



Heterogeneous life-cycle profiles, income risk and consumption inequality[☆]

Giorgio E. Primiceri, Thijs van Rens^{*}

^a Northwestern University, NBER and CEPR, USA

^b CREI, Universitat Pompeu Fabra and CEPR, Ramon Trias Fargas 25, 08005 Barcelona, Spain

ARTICLE INFO

Article history:

Received 1 May 2008

Received in revised form

7 October 2008

Accepted 7 October 2008

Available online 18 October 2008

JEL classification:

D12

D31

D52

D91

E21

Keywords:

Consumption

Inequality

Risk

Incomplete markets

Heterogeneity

ABSTRACT

Was the increase in income inequality in the US due to permanent shocks or merely to an increase in the variance of transitory shocks? The implications for consumption and welfare depend crucially on the answer to this question. We use Consumer Expenditure Survey (CEX) repeated cross-section data on consumption and income to decompose idiosyncratic changes in income into predictable life-cycle changes, transitory and permanent shocks and estimate the contribution of each to total inequality. Our model fits the joint evolution of consumption and income inequality well and delivers two main results. First, we find that permanent changes in income explain all of the increase in inequality in the 1980s and 1990s. Second, we reconcile this finding with the fact that consumption inequality did not increase much over this period. Our results support the view that many permanent changes in income are predictable for consumers, even if they look unpredictable to the econometrician, consistent with models of heterogeneous income profiles.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

This paper evaluates the nature of increased income inequality in the US over the 1980–2000 period. This is important because income inequality originating from different sources may have different implications for consumption inequality and welfare. For example, under standard models of consumption smoothing, households do not adjust their consumption much in response to transitory shocks to their income. Hence, increases in income inequality generated by transitory shocks will have only very small effects on consumption inequality and welfare. Similarly, consumption does not respond to permanent changes in income that are insured or foreseen in advance. On the other hand, unexpected and uninsurable permanent income shocks will translate almost one-for-one into changes in consumption and will, therefore, have strong welfare effects.

[☆] This paper was prepared for the April 2008 Carnegie-Rochester Conference on Public Policy. We are grateful to Pierre-Olivier Gourinchas and Jonathan Parker for many helpful discussions and to Richard Blundell, Chris Carroll, Guido Lorenzoni, Alberto Martin, Claudio Michelacci, Josep Pijoan-Mas, Chris Telmer, our discussants Miklós Koren, Michele Pellizzari and especially Jonathan Heathcote, and the editor Mark Bilal for helpful suggestions. Thijs van Rens gratefully acknowledges financial support from the Spanish Ministry of Education and Science (Grants Juan de la Cierva, SEJ2005-01124 and SEJ2006-02235); the Generalitat de Catalunya, DURSI (Grants Beatriu de Pinós and 2005SGR00490) and the Barcelona GSE Research Network.

^{*} Corresponding author. Tel.: +34 93 542 2598; fax: +34 93 542 1860.

E-mail address: tvans@crei.cat (T. van Rens).

We use repeated cross-section data on income and consumption from the Consumer Expenditure Survey (CEX) to estimate the extent to which different types of income shocks have contributed to the evolution of inequality. In order to extract this information, we need to put some structure on the data. More precisely, we make assumptions on the form of the stochastic process governing the evolution of individual income and postulate a model of consumption choice. These assumptions allow us to map cross-sectional variances of income and consumption within a cohort (inequality) into variances of permanent and transitory shocks (risk).

In our model, income follows an exogenous stochastic process driven by permanent and transitory shocks. We assume that consumers can self-insure against transitory shocks. In addition, we allow for permanent changes in income that do not translate into changes in consumption. We model these permanent income ‘shocks’ that do not affect consumption as heterogeneity: changes in income over the life-cycle that are predictable to the consumer. If in reality there are other reasons why changes in consumption do not reflect permanent changes in income, then we will overestimate the contribution of heterogeneity to inequality. We discuss this issue, in particular the possibility that there exist insurance markets that allow for risk sharing between consumers, and argue that there is a role for heterogeneity over and above risk sharing.

Our study delivers two main results. First, essentially all of the increase in income inequality over the sample period is due to an increase in the cross-sectional variance of *permanent* shocks to income. Second, most of these permanent income shocks were *not* of the kind that gets transmitted to consumption. Therefore, our estimates point to heterogeneity as a major source of the increase in inequality in the 1980s. The variance of transitory and unpredictable permanent shocks to income also increased in the early 1980s, but the increase was small compared to the total increase in inequality and got reversed by the end of the 1990s.

The intuition behind these results is straightforward. The trends in the data can be characterized by three salient features: (i) individual income is highly persistent over the whole sample period, (ii) income inequality rose substantially in the 1990s and, particularly, in the 1980s but (iii) over the same period consumption inequality did not increase much. If the evolution of income inequality were driven by transitory shocks, we should see much lower autocorrelation in individual income processes. If unexpected and uninsurable permanent shocks were the driving force, we would expect a rise in consumption inequality accompanying the increase in income inequality. This leaves only the third candidate, heterogeneity, able to explain all aspects of the data.

We are not the first to notice that consumption does not respond to permanent income shocks as much as standard models would predict. This finding is typically interpreted as evidence that consumers have access to markets that allow them to share risks with other consumers, insuring some or all of their idiosyncratic shocks (Krueger and Perri, 2006; Storesletten et al., 2004b; Primiceri and Van Rens, 2004; Pistaferri et al., forthcoming; Heathcote et al., 2006).¹ In this paper, we offer an alternative explanation. If there is heterogeneity in life-cycles across consumers, as Lillard and Weiss (1979) and more recently Guvenen (2005, 2007) have argued, then consumption may not reflect changes in income, even if they are permanent, because these changes are predictable to the household in advance.

Partial risk sharing and heterogeneity are observationally equivalent in our model. Nevertheless, we argue that it is unlikely that risk sharing is the sole mechanism responsible for the muted response of consumption to permanent shocks. First, the degree of risk sharing necessary to match the data would have to be substantially higher than what other studies have found (Attanasio and Davis, 1996). Second, we test a number of predictions of the risk sharing hypothesis (some risk sharing happens through government taxes and transfers or through markets for financial assets) and do not find convincing evidence for any of these. Finally, our interpretation that heterogeneity is an important driver of inequality is consistent with a number of recent papers decomposing inequality in heterogeneity and uncertainty, using schooling choices (Cunha et al., 2005; Cunha and Heckman, 2006; Huggett et al., 2006). The identifying assumption in these papers as well as in ours is that heterogeneity, even if unobservable to econometrician, is forecastable to the consumer and therefore affects her choices. Then, using an observable outcome of those choices, one can identify heterogeneity from risk. The main difference is that in our case the observable is not the individual’s education level but her consumption choice. The fact that both ‘instruments’ yield similar conclusions about the sources of inequality provides additional support for our interpretation.

Earlier investigations of the sources of increase in income inequality have followed either of two alternative approaches. Carroll (1992), Gottschalk and Moffitt (1994) and Moffitt and Gottschalk (1995, 2002) use only data on income, thus avoiding having to model consumer behavior and arguing that the autocovariance structure of income growth is informative about the relative importance of permanent and transitory shocks. In particular, Gottschalk and Moffitt (1994) exploit the long panel dimension of the Michigan PSID. They emphasize the contribution of transitory inequality, but nevertheless conclude that approximately two-thirds of the increase in inequality between the 1970–1978 and 1979–1987 periods was due to permanent shocks and only one-third to transitory shocks. Moreover, Moffitt and Gottschalk (2002) find that transitory shocks contributed negatively to the overall evolution of income inequality in the 1990s.

¹ Recently, the basic finding that the increase in the volatility of shocks to income in the 1980s did not translate into an increase in consumption risk has been questioned. Attanasio et al. (2004) use the CEX diary survey to show that the increase in consumption inequality was more pronounced for frequently purchased items like food. Gorbachev (2007) uses Panel Study on Income Dynamics (PSID) data and confirms that the volatility of annual changes in individual food consumption increased substantially over this period.

On the other hand, [Blundell and Preston \(1998\)](#) investigate a similar issue using consumption data and a simple model of consumption behavior. Their identifying assumption is the permanent income hypothesis (PIH) in its pure form, which implies that consumption responds to permanent but not to transitory shocks to income. Since consumption inequality did not increase (much) over the sample period, Blundell and Preston conclude that the increase in income inequality must have been mainly due to transitory shocks.

In this paper, we use the information in both the autocovariance structure of income and the comovement between consumption and income inequality. As documented by [Gottschalk and Moffitt](#) on the one hand and [Blundell and Preston](#) on the other, these two pieces of information seem to contradict each other.² In order to reconcile them, we need to allow for income shocks that are permanent, but are not transmitted to changes in consumption. Predictable permanent changes in income, capturing heterogeneity in life-cycle profiles, deliver this property.

The paper most closely related to ours is [Pistaferri et al. \(forthcoming\)](#), who use individual income and consumption data to estimate the extent to which households are able to insure against income shocks. [Pistaferri et al. \(forthcoming\)](#) use income data from the PSID and adopt an imputation procedure to construct a measure of total non-durable consumption for households in the PSID, given food expenditure data and a demand function for food, estimated from the CEX. One advantage of our approach is that we measure consumption and income for the same household and do not need to worry about potential weaknesses of the imputation procedure. Consistent with our estimates, [Pistaferri et al. \(forthcoming\)](#) find that consumption is insulated from most income shocks, but they interpret this result as evidence for a substantial degree of risk sharing. We show that heterogeneity can explain the same patterns in the data as partial risk sharing and argue in favor of the former interpretation.

In this respect, our paper is related to the work of [Guisen \(2005, 2007\)](#), who shows that heterogeneity in income profiles accounts for a large part of the increase in income inequality for a given cohort with age. The predictable and unpredictable shocks in Guisen's work have different statistical properties, which allows for their identification using income data only. In this paper, identification relies on the comovement of consumption with income. It is therefore reassuring that our results are broadly consistent.

This paper is organized as follows. In the next section, we describe the structure we impose on the stochastic process for income. We also set out a simple model of consumption and discuss how this model can be used to decompose income changes into predictable life-cycle shocks and permanent and transitory income risk. Section 3 describes the dataset and discusses the evolution of income and consumption inequality in the raw data. In Section 4, we discuss how we use the information in these data to estimate our model and describe the estimation procedure. Finally, in Section 5 we provide some evidence that the estimated model gives an accurate description of the joint evolution of income and consumption inequality and present our results. Section 6 concludes.

2. Model

In this section, we discuss the model that we employ to relate the evolution of income and consumption inequality to income risk. Consider a stochastic process for log income y_{it} of an individual consumer i of age a at time t , where we omit the cohort index a for simplicity. Income consists of a permanent and a transitory component and is subject to three types of shocks,

$$y_{it} = y_{it}^p + u_{it} \quad (1)$$

$$y_{it}^p = y_{it-1}^p + v_{it} + \alpha_{it} \quad (2)$$

where u_{it} is a transitory shock and v_{it} and α_{it} are permanent shocks. The shocks u_{it} and v_{it} are unpredictable to the consumer and thus represent income risk. We assume these shocks have zero mean and are uncorrelated over time and with each other. The shock α_{it} looks unpredictable to the econometrician, but is predictable to the consumer. Thus, α_{it} will contribute to inequality but not to risk. Conditional on the information set of the econometrician, we assume α_{it} has zero mean, is serially uncorrelated as well as uncorrelated with other shocks.³ The variances of the shocks are assumed to be constant across individuals but may vary over time. These time-varying variances represent transitory and permanent risk and the contribution of predictable shocks to inequality.

The decomposition of income into a permanent component that follows a random walk, and a transitory component that is serially uncorrelated, is both convenient and fairly general, and has been widely used in the literature. [Moffitt and Gottschalk \(1995\)](#) test a more general process allowing the transitory component of income to follow an ARMA process. They find that an ARMA(1,1) describes the data best, but the autocorrelation in the transitory shocks is close to 0.

² Blundell and Preston use the UK Family Expenditure Survey whereas Gottschalk and Moffitt use US data. However, [Dickens \(2000\)](#) shows that the autocovariance structure in the UK data is very similar to that documented by Gottschalk and Moffitt although transitory shocks seem to be somewhat more important in the UK. In this paper, we show that the US CEX Survey shows a very similar pattern for consumption inequality as documented by Blundell and Preston for the UK.

³ The zero mean assumption is true by construction because we remove the average life-cycle profile from our data, see Appendix A. For simplicity, we also assume that there are no aggregate shocks. In the earlier work we allow for aggregate shocks and find that these have a negligible effect on the trends in inequality ([Primiceri and Van Rens, 2004](#)).

Storesletten et al. (2004a) allow for the persistent component of income to have an autocorrelation coefficient smaller than unity. Their point estimate for the autocorrelation lies between 0.98 and unity (for annual time series) and they cannot reject the hypothesis that the persistent income shocks are permanent.

Substituting out y_{it}^p from expression (1) we get the following expression for the innovations to income:

$$y_{it} = y_{it-1} + v_{it} + \alpha_{it} + \Delta u_{it} \quad (3)$$

Income changes either because of shocks to permanent income, or because of *changes* in the transitory component. The intuition for this is simply that the effect of a transitory shock dies out in one period, so ceteris paribus a shock to transitory income at time t raises income at time t and decreases it again at time $t + 1$. Notice that α_{it} and v_{it} are clearly not separately identified using income data alone, which is why we need a model of consumption behavior and data on consumption to identify these shocks.

In its simplest form, the PIH predicts that consumption follows a random walk, and that only shocks to permanent income (i.e. expected life-time income) translate into changes in consumption. Following Blundell and Preston (1998), we use this prediction to separate permanent from transitory shocks to income. Consumption should not respond to transitory shocks since these shocks (almost) do not affect the net present value of life-time income.

The behavior of consumption under the PIH can be summarized by the following Euler equation for log consumption⁴:

$$c_{it} = c_{it-1} + v_{it} \quad (4)$$

Comparing Eq. (4) with (3) reveals the source of identification of life-cycle heterogeneity from unpredictable permanent shocks. Permanent changes in income that do not translate into changes in consumption are attributed to heterogeneity. If there are other reasons why consumption does not respond one-for-one to changes in permanent income, we will spuriously overestimate the contribution of heterogeneity to income inequality. For example, in the simple model we use here, we cannot separately identify predictable changes in income from unexpected, but insurable, income shocks. In Section 5.3, we discuss this issue in more detail.

The evolution of income and consumption inequality follows by taking a cross-sectional variance of Eqs. (3) and (4).⁵

$$\Delta var_t(y) = var_t(v) + var_t(\alpha) + \Delta var_t(u) \quad (5)$$

$$\Delta var_t(c) = var_t(v) \quad (6)$$

It is important to realize that the above expressions hold for individuals in the same cohort of consumers that are born around the same time. This reconciles the prediction put forward by Deaton and Paxson (1994) that shocks to permanent income unambiguously and irreversibly increase consumption inequality, as in Eq. (6), with the observation that aggregate inequality does not always increase in the long run.

3. Data

For our empirical analysis, we use data on US household income and consumption from the CEX (US Department of Labor, Bureau of Labor Statistics, 1999). This survey is conducted on an annual basis from 1980. Notice that although the CEX data on income are not of the best quality, the CEX is the only US dataset that has acceptable consumption as well as income data for the same individuals.⁶

3.1. The microdata

In Appendix A, we discuss the construction of the dataset and the way we control for inflation, seasonality, age effects, attrition bias and family composition. The final dataset contains 42,325 urban households with complete income and consumption data and a reference person (the person or one of the persons who owns or rents the home) who is not retired, not a student nor living in student housing and between 20 and 65 years old. This sample is representative for the full CEX sample of households aged between 20 and 65 (see the Appendix for a more extensive discussion).

From the individual level data on consumption and income, we construct a synthetic panel dataset of second moments for five 10-year cohorts. Households are assigned to a cohort based on the age of the reference person in 1980. For example, cohort 45 consists of households with a reference person between 41 and 50 years old in 1980. In 1990, the average age of this cohort was 55 years. Table 1 presents the structure of the dataset and reports cell sizes by cohort-year cells. To avoid sample sizes that are too small to get a good estimate of the second moments, we eliminate cells with average age below 25 or above 60 (these cells are shaded in the table). This implies that over time the oldest cohorts exit and younger cohorts

⁴ As in Blundell and Preston (1998), this equation can be derived in a stylized model with infinitely lived consumers with time-separable CRRA preferences over consumption, who have unconstrained access to a risk-free bond for borrowing and lending but otherwise face incomplete asset markets. To obtain the martingale property of log consumption, one needs to log-linearize the Euler equation and the life-time budget constraint.

⁵ Notice that $var(\Delta u_{it}) = var_t(u) + var_{t-1}(u)$ and $2cov(\Delta u_{it}, y_{it-1}) = -2cov(u_{it-1}, y_{it-1}) = -2var_{t-1}(u)$.

⁶ The only alternative would be the PSID, which has better income data and a longer panel dimension, but only a rough proxy for consumption (expenditures on food).

Table 1

Cell sizes by cohort and year.

Year	Cohorts (average age of reference person in 1980)								
	5	15	25	35	45	55	65	75	85
1980		21	205	179	136	126	82	23	5
1981		93	953	913	622	610	290	84	24
1982		56	260	200	143	134	61	19	4
1983		151	778	658	439	482	163	44	11
1984		201	812	731	453	440	153	43	11
1985		359	1168	1028	693	568	185	52	4
1986	1	116	241	162	118	85	37	12	3
1987	0	441	933	781	588	390	124	37	4
1988	1	422	777	637	416	298	87	22	5
1989	8	505	759	653	443	254	86	21	2
1990	32	550	811	678	421	244	61	20	2
1991	48	604	766	640	360	219	62	22	
1992	83	633	831	574	388	171	50	16	
1993	106	660	819	581	364	164	44	13	
1994	151	692	811	602	341	124	32	7	
1995	221	773	938	736	386	122	35	4	
1996	137	355	368	253	133	31	9	4	
1997	289	643	708	483	192	60	16	2	
1998	352	685	691	508	215	70	19	2	
1999	453	693	726	499	200	69	16	3	
2000	342	565	566	384	144	52	10		

Note: Shaded cells are not used in the estimation.

enter the sample and guarantees that the average age of the sample is roughly constant over time, although the evolution of age shows a 'clunky' pattern, gradually increasing each year and sharply decreasing in years when cohorts enter or exit the sample. The secondary dataset contains 75 cohort-year cells with a median cell size of 602 households.⁷

Our measure of inequality within a cohort is the cross-sectional variance of consumption or income. Other second moments we use in the estimation are the covariance between consumption and income and the autocovariance of income. For all moments, we use consumption and income in logs. In the remainder of this section, we discuss two data problems that may affect these moments: measurement error and the timing of the questions on income in the CEX. Section 3.2 presents some descriptive evidence from the raw data on the evolution of income and consumption inequality over the sample period.

Both income and consumption are measured with error. Our estimation results, however, are likely not to be affected by this problem. Assuming that the measurement error is uncorrelated with the true levels of income and consumption, then it adds an additive term to the variance of income and consumption. If we further assume that the cross-sectional variance of the measurement error is constant over time, then this additive bias term will drop out when we take first differences for a cohort over time, so the evolution of inequality is unaffected, even if the level of inequality is biased. Throughout the paper, we refrain from interpreting the levels of inequality and only use changes in inequality for our estimation.⁸

A more serious data problem is the timing of the questions on income and consumption in the CEX (Gervais and Klein, 2005). Questions on consumption are asked in four quarterly interviews and refer to the quarter preceding the interview. Therefore, the four observations for consumption can be added up to obtain one observation for annual consumption in the year preceding the last interview. Questions about income are asked only in the first and last quarter and refer to income in the year preceding the interview. Therefore, annual income from the last interview corresponds to the same period as annual consumption and neither consumption nor income inequality are affected by this timing problem. However, annual income from the first interview does not refer to the preceding year, but overlaps income from the last interview by one

⁷ In some years, the sample sizes are substantially below the median cell size, see Table 1. In 1986 and 1996, there are no expenditure data for interviews held in January because of changes in the sample design. In addition, in 1986 the BLS changed the numbering of the household identifiers so that households cannot be matched across the 1986 and 1985 files, leading to a particularly large drop in the number of observations in that year. Sample sizes in 1982 are lower partly because the survey sample was smaller in the earlier years and partly because in 1982 and 1983 the interview family files do not contain the summary expenditure variables that we use to construct our measure of consumption, so that we need to aggregate these data from the detailed expenditure files, which were not available for all households.

⁸ Apart from the measurement error problem, the levels of inequality are also very sensitive to outliers. Trimming the top and bottom 0.1% of the income distribution in each year, reduces the sample average of the variance of log income from 0.74 to 0.68. Trimming the top and bottom 1% reduces inequality further to as little as 0.55. Because the level of trimming is arbitrary, we do not trim outliers but rather refrain from interpreting the levels of inequality.

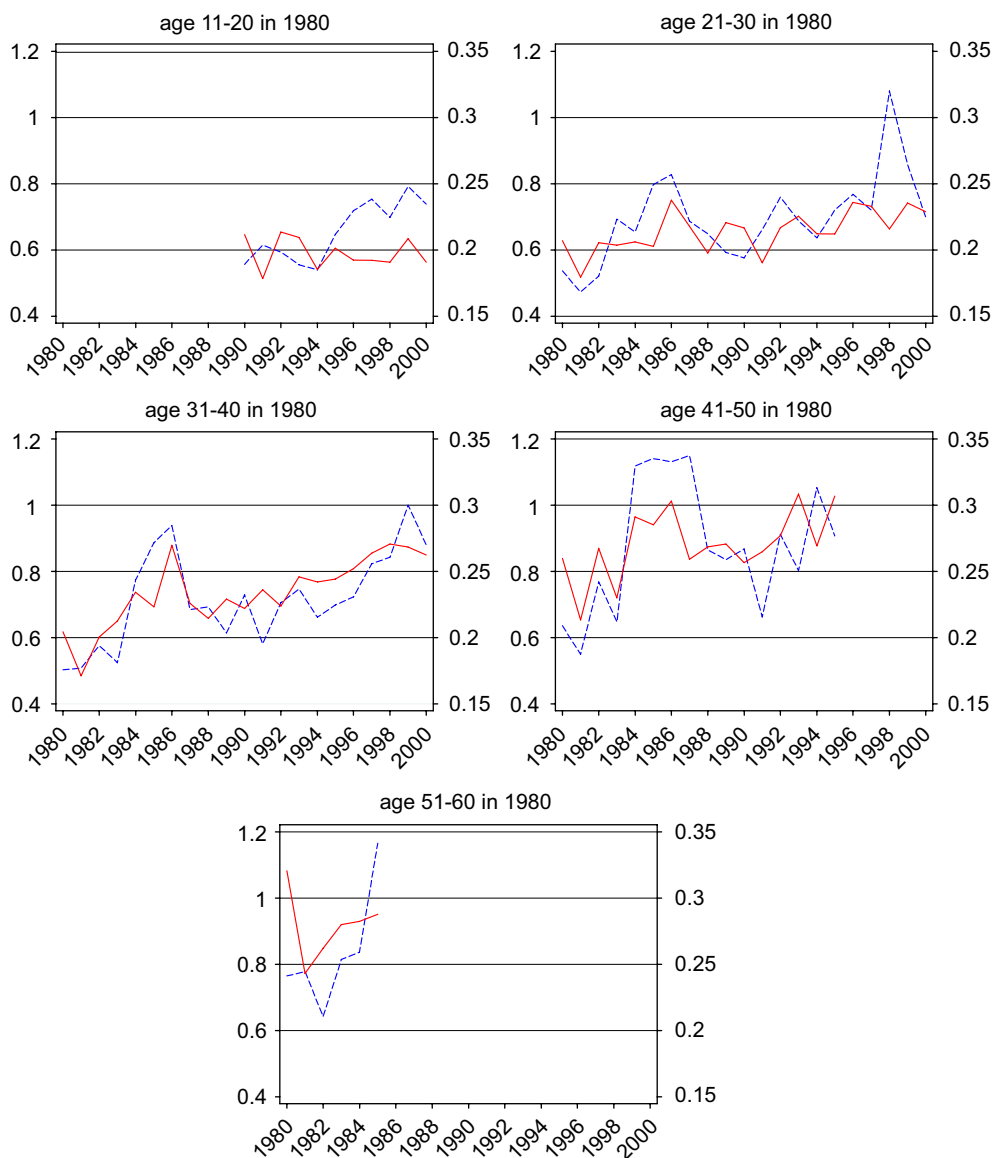


Fig. 1. Income and consumption inequality by cohort. The blue dashed line is income inequality (the variance of log income), plotted on the left scale. The red solid line is consumption inequality, plotted on the right axis. The five different graphs represent five different cohorts, identified by their age in 1980. For the sample sizes used to calculate inequality in each cohort, see Table 1.

quarter. This biases the estimated covariance of income growth with past levels of income, one of the moment conditions we use to estimate the model (see Section 4.1). We deal with this problem by assuming that income changes only at the beginning of the year, so that observed income in the previous year \tilde{y}_{it-1} is a linear combination of the true income in the previous year and income in this year, $\tilde{y}_{it-1} = \frac{3}{4}y_{t-1} + \frac{1}{4}y_t$, and correct the moment condition accordingly.⁹

3.2. Income and consumption inequality

Fig. 1 shows consumption and income inequality (the variance of log real consumption and income) for the five cohorts over the sample period. The consumption graphs are comparable to Deaton and Paxson (1994, Fig. 2) although our sample period is twice as long. We would expect to see two stylized facts in these data. First, as shown by Deaton and Paxson (1994), inequality should rise within a cohort with age (so therefore over time) both for income and for consumption, with

⁹ Under this assumption, $\frac{3}{4}\text{cov}(y_t - \tilde{y}_{t-1}, \tilde{y}_{t-1}) - \frac{1}{4}\Delta\text{var}(y_t)$ is a consistent estimator for $\text{cov}(\Delta y_t, y_{t-1})$. We also estimated the model under the 'naive' assumption that $\tilde{y}_{t-1} = y_{t-1}$ and found that this makes very little difference in the results.




Fig. 2. Age effects in inequality. The blue solid line plots the coefficient estimates of a regression of income and consumption inequality on age dummies, controlling for cohort effects (five cohort dummies). For the raw data by cohort, see Fig. 1. The red dotted lines represent the two standard error bands for these estimates. The green dashed line controls for time effects (seven time dummies for 3 year periods: 1980–1982, 1983–1985, etc.), rather than cohort effects.

the increase being less pronounced for consumption because of smoothing. Second, there should be an increase in inequality common to all cohorts in the 1980s, which then flattens out in the 1990s. Both ‘facts’ are not easy to see, partly because noise clouds the picture, and partly because both age and time effects are interacting in the same graphs.

In Fig. 2, we plot the coefficient estimates of a regression of income and consumption inequality on age dummies, as in Deaton and Paxson (1994, Fig. 4). The solid line shows the evolution of inequality with age controlling for cohort effects, the dashed line controlling for time effects. We document a significant and approximately linear increase in within-cohort consumption inequality with age, although the effect is substantially smaller than in Deaton and Paxson. As pointed out by Slesnick and Ulker (2004) and Heathcote et al. (2005), the large increase in inequality over the life-cycle that Deaton and Paxson find is partly due to the fact that their sample period covers only the 1980s.

Fig. 3 focuses on the time effects. The upper solid lines show the evolution of inequality over time for the average cohort in our sample. This line is constructed by regressing income and consumption inequality on year dummies and plotting the coefficient estimates. The lower solid line controls for a linear trend in inequality due to the fact that the cohort ages over time (the straight dotted line plots this estimated age effect). Controlling for the age effect, the evolution of average within-cohort inequality should be similar to that of aggregate inequality, the cross-sectional variance of log income or consumption for the whole sample at a given point in time. This is indeed the case as can be seen from the graph by comparing the lower solid line to the dashed line, which plots aggregate inequality.¹⁰ In the remainder of this paper, we will

¹⁰ The (small) differences come from changes in the average age of the sample over time. Because inequality increases with age, changes in the average age of the population will affect aggregate inequality. For example, the ‘baby-boomers’ entered the labor market around 1973 when they were 23 years old. Ceteris paribus we would expect consumption and income inequality to decrease around this time. Inequality would also decrease around the year 2045, when the baby-boomers retire, and would gradually increase in between due to the aging of the labor force.




Fig. 3. Aggregate and average within-cohort inequality. The blue solid lines show the evolution of inequality over time for the average cohort in our sample. These lines are constructed by regressing income and consumption inequality on year dummies and plotting the coefficient estimates. The lower line controls for a linear trend in age and the straight green dotted line plots the estimated age effect. The red dashed line is aggregate inequality, the cross-sectional variance of log income or consumption for the whole sample in each year.

present average within-cohort inequality and interpret it as aggregate inequality plus an approximately linear increase in inequality due to the age effect.

Income inequality rose sharply in the early 1980s and remained high during the second half of the 1980s and all of the 1990s. Consistent with other studies, we also find a temporary peak in inequality in the mid 1980s, which seems to be specific to the CEX data (Attanasio et al., 2004). We argue that this peak is partly driven by sampling error, which even with relatively large cell sizes may lead to large swings in the variance since the variance is sensitive to outliers. The peak becomes less pronounced if we trim the income distribution for outliers, see Footnote 8, or if we use a robust estimator for the variance, see the discussion of Fig. 4 in Section 5. Because any procedure to deal with potential outliers is, to some extent, arbitrary, we use the raw data series and deal with the sampling error in the estimation procedure, see Section 4.3.

Consumption inequality did not increase much over the sample period. This is also consistent with what other studies have found (Krueger and Perri, 2006).

4. Empirical approach

The raw data are very noisy due to the relatively small number of households in a cohort-year cell. In this section, we discuss our estimation procedure, which is designed to extract slow moving trends from these noisy data. First, we present a set of moment conditions that represent the information available in the data. Then, we discuss a likelihood based, Bayesian procedure that treats the time-varying variances of the idiosyncratic shocks as unobservable components. Because this procedure imposes smoothness on movements in the time-varying variances, it performs well in distinguishing low frequency trends from noise. Moreover, the Gibbs sampler used to evaluate the likelihood has more robust convergence properties than the high dimensional minimization routine needed to estimate the model by minimum distance methods.

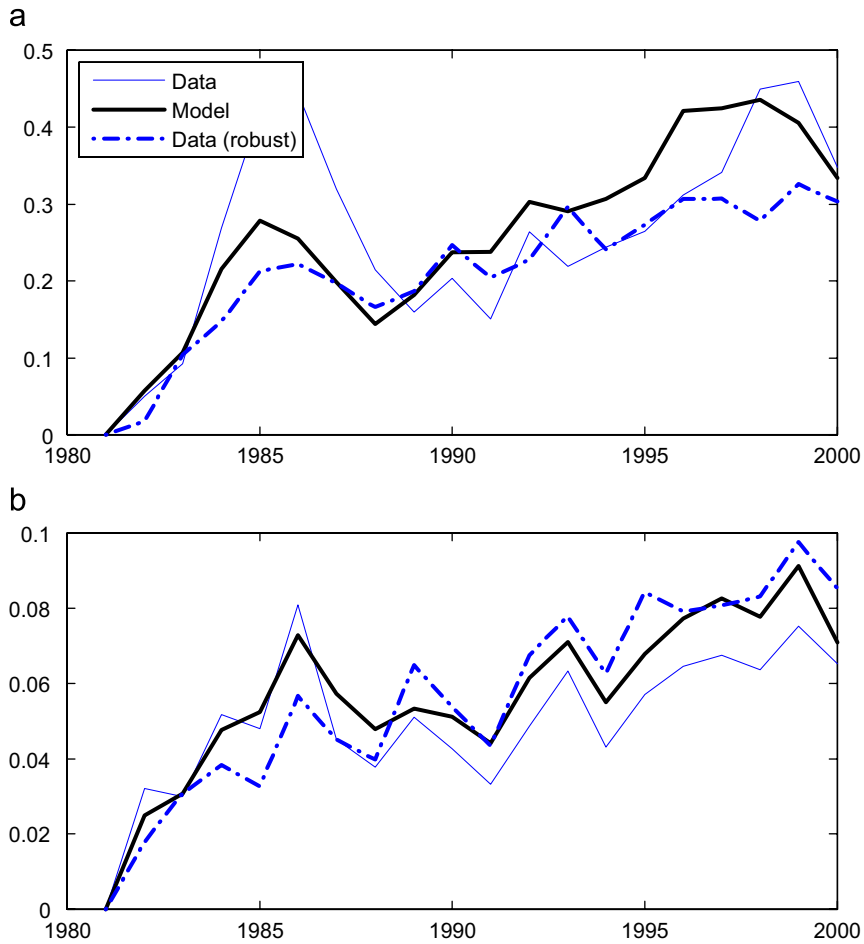


Fig. 4. Income and consumption inequality: data and model predicted values. (a) Actual (robust and non robust) and model predicted income inequality; (b) actual (robust and non robust) and model predicted consumption inequality.

4.1. Moment conditions

We use expressions (3) and (4) to calculate moments that we can measure from the data. Following [Blundell and Preston \(1998\)](#), first of all we use changes in the variances of log income and log consumption, which represent the evolution of income and consumption inequality, in which we are primarily interested. These moment conditions are given in Eqs. (5) and (6).

But there is more information in the data than just those two moment conditions. First of all, we also use the change in the covariance between log income and log consumption. Calculating the evolution of the covariance from Eqs. (3) and (4), we get

$$\Delta cov_t(y, c) = var_t(v) = \Delta var_t(c) \quad (7)$$

The evolution of the covariance of income and consumption contains the same information as the evolution of the variance of consumption under the model. Using both moment conditions should improve the efficiency of our estimates. The overidentifying restriction also allows us to test the model specification.

A fourth moment condition is found in the autocovariance of income. Using the information in the time series properties of income is attractive, because it corresponds to the methodology in [Gottschalk and Moffitt \(1994\)](#) and [Moffitt and Gottschalk \(1995, 2002\)](#). Because $cov_t(y, y_{t-1}) = cov_t(\Delta y, y_{t-1}) + var_{t-1}(y)$ and we are already using the information contained in the variance of income, we use $cov_t(\Delta y, y_{t-1})$. From (3), we get

$$cov_t(\Delta y, y_{t-1}) = -var_{t-1}(u) \quad (8)$$

Moment conditions (5)–(8) contain all information in the second moments of the joint evolution of income and consumption that we can retrieve from the data.

4.2. Identification

Consider Eqs. (5)–(8), which hold for every cohort and every time period and therefore represent $4JT$ moment conditions, where J is the number of cohorts and T the number of time periods. These moment conditions need to identify $3T + 1$ parameters: $var_t(v)$ and $var_t(\alpha)$ for $t = 1$ to T and $var_t(u)$ for $t = 0$ to T . The autocovariance of income (8) provides an estimate for $var_t(u)$ for $t = 0$ to $T - 1$. Similarly, the moment conditions for the variance of consumption (6) or the covariance (7) pin down $var_t(v)$ for $t = 1$ to T . Finally, given $var_t(v)$ and $\Delta var_t(u)$, the variance of income (5) can be used to retrieve $var_t(\alpha)$ for $t = 1$ to $T - 1$. The variance of the transitory shocks and therefore also the variance of the heterogeneity in life-cycle profiles in the last period, $var_T(u)$ and $var_T(\alpha)$, are identified from a smoothness assumption on the time variation in the variances of shocks to income.¹¹

Measurement error does not affect the moment conditions for $\Delta var_t(y)$, $\Delta var_t(c)$ and $\Delta cov_t(y, c)$ as we argued in Section 3.1. In the moment condition for $cov_t(\Delta y, y_{t-1})$, the variance of classical measurement error in income enters as an additive constant. This constant is not separately identified from the level of the variance of the transitory shocks, $var_t(u)$. However, since the variance of transitory shocks enters only as a first difference, its level is not important for the evolution of income inequality.

Finally, we note that the assumption that u_{it} , v_{it} and α_{it} are uncorrelated with past values of income and consumption in the cross-section is crucial for the identification strategy because it allows us to use the change in the variances and covariance, rather than the variances and covariance of the changes in income and consumption. This assumption is not an implication of the PIH (which holds for an individual consumer) but follows from assuming that lagged aggregate consumption is in each individual consumer's information set (Chamberlain, 1984; Deaton and Paxson, 1994). As shown by Blundell and Preston (1998), testing the overidentifying restriction that the covariance between income and consumption contains the same information as the variance of consumption can be interpreted as a test for this assumption as well as for the specification of the income process more generally (p. 615). A likelihood ratio test of this restriction against an unrestricted version of the model, in which the covariance between income and consumption is left completely unconstrained, gives a $\chi^2(19)$ statistic of 21.1, so that we cannot reject the null hypothesis that the restriction is satisfied in the data (p -value 0.33). This conclusion is confirmed by a Hausman test that our estimates are the same whether or not we use the moment condition for the covariance (p -value 0.99).

4.3. Estimation

To estimate the model we take a Bayesian, likelihood based approach, treating the time-varying variances $var_t(v)$, $var_t(\alpha)$ and $var_t(u)$ as unobservable states. Since we need to specify a law of motion for the time-varying variances, we assume that these variances follow independent random walk processes. Of course, variances cannot be negative and, at first sight, the random walk assumption may seem inadequate. However, because the time dimension of the sample is short, the random walk can be thought as a (good) first order approximation of a more complicated and theoretically justifiable process for the two variances.¹² The assumption has the advantage that it imposes smoothness on the movements in the variances. Since we want to capture low frequency time variation, the smoothness helps to identify signal from noise.

A Bayesian approach is natural in estimating unobservable components. Even more so in a panel context, where the distinction between parameters and shocks is less clear than in other situations. Moreover, because we use flat and uninformative priors, the Bayesian procedure has a likelihood interpretation. With flat priors, the posterior modes of the parameters correspond exactly to the maximum likelihood estimates. Finally, and particularly important in this case, the Bayesian approach allows to split up the high dimensional problem into a series of simpler and lower dimensional ones. This has the advantage that the numerical procedure is more robust and that it is easier to calculate standard errors that are correct for finite sample inference instead of relying on asymptotic theory. Appendix B describes the Markov Chain Monte Carlo (MCMC) algorithm for the numerical evaluation of the posterior of the parameters of interest.

5. Results

Fig. 4 plots the actual data (thin solid line) and the fitted values of our model (thick solid line) for the evolution of inequality over the sample period. The upper panel displays income inequality, the lower panel consumption inequality. The model captures the overall trend in both income and consumption inequality very well, as well as some of the high frequency fluctuations in the data. The random walk assumption on the law of motions for the time-varying variances imposes some smoothness on these estimates. As a consequence, the large peak in income inequality from 1984 to 1988 for instance (which is not present in other datasets), is not captured.

¹¹ See Section 4.3 for details. The weaker identification scheme affects *only* the estimates in the last period of our sample. Moreover, a model with time-invariant heterogeneity, $var_t(\alpha) = var(\alpha)$ for all t , which is identified without the smoothness assumption, gives virtually identical results, see Section 5.4.

¹² The estimation algorithm allows to restrict the variances to be positive at all points in time. However, because the point estimates turn out to be positive, the normality assumption does not affect the results. We also reestimated our model assuming the variances follow autoregressive processes and found very similar results.

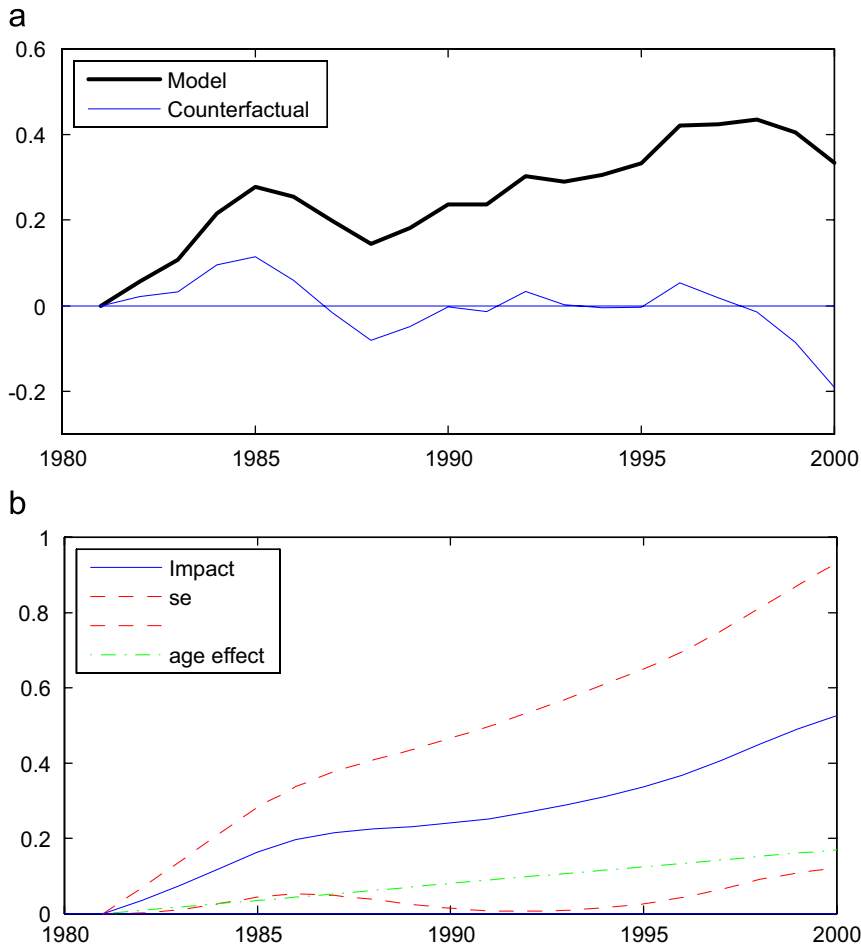


Fig. 5. Contribution of predictable permanent shocks to income inequality. (a) Income inequality without permanent predictable inequality. (b) Impact of permanent predictable inequality.

We argue that the deviation of the actual data from the fitted values is largely attributable to measurement and sampling error. To support this argument, the third line in the graphs (dash-dotted) presents the raw data again, now using a robust estimator for inequality.¹³ As is clear from the graph, the model predicted values are very close to the robust series. We did not use these series in the estimation procedure so that the fit is quite remarkable. We conclude that the estimation procedure manages well to distinguish noise from signal and the fitted values provide a good description of the joint evolution of income and consumption inequality.

5.1. Sources of inequality

In order to assess the contribution of the different shocks to changes in inequality, we ask the question how income inequality would have evolved without each shock. Figs. 5–7 present the counterfactual evolution of income inequality if predictable changes α_{it} , unpredictable permanent shocks v_{it} or transitory shocks u_{it} would have been zero for all individuals in every period. The upper panels of these graphs show the predicted values for income inequality for the counterfactual exercise (thin solid lines) as well as for the full model (thick solid lines). The lower panels plot the difference between the two lines, which represents the contribution of each type of shock, with one standard error bands. In order to be able to compare those graphs to the evolution of aggregate inequality, we have also plotted a straight line representing the average increase of within-cohort income inequality with age, which we refer to as the age effect.

It is clear from Fig. 5 that predictable permanent shocks explain the vast majority of changes in income inequality. Without these predictable shocks, income inequality would actually have gone down over the sample period. This result is consistent with Guvenen (2005), who finds that heterogeneous life-cycle changes make up 65–80% of the life-time increase

¹³ Assuming the logs of income and consumption are normally distributed in the cross-section, the robust estimator for the standard deviation equals the median absolute deviation from the median divided by 0.6745 (Huber, 1981).

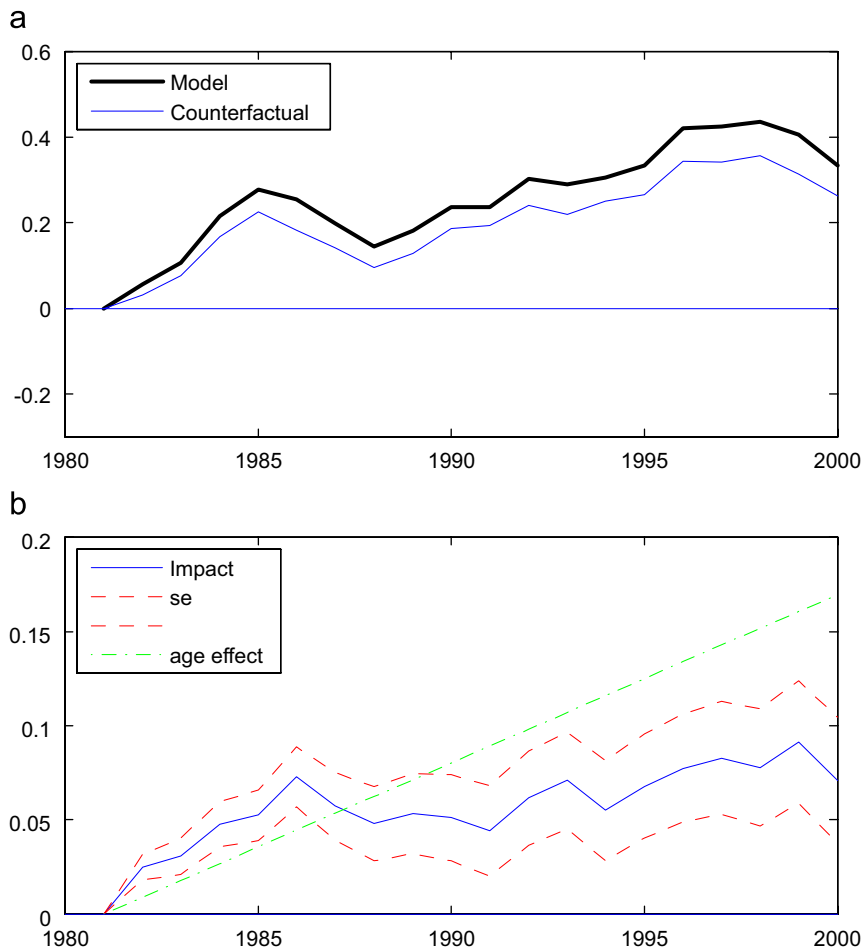


Fig. 6. Contribution of unpredictable permanent shocks to income inequality. (a) Income inequality without permanent unpredictable inequality. (b) Impact of permanent unpredictable inequality.

in income inequality within a cohort. We show that, in addition, changes in the amount of this heterogeneity can account for the increase in aggregate income inequality over the period 1980–2000.

The variance of unpredictable permanent shocks went up as well, see Fig. 6, so that part of the increase in inequality in the early 1980s can be attributed to increased permanent income risk. However, this contribution is a factor 3 smaller than the increase in inequality due to predictable shocks. Moreover, from the second half of the 1980s onwards, permanent risk seems to have gone down again and, at the end of the sample, the increase is smaller than the age effect, so that aggregate inequality would have gone down if permanent risk were the only source of inequality over this period.

As shown in Fig. 7, transitory inequality also increased in the early 1980s. But this increase is very small, much smaller than the increase in transitory inequality found by Gottschalk and Moffitt (1994); Moffitt and Gottschalk (2002) and more in line with the results of Pistaferri et al. (forthcoming). On the other hand, the evolution of transitory inequality is consistent with Moffitt and Gottschalk (2002). Like them, we find a reversal of the increase in transitory inequality, with inequality due to transitory shocks decreasing in the late 1980s and throughout the 1990s. If transitory shocks were the only source of inequality, by 2000 income inequality would have decreased substantially compared to 1980.

So was the increase in income inequality in the 1980s due to permanent or transitory shocks? Our estimates clearly point towards the importance of permanent sources of inequality. However, since we estimate most of the permanent shocks to be predictable to consumers, we do not find evidence for an increase in permanent income risk (the variance of unexpected permanent shocks) over this period. Based on our estimates, the evolution of risk shows a markedly different picture than the evolution of inequality. Whereas inequality increased in the 1980s and remained high, the increase in risk seems to have been temporary. By 2000, permanent income risk was as high as it was in 1980 and transitory risk had substantially decreased.

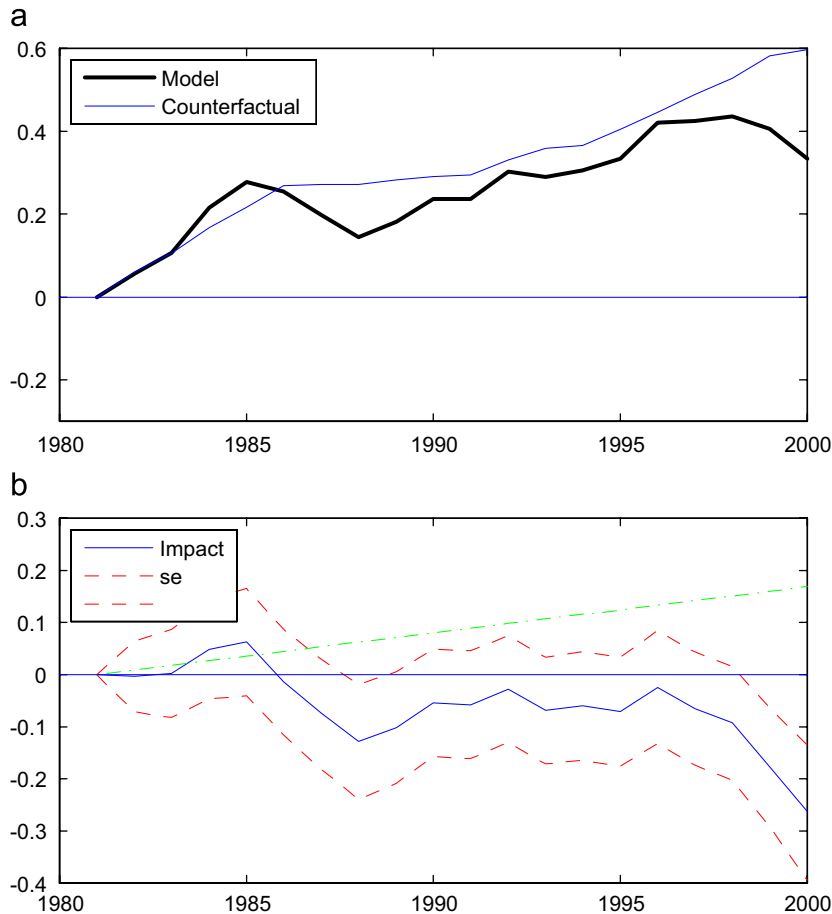


Fig. 7. Contribution of transitory shocks to income inequality. (a) Income inequality without transitory inequality. (b) Impact of transitory inequality.

5.2. The joint evolution of income and consumption inequality

We have shown that our model manages to capture the joint evolution of consumption and income inequality well and that a large part of the evolution of income inequality is explained by changes in income that look unpredictable to the econometrician but are predictable for consumers. The reason is simple. The time series properties of income (its autocovariance) suggest that most income changes are permanent. However, the evolution of consumption inequality shows that consumption nevertheless did not respond much to these changes in income. Therefore, in the context of our simple model of consumption behavior, the vast majority of permanent shocks are estimated to be predictable.

As pointed out in the Introduction, this simple insight reconciles two seemingly contradictory branches of the literature (Carroll, 1992; Gottschalk and Moffitt, 1994; Moffitt and Gottschalk, 1995, 2002; Blundell and Preston, 1998). In order to understand what drives this result, we re-estimated our model several times, imposing different sets of restrictions in order to reproduce either Gottschalk and Moffitt's or Blundell and Preston's results. In Fig. 8 we plot fitted values for income and consumption inequality for these alternative models and Fig. 9 presents the contribution of permanent shocks for each. The thin and thick solid lines in Fig. 8 are the same as in Fig. 4 and represent the evolution of inequality in the data and in the baseline model, respectively. In Fig. 9, the thick solid line presents our estimate for the total contribution of predictable and unpredictable permanent shocks on income inequality.

First, consider the dashed line in Fig. 8, which represents the estimates of a model in the spirit of Moffitt and Gottschalk (1995). To obtain these estimates, we simplified the income process by no longer distinguishing between predictable and unpredictable permanent shocks. Then, we estimated this model on a subset of the moment conditions we use in the baseline, removing all information about consumption inequality and using only the moment conditions for income inequality (5) and the autocovariance of income (8). To obtain fitted values for consumption inequality, we assume all permanent shocks are unexpected to consumers, the assumption that Gottschalk and Moffitt make implicitly when interpreting their results. Unsurprisingly, this model fits the evolution of income inequality well but completely fails to explain the evolution of consumption inequality. The reason is that the autocovariance of income suggests that the increase

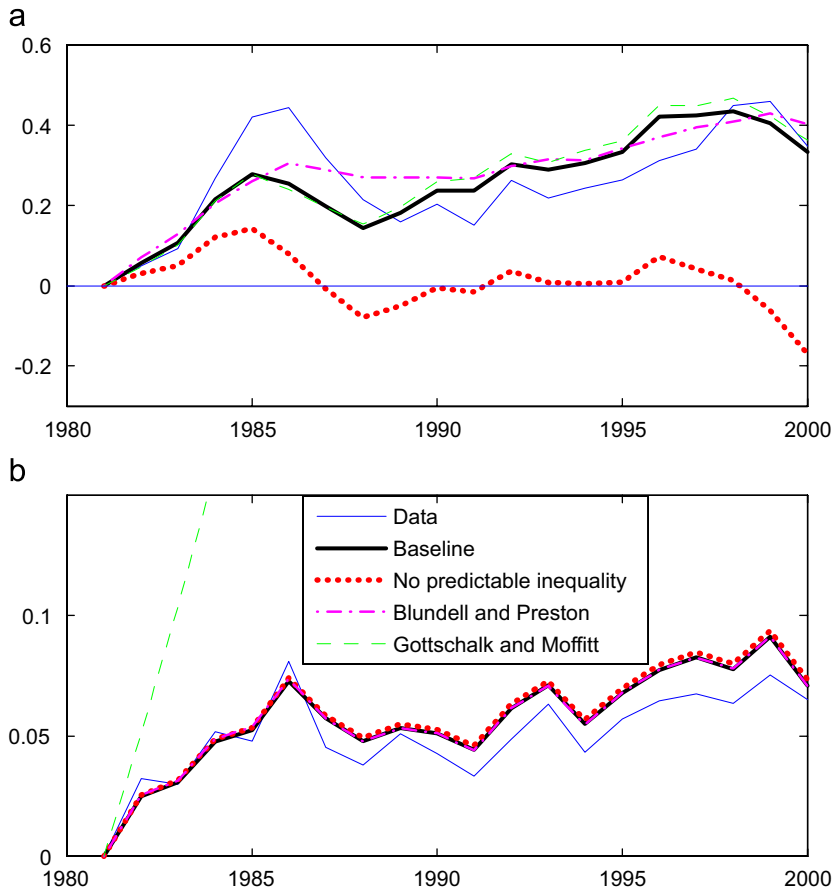


Fig. 8. Model predicted evolution of inequality for different models. (a) Income inequality: actual and predicted by different models. (b) Consumption inequality: actual and predicted by different models.




Fig. 9. Contribution of permanent shocks to inequality for different models.

in income inequality is due to permanent shocks (Fig. 9). Therefore, consumption inequality should have increased substantially over the whole sample period, with the strongest increase in the early 1980s.

Next, consider the dash-dot line, which replicates the estimates in Blundell and Preston (1998). To estimate this model, we again use the simplified income process and consumption model, but now we estimate it using the consumption moment conditions (6) and (7) in addition to the moment condition for income inequality (5), but not condition (8) for the

autocovariance of income. By assumption, this model captures the evolution of both income and consumption inequality well. Blundell and Preston consider all permanent shocks to be unpredictable and identify these shocks as shocks to which consumption inequality responds. Consequently, because consumption inequality did not increase much over the sample period, they find a very small contribution of permanent shocks to income inequality (see Fig. 9).

Then we re-estimate the simplified model without predictable shocks, but now using all moment conditions. These estimates are presented as the dotted line in Figs 8 and 9. Now, there is a conflict between the information in the consumption data, represented in moment conditions (6) and (7), and the information in moment condition (8) about the time series properties of income. As a result, the estimate for the contribution of permanent shocks as well as the fitted values for consumption inequality are very close to those of Blundell and Preston, but the model can no longer match the evolution of income inequality as the estimation procedure tries to find a ‘compromise’ between conflicting sets of moment conditions.

Finally, in the full model we allow for predictable permanent shocks. These shocks are permanent insofar as the autocovariance of income is concerned, but they are also ‘transitory’ in the definition of Blundell and Preston, in the sense that consumption inequality does not increase because of these shocks. With this extension, we match the joint evolution of income and consumption inequality as well as Blundell and Preston do, but we find a contribution of permanent shocks (predictable plus unpredictable) to income inequality that is close to Gottschalk and Moffitt’s estimates.

5.3. Risk sharing

What alternative explanations are consistent with the evolution of income as well as consumption inequality? Several authors have proposed that consumers may be able to partially insure their consumption against permanent income shocks (Storesletten et al., 2004b; Primiceri and Van Rens, 2004; Pistaferri et al., forthcoming). In the context of the model in this paper, the heterogeneity and partial insurance interpretations are observationally equivalent. Our consumption model is an incomplete markets model and insurance markets are non-existing. By investing in a risk-free bond, consumers can save and borrow freely, but they cannot pool risks with other consumers so that they cannot insure their consumption path against unexpected permanent shocks. If in reality insurance markets do exist, it is possible that our estimate for the variance of α_{it} includes not only predictable shocks but also unpredictable but insurable permanent shocks. Although this distinction is not the main focus of this paper, in this section we argue that heterogeneity provides an explanation for the data that is at least as plausible as partial risk sharing, which has received far more attention in the literature.

Storesletten et al. (2004b) evaluate how consumption and income inequality change with age (over the life cycle). They find that consumption inequality predicted by their model is about 20% higher than in the data and explain this discrepancy by implicit risk sharing through the social security system. Pistaferri et al. (forthcoming) evaluate the evolution of inequality over time in a framework similar to the one in this paper. They use food consumption and household characteristics to impute total non-durable consumption in the PSID data and estimate the fraction of permanent income shocks that are not insured. They find consumption growth responds by a factor 0.64 to permanent shocks, which implies that the fraction of the variance of permanent income shocks that are insured based on their estimates is 0.6, roughly similar to our estimates.

How can we distinguish the risk sharing explanation from the heterogeneous income profiles explanation? If consumption does not respond to income shocks because of risk sharing, we would expect part of that risk sharing to happen through the government, through taxes and transfers, and part through markets for financial assets. We test this prediction of the risk sharing hypothesis by re-estimating the model for different measures of income. In each case, we report the results as the fraction of permanent shocks that are predictable or insured in three broad time periods: the early 1980s, all of the 1980s and the whole sample period (1980s and 1990s).

These estimates are reported in Table 2. In the baseline estimates, income is defined as disposable income after taxes and transfers. Rows 2, 3 and 4 of the table present estimates when income is measured as gross income before taxes (but including financial income and transfers); gross income before taxes excluding income from financial markets; and earned income (before taxes and transfers and excluding all sources of income other than wage and salary payments). The estimates in rows 2 and 3 are indistinguishable from the baseline. Neither the tax system nor financial markets seem to contribute to risk sharing. Transfers seem to provide some insurance, with the fraction of shocks to which consumption is insured going up as we would expect. However, the difference in the estimates is very small and not significant.

Thus, the evidence is not inconsistent with the risk sharing hypothesis. However, we find weak quantitative support for its basic predictions. This suggests that it is unlikely that risk sharing constitutes the whole story and thus that heterogeneity plays an important role. Moreover, as we argued in the introduction, the interpretation that most inequality is driven by heterogeneity is consistent with a number of recent papers decomposing inequality in heterogeneity and uncertainty using schooling choices (Cunha et al., 2005; Cunha and Heckman, 2006; Huggett et al., 2006).

5.4. Robustness

In this section, we explore the robustness of our results to a number of modifications to the dataset and model. The estimates are summarized in Table 3. First, we evaluate to what extent the results are sensitive to the choices we made in

Table 2
Contribution of predictable permanent shocks for different measures of income.

	Share of permanent inequality that is predictable		
	1980–1985	1980–1990	1980–2000
(1) Baseline	0.76 (0.17)	0.80 (0.15)	0.83 (0.12)
(2) Income before taxes	0.76 (0.18)	0.80 (0.16)	0.83 (0.12)
(3) Income (before tax), excl financial income	0.76 (0.18)	0.80 (0.16)	0.84 (0.12)
(4) Earned income (before tax and transfers)	0.81 (0.21)	0.84 (0.16)	0.89 (0.11)

Note: Standard errors in parentheses.

Table 3
Robustness checks.

	Share of permanent inequality that is predictable		
	1980–1985	1980–1990	1980–2000
(1) Baseline	0.76 (0.17)	0.80 (0.15)	0.83 (0.12)
(2) Including topcoded incomes	0.74 (0.17)	0.79 (0.15)	0.83 (0.12)
(3) Different equivalence scale	0.75 (0.20)	0.79 (0.20)	0.79 (0.17)
(4) CPI: item specific	0.74 (0.17)	0.78 (0.15)	0.81 (0.12)
(5) CPI: total expenditures	0.76 (0.17)	0.80 (0.15)	0.83 (0.11)
(6) Consumption: food only	0.66 (0.23)	0.73 (0.21)	0.79 (0.15)
(7) Consumption: total expenditures	0.62 (0.22)	0.65 (0.21)	0.74 (0.16)
(8) Sampling error: robust estimators	0.70 (0.19)	0.71 (0.17)	0.76 (0.14)
(9) Credit constraints: $\lambda = 0.031(0.023)$	0.80 (0.17)	0.83 (0.14)	0.85 (0.11)
(10) Interest rate: $r = 5\%$	0.76 (0.16)	0.80 (0.15)	0.83 (0.11)
(11) No time variation predictable shocks	0.67 (0.20)	0.80 (0.22)	0.82 (0.22)

Note: Standard errors in parentheses.

constructing our secondary dataset of variances and covariances from the microdata (see Appendix A for details). For the estimates in row 2, we do not exclude households for which income is topcoded from the dataset; in row 3, we try a different equivalence scale to convert household consumption into per capita consumption equivalents¹⁴; and in rows 4 and 5 we try different deflators to convert nominal consumption into real terms.¹⁵ In all cases, the results are very close to and insignificantly different from the baseline.

Second, we explore how the results are affected by our choice about what kind of expenditures to include in non-durable consumption. Rows 6 and 7 present estimates if we use only expenditures on food and beverages or all expenditures (including durables), respectively. It is quite remarkable that the results are virtually unaltered even for these large deviations from the baseline.

Next, in order to determine whether sampling error might be affecting the results, we use a dataset in which we use a robust estimator for the variances of consumption and income. Even though the data are quite noisy and the raw data series for the robust estimators are rather different (see Fig. 4), the estimation method performs well at extracting signal from noise and the estimates in row 8 are very close to the baseline.

Finally, we explore sensitivity to modifications of the model. Our baseline model does not capture excess sensitivity to (transitory) changes in income. As a rough control for excess sensitivity, we assume that a fraction λ of consumers is 'hand-to-mouth' and consumes all of current income. This parameter λ can be interpreted as the fraction of consumers that are credit constrained or as a proxy for precautionary savings.¹⁶ The modified model equations and moment conditions are given in Appendix C and row 9 presents the estimates for this model. Because our estimate for λ is very low (3.1% with a

¹⁴ We divide income by the number of people in the household and consumption by the number of adults plus 0.4 times the number of children as in Parker and Preston (2005), instead of regressing consumption and income on the number of adults and the number of children and taking residuals.

¹⁵ In the baseline, we use the CPI for non-durables for consumption and the CPI for total expenditures for income. In row 4, we use item specific CPI indices for the different categories of expenditures that constitute non-durable consumption. In row 5, we use the CPI for total expenditures for consumption as well as for income.

¹⁶ Precautionary saving is closely related to liquidity constraints, both theoretically and empirically (Carroll, 2001). Gourinchas and Parker (2001) non-parametrically estimate the consumption policy rule and find that consumption does not respond to cash-on-hand, for consumers with liquid wealth above a certain level \bar{A} which would be in line with the PIH. If wealth is below \bar{A} , the marginal propensity to consume out of extra cash-on-hand is close to 1.

standard error of 2.3%), allowing for credit constraints does not change the results much.¹⁷ In row 10 we present estimates if we assume a non-zero interest rate, so that the marginal propensity to consume out of transitory shocks is small but not zero and in row 11, we restrict the variance of predictable life-cycle shocks to be constant over time. Again, these changes affect the results very little.

6. Conclusions

In this paper, we used repeated cross-section data on income and consumption from the CEX to evaluate the nature of the increase in income inequality in the US over the last two decades. The stochastic process for income that we assume includes predictable life-cycle changes and unexpected permanent and transitory shocks. We estimate the contribution of each of these three shocks to total inequality. The model fits the joint evolution of income and consumption inequality well. Almost all of the increase in income inequality was due to predictable life-cycle shocks. The variances of both permanent and transitory unexpected shocks also increased in the early 1980s, but these increases were small and got reversed in the 1990s.

Our set of moment conditions summarizes all information available from the CEX data. In particular, we use information both on the autocovariance structure of income and on the comovement of income and consumption. By allowing for predictable changes in permanent income (heterogeneity), we reconcile the seemingly contradictory findings that the increase in income inequality was due mainly to permanent shocks (Moffitt and Gottschalk, 2002), yet consumption inequality did not increase much over the same period (Blundell and Preston, 1998; Krueger and Perri, 2006).

Appendix A. Data description

The CEX is a rolling panel. Each month a new group of about 500 new households enters the survey (annual sample size is about 5,870 households in the later years). A household or 'consumer unit' is a group of individuals living together as a family. These households are interviewed each quarter, for five quarters in a row. The first meeting is an introductory interview where respondents are asked about family characteristics and are given information about how to gather their expenditure information. In the second through fifth interview households report expenditures over the previous quarter. Expenditures are coded by the Bureau of Labor Statistics, assigned a Universal Classification Category (UCC) number, and aggregated into several broader categories. The BLS gathers these data primarily in order to calculate the Consumer Price Index. Questions about income are asked in the second and fifth interview only, and refer to the preceding 12 months, see Section 3.1 in the main text.

As our measure of consumption we use non-durable consumption, consisting of expenditures on food and beverages, utilities, gas and motor oil, public transportation, reading materials, tobacco products, personal care and apparel. For income we use family income after tax. Nominal income and consumption are converted to real values using the CPI-U indices (all urban consumers) for total expenditure and non-durable expenditures, respectively.

We limit our sample to households with a reference person (the person or one of the persons who owns or rents the home) between 20 and 65 years old; that live in urban areas (data on non-urban households is not available in 1982 and 1983); report non-zero expenditures on food; for which complete data on income are available and not topcoded (drops 18% of the sample between 20 and 65 years old) and that are not retired (4%), student (1.1%) or living in student housing (0.5%). Because households need to be matched across quarterly surveys in order to obtain a measure of annual consumption, we run a series of checks to identify mismatches.¹⁸ The resulting sample contains information of about 2,100 households per year.

Comparing our sample with the full sample of urban households with reference person between 20 and 65 years of age, the households in our sample are slightly younger (40.14 instead of 41.27 years old) because we dropped households with a retired reference person, and have somewhat higher income (\$31,804 versus \$27,908 per year) because we removed incomplete income reporters. The two samples are very similar in terms of family size, the fraction of married and single reference persons, the number of adults and children, the number of earners and average hours worked by the reference person and her or his partner. The samples are also very similar in terms of three different measures of consumption: expenditures on food and beverages, non-durable consumption as defined above and total expenditures.

We make a number of adjustments to the raw data in order to make them comparable to their corresponding theoretical concepts. First, we purge individual income from its predictable dynamics because of seasonality, attrition bias and life-cycle changes. To do this, we regress log real income on a set of month dummies, interview dummies and a fourth order

¹⁷ The low estimate for λ might be surprising, particularly in light of the fact that previous estimates from aggregate data (Campbell and Mankiw, 1990) point towards a much larger fraction of hand-to-mouth consumers (about 50%). Our estimate is consistent, however, with Attanasio and Weber (1995), who show that the Campbell–Mankiw result is driven by aggregation problems and by the effect of demographics and labor supply variables on the marginal utility of consumption.

¹⁸ In particular, we suspect a 'mismatch' if the household changes cohort or because any of the six categories of family composition (male and female members under 2 years old, between 2 and 15 years old and over 15) changes by more than two people. Clearly not all of those are actually mismatches. In particular, a household can change cohort if the title for the house moves from mother to daughter for instance. However, these changes invalidate the link between an observed and a theoretical household.

polynomial in age and take the residuals. A fourth order polynomial captures the shape of the average age profile in income well.¹⁹ In order not to introduce spurious (lack of) correlation, we apply the same procedure to consumption.²⁰

Second, we control for family composition to translate consumption and income per household into per capita terms. If there are returns to scale from living together with other consumers in a household, then family composition may directly affect the marginal utility of consumption. Typically, the literature uses an approximate equivalence scale to address this problem. We follow this practice and regress consumption on the number of adults and the number of children in the household. The estimates indicate that consumption is higher by about 27% for each extra adult, and by 4% for each additional kid.²¹ These coefficients are similar for income, although the coefficient on the number of children has the opposite sign.²²

We do not control for other potential preference shifters like education or hours worked (to allow for non-separabilities between consumption and leisure). These variables are highly correlated with income so that we risk removing exactly the variation we are interested in. Conditioning on education has the additional problem that it would remove all changes in inequality because of changes in the skill premium, an important source of earnings inequality.

Appendix B. Estimation method

This appendix describes the MCMC algorithm for the numerical evaluation of the posterior. The parameters of interest are the unobservable states, $var_t(v)$, $var_t(\alpha)$ and $var_t(u)$ and the so called hyperparameters, which are divided in two blocks: Σ contains the variances of the innovations to the unobservable states and the variances of the error terms in the moment conditions and λ represents the excess sensitivity parameter. All the shocks are assumed to be jointly normal, with a block diagonal covariance matrix.

The estimation algorithm is based on Gibbs sampling. Gibbs sampling is a particular variant of MCMC methods and consists of stepwise drawing from lower dimensional conditional posteriors instead of from the high dimensional joint posterior of the whole set of parameters. In this application, Gibbs sampling is carried out in three steps:

1. *Drawing the variances of the innovations to the unobservable states and the variances of the error terms in the moment conditions.* Conditional on the data and the rest of the parameters, the residuals of the model are observed. Therefore, the posterior of each element of Σ is inverse-gamma with T degrees of freedom and scale parameters given by the sum of squared residuals (details can be found in Gelman et al., 1995). We use a loose, but non-flat prior for the variances of the innovations to the unobservable states. The prior we use is an inverse-gamma with 2 degrees of freedom and scale parameters equal to 0.0005 for the permanent shocks and 0.005 for the transitory shocks. The reason we use a non-flat prior here is to avoid the so called pile-up problem, which is common in time-varying parameter models (see, for instance, Stock and Watson, 1998). Notice that the prior favors time variation in the variance of the transitory shocks. Therefore, if anything, it strengthens our result that transitory shocks did not matter for the increase in inequality of the 1980s.
2. *Drawing the unobservable states.* Conditional on the data and the rest of the parameters, Eqs. (5)–(8) form a system of observation equations. Together with the random walk assumption for the evolution of the time-varying variances, this is a linear and Gaussian state space model. We use a standard simulation smoother (see, for instance, Carter and Kohn, 1994) to make draws from the posterior of the unobservable states (the time-varying variances).
3. *Drawing the excess sensitivity parameter (this step is not implemented for the model without λ).* Conditional on the data and the rest of the parameters, λ appears as a regression coefficient in a system of linear equations. Therefore, its posterior distribution is Gaussian, with mean and variances given by the SUR estimate and the variance of the SUR estimator.

Our estimates are based on 30,000 iterations of the Gibbs sampler, discarding the first 5,000 to allow the system to converge to its ergodic distribution. The sample autocorrelation functions of the draws decay fast and the convergence checks are fully satisfactory.

¹⁹ These 'age effects' in the levels of income should not be confused with the age effects in the variance, which we take into account explicitly in our estimation procedure.

²⁰ In part, the hump-shaped age profile in consumption reflects the predictions of the model, e.g. precautionary savings. However, demographics and other—potentially unobservable—preference shifters also affect the age profile. Because it is impossible to disentangle the two effects, we purge the age profile from consumption as well as income, preferring to remove valid variation from the data rather than not to remove variation that has no bearing on the model.

²¹ We also tried more flexible specifications, allowing for extra persons to have different effects on household consumption depending on their age and gender, but these differences were insignificant.

²² There is an issue whether we want to use separately estimated coefficients for income, or use the coefficients from the consumption regression. Using the estimates from the consumption regression leaves the savings rate unaltered (because income and consumption are adjusted by the same percentage) but may not remove spurious changes in income due to family composition, see Gourinchas and Parker (2002) for a discussion.

Appendix C. Credit constraints and precautionary saving

With a fraction λ of consumers behaving hand-to-mouth, either because they are credit constrained or because of precautionary savings, the response of consumption to income shocks is given by

$$\Delta c_{it} = \begin{cases} \Delta y_{it} & \text{for a fraction } \lambda \text{ of the individuals} \\ \Delta c_{it}^* & \text{for the remaining fraction } 1 - \lambda \end{cases} \quad (9)$$

where Δc_{it}^* is the change in consumption for consumers that behave as if they follow the permanent income model as in (4).

The variance of changes in consumption within a cohort over time, is given by a weighted average of the within-group variances for credit constrained and PIH consumers.²³

$$\text{var}_t(\Delta c) = (1 - \lambda)\text{var}_t(\Delta c^*) + \lambda\text{var}_t(\Delta y) \quad (10)$$

The change in the variance of the level of consumption depends not only on the variance of the changes in consumption, but also on their covariance with past levels of consumption.

$$\Delta \text{var}_t(c) = \text{var}_t(\Delta c) + 2\text{cov}(\Delta c_{it}, c_{it-1}) \quad (11)$$

We need to take a stance on whether consumption last period was set according to the PIH or credit constraints were binding in the last period. Let p be the probability that if the consumer is constrained this period, she was also constrained in the previous period. Then, the evolution of consumption inequality is given by

$$\Delta \text{var}_t(c) = \text{var}_t(v) + \lambda\text{var}_t(\alpha) + \lambda\text{var}_t(u) + \lambda(1 - 2p)\text{var}_{t-1}(u) \quad (12)$$

Assuming $p = 1$, expression (12) simplifies to

$$\Delta \text{var}_t(c) = \text{var}_t(v) + \lambda\text{var}_t(\alpha) + \lambda\Delta \text{var}_t(u) \quad (13)$$

Compared to the evolution of consumption inequality under the PIH in Eq. (6), expression (13) tells us that consumption inequality may increase because of an increase in transitory income inequality, because a fraction λ of consumers displays excess sensitivity to income changes.

In terms of the remaining moment conditions, like in the baseline model in Section 4.1, the covariance of income and consumption contains the same information as the variance of consumption. The moment condition for the autocovariance of income (8) is unaffected by credit constraints.

Identification in model with credit constrained consumers is weaker than in the simpler model in the main text. Conditional on λ , identification of the time-varying variances of the shocks is the same, see Section 4.2. The parameter λ is identified from the correlation between $\Delta \text{var}_t(c)$ and $\Delta \text{var}_t(u)$, see Eq. (13), where $\Delta \text{var}_t(u)$ is identified from the autocovariance structure of income (8). The additional assumption required to estimate λ from Eq. (13) is orthogonality between $\Delta \text{var}_t(u)$ and $\text{var}_t(v) + \lambda\text{var}_t(\alpha)$, which follows from our assumption that the variances of the shocks follow independent laws of motion, see Section 4.3. If the various sources of risk are in fact positively correlated over time, then our estimate for λ is an upper bound for the true fraction of hand-to-mouth consumers in the economy. Since our estimate for λ is very low, we are confident that our results are not driven by this assumption.

References

- Attanasio, O.P., Weber, G., 1995. Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey. *Journal of Political Economy* 103 (6), 1121–1157.
- Attanasio, O.P., Davis, S.J., 1996. Relative wage movements and the distribution of consumption. *Journal of Political Economy* 104 (6), 1227–1262.
- Attanasio, O.P., Battistin E., Ichimura, H., 2004. What really happened to consumption inequality in the US? NBER Working Paper 10338.
- Blundell, R., Preston, I., 1998. Consumption inequality and income uncertainty. *Quarterly Journal of Economics* 113 (2), 603–640.
- Campbell, J.Y., Mankiw, N.G., 1990. Consumption, income, and interest rates: reinterpreting the time series evidence. NBER Working Paper No. w2924*.
- Carroll, C.D., 1992. The buffer-stock theory of saving: some macroeconomic evidence. *Brookings Papers on Economic Activity* 1992 (2), 61–156.
- Carroll, C.D., 2001. A theory of the consumption function, with and without liquidity constraints (expanded version). NBER Working Paper 8387.
- Carter, C.K., Kohn, R., 1994. On gibbs sampling for state space models. *Biometrika* 81, 541–553.
- Chamberlain, G., 1984. Panel data. In: Griliches, Z., Intriligator, M.D. (Eds.), *Handbook of Econometrics*, vol. 2. North-Holland, Amsterdam.
- Cunha, F., Heckman, J., 2006. The Evolution of Labor Earnings Risk in the U.S. Economy. Mimeo, University of Chicago.
- Cunha, F., Heckman, J., Navarro, S., 2005. Separating uncertainty from heterogeneity in life cycle earnings. *Oxford Economic Papers* 57 (2), 191–261.
- Deaton, A., Paxson, C., 1994. Intertemporal choice and inequality. *Journal of Political Economy* 102 (3), 437–467.
- Dickens, R., 2000. The evolution of individual male earnings in Great Britain: 1975–1995. *Economic Journal* 110 (460), 27–49.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 1995. *Bayesian Data Analysis*. Chapman and Hall, London.
- Gervais, M., Klein, P., 2005. Measuring consumption smoothing in CEX data. Mimeo.
- Gorbachev, O., 2007. Did household consumption become more volatile? Mimeo, Columbia University.
- Gottschalk, P., Moffitt, R.A., 1994. The growth of earnings instability in the U.S. labor market. *Brookings Papers on Economic Activity* 1994 (2), 217–254.

²³ This is a special case of the decomposition of an unconditional variance into the expectation of a conditional variance plus the variance of the conditional expectation. The derivation relies on an assumption that the probability that a consumer is credit constrained is uncorrelated with recent shocks to her income. If this assumption is violated, $\text{var}_t(\Delta c)$ includes a between-group variance term $\lambda(1 - \lambda)[E(\Delta c_{it}^* | X = 0) - E(\Delta y_{it} | X = 1)]^2$, where X is a binary variable indicating whether the consumer is constrained, i.e. $P[X = 1] = \lambda$. Under the assumption that X is uncorrelated with v_{it} , α_{it} , u_{it} and u_{it-1} , the cross-sectional conditional means of Δc^* and Δy are equal to zero.

- Gourinchas, P.-O., Parker, J.A., 2001. The empirical importance of precautionary saving. *American Economic Review* 91 (2), 406–412.
- Gourinchas, P.-O., Parker, J.A., 2002. Consumption over the lifecycle. *Econometrica* 70 (1), 47–91.
- Guvenen, F., 2005. An empirical investigation of labor income processes. Mimeo.
- Guvenen, F., 2007. Learning your earning: are labor income shocks really very persistent. *American Economic Review* 97 (3), 687–712.
- Heathcote, J., Storesletten, K., Violante, G.L., 2005. Two views of inequality over the life-cycle. *Journal of the European Economic Association* 3 (2–3), 765–775.
- Heathcote, J., Storesletten, K., Violante, G.L., 2006. Consumption and labor supply with partial insurance. Mimeo.
- Huggett, M., Ventura, G., Yaron, A., 2006. Sources of lifetime inequality. Mimeo.
- Huber, P.J., 1981. *Robust Statistics*. Wiley, New York.
- Krueger, D., Perri, F., 2006. Does income inequality lead to consumption inequality? Evidence and theory. *Review of Economic Studies* 73 (1), 163–193.
- Lillard, L.A., Weiss, Y.A., 1979. Components of variation in panel earnings data: American scientists, 1960–70. *Econometrica* 47, 437–454.
- Moffitt, R.A., Gottschalk, P., 1995. Trends in the autocovariance structure of earnings in the U.S.: 1969–1987. *John Hopkins Working Papers in Economics*, June 1993, revised July 1995.
- Moffitt, R.A., Gottschalk, P., 2002. Trends in the transitory variance of earnings in the United States. *Economic Journal* 112, C68–C73.
- Parker, J.A., Preston, B., 2005. Precautionary savings and consumption fluctuations. *American Economic Review* 95 (4), 1119–1143.
- Pistaferri, L., Blundell, R., Preston, I., forthcoming. Consumption inequality and partial insurance. *American Economic Review*. See (<http://www.aeaweb.org/aer/contents/accepted.html>).
- Primiceri, G.E., Van Rens, T., 2004. Inequality over the business cycle: estimating income risk using micro-data on consumption. Mimeo.
- Slesnick, D.T., Ulker, A., 2004. Inequality and the life cycle: age, cohort effects, and consumption. Mimeo.
- Stock, J.H., Watson, M.W., 1998. Asymptotically median unbiased estimation of coefficient variance in a time varying parameter model. *Journal of the American Statistical Association* 93 (441), 349–358.
- Storesletten, K., Telmer, C.I., Yaron, A., 2004a. Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy* 112 (3), 695–717.
- Storesletten, K., Telmer, C.I., Yaron, A., 2004b. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics* 51 (3), 609–633.
- US Department of Labor, Bureau of Labor Statistics, 1999. *Consumer Expenditure Survey, Interview Survey, 1980–1999*.