

Predictable Life-Cycle Shocks, Income Risk and Consumption Inequality*

Giorgio E. Primiceri

Northwestern University and NBER

Thijs van Rens

CREI and Universitat Pompeu Fabra

February 2006

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 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 90s. 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.

Keywords: consumption, inequality, risk, incomplete markets, heterogeneous income profiles

JEL classification: D12 D31 D52 D91 E21

*Some of the results of this paper are contained in an earlier paper "Inequality over the Business Cycle: Estimating Income Risk Using Micro-Data on Consumption". We are grateful to Pierre-Olivier Gourinchas and Jonathan Parker for many helpful discussions. Thijs van Rens gratefully acknowledges financial support from the Spanish Ministry of Education and Sciences (grant SEJ2005-01124) and the Generalitat de Catalunya (DURSI, 2005SGR 0049).

1 Introduction

This paper evaluates the nature of increased income inequality in the US over the 1980-2000 period. Understanding the sources of inequality is important because of their impact on consumption 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. On the other hand, permanent income shocks will translate almost one-for-one into changes in consumption and will, therefore, have strong welfare effects.

We use repeated cross-section data on income and consumption from the Consumer Expenditure Survey (CEX) to estimate the extent to which changes in transitory and permanent income risk 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 need to postulate a model of consumption choice and make assumptions on the form of the stochastic process governing the evolution of individual income. 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. Following Lillard and Weiss (1979) and more recently Guvenen (2005a, 2005b), we also allow for predictable changes in income, driven by heterogeneity in the life-cycles of different consumers. Since these life-cycle shocks are predictable for the consumer at the beginning of her working life, they contribute to inequality but not to risk. Our model of consumption choice is a simple linearized incomplete markets model, modified to take into account that some consumers may be credit constrained or exhibit precautionary savings behavior.

Our results provide two main conclusions. First, we find that 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. Second, we find that most of these permanent shocks were predictable to consumers. The variance of *unpredictable* permanent and transitory shocks also increased in the early eighties, but the increase was small compared to the total increase in inequality and got reversed by the end of the nineties. Thus, the largest part of the increase in inequality, and therefore the effect on welfare, was not driven by earnings risk during workers' careers but realized before consumers entered the labor market.

A brief summary of the literature helps to understand the intuition behind our results and clarifies the contribution of the paper. Previous studies have followed either of two alternative approaches. In a series of papers, Gottschalk and Moffitt (1994, 1995, 2002) identify the contribution of permanent and transitory shocks to the increase in inequality using only data on income. They argue that the autocovariance structure of income growth is informative about the relative importance of permanent and transitory shocks. They exploit the long panel dimension of the Michigan Panel Study on Income Dynamics (PSID) and find that the increase in inequality was mainly due to permanent shocks.

Adopting a very different methodology, Blundell and Preston (1998) use consumption data and a simple model of consumption behavior similar to ours to learn something about the nature of income uncertainty. Their approach reverses the original idea of Deaton and Paxson (1994), who focused on testing predictions of the standard consumption models, taking as given a particular stochastic process for income. Blundell and Preston's identifying assumption is that under the permanent income hypothesis, consumption responds to permanent but not to transitory shocks to income. Since consumption inequality did not increase (much) over the sample period, they conclude that the increase in income inequality must have been due to transitory shocks.

In this paper, we use both the information in the autocovariance structure of income and the information in the comovement of 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 are in striking contradiction with each other. In order to reconcile them, we need to allow for income shocks that are permanent, but are not transmitted to changes in consumption. Heterogeneity in life-cycle profiles delivers this property. The econometrician observes idiosyncratic shocks that change income permanently. To consumers however, these are not shocks but predictable changes in income that have already been incorporated in the level of consumption. This explains why our estimate of the contribution of predictable shocks is so large: predictable shocks explain the observation that income inequality increased permanently but consumption inequality did not increase much over the same period (Blundell and Preston 1998, Krueger and Perri 2002).

There may be other reasons why consumption does not respond to some permanent income shocks. In particular, if consumers have access to insurance markets, they could share risks with other consumers, insuring some or all of their idiosyncratic shocks, depending on the degree of market incompleteness. This interpretation has been suggested by Krueger and Perri (2002), Blundell, Preston and Pistaferri (2005) and by us in an earlier paper (Primiceri and Van Rens 2004). We investigate the extent to which partially complete insurance markets could explain the joint evolution of income and consumption inequality. Although the evidence is not conclusive, we find 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, highly educated consumers have better access to insurance markets, there is a larger incentive to insure against large shocks) and do not find convincing evidence for any of these.

The finding that heterogeneity in life-cycle profiles is important, relates our work to that of Guvenen (2005a, 2005b). While most of the consumption literature has assumed that, conditional on some observable characteristics, all households face the same life-cycle profiles, Guvenen (2005b) shows that heterogeneity in income profiles accounts for a large part of the increase in income inequality within cohorts. Moreover, in a

second paper (2005a) he shows that heterogeneous income profiles are also consistent with a number of features of consumption data. In this paper, we show that in an incomplete markets model, allowing for heterogeneity in life-cycle profiles, which gives rise to permanent but predictable changes in income, is not only consistent with these features, but also necessary to explain the joint evolution of income and consumption inequality.

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 data we use and discusses the evolution of income and consumption inequality in the raw data. In section 4, we discuss how to 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. We assume income is subject to heterogeneous but predictable life-cycle changes as well as to permanent and transitory shocks. The model for consumption is a simple linearized incomplete markets model, extended to account for precautionary savings or credit constraints.

2.1 Income process

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,

$$\begin{aligned} y_{it} &= y_{it}^p + u_{it} \\ y_{it}^p &= y_{it-1}^p + v_{it} + \alpha_{it} \end{aligned} \tag{1}$$

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.¹

¹For simplicity, we also assume that there are no aggregate shocks. In earlier work we allow for aggregate shocks and find their contribution to inequality to be negligible (Primiceri and Van Rens 2004).

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. Notice that under these assumptions, α_{it} and v_{it} are clearly not separately identified from income data alone, which is why we use consumption data to identify these shocks.

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 zero. Storesletten et.al. (2000a) 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.

Recently however, Guvenen (2005b) showed that failing to take heterogeneity in life-cycle profiles into account may explain estimates of the autocorrelation in income close to one even if the true persistence is much lower. In his estimates, heterogeneity in income profiles explains 65-80% of the life-time increase in income inequality within a cohort and the autocorrelation of persistent unpredictable shocks is around 0.8. In our framework, the predictable shocks α_{it} capture heterogeneity in income profiles. We will refer to α_{it} as predictable shocks or as life-cycle shocks interchangeably. Consistent with Guvenen's results, we find these predictable idiosyncratic shocks to be an important component of inequality.

Different from Guvenen, we do not allow for persistent shocks that are not permanent. This reflects a difference in approach between our paper and Guvenen's. We use consumption data to distinguish life-cycle changes from permanent shocks, whereas Guvenen use income panel data with a long time series dimension. Both approaches have advantages and disadvantages. The disadvantage of our income process is that if the true autocorrelation of persistent income shocks is 0.8 instead of 1, we would expect our estimated decomposition between transitory and permanent unexpected shocks to be biased. On the other hand, Guvenen can not identify permanent shocks from life-cycle changes, so if there are in fact permanent shocks to income, his estimates both for the autocorrelation of persistent (but not permanent) shocks as for the contribution of life-cycle changes would be biased. Largely however, the difference between this paper and Guvenen's is a matter of focus. Guvenen focuses on age effects and estimates the contribution of heterogeneity to the increase in within cohort inequality over the life-cycle. We allow for time variation in the degree of this heterogeneity and ask how much life-cycle shocks have contributed to changes in aggregate inequality.

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

$$\Delta y_{it} = v_{it} + \alpha_{it} + \Delta u_{it}$$

It is important to realize that income changes because of a shock to permanent income,

or because of a *change* in the shock to transitory income. 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 then decreases it again at time $t + 1$.

2.2 Consumption and income inequality under the PIH

In its simplest form, the permanent income hypothesis (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. In addition, under the PIH consumption should not respond to predictable shocks since these do not affect the expected net present value of life-time income. It is crucial for our estimation strategy that, under the PIH, consumption follows a random walk, so that we can interpret first differences as expectational innovations in consumption and therefore as the response to unexpected changes in income. As in Blundell and Preston, we obtain the random walk property by using CRRA preferences with non-stochastic asset returns and log-linearizing the Euler equation.²

Combining the Euler equation with the budget constraint, we can write changes in consumption as changes in the expected value of life-time income plus additive terms for precautionary savings and relative impatience.³ Using the income process specified above, the expression for the change in log consumption reduces to

$$\Delta c_{it} = \rho_t v_{it} + r_t u_{it} + b_{it} \quad (2)$$

where

$$\begin{aligned} b_{it} &= \frac{1}{2}\theta_i V_{t-1} [c_{it}] + \frac{1}{\theta_i} \log \beta_i R_t \\ \rho_t &= \frac{1-(1/R)^{T_w-t}}{1-(1/R)^{T-t}} \quad \text{and} \quad r_t = \frac{R-1}{R} \frac{1}{1-(1/R)^{T-t}} \end{aligned}$$

where θ_i is the coefficient of relative risk aversion and β_i is the discount rate. Consumption does not respond to life-cycle shocks because $E_{t-1} \alpha_{it} = \alpha_{it}$: the realization of these shocks does not reveal any information and therefore does not affect the expected net present value of future income.

The marginal propensity to consume out of a permanent shocks is ρ_t and the marginal propensity to consume out of a transitory shock r_t . Notice that r_t approximately equals the interest rate $R - 1$ for $T \rightarrow \infty$ or for $t \ll T$, and that $0 < r_t < 1$ for $t \leq T - 2$. In the remainder of the paper we will assume that the interest rate is constant and close to zero and that $T_w = T$ (no retirement savings), so that $r_t \equiv 0$ and $\rho_t \equiv 1$. The

²The alternative is to use quadratic utility as in Hall (1978). We choose to work in logarithms to avoid scaling problems. This has the additional advantage that we can allow for a more general specification of the utility function.

³Because the log-linear Euler equation, $c_{it} = E_t [c_{it+1}] - \frac{1}{2}\theta_i V_t [c_{it+1}] - \frac{1}{\theta_i} \log \beta_i R_{t+1}$, includes an additive variance term reflecting precautionary savings, we need to assume that ‘revisions to variance forecasts’ are zero, i.e. $V_t [c_{it+j}] = V_{t+1} [c_{it+j}]$ for $j > 1$.

implication is that consumers react to an income shock in the same manner, regardless of how far they are from retirement age. In the context of a life-cycle model, this is clearly not a very realistic assumption, since just before retirement there is no difference between a transitory and a permanent shock to income, which is captured by the fact that $\rho_t = r_t$ for $t = T_w - 1$. However, the decomposition of income in a transitory and a permanent part with $r = 0$ corresponds to the original paper by Friedman, who defines transitory income shocks as shocks out of which the MPC is zero (see Carroll 2001 for a discussion). Furthermore, we will assume there is no heterogeneity in preferences or discount rates, $\theta_i = \theta$ and $\beta_i = \beta$ for all i . Combined with our alternative proxy for time-variation in precautionary saving (see section 2.3), this assumption allows us to set $b_{it} = 0$ for all i and t .

Then, the following two equations summarize the model

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

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

These expressions simply state that consumption changes one-for-one in response to an unexpected permanent shock to income, but does not respond at all to transitory or predictable income shocks. The evolution of income and consumption inequality follows by taking a cross-sectional variance.⁴

$$\Delta \text{var}_t(y) = \text{var}_t(v) + \text{var}_t(\alpha) + \Delta \text{var}_t(u) \quad (5)$$

$$\Delta \text{var}_t(c) = \text{var}_t(v) \quad (6)$$

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 a shock to permanent income unambiguously and irreversibly increases consumption inequality, as in equation (6), with the observation that aggregate inequality does not (always) increase in the long run.

2.3 Credit constraints and precautionary saving

The model we described so far does not capture excess sensitivity to (transitory) changes in income. As a rough control for excess sensitivity, we will assume that a fraction λ of consumers is ‘hand-to-mouth’ and consumes all of current income.⁵ Clearly λ can be interpreted as the fraction of consumers that are credit constrained. We argue that it is also a good proxy for precautionary savings behavior. Precautionary saving is theoretically closely related to liquidity constraints and empirically virtually indistinguishable (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 \tilde{A} which would be in line with the

⁴Notice that $\text{var}(\Delta u_{it}) = \text{var}_t(u) + \text{var}_{t-1}(u)$ and $2\text{cov}(\Delta u_{it}, y_{it-1}) = -2\text{cov}(u_{it-1}, y_{it-1}) = -2\text{var}_{t-1}(u)$.

⁵Campbell and Mankiw (1990) show that a model where about half the consumers is forward looking and the other half hand-to-mouth describes the aggregate data better than the PIH.

permanent income hypothesis. If wealth is below \tilde{A} , the marginal propensity to consume out of extra cash-on-hand is close to one.

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,

$$\Delta c_{it} = \begin{cases} \Delta y_{it} & \text{for a fraction } \lambda \text{ of the individuals} \\ \Delta c_{it}^* & \text{for the remaining fraction } 1 - \lambda \end{cases} \quad (7)$$

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.⁶

$$var_t(\Delta c) = (1 - \lambda) var_t(\Delta c^*) + \lambda var_t(\Delta y)$$

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.

$$\Delta var_t(c) = var_t(\Delta c) + 2cov(\Delta c_{it}, c_{it-1})$$

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

$$\Delta var_t(c) = var_t(v) + \lambda var_t(\alpha) + \lambda var_t(u) + \lambda(1 - 2p) var_{t-1}(u) \quad (8)$$

In the following we will assume that $p = 1$, so that (8) simplifies to

$$\Delta var_t(c) = var_t(v) + \lambda var_t(\alpha) + \lambda \Delta var_t(u) \quad (9)$$

Compared to the evolution of consumption inequality under the PIH in equation (6), expression (9) 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.

3 Data

For our empirical analysis, we use data on US household income and consumption from the CEX, the Consumer Expenditure Survey (U.S. 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

⁶This 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. Notice that the between-group variance term $\lambda(1 - \lambda)(E[\Delta c_{it}^*] - E[\Delta y_{it}])^2$ evaluates to zero because the cross-sectional mean of both Δc^* and Δy equals zero.

only US dataset that has acceptable consumption as well as income data for the same individuals.⁷

3.1 The microdata

In appendix A we discuss the dataset and our procedure to control for inflation, seasonality, age effects, attrition bias and family composition. The final dataset contains 42,325 households between 20 and 65 years old, whose reference person is not retired nor a student or living in student housing. The sample is representative for the full CEX sample of households aged between 20 and 65 (see the appendix for a more extensive discussion). These households are assigned to five 10-year cohorts by the age of the reference person in 1980. Our secondary dataset contains 75 cohort-year cells with a median cell size of 602 households, see table 1. In this section, we discuss two problems with the data: measurement error and the timing of the questions on income in the CEX. Section 3.2 presents some evidence from the raw data on the evolution of income and consumption inequality over the sample period.

Clearly, both income and consumption are measured with error. However, the evolution of inequality and therefore our estimation results are largely unaffected by this measurement error. Assuming that the measurement error is uncorrelated with the true levels of income and consumption, then measurement error 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 as in (10) and (11) is unaffected, even if the level of inequality is biased.

A more serious problem with the data 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 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 in this year, $\tilde{y}_t = \frac{3}{4}y_{t-1} + \frac{1}{4}y_t$, and correct the moment condition accordingly.⁸

⁷The only alternative would be the Panel Study of Income Dynamics (PSID), which has better income data and a longer panel dimension, but only a rough proxy of consumption (expenditures on food).

⁸Under the assumption, $\frac{8}{3}cov(\Delta\tilde{y}_t, \tilde{y}_{t-1}) - \frac{1}{2}\Delta var(y_t)$ is a consistent estimator for $cov(\Delta y_t, y_{t-1})$. We also estimated the model under the ‘naive’ assumption that $\tilde{y}_{t-1} = y_{t-1}$ and find that this makes

3.2 Income and consumption inequality

Figure 1 shows consumption and income inequality for the five cohorts over the sample period. Consumption and income are logs of real data, and have been adjusted for seasonality, attrition bias, family composition and age profile as described in appendix A, but otherwise these are the raw data. The consumption graphs are comparable to Deaton and Paxson (1994, figure 2) although our sample period is twice as long. We would expect to see two stylized facts in these data. First, we know since Deaton and Paxson (1994) that inequality should rise within a cohort with age (so therefore over time) and that this effect should be present for consumption as well as for income, although less pronounced for consumption because of smoothing. Second, there should be an increase in inequality in the 1980s, which then flattens out in the 90s. Both ‘facts’ are not easy to see, partly because noise clouds the picture, and because both effects are interacting in the same graphs.

In figure 2 we plotted coefficient estimates of a regression of income and consumption inequality on age dummies. The graph is comparable to figure 4 in Deaton and Paxson, and indeed remarkably similar, even though our sample period is longer, we removed households with heads over 65 and we removed the age profile from the levels of consumption. In line with Deaton and Paxson’s conclusion, we find that the data on consumption inequality seem consistent with the prediction of the PIH that within cohort inequality should increase with age, although the effect is not very strong.

The fact that within-cohort inequality always tends to increase, implies that it is important to take demographics into account when describing aggregate inequality. If inequality increases with age, as Deaton and Paxson show it does, then changes in the age distribution of the workforce will have effects on aggregate inequality. Suppose the ‘baby-boomers’ entered the labor market around 1980 when they were thirty. *Ceteris paribus* we would expect consumption and income inequality to decrease around this time. Inequality would also decrease around in 2045, when the baby-boomers retire, and would gradually increase in between due to the age effect. To isolate aggregate inequality and to see how much it is affected by demographic changes, we plotted the variance of log consumption and income for the whole sample as well as average within-cohort inequality over time. The two graphs turn out to be very similar, see figure 3. In the remainder of this paper, we will present average within-cohort inequality, controlled for age effects, and interpret it as aggregate inequality.

Income inequality sharply increased in the early eighties and then remained high over the late eighties and all of the nineties. Consistent with other studies, we also find a temporary peak in inequality in the mid eighties, which seems to be specific to the CEX data (Attanasio et.al. 2004). Consumption inequality did not increase much over the sample period. This is also consistent with what other studies have found (Krueger and Perri 2002).

very little difference in the results.

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 all 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.

4.1 Moment conditions

We use expressions (3) and (7) 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 (5) and (9) and reproduced here for convenience.

$$\Delta var_t(y) = var_t(v) + var_t(\alpha) + \Delta var_t(u) \quad (10)$$

$$\Delta var_t(c) = var_t(v) + \lambda var_t(\alpha) + \lambda \Delta var_t(u) \quad (11)$$

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 (3) and (7), we get

$$\Delta cov_t(y, c) = \Delta var_t(c) = var_t(v) + \lambda var_t(\alpha) + \lambda \Delta var_t(u) \quad (12)$$

Clearly, the evolution of the covariance of income and consumption contains very similar information as the evolution of the variance of consumption under the model. Using both moment conditions should improve the efficiency of our estimates and help distinguish signal from noise.

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, 1995, 2002). Because $cov(y_{it}, y_{it-1}) = cov(\Delta y_{it}, y_{it-1}) + var(y_{t-1})$ and we are already using the information contained in the variance of income, we use $cov(\Delta y_{it}, y_{it-1})$. From (3), we get

$$cov_t(\Delta y, y_{t-1}) = -var_{t-1}(u) \quad (13)$$

Moment conditions (10), (11), (12) and (13) contain all information in the second moments of the joint evolution of income and consumption that we can retrieve from the data.

4.2 Age effects

Because our dataset contains 5 cohorts (at least 3 in every year), we can take first differences across cohorts as well as over time. This procedure adds three more equations to our vector of moment conditions, which makes the estimation more efficient and helps to identify some of the model parameters. To see this, consider the evolution of consumption inequality from (11), explicitly taking into account that each year the cohort grows a year older as well.

$$var_{at}(c) = var_{a-1,t-1}(c) + var_t(v) + \lambda var_t(\alpha) + \lambda \Delta var_t(u)$$

where a is age. Iterating back this expression to the ‘birth year’ of the cohort (in practice we use the year in which the cohort entered the labor force), we get

$$var_{at}(c) = var_{0,t-a}(c) + \sum_{s=0}^{a-1} [var_{t-s}(v) + \lambda var_{t-s}(\alpha)] + \lambda [var_t(u) - var_{t-a}(u)]$$

Now consider two cohorts in year t , one with age a and another one with age $a - 10$. We will assume that there are no cohort effects, in the sense that all cohorts start with the same inequality at birth: $var_{0,t-a}(c) = var_{0,t-a+10}(c)$. We need to make this assumption because it is impossible to separately identify cohort, age and time effects and follow Heathcote, Storesletten and Violante (2005) who show that cohort effects can be safely abstracted from. Then, taking a first difference across cohorts, we get the following moment condition.

$$\begin{aligned} \Delta_a var_{at}(c) &= var_{at}(c) - var_{a-10,t}(c) \\ &= \sum_{s=a-10}^{a-1} var_{t-s}(v) + \lambda \sum_{s=a-10}^{a-1} var_{t-s}(\alpha) + \lambda \Delta_a var_{t-a+10}(u) \end{aligned}$$

Inequality between the two cohorts differs, because of shocks that were realized between $t - a + 10$ and $t - a$, when the older cohort was already alive but the younger one was not yet. Age moment conditions for the variance of income and the covariance can be similarly derived.

The age moment conditions allow us to identify which part of the increase in within-cohort inequality is caused by age and which part by time effects. For this purpose, we make two assumptions. First, we do not allow for time variation in the pre-sample period and only estimate the average variances of the permanent shocks. Let A denote the average variance of the predictable permanent shocks α_{it} from 1925 to 1965 and V the average variance of the unpredictable permanent shocks v_{it} over the same period.⁹ This assumption gives a tighter estimate of the average age effects but otherwise does not affect the results since we are not interested in the evolution of inequality before our sample period (unlike Storesletten, Telmer and Yaron 2000b). Second, we assume that the contribution to inequality due to transitory shocks in the pre-sample period averages out so that we can set it to zero. This assumption is saying that because the contribution

⁹The oldest cohort in our sample was born in 1925 the youngest in 1965.

of permanent shocks cumulates over time and the contribution of transitory shocks does not, for large age differences, differences in inequality across cohorts are caused only by permanent shocks.

Then, the age moment conditions for the variance of income, the variance of consumption and their covariance, can be written as,

$$\Delta_a var_{at}(y) = 10(V + A) \quad (14)$$

$$\Delta_a var_{at}(c) = 10(V + \lambda A) \quad (15)$$

$$\Delta_a cov_{at}(y, c) = \Delta_a var_{at}(c) = 10(V + \lambda A) \quad (16)$$

where V and A are age effects, or the average contribution of unpredictable and predictable permanent shocks to within cohort inequality because of age. Whenever we present our estimates for the time effects of the variance of permanent shocks, we also present estimates for the age effects so that the difference can be interpreted as aggregate inequality.

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.¹⁰ 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 algorithm for the numerical evaluation of the posterior of the parameters of interest.

¹⁰The 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, we conclude that the normality assumption does not affect the results.

5 Results

Figure 4 plots the actual evolution of average within cohort inequality (thin solid line) and the fitted values of the model (thick solid line). 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. These estimates are obtained using a larger set of moment conditions than the two represented in the figure. For example, the fitted values for the variance of consumption and the covariance of consumption and income are the same and represent an ‘average’ of the data for those two sets of moment conditions. The random walk assumption on the law of motions for the time varying variances imposes further smoothness. 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.

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 the variances.¹¹ As is clear from the graph, the model predicted values are very close to the robust series. We should stress that 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.

The point estimate for the fraction of credit constrained consumers λ is consistently very low. In the baseline model, we estimate λ at 2.2% with a standard error of 2.0% hence insignificantly different from zero. The estimate is robust to various modifications of the dataset, see table 2. This finding might appear surprising, particularly in light of the fact that previous estimates of λ from aggregate data (Campbell and Mankiw 1990) point towards a much larger fraction of hand-to-mouth consumers of about 50%. It 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. Since we use micro-data and account for household characteristics such as family composition, we find a substantially lower estimate for excess sensitivity of consumption to income.

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. Figures 5, 6 and 7 present the counterfactual evolution of income inequality if unpredictable permanent shocks v_{it} , predictable life-cycle shocks α_{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

¹¹ 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).

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. The lower panels can be interpreted as the evolution of income inequality if there would have been only one type of shocks to income. Because these graphs present average within cohort inequality, in figures 5 and 6, we have also plotted the increase in inequality because of the age effects V and A (the age effect of transitory shocks in figure 7 is zero by assumption).

It is clear from figure 6, 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 (2005b), in a rather different context. Guvenen finds that heterogeneous life-cycle changes make up 65 to 80% of the life-time increase in income inequality within a cohort. We assert that, in addition, changes in the amount of this heterogeneity accounts for more than 100% of the increase in aggregate income inequality over the period 1980-2000.

The variance of unpredictable permanent shocks went up as well, see figure 5. Part of the increase in inequality in the early eighties can be attributed to increased permanent income risk, but this increase is a factor 3 smaller than the increase in inequality due to predictable shocks. Moreover, from the second half of the eighties onwards however, permanent risk seems to have gone down again and at the end of the sample the increase is close to and insignificantly different from the increase that can be explained by the within cohort age effect. Transitory inequality also increased in the early eighties with the peak in the same order of magnitude as that in the variance of unpredictable permanent shocks and not significantly different from zero. Here, the reversal in the late eighties and throughout the nineties is even more pronounced: by 2000 income inequality would have decreased substantially if transitory shocks were the only shocks to income.

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, this finding cannot be interpreted to mean that permanent income risk (the variance of *unexpected* permanent shocks) has increased 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 eighties and stayed high, risk seems to have temporarily increased in the early eighties and then come down again. By 2000, permanent risk was as high as it was in 1980 whereas transitory risk did substantially decrease.

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 hitherto disparate, and seemingly contradictory, branches of literature (Gottschalk and Moffitt 1994, 1995, 2002; Blundell and Preston 1998). In order to understand to what drives this result, we re-estimated our model several times, imposing different sets of restrictions in order to reproduce either Moffitt and Gottschalk's or Blundell and Preston's results. In figure 8 we plot fitted values for income and consumption inequality for these models and figure 9 presents the contribution of permanent shocks for each model. The thin and thick solid lines in figure 8 present the data and the fitted values for the full model and are the same as in figure 4. In figure 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, which represents the estimates of a model in the spirit of Gottschalk and Moffitt (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 (10) and (14) and the autocovariance of income (13). 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 in income inequality is due to permanent shocks (figure 9). Therefore, consumption inequality should have increased substantially over the whole sample period, with the strongest increase in the early eighties.

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 (11), (12), (15) and (16) in addition to the moment conditions for income inequality (10) and (14), but not condition (13) 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 figure 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 figures 8

and 9.¹² Now, there is a conflict between the information in the consumption data, represented in moment conditions (11) and (12), and the information in moment condition (13) 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

Our consumption model is an incomplete markets model in the sense that 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. Saving and borrowing allows non-credit constrained consumers to self-insure against transitory shocks as well as predictable permanent shocks, but 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 shocks. In the context of our baseline estimates, the two interpretations are observationally equivalent: both types of shock are permanent in terms of the autocovariance of income and consumption does not respond to either type.

Consider a model in which every period a fraction of consumers receives a permanent shock to which they had previously insured. One can think of this either as a fraction of consumers that have access to markets for state-contingent assets or as a fraction of shocks that are insurable if all consumers receive a single shock in each period. Then, $var_t(\alpha) / (var_t(v) + var_t(\alpha))$ provides an estimate for this fraction in every time period and can be seen as a measure of market completeness or the degree of risk sharing in the economy. In our baseline estimates, the degree of risk sharing such defined is about 0.85 (with a standard error of 0.11) averaged over the whole sample and 0.66 (0.10) based on the information in the age moment conditions about the pre-sample period.

The partial risk sharing interpretation relates our analysis to our earlier paper (Primiceri and Van Rens 2004) and to Blundell, Preston and Pistaferri (2005). Blundell et.al. use food consumption and some household characteristics to impute total nondurable consumption in the PSID data. Then, in a framework similar to the one in this paper, they estimate the fraction of permanent income shocks that are not insured, in the sense

¹²We also allow for a fraction λ of credit constrained consumers in these estimates, but given the very small point estimate this makes no difference in the results.

that consumption responds to these shocks. They find consumption growth responds by a factor 0.61 to permanent shocks, which implies that the fraction of the variance of permanent income shocks that are insured based on their estimates is 0.63, roughly similar to our estimate.

How can we distinguish the risk sharing explanation from the predictable shocks or heterogeneous income profiles explanation? In this section, we test a number of predictions of the risk sharing hypothesis that may shed some light on the difference. In particular, we re-estimate the model on a number of samples that differ in the way they have been constructed from the microdata. For each dataset, we report the results as the fraction of permanent shocks that are predictable or insured in three broad time periods: the early eighties, all of the eighties and the whole sample period (eighties and nineties). In addition, we present the fraction of predictable permanent shocks estimated from the age moment conditions (14), (15) and (16), which represents the average fraction of predictable permanent shocks in the pre-sample period 1925-1965. These estimates are reported in table 2.

First, we explore how the results change using different definitions of income. 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. Rows 2, 3 and 4 of the table presents estimates for the model when income is not defined as disposable income after taxes and transfers, as in the baseline, but 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). Rows 2 and 3 are indistinguishable from the baseline estimates. 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.

Secondly, we would expect that highly educated consumers, who have better access to insurance markets, are able to insure more idiosyncratic risk than consumers with little education. To test this prediction, we split up the sample in two, including only households for which the reference person is in the lower or upper half of the education distribution respectively. These estimates are presented in rows 5 and 6. The difference between the estimates in the two rows goes in the direction that the risk sharing hypothesis would predict, with more highly educated households able to insure a somewhat larger fraction of shocks. However, the difference is small and insignificant. Moreover, it is puzzling that risk sharing within *each* subsample, seems to be higher than in the full sample, suggesting that the amount of risk sharing between consumers with different education levels is particularly low. Note also that the finding that consumers with higher education level experience more shocks to income that consumption does not respond to, is consistent not only with risk sharing, but also with heterogeneity in life-cycle profiles (see Guvenen 2005b).

Finally, if risk sharing is important, we would expect that large income shocks are

mainly responsible for this result. In the presence of small financial transaction costs, there is an incentive to insure large shocks but not small ones. Anecdotal evidence from formal insurance markets suggests that insurance is indeed available only against major events. In rows 7 and 8 we present estimates of the model on a dataset from which we have eliminated the 5 and 10% percent of the households that experience the largest shocks. Clearly, this takes away a large part of the identifying variation and in fact the standard errors of our estimates increase substantially. The point estimates however, are still very close to the baseline. The earlier periods suggest somewhat less risk sharing but on average over the 1980-2000 sample period the the estimates are identical to those of the baseline sample.

Concluding, we find that the evidence is not inconsistent with the risk sharing hypothesis. However, the very small magnitudes of the effects make it impossible to reject the null that its predictions are false. These small magnitudes also suggest that it is unlikely that risk sharing constitutes a large part of the story. Heterogeneous life-cycle changes, as in Guvenen (2005a, 2005b), are also consistent with the results in this section and seem to be the more likely explanation for our results.

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 had to make in 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 equivalents;¹³ and in rows 4 and 5 we try different deflators to convert nominal consumption into real terms.¹⁴ 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 nondurable 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 figure 4), the estimation method performs well at extracting signal from noise and the estimates in row 8 are very close to the baseline.

¹³In the baseline we regress consumption and income on the number of adults and the number of children and take the residuals. In row 3, 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.

¹⁴In the baseline, we use the CPI for nondurables 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 nondurable consumption. In row 5, we use the CPI for total expenditures to see whether the results might be driven by the fact that we use different deflators for consumption and income.

Finally, we explore sensitivity to modifications of the model. In row 9 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; in row 10, we restrict the variance of predictable life-cycle shocks to be constant over time; and in rows 11 and 12 we relax the assumption that a consumer that is credit constrained this period was also constrained last period with probability one, see section 2.3.¹⁵ Again, none of these changes affects the results.

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 eighties, but these increases were small and got reversed in the nineties.

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 heterogeneity in consumers' life-cycle profiles, following Guvenen (2005a, 2005b), 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 sample period (Blundell and Preston 1998, Krueger and Perri 2002).

We explore to what extent our results might be driven not by heterogeneity in life-cycle profiles, but by risk sharing. To match the observed evolution of income and consumption inequality, insurance markets would have to be close to complete. Moreover, we test a number of predictions from a partial insurance model and do not find conclusive evidence for any. We conclude that it is unlikely that risk sharing is the main explanation that consumption inequality did not increase more.

A Data description

The Consumer Expenditure Survey (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). These households are then interviewed each quarter, for five quarters in a row. The first meeting is an introductory interview where respondents

¹⁵The probability p that a currently credit constrained consumer was constrained also last period is not identified from our data. Therefore, we set it to two different values, $p = 0.75$ and 0.5 compared to $p = 1$ in the baseline.

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, and CPIs are available for different expenditure categories corresponding to those in the CEX. The unit of observation is a ‘consumer unit’ (CU), which is a group of individuals living together as a family. Questions about income are asked in the first and fifth interview only, and refers the previous 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 (clothing). 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 nondurable expenditures respectively.

We run a series of checks to identify mismatches, including large changes in family composition.¹⁶ In addition, we drop the following households: non-urban households because they are not available in 1982 and 1983; households that report zero expenditures on food; households for which some categories of income are flagged as incompletely reported or topcoded (18% of all urban households with reference person between 20 and 65 years of age); households of which either the reference person or his or her partner is retired (4%); and households of which either the reference person or his or her partner is either a student or living in student housing (1.1 and 0.5%). Finally, we remove the bottom 0.05 and top 0.01 percentile of the consumption and income distribution as well as of the first differences, to avoid the variances to be driven by outliers. In total we drop about 30,000 quarterly observations out of 260,000 corresponding to 8000 out of 50,000 households. Finally, dropping households with a reference person under 20 or over 65 years old, we are left with about 2100 out of 6000 households per year (many households have a reference person over 65 years old).

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 (41.27 instead of 40.14 years old) because we dropped households with a retired reference person, and have somewhat higher income (\$31804 versus \$27908 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

¹⁶In particular, we suspect a ‘mismatch’ if the household changes cohort or because any of the 6 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 2 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, we feel these changes invalidate the link between an observed and a theoretical household. Notice also that we do not drop these households from the sample, but split them up into two ‘households’ in order to preserve as many observations as we can.

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 adjustment to the raw data in order to make them more comparable to the theoretical concepts described by the model. In particular, we control for seasonality, family composition, attrition bias and age effects. Entering these variables as controls rather than purging the data of these effects is not feasible because our estimation works off a secondary dataset of means, variances and covariances. To control for seasonality and attrition bias, we simply run a regression of log real consumption and income on a set of month dummies as well as interview dummies and take the residuals adding the constant back up.

Controlling for family composition is necessary because the model refers to consumption per capita whereas the dataset only contains consumption per households. If there are returns to scale from living together with other consumers in a household, then the family composition may shift preferences over consumption goods. Typically, the literature either estimates or uses a rough approximation of a simplified equivalence scale. 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.¹⁷ Notice that if we would not control the data for family composition, the variance in family size between households would introduce spurious variance in consumption.¹⁸ We control for family composition effects in income similarly. The coefficient estimates in this regression are similar to those for consumption (although the coefficient on the number of children has the opposite sign).¹⁹

Graphing log real consumption and income, controlled for seasonality, attrition bias and family size by cohort over time it is clear that both income and consumption increase for the younger cohorts and decrease for the older cohorts. We take out this ‘hump shaped’ life-cycle profile of consumption and income by taking the residuals of

¹⁷We 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. The estimates intuitively make sense, although the coefficients seem rather small. Other adjustments that have been used in the literature are dividing household consumption by the number of adults plus 0.4 times the number of children (Parker and Preston 2001) or by a more elaborate equivalence scale like Blundell and Preston (1998) who calculate the first adult as 0.55 couple, the second and third adults 0.45 times, and subsequent adults 0.4 times, and give children different weights depending on their age from 0.07 (under 2 years) to 0.38 (aged 17 or 18).

¹⁸In order not to introduce spurious changes in consumption because of this procedure, we first run a cross-sectional regression on family composition and take the residuals including the constant. Then, we run a fixed effects regression on the same explanatory variables and again take the residuals including the family specific fixed effect. Notice that this second step will only change consumption and income for families that experience a within sample change in family composition.

¹⁹There 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 (dad may start working less because he wants to be with the baby more often, see Gourinchas and Parker (2002) for a discussion). We use the estimated coefficients for income.

a regression on a fourth order polynomial in age, which captures the shape of the age profiles well.²⁰ This has the additional advantage that by construction, consumption and income will have mean zero trend across individuals.²¹

We use these data to construct a synthetic panel dataset where we follow groups of individuals born around the same time over the whole sample period. We use 10 year cohorts which are labeled by their average age in 1980, i.e. cohort 45 consists of consumers that were aged 41 through 49 in 1980. The age of a household is defined as the age of the reference person (the person or one of the persons who owns or rents the home) at the last interview. Because the cell sizes are too small for monthly time series, we use annual time periods. Table 1 reports the cell sizes for the cohort-year cells. We eliminate cells with average age below 25 or above 60 (these cells are shaded in the table).

B Estimation method

This appendix describes the Markov Chain Monte Carlo (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 three blocks: Σ contains the variances of the innovations to the unobservable states and the variances of the error terms in the moment conditions, $[A, V]'$ is the vector of age effects 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 four steps:

1. Drawing the age effects.

Conditional on the data and the rest of the parameters, A and V appear as regression coefficients in a system of linear equations. Therefore, their posterior distribution is Gaussian, with mean and variances given by the SUR estimate and the variance of the SUR estimator.

2. Drawing the excess sensitivity parameter.

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.

²⁰These ‘age effects’ in the levels of consumption and income should not be confused with the age effects in the variances, which we take into account explicitly in our estimation procedure.

²¹We considered other possible preference shifters as well, the most important being hours worked by the first and second earner, in order to allow for non-separabilities between consumption and leisure (families with a ‘working mum’ may eat out more often, which leads to higher expenditures on food than families that eat at home). However, since these variables are highly correlated with income we risk removing exactly the variation we are interested in: the comovement between consumption and income.

3. 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. 1994). 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.

4. Drawing the unobservable states.

Conditional on the data and the rest of the parameters, equations (10), (11), (12), (13) 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).

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

References

- Attanasio, Orazio, Erich Battistin and Hidehiko Ichimura (2004). What really happened to consumption inequality in the US? mimeo, UCL-IFS.
- Attanasio, Orazio, Steven J. Davis (1996). Relative Wage Movements and the Distribution of Consumption. *Journal of Political Economy*, 104(6), pp.1227-1262.
- Attanasio, Orazio P. and Guglielmo Weber (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, 103(6), pp. 1121-1157.
- Blundell, Richard and Ian Preston (1998). Consumption Inequality and Income Uncertainty. *Quarterly Journal of Economics*, 113(2), pp.603-40.
- Blundell, Richard, Ian Preston and Luigi Pistaferri (2005). Consumption Inequality and Partial Insurance. mimeo UCL-IFS/Stanford.

- Campbell, John Y., N. Gregory Mankiw (1990). Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. NBER Working Paper No.w2924*
- Carroll, Christopher D. (2001). A Theory of the Consumption Function, With and Without Liquidity Constraints (expanded version). NBER Working Paper 8387.
- Carter, C. K. and R. Kohn (1994), "On Gibbs Sampling for State Space Models", *Biometrika*, 81, 541-553.
- Deaton, Angus and Christina Paxson (1994). Intertemporal Choice and Inequality. *Journal of Political Economy*, 102(3), pp. 437-467.
- Gelman, A., J. B. Carlin, H. S. Stern and D. B. Rubin (1995), "*Bayesian Data Analysis*", London: Chapman and Hall.
- Gervais, Martin and Paul Klein (2005). Measuring Consumption Smoothing in CEX Data. mimeo, University of Western Ontario.
- Gottschalk, Peter and Robert Moffitt (1994). The Growth of Earnings Instability in the U.S. Labor Market. *Brookings Papers on Economic Activity*, 1994(2), pp. 217-254.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker (2001). The Empirical Importance of Precautionary Saving. *American Economic Review*, AEA Papers and Proceedings January 2001.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002). Consumption over the Lifecycle. *Econometrica*, 70(1), pp.47-91.
- Guvenen, Fatih (2005a). Learning Your Earning: Are Labor Income Shocks Really Very Persistent. mimeo, University of Rochester.
- Guvenen, Fatih (2005b). An Empirical Investigation of Labor Income Processes. mimeo, University of Rochester.
- Heathcote, Jonathan, Kjetil Storesletten and Gianluca Violante (2005). Two Views of Inequality Over the Life-Cycle. *Journal of the European Economic Association*, 3(2-3), pp.765-775.
- Huber, Peter J. (1981). *Robust statistics*. New York: Wiley.
- Kim, C. and C. R. Nelson (1999), *State-Space Models with Regime Switching*, MIT Press: Cambridge, Massachusetts.
- Krueger, Dirk and Fabrizio Perri (2002). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. NBER Working Paper 9202.
- Lillard, Lee A. and Yoram A. Weiss (1979). Components of Variation in Panel Earnings Data: American Scientists, 1960-70. *Econometrica*, 47, pp.437-454.

- Moffitt, Robert A. and Peter Gottschalk (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, Robert A. and Peter Gottschalk (2002). Trends in the Transitory Variance of Earnings in the United States. *Economic Journal*, 112, pp.C68-C73.
- Primiceri, Giorgio E. and Thijs van Rens (2004). Inequality over the Business Cycle: Estimating Income Risk Using Micro-Data on Consumption. CREI Working Paper.
- Stock, James H. and Mark W. Watson (1998), Asymptotically Median Unbiased Estimation of Coefficient Variance in a Time Varying Parameter Model, *Journal of the American Statistical Association*, 93, pp.349-358.
- Storesletten, Kjetil, Chris I. Telmer, Amir Yaron (2000a). Consumption and Risk Sharing Over the Life Cycle. NBER Working Paper 7995.
- Storesletten, Kjetil, Chris I. Telmer, Amir Yaron (2000b). The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk. NBER Working Paper 8040.
- US Department of Labor, Bureau of Labor Statistics (1999). Consumer Expenditure Survey, Interview Survey, 1980-1999.

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		

Table 2
Contribution of predictable permanent shocks to inequality in different samples

	Lambda	Share of permanent inequality that is predictable			
		pre-sample	1980-1985	1980-1990	1980-2000
1) Baseline	0.022 (0.020)	0.66 (0.10)	0.79 (0.17)	0.83 (0.15)	0.85 (0.11)
2) Income before taxes	0.019 (0.018)	0.63 (0.11)	0.79 (0.18)	0.83 (0.16)	0.85 (0.12)
3) Income (before tax), excl financial income	0.019 (0.017)	0.64 (0.10)	0.79 (0.18)	0.83 (0.15)	0.86 (0.11)
4) Earned income (before tax and transfers)	0.029 (0.025)	0.82 (0.08)	0.83 (0.22)	0.87 (0.17)	0.91 (0.12)
5) Low education levels	0.029 (0.026)	0.71 (0.10)	0.78 (0.23)	0.84 (0.21)	0.91 (0.13)
6) High education levels	0.048 (0.039)	0.80 (0.28)	0.85 (0.23)	0.92 (0.16)	0.94 (0.12)
7) Without 5% largest income shocks	0.035 (0.022)	0.43 (0.41)	0.75 (0.24)	0.79 (0.22)	0.86 (0.15)
8) Without 10% largest income shocks	0.035 (0.022)	0.42 (0.42)	0.73 (0.25)	0.77 (0.23)	0.85 (0.15)

Table 3
Robustness checks

	Lambda	Share of permanent inequality that is predictable			
		pre-sample	1980-1985	1980-1990	1980-2000
1) Baseline	0.022 (0.020)	0.66 (0.10)	0.79 (0.17)	0.83 (0.15)	0.85 (0.11)
2) Including topcoded incomes	0.025 (0.021)	0.63 (0.10)	0.77 (0.18)	0.82 (0.15)	0.85 (0.12)
3) Different equivalence scale	0.033 (0.027)	1.01 (0.26)	0.79 (0.20)	0.83 (0.19)	0.83 (0.17)
4) CPI: Item specific	0.022 (0.020)	0.66 (0.10)	0.77 (0.18)	0.81 (0.15)	0.83 (0.12)
5) CPI: Total expenditures	0.021 (0.019)	0.66 (0.10)	0.79 (0.17)	0.83 (0.15)	0.85 (0.11)
6) Consumption: Food only	0.045 (0.035)	0.62 (0.12)	0.70 (0.26)	0.77 (0.22)	0.82 (0.15)
7) Consumption: Total expenditures	0.033 (0.029)	0.64 (0.11)	0.64 (0.23)	0.68 (0.21)	0.75 (0.16)
8) Sampling error: Robust estimators	0.058 (0.041)	0.53 (0.14)	0.73 (0.20)	0.75 (0.18)	0.79 (0.14)
9) Interest rate: 5%	0.024 (0.022)	0.66 (0.10)	0.81 (0.17)	0.84 (0.15)	0.86 (0.10)
10) No time variation predictable shocks	0.022 (0.020)	0.66 (0.10)	0.70 (0.20)	0.82 (0.21)	0.84 (0.21)
11) Credit constraints: $p = 0.75$	0.011 (0.010)	0.65 (0.10)	0.80 (0.16)	0.83 (0.15)	0.86 (0.11)
12) Credit constraints: $p = 0.50$	0.006 (0.005)	0.65 (0.10)	0.80 (0.16)	0.83 (0.15)	0.85 (0.11)

Figure 1

Income (dotted, left scale) and consumption (solid, right scale) inequality by cohort

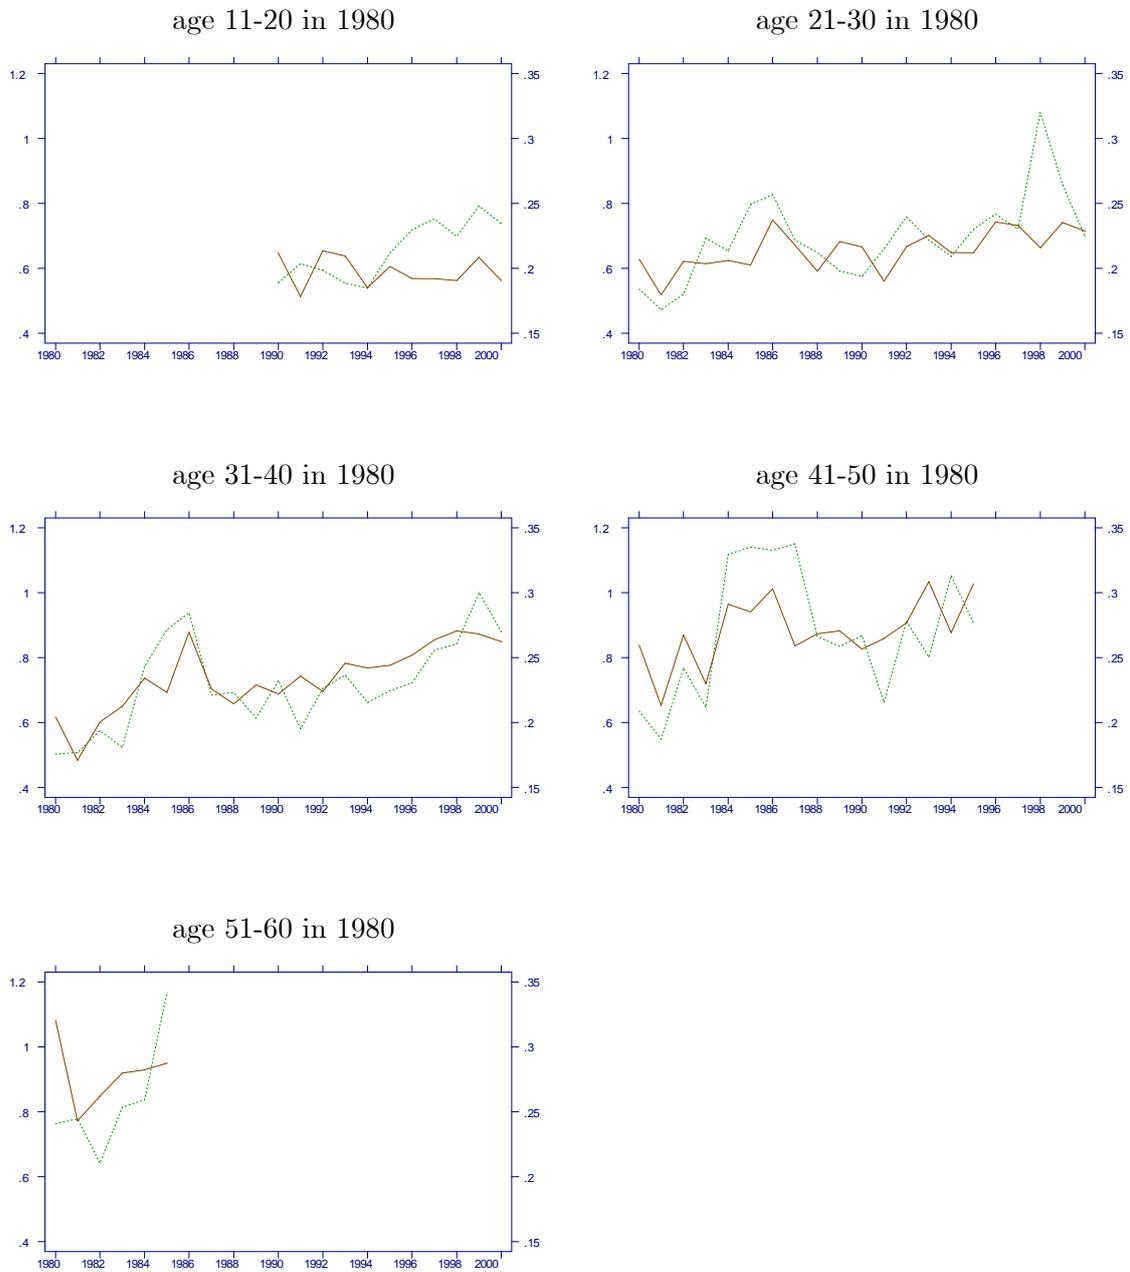
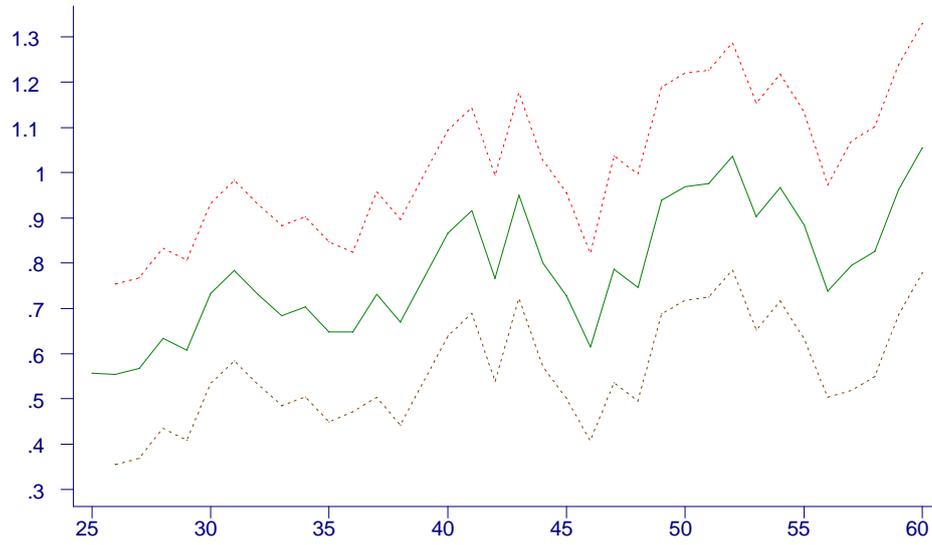


Figure 2
Age effects in inequality (2 standard error bands)

Income



Consumption

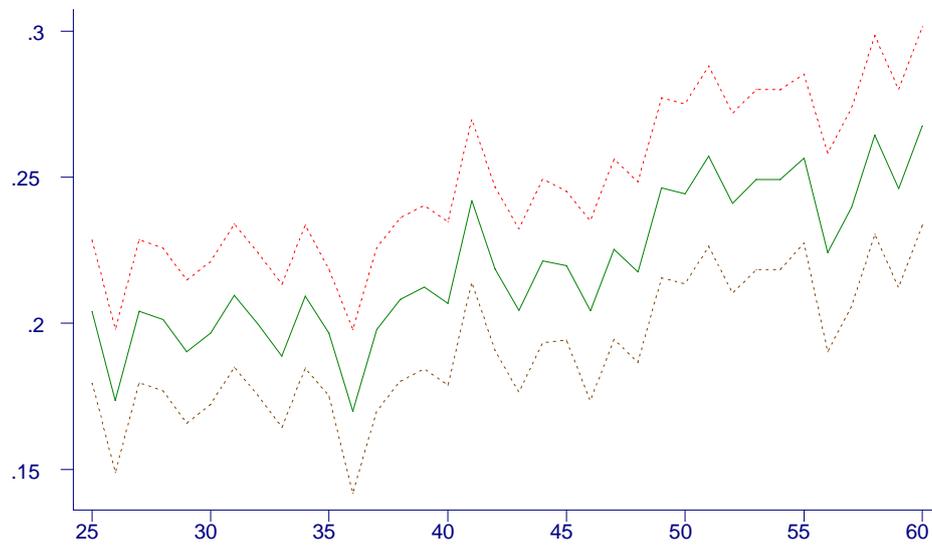
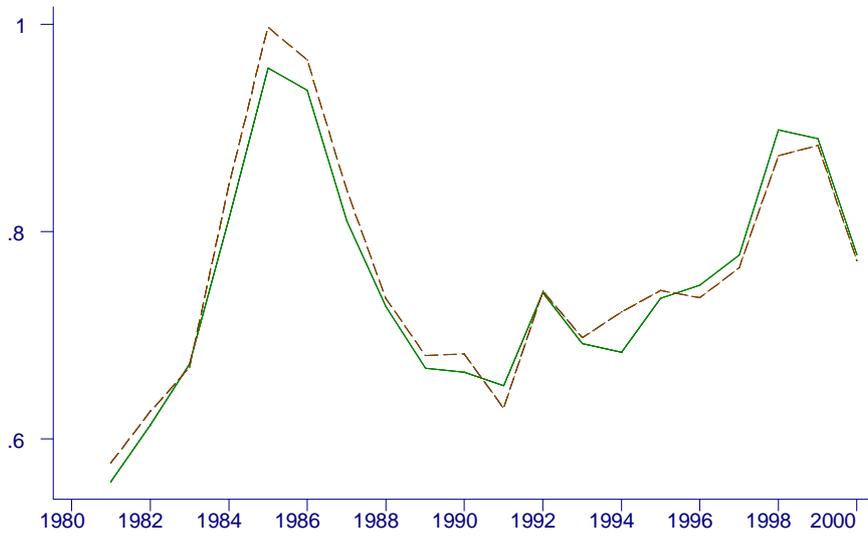


Figure 3
Aggregate (solid) and average within cohort (dashed) inequality

Income



Consumption

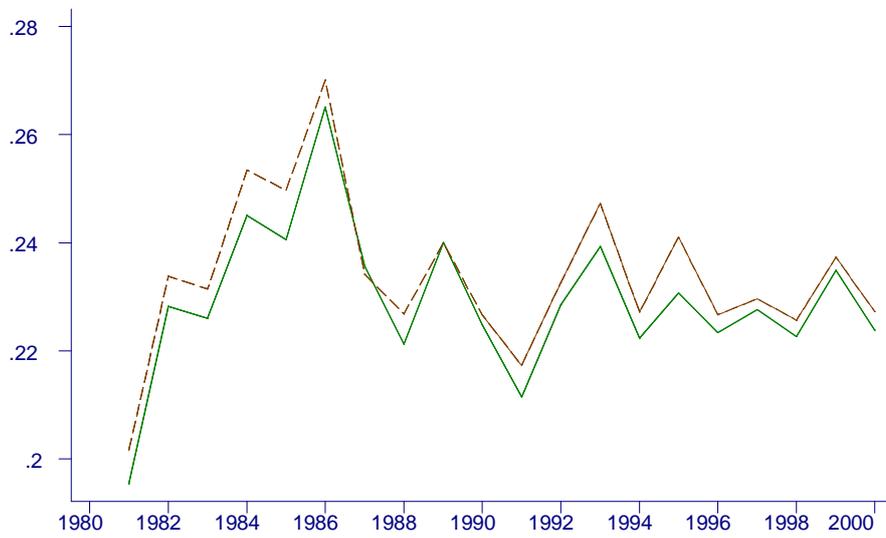


Figure 4
Income and consumption inequality: data and model predicted values

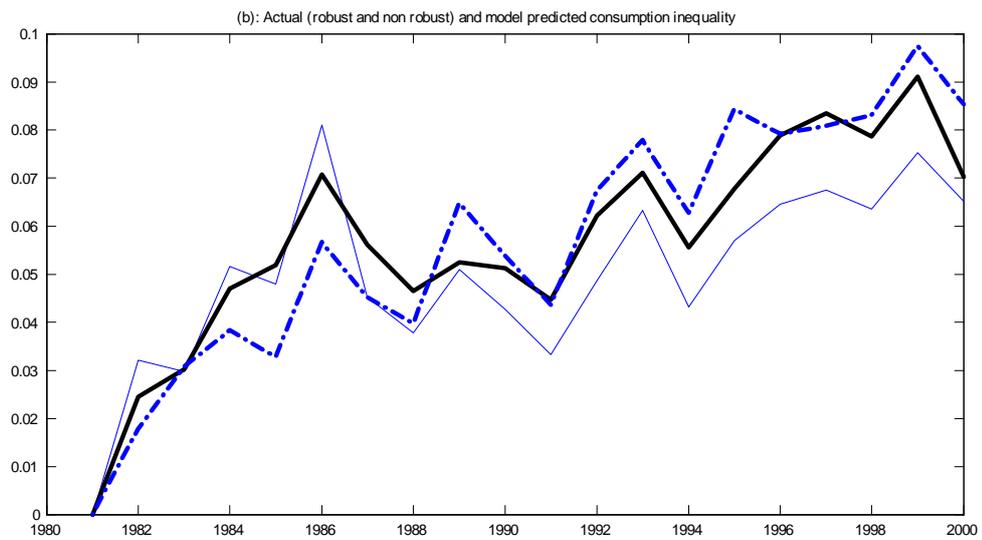
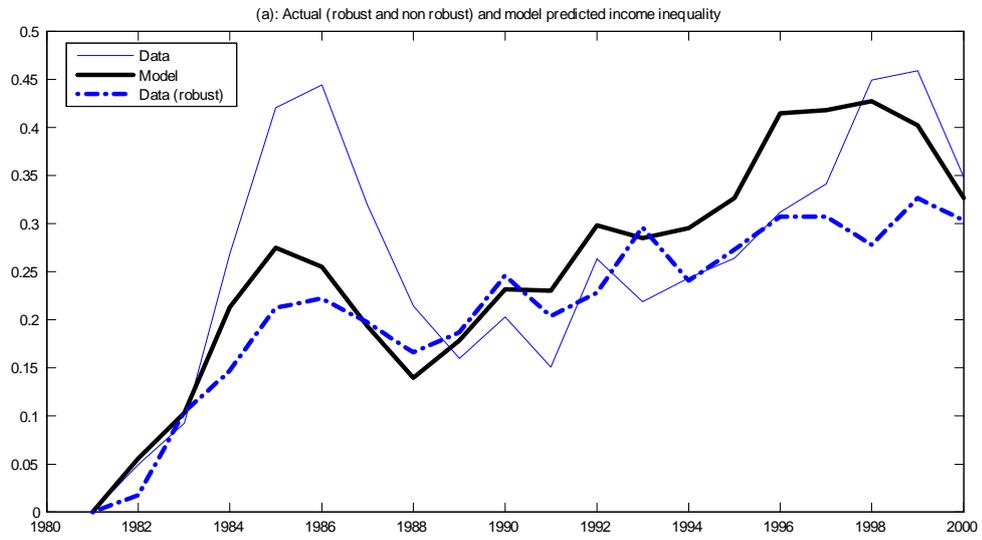


Figure 5
Contribution of unpredictable permanent shocks to income inequality

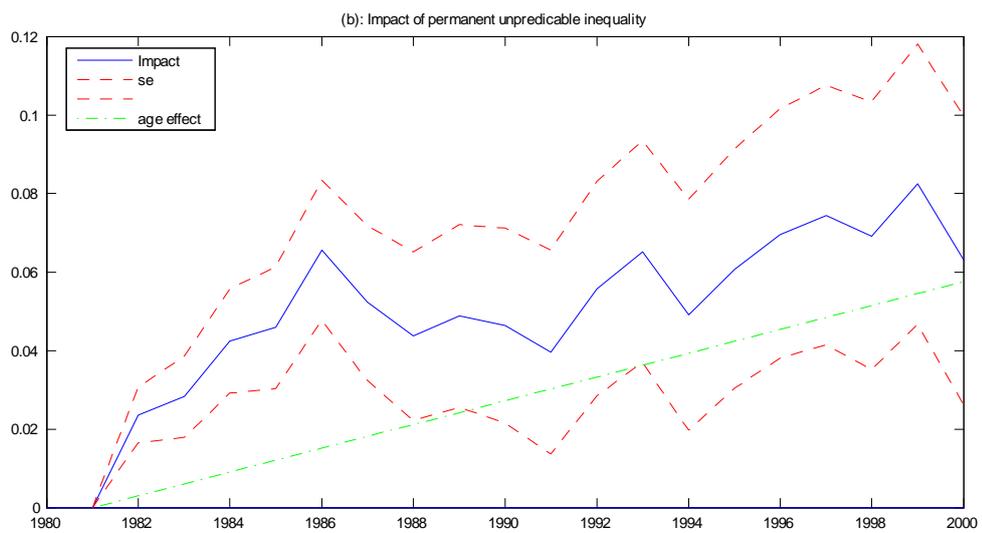
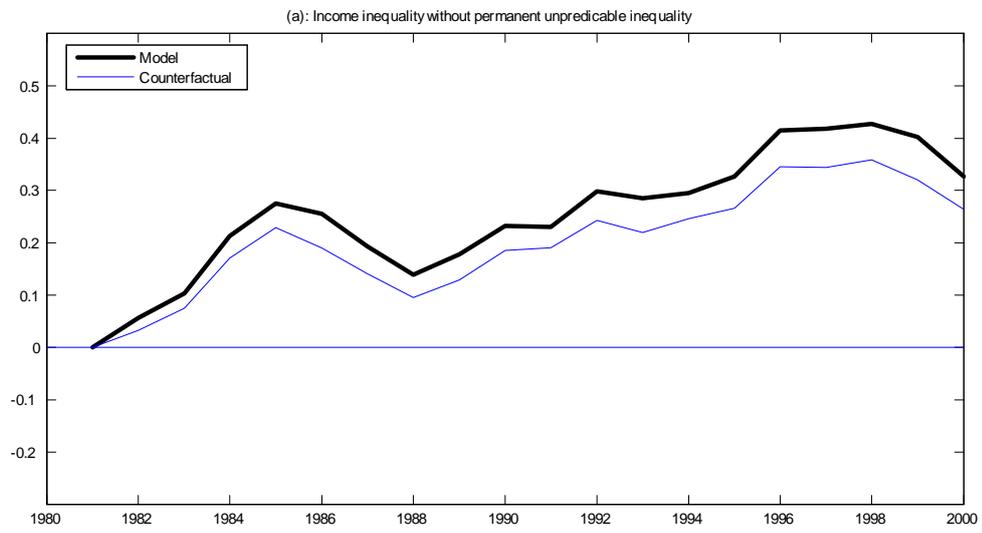


Figure 6
Contribution of predictable permanent shocks to income inequality

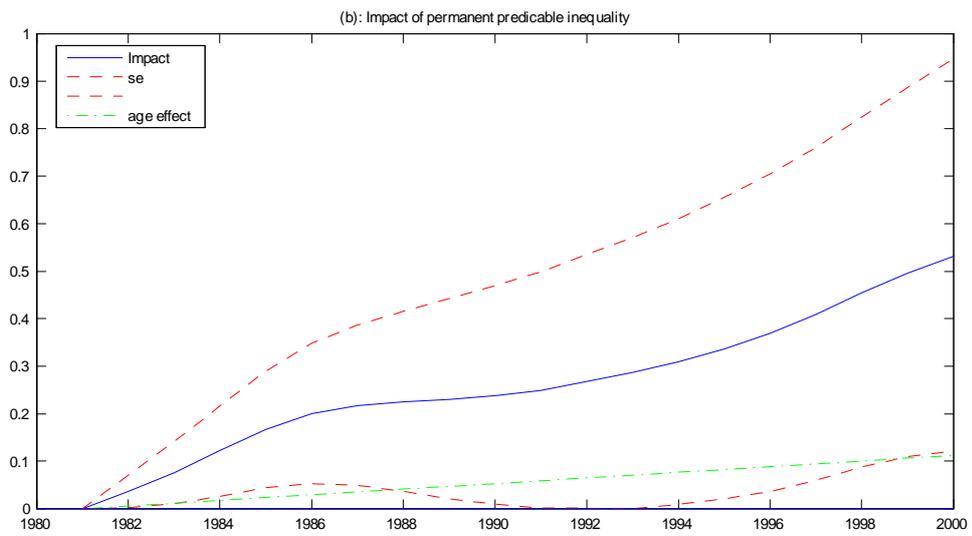
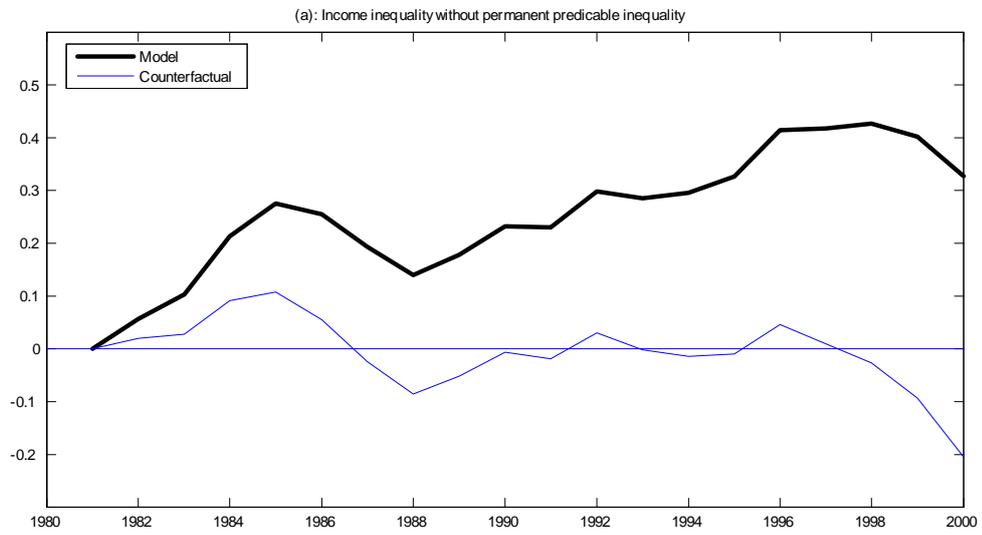


Figure 7
Contribution of transitory shocks to income inequality

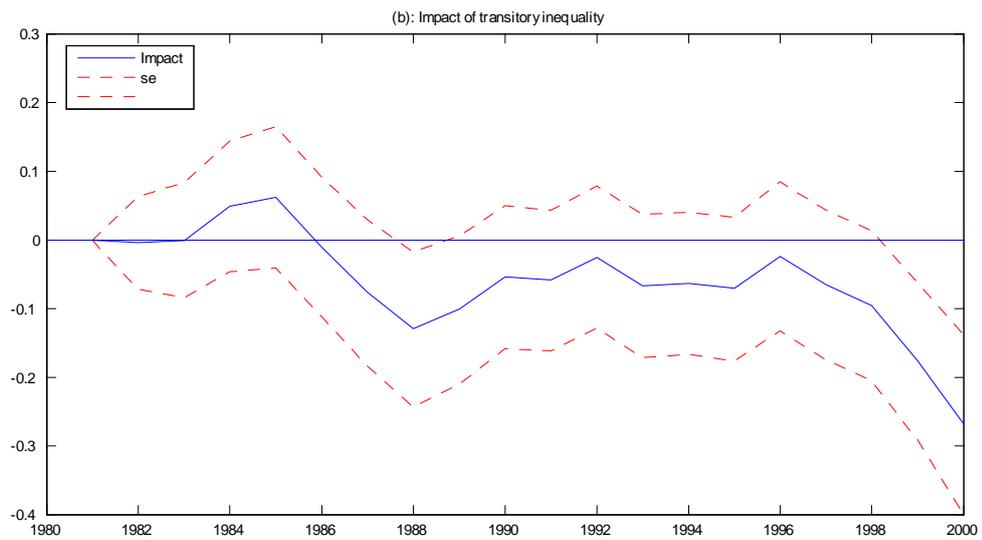
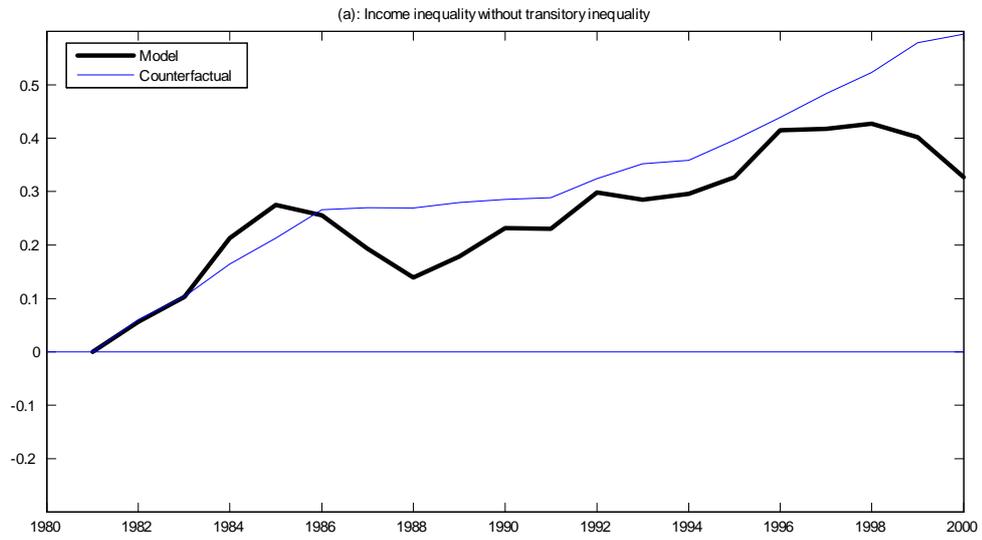


Figure 8
 Model predicted evolution of inequality for different models

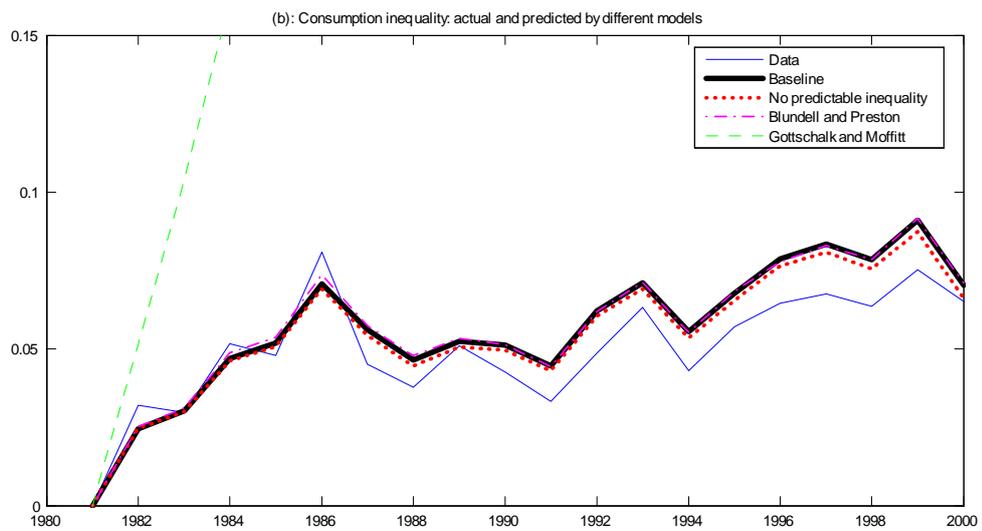
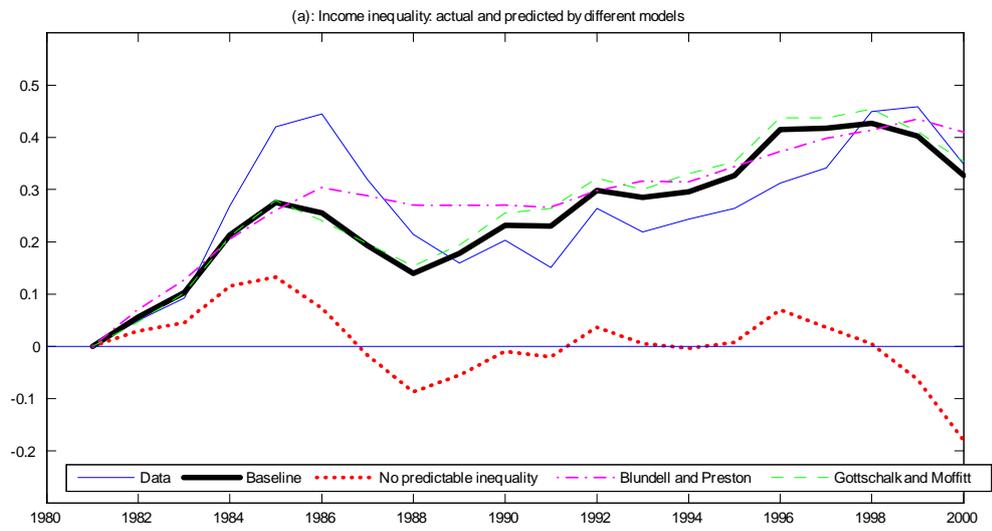


Figure 9
Contribution of permanent shocks to inequality for different models

