

Inequality in Russia over Time and over the Life Cycle*

Maksym Bryukhanov[†] Dmytro Hryshko[‡]

Abstract

Using Russian longitudinal data for 1994–2018, we document a secular decline in consumption and income inequality. Although within-cohort inequality is also declining, the lifecycle inequality profiles of income and consumption are surprisingly flat. A calibrated lifecycle model with incomplete markets, high initial variance of the persistent income component, and moderately persistent income shocks is consistent with nearly flat lifecycle inequality profiles and the puzzlingly large insurance role of assets found in the Russian data. This is in contrast to the standard calibrations that fail to match the lifecycle inequality profiles and the panel-data evidence on consumption insurance.

Keywords: consumption insurance, emerging markets.

JEL Classification: D12, D15, D31, E21, P36

*Support from the SSHRC IG Grant No. 435-2018-1275 and of the Basic Research Program (“Project 5-100”) of the National Research University Higher School of Economics are gratefully acknowledged. We thank Giulio Fella for fruitful conversations at an early stage of the project.

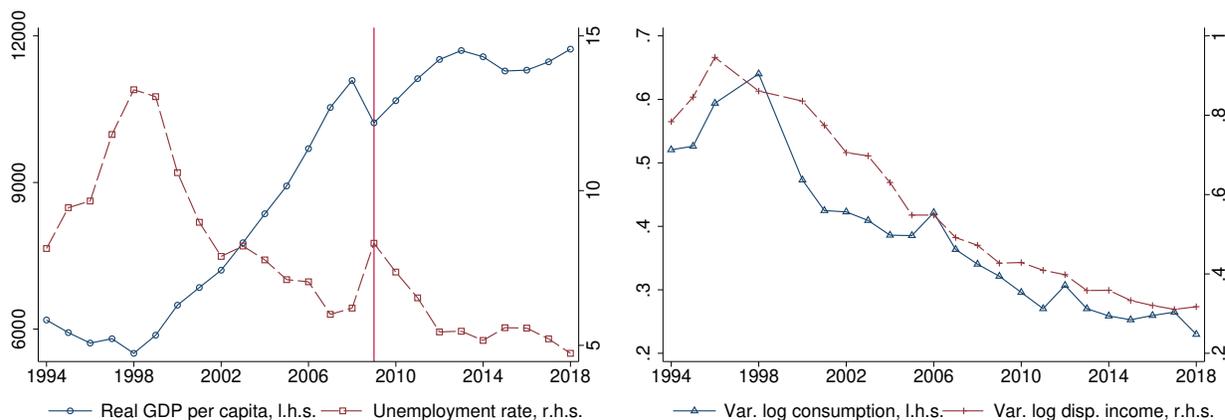
[†]E-mail: Mbryukhanov@gmail.com.

[‡]Department of Economics, University of Alberta, 8–14 HM Tory Building, T6G 2H4, Edmonton, AB, Canada. E-mail: dhryshko@ualberta.ca.

1 Introduction

Income inequality has risen steeply during the first few years of the transition of Russia to a market economy.¹ Starting in 1998, the Russian economy embarked on a path of a decade-long, robust economic growth accompanied by a substantial reduction in unemployment; see Figure 1, left panel. The financial crisis and the collapse of oil prices in 2009 have ended the period of high economic growth, but the Russian economy resumed its growth in 2010, albeit at a slower pace. The gains in the aggregate have been mirrored by a sustained reduction in income and consumption inequality among Russian households; see Figure 1, right panel.

FIGURE 1: AGGREGATE STATISTICS AND INEQUALITY IN RUSSIA



This paper provides an integrated view of inequality in Russia over time, for all households combined, and over the life cycle, for a typical cohort as it ages.

We proceed in the following steps. We first document inequality trends over time using panel data for 1994–2018 from the Russian Longitudinal Monitoring Survey (RLMS) and repeated cross-sectional data on household consumption from the Household Budget Survey (HBS) conducted by Rosstat, the official statistical agency in Russia. Since the average age in the data does not change significantly during the period (see Online Appendix Tables A-1–A-2), inequality trends must be driven by cohort and/or time effects.

¹The Gini coefficient for income more than doubled in a matter of several years, from 0.22 in 1989 to 0.54 in 1992, right after the dissolution of the USSR; see, e.g., [Guriev and Rachinsky \(2008\)](#).

The observed fall in income inequality may be due to a fall in the income-shock variances, whereas consumption inequality may fall, in addition, due to improved insurability of income shocks, where both the income-shock variances and insurance against the shocks could be cohort- and/or time-specific.

We show that cohort effects are negligible, whereas time effects are the essential drivers of inequality trends in two different ways. First, due to the relatively large cross-sectional dimension of HBS data, we can zoom in on the inequality patterns for particular birth cohorts. We find the secular decline in consumption inequality for various cohorts, with similar variances across the cohorts. We reach the same conclusion when analyzing within-cohort variances in income in our much smaller RLMS data. Second, we rely on the methodology of [Heathcote, Storesletten, and Violante \(2005\)](#) to confirm negligible cohort effects and essential time effects behind the inequality trends.

Guided by this finding, we use panel data from the RLMS for 1994–2018 to measure the extent of consumption insurance against permanent and transitory shocks to net family incomes, assuming that the variances of permanent and transitory shocks change with time and examine if this insurance differed for the periods of fast versus slow economic growth. For this measurement, we rely on the state-of-the-art methodology developed by [Blundell, Pistaferri, and Preston \(2008\)](#), where insurance is defined as the fraction of permanent and transitory shocks to net family incomes that do not pass through to household consumption. We find a substantial role of assets in insuring permanent income shocks, at around 40%, and a significant sensitivity of nondurable consumption to transitory shocks, with a marginal propensity to consume out of the transitory shock equal to about 0.12. We further find that the insurance of permanent shocks during the period of slow aggregate growth, 2009–2018, somewhat worsened, while the insurance of transitory shocks somewhat improved relative to their respective values for the period of fast economic growth prior to 2009. However, these changes were minor (and statistically insignificant), so most of the reduction in consumption inequality over time was due to a reduction in the variances of both permanent and transitory

shocks to net family incomes.

Our estimate of consumption insurance against permanent income shocks is similar to the estimates in [Gorodnichenko et al. \(2010\)](#) for an earlier period in Russia and in [Blundell et al. \(2008\)](#) for the U.S. As financial development in Russia and the U.S. is arguably different, [Gorodnichenko et al. \(2010\)](#) labelled the finding of nearly the same insurance against permanent income shocks in the two countries as a puzzle and called for more research.

To shed light on the puzzle, we next analyze income and consumption inequality for a typical cohort as it ages within a structural lifecycle model. To this end, we need to extract the age profiles from the observed inequality patterns.

Guided by our previous findings on time and cohort effects, we extract the age profiles by setting cohort effects to zero and find that the resulting paths for the variances of income and consumption over the life cycle are nearly parallel and virtually flat. A number of immediate implications of these findings follow. First, the flat lifecycle profile of income inequality points to the absence of a random walk in household incomes. Second, the available evidence points against the permanent income hypothesis as a plausible explanation of the observed data patterns since, regardless of the income process, it predicts an increasing consumption inequality for a typical cohort as it ages; see, e.g., [Deaton and Paxson \(1994\)](#). A model with precautionary motives coupled with a nontrivial income risk and the lack of insurance against it would go a long way in explaining why households might postpone consumption early in the life cycle, thereby dampening the forces that may raise consumption inequality in early stages. A properly calibrated model is required to explore if such a model is capable of explaining the full life cycle of income and consumption inequality.

Calibration of an incomplete-markets model is the next task we undertake in the paper. In the model, consumption and savings are endogenous, and income risk is exogenous. Our choice of the income process, the elements of which will be calibrated, is guided by the following considerations. The flat lifecycle profile of income inequality for a typical Russian cohort is incompatible with standard calibrations of incomplete-markets lifecycle models where the

variance of the long-lasting component is small or zero for everyone early in the life cycle, and the variance of income fans out over time due to the accumulation of idiosyncratic persistent shocks; see, e.g., [Storesletten, Telmer, and Yaron \(2004\)](#). However, it can be justified by an autoregressive component with a finite persistence and a sufficiently high initial variance.² We calibrate the persistence of the long-lasting component within a model jointly with the other parameters by targeting the income moments and the amount of insurance against long-lasting income shocks we found in the RLMS data.³ Although this insurance estimate was obtained assuming that household incomes contain random walks, using it as a calibration target is justified by the well-known result that the measured insurance of long-lasting shocks is not very sensitive to the misspecification of the persistent component of income (see [Kaplan and Violante, 2010](#) and [Hryshko and Manovskii, 2022](#) for details).

Calibrating the model, the persistence of the long-lasting component is found to be about 0.94 and is consistent with the amount of insurance we observe in the Russian panel data. Thus, the large insurance role of long-lasting shocks in Russia is puzzling only relative to the assumption of a random walk in household incomes maintained in our empirical estimation but overturned by our calibration results. Moreover, although we do not explicitly target the lifecycle profiles of income and consumption inequality in the baseline calibration, we are able to fit them reasonably well, with consumption inequality late in the life cycle being somewhat higher in the model than in the data.

To improve on the fit of the consumption inequality late in life, we add warm-glow bequest motives to the baseline model as in [De Nardi, French, and Jones \(2016\)](#). Calibrating the augmented model, the model-implied consumption and income inequality profiles are within the 95% confidence interval of the respective data profiles. Thus, a lifecycle model with persistent, but not fully permanent, income shocks, high initial variance of the shocks, and

²High dispersion of the persistent component early in the life cycle is consistent with many frictions of an economy with drastic changes in the economic environment. E.g., RLMS data available for our estimation period reveals that nearly 50% of males aged 25–34 are employed in a job that does not match their education/training.

³Such an approach that uses model-based inference on the income process parameters is similar in spirit to [Guvenen and Smith \(2014\)](#) and [Hryshko \(2007\)](#).

bequest motive is consistent both with the lifecycle profiles of income and consumption inequality and the panel-data evidence on consumption insurance.

In a recent paper, [Schulhofer-Wohl \(2018\)](#) argues that the standard practice of placing restrictions on time or cohort effects to identify the age effects might potentially result in the misspecified lifecycle profiles. He proposes identifying the lifecycle profiles jointly with the other parameters within a structural model. To examine the validity of our calibration results, we follow the procedure outlined in [Schulhofer-Wohl \(2018\)](#) where we explicitly target the lifecycle consumption and income inequality profiles obtained from the data using the standard restrictions on i) cohort effects (as we have done in the empirical analysis) and ii) time effects (following [Deaton, 1997](#) and [Aguiar and Hurst, 2013](#)). We find that the calibrated income-process parameters, preference parameters, and corrected inequality profiles are similar to those in the benchmark calibration augmented with bequest motives.

Our findings underscore the usefulness of the standard incomplete-markets lifecycle model in identifying the income process parameters and providing an integrated view of the panel-data evidence on consumption insurance against income shocks and income and consumption inequality over the life cycle in Russia. The model helps us pin down the persistence of long-lasting income shocks, the precise measurement of which is key for evaluating the insurance role of the tax and transfer system (e.g., [Blundell, Graber, and Mogstad, 2015](#)), and the lifecycle profiles of inequality, which are the crucial inputs of structural models seeking to uncover its sources (e.g., [Huggett, Ventura, and Yaron, 2011](#)).

Our paper is most closely related to [Gorodnichenko et al. \(2010\)](#), who use the RLMS to document the inequality patterns for the period 1994–2005.⁴ We consider the more extended period of 1994–2018 that featured a slowdown in the aggregate economic growth and build a model that reconciles the time-series and lifecycle inequality patterns.

[Santaaulàlia-Llopis and Zheng \(2018\)](#) apply the methodology of [Blundell et al. \(2008\)](#) to

⁴That study is part of the *Review of Economic Dynamics* special issue devoted to the analysis of inequality trends over time and the life cycle for several industrialized and developing economies. [Krueger, Perri, Pistaferri, and Violante \(2010\)](#) provide a summary of the issue. See also [Lise, Sudo, Suzuki, Yamada, and Yamada \(2014\)](#) and [Ding and He \(2018\)](#) for a comparable set of facts for Japan and China, respectively.

Chinese household data to show that consumption insurance of permanent shocks significantly worsened in the post-reform period of 1998–2009. We cannot compare our findings to the pre-transition years in Russia due to the lack of data to make a proper parallel to [Santaaulàlia-Llopis and Zheng \(2018\)](#), but we provide an additional focus on the lifecycle dimension of inequality. Note also that there was an increase in income and consumption inequality in China since 1989, both in the pre- and post-reform periods, rather than a secular decline observed in Russia since 1998. Thus, economic growth was accompanied by a fall in inequality in Russia and an increase in China. Russian experience is not unique, though, as many countries in Latin America experienced reductions in income inequality in the 2000s. [López-Calva and Lustig \(2010\)](#) list lower returns to education and more generous public transfers as the major factors behind the fall in income inequality in those countries. [Calvo, López-Calva, and Posadas \(2015\)](#) highlight the very same factors as the drivers of the income-inequality fall in Russia. [Engbom and Moser \(2022\)](#) find that a large increase in the real minimum wage in Brazil in the 2000s was responsible for a significant reduction in earnings inequality. None of these papers, however, link income and consumption inequality as we do in this paper.

The rest of the paper is structured as follows. Section 2 describes the data we use. Section 3 presents the trends of consumption and income inequality over time and examines the importance of time and cohort effects for the observed trends using the methodology of [Heathcote et al. \(2005\)](#). This section also estimates consumption insurance, the variances of permanent and transitory income shocks, and analyzes how the variances and insurance changed with time. Section 4 extracts income and consumption inequality profiles over the life cycle and calibrates the standard incomplete-markets model. Section 5 concludes.

2 Data

Our main source of consumption and income data is the RLMS, 1994 to 2018 waves.⁵ We also use consumption data from the HBS conducted by the national statistical agency, Rosstat, for the years 2003–2014. Both the RLMS and HBS are used in recent studies of inequality in Russia; see, e.g., [Bussolo and Luongo \(2020\)](#) and [Lisina and Van Kerm \(2022\)](#).

The RLMS is a longitudinal, nationally representative survey of Russian households conducted annually except for 1997 and 1999, when the data was not collected. It contains detailed information on demographics, consumption, and income and allows us to explore the evolution of income and consumption inequality over time. The survey is run by the Carolina Population Center at the University of North Carolina at Chapel Hill and an independent research center “Demoscope.”⁶ It produces a multistage probability sample of dwellings with 98 primary sampling units designed to represent the Russian Federation. About 4,000 households were interviewed in 1994. Although the original design had no intention of following households moving out of their original dwelling units in 1994, after 1996, all individuals and households were followed when they moved out of the household units. This created a panel of households that we use in our paper.

The initial survey response rate was high, at about 88%, and attrition rates are moderate partly due to low mobility in Russia.⁷ The sample was replenished several times in Moscow and St. Petersburg, geographical areas with the greatest nonresponse and attrition rates in

⁵We do not use data from the earlier waves of the survey because it was conducted by a different institution and utilized a different methodology; see [Kozyreva, Kosolapov, and Popkin \(2016\)](#).

⁶“Demoscope” is a Russian private company specializing in analyzing public opinion. Since 2010, the Higher School of Economics (HSE) began providing funding for the survey. The survey was officially renamed to the RLMS-HSE; <https://www.hse.ru/rlms>.

⁷As in [Hryshko and Manovskii \(2022\)](#), we calculate attrition for individuals ages twenty-four to fifty in the first year of the RLMS who are potentially observed for at least fifteen years in the dataset by age sixty-five; see their Supplementary Appendix Table A-11. This attrition rate is 49% for females and 61% for males. Female attrition rate is comparable to that observed in the German Socio-Economic Panel (GSOEP) and British Household Panel Survey (BHPS) and somewhat lower than in the Household, Income and Labour Dynamics in Australia (HILDA)—49%, 48%, and 56%, respectively. Male attrition is higher than in the GSOEP and BHPS (51% and 56%, respectively) and comparable to the HILDA value of 58%. Attrition rates in the Panel Study of Income Dynamics (PSID) in the U.S. are lower than in any other dataset, at 41% for men and 37% for women.

the dataset, and also once in a few other cities. In 2010, there was a 50% increase in sample size following the original selection approach.

Our measure of nondurable consumption comprises expenditures on food at home and away from home, alcohol and tobacco, clothing, fuel and gas, utilities, transportation, renovation and communication services, entertainment and vacation, personal care, and lawyer, notary, and realtor fees. We also add the value of consumed homegrown produce (such as, e.g., vegetables, nuts, fruit, berries, poultry, eggs, milk, and honey) to the consumption measure.⁸ Our consumption measure excludes expenditures on health and education, which is typical in the literature, as they are concentrated late or early in the life cycle; see, e.g., [Aguiar and Hurst \(2013\)](#). Our main income measure is net family income from various sources that includes private and public transfers.⁹ For each year, it is calculated as a sum of net incomes of all household members reported in the annual individual files. We also added the value of goods produced at home and sold in the market to the income measure.¹⁰ Income and consumption measures are deflated by the overall CPI with the base year of 2002. Our RLMS samples include households with married or cohabiting couples whose male is of age 25–59. Among those, we exclude households with gaps in the survey presence of more than two years and households with low incomes and income growth outliers. Low incomes are considered to be incomes below 100 real rubles (in 2002 prices), whereas income growth outliers are defined as observations with income growth rates above 500% or below –80%. Our final RLMS sample contains 6,733 households.

The HBS is cross-sectional but features a larger number of households and allows us to check the robustness of the RLMS consumption inequality trends to attrition that is characteristic of virtually any longitudinal dataset.¹¹ Similarly to the RLMS, it is nationally

⁸We calculate the values using the median prices of each homegrown item in each locality. If the number of observations per locality is less than a hundred, we use median prices for (geographically bigger) primary sampling units instead.

⁹All income measures recorded in the RLMS are net of taxes.

¹⁰For this calculation, we use the median local prices calculated as described above.

¹¹We apply survey household weights designed to correct sample statistics for attrition. For more information on attrition, weighting, and replenishment of the RLMS see <https://rlms-hse.cpc.unc.edu/project/samprep/>.

representative. The HBS is conducted quarterly by Rosstat, the official statistical agency in Russia, publicly available since 2003, and interviews about 50,000 households each quarter. Its history began in 1952, but due to various methodological deficiencies, it was substantially redesigned in 1994. Households in the survey are present for one year, interviewed each quarter (quarterly questionnaire of expenditures and in-kind consumption) and annually (questionnaire of living conditions, education, etc.). Stratified sampling is based on Census data and produces samples representative of each region in Russia for urban and rural locations; see [Yemtsov \(2008\)](#) for more details on the data.

We use information from 2003–2014 surveys for the fourth quarter when the majority of RLMS interviews are conducted. We select households whose heads are of ages 25–59 and drop consumption outliers defined as the records above the ninety-ninth and below the first percentiles of annual consumption distributions. Our cross-sectional samples contain around 16,000 households on average. Since the HBS does not have direct questions about incomes, we do not use the income information provided in the survey (it is calculated as a sum of total expenditures and an estimate of household savings provided in the survey).

Sample means for selected variables in our RLMS and HBS data are presented in Online Appendix Tables [A-1](#) and [A-2](#), respectively. Our samples feature similar average age. Households in the HBS are somewhat more educated, more likely to be located in big towns, have a smaller family size and number of kids, although the discrepancies are not large. For comparison, according to the official data published by Rosstat, the share of male urban population is 73% in 2002, 73.8% in 2010, and 74.1 in 2015—these numbers are closely replicated by our HBS data.¹² The Census share of college-educated males in our age group is 19.7% in 2002, 25.4% in 2010, and 29.3% in 2020, and compares favourably with our RLMS sample.¹³

¹²See <https://rosstat.gov.ru/folder/210/document/13204>.

¹³Go to https://rosstat.gov.ru/vpn_popul for Census data.

3 Inequality over time

In this section, we present inequality trends over time; estimate the extent to which non-durable household consumption is shielded from permanent and transitory shocks to household net incomes; and study how consumption insurance changed over time.

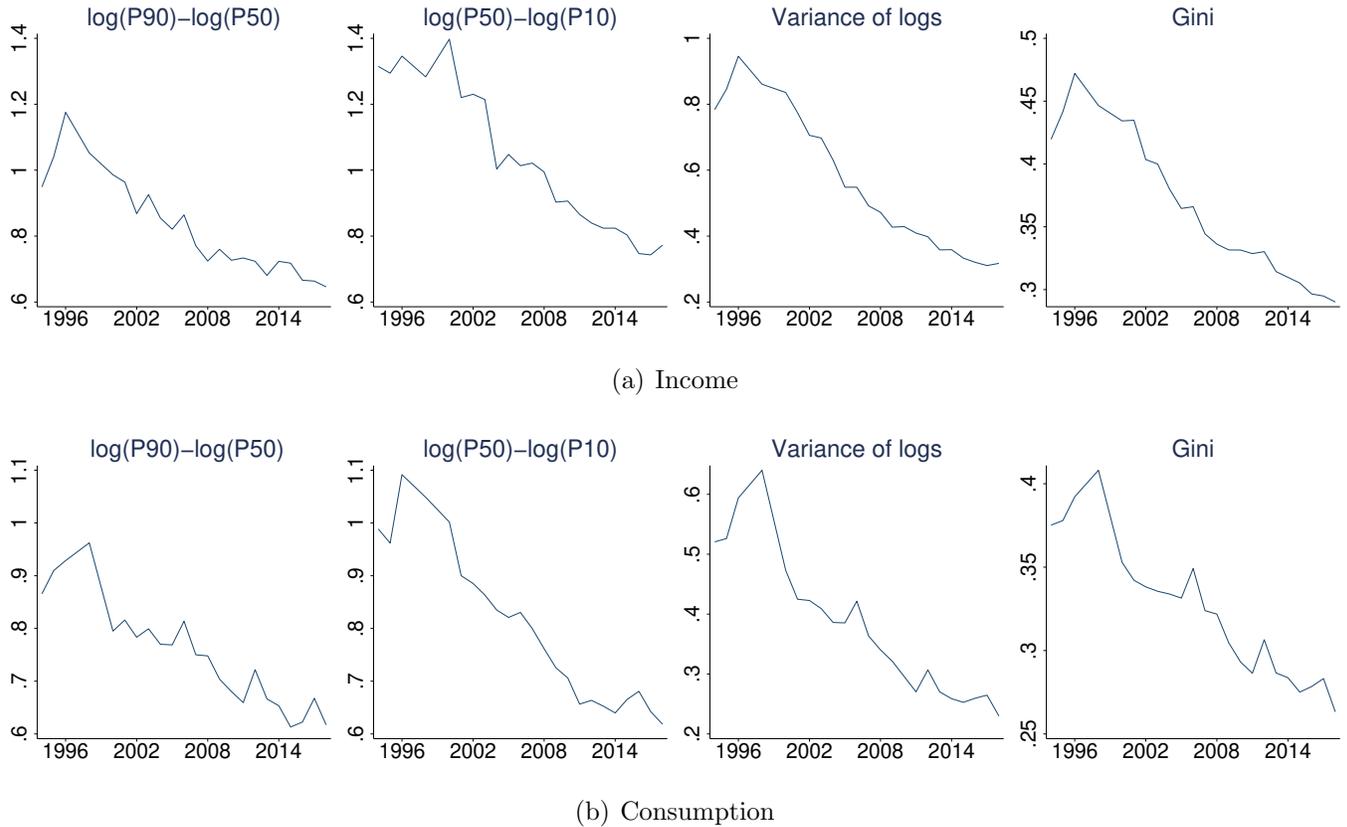
3.1 Trends in the RLMS

Figure 2 plots the trends in various inequality measures for equivalized household disposable income and nondurable consumption in Russia during the period 1994–2018. To equivalize household income and consumption, we divide them by the square root of family size (see, e.g., Johnson, Smeeding, and Torrey, 2005). We apply household weights provided in the survey to analyze the trends.

Since 1998, all measures—relative log percentiles, variances of logs, and Gini coefficients—point to a decline of inequality over time, and the two leftmost graphs indicate the process of catching up of lower income and consumption percentiles to the upper percentiles. Figure 3, where each log percentile subtracted its respective value in 1998, highlights that all income percentiles considered were growing since 1998, with the economywide growth benefiting the lower income percentiles more. Consumption has grown slower than income since 1998, reflecting a consumption-smoothing motive.

Figures A-1–A-2 in the Online Appendix show Lorenz curves for income and consumption distributions in 1998, 2008, and 2018. In each decade, income and consumption distributions become more equal, with more significant drops in inequality happening between 1998 and 2008. Online Appendix Figure A-3(a) shows concentration curves for consumption distributions in two years separated by one decade, years t and $t + 10$, ordered by average income in year $t - 1$ (calculated as net family incomes averaged over the period $t - 4$ to $t - 1$). In the leftmost graph, $t = 1998$, whereas $t = 2008$ in the rightmost graph. In each decade, the consumption of income-poorer families is catching up with their richer counterparts, with

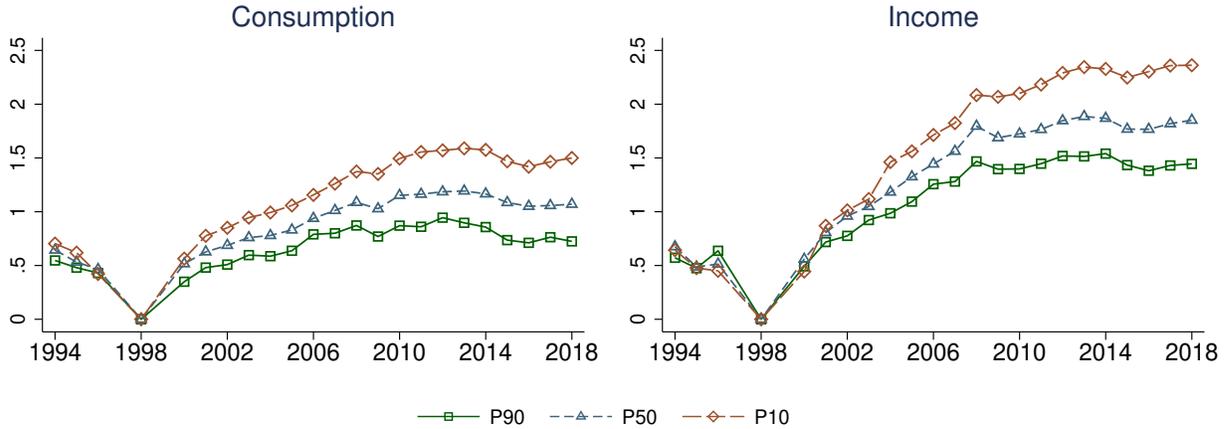
FIGURE 2: INEQUALITY OVER TIME IN RUSSIA. RLMS DATA, 1994–2018



Note: The figure plots various inequality statistics in the RLMS weighted using household weights. P90 (P10) stands for the 90th (10th) percentile of the cross-sectional distribution of income in panel (a) and consumption in panel (b).

more consumption gains occurring between 1998 and 2008. Online Appendix Figure A-3(b) plots concentration curves for consumption distributions in the same years ordered now by past average household consumption, a proxy for household permanent income. The plots show that, over one decade, the consumption of households with lower permanent incomes is getting closer to the consumption of households with higher permanent incomes, with the speed of catching up slowing down in the most recent decade.

FIGURE 3: CONSUMPTION AND INCOME PERCENTILES OVER TIME. RLMS DATA, 1994–2018



Note: The figure plots logs of the 90th, 50th, and 10th percentiles of consumption and income distributions in the RLMS relative to their respective values in 1998.

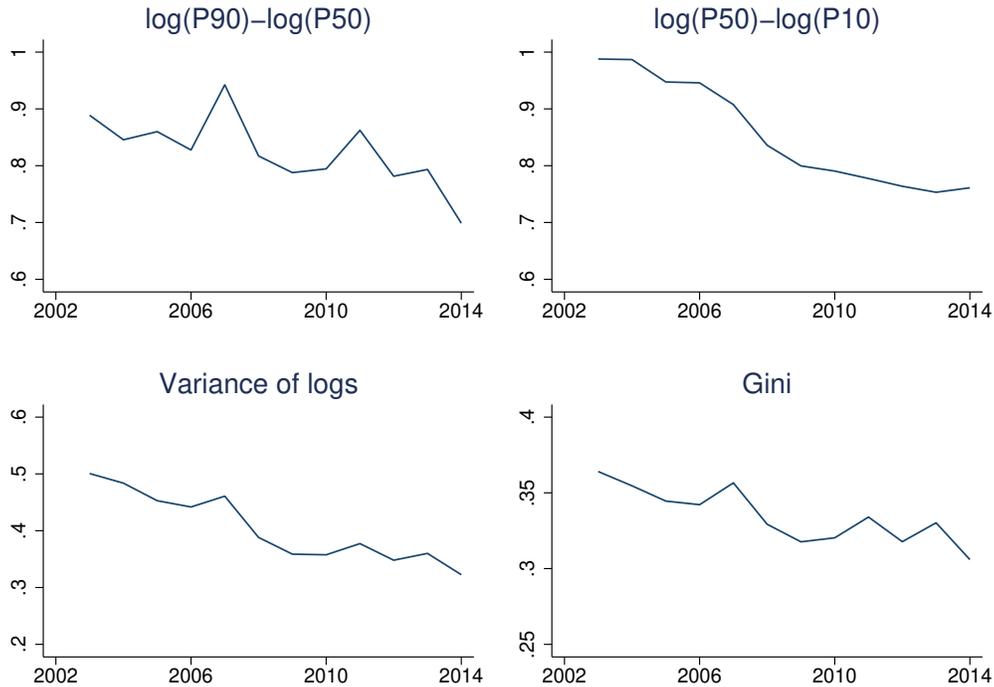
3.2 Trends in the HBS

As in the RLMS, we use a nondurable consumption measure equalizing it by the square root of family size and apply household weights when analyzing the trends. Figure 4 shows a secular decline of all inequality measures, similar to the findings above for the RLMS. Importantly, HBS data is not subject to dynamic attrition as the RLMS might be and features much larger sample sizes. Yet, HBS data shows qualitatively similar patterns, with somewhat flatter trends in inequality than in the RLMS.

Admittedly, both the RLMS and HBS do not contain the super-rich whose absence will lower the estimated inequality measures,¹⁴ more so for income than consumption. The trends we consider are still informative as they reflect the evolution of inequality for the overwhelming majority of the Russian population.

¹⁴See [Novokmet et al. \(2018\)](#) and [Gurieva and Rachinsky \(2008\)](#) for an attempt to assess the importance of top-earners for measuring income inequality in Russia.

FIGURE 4: CONSUMPTION INEQUALITY OVER TIME IN RUSSIA. HBS DATA, 2003–2014



Note: The figure plots various inequality statistics in the HBS weighted using survey household weights. P90 (P10) stands for the 90th (10th) percentile of the cross-sectional distribution of consumption.

3.3 The importance of time effects for the trends in inequality

Since average age did not change substantially during 1994–2018 (see Online Table A-1), the secular decline in inequality must be driven by time and/or cohort effects. In Online Appendix I, we evaluate the importance of time versus cohort effects for the trends in income and consumption inequality following the methodology of Heathcote et al. (2005). For both trends, we find strong evidence that the decline in cross-sectional variances is driven by the time effects, while the cohort effects appear to have little to no influence on the trends.

3.4 Consumption insurance and income-shock variances

The decline in income inequality can be attributed to the decline in the variances of income shocks of different durability, whereas the decline in consumption inequality, in addition,

can be attributed to the improvement of household insurance against those shocks. In this section, we estimate a permanent-transitory decomposition of net family income and consumption insurance against permanent and transitory income shocks. Guided by our finding above that the trend in income inequality is driven by time effects, we assume that the variances of permanent and transitory shocks to family income are time-specific.¹⁵ Since consumption inequality may fall due to more insurance against the shocks and/or lower variances of the shocks, we also examine if the former or the latter mechanism is more important for the secular decline in consumption inequality.

We follow the state-of-the-art methodology of [Blundell et al. \(2008\)](#) for estimating consumption insurance against long-lasting and transitory shocks to net family incomes and variances of income shocks. We estimate, by a diagonally-weighted minimum distance, the following model using RLMS data:

$$y_{it} = \alpha_i + p_{it} + \epsilon_{it} \tag{1}$$

$$p_{it} = \rho p_{it-1} + \xi_{it} \tag{2}$$

$$\Delta y_{it} = \xi_{it} + \epsilon_{it} - \epsilon_{it-1} \tag{3}$$

$$\Delta c_{it} = \phi \xi_{it} + \psi \epsilon_{it} + \zeta_{it} + \Delta u_{it}, \tag{4}$$

where c_{it} is log idiosyncratic consumption, y_{it} is log idiosyncratic income, p_{it} is its log persistent stochastic component, α_i is the fixed effect, ξ_{it} and ϵ_{it} are persistent and transitory shocks to net family incomes, respectively, ϕ and ψ are the transmission of those shocks to consumption, ζ_{it} is the permanent shock to consumption due to sources other than income shocks, and u_{it} is measurement error in consumption.

Following [Gorodnichenko et al. \(2010\)](#), we assume that the the shocks to p_{it} are permanent, i.e. $\rho = 1$, and that transitory shocks are not persistent. Idiosyncratic consumption

¹⁵In principle, the trend in income-shock variances and consumption insurance could be due to time and/or cohort effects. See [Browning, Ejrnaes, and Alvarez \(2010\)](#) for a discussion in the context of modeling the time series of the cross-sectional income variance in the U.S.

and income are calculated as residuals from a regression of family consumption and income on the full set of head’s year of birth, family size, and kids dummies, and interactions of the region, education, big city, and employment status dummies with the full set of year dummies. These regressions, extracting the idiosyncratic components of income and consumption, are standard in the literature; see, e.g., [Blundell et al. \(2008\)](#).

Since our main income measure is household earnings net of taxes and includes public and private transfers and the value of homegrown produce sold in the market, while the consumption measure includes homegrown goods used for household needs, the measured insurance reflects the role of assets in insuring the shocks to net family incomes.¹⁶ Accounting for home production is necessary for the precise measurement of the role of assets in insuring consumption from the shocks to net family incomes. A measure of consumption excluding home produce would be more volatile and indicate a smaller role of assets in insuring family consumption from the shocks to net family incomes. Excluding income flows from the homegrown produce sold in the market would make net family incomes more volatile and result in an inflated estimate of the role of assets in insuring consumption from income shocks.

Column (1) in Table 1 shows the results for the full period where we estimate the transmission coefficients ϕ , ψ , the variance of ζ_{it} , the variances of permanent and transitory shocks, and the variance of measurement error in consumption for various years.¹⁷ We find that nearly 40% of permanent shocks do not pass through to family consumption. This is to be compared with an estimate of 36% for the U.S. in [Blundell et al. \(2008\)](#). [Gorodnichenko et al. \(2010\)](#) find similarly high insurance of permanent shocks for the period 2000–2005. This is somewhat odd as an average U.S. household is more income- and wealth-rich than

¹⁶Using different income measures as a source of risk to household budgets will allow for the measurement of different insurance concepts. Using male earnings will allow one to study the role of spousal labor supply, taxes and transfers, and assets in insuring consumption from permanent and transitory shocks to *male earnings*; e.g., [Blundell, Pistaferri, and Saporta-Eksten \(2016\)](#). Using pretax family earnings will allow evaluating the role of tax and transfer system and assets in insuring consumption from the shocks to *family earnings*; e.g., [Arellano, Blundell, and Bonhomme \(2017\)](#) and [Hryshko and Manovskii \(2023\)](#).

¹⁷Since the transitory shock and measurement error in incomes are not separately identified, the estimated variance of transitory shocks also reflects transitory variation in incomes due to measurement error.

its Russian counterpart. [Hryshko and Manovskii \(2022\)](#) show, however, that high insurance of permanent shocks is consistent with the autoregressive component driven by moderately persistent shocks. We will return to this issue below, where we will calibrate the persistence of the autoregressive component that will be consistent with our estimate of ϕ .¹⁸ Our estimate for the insurance of transitory shocks is close to 90% and is statistically significant. It indicates a nontrivial sensitivity of Russian households to transitory shocks to income.

TABLE 1: CONSUMPTION INSURANCE OVER TIME. RLMS DATA

	1994–2018	1994–2008	2009–2018
	full period	rapid aggregate growth	slow aggregate growth
	(1)	(2)	(3)
ϕ , transmission of perm. shock	0.61 (0.04)	0.56 (0.06)	0.65 (0.06)
ψ , transmission of trans. shock	0.12 (0.02)	0.16 (0.03)	0.09 (0.03)

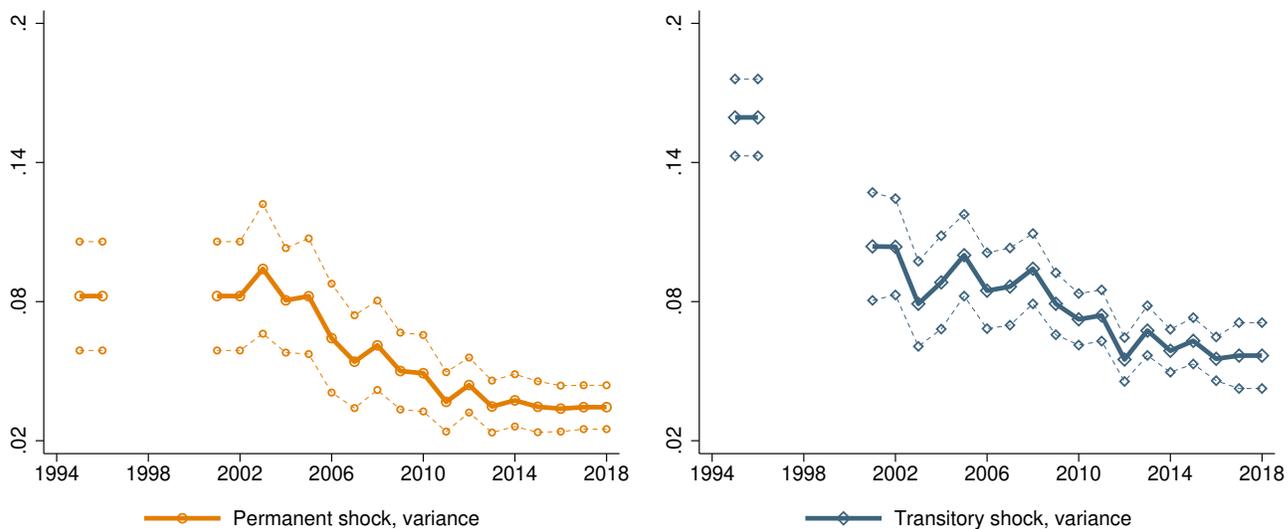
Notes: Standard errors in parentheses. Results from the model in Eqs. (1)–(4) estimated using diagonally-weighted minimum distance. Estimated variances of permanent and transitory shocks by year are plotted separately in Figure 5. p-value for the test of equal ϕ (ψ) in columns (2) and (3) is 22.7% (11.7%).

Columns (2) and (3) estimate the model in Eqs. (1)–(4), allowing for different values of insurance during the periods of fast and slow aggregate economic growth in Russia. We find a slight worsening of consumption insurance against permanent shocks and somewhat better insurance against transitory income shocks during the period of slower growth (with both changes being statistically insignificant). The latter is not surprising as during the period 2009–2018 there was a substantial increase in the minimum wage and pensions in Russia (see a discussion below). Since insurance of both permanent and transitory shocks did not

¹⁸Note that the minimum-distance estimate of ϕ does not deviate much from its true value when the estimated permanent component is misspecified; see [Kaplan and Violante \(2010\)](#) and [Hryshko and Manovskii \(2022\)](#).

change drastically over time, the decline in consumption inequality must be driven by the decline in the variance of the income shocks. This is confirmed in Figure 5, where we plot the estimated variances of permanent and transitory shocks from the model in columns (2)–(3). The figure shows that the size of transitory shocks is substantially larger and that the variances of permanent and transitory shocks are both falling at about the same pace during the estimation period.

FIGURE 5: VARIANCES OF SHOCKS OVER TIME. RLMS DATA



Note: Estimated variances from the model in column (1) of Table 1.

What could be the forces for the secular decline in inequality? Inequality could be driven by changes in household characteristics (e.g., higher female labor force participation, more families without kids, etc.) and returns to those characteristics (e.g., lower wages paid to more educated workers). Calvo, López-Calva, and Posadas (2015) find that composition effects have little impact on the observed reduction in wage inequality in Russia. Instead, it was driven by a lower return to secondary education, a relatively higher increase in pay for work in rural vs. urban areas, public vs. private firms, smaller vs. bigger firms, and the interregional convergence in pay. Increases in minimum wages significantly raised incomes

of those at the bottom of the distribution and were inequality-reducing.¹⁹ The real federal minimum wage was nearly flat since the start of the millennium, up until 2007, when it increased by about 90%, with a further 60% increase in 2009; see [Kapelyuk \(2015\)](#). Real pensions grew by about 80% from 1997 to 2008 and by more than 50% from 2008 to 2014. [Lisina and Van Kerm \(2022\)](#) find that increases in the pension levels contributed to the secular reduction in inequality of net family incomes. There were substantial changes in the economy’s industrial and occupational composition during the period. [Calvo et al. \(2015\)](#) argue that occupational shifts in Russia were slowing down the decline in wage inequality. [Gimpelson \(2016\)](#) finds large employment gains in construction, services and finance industries and losses in agriculture and processing industries. At the same time, wages in the industries paying the most, e.g., finance, grew at a slower pace than wages in the least-paying industries, e.g., education, health, and agriculture. Thus, lower employment in the low-pay industries and slower wage growth in the high-pay industries both contributed to a reduction in wage inequality. Inequality was also falling within industries.

Our results are not informative on the importance of any particular factor for the observed decline in inequality. Still, the above changes could drive reductions in both permanent and transitory inequality.

4 Inequality over the life cycle

This section complements our analysis by examining the lifecycle dimension of inequality. We first extract the lifecycle profiles of income and consumption inequality from the trends we documented in Section 3. We then calibrate a lifecycle model used as a lens for interpreting the evidence on consumption insurance and lifecycle trends in inequality.

¹⁹[Engbom and Moser \(2022\)](#), using administrative microdata from Brazil, show that a substantial increase in minimum wages can spill over to increases in wages high up in the distribution, explaining nearly half of the large decline in earnings inequality between 1996 and 2018 in Brazil.

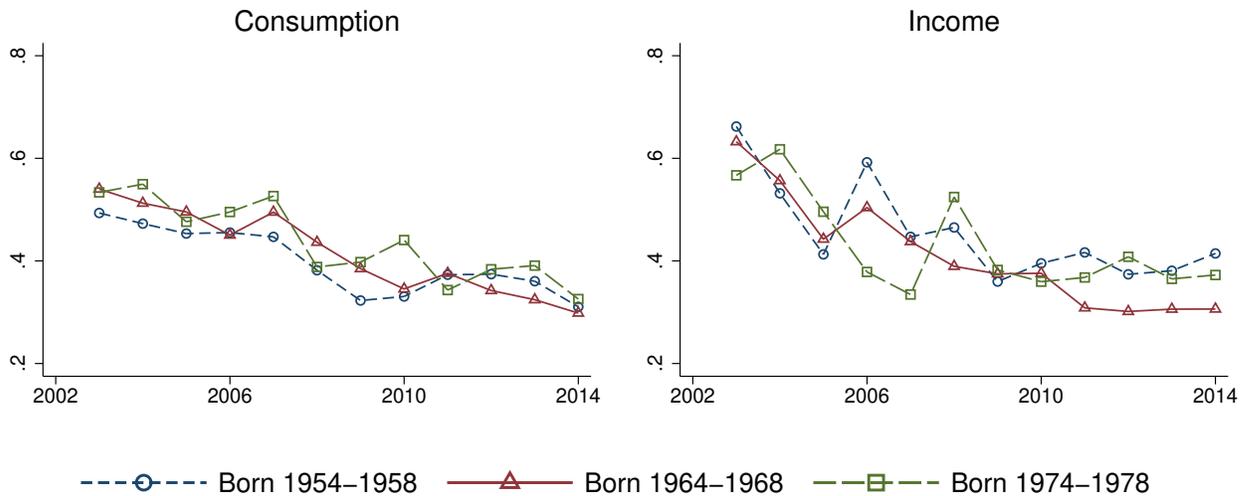
4.1 Within-cohort inequality

The HBS is large enough cross-sectionally to zoom in on the trend in consumption inequality for particular birth cohorts. Figure 6, left panel, plots the variance of consumption during the available period of 2003–2014 for households whose head is born in 1954–1958, 1964–1968, and 1974–1978. For any given cohort, the variances of consumption are declining over time. Moreover, the variances do not differ much across the cohorts considered. This is consistent with the presence of significant time effects driving down the variance of consumption along with stable variances over the life cycle, a conjecture that we will substantiate more formally shortly. In the right panel of the figure, we plot variances of incomes from the RLMS for the same cohorts of households. The income trends are noisier than the consumption trends since our RLMS samples are much smaller cross-sectionally. Yet, the patterns for income data are qualitatively similar. This visual evidence further corroborates our analysis of Section 3 where we concluded that cohort effects are inessential in driving the inequality trends.

There are some other important implications of the evidence in Figure 6. First, the left panel is inconsistent with the permanent income hypothesis, which predicts an increasing within-cohort consumption inequality regardless of the income process; see [Deaton and Paxson \(1994\)](#). Second, the right panel is inconsistent with models of incomes containing random walks, which predict fanning out of the income distribution as cohorts age and hence an increasing within-cohort income inequality. This implies that the income process we used in the analysis of consumption insurance above is misspecified. The measures of consumption insurance are nonetheless robust to this misspecification and are informative about the true persistence of longer-lasting income shocks; see, e.g., [Kaplan and Violante \(2010\)](#) and [Hryshko and Manovskii \(2022\)](#). Since it is hard to distinguish statistically a random-walk permanent component from an autoregressive component with finite persistence of shocks in small samples, we will use a lifecycle model below in which we will calibrate the persistence of long-lasting shocks by targeting, besides the income moments, the amount of consumption insurance we estimated using RLMS data. For that, we need to extract the lifecycle profiles

of income and consumption inequality from the observed trends.

FIGURE 6: WITHIN-COHORT INEQUALITY IN RUSSIA. 2003–2014

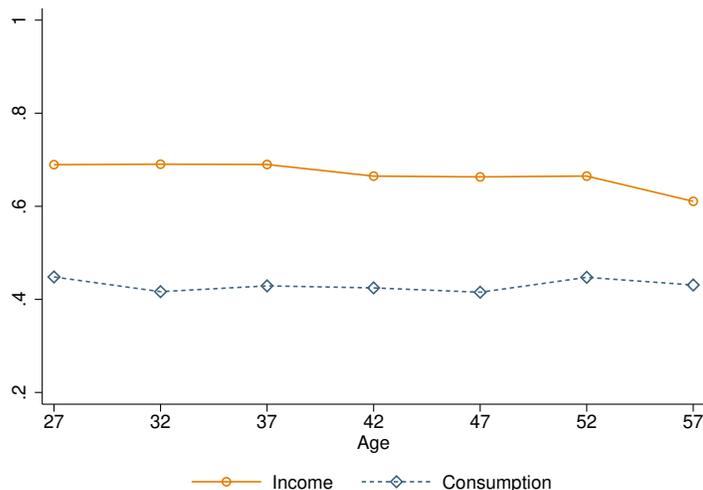


Note: The figure shows cross-sectional variances of equivalized nondurable consumption (in the left panel) and equivalized net family income (in the right panel) for selected cohorts using the HBS and RLMS data, respectively. See Section 2 for the sample construction details.

4.2 Extracting lifecycle profiles: setting cohort effects to zero

Figure 7 plots the age effects for the variance of log equivalized nondurable consumption and net family income obtained from a regression of each of those measures on age dummies and the full set of time dummies. Each age in the figure corresponds to the midpoint of a five-year age range; e.g., age 27 is the midpoint of the age range 25–29. As is clear from the figure, controlling for time effects and restricting cohort effects to zero produces nearly flat lifecycle inequality profiles for nondurable consumption and income.

FIGURE 7: INEQUALITY OVER THE LIFE CYCLE. RLMS DATA



Notes: The figure plots age effects in the variances of equivalized income (solid line) and equivalized consumption (dashed line) from a regression of the respective variances on age dummies and time dummies.

4.3 Modeling consumption inequality over the life cycle

A nearly flat lifecycle profile of consumption is hard to reconcile with the standard intertemporal theory of consumption when households build a buffer of assets to insure the risk to their disposable incomes; e.g., [Deaton \(1991\)](#) and [Carroll \(1997\)](#). Standard calibrations of lifecycle models—e.g., [Guvenen and Smith \(2014\)](#) and [Storesletten et al. \(2004\)](#)—assume that the variance of the autoregressive component at the start of working life is zero so that the variance of idiosyncratic incomes increases over the life cycle due to accumulation of independent persistent or permanent income shocks. We have to depart from the assumption: A flat lifecycle income profile can be generated by a combination of stationary but persistent incomes and a relatively high initial income variance.²⁰ In the following, we will, therefore, assume that $\rho < 1$ and $\text{var}(p_{i0}) = \frac{\sigma_\xi^2}{1-\rho^2}$, that is the variance of the initial value of the autoregressive component equals its long-run variance.²¹ Next, we describe the remaining elements

²⁰A decreasing lifecycle profile of income inequality may be generated by an income process with the variance of initial incomes exceeding the long-run variance of income.

²¹[De Nardi, Fella, and Paz-Pardo \(2019\)](#) argue that the canonical income process should be modified to allow for age-dependent persistence and variance of shocks to fit nonlinear consumption and income lifecycle

of the model used to interpret the lifecycle profiles of consumption and income inequality.

4.3.1 The model

We assume that households supply labor inelastically, value consumption using a CRRA utility function, and discount future utility at the rate of $\beta^{-1} - 1$. They start working at age t_0 , retire at age T_R , and face survival risk from age T_{R+1} until they die at age T with certainty. Households solve the following problem:

$$\max_{\{C_{it}\}_{t=t_0}^T} E_{i,t_0} \sum_{t=t_0}^T \beta^{t-t_0} s_t \frac{C_{it}^{1-\gamma} - 1}{1-\gamma},$$

subject to

$$W_{it+1} = (1+r)(W_{it} + Y_{it} - C_{it}), \quad t = t_0, \dots, T$$

$$Y_{it} = \mu_t \exp(\alpha_i) P_{it} \exp(\epsilon_{it}), \quad t = t_0, \dots, t_R$$

$$P_{it} = P_{it-1}^\rho \exp(\xi_{it}), \quad t = t_0, \dots, t_R$$

$$Y_{it} = \kappa \exp(\alpha_i) P_{it_R}, \quad t = t_R + 1, \dots, T$$

$$W_{it} \geq 0, \quad t = t_0, \dots, T.$$

C_{it} , W_{it} , Y_{it} are, respectively, household i 's consumption, wealth, and income at age t . Income is subject to the persistent and transitory risk up to retirement and is deterministic after that. μ_t is the common lifecycle income profile, α_i is the individual fixed effect, P_{it} is the persistent component, and ϵ_{it} is the transitory shock. The persistent component is an AR(1) process in logs with persistence ρ , and ξ_{it} is the shock to the autoregressive component. κ is the replacement rate, a fraction of last period's earnings net of the transitory component paid at retirement, γ is the coefficient of relative risk aversion, and r is the net real interest

inequality profiles in the U.S. In our data, both profiles are flat, and these modifications are not necessary to fit the data, as we will demonstrate below.

rate on a riskfree asset. We assume that households cannot borrow.

4.3.2 Baseline calibration

Households start their life at age 25 with zero assets, retire at age 59 (the official retirement age of males in Russia during the period we analyze), and die at age 85, that is, $t_0 = 25$, $t_R = 59$, and $T = 85$. The rest of the parameters used in our calibrations are listed in Table 2.

The common lifecycle income profile is estimated from RLMS data. For the income profile, we use age effects from a regression of log net family income on age dummies, the full set of year dummies, and also dummies for primary sampling units, big city, family size, and the number of kids dummies, all interacted with the year dummies. Survival probabilities are for females, taken from the Human Mortality Database for the years 2005–2009.²² The replacement rate of 50% is taken from the OECD.²³ During our estimation period of 1994–2018 inflation outpaced nominal interest rates on average. We, therefore, set the real interest rate to a low negative value of -2% . The coefficient of relative risk aversion, γ , is set to 3. We will calibrate the persistence of long-lasting shocks, ρ , internally. For this reason, we do not use the estimated variances of shocks from Table 1 but rather calibrate them, as those estimates are obtained when the autoregressive component is assumed to be a random walk.

The bottom five parameters in Table 2 are calibrated by minimizing the sum of squared log deviations of the model moments from the data moments listed in the first column of Table 3. Our motivation behind the use of those moments is the following. The variance of idiosyncratic log incomes, y_{it} , at early ages is informative about the variance of fixed effects, whereas the variance of income growth is indicative of the overall income risk during the life cycle and disciplines the combined size of long-lasting and transitory income risk. The two extra moments for the variance of log incomes at longer horizons are sufficient for the identification of the variance and persistence of long-lasting shocks, and the variance of

²²https://www.mortality.org/cgi-bin/hmd/hmd_download.php

²³<https://data.oecd.org/pension/net-pension-replacement-rates.htm#indicator-chart>

TABLE 2: BASELINE CALIBRATION

Parameter	Notation	Value	Data source
<i>Externally calibrated</i>			
Income age profile	μ_t	various	RLMS
Survival probabilities	s_t	various	Human Mortality Database
Replacement rate	κ	0.50	OECD
Interest rate	r	-2%	OECD
CRRA	γ	3	
<i>Internally calibrated</i>			
Variance of persistent income shocks	σ_ξ^2	0.055	
Variance of transitory income shocks	σ_ϵ^2	0.082	
Variance of fixed effects	σ_α^2	0.119	
Persistence of long-lasting shocks	ρ	0.940	
Time discount factor	β	0.939	

transitory shocks. Transmission coefficients for persistent and transitory shocks in the model are calculated as $\frac{E[\Delta c_{it} \sum_{j=1} \Delta y_{it\pm j}]}{E[\Delta y_{it} \sum_{j=1} \Delta y_{it\pm j}]}$ and $\frac{E[\Delta c_{it} \Delta y_{it+1}]}{E[\Delta y_{it} \Delta y_{it+1}]}$, where c_{it} is family i 's log consumption at age t . These moments recover the transmission coefficients under the assumption of a random walk in incomes, the assumption we maintained when calculating the empirical targets. By targeting these moments, we match the amount of consumption insurance available to households in the data. We do not observe wealth in RLMS data and target the wealth-to-income ratio of three in our calibration.²⁴ This value is conservative—it constrains the

²⁴Kaplan and Violante (2014) develop a model where households save into liquid and illiquid assets to explain a high marginal propensity to consume out of transitory income transfers observed in U.S. data. In the model, saving into an illiquid as opposed to a liquid asset involves a tradeoff between a higher consumption in the future from holding the former and a worsened consumption smoothing due to its high transaction costs. As a result, not only poor (with little liquid and no illiquid assets) but also wealthy (with little liquid and significant illiquid assets) hand-to-mouth households are highly sensitive to transitory income transfers. We do not consider this extension of the standard lifecycle model because its calibration requires detailed information on household asset portfolios, which is not available for Russia. Moreover,

amount of wealth available in the model for insurance of the shocks and helps discipline the income-process parameters. Without this target, one could easily match the transmission coefficient for permanent shocks observed in the data under the maintained assumption of a random walk in incomes.

TABLE 3: DATA AND BASELINE MODEL’S MOMENTS

	Data	Model
$\text{var}[y_{it}]$, ages 25–29	0.69	0.69
$\text{var}[\Delta y_{it}]$, ages 25–59	0.217	0.220
$\text{var}[\Delta^2 y_{it}]$, ages 25–59	0.289	0.271
$\text{var}[\Delta^3 y_{it}]$, ages 25–59	0.305	0.320
Transmission of persistent shocks	0.61	0.62
Transmission of transitory shocks	0.12	0.12
Wealth-to-income ratio	3.00	3.00

Notes: The growth of idiosyncratic income over a k -year horizon is measured as $\Delta^k y_{it} = y_{it} - y_{it-k}$, $k = 2, 3$.

Next, we simulate our model for a large number of households and estimate the lifecycle profiles of income and consumption. The leftmost plot in Figure 8 reproduces the data profiles in Figure 7 and shows the profiles from our baseline calibration. We match the income profile rather well. Since the model replicates the amount of insurance in the data, we have a reasonably good fit to the lifecycle profile of consumption. Our calibration did not target the shape of the consumption profile, but the model features reproduce the flatness of the profile up to the late stages of working life. Inequality of consumption starts rising in the model after age fifty when most of the lifetime income risk is resolved and precautionary savings motives are weakened.

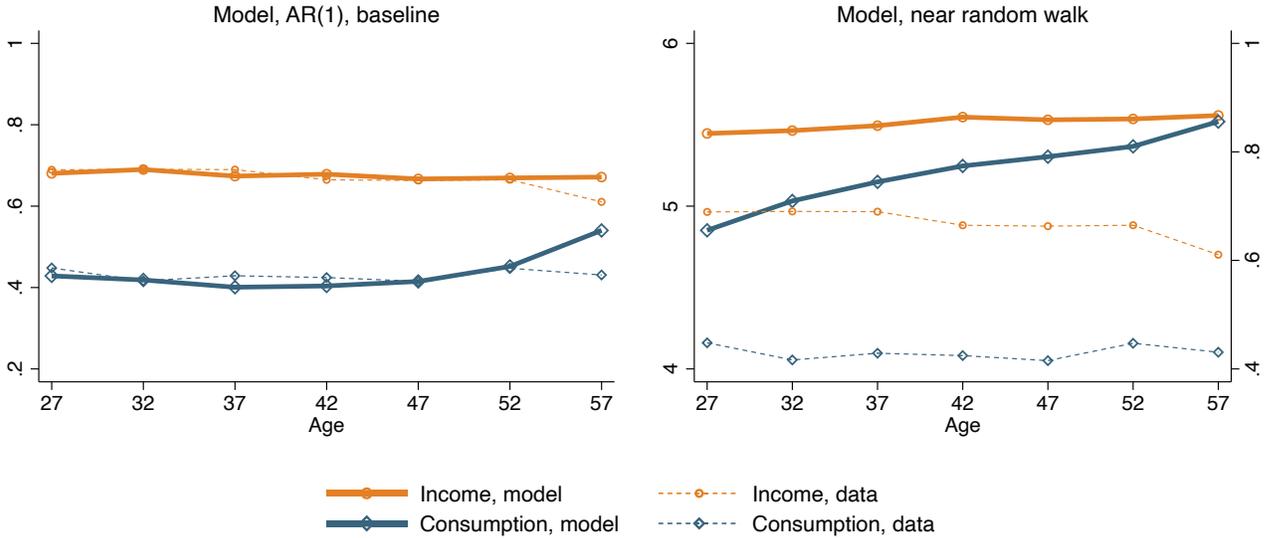
The rightmost plot shows the profiles from a calibration, in which we assume that the mechanism of Kaplan and Violante (2014) is less relevant in the Russian context since housing, the main source of illiquidity, is saved for in the U.S. and other developed economies but owned outright by the majority of Russian households due to free privatization of the 1990s; e.g., 90% percent of households in Russia owned their dwelling in 2013 (<https://33.rosstat.gov.ru/storage/mediabank/vid+qilja.htm>).

autoregressive component is nearly a random walk: $\rho = 0.995$. For this calibration, we choose the persistence of 0.995 for the following reasons. First, we need a finite persistence as a random walk would imply an increasing lifecycle profile of income inequality for any initial conditions of the autoregressive component. Second, a persistence of 0.995 would be impossible to distinguish from one in the data. We further take the variances of persistent and transitory shocks from Table 1 and assume that the variance of fixed effects is zero (there is no need for this income component, as will be seen momentarily). We then calibrate the time discount factor β by matching the wealth-to-income ratio of 3.²⁵

First, note that the scale of variances of income and consumption over the life cycle in the model (left y-axis) exceeds by more than an order of magnitude their scale in the data (right y-axis). This is due to the large variance of persistent shocks estimated in the data. Second, we are not able to reproduce the relatively flat inequality profile of consumption in the early and prime stages of the life cycle if we assume that the autoregressive component is nearly a random walk. Thus, our baseline calibration is by far superior to the (near) random-walk calibration.

²⁵The random-walk calibration is incapable of matching as low value for the transmission coefficient of permanent shocks as observed in the data when the wealth-to-income target is set to a reasonable value; see, e.g., Kaplan and Violante (2010) and Fella, Frache, and Koeniger (2020). This result is known in the literature as the excess insurance puzzle.

FIGURE 8: INEQUALITY OVER THE LIFE CYCLE. MODEL VS. DATA



Notes: The leftmost figure plots the inequality profiles implied by the baseline model’s calibration with the persistence of long-lasting shocks of 0.94, whereas the rightmost figure plots the inequality profiles implied by the calibration that assumes the persistence of long-lasting shocks equal to 0.995. Dashed lines in both figures represent the data inequality profiles plotted in Figure 7. Both calibrations do not target consumption and income inequality profiles.

The wealth-to-income ratio of 3 that we targeted in the baseline calibration is consistent with the estimates of net private wealth-to-national income ratio in [Novokmet, Piketty, Yang, and Zucman \(2018\)](#) for our estimation period; see their Figure 3. However, we also experimented with lower values for the wealth-to-income target to address the possible concern of the under-representation of income and wealth-rich households in the RLMS data. Those alternative calibrations resulted in an inferior fit of the insurance of persistent shocks observed in the data and a poorer fit of the consumption inequality profile; see Online Appendix Table B-1 and Figure B-1.

To fit the nearly flat income profile, we assumed that the initial variance of the persistent income component equals its long-run variance. To examine the sensitivity of our results to the assumption, we calibrated the model, assuming instead that the initial variance of the persistent component equals half the size of its long-run variance. The results are in Online Appendix Table B-2 and Figure B-2. Although the model fits well the income-

growth moments, the variance of income at early ages and consumption insurance, it fails at producing the nearly flat consumption and income inequality profiles in the data. In another experiment (unreported), we introduced an extra parameter θ to calibrate the initial variance of the persistent component as $\theta \frac{\sigma_\xi^2}{1-\rho^2}$. The estimated value of θ was very close to one, the benchmark model’s value. Interestingly, all these alternative calibrations yielded a similar estimated persistence of long-lasting shocks of 0.94.

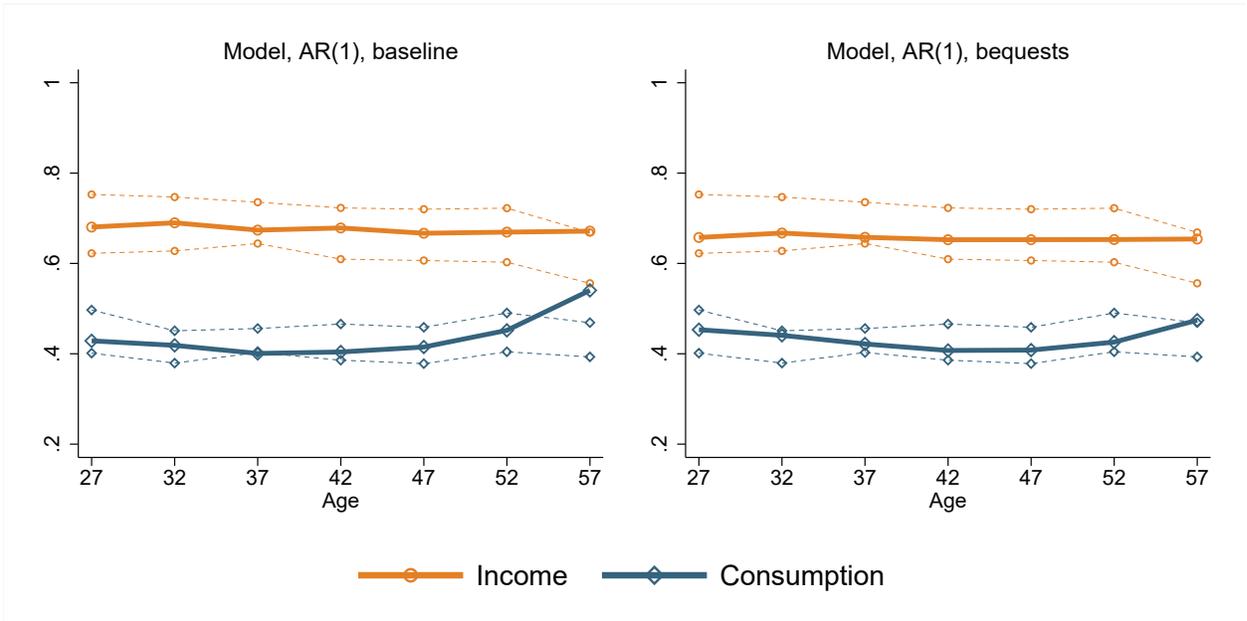
4.3.3 Adding bequest motives to the model

To improve the model fit of consumption inequality late in the life cycle, we include warm-glow bequest motives as in [De Nardi, French, and Jones \(2016\)](#).²⁶ Specifically, we assume that after retirement, households also value leaving bequest at age t upon death using function $b_1 \frac{(b_2 + W_{it})^{1-\lambda}}{1-\lambda}$, where b_1 measures the intensity of bequest motive and b_2 measures the curvature of the bequest function and ensures finite marginal utility from leaving zero bequests for low-wealth households.²⁷ We calibrate the parameters of the bequest function, adding the inequality consumption profile to the calibration targets of [Table 3](#). Since we have extra moments used for calibration, we also calibrated the coefficient of relative risk aversion internally.

²⁶An alternative way to flatten the consumption inequality profile late in the life cycle is to introduce uncertain medical expenses after retirement. We have experimented with this option but achieved a better fit to the data by adding bequests. Importantly, that calibration yielded similar estimates of the income-process parameters, with the persistence of long-lasting shocks calibrated to a value close to 0.94.

²⁷We also considered a possibility of receiving bequests from ages 30 to 55—when parents are of retirement age up to the age of death at 85 assuming that the age distance between adjacent generations is thirty years—but it did not help in fitting our calibration targets, so we assumed in this calibration that households do not receive bequests.

FIGURE 9: LIFECYCLE INEQUALITY. ADDING BEQUESTS TO THE BASELINE MODEL



Notes: The left figure plots, in solid lines, the inequality profiles implied by the baseline model’s calibration, whereas the right figure plots, in solid lines, the inequality profiles implied by the calibration that adds bequest motives to the baseline model. Dashed lines in both figures represent the 95% confidence intervals for the data inequality profiles plotted in Figure 7. The baseline model’s calibration does not target consumption and income inequality profiles. Calibration of the model with bequest motives added does not fit the income inequality profile.

The values for the calibrated parameters are reported in Table A-3, and the fit to the data moments is shown in Table A-4 in the Online Appendix. There are some differences in the calibrated parameters relative to their values in the baseline calibration, e.g., the time discount factor is somewhat lower, the variance of persistent shocks is lower, and the variance of transitory shocks is higher. However, the persistence of long-lasting shocks is calibrated at a similar value. Since the original moments have a lower effective weight in the calibration, we overfit the wealth-to-income ratio; see Table A-4. The fit to the inequality profiles is plotted in Figure 9. Solid lines with circles and diamonds plot the inequality profiles implied by the model for income and consumption, respectively, while the dashed lines around them represent the 95% confidence bands for the profiles observed in the data. The lifecycle profile for consumption inequality in our new calibration is explicitly targeted, whereas the income inequality profile is not. For the baseline calibration, reproduced in the left panel of

the figure, the income inequality profile implied by the model is within the 95% confidence interval, whereas consumption inequality at late ages exceeds the upper bound of the data confidence interval. Both income and consumption inequality profiles from our calibrated model with bequests are now within the 95% confidence intervals of their respective data profiles; see the right panel of Figure 9.

Online Appendix Figure B-3 plots, in solid lines, the relative percentiles for income and consumption implied (and not targeted) by the calibration of the model with bequests. Dashed lines in both figures represent the 95% confidence intervals, calculated by bootstrap, for the relative percentiles in the data. The log 90/50 and 50/10 ratios of consumption over the life cycle are within the 95% confidence interval of their data values. For income, the log 90/50 (50/10) ratio is at the upper (lower) bound of its 95% confidence interval in the data. However, these differences are not large. The log 90/50 (50/10) ratio for income in the data is 0.93 (1.23) and 1.09 (1.09) in the model. This implies the relative ratios of 2.5 (3.4) in the data and 3.0 (3.0) in the model.

Our model takes the income process as an exogenous input. What could be a mechanism generating the nearly flat lifecycle income inequality profile? [Huggett, Ventura, and Yaron \(2011\)](#) provide an example within a model which endogenously generates a persistent income process. In the model, individuals generate earnings used for consumption and saving by supplying human capital to competitive firms; human capital differs across agents at the start of the life cycle and over the life cycle due to different learning abilities and shocks. [Huggett et al. \(2011\)](#) show that a model where agents have the same learning ability generates a flat lifecycle earnings inequality profile. It would be, perhaps, far-stretched to make an argument of similar abilities across the Russian population of workers, but an alternative mechanism within the model could potentially make the lifecycle earnings inequality profile flat. Consider the Ben-Porath production function of human capital in [Huggett et al. \(2011\)](#) but also standard in the literature: $h_{it+1} = [h_{it} + a_i(h_{it}s_{it})^\alpha]$, where h , a , s , and α stand for the stock of human capital, learning ability, time allocated to investment into human capital,

and the curvature of the production technology of new human capital, respectively. Even with varying learning ability, a_i , agents will invest little into human capital over the life cycle if α is close to zero, perpetuating human capital and earnings inequality at the start of life for a cross-section of workers. There is some evidence that α is low or even zero/negative in Russia and other less-developed economies. [Gorodnichenko et al. \(2010\)](#), using RLMS data for 1994–2005, and [Gimpelson, Kapeliushnikov, and Oshchepkov \(2016\)](#), for a later period, show that the experience premium in Russia is negative. [Lagakos, Moll, Porzio, Qian, and Schoellman \(2018\)](#) and [Jedwab, Romer, Islam, and Samaniego \(2023\)](#) argue, respectively, that returns to experience are twice as high in rich than in poor economies, and that workers in rich economies accumulate twice as much human capital. [Donovan, Lu, and Schoellman \(2023\)](#), using cross-country micro data on labor market dynamics, argue that more severe search frictions could also be an impediment to lifecycle wage growth in poor economies. All this evidence is consistent with low α in Russia (and other less-developed economies) and a flat lifecycle profile of earnings inequality.

4.3.4 Extracting lifecycle profiles of inequality using the model

To isolate the lifecycle inequality profiles, we set the cohort effects to zero when decomposing the variances of income and consumption into the sum of time, cohort, and age effects. However, [Schulhofer-Wohl \(2018\)](#) argues that placing restrictions on cohort or time effects might result in misspecified lifecycle profiles used as inputs into and/or targets in structural models. This, in turn, might lead to incorrect inferences about the structural parameters. Although we did not fit the lifecycle inequality profiles explicitly, we compared the empirical profiles and their model counterparts to judge the model’s fit to the moments that are untargeted in our calibration. To sharpen the estimated values of the calibrated parameters and to explore the validity of our conclusion above on the importance of time effects for identification of the age effects in cross-sectional variances of consumption and income, we adopt the procedure in [Schulhofer-Wohl \(2018\)](#) that allows for joint identification of the

model parameters and the lifecycle profiles of income and consumption inequality. In this section, we will first briefly describe the procedure and then present our results based on it.

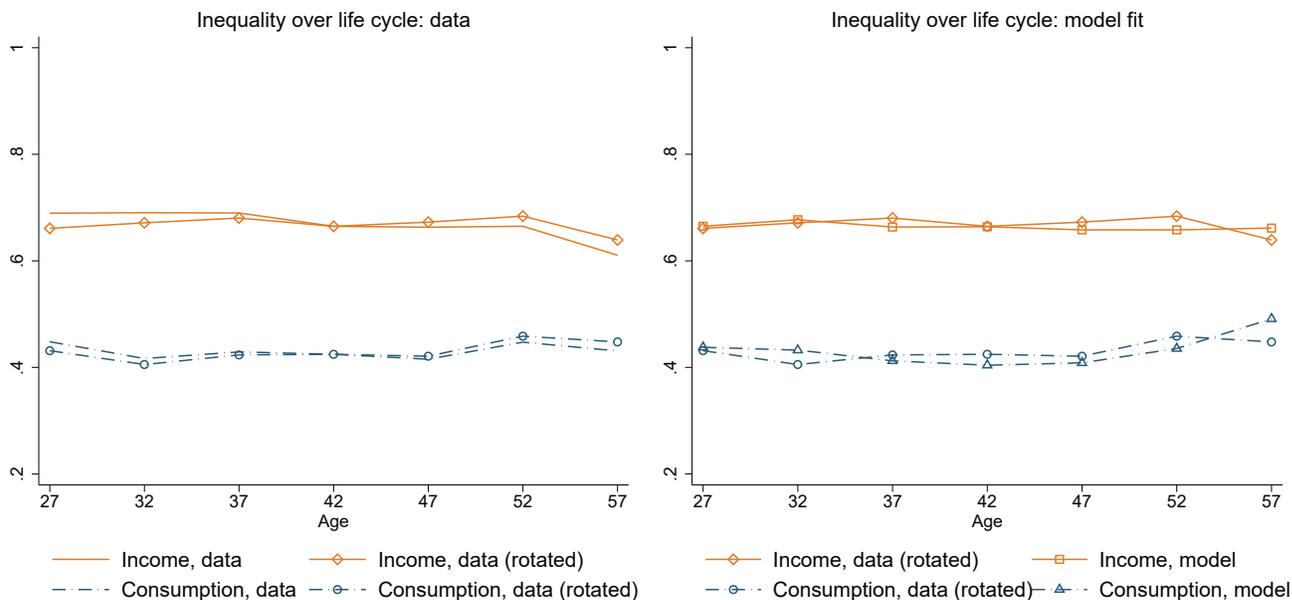
Let the empirical lifecycle profile, identified by placing restrictions on time or cohort effects, be $\hat{\alpha}_j$ and the model-based profile be $q_j(a; \theta)$ for $j = (c, i)$, where a is the age vector $[1, \dots, A]$,²⁸ θ is a vector of model parameters, and (c, i) stands for consumption and income, respectively. Schulhofer-Wohl (2018) proposes to minimize a weighted distance between $q_j(a; \theta)$ and $\hat{\alpha}_j + k_j \mathbf{a}$, where \mathbf{a} is a column vector $[1 - \bar{a}, \dots, A - \bar{a}]'$, with \bar{a} being the average age. Note that the empirical lifecycle profile will coincide with the profile identified structurally only if k_j is estimated to be zero. Following Schulhofer-Wohl (2018), we will refer to $\hat{\alpha}_j + k_j \mathbf{a}$ below as the rotated data profile.²⁹

In Online Appendix Table A-5, we present the results of a calibration of the model with bequest motives added where, in addition to the data moments of our benchmark calibration, we also target the rotated lifecycle profiles. We have added seven moments to be matched (income inequality profile) with only two extra parameters, k_i and k_c , that have to be calibrated. As can be seen in Table A-5, the fit to the data moments is fairly good, the calibrated parameters are very close to their respective values from the previous calibrations, and the rotation parameters k_i and k_c are estimated to be close to zero. As a result, the rotated inequality profiles are very close to the data profiles obtained by assuming that the cohort effects are zero; see Figure 10.

²⁸The model ages 1 and A correspond to ages 25 and 59 in the data.

²⁹Rotated data profile is a result of adding a linear trend $k_j \mathbf{a}$ to the data profile using the estimated trend coefficient k_j so that the slope of the data profile at different ages corresponds to the lifecycle profile predicted by the model.

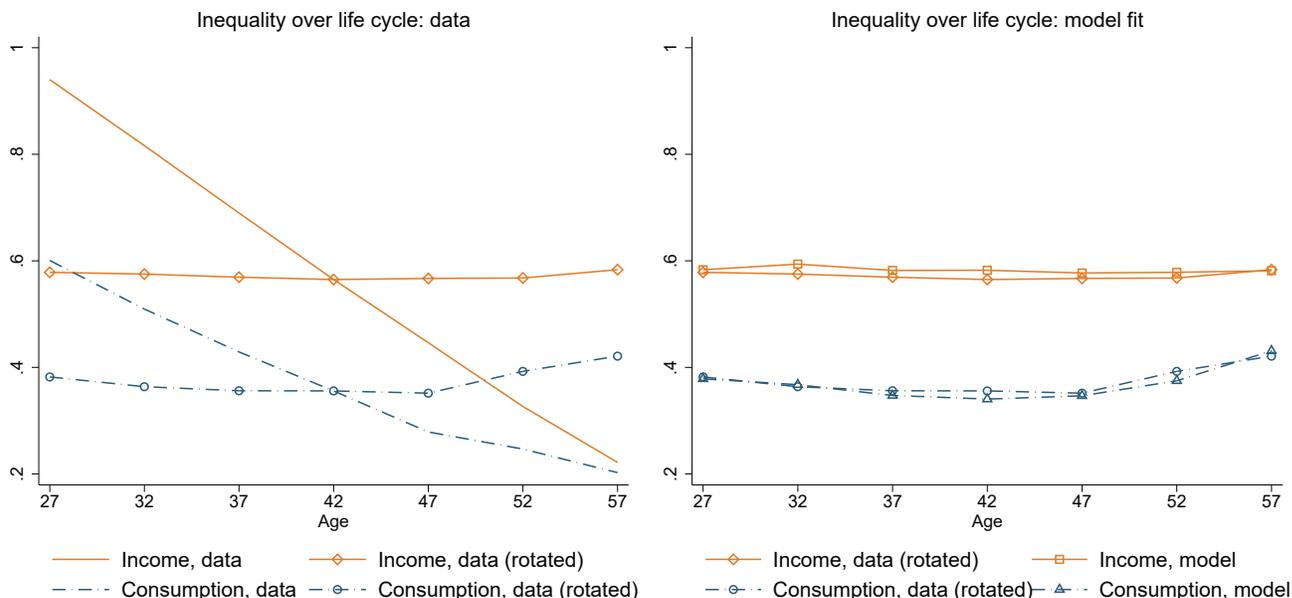
FIGURE 10: LIFECYCLE INEQUALITY. MODEL VS. DATA PROFILES OBTAINED BY SETTING COHORT EFFECTS TO ZERO



Notes: The left figure plots in solid and dashed lines lifecycle income and consumption inequality profiles, respectively, estimated from the RLMS data by setting cohort effects to zero. The rotated income and consumption data profiles are in solid line with hollow diamonds and dashed line with hollow circles, respectively. The right panel shows the model-implied income and consumption inequality profiles in a solid line with hollow squares and a dashed line with hollow triangles, juxtaposed against the rotated data profiles. See Section 4.3.4 for details on the construction of the rotated data and model profiles.

Although our choice of restricting cohort effects to zero was motivated in Sections 3.3 and 4.1, the resulting lifecycle profiles could still be misspecified. For robustness, we next use the lifecycle inequality profiles identified by normalizing the time effects and placing no restrictions on the cohort effects, as in Deaton (1997) and Aguiar and Hurst (2013). The time dummies are restricted to sum to zero and be orthogonal to the time trend. These data profiles steeply decline with age—see the solid and dash-dot lines in the left panel of Figure 11.

FIGURE 11: LIFECYCLE INEQUALITY. MODEL VS. DATA PROFILES OBTAINED BY RESTRICTING TIME EFFECTS



Notes: The left figure plots in solid and dashed lines lifecycle income and consumption inequality profiles, respectively, estimated from the RLMS data by estimating cohort effects as in Deaton (1997). The rotated income and consumption data profiles are in a solid line with hollow diamonds and a dashed line with hollow circles, respectively. The right panel shows the model-implied income and consumption inequality profiles in a solid line with hollow squares and a dashed line with hollow triangles, respectively, juxtaposed against the rotated data profiles. See Section 4.3.4 for details on the construction of the rotated data and model profiles.

The rotated profiles identified off the model are, however, nearly flat and similar to the profiles from the benchmark calibration that assumed away cohort effects. The fit of the model to the rotated data profiles is at least as good as in the benchmark calibration, with a somewhat superior fit of the consumption profile late in the life cycle—see the right panel of Figure 11 and the bottom panel of Table A-6 in the Online Appendix. The calibrated income process parameters, the time discount factor, the bequest parameters, and the relative risk aversion parameters are similar to their respective values in the calibration that used lifecycle inequality profiles obtained by setting cohort effects to zero as the calibration inputs.

Summing up, the lifecycle model helped us pin down the income process parameters, the persistence of long-lasting shocks, in particular, which is hard to distinguish from unity in

small samples using statistical methods, and the inequality profiles, which may be potentially misspecified when estimated using the standard methods. Moreover, it provides an integrated view of the panel-data evidence on consumption insurance against income shocks and income and consumption inequality over the life cycle in Russia—the lifecycle inequality profiles we uncover are consistent with the estimated insurance of the shocks in panel data.

5 Conclusion

In this paper, we study income and consumption inequality in Russia over time and over the life cycle using household data for 1994–2018. A secular decline in inequality since 1998 that we find was primarily driven by a reduction in the variances of permanent and transitory income shocks. Although consumption insurance against permanent shocks to net family incomes during the period of slow aggregate growth was somewhat lower than during the period of fast growth (the periods prior and post-2008, respectively), the insurance of permanent shocks remained high. Similarly to [Gorodnichenko et al. \(2010\)](#), we find that nearly forty percent of permanent shocks to household net incomes do not pass through to nondurable consumption in Russia. Since this level of insurance is similar to that found in the U.S. by [Blundell et al. \(2008\)](#), in their concluding remarks, [Gorodnichenko et al. \(2010\)](#) label it as a puzzle that calls for more research.

To provide a comprehensive view of Russian inequality and shed light on the puzzle, we further study the age profiles of income and consumption inequality. Due to perfect collinearity between time, cohort, and age effects, isolating age effects from the inequality trends requires placing restrictions on time or cohort effects. Controlling for time effects, found to be the important drivers of inequality, the estimated lifecycle profiles in income and consumption inequality are nearly flat. This is inconsistent with the standard view that household incomes contain a random walk and that income inequality early in the life cycle is low and builds up over time due to different histories of long-lasting shocks across

households. Whereas the measured consumption insurance in panel data is insensitive to the misspecification of the autoregressive component, it is hard to distinguish statistically the persistence of the autoregressive component from unity in small samples. We, therefore, resort to a lifecycle model where, jointly with the other parameters, we calibrate the autoregressive parameter of the income process using income moments and the amount of insurance estimated from the panel data. Calibrating the model, the insurance of persistent shocks and the nearly flat lifecycle profile of consumption inequality we find in the data are consistent with an autoregressive persistence of long-lasting shocks of about 0.94 and high initial variance of the persistent component. Thus, the high insurance of permanent shocks in the panel data is puzzling only relative to the assumption of a random walk in household incomes.

In this paper, we treated the variances of income shocks as exogenous parameters. An interesting avenue for future research would be to uncover the structural determinants of a more than decade-long decline in income inequality and a substantial income risk that Russian households face early in the life cycle.

References

- AGUIAR, M. AND E. HURST (2013): “Deconstructing Life Cycle Expenditure,” *Journal of Political Economy*, 121, 437–492.
- ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): “Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework,” *Econometrica*, 85, 693–734.
- BLUNDELL, R., M. GRABER, AND M. MOGSTAD (2015): “Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family,” *Journal of Public Economics*, 127, 58–73.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98, 1887–1921.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): “Consumption Inequality and Family Labor Supply,” *American Economic Review*, 106, 387–435.

- BROWNING, M., M. EJRNE, AND J. ALVAREZ (2010): “Modelling Income Processes with Lots of Heterogeneity,” *Review of Economic Studies*, 77, 1353–1381.
- BUSSOLO, M. AND P. LUONGO (2020): “The Distributive Impact of Terms of Trade Shocks: The Sase of the Oil Price Changes in Russia,” *Economics of Transition and Institutional Change*, 28, 487–513.
- CALVO, P. A., L. F. LÓPEZ-CALVA, AND J. POSADAS (2015): “A Decade of Declining Earnings Inequality in the Russian Federation,” World Bank Policy Research Working Paper 7392.
- CARROLL, C. D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112, 1–55.
- DE NARDI, M., G. FELLA, AND G. PAZ-PARDO (2019): “Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare,” *Journal of the European Economic Association*, 18, 890–926.
- DE NARDI, M., E. FRENCH, AND J. B. JONES (2016): “Savings After Retirement: A Survey,” *Annual Review of Economics*, 8, 177–204.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1121–1248.
- (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Baltimore, Maryland: The Johns Hopkins University Press.
- DEATON, A. AND C. PAXSON (1994): “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 102, 437–467.
- DING, H. AND H. HE (2018): “A Tale of Transition: An Empirical Analysis of Economic Inequality in Urban China, 1986–2009,” *Review of Economic Dynamics*, 29, 106–137.
- DONOVAN, K., W. J. LU, AND T. SCHOELLMAN (2023): “Labor Market Dynamics and Development,” *The Quarterly Journal of Economics*, 138, 2287–2325.
- ENGBOM, N. AND C. MOSER (2022): “Earnings Inequality and the Minimum Wage: Evidence from Brazil,” *American Economic Review*, 112, 3803–47.
- FELLA, G., S. FRACHE, AND W. KOENIGER (2020): “Buffer-stock Saving and Households’ Response to Income Shocks,” *International Economic Review*, 61, 1359–1382.
- GIMPELSON, V. (2016): “Structural Change and Inter-industry Wage Differentiation,” *Journal of the New Economic Association*, 31, 186–197.
- GIMPELSON, V., R. KAPELIUSHNIKOV, AND A. OSHCHEPKOV (2016): “Return to Tenure Revisited,” *HSE Economic Journal*, 20, 553–587.
- GORODNICHENKO, Y., K. SABIRIANOVA PETER, AND D. STOLYAROV (2010): “Inequality and Volatility Moderation in Russia: Evidence from Micro-level Panel Data on Consumption and Income,” *Review of Economic Dynamics*, 13, 209–237.

- GURIEV, S. AND A. RACHINSKY (2008): “The Evolution of Personal Wealth in the Former Soviet Union and Central and Eastern Europe,” in *Personal Wealth from a Global Perspective*, ed. by J. B. Davies, Oxford: Oxford University Press, 134–149.
- GUVENEN, F. AND A. SMITH (2014): “Inferring Labor Income Risk and Partial Insurance from Economic Choices,” *Econometrica*, 82, 2085–2129.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2005): “Two Views of Inequality Over the Life Cycle,” *Journal of the European Economic Association*, 3, 765–775.
- HRYSHKO, D. (2007): “Excess Smoothness of Consumption in an Estimated Life Cycle Model,” Mimeo, University of Alberta.
- HRYSHKO, D. AND I. MANOVSKII (2022): “How Much Consumption Insurance in the U.S.?” *Journal of Monetary Economics*, 130, 17–33.
- (2023): “Income Dynamics and Consumption Insurance,” Tech. rep., Mimeo, University of Pennsylvania.
- HUGGETT, M., G. VENTURA, AND A. YARON (2011): “Sources of Lifetime Inequality,” *American Economic Review*, 101, 2923–2954.
- JEDWAB, R., P. ROMER, A. M. ISLAM, AND R. SAMANIEGO (2023): “Human Capital Accumulation at Work: Estimates for the World and Implications for Development,” *American Economic Journal: Macroeconomics*, 15, 191–223.
- JOHNSON, D. S., T. M. SMEEDING, AND B. B. TORREY (2005): “Economic Inequality through the Prisms of Income and Consumption,” *Monthly Labor Review*, 128, 11–24.
- KAPELYUK, S. (2015): “The Effect of Minimum Wage on Poverty: Evidence from RLMS-HSE data,” *Economics of Transition*, 23, 389–423.
- KAPLAN, G. AND G. L. VIOLANTE (2010): “How Much Consumption Insurance Beyond Self-Insurance?” *American Economic Journal: Macroeconomics*, 53–87.
- (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82, 1199–1239.
- KOZYREVA, P., M. KOSOLAPOV, AND B. M. POPKIN (2016): “Data Resource Profile: The Russia Longitudinal Monitoring Survey—Higher School of Economics (RLMS-HSE) Phase II: Monitoring the Economic and Health Situation in Russia, 1994–2013,” *International Journal of Epidemiology*, 45, 395–401.
- KRUEGER, D., F. PERRI, L. PISTAFERRI, AND G. L. VIOLANTE (2010): “Cross-sectional Facts for Macroeconomists,” *Review of Economic Dynamics*, 13, 1–14.
- LAGAKOS, D., B. MOLL, T. PORZIO, N. QIAN, AND T. SCHOELLMAN (2018): “Life Cycle Wage Growth across Countries,” *Journal of Political Economy*, 126, 797–849.

- LISE, J., N. SUDO, M. SUZUKI, K. YAMADA, AND T. YAMADA (2014): “Wage, Income and Consumption Inequality in Japan, 1981–2008: From Boom to Lost Decades,” *Review of Economic Dynamics*, 17, 582–612.
- LISINA, A. AND P. VAN KERM (2022): “Understanding Twenty Years of Inequality and Poverty Trends in Russia,” *Review of Income and Wealth*, 68, S108–S130.
- LÓPEZ-CALVA, L. F. AND N. LUSTIG (2010): *Explaining the Decline in Inequality in Latin America: Technological Change, Educational Upgrading, and Democracy*, Brookings Institution Press, 1–24.
- NOVOKMET, F., T. PIKETTY, L. YANG, AND G. ZUCMAN (2018): “From Soviets to Oligarchs: Inequality and Property in Russia, 1905–2016,” *Journal of Economic Inequality*, 16, 189–223.
- SANTAEULÀLIA-LLOPIS, R. AND Y. ZHENG (2018): “The Price of Growth: Consumption Insurance in China 1989–2009,” *American Economic Journal: Macroeconomics*, 10, 1–35.
- SCHULHOFER-WOHL (2018): “The Age-time-cohort Problem and the Identification of Structural Parameters in Life-cycle Models,” *Quantitative Economics*, 9, 643–658.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): “Consumption and Risk Sharing Over the Life Cycle,” *Journal of Monetary Economics*, 51, 609–633.
- YEMTSOV, R. (2008): “Through the Looking-Glass: What is Behind Official Data on Inequality in Russia over 1992–2003?” Tech. rep.

Online Appendix to “Inequality in Russia over Time and over the Life Cycle” by Maxym Bryukhanov and Dmytro Hryshko

TABLE A-1: SAMPLE MEANS FOR SELECTED YEARS. RLMS DATA

Year	Age	College	Number of kids	Family size	Lives in town
1994	41.68 (0.22)	0.26 (0.01)	1.00 (0.02)	3.63 (0.03)	0.77 (0.01)
1998	42.10 (0.23)	0.24 (0.01)	0.96 (0.02)	3.69 (0.03)	0.75 (0.01)
2002	42.48 (0.23)	0.25 (0.01)	0.83 (0.02)	3.72 (0.03)	0.73 (0.01)
2006	43.45 (0.22)	0.23 (0.01)	0.71 (0.02)	3.65 (0.03)	0.72 (0.01)
2010	43.35 (0.18)	0.25 (0.01)	0.77 (0.02)	3.63 (0.02)	0.70 (0.01)
2014	43.49 (0.22)	0.27 (0.01)	0.79 (0.02)	3.59 (0.03)	0.71 (0.01)
2018	43.88 (0.21)	0.27 (0.01)	0.85 (0.02)	3.63 (0.03)	0.72 (0.01)

Notes: Statistics are based on our RLMS sample that includes households with married or cohabiting couples whose male is aged 25–59. We apply survey household weights when calculating the means. Bootstrap standard errors in parentheses. “College” is a dummy variable that equals one if an individual finished at least a college degree and zero otherwise. See Section 2 for more details on sample construction.

TABLE A-2: SAMPLE MEANS FOR SELECTED YEARS. HBS DATA

Year	Age	College	Number of kids	Family size	Lives in town
2003	43.51 (0.15)	—	0.65 (0.01)	3.48 (0.02)	0.74 (0.01)
2006	44.23 (0.17)	0.29 (0.01)	0.65 (0.01)	3.48 (0.02)	0.74 (0.01)
2010	43.87 (0.18)	0.29 (0.01)	0.70 (0.01)	3.36 (0.02)	0.73 (0.01)
2014	43.81 (0.14)	0.35 (0.01)	0.82 (0.02)	3.37 (0.02)	0.74 (0.01)

Notes: Statistics are based on our HBS sample that includes households whose male head is aged 25–59. We apply survey household weights when calculating the means. Bootstrap standard errors in parentheses. “College” is a dummy variable that equals one if an individual finished at least a college degree and zero otherwise. See Section 2 for more details on sample construction. Data on education in 2003 is not available.

TABLE A-3: CALIBRATION OF A MODEL WITH BEQUESTS

Parameter	Notation	Value	Internally calibrated	Data source
Income age profile	μ_t	various	N	RLMS
Survival probabilities	s_t	various	N	Human Mortality Database
Replacement rate	κ	0.50	N	OECD
Interest rate	r	−2%	N	OECD
CRRA	γ	3.355	Y	
Variance of persistent income shocks	σ_ξ^2	0.044	Y	
Variance of transitory income shocks	σ_ϵ^2	0.091	Y	
Variance of fixed effects	σ_α^2	0.172	Y	
Persistence of long-lasting shocks	ρ	0.941	Y	
Time discount factor	β	0.918	Y	
Bequest motive, intensity	b_1	0.122	Y	
Bequest motive, curvature	b_2	0.285	Y	
Bequest motive, power	λ	1.933	Y	

TABLE A-4: MOMENTS FROM THE DATA AND FROM THE MODEL WITH BEQUESTS

	Data	Model
$\text{var}[y_{it}]$, ages 25–29	0.69	0.66
$\text{var}[\Delta y_{it}]$, ages 25–59	0.217	0.228
$\text{var}[\Delta^2 y_{it}]$, ages 25–59	0.289	0.270
$\text{var}[\Delta^3 y_{it}]$, ages 25–59	0.305	0.310
Transmission of persistent shocks	0.61	0.64
Transmission of transitory shocks	0.12	0.12
Wealth-to-income ratio	3.00	3.17

Notes: The growth of idiosyncratic income over a k -year horizon is measured as $\Delta^k y_{it} = y_{it} - y_{it-k}$, $k = 2, 3$. In addition to the data moments listed in column 1, we also targeted the lifecycle profile of consumption inequality. See Figure 9 for the fit of the model-implied income and consumption inequality profiles to the data.

TABLE A-5: CALIBRATION TARGETING INEQUALITY PROFILES OVER THE LIFE CYCLE
OBTAINED BY SETTING COHORT EFFECTS TO ZERO

Panel A: Calibrated parameters		
CRRA, γ		3.297
Variance of persistent income shocks, σ_{ξ}^2		0.052
Variance of transitory income shocks, σ_{ϵ}^2		0.084
Variance of fixed effects, σ_{α}^2		0.115
Persistence of long-lasting shocks, ρ		0.941
Time discount factor, β		0.913
k_i		0.001
k_c		0.002
Bequest motive, intensity, b_1		0.192
Bequest motive, curvature, b_2		0.283
Bequest motive, power, λ		1.984

Panel B: Data targets and model fit		
	Data	Model
$\text{var}[\Delta y_{it}]$, ages 25–59	0.217	0.222
$\text{var}[\Delta^2 y_{it}]$, ages 25–59	0.289	0.272
$\text{var}[\Delta^3 y_{it}]$, ages 25–59	0.305	0.319
Transmission of persistent shocks	0.61	0.62
Transmission of transitory shocks	0.12	0.13
Wealth-to-income ratio	3.00	3.24
Income variances over the life cycle		
Ages 25–29	0.662	0.665
Ages 30–34	0.672	0.677
Ages 35–39	0.681	0.664
Ages 40–44	0.665	0.664
Ages 45–49	0.672	0.658
Ages 50–54	0.683	0.659
Ages 55–59	0.637	0.662
Consumption variances over the life cycle		
Ages 25–29	0.433	0.438
Ages 30–34	0.406	0.432
Ages 35–39	0.424	0.412
Ages 40–44	0.425	0.404
Ages 45–49	0.420	0.409
Ages 50–54	0.458	0.435
Ages 55–59	0.446	0.491

Notes: The growth of idiosyncratic income over a k -year horizon is measured as $\Delta^k y_{it} = y_{it} - y_{it-k}$, $k = 2, 3$. See Figure 10 for the fit of the model-implied income and consumption inequality profiles to the data.

TABLE A-6: CALIBRATION TARGETING INEQUALITY PROFILES OVER THE LIFE CYCLE
OBTAINED BY NORMALIZING TIME EFFECTS

Panel A: Calibrated parameters		
CRRA, γ		3.304
Variance of persistent income shocks, σ_ξ^2		0.043
Variance of transitory income shocks, σ_ϵ^2		0.091
Variance of fixed effects, σ_α^2		0.111
Persistence of long-lasting shocks, ρ		0.941
Time discount factor, β		0.923
k_i		0.024
k_c		0.014
Bequest motive, intensity, b_1		0.204
Bequest motive, curvature, b_2		0.283
Bequest motive, power, λ		2.099
Panel B: Data targets and model fit		
	Data	Model
var $[\Delta y_{it}]$, ages 25–59	0.217	0.224
var $[\Delta^2 y_{it}]$, ages 25–59	0.289	0.265
var $[\Delta^3 y_{it}]$, ages 25–59	0.305	0.304
Transmission of persistent shocks	0.61	0.64
Transmission of transitory shocks	0.12	0.12
Wealth-to-income ratio	3.00	3.07
Income variances over the life cycle		
Ages 25–29	0.578	0.583
Ages 30–34	0.575	0.593
Ages 35–39	0.569	0.582
Ages 40–44	0.564	0.582
Ages 45–49	0.566	0.577
Ages 50–54	0.567	0.578
Ages 55–59	0.583	0.581
Consumption variances over the life cycle		
Ages 25–29	0.382	0.378
Ages 30–34	0.363	0.367
Ages 35–39	0.356	0.347
Ages 40–44	0.355	0.340
Ages 45–49	0.351	0.346
Ages 50–54	0.392	0.374
Ages 55–59	0.421	0.431

Notes: The growth of idiosyncratic income over a k -year horizon is measured as $\Delta^k y_{it} = y_{it} - y_{it-k}$, $k = 2, 3$. The raw inequality profiles are obtained by normalizing time effects as in Deaton (1997). See Figure 11 for the fit of the model-implied income and consumption inequality profiles to the data.

I Trends in inequality: time or cohort effects?

In this appendix, we evaluate the importance of time versus cohort effects for the trends in income and consumption inequality following the methodology of [Heathcote et al. \(2005\)](#). They decompose the variance of income, hours, and consumption in the U.S. into a sum of age, time, and cohort effects. Let a denote age, t year, and k cohort (year of birth) so that $k = t - a$. The variance of x , where in our case x is either log consumption or log income, can be decomposed as

$$\text{var}[x(a, t, t - a)] = g_1^x(a) + g_2^x(t) + g_3^x(t - a).$$

[Heathcote et al. \(2005\)](#) make use of the following quantities:

$$\Delta \text{var}[x_{t,t+1}^a] = \Delta g_2^x(t + 1) + \Delta g_3^x(t + 1 - a) \tag{A1}$$

$$\Delta \text{var}[x_{t,t+1}^k] = \Delta g_1^x(a + 1) + \Delta g_2^x(t + 1) \tag{A2}$$

$$\Delta \text{var}[x_{a,a+1}^{t+1}] = \Delta g_1^x(a + 1) - \Delta g_3^x(t + 1 - a). \tag{A3}$$

where $g_j^x(z+1) \equiv g_j^x(z+1) - g_j^x(z)$, $j = 1, 2$, $z = a$ if $j = 1$ and $z = t$ if $j = 2$, $\Delta g_3^x(t+1-a) = g_3^x(t+1-a) - g_3^x(t-a)$.

Eq. (A1) measures the change in cross-sectional variances for individuals of the same age across time. It, therefore, differences out age effects and reflects the influence of time and cohort effects only. Eq. (A2) measures the change in cross-sectional variances for individuals of the same cohort across time. Differencing cancels out cohort effects and results in a function of age and time effects only. Eq. (A3) measures the change in cross-sectional variances for individuals of ages a and $a + 1$ at a point in time and therefore cancels out time effects and is a function of age and cohort effects only. Taking the time $t + 1$ averages

of Eqs. (A1) and (A3) across ages, and Eq. (A2) across cohorts gives:

$$\bar{\Delta}\text{var}[x_{t,t+1}^a] = \Delta g_2^x(t+1) + \bar{\Delta}g_3^x(t+1) \quad (\text{A4})$$

$$\bar{\Delta}\text{var}[x_{t,t+1}^k] = \bar{\Delta}g_1^x + \Delta g_2^x(t+1) \quad (\text{A5})$$

$$\bar{\Delta}\text{var}[x_{a,a+1}^{t+1}] = \bar{\Delta}g_1^x - \bar{\Delta}g_3^x(t+1). \quad (\text{A6})$$

The following implications can be tested.

1. If cohort effects are small—that is, $\bar{\Delta}g_3^x(t+1) \approx 0$:
 - (a) the correlation of (average) within-cohort and within-age changes in variances should be close to one, all due to time effects (see Eqs. A4 and A5);
 - (b) within-time changes in variances should not vary with time (Eq. A6);
 - (c) the correlation of (average) within-age and within-time changes in variances should be indistinguishable from zero (see Eqs. A4 and A6).
2. If time effects are small—that is, $\bar{\Delta}g_2^x(t+1) \approx 0$:
 - (a) within-cohort changes in variances should not vary with time (Eq. A5);
 - (b) (average) within-age and within-time changes in variances must be strongly negatively correlated due to cohort effects (see Eqs. A4 and A6).

Since our HBS data is cross-sectionally large, we can obtain relatively precise estimates of the variance changes in Eqs. (A1)–(A3) for each a , k , and t and their time averages. Table A-7 reports the averages for nondurable consumption from the HBS for the years 2003–2014.

Implication 1(a) is supported by our data, as can be readily seen in the next-to-last row of Table A-7, where the estimated correlation equals 0.996. This high correlation is also possible in the absence of time effects, $\Delta g_2^x(t+1) \approx 0$, when age effects are a multiple

of cohort effects. This would imply that the (average) within-cohort changes in variances, measured by Eq. (A5), are constant over time. We provide a formal test of this hypothesis, 2(b), in the bottom row of Table A-7. It can be rejected at about 4% significance level. Although the correlation of within-age and within-time changes in variances is on average 0.18 and different from zero, its standard error, calculated by bootstrap, is large, so we cannot reject the null that it equals zero. Thus, 1(c) is also supported by the data. Since these variance changes are not strongly negatively correlated, 2(b) is not supported by the data. Finally, 1(b) cannot be rejected at any plausible significance level.

Piecing the evidence together, one may conclude that the time effects strongly affect the evolution of variances of nondurable consumption over time, while cohort effects appear to have little to no influence on the trends in the variances.

TABLE A-7: CHANGES IN CROSS-SECTIONAL VARIANCES OF HOUSEHOLD NONDURABLE CONSUMPTION. HBS DATA

Year	$\bar{\Delta}\text{var}[x_{t,t+1}^a]$ (1)	$\bar{\Delta}\text{var}[x_{t,t+1}^k]$ (2)	$\bar{\Delta}\text{var}[x_{a,a+1}^t]$ (3)
2003	-1.062 (-0.598)	-1.043 (-0.973)	-0.190 (-0.105)
2004	-3.483** (-2.458)	-3.512*** (-3.444)	-0.127 (-0.076)
2005	-0.975 (-0.740)	-1.317 (-0.923)	0.119 (0.088)
2006	0.931 (0.823)	1.100 (0.850)	0.023 (0.016)
2007	-6.118*** (-4.958)	-6.295*** (-5.091)	-0.278 (-0.184)
2008	-2.819*** (-3.120)	-2.972*** (-3.727)	-0.126 (-0.103)
2009	-0.547 (-0.480)	-0.310 (-0.299)	0.141 (0.139)
2010	1.906 (1.146)	2.267 (1.369)	0.350 (0.253)
2011	-2.841* (-1.915)	-3.053** (-2.110)	-0.113 (-0.073)
2012	1.098 (0.668)	1.057 (0.666)	-0.213 (-0.150)
2013	-3.245** (-2.510)	-3.194* (-1.955)	0.326 (0.158)
2014	— —	— —	-0.240 (-0.230)
Correlation with (1): (s.e.)		0.998 (0.004)	0.177 (0.235)
Test of equal coeff., p-value:		3.6%	100%

Notes: $\bar{\Delta}\text{var}[x_{t,t+1}^a]$ is the within-age change in variances averaged across all ages in a given year. $\bar{\Delta}\text{var}[x_{t,t+1}^k]$ is the within-cohort change in variances averaged across all cohorts in a given year. $\bar{\Delta}\text{var}[x_{a,a+1}^t]$ is the between-age change in variances averaged across all ages in a given year. All entries are multiplied by 100. We used survey household weights to calculate the variances. t-statistics in parentheses. Standard errors for the correlation coefficients calculated by bootstrap. *** (**) [*] denotes significance at the 1% (5%) [10%] level.

Next, we perform the same analysis for net family incomes using the RLMS since the HBS data do not contain reliable income information. Because the RLMS is much smaller cross-

TABLE A-8: CHANGES IN CROSS-SECTIONAL VARIANCES OF NET FAMILY INCOME. RLMS
DATA

Year	$\bar{\Delta}\text{var}[x_{t,t+1}^a]$ (1)	$\bar{\Delta}\text{var}[x_{t,t+1}^k]$ (2)	$\bar{\Delta}\text{var}[x_{a,a+1}^t]$ (3)	Year	$\bar{\Delta}\text{var}[x_{t,t+1}^a]$ (1)	$\bar{\Delta}\text{var}[x_{t,t+1}^k]$ (2)	$\bar{\Delta}\text{var}[x_{a,a+1}^t]$ (3)
1994	-5.37 (-1.31)	-6.65* (-1.76)	-6.05 (-1.05)	2007	0.57 (0.16)	-0.16 (-0.03)	0.93 (0.16)
1995	14.54*** (3.11)	13.18*** (2.77)	-4.99 (-0.62)	2008	-5.37** (-2.19)	-4.33 (-1.61)	1.26 (0.36)
1996	—	—	-4.08 (-0.75)	2009	2.39 (0.79)	2.63 (0.81)	-0.90 (-0.33)
1998	—	—	-2.42 (-0.75)	2010	-4.19* (-1.73)	-4.21** (-2.33)	-0.90 (-0.22)
2000	-7.42 (-1.53)	-8.47 (-1.49)	-2.26 (-0.24)	2011	0.25 (0.17)	2.48 (0.68)	0.02 (0.01)
2001	-2.93 (-0.49)	4.00 (0.34)	-5.06 (-0.63)	2012	-1.13 (-1.00)	-2.87 (-1.03)	0.07 (0.02)
2002	-1.71 (-0.28)	-6.06 (-0.70)	-6.88 (-1.10)	2013	-1.66 (-0.85)	-1.32 (-1.30)	0.73 (0.26)
2003	-7.51** (-2.31)	-5.70 (-1.18)	0.90 (0.89)	2014	-3.15 (-1.48)	-3.24 (-1.56)	0.87 (0.21)
2004	-8.44** (-2.14)	-9.10** (-2.20)	-2.88 (-0.57)	2015	1.32 (0.99)	1.66 (1.35)	1.66 (1.29)
2005	2.21 (0.61)	2.11 (0.57)	-1.57 (-0.49)	2016	-1.59 (-1.20)	-2.02 (-1.51)	0.77 (0.44)
2006	-4.92 (-1.33)	-3.30 (-1.65)	1.00 (0.20)	2017	-0.70 (-0.98)	-1.17 (-0.74)	-0.41 (-0.24)
				2018	—	—	-0.03 (-0.03)
Corr. with (1): (s.e.)						0.909 (0.064)	0.083 (0.229)
Test of equal coeff., p-value:						3.6%	99.7%

Notes: $\bar{\Delta}\text{var}[x_{t,t+1}^a]$ is the within-age change in variances averaged across all age groups in a given year. $\bar{\Delta}\text{var}[x_{t,t+1}^k]$ is the within-cohort change in variances averaged across all cohorts in a given year. $\bar{\Delta}\text{var}[x_{a,a+1}^t]$ is the between-age change in variances averaged across all age groups in a given year. All entries are multiplied by 100. We used survey household weights to calculate the variances. t-statistics in parentheses. Standard errors for the correlation coefficients calculated by bootstrap. *** (**) [*] denotes significance at the 1% (5%) [10%] level.

sectionally, we group individuals into age and cohort bins so that age a represents individuals of ages $a - 2$ through $a + 2$ and individuals from cohort k represent individuals born in years $k - 2$ through year $k + 2$. The results are presented in Table A-8. The bottom two rows

of the table point to the same conclusion that we reached with the consumption data—the decline in cross-sectional variances of incomes appears to be due to the time effects, with the cohort effects having little to no influence on the trend.

FIGURE A-1: LORENZ CURVES FOR INCOME DISTRIBUTIONS. RLMS DATA

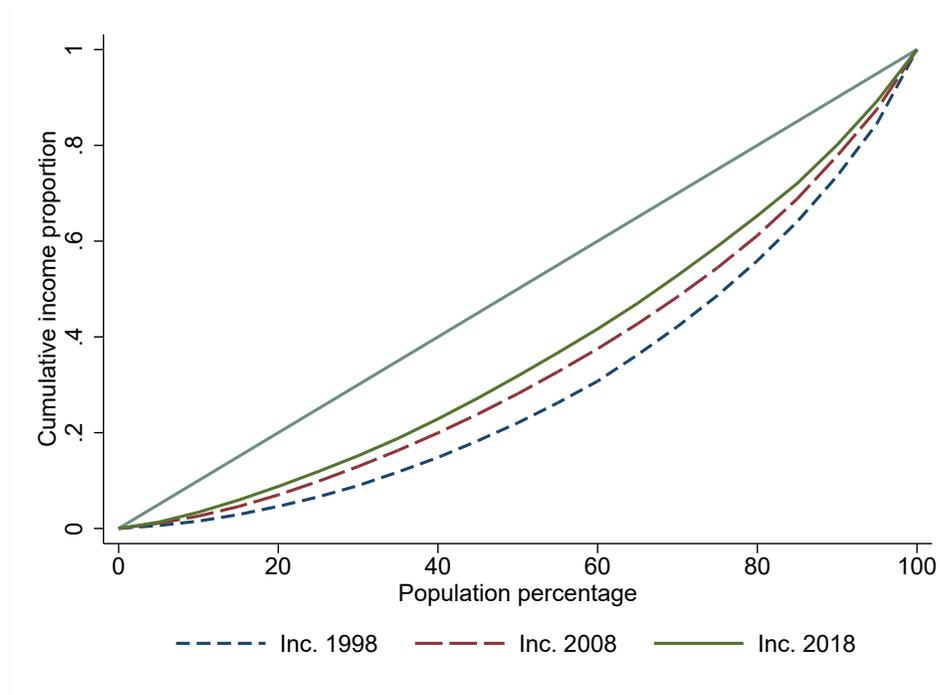


FIGURE A-2: LORENZ CURVES FOR CONSUMPTION DISTRIBUTIONS. RLMS DATA

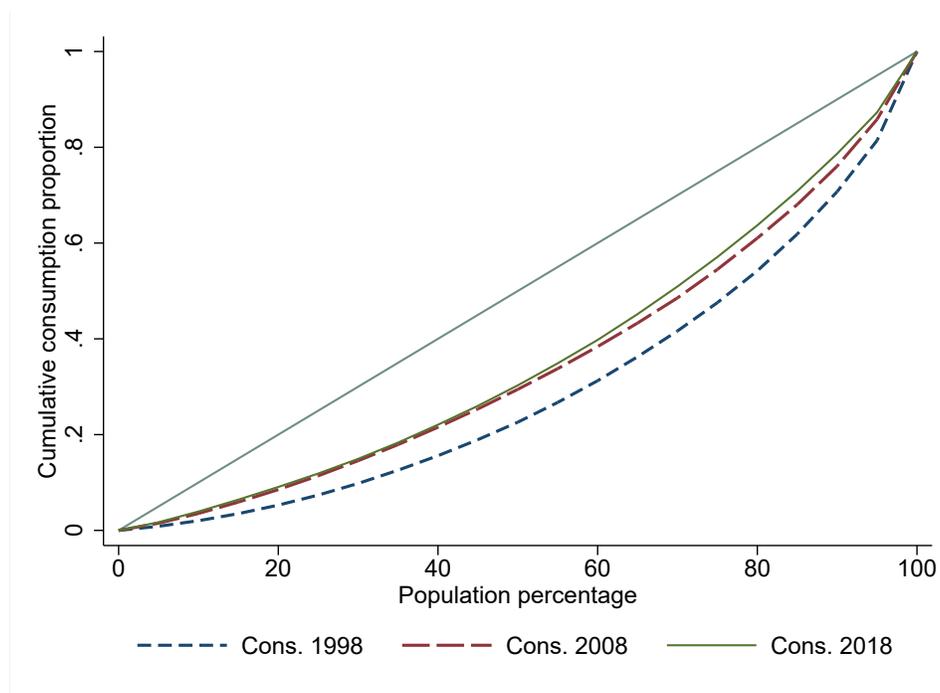
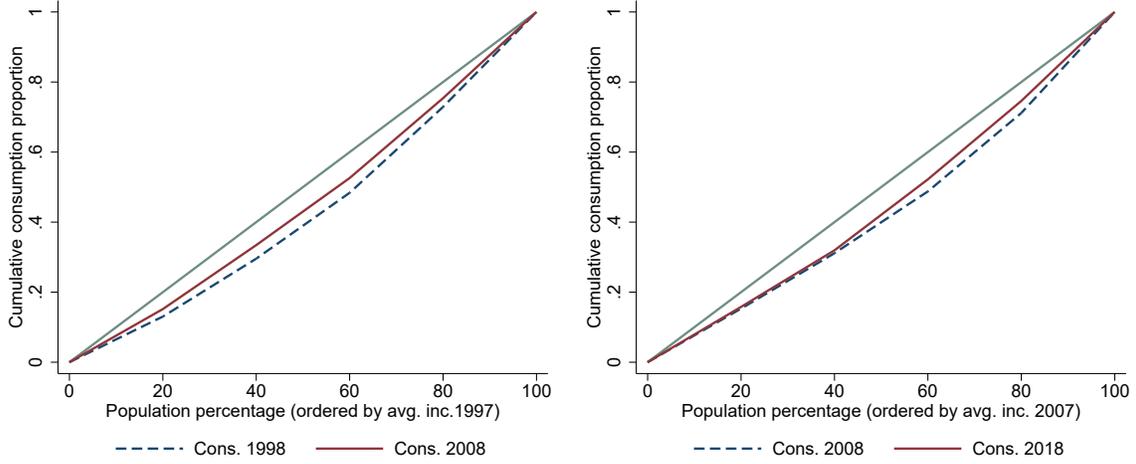
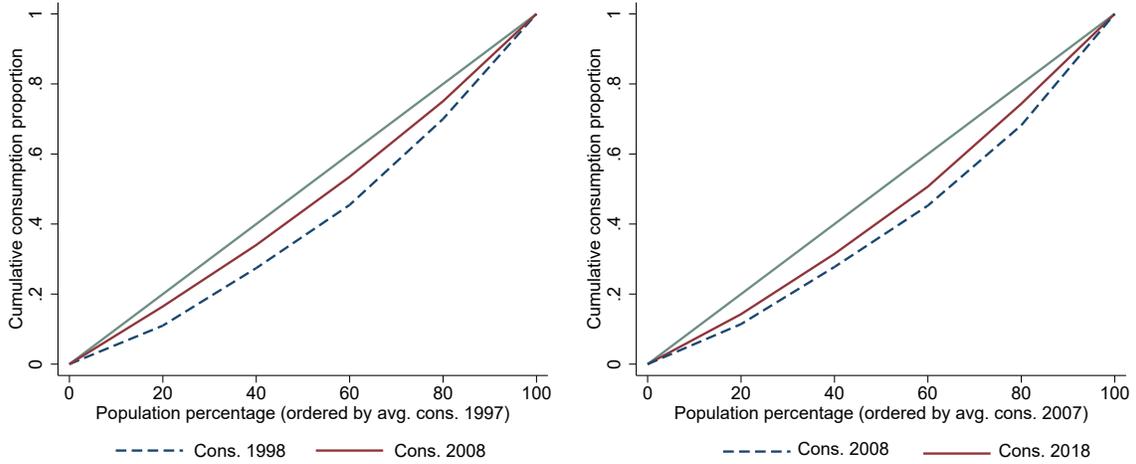


FIGURE A-3: CONCENTRATION CURVES FOR CONS. DISTRIBUTIONS. RLMS DATA



(a) Ordered by past average income

Notes: We plot consumption distributions for years t and $t + 10$ ordered by average income in year $t - 1$. In the left (right) panel $t = 1998$ (2008). Average income is calculated as the average over the years $t - 4$ to $t - 1$.

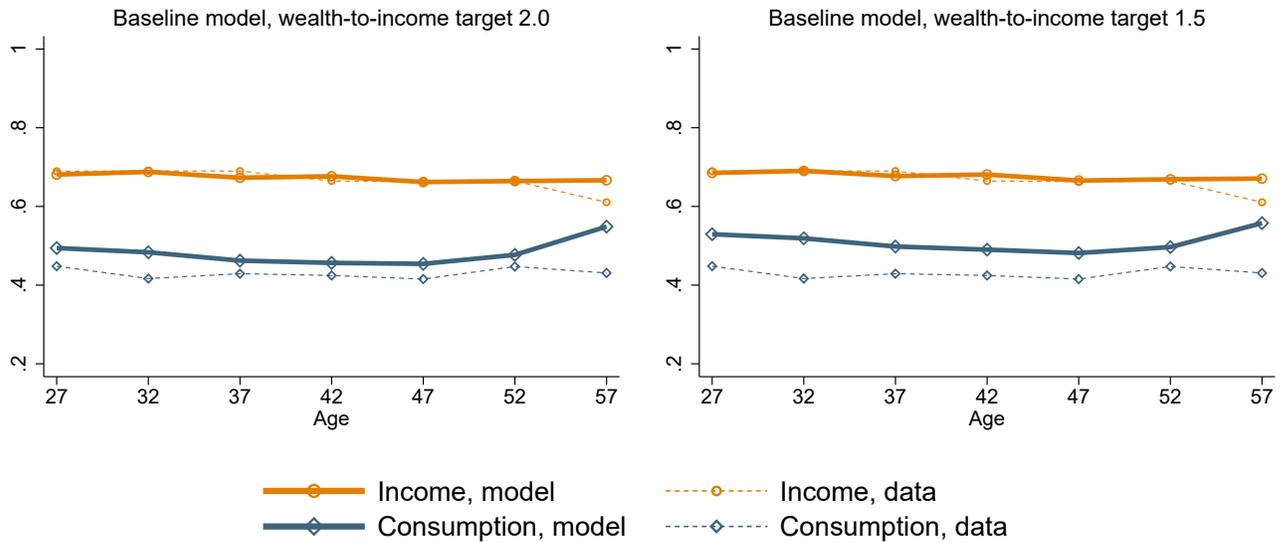


(b) Ordered by past average consumption

Notes: We plot consumption distributions for years t and $t + 10$ ordered by average household consumption in year $t - 1$. In the left (right) panel $t = 1998$ (2008). Average consumption is calculated as the average over the years $t - 4$ to $t - 1$.

II Additional calibration results

FIGURE B-1: LIFECYCLE INEQUALITY. VARIOUS WEALTH-TO-INCOME TARGETS



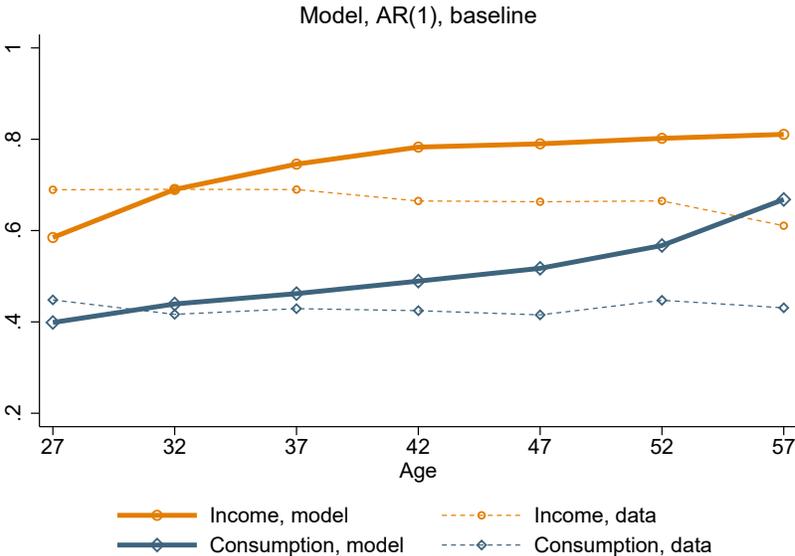
Notes: The left figure plots, in solid lines, the inequality profiles implied by a calibration targeting the wealth-to-income ratio of 2, whereas the right figure plots, in solid lines, the inequality profiles implied by the calibration that targets the wealth-to-income ratio of 1.5. Dashed lines in both figures represent the data inequality profiles plotted in Figure 7. Both calibrations do not target consumption and income inequality profiles.

TABLE B-1: BASELINE MODEL. VARIOUS WEALTH-TO-INCOME TARGETS

<i>Wealth-to-income target:</i>	3.0	2.0	1.5
Panel A: Calibrated parameters			
CRRA, γ	3.0	3.42*	3.35*
Variance of persistent income shocks, σ_{ξ}^2	0.055	0.041	0.033
Variance of transitory income shocks, σ_{ϵ}^2	0.082	0.095	0.101
Variance of fixed effects, σ_{α}^2	0.119	0.231	0.297
Persistence of long-lasting shocks, ρ	0.940	0.940	0.940
Time discount factor, β	0.939	0.916	0.917
Panel B: Data and model moments			
var[y_{it}], ages 25–29 [0.69]	0.69	0.69	0.69
var[Δy_{it}], ages 25–59 [0.217]	0.220	0.229	0.234
var[$\Delta^2 y_{it}$], ages 25–59 [0.289]	0.271	0.267	0.264
var[$\Delta^3 y_{it}$], ages 25–59 [0.305]	0.320	0.304	0.293
Transmission of persistent shocks [0.61]	0.62	0.69	0.76
Transmission of transitory shocks [0.12]	0.12	0.13	0.14
Wealth-to-income ratio	3.0	2.25	1.85

Notes: *Calibrated together with the other parameters listed in Panel A. In Panel B, values for the data moments are in square brackets. The growth of idiosyncratic income over a k -year horizon is measured as $\Delta^k y_{it} = y_{it} - y_{it-k}$, $k = 2, 3$. See Figure B-1 for the fit of the model-implied income and consumption inequality profiles to the data when the model targets the wealth-to-income ratio of 1.5 and 2.0.

FIGURE B-2: INEQUALITY. LOW INITIAL VARIANCE OF THE PERSISTENT COMPONENT



Notes: The calibration targets the baseline model’s moments listed in Table 3. In this calibration, we assume that the initial variance of the persistent component equals half the size of its long-run variance, $0.5 \cdot \frac{\sigma_{\xi}^2}{1-\rho^2}$.

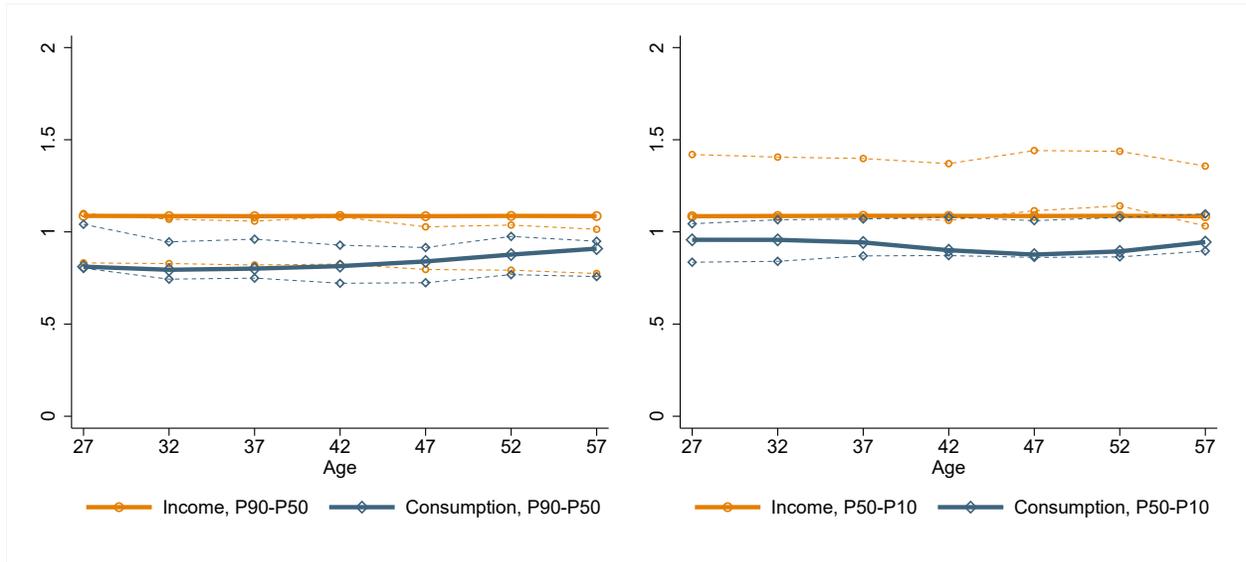
TABLE B-2: BASELINE MODEL. $\text{VAR}(P_{i0}) = 0.5 \cdot \frac{\sigma_\xi^2}{1-\rho^2}$

Panel A: Calibrated parameters		
CRRA, γ		3.0
Variance of persistent income shocks, σ_ξ^2		0.061
Variance of transitory income shocks, σ_ϵ^2		0.077
Variance of fixed effects, σ_α^2		0.217
Persistence of long-lasting shocks, ρ		0.940
Time discount factor, β		0.936

Panel B: Data and model moments		
	Data	Model
$\text{var}[y_{it}]$, ages 25–29	0.69	0.69
$\text{var}[\Delta y_{it}]$, ages 25–59	0.217	0.216
$\text{var}[\Delta^2 y_{it}]$, ages 25–59	0.289	0.273
$\text{var}[\Delta^3 y_{it}]$, ages 25–59	0.305	0.327
Transmission of persistent shocks	0.61	0.60
Transmission of transitory shocks	0.12	0.12
Wealth-to-income ratio	3.00	2.93

Notes: The growth of idiosyncratic income over a k -year horizon is measured as $\Delta^k y_{it} = y_{it} - y_{it-k}$, $k = 2, 3$. See the right panel of Figure B-2 for the fit of the model-implied income and consumption inequality profiles to the data.

FIGURE B-3: INCOME AND CONSUMPTION PERCENTILES IN THE DATA AND MODEL



Notes: The figure plots, in solid lines, the log 90/50 and 50/10 ratios for income and consumption implied (and not targeted) by the baseline model's calibration. Dashed lines in both figures represent the 95% confidence intervals, calculated by bootstrap, for the relative log percentile ratios in the data.

Appendix References

DEATON, A. (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Baltimore, Maryland: The Johns Hopkins University Press.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2005): “Two Views of Inequality Over the Life Cycle,” *Journal of the European Economic Association*, 3, 765–775.