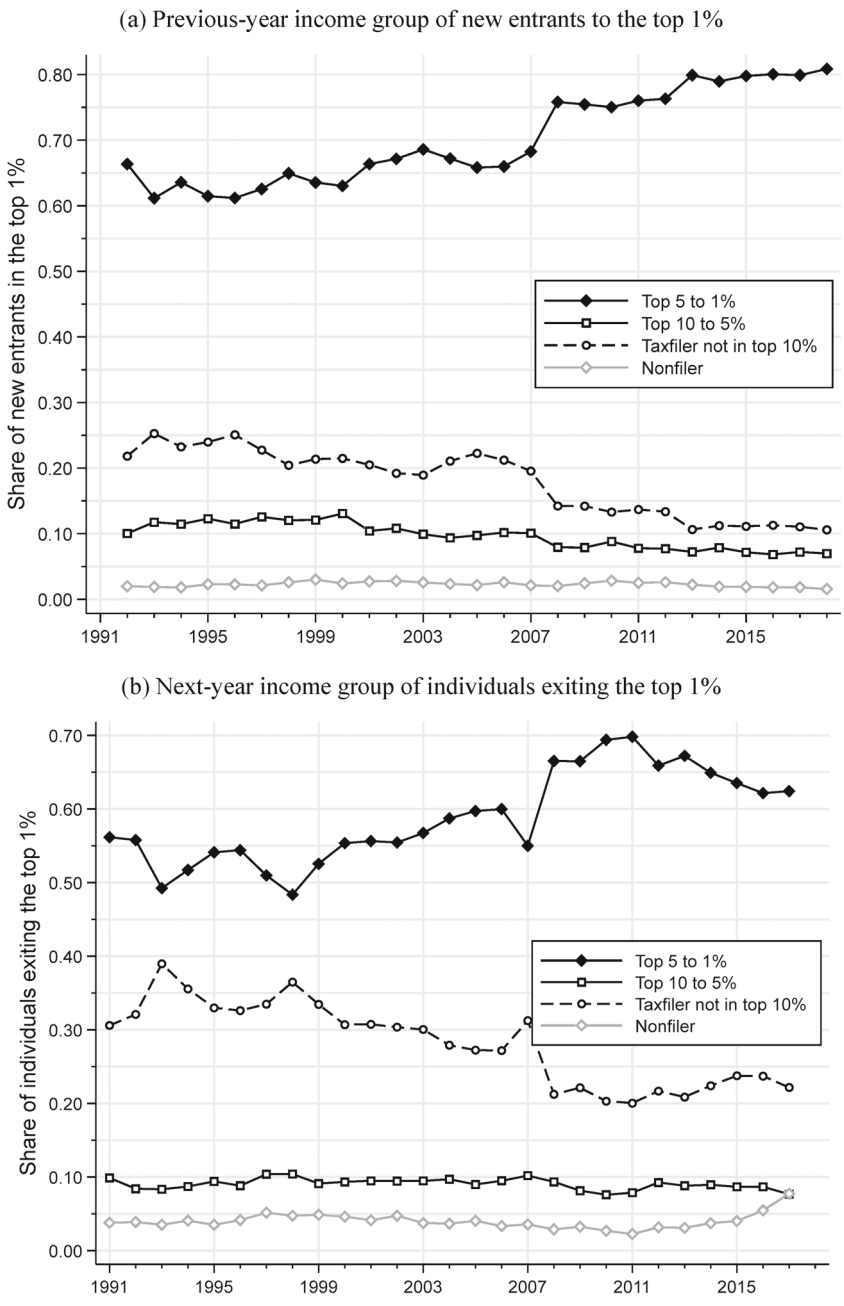

Figure 4. Origins and destinations of top 1% income group members. The individual is the unit of analysis. Estimates are based on the adult population (aged 15 or above).

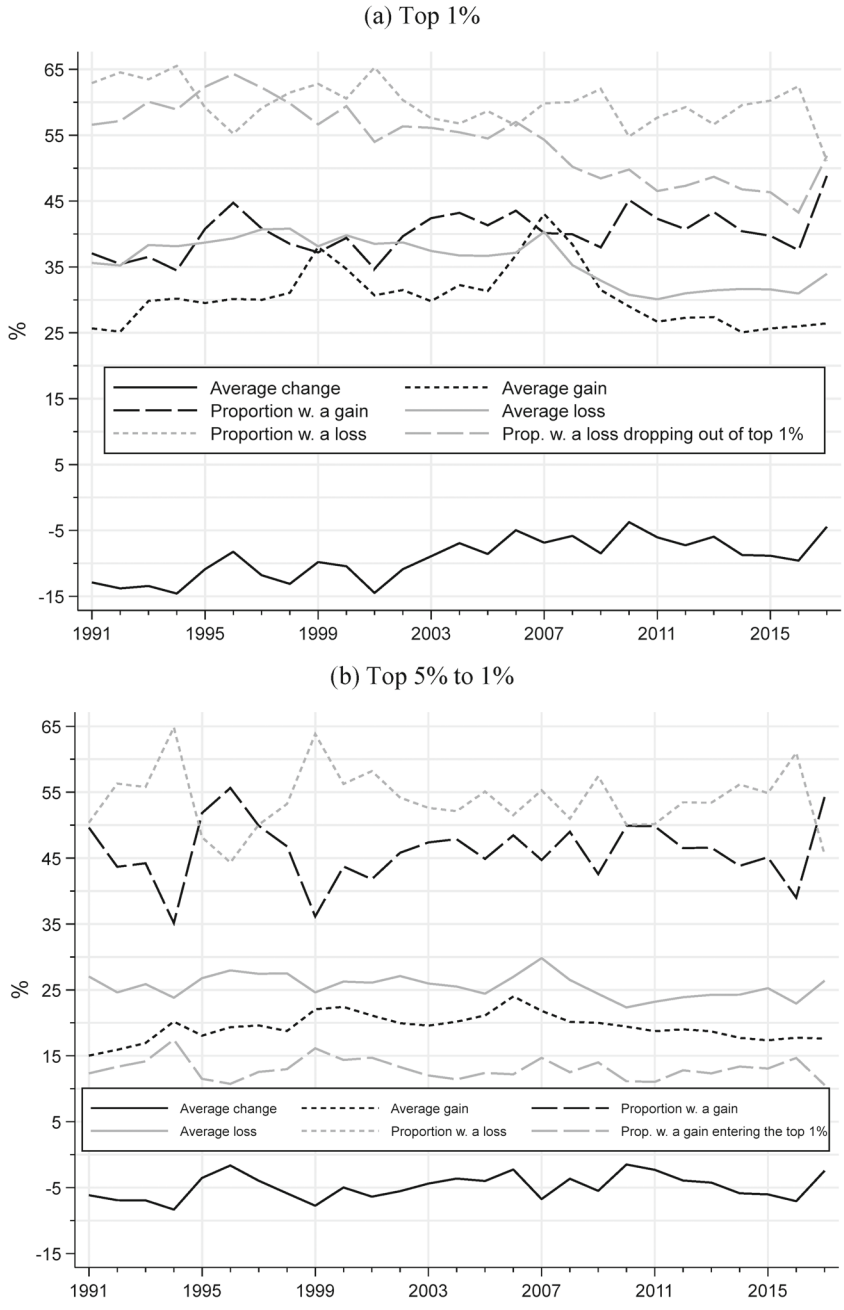

Figure 5. Average yearly income changes. Estimates for year t refer to income changes between year t and $t+1$ for those in the top 1% income group in year t. Average loss (gain) refers to the mean income loss (gain) among those with a loss (gain).

the poverty line. Analogously, if the incomes of the individuals leaving the top 1 percent are just below the threshold defining the top 1 percent rather than well below (or those entering the top 1 percent come from just below the top 1 percent), then the nature of top-income persistence is rather different.

Figure 4 shows the income groups that entrants to the top 1 percent belonged in the year prior to entry, and the income groups that those leaving the top 1 percent went in the year of exit. Consistent with our earlier evidence on increasing persistence, we see that those entering and leaving the top 1 percent are increasingly coming from and going to the top 5 percent to 1 percent group.

Panel (a) of Figure 4 shows that the proportion of top 1 percent entrants coming from the top 5 percent to 1 percent income group has risen from approximately 60 percent in 1993 to approximately 80 percent since 2013. (The years cited refer to the years of entry into the top 1 percent.) There was a particularly sharp rise between 2007 and 2008, from approximately 64 percent to 75 percent. The proportion entering from the top 10 percent to 5 percent has fallen from 25 percent in 1993 to just over 10 percent since 2013, while the proportion entering from lower down the income distribution but still lodging a tax return has declined from approximately 14 percent in 2000 to approximately 7 percent since 2011. The proportion of entrants who did not lodge a tax return in the previous year has remained relatively stable at approximately 2 percent.

Panel (b) of Figure 4 shows that the proportion of those leaving the top 1 percent who move to the top 5 percent to 1 percent rose between 1998 and 2011 (where the year refers to the last year before exit from the top 1 percent), from approximately 48 percent to 70 percent. However, since 2011 the proportion moving to the top 5 percent to 1 percent has declined somewhat, to approximately 63 percent in 2018. Similar to the finding for entry to the top 1 percent, there was a sharp rise in the proportion of those who exited to the top 5 percent to 1 percent between 2007 and 2008. Most of the rise in the proportion going to the top 5 percent to 1 percent has been at the expense of people moving to the top 10 percent to 5 percent group.

In sum, we find that not only has the probability of exit from the top 1 percent declined since 1991, but also that those entering the top 1 percent are increasingly coming from near the top of the income distribution, and those leaving the top 1 percent are increasingly staying near the top of the income distribution. This reinforces our general finding of a rise in top-income persistence, but also highlights the relevance of examining changes in real incomes in addition to changes in ranks. We look at income changes in more detail in the next section.

4. The Income Changes Associated with Changes in Top-Income Group Membership

We have documented a clear increase in top-income persistence. We have focused on positional (im)mobility but also pointed out that absolute income changes are relevant too. Building on this point, this section describes how the income changes affecting the top 1 percent have evolved over time, focusing on one-year income changes.

Figure 5 displays various measures of income changes from 1 year to the next of those in the top 1 percent and those in the top 5 percent to 1 percent. Panel (a) shows that, 1 year after being observed in the top 1 percent income group, individuals experienced an income loss of between 10 percent and 15 percent in the 1990s on average (conditional on having an income loss), but only between 5 percent and 10 percent in the 2000s and 2010s. This average loss is the net result of 50 percent to 65 percent of the top 1 percent experiencing an income loss (of 30 percent to 40 percent on average) and the remaining 35 percent to 50 percent experiencing an income gain (of 25 percent to 35 percent on average, among gainers). Some important fluctuations in these proportions are apparent. First, the average gain fluctuated between 25 percent and 43 percent over the period but it was back to its 1991 value of 26 percent by 2017. Second, the average income loss decreased, especially from 2007 to 2011, when it went from 40 percent to 30 percent.

Changes in the average one-year income decline among those in the top 1 percent experiencing an income decline appear to correlate with the fluctuations in the one-year persistence rates presented in Figure 1. Figure 1 shows that most of the increase in the persistence rates occurred between 2000 and 2004 and between 2007 and 2011. Figure 5 shows that both periods recorded a decrease in the average income decline of those experiencing a decline. Taken together, these findings suggest that reductions in the average income decline experienced by the top 1 percent have driven the increased persistence rate in that group to some extent. This effect can be mitigated (reinforced) by reductions (increases) in the proportion facing a decline, but Figure 5 shows that the proportion with a loss was in 2016 close to its 1991 level. Indeed, the figure also shows that the proportion of people in the top 1 percent experiencing an income decline that drops them out of the top 1 percent has steadily declined since 2006, from 57 percent to 43 percent in 2016 (though it picked up again in 2017 to 46 percent).

Another factor potentially contributing to higher persistence rates is that individuals below the 99th percentile ($p99$) experienced smaller and/or fewer income gains. Panel (b) of Figure 5 presents evidence for the top 5 percent to 1 percent income group. The findings for the 1991–2017 period as a whole are not conclusive. There was a decline in the average increase in income among those experiencing an increase since 2006 and the proportion of the top 5 percent to 1 percent experiencing an increase in income also declined between 2011 and 2016. However, there has been little net change in the proportion of this income group experiencing an income increase putting them in the top 1 percent.

Panel (a) of Figure 6 shows the distribution of income within the top 1 percent, plotting the density of relative distance to $p99$ in 1991, 2001, 2011, and 2018. (Relative distance is the percentage difference between observed income and $p99$.) A large proportion of top 1 percent income recipients have incomes just above $p99$, with the overwhelming majority having incomes that are at most 50 percent larger than $p99$. Given the increase in top-income persistence, we would expect these densities to have shifted to the right, indicating a move away from $p99$. Indeed, this type of shift occurred between 1991 and 2001, but there was no substantial shift thereafter, when persistence was increasing.

Together with Figure 1, these findings suggest that the increased persistence rates are unlikely to be explained by a "rungs of the ladder growing further apart"

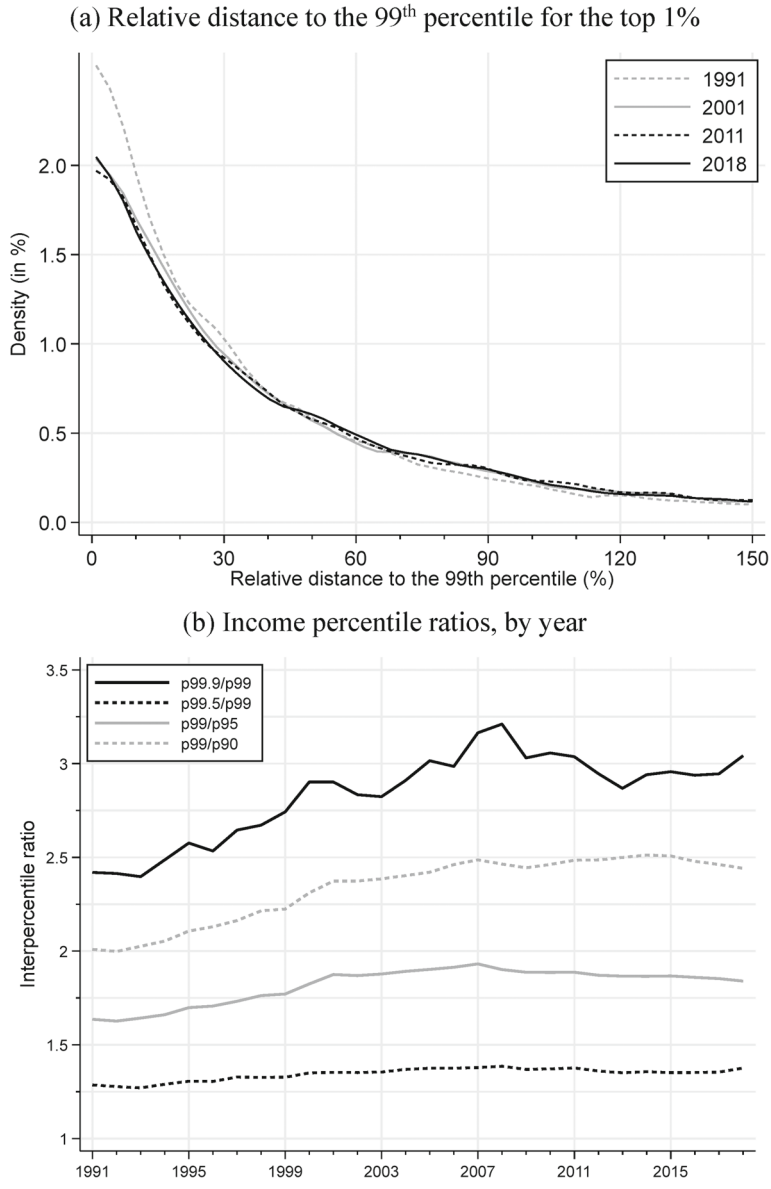

Figure 6. Relative income distance to the 99th percentile for the top 1%, and income percentile ratios. In panel (a), kernel density estimates are calculated for the distribution of relative income distance to $p99$ using an Epanechnikov kernel and bandwidth of 1. The densities are truncated at incomes 150% above $p99$. Relative income distance is the percentage gap between observed income and $p99$.

story but, instead, by a change in the nature of income dynamics. We look at the ladder rungs story further shortly.

Appendix S1 E provides evidence of a similar story for the top 0.1 percent, with Figure E5 showing a reduction in the average income loss and no major shift to the right for the income densities besides the 1991–2001 shift (Figure E6).[15] The story for the top 10 percent is somewhat different. Appendix S1 Figure F5 shows no clear downward trend in the average loss. These somewhat more stable income dynamics patterns are accompanied by a series of shifts to the right of the income densities (Figure F6).

Panel (b) of Figure 6 further probes whether the "rungs of the ladder" have been moving further apart. It presents the evolution of percentile ratios $p99.9/p99$, $p99.5/p99$, $p99/p95$, and $p95/p90$. Although all ratios have increased since 1991, the increases have been larger as we move closer to the top. Such growing disparities make positional top-income mobility more difficult as they increase the income changes required to move across top income groups, for example, from the top 1 percent income group to the top 0.1 percent income group. However, changes in the $p99/p95$ and $p99/p90$ ratios—the most relevant for top 1 percent persistence—have been small since 2000, and thus over the period when the increase in top-income persistence occurred. This is further evidence that the "rungs of ladder growing further apart" story is not the main explanation for the increased top-income persistence and documented in the preceding section.

5. Differences in Top-Income Persistence Across Population Subgroups

Having demonstrated a rise in top-income persistence, we wish to account for it. In this section, we examine differences in top-income persistence levels and trends across population subgroups. (Trends in the subgroup composition of the top 1 percent are shown in Appendix S1 C.) ALife contains a relatively limited set of characteristics but we are able to report breakdowns for subgroups defined by the characteristics most commonly used in top-income analysis, namely sex, main income source, age, occupation, and region (federal state of residence). In this section, we consider the extent to which the rise in top-income persistence can be accounted for by initial subgroup persistence rates, the initial subgroup compositions of the top-income group, and how each of these components changes over time.

To be specific, suppose the population is exhaustively partitioned into two subgroups. The top-income persistence rate (defined as for Figure 1) for the population is simply the weighted average of the persistence rates of the two subgroups where the weights are the subgroup fractions of the top income group. An increase in either subgroup's top-income persistence rate between 2 years will increase the population persistence rate (with the contribution of the increase contingent on the subgroup's share in the top income group). However, an increase in a subgroup's share in the top income group increases the overall persistence rate only if that

[15]The large increase in average income gains apparent in 2008 in Figure E7 is due to three outliers, each experiencing an income gain of more than $15 million.

subgroup's persistence rate is greater the persistence rate for the other subgroup; if not, the change in top-income composition has an offsetting effect on the overall persistence rate.[16]

Panels (a) and (b) of Figure 7 display persistence rates separately for individuals distinguished by whether their main source of income is labor income or non-labor income. Labor income comprises wages and salaries only, while non-labor income comprises all other income including income from investments, businesses, and government benefits. For this exercise, we treat the incomes of self-employed individuals as 70 percent labor income and 30 percent non-labor income following the Distributional National Accounts guidelines of Alvaredo *et al.* (2021, p. 52). We define the individuals with main income source as labor income as those with a labor income share of total income that is greater than one half in the base year of the T-year persistence rate calculation. The remaining individuals form the non-labor income group. This group is substantially smaller than the labor income group throughout the period, making up just under 25 percent of all adults in the top 1 percent in the early 1990 s and in the mid-2010 s, increasing to nearly 28 percent of the top 1 percent in 2007 (Appendix S1 Figure C1).

Panel (a) of Figure 7 shows persistence rates for individuals in the top 1 percent in base year t in the labor income group, whereas Panel (b) shows persistence rates for individuals in the non-labor income group. There is more volatility in persistence rates among individuals whose main source is non-labor income, reflecting greater volatility in business incomes and greater cyclicality in investment and business incomes than labor incomes. In particular, there were greater falls in persistence rates for individuals whose main income source was non-labor income around the time of the GFC (mid- to late-2000s). However, both groups experienced rising persistence rates over the 1991–2018 period as a whole.

The magnitude of the increase in persistence rate was greater for the non-labor income group if one looks at the period as whole. For example, the one-year persistence rate for the labor income group was 67 percent in 1991, rising to 75 percent in 2015, an increase of 8 percentage points (pat) or 12 percent. In contrast, for the non-labor income group, the 1-year persistence rate rose from 57 percent to 68 percent, an increase of 14 ppt or 20 percent. (Similar cross-group differentials are apparent for the one-year and five-year persistence rates.) However, looking at the period as a whole is potentially misleading because of the volatility of the non-labor income group series. For example, over the period between 2000 and 2018, the increase in the one-year persistence rate was around 10 ppt for the labor income group but, smaller, only around 4 ppt, for the labor income group.

Whatever the case, the increase in the non-labor income group's persistence rate makes a relatively small contribution to the increase in the population persistence rate because this group remains relatively small, never greater than 28 percent (Appendix S1 Figure C1).

[16]The population top-income persistence rate for a specific base year, $r = pr_1 + (1-p)r_2$ where p is the proportion of the top-income group from subgroup 1, $1 - p$ is the proportion of the top-income group from subgroup 2, and r_j is the persistence rate for subgroup $j = 1, 2$. Hence, the change over time in r, $\dot{r} = p\dot{r}_1 + (1-p)\dot{r}_2 + \dot{p}\left(r_1 - r_2\right)$.

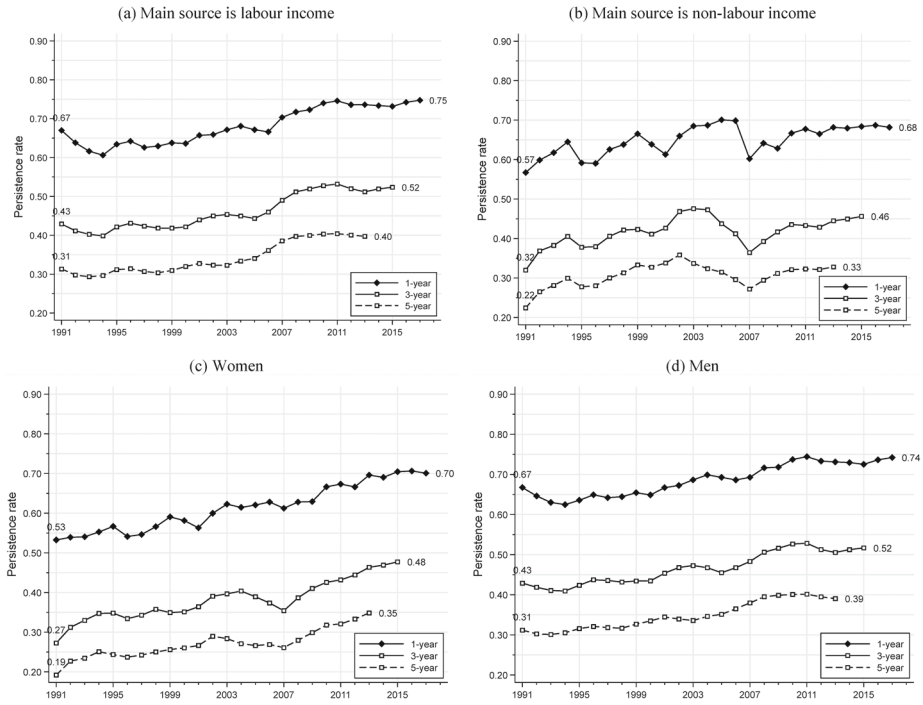

Figure 7. Persistence rates for the top 1%, by main source of income and by sex. The persistence rate is the share of individuals in the top 1% in the base year who appear every year in top 1% between the base year and then in the subsequent 1, 3 and 5 years. The definition of "main source" is in the main text.

Panels (c) and (d) of Figure 7 compare top 1 percent persistence rates for women and men, respectively. In the early 1990s, persistence rates were markedly higher for men than women, whether one looks at 1-, 3-, or 5-year rates. Thereafter, persistence rates grew both for men and women, but growth was greater for women, so sex differentials in persistence rates narrowed between 1991 and 2018. For example, the one-year persistence rate for men in the top 1 percent was 67 percent in 1991, rising to 74 percent in 2017, an increase of 7 ppt (or 21 percent). In contrast, for women in the top 1 percent, the one-year persistence rate rose from 53 percent to 70 percent, an increase of 17 ppt (78 percent). Similar cross-group differentials are apparent for the 3-year and 5-year persistence rates. Boschini *et al.* (2020) show that Sweden's experience was similar to Australia's: top-income persistence rates initially lower for women but converging toward the rates for men.

The difference between the sexes in increases in Australian top-income persistence rates does not account for the secular rise in the population persistence rate. The proportion of the top 1 percent that is female rose from around 17 percent in the early 1990s to almost 24 percent in 2018, with the greatest increase in the decade and a half prior to 2007 (Appendix S1 Figure C1), but remained substantially smaller than the fraction for men. This has two consequences. First, although the rise in top-income persistence rates for women was greater for women than men, it gets a low weight (relative to the weight given to the rise in the men's persistence

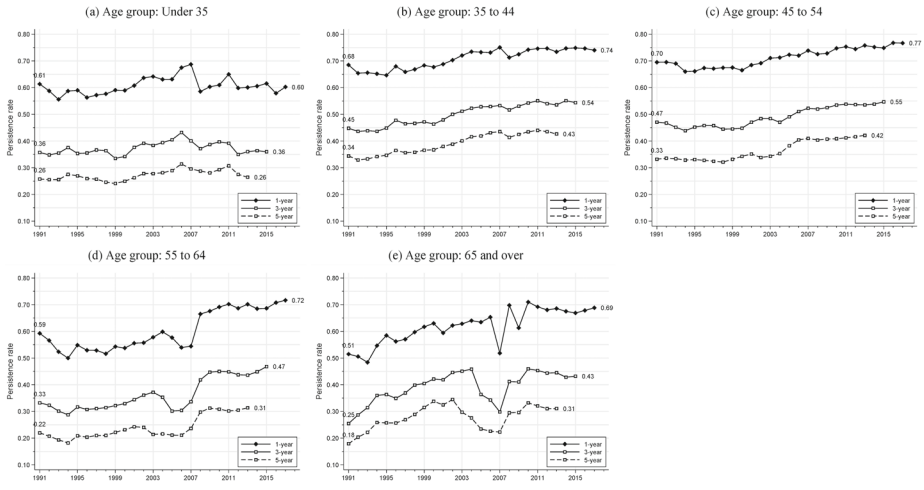

Figure 8. Persistence rates for the top 1% income group, by age group. The persistence rate is the share of individuals in the top 1% in the base year who appear every year in top 1% between the base year and then in the subsequent 1, 3 and 5 years.

rate). Second, because the top 1 percent persistence rate for women remained below that for men, the increasing fraction of women in the top 1 percent contributed to a reduction in the population persistence rate (see Footnote 15).

Figure 8 compares top 1 percent persistence rates across five age groups: under 35, 35–44, 45–54, 55–64, and 65+ years. Persistence rates rose between the early-1990 s and late-2010 s for all age groups with the notable exception of the youngest one. For the 35–44 and 45–54 age groups, the increase in one-year persistence rate over the period was 4 ppt and 7 ppt (around 10 percent in each case), but 13 ppt (22 percent) for the 55–64 year olds, and 18 ppt (35 percent) for those aged 65+ years. The first two groups had higher persistence rates than the two older groups at the start of the period, but the greater increase for the latter two groups led to a convergence of persistence rates across the oldest four age groups by the end of the period.

Figure 8 also shows that summarizing trends by comparing the beginning and end of the 1991–2018 period is potentially misleading. For the two oldest groups, there were relatively large decreases in persistence rates in the years just preceding the GFC (before 2007) followed by relatively large increases. There are also interesting changes in the age group composition of the top 1 percent before and after 2007 (Appendix S1 Figure C2). For example, in the years prior to 2007, the 35–44 and 45–54 age groups formed the highest shares of the top 1 percent (both with shares around 30 percent) but, subsequently, the share for the 35–44 group declined by more than 5 ppt over the subsequent decade, whereas the share for the 45–54 group rose by almost the same amount. The 55–64 age group had the third largest share of the top 1 percent throughout the period, rising steadily from around 15 percent to 23 percent (albeit with some fluctuation just before and after 2007). Hence, by contrast with the breakdowns by sex, we see that the changes in age group composition of the top 1 percent are such that the age groups with the largest increases

in top 1 percent persistence rates are given the largest weights in the calculation of (changes in) the persistence rate for all individuals.

We have also undertaken breakdowns by occupation and state of residence (available from the authors on request), and these reveal no clear differences in top-income persistence trends across subgroups. Overall, this section suggests that factors affecting the middle three age groups are the main correlates of the rise in the overall top-income persistence rate. We examine this hypothesis further in the next section.

6. The Rise of Top Income Persistence: A Decomposition Analysis

We now turn to regression analysis to examine in more detail the factors associated with the rise in top-income persistence. In the previous section, we looked at characteristics, variously defined, in turn, rather than simultaneously as we can with regression. We found that for most subgroups with increasing persistence rates there were concomitant subgroup changes in top 1 percent composition that offset them. The exceptions were the breakdowns by age group.

We use an Oaxaca-Blinder style decomposition approach extended to non-linear regressions to account for the rise in the probability of staying in the top 1 percent, and the fall in the probability of joining the top 1 percent. We carry out the decompositions separately for two subperiods, 1991/1992 to 2000/2001 and 2001/2002 to 2017/2018, because we have shown that top income persistence was mostly stable in the 1990 s but increased in the 2000 s.[17]

In the top 1 percent stayer regressions, we are looking to see the extent to which changes over time in the stayer probability is accounted for by (i) changes in the characteristics of the stayers, or by (ii) changes in the "returns" to these characteristics (changes in coefficient estimates), i.e. changes in conditional persistence rates, broadly speaking (see the discussion of average partial effects [APEs] below). For the top 1 percent joiner regressions, there is an analogous breakdown.

6.1. *Methods*

We suppose that the probability of each individual i in the top 1 percent in year $t-1$ of staying in the top 1 percent in year t, P_i^t, is a non-linear function of individual i's characteristics, i.e.

$$P_i^t = F\left(X_i^t \widehat{\beta}^t\right) \tag{1}$$

where X_i^t is a vector of year t characteristics for i, and $\widehat{\beta}^t$ is the vector of estimated returns associated with each of those characteristics.[18] Similarly, for the joiner

[17]Estimates for the 1991/1992 to 2017/2018 period as a whole are available upon request. Oaxaca-Blinder style decompositions are descriptive in nature and estimates are easier to interpret when the period analyzed comprises one unidirectional change in the outcome variable—the case for our two subperiods.

[18]Because staying probabilities refer to persistence between years $t-1$ and t, we measure "year t" characteristics using year $t-1$ values. Referring to t rather than $t-1$ simplifies notation. We use the same notational convention for the decomposition of changes in joiner probabilities.

model, P_i^t, is the probability of each individual i not in the top 1 percent in year $t-1$ of entering the top 1 percent in year t.[19]

Taking averages over individuals, we write the change in the top 1 percent entry and exit probabilities between years s and t as $\Delta\overline{P} = \overline{P}^t - \overline{P}^s$. From equation (1), this change can be written as:

(2a)

$$\overline{P}^t - \overline{P}^s = \left[\sum_{i=1}^{N^t} \frac{F\left(X_i^t \widehat{\beta}^s\right)}{N^t} - \sum_{i=1}^{N^s} \frac{F\left(X_i^s \widehat{\beta}^s\right)}{N^s}\right] + \left[\sum_{i=1}^{N^t} \frac{F\left(X_i^t \widehat{\beta}^t\right)}{N^t} - \sum_{i=1}^{N^t} \frac{F\left(X_i^t \widehat{\beta}^s\right)}{N^t}\right], \text{ or}$$

(2b)

$$\overline{P}^t - \overline{P}^s = \left[\sum_{i=1}^{N^t} \frac{F\left(X_i^t \widehat{\beta}^t\right)}{N^t} - \sum_{i=1}^{N^s} \frac{F\left(X_i^s \widehat{\beta}^t\right)}{N^s}\right] + \left[\sum_{i=1}^{N^s} \frac{F\left(X_i^s \widehat{\beta}^t\right)}{N^s} - \sum_{i=1}^{N^s} \frac{F\left(X_i^s \widehat{\beta}^s\right)}{N^s}\right].$$

$\Delta\overline{P}$ = changes in characteristics + changes in returns to characteristics

Thus, the change in the top-income persistence rate between years s and t is decomposed into two components: the first reflects changes in characteristics; the second reflects changes in the returns to characteristics, i.e. coefficients. (An analogous decomposition is applied to changes in the probability of entering the top 1 percent.) The decomposition expression is not unique. Equation (2a) states that changes in characteristics are weighted by initial-year coefficients and changes in coefficients by final-year average characteristics. Equation (2b) states that the changes in characteristics are weighted by final-year coefficients and changes in coefficients by initial-year average characteristics. We check the sensitivity of our decomposition estimates by undertaking our calculations separately for the two variants. The initial and final years (s and t) depend of course on which of the two subperiods we are considering.

The specific decomposition method we use is that of Fairlie (2005) because it straightforwardly allows us to assess not only the total contributions to trends of characteristics and of returns to characteristics but also the separate contributions of subsets of characteristics. We take $F(.)$ to be the logit function, randomize the order of the variables in the detailed decompositions, and use 100 replications.

Our choice of characteristics (X) is constrained because ALife contains a limited number of personal characteristics. We include sex, age groups (five categories), state of residence (six categories), self-employment status and whether labor income is the main source of income. For the top 1 percent stayer model we also include normalized income rank within the top 1 percent, and for the joiner model we include the relative income distance to $p99$. (We use relative income distance for the joiner model because relative rank cannot be calculated—we do not have the full population in ALife.) Characteristic means for the initial and final years of the two

[19]The estimation sample for the top-income joiner model excludes individuals who do not have ALife records in year $t-1$, i.e. non-tax filers.

subperiods are in Appendix S1 Tables G1 and G2. The logit coefficient estimates are in Appendix S1 Tables G3 and G6 for the stayer and joiner models, respectively. (We report robust standard errors throughout.) Because logit regression coefficients are difficult to interpret, we also present the corresponding APEs. These are conditional probabilities of one-year top-income persistence (and top-income entry). See Appendix S1 Tables G4 and G5, and G7 and G8, for the stayer and joiner regression models, respectively. It is conventional to evaluate APEs in a year t outcome regression using the characteristics corresponding to that year. We also present year t APEs calculated using characteristics defined at year s so that we have APEs corresponding to the predicted probability expressions implied by the decomposition expressions in equation (2a) and (2b). This also provides an additional sensitivity check to the choice of decomposition weights.

6.2. *Decomposition estimates*

We report our Fairlie decomposition estimates for the one-year stayer probability in Table 3. Panel (a) refers to the first subperiod and panel (b) to the second. There are two sets of estimates for each subperiod, reflecting whether we use formula (2a) or (2b).

Over the first subperiod, the one-year persistence rate remained much the same, around 64 percent. Panel (a) shows that this is the outcome of small changes in coefficients and (average) characteristics, with the relative size and sign of these components sensitive to the weighting scheme. Nonetheless we also see that changes in coefficients on the age group variables worked to increase persistence regardless of the weighting variant.

The second subperiod is more interesting because it covers when the one-year persistence rate increased markedly, from 64.7 percent in 2001/02 to 73.3 percent in 2017/2018. Of this 8.5 ppt increase, changes in characteristics explain between 0.904 ppt and 0.002 ppt only, depending on whether equation (2a) or (2b) is used for the calculation. Hence, changes in coefficients account for virtually all the increase in the persistence rate. More specifically, changes in the estimated coefficients on the age group dummies alone explain between 7.4 ppt and 8.2 ppt, depending on which variant of the decomposition formula we use.

An obvious follow-up question is: are there specific age (or other) groups for which coefficients changed over the second subperiod or were all groups equally affected? To answer this question, we calculate APEs, which are in effect group-conditional persistence rates. We present the estimates in Appendix S1 Table G4 for the first subperiod and G5 for the second subperiod. (Results are unaffected by which variant we calculate.) For brevity, we focus on the second subperiod. In our regressions, the reference age group is "aged less than 35" and, for this group, the one-year persistence rate changed little over the second subperiod (Figure 8). Table G5 shows that, relative to this baseline, APEs for all older age groups increased markedly. For the 35–44 group, the increase was just over 5 ppt and for the 45–54 group just over 7 ppt. Around the start of the second subperiod (2000/2001), the 55–64 group's one-year persistence rate was below that for the "less than 35" group but was higher by the end of the subperiod. Correspondingly, there is a substantial increase in the 55–64 group's APE of just

Table 3
Decomposition of the Change in the Probability of Staying in the Top 1% in Year $t+1$

(a) 1991/1991–2000/2001	Equation (2a)		Equation (2b)	
Probability of staying in the top 1% in year $t+1$ (in %)	64.3		63.7	
	Contribution (ppt)	SE	Contribution (ppt)	SE
Sex	0.027	0.011	0.022	0.008
Age	−0.539	0.076	0.117	0.078
State of residence	−0.671	0.089	−0.336	0.117
Employment*	−0.343	0.051	−0.164	0.045
Normalized income rank	−0.047	0.034	0.003	0.019
Changes in characteristics (all)	−1.573		−0.358	
Changes in coefficients (all)	0.884		−0.332	
Change in age coefficients only	0.469		1.130	
Total change (ppt)	−0.689		−0.689	
(b) 2001/2002–2017/2018	Equation (2a)		Equation (2b)	
Probability of staying in the top 1% in year $t+1$ (in %)				
	0.241	0.042	0.044	0.027
Sex	0.531	0.098	−0.482	0.102
Age	0.126	0.094	0.327	0.068
State of residence	0.025	0.015	0.065	0.025
Employment*	−0.019	0.022	0.047	0.030
Normalized income rank	0.904		0.002	
Changes in characteristics (all)	7.621		8.523	
Changes in coefficients (all)	8.175		7.404	
Change in age coefficients only	8.525		8.525	

Notes: Oaxaca-Blinder decomposition extended to a logit model by Fairlie (2005). *Includes the dummy for self-employment status and the dummy for labor as the main source of income. See Appendix tables for estimates of logit coefficient and average partial effects. *Source*: Authors' calculations based on ALife data.

over 15 ppt. Table G5 shows that APEs for no other characteristics changed as much as those for age.

In sum, the regression decomposition estimates confirm the conclusions of the univariate analysis, i.e. to account for the increase in one-year top-1 percent persistence rates over the second subperiod, we need to look at factors that were particularly relevant for the age groups other than the youngest or oldest ones.

Before engaging in this discussion, we report our decomposition estimates for the top 1 percent joiner probability. See Table 4, organized by subperiod and calculation variant as Table 3 is. Like the stayer probability, the joiner probability hardly changed between 1991/1992 and 2000/2001. However, there was a slight fall in the joiner probability in the second subperiod, from 0.50 percent in 2001/2002 to 0.38 percent in 2017/18, consistent with increasing top-1 percent persistence. In line with the change in stayer probability decomposition, changes in individual characteristics explain little of this drop. Instead, the 0.12 ppt decrease in the probability of entering the top 1 percent is explained by changes in coefficients almost entirely, with changes in the age coefficients explaining about one third of the decrease (according to both calculation formulae).

TABLE 4
DECOMPOSITION OF THE CHANGE IN THE PROBABILITY OF JOINING THE TOP 1% IN YEAR $t+1$ FOR NON-TOP 1% TAX FILERS IN YEAR $t-1$

(a) 1991/1991–2000/2001	Equation (2a)		Equation (2b)	
Probability of joining the top 1% in year $t+1$ (in %)	0.505 Contribution (ppt)	SE	0.522 Contribution (ppt)	SE
Sex	0.003	0.001	0.002	0.001
Age	−0.008	0.001	−0.006	0.002
State of residence	0.007	0.001	0.004	0.002
Employment*	−0.009	0.002	−0.007	0.003
Relative distance to $p99$	0.132	0.003	0.180	0.004
Changes in characteristics (all)	0.125		0.172	
Changes in coefficients (all)	−0.108		−0.155	
Change in age coefficients only	−0.025		−0.019	
Total change (ppt)	0.017		0.017	
(b) 2001/02–2017/18	Equation (2a)		Equation (2b)	
Probability of joining the top 1% in year $t+1$ (in %)	0.504 Contribution (ppt)	SE	0.384 Contribution (ppt)	SE
Sex	0.008	0.001	0.003	0.001
Age	−0.011	0.002	−0.002	0.001
State of residence	0.002	0.001	0.006	0.001
Employment*	−0.007	0.002	−0.003	0.001
Relative distance to $p99$	0.010	0.003	0.001	0.001
Changes in characteristics (all)	0.002		0.005	
Changes in coefficients (all)	−0.122		−0.125	
Change in age coefficients only	−0.034		−0.033	
Total change (ppt)	−0.120		−0.120	

Notes: Oaxaca-Blinder decomposition extended to a logit model by Fairlie (2005). * Includes the dummy for self-employment status and the dummy for labor as the main source of income. See Appendix tables for estimates of logit coefficient and average partial effects. *Source*: Authors' calculations based on ALife data.

APEs for the joiner model in the second subperiod are in Appendix S1 Table G8 (G7 for the first subperiod). The most distinct age group-related decline in APE is for the 55–64 age group (around 0.23 ppt) but there are also some other notable declines. For example, the APE for "main income source is labor income" increased by around 0.30 ppt.

In sum, there are echoes in these results of the stayer probability decompositions, including the relevance of age-related changes. Because the changes in joiner probabilities over the second subperiod are so small, we concentrate the remainder of our discussion in this section on top-income persistence per se.

6.3. *What underlies the age-related changes in top-income persistence?*

Our answers to this question are inevitably speculative; it is hard to identify cause and effect, and ALife is limited in the information it contains. To us, the explanation with greatest plausibility as a "smoking gun" is the marked change in the effective age of retirement in Australia over the last two decades, i.e. the average age at which individuals withdraw from the labor force. Swoboda (2016, Figure 1) reports that the effective retirement age for both men and women fell substantially

over the 1970s and 1980s, flattened out over the 1990s, increased around 2000, and flattened out again from around 2010. The timing of these changes matches up with the patterns of trends in top-income persistence: see, e.g. Figure 1. Relatedly, Wilkins *et al.* (2021) show using Household, Income and Labour Dynamics in Australia (HILDA) data that, between 2001 and 2019, the mean age at retirement increased from 61.6 to 66.6 for men and from 61.0 to 65.1 for women. Perhaps more striking is that the proportion retiring between the ages of 55 and 64 fell from 54 percent to 25 percent for men, and from 56 percent to 36 percent for women. For men, the increase in mean age at retirement is steady whereas, for women, the increase in mean age at retirement of recent retirees is only evident from 2007.

Our argument is that retirement typically moves better-off workers from reliance on labor market income to non-labor income, notably private pension income (called "superannuation" in Australia), and is associated with a decrease in total taxable income as well as a possible increase in volatility depending on the asset base of an individual's superannuation. Such decreases in income levels and increases in volatility are likely to lower the chances of top-income earners remaining in the top-income group. Putting the argument in reverse, if more individuals defer retirement to a later age, we would expect rising top-income persistence (as we have found). This story tallies with our finding that rising persistence was associated specifically with increases in persistence for individuals aged 45–54 and especially for those aged 55–64 (and no effects for those aged less than 35 or already over retirement age, then 65). However, the story does not explain why we also find rising persistence in the second subperiod for those aged 35–44.

Arguing for the influence of rising average retirement ages raises the further question of why this trend occurred. One argument is that the Australian labor market has become more accommodating of older workers—the types of jobs that people work in is important to them and there has been a shift away from physically taxing jobs—and the health of older people has been improving. Consistent with this argument, Wilkins *et al.* (2021, p. 124) report that the reasons for retirement given by HILDA survey respondents have shifted dramatically from "forced to" (by employer or poor health) to "wanted to" (financially ready, etc.).

A second potential influence is that Australia has been progressively increasing the age at which individuals can access their superannuation (the "preservation age"). For example, for those born before July 1, 1960, the preservation age is 55 years, and increases by 1 year for each succeeding year of birth to reach 60 years for those born after July 1, 1964 (Australian Tax Office 2022). However, the effects may not be large over the period we examine because the preservation age has been increasing only since 2016 (only in 2016 did someone need to be at least 56 instead of 55 to have reached the preservation age).

Third, there have been changes to Australia's public retirement pension (the "age pension"). Although this is income- and asset-tested, and so not obviously relevant for top-income earners, there may be an indirect effect on their retirement decisions if raising the eligibility age changes general perceptions of what an appropriate or "normal" retirement age is. Indeed, many people refer to the age pension eligibility age as "the retirement age". The age at which individuals may access the age pension is 65 for those born before June 30, 1952 but increases in steps through to 67 for those born in 1957 or later (Department of Social Services, 2022). In the past, women could access the age pension at age 60 but, since the 1990s, the

eligibility age has been increased. By 2013, it was the same age as for men (Super-Guide 2022). This is consistent with the steeper increase in top-income persistence rates for women compared to men, accompanied by their convergence (Figure 7).

The GFC may also have affected average retirement ages. (Note the marked increase in top-income persistence rates in the 5 years after 2007: see Figure 1c.) The GFC did not have a large impact on the labor market in Australia, with the seasonally-adjusted unemployment rate only rising from 4 percent in August 2008 to a peak of 5.9 percent in May 2009 and thereafter steadily declining until mid-2011 (Australian Bureau of Statistics (ABS), 2022). However, the GFC did have a large negative impact on capital asset values, including the value of superannuation assets, and spurred all income earners, including the richest ones, to remain in the labor force. By hitting non-labor incomes harder than labor incomes, the GFC may also have led to an increase in the fraction of individuals whose principal income source is labor income (note the 2007 peak in the fraction whose principal income source is non-labor income, shown in Appendix S1 Figure C1).

7. Summary and Conclusions

We have documented top-income persistence levels and trends for Australia using multiple perspectives, exploiting the advantages of newly available longitudinal income tax data. We find a distinct rise in top-income persistence over the 1991–2018 period with most of the increase occurring in the mid-2000s and late-2010s. That is, over a period in which yearly top-income shares were rising, there was also a growing inequality in the opportunities to have a top income.

Around the mid- to late-2010s, top-income persistence in Australia was towards the top of the range found for other countries with estimates derived from tax data—though we would reiterate that cross-national differences are difficult to assess because of comparability issues. Improving comparability in terms of definitions, samples, and time periods covered is an important task for future cross-national research.

Also important for future research is further analysis of the drivers of top-income persistence. Our decomposition analyses show that changes in the subgroup composition of the top 1 percent do not account for the rise in top-income persistence; rather, it is accounted for by changes in subgroup persistence rates, notably those for individuals aged 35–64, especially 55–64 years. We suggest that these effects are partially related to increases in the effective retirement age over the relevant period but acknowledge that this conclusion is speculative and does not explain rising persistence for younger age groups. An in-depth account of why top-income persistence has risen in Australia is a story yet to be told.

References

Aaberge, R., Atkinson, A. B., and Modalsli, J., 2013 The Ins and Outs of Top Income Mobility, Discussion Paper No. 7729, IZA, Bonn.

Alvaredo, F., A. B. Atkinson, T. Blanchet, L. Chancel, L. Bauluz, M. Fisher-Post, I. Flores, B. Garbinti, J. Goupille-Lebret, C. Martínez-Toledano, M. Morgan, T. Neef, T. Piketty, A.-S. Robilliard,E. Saez, L. Yang and G. Zucman, *Distributional National Accounts Guidelines: Methods and Concepts Used in the World Inequality Database*, World Inequality Lab, Paris School of Economics, Paris, 2021.

Atkinson, A. B. and A. Leigh, "The Distribution of Top Incomes in Australia," *Economic Record*, 83, 247–61, 2007.

Atkinson, A. B., T. Piketty and E. Seaz, "Top Incomes in the Long Run of History," *Journal of Economic Literature*, 49, 3–71, 2011.
Australian Bureau of Statistics (ABS). (2022) *Labour Force, Australia, May 2022*, ABS Catalogue Number 6202.0, ABS, Canberra, 2022. https://www.abs.gov.au/statistics/labour/employment-and-unemployment/labour-force-australia/latest-release, [accessed on 24 June 2022].
Australian Tax Office. (2022) *When you can access your super*. https://www.ato.gov.au/individuals/super/in-detail/withdrawing-and-using-your-super/withdrawing-your-super-and-paying-tax/?page=2, [accessed on 24 June 2022].
Auten, G. and G. Gee, "Income Mobility in the United States: New Evidence from Income Tax Data," *National Tax Journal*, 57, 301–28, 2009.
Auten, G., G. Gee and N. Turner, "Income Inequality, Mobility, and Turnover at the Top in the US, 1987–2010," *American Economic Review: Papers & Proceedings*, 103, 168–72, 2013.
Boschini, A., K. Gunnarsson and J. Roine, "Women in Top Incomes: Evidence from Sweden 1974–2013," *Journal of Public Economics*, 181, 104115, 2020.
Burkhauser, R. V., M. H. Hahn and R. Wilkins, "Measuring Top Incomes using Tax Record Data: a Cautionary Tale from Australia," *Journal of Economic Inequality*, 13, 181–205, 2015.
Burkhauser, R. V., M. H. Hahn and R. Wilkins, "Transitioning from an Historical to a Contemporary Use of Tax Record Data for Measuring Top Incomes in Australia," *Economic Papers*, 37, 1–33, 2018.
Department of Social Services. 2022. *Australian Government Age Pension*. https://www.dss.gov.au/seniors/benefits-payments/age-pension, [accessed on 24 June].
Drechsel-Grau, M., Peichl, A., Schmieder, J., Schmid, K.D., Walz, H. and Wolter, S.. 2022 Inequality and income dynamics in Germany (No. w29818). National Bureau of Economic Research.
Fairlie, R.W., "An Extension of the Blinder-Oaxaca Decomposition Technique to Logit and Probit Models," *Journal of Economic and Social Measurement*, 30, 305–316, 2005.
Jenderny, K., "Mobility of Top Incomes in Germany," *Review of Income and Wealth*, 62, 245–65, 2016.
Joyce, R., Pope, T., and Roantree, B. (2019) The Characteristics and Incomes of the Top 1%, Briefing Note BN254, Institute for Fiscal Studies, London.
Kopczuk, W., E. Saez and J. Song, "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937," *Quarterly Journal of Economics*, 125, 92–128, 2010.
Landais, C., *Top Incomes in France: Booming Inequalities?*" Unpublished paper, Paris School of Economics, Paris, 2008.
Martínez, I. Z. (2017) *Who is a Top Earner and for How Long? Top Income Mobility in Switzerland*, World Conference, Paris.
Morelli, S., T. Smeeding and J. Thompson, "Post-1970 Trends in Within-Country Inequality and Poverty: Rich and Middle Income Countries," in A. B. Atkinson, and F. Bourguignon (eds), *Handbook of Income Distribution*, Elsevier-North Holland, Amsterdam, 593–696, 2015.
Saez, E. and M. R. Veall, "The Evolution of High Incomes in Northern America: Lessons from Canadian Evidence," *American Economic Review*, 95, 831–49, 2005.
Shorrocks, A. F., "Income inequality and income mobility," *Journal of Economic Theory*, 19, 376–93, 1978.
SuperGuide. (2022) *What is the Age Pension Age?*. https://www.superguide.com.au/in-retirement/current-age-pension-age, [accessed on 24 June 2022].
Swoboda, K., "Influencing Decisions to Retire," *Key Issues for the 45th Parliament*, Department of Parliamentary Services, Parliament of Australia, Canberra, 32–3, 2016.
Wilkins, R., Vera-Toscano, E., Botha, F., and Dahmann, S. C. (2021) The Household, Income and Labor Dynamics in Australia Survey: Selected Findings from Waves 1 to 19. Report commissioned by the Australian Government Department of Social Services, Melbourne Institute of Applied Economic and Social Research, Melbourne.

Supporting Information

Additional supporting information may be found in the online version of this article at the publisher's web site:

Appendix S1. Supporting Information.