Inequality over the Business Cycle: Estimating Income Risk using Micro-Data on Consumption

Giorgio Primiceri and Thijs van Rens[*]

Preliminary draft: July 2002

Abstract

We use CEX repeated cross-section data on consumption and income, to evaluate the nature of increased income inequality in the 1980s and 90s. We decompose unexpected changes in family income into transitory and permanent, and idiosyncratic and aggregate components, and estimate the contribution of each component to total inequality. The model we use is a linearized incomplete markets model, enriched to incorporate risk-sharing while maintaining tractability. Our estimates suggest that taking risk sharing into account is important for the model fit; that the increase in inequality in the 1980s was mainly permanent; and that inequality is driven almost entirely by idiosyncratic income risk. In addition we find no evidence for cyclical behavior of consumption risk, casting doubt on Constantinides and Duffie's (1995) explanation for the equity premium puzzle.

Keywords: consumption, inequality, risk, incomplete markets, business cycle
JEL classification: D12 D31 D52 D91 E21 E32 G10

Contents

1 Introduction

There is much and uncontroversial evidence that earnings and income inequality in the US increased over the last two decades.[1] In this paper we try to gain some insight in the type of income risk that was mainly responsible for this increase. In a series of papers, Gottschalk and Moffitt (1994, Moffitt and Gottschalk 1995, 2002) try to measure the extent to which the increase in earnings inequality could have been merely an increase in 'churning' or transitory movement in earnings. The welfare implications of such an increase in transitory inequality would be very different from the implications of an increase in the variance of permanent income.

Using a very different approach, Blundell and Preston (1998) elaborate on this last issue, using consumption inequality as a better measure for welfare comparison than income or earnings inequality. The literature on consumption inequality, starting with Deaton and Paxson (1994), initially focused on testing predictions of the standard consumption models on the data, assuming a particular stochastic process for income. Blundell and Preston reversed this, and ask the question how we can use consumption models to learn something about the nature of income uncertainty. They use the joint evolution of income and consumption inequality, to decompose idiosyncratic shocks to income into transitory and permanent components.

Distinguishing transitory from permanent shocks to income is an important part of the decomposition of income risk, because consumption can be smoothed against the former but not the latter through borrowing and lending of a risk-free bond. However, self-insurance is only one instrument to smooth consumption. The literature so far has ignored the possibility that consumers are able to smooth consumption by sharing risks, through the market for risky financial assets, family ties, formal insurance policies etc. To explore the implications of insurance through this channel, we need to distinguish idiosyncratic shocks with a cross-sectional mean of zero, which can be insured through risk-sharing across individuals, from aggregate shocks that affect the income of all consumers and cannot be insured. Thus, we want to distinguish four components of income risk: transitory versus permanent, as well as aggregate versus idiosyncratic shocks. Notice that only in the case of a transitory idiosyncratic shock, self-insurance and risk-sharing are substitutes. Permanent idiosyncratic shocks can be insured through risk-sharing arrangements, but not by borrowing and lending; whereas aggregate permanent shocks are uninsurable through either mechanism.

In this paper, we extend the existing literature in two directions. First, we use a model of consumption behavior which allows for the possibility that consumers are able to insure part of the idiosyncratic permanent shocks to income through risk sharing. Second, we extend the stochastic process for income to account for aggregate sources of uncertainty. Both extensions are important as ignoring them may bias the estimates for

---

[1]See Autor, Katz and Krueger (1997) for a survey.

the components of income risk. Not taking risk-sharing into account would spuriously attribute any insurable increase in inequality, i.e. any increase in income inequality that does not translate into an increase in consumption inequality, to an increase in the variance of transitory shocks.

The second extension to take aggregate shocks into account, is important for at least two reasons. First of all, we may be interested in the direct distributional impact of aggregate shocks. In other words, which part of changes in inequality is attributable to a deterministic relation between inequality and business cycle fluctuations? Of course, if an aggregate shock hits all consumers equally, these shocks will affect the mean but not the shape of the income distribution. However, if the sensitivity of income to the business cycle differs across consumers, then aggregate shocks will have an impact on the higher moments of the income distribution.

A second reason why it is attractive to have aggregate shocks and idiosyncratic risk in the same model, is that it will allow us to make statements about the comovement of risk with the business cycle. Many people have speculated about the cyclical behavior of the variance of income shocks. Constantinides and Duffie (1996) provide an explanation for the equity premium puzzle based on consumption risk being countercyclical. They hypothesize that periods in which stocks offer high returns coincide not only with times of high average consumption levels, but also with times of low consumption risk, thus further deflating their value as insurance instruments.

Our contribution is methodological as well as substantive. We use a simple model of the joint evolution of consumption and income to derive a large set of moment conditions that capture all the available information in the data. We then estimate the components of income risk, using both GMM and a likelihood-based estimation technique treating the aggregate shocks and time-varying variances of the idiosyncratic shocks as unobservable components. We argue that both estimation methods are more efficient and more formal than a decomposition using an exactly identified model, and that the likelihood based estimation solves many of the practical difficulties associated with GMM when the parameter space is high dimensional.

We find that almost all inequality is driven by idiosyncratic risk. Aggregate shocks have a negligible impact on the variance of the income distribution, implying that there is no systematic relation between inequality and the business cycle. We also find no evidence that idiosyncratic consumption risk is countercyclical. Our estimates for the transitory and permanent parts of idiosyncratic shocks to income correspond well to Moffitt and Gottschalk's (2002) results, even though their methodology is entirely different from ours. Risk sharing as an insurance mechanism against idiosyncratic permanent shocks is important to reconcile the evolution of consumption inequality with the decomposition between transitory and permanent inequality using income data only (like Moffitt and Gottschalk do).

This paper is organized as follows. In the next section, we set out a simple model

4

of consumption smoothing and discuss how this can be used to decompose income risk into its four components. We also discuss an ad-hoc adjustment of the basic model to incorporate the possibility or risk-sharing through insurance markets. Section 3 describes the data we use and presents some of the main results that can be seen in the raw data. In section 4 we discuss the estimation procedure, and provide some evidence that our estimation model performs better than a simple decomposition in capturing the actual evolution of inequality from the data. Finally, in section 5 we present our results, and section 6 concludes.

2   Income risk and inequality

In this section we discuss the standard life-cycle model of consumption and the stochastic process for income that we employ to relate the evolution of income and consumption inequality to income risk, largely following Blundell and Preston (1998). The basic model is a linearized incomplete markets model, but we extend it to take risk sharing into account.

2.1   The consumption Euler equation

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. It is crucial for the estimation strategy that consumption follows a random walk, so we can interpret first differences as expectational innovations in consumption, and therefore as the response to unexpected changes in income. To obtain the random walk property, we can either use quadratic utility as in Hall's (1978) classic paper, or use CRRA preferences with non-stochastic asset returns and log-linearize the Euler equation. We will take the second approach which has obvious advantages in terms of generality and therefore makes explicit the simplifying assumptions we need to make.[2]

Consider the standard intertemporal consumption problem, with stochastic labor income and CRRA preferences. We then get the familiar Euler equation, where we allow for heterogeneity in the discount factor and coefficient of relative risk aversion.

$$c_{it}^{-\theta_i} = \beta_i E_t \left[R_{t+1} c_{it+1}^{-\theta_i} \right].$$

Assuming asset returns are non-stochastic,[3] we can make a log-normal approximation

---

[2]The disadvantage is that we need to log-linearize the budget constraint.

[3]Because we will use our estimation results to make inferences about risk sharing through risky assets, it is awkward that all uncertainty is in labor income. However, one can think of the problem as a portfolio choice, where we use only the pricing kernel for the riskfree bond.

to obtain the result that consumption has a unit root[4]

$$\log c_{it} = E_t\left[\log c_{it+1}\right] - \underbrace{\tfrac{1}{2}\theta_i V_t\left[\log c_{it+1}\right]}_{\text{Precautionary savings}} - \underbrace{\tfrac{1}{\theta_i}\log\beta_i R_{t+1}}_{\text{Relative impatience}}.$$

where the term $\tfrac{1}{2}\theta_i V_t\left[\log c_{it+1}\right]$ expresses precautionary savings (consumption today is lower than consumption tomorrow because consumption tomorrow is uncertain), and the term $\tfrac{1}{\theta_i}\log\beta_i R_{t+1}$ represents the relative impatience of individual i with respect to the market.

Iterating, we can express expectations about the future consumption path as

$$\log c_{it} = E_t\left[\log c_{it+s}\right] - \sum_{j=1}^{s}\left(\tfrac{1}{2}\theta_i E_t V_{t+j-1}\left[\log c_{it+j}\right] + \tfrac{1}{\theta_i}\log\beta_i R_{t+j}\right).$$

If 'revisions to variance forecasts' are zero, i.e. if $E_t V_{t+j-1}\left[\log c_{it+j}\right] = V_{t+j-1}\left[\log c_{it+j}\right]$ for $j > 1$, then we get the following useful expression for the change in log consumption (cf. Blundell and Preston 1998):

$$\Delta\log c_{it} = (E_t - E_{t-1})\left[\log c_{it+s}\right] + \tfrac{1}{2}\theta_i V_{t-1}\left[\log c_{it}\right] + \tfrac{1}{\theta_i}\log\beta_i R_t.$$

From the (log-linearized) budget constraint, we get that expectational revisions to consumption equal changes in the expected value of life-time income

$$(E_t - E_{t-1})\left[\log c_{it+s}\right] = \tfrac{1-1/R}{1-(1/R)^{T-t}}\sum_{s=0}^{T_w-t} R^{-s}\left(E_t - E_{t-1}\right)\left[\log y_{it+s}\right],$$

where R is a function of future interest rates over the remaining life, T is the time of death and T_w the time of retirement. Notice that assets drop out of this expression because they are predetermined (i.e. they are not subject to expectation revisions) by the assumption of non-stochastic interest rates.

2.2   Income process

Consider an individual income process that consists of a permanent and a transitory component, where shocks can be either aggregate or idiosyncratic.

$$\log y_{it} = y_{it}^p + \underbrace{u_{it} + \gamma_i e_t}_{\text{transitory shocks}} \tag{1}$$

$$\log y_{it}^p = \alpha_i + y_{it-1}^p + \underbrace{n_{it} + \delta_i v_t}_{\text{permanent shocks}}$$

where u_{it} is a transitory idiosyncratic shock, n_{it} a permanent idiosyncratic shock, α_i a deterministic drift, e_t a transitory aggregate shock, and v_t a permanent aggregate shock

---

[4]If x is normal (so e^x is log-normal), then $\log E\left[e^x\right] = E\left[x\right] + \tfrac{1}{2}V\left[x\right]$. This approximation is very similar to the expectation of a second order Taylor approximation of e^x around $E\left[x\right]$.

to income. Individuals respond differently to aggregate shocks, maybe because some people work in more cyclical industries. An implicit assumption in the way we specified the income process is that the responses to aggregate shocks for each individual are constant over time.

All shocks are assumed to have zero mean. The idiosyncratic responses to aggregate shocks γ_i and δ_i are normalized to have cross-sectional mean 1, and the idiosyncratic shocks are normalized to have cross-sectional mean zero. The deterministic drift α_i has a cross sectional mean of zero because we remove the average age effects from the data. We also assume that all shocks are uncorrelated over time and uncorrelated with each other; the variances of the idiosyncratic shocks u_{it} and n_{it} may vary over time but are constant across individuals; and the idiosyncratic responses to aggregate shocks γ_i and δ_i are constant over time and in the cross-section uncorrelated with u_{it} and n_{it}.

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 auto-correlation 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. *ADD: references to MaCurdy (1982) and Quah (1992, 1998).*

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

$$\Delta \log y_{it} = \alpha_i + \underbrace{\Delta u_{it} + \gamma_i \Delta e_t}_{\text{change transitory shocks}} + \underbrace{n_{it} + \delta_i v_t}_{\text{permanent shocks}} .$$

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 $t + 1$.

2.3 Estimation model

Substituting the income process specified above into the log-linearized Euler equation, we get the following expression for the change in log consumption[5]

$$\Delta \log c_{it} = \underbrace{r_t \ (u_{it} + \gamma_i e_t)}_{\text{transitory shocks}} + \underbrace{\rho_t \ (n_{it} + \delta_i v_t)}_{\text{permanent shocks}} + \tfrac{1}{2}\theta_i V_{t-1} \left[\log c_{it}\right] + \tfrac{1}{\theta_i} \log \beta_i R_t$$

where

$$r_t = \tfrac{R-1}{R} \tfrac{1}{1-(1/R)^{T-t}} \quad \text{and} \quad \rho_t = \tfrac{1-(1/R)^{T_w-t}}{1-(1/R)^{T-t}}.$$

The marginal propensity to consume out of a permanent shocks is ρ_t, out of a transitory shock r_t. Notice that r_t approximately equals the interest rate $R - 1$ for $T \to \infty$ or for $t << T$, and that $0 < r_t < 1$ for $t \leq T - 2$.

In the rest 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 $\rho_t \equiv 1$. Of course we make these assumptions for the sake of simplicity. The implication is that consumers react to an income shock in the same manner, regardless of how far they are from retirement age. This is clearly not a very realistic assumption, as 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 2001a p.6 for a discussion). As long as we keep this difference into account when interpreting the results, the assumption is innocuous.[6]

Then, the following two equations summarize our estimation model

$$\Delta \log y_{it} = \alpha_i + \Delta u_{it} + \gamma_i \Delta e_t + n_{it} + \delta_i v_t \quad (2)$$

$$\Delta \log c_{it} = b_i + n_{it} + \delta_i v_t \quad (3)$$

where

$$b_i = \tfrac{1}{2}\theta_i V_{t-1} \left[\log c_{it}\right] + \tfrac{1}{\theta_i} \log \beta_i R.$$

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 model prediction that a shock to permanent income unambiguously and irreversibly increases

---

[5] $y_{it+s} = y_{it-1} + (s+1)\alpha_i + \Delta\tau_{it-1} + ... + \Delta\tau_{it+s-1} + \pi_{it-1} + ... + \pi_{it+s-1}$ so that (note that $E_t \Delta \tau_{it+1} = -\tau_{it}$): $(E_t - E_{t-1})[y_{it+s}] = 0$ if $s = 0$, $\tau_{it} + \pi_{it}$ if $s = 1$, and π_{it} for $s \geq 2$. Substituting this into the sum gives the expression in the text.

[6] Although the expressions look nicer with $r_t = 0$ and $\rho_t = 1$, it is not clear that the model is easier to estimate with this assumption. In particular, if r_t and ρ_t are allowed to vary across cohorts this introduces an extra dimension of identifying variation. However, we would be suspiciuous of any parameters identified off this variation alone.

inequality, first put forward by Deaton and Paxson (1994), with the observation that aggregate inequality seems not (always) to increase in the long run.

The fact that within-cohort inequality will always tend 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. We will briefly return to the distinction between age and time effects when we discuss the data, but in this paper we are generally concerned with the evolution of within-cohort inequality.

We use expressions (2) and (3) to calculate means, variances and covariances for each cohort in each year. We then use these moment conditions to estimate the components of income risk we are interested in. In addition to the time varying variances of the idiosyncratic shocks $Var_t(n)$ and $Var_t(u)$ which represent permanent and transitory idiosyncratic risk, we estimate the actual aggregate shocks e_t and v_t, as well as the model parameters $Var(\gamma)$, $Var(\delta)$, $Cov(\gamma, \delta)$, $Cov(\gamma, y_0)$, $Cov(\delta, y_0)$ and $Cov(\delta, c_0)$ that measure the distributional impact of the business cycle. The full set of moment conditions is given in appendix A. These include the mean changes in income and consumption, the variances and covariance of the changes in income and consumption, the changes in the variances and covariance of income and consumption, as well as the autocovariance of income. In the next subsection we discuss one of these moment conditions into detail, to illustrate the distributional impact of the aggregate shocks.

2.4   Impact aggregate risk on inequality

Consider the change in the variance of log consumption, calculated from equation (3)

$$\underbrace{\Delta Var_{jt}(\log c)}_{\text{growth consumption inequality}} = \underbrace{Var_j(b) + 2Cov(b_i, \log c_{it-1}) +}_{\text{cohort-specific deterministic trend}} \underbrace{Var_t(n)}_{\text{idiosyncratic (permanent) risk}}$$
$$\underbrace{+v_t^2 Var(\delta) + 2v_t Cov(\delta_i, \log c_{it-1}) + \xi_{jt}^c}_{\text{distributional impact aggregate shocks}}$$

where $\Delta Var_{jt}(\log c) = Var(\log c_{ijt}) - Var(\log c_{ijt-1})$ denotes the change in the variance of consumption in cohort j and year t, and the disturbance term ξ_{jt}^y represents the difference between the sample and population variances.

The impact of aggregate shocks on inequality consists of two parts. The first part is always positive. It represents the cross-sectional variance of changes in individual consumption in response to an aggregate shock $Var(\delta_i v_t)$. Because income of some consumers is more sensitive to the business cycle than of others, this effect increases the divergence in income because of an aggregate shock, whether it is a positive or

negative shock. Nevertheless the distributional impact of aggregate shocks is not always to increase inequality. This is what the second term stands for. If the covariance between δ_i and c_{it-1} is positive, meaning that consumers with high initial consumption are sensitive to business cycle shocks, then a negative aggregate shock will decrease inequality if the size of the shocks is not too large. Intuitively, if the rich are affected more strongly by aggregate shocks than the poor, their consumption will decline more due to a negative shocks, thus bringing them closer to the consumption level of the poor. Similarly, if $Cov\left(\delta_i, \log c_{it-1}\right)$ is negative, then negative aggregate shocks will increase inequality, whereas positive aggregate shocks will narrow the income distribution, at least if the size of the shocks is not too large.

Why is it important to allow for aggregate shocks in this framework? First of all, we may be interested in the direct distributional impact of the business cycle, as explained above. Second, not taking these effects into account may bias the decomposition of income shocks into transitory and permanent shocks. This is because of the idiosyncratic responses to aggregate shocks, which make innovations in income and consumption correlated with past levels of income and consumption, e.g. $Cov\left(\Delta \log c_{it}, \log c_{it-1}\right) = Cov\left(n_{it} + \delta_i v_t, \log c_{it-1}\right) \neq 0$ even if we assume that permanent *idiosyncratic* shocks are uncorrelated to any information in the past. Similarly, because of these idiosyncratic sensitivities to the cycle, the assumption that transitory and permanent shocks are uncorrelated to each other breaks down, as $Cov\left(\Delta u_{it} + \gamma_i \Delta e_t, n_{it} + \delta_i v_t\right) = v_t \Delta e_t Cov\left(\gamma, \delta\right) \neq 0$ generally. We will take all these possible correlations into account when we estimate the components of income risk. Finally, it is attractive to have aggregate shocks and idiosyncratic risk in the same model, as it will allow us to make statements about the cyclical movement of risk.

2.5   Risk sharing

The model we presented so far is an 'incomplete markets model' in the sense that insurance markets are non-existing: consumers can save and borrow a risk-free bond freely, but they cannot pool risks with other consumers. We saw that by borrowing and lending consumers could fully insure against transitory shocks, but could not insure permanent shocks at all. If instead consumers can partly insure against permanent shocks by risk-sharing through insurance markets or markets for risky financial assets, our estimates for the transitory and permanent components of income will be biased (insurable permanent shocks will be estimated as if they were transitory shocks). Therefore, we want to extend the model to take this kind of risk-sharing between consumers into account.

Consider the opposite polar case, where markets are complete. Then, consumers can fully insure against idiosyncratic shocks, so that a shock with cross-sectional mean zero will have no effect on consumption at all. However, in our specification of the income process, there are two types of idiosyncrasies: true idiosyncratic shocks with

10

cross-sectional mean zero, and idiosyncratic responses to aggregate shocks. When markets are complete, consumers can insure against the idiosyncratic components of the response of their income to an aggregate shock, but unlike in the case of true idiosyncratic shocks, insurance is not 'free': some individuals face more risk (are more sensitive to aggregate shocks) and have to compensate other consumers for participating in a risk-sharing scheme.[7]

To formalize this, consider a standard Arrow-Debreu economy where we get the familiar result that marginal utilities of consumption move in lock-step and every individual consumes a fixed fraction of aggregate income. With CRRA preferences:

$$\frac{u'\left(c_{jt}\right)}{u'\left(c_{it}\right)} = \left(\frac{c_{jt}}{c_{it}}\right)^{-\theta} = \frac{\mu_j}{\mu_i} \Rightarrow c_{it} = \mu_i \overline{c}_t = \mu_i \overline{y}_t$$

where the last step (aggregate consumption equals aggregate income) follows from the assumption that the consumption good is non-storable (i.e. no self-insurance is possible; this assumption is innocuous since we are only considering permanent shocks to income). In general, the fraction μ_i of total income that goes to consumer i will depend on her sensitivity to aggregate income shocks δ_i, but if we work in first-differences of the log of consumption, the above conveniently simplifies to

$$\frac{c_{it}}{c_{it-1}} = \frac{\overline{y}_t}{\overline{y}_{t-1}} \Rightarrow \Delta \log c_{it} = \Delta \log \overline{y}_t$$

so in percentage terms, there is still full insurance. The only difference with the standard model is that a consumer's initial endowment depends not only on initial wealth, but also on 'initial risk', her sensitivity to the business cycle.

With our specification of the income process we saw that with incomplete insurance markets, consumption responded one-to-one to permanent shocks to income $\Delta \log c_{it} =$

---

[7]Consider the following very stylized example. Two individuals get a fixed income of 2, plus a part that depends on the business cycle. Aggregate shocks v_t can be either good $(+1)$ or bad (-1). Consumer one's income is very sensitive to the business cycle $\delta_1 = 2$, but consumer two's income does not respond at all to aggregate shocks $\delta_2 = 0$. Risk-sharing takes the usual form, where both consumers commit to transfer their income to a mutual fund and get an ex ante agreed upon fraction from the amount in the mutual fund, regardless of the state of the economy. In this example, the fund will contain 4 units of income in a boom and 2 units in a recession. Although both consumers have identical expected incomes of 2 units, it is clear that a $50 - 50$ division of the fund will not work, as consumer 1 will have no incentive to participate in the fund: if she does not participate she gets an income of 2 in every period; if she participates, she gets sometimes 1 and sometimes 3 which yields lower utility given that she is risk-averse. Because consumer 2 ex ante faces more risk than consumer 1, she needs to compensate consumer 1 for participating in the fund. A possible solution would be that consumer one gets three quarters of the total amount in the fund, and consumer 2 gets one quarter. In this case consumer one will prefer to have an income of 3 or 1.5 to having 2 in every period because the expected value of her income stream increases, whereas consumer two will prefer to have 1 or 0.5 over having 2 or 0 because she will be able to insure some of the risk she faces.

$n_{it} + \delta_i v_t$. Here we derived that in the complete markets model we get $\Delta \log c_{it} = \Delta \log \overline{y}_t = v_t$. Therefore, we define the degree of risk sharing or 'insurance markets completeness' ϕ as the weight in a linear combination of both solutions[8]

$$\Delta \log c_{it} = (1 - \phi)(n_{it} + \delta_i v_t) + \phi v_t = (1 - \phi) n_{it} + \tilde{\delta}_i v_t$$

where $\tilde{\delta}_i = (1 - \phi) \delta_i + \phi$. If $\phi = 1$ then insurance markets are complete. All risk except permanent aggregate shocks are perfectly insured. If $\phi = 0$ insurance markets are non-existing, and the only channel for consumption smoothing is self-insurance.

Clearly, this is a very rough proxy for risk-sharing through insurance markets. In particular, we are assuming that the degree of risk-sharing is the same for all consumers. It may be more realistic to assume that ϕ varies across consumers, since only a fraction of consumers participates on the market for risky assets. For reasons of tractability we consider only a fixed ϕ, although in appendix B we consider a simple ad-hoc adjustment for credit constraints and limited access to insurance markets. More fundamentally, ϕ is an ad-hoc adjustment to the reduced form model and care should be taken in giving it any structural interpretation. It is an ordinal measure for market completeness, but the units are unclear. If ϕ is high, say 0.75, this means only that consumers manage to insure a large part of idiosyncratic permanent shocks. It does not imply anything about the number of assets they use to achieve this insurance, nor does it give any insight into which shocks are insured.[9]

3   Data: CEX 1980-2000

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, conducted on an annual basis since 1980, follows households across five quarters so that by matching households across the quarterly data files, we can calculate first differences in income and consumption. Notice that although the CEX has only a very short panel dimension and the income data are not of the best quality, the CEX is the only US dataset that has acceptable consumption as well as income data, which we need in order to be able to calculate the covariance between income and consumption.[10] In addition, the CEX contains good data on assets, which allows flexibility for further research.

[8]Luigi Pistaferri, Richard Blundell and Ian Preston are working with a similar model, which they presented at the SED 2002 meeting. We were not able to get the paper.

[9]Because risk sharing is identified from the variances, insurance against some large shocks (e.g. unemployment risk) is empirically indistinguishable from insurance against many small shocks or partial insurance of all shocks.

[10]The 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).

In appendix C we discuss the dataset and our procedure to clean the data for seasonality, age effects, attrition bias and family composition. The final dataset contains 60,000 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, as illustrated in table 1 (see the appendix for a more extensive discussion). These households are divided over 192 cohort-year cells with a median cell size of 247 households, see table 2.

In this section, we first discuss how measurement error in the data will and will not bias our estimates. Then we present some evidence from the raw data on the evolution of income and consumption inequality over the sample period.

3.1 Measurement error

Clearly, both income and consumption are measured with error. This problem is exacerbated by taking first differences if the measurement error is less autocorrelated than the signal. Notice that this is the case by definition, because measurement error cannot have a unit root, whereas consumption does according to the PIH.[11] However, because we work with a secondary dataset and are only really interested in the means and variances of changes in consumption and income, there is an easy way to control for measurement error, if we are willing to make some assumptions on the error structure.

First, consider the moment conditions for the mean changes in income and consumption within a cohort. Under the reasonable assumption that measurement error is uncorrelated across individuals, taking means should virtually eliminate the measurement error as it is divided by the number of observations in the cohort-time cell.

Next, consider the variances and covariance of Δy and Δc. Now, assuming the measurement error is uncorrelated with true changes in income and consumption, the only thing that changes is that an additive term is added to the moment conditions, representing the cross-sectional variance of the measurement error. The size of this additive term will depend on the autocorrelation in the measurement error as well; the more highly autocorrelated, the more measurement error is eliminated by taking first

---

[11] A reinterview program for the CEX is available (a subsample of approximately 6 percent of households are reinterviewed by a supervisor on an ongoing basis), which in principle could be used to assess the amount of measurement error.

differences.[12]

Finally, when we calculate first differences in the variances and covariance of y and c, the variance of the measurement error will drop out, assuming the cross-sectional variance of the measurement error is constant over time. Therefore, we will be able to identify this variance from the difference in the constant in the moment conditions for the variances of the changes and the changes in the variances of income and consumption.

3.2   Income and consumption inequality

Figure 1 shows consumption and income inequality for nine 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 C, 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. Notice that income is on the left-hand axis, and consumption on the right hand axis. 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 it should be 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 (with 2 standard error bands) 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.

To isolate aggregate inequality, we plotted the variance of log consumption and income for the whole sample over time. Potentially demographic changes cloud the picture of the relevant measure of inequality, but the graphs for aggregate and average within

---

[12]The model we have in mind is the following. If measurement error, say in consumption, is additive in logs and uncorrelated with the signal,

$$c_{it} = c_{it}^* + \varepsilon_{it}^c \quad \text{where} : \varepsilon_{it}^c = \rho_c \varepsilon_{it-1}^c + \eta_{it}^c$$

where the η's are mean zero, uncorrelated in the cross-section and over time, and uncorrelated with c_{it}^*. Then

$$Var\left(\Delta c_{it}\right) = Var\left(\Delta c_{it}^*\right) + \frac{2}{(1+\rho_c)} Var\left(\eta_{it}^c\right)$$

cohort inequality turn out to be very similar. The graphs for aggregate inequality are comparable to Krueger and Perri (2001). Like them, we find that the increase in income inequality in the eighties is less pronounced in the CEX data then in other datasets. The graph for aggregate consumption also shows Krueger and Perri's main empirical finding (and the basic observation they use to justify their paper), that consumption inequality seems not to have increased at all, and may even have decreased.

4   Estimation

Looking at the raw data is informative. It shows that our data are in line with commonly accepted stylized facts about the development of inequality. The graphs also show that the data are very noisy, due to the relatively small number of households in a cohort-year cell, and this is something the estimation technique should take into account. In the next subsections we discuss two different procedures we use to estimate the model: GMM and a likelihood based, Bayesian procedure that treats the aggregate shocks and time-varying variances of the idiosyncratic shocks as unobservable components.[13] We will argue that either estimation method is more efficient and more formal than a decomposition using an exactly identified model, and that the Bayesian estimation procedure performs slightly better to distinguish between signal and noise in the data and has more robust convergence properties.

4.1   Generalized method of moments

We use all moment conditions derived in appendix A.[14] The error terms in those equations are assumed to have mean zero, but no other assumptions are imposed. In particular, we do not make any assumptions on the correlation between the error term and the time varying variances of the idiosyncratic shocks, so we basically treat these time varying variances as time-specific fixed effects or time dummies. The estimates are one step asymptotically efficient estimates, using the optimal weighting matrix calculated manually from the moment conditions. This is possible because all conditions are linear in the data which are the only source of variation.

For reasons of convergence, we estimate the model in two stages. First, the aggregate shocks e_t and v_t and the cohort specific mean trends α_j and b_j are estimated from the means of changes in income and consumption. In the second stage all other parameters are estimated from the remaining moment conditions, treating the aggregate shocks as known. The (asymptotic) standard errors of these parameters do not take the uncer-

---

[13]Matlab programs that implement both methods are available from the authors on request.

[14]Instead of the variance of the change in consumption we use $\Delta Var\left(\log c_{it}\right) - Cov\left(\Delta \log c_{it}, \log c_{it-1}\right) = Var\left(\Delta \log c_{it}\right)$ because of the timing problem in the CEX (see section C.1 in the data appendix for details).

tainty in the estimates of the aggregate shocks into account. Because e_0 is not identified, we use an initial condition setting e_0 equal to zero, its unconditional mean.[15] The estimates for the aggregate shocks and the time varying variances of the idiosyncratic shocks are quite robust, but the scalar parameters are very imprecisely estimated.

4.2 Unobservable components model

Differently from the GMM procedure, in the likelihood based estimation technique we treat v_t, e_t, $Var_t(n)$ and $Var_t(u)$ as unobservable states. The aggregate shocks v_t and e_t are by assumption i.i.d. processes, while in order to specify a law of motion for $Var_t(n)$ and $Var_t(u)$ we make the additional assumptions that $Var_t(n)$ and $Var_t(u)$ follow two independent random walk processes:

$$
\begin{aligned}
Var_t(n) &= Var_{t-1}(n) + w_{nt} \\
Var_t(u) &= Var_{t-1}(u) + w_{ut}.
\end{aligned}
$$

Of course, $Var_t(n)$ and $Var_t(u)$ cannot become negative and, at first sight, the random walk assumption may seem inadequate. However, because the time dimension of the sample is short (only twenty time periods), the random walk can be thought as a (good) first order approximation of a more complicated and theoretically justifiable process for the two variances.[16]

Notice that, dealing with unobservable components, a Bayesian approach is the natural one. Even more so in a panel context, where the distinction between parameters and shocks is less clear than in other situations. Finally, and particularly important in this case, as we will see 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 is easier to calculate standard errors which are correct for finite sample inference, instead of relying on asymptotic distribution theory.

Because we use flat or uninformative priors, the Bayesian procedure has a likelihood interpretation. With flat priors the posterior modes of the parameters would exactly correspond to the maximum likelihood estimates. For reasons of robustness, we use posterior means as is common in the literature. With well-behaved posterior distributions the difference is negligible.

All the shocks of the model are assumed to be jointly normal with zero mean and a block diagonal variance-covariance matrix. The only covariance block that is left un-

---

[15]We also tried setting $e_T = 0$ and $\sum_{t=0}^{T} e_t = 0$ and the results are not sensitive to the exact identifying assumption used.

[16]The reason why the random walk is preferred to a more general AR specification derives from the fact that we want to capture low frequency time variation. Furthermore, with a sample of only twenty time periods, a random walk would not be easily distinguishable from many AR processes. We also estimated the model assuming the variances are i.i.d. and the results were very similar.

constrained is the covariance between the innovations to the idiosyncratic permanent variance and the business cycle, represented by the aggregate shocks. Again, the normality assumption for the disturbances in the moment conditions for the variances of changes in consumption and income should be interpreted only as an approximation around the peak of the true distribution. We could restrict the variances of the idiosyncratic shocks to be positive, but because the point estimates turn out to be positive, we conclude that the normality assumption does not affect the results.

The objective here is providing a procedure to evaluate the posterior distribution of the parameters of the model. The parameters of interest are the unobservable states, v_t, e_t, $Var_t(u)$ and $Var_t(n)$ and the so called hyperparameters, which are divided in three blocks: Σ contains the variances of the aggregate shocks and all error terms in the moment conditions, θ is a vector with the variances and covariances of interest

$$\begin{aligned} \theta &= [Var\left(\alpha\right), Var\left(b\right), Cov(b, c_0), Cov(\alpha, y_0), Var\left(\gamma\right), Var(\delta), Cov\left(\gamma, \delta\right), \\ &\quad Cov\left(\gamma, y_0\right), Cov\left(\delta, y_0\right), Cov(\gamma, c_0), Cov(\delta, c_0)]', \end{aligned}$$

and ϕ is a scalar parameter representing the degree of risk sharing as defined above.

The parameter space is very high dimensional and analytical maximization of the likelihood or posterior is not feasible. Numerical methods and in particular Gibbs sampling are needed to tackle the problem. Gibbs sampling is a particular variant of Markov Chain Monte Carlo 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 five steps, drawing in turn the aggregate shocks, the hyperparameters θ, the time varying variances of the idiosyncratic shocks, the hyperparameter ϕ and the hyperparameters Σ, in each step conditional on the observed data and the other parameters. A more detailed description of the sampling procedure is provided in appendix D.

Our simulations are based on 6,000 iterations of the Gibbs sampler, discarding the first 1,000 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.

4.3 Estimation versus decomposition

In figure 3 we have graphed the actual evolution of average within cohort income and consumption inequality and the fitted values of the model estimated by GMM and as an unobservable components model. At first sight it seems that neither model manages to capture the evolution of inequality, particularly income inequality, well. However, we use a set of moment conditions much wider than just the change in the variance of income and consumption, which are plotted in this figure. The fact that the estimates do not capture all of the large peak in income inequality from 1983 to 1988 for instance,

is because this peak is less pronounced in some of the other moment conditions. We attribute this to the fact that the data are very noisy. Since we work with variances these series are sensitive to outliers, and because the cohort-year cells are fairly small, this introduces large swings in inequality.

The GMM and likelihood based estimates for the aggregate shocks and the time varying variances of the idiosyncratic shocks are very similar. In the GMM estimates, we restrict $Cov(b, c_0)$ and $Var(b)$ to be zero and set the degree of risk sharing ϕ equal to its estimated value from the unobservable components model. This was necessary in order to make the GMM program converge. The estimated values for $Cov(b, c_0)$ and $Var(b)$ in the unobservable components model explain almost all of the difference in the predicted values for both models.

The small remaining differences are due to the estimation procedure. GMM minimizes the variance of the error term in the moment conditions, and will thus by definition try to minimize the distance of the model predicted values from the actual income and consumption inequality. A likelihood based estimator however, sets the variance of the noise and the innovations in the unobservable states in order to maximize the likelihood of the model, without a direct attempt to minimize the variance of the noise. Of course asymptotically GMM is efficient and equivalent to maximum likelihood. However, our sample is fairly small compared to the number of parameters we are trying to estimate. And since we know that the data are noisy, we feel the estimation procedure should take this noise into account explicitly.

To further support our view that the model predicted values are a better description of the actual evolution of inequality than the raw data series for the variances of income and consumption, in figure 4 we plotted again the raw data on income inequality and the fitted values from the unobservable component model, but now we added a robust estimator for the variance of income, as well as a comparable time series from the Current Population Survey.[17] The robust estimator is calculated from the median absolute deviation from the median and translated back into a variance assuming normality of log income.[18] Both series have been shifted upward by a constant. As is clear from the graph, the model predicted values are very close to the robust series from the CEX, and also much closer to the CPS series than the (non-robust) raw CEX data. We should stress that we did not use either series in the estimation procedure, so the fit is quite remarkable, and we conclude that the estimation procedure manages quite well

---

[17]We use the CPS March supplement (US Department of Commerce, Bureau of the Census and US Department of Labor, Bureau of Labor Statistics 2001), which has detailed income data comparable to the CEX . These data were treated in the same way as the CEX data to control for age effects, family composition and attrition bias. We use the household records and merge these with the person records for the householder and spouse. The resulting dataset has a much larger sample size than the CEX, with typically about 3,000 households in a cohort-year cell.

[18]For a normal distribution the median absolute deviation from the median equals 0.6745 times the standard deviation (Huber 1981).

to distinguish noise from signal and the model provides a good description of the joint evolution of income and consumption.

To summarize, our estimation procedure differs fundamentally from the decomposition that has so far been used in the literature in two ways. First, we use all the available information in the joint evolution of income and consumption, by using as many moment conditions as we can calculate from the two series. This provides a substantial efficiency gain and makes the estimates less sensitive to outliers and data problems. Secondly, we propose a formal, likelihood based estimation method that has much better convergence properties than the more commonly used GMM. It seems this method could be useful more generally for high dimensional minimum distance problems.

4.4   Why we need data on both income and consumption?

In this paper we use the joint evolution of consumption and income inequality to estimate the components of income risk. A natural question to ask is why we use consumption data at all. In fact there is a literature that tries to distinguish transitory and permanent components to earnings variability, using data on income only.

Gottschalk and Moffitt (1995) and Moffitt and Gottschalk (2002) exploit the long panel dimension of the PSID. They estimate the permanent component of earnings by a 5 year moving average, and the transitory component as the deviation of income from that smoothed series. The identifying assumption is that the 5 year average of the transitory shocks is very close to its unconditional mean of zero. Of course it is not possible to reproduce Gottschalk and Moffitt's methodology using the CEX, because only two annual observations for income are observed for each individual. However, it is still possible to identify the variance of transitory shocks from income data only, using the autocovariance of income (or -equivalently- the covariance between the $\Delta \log y_{it}$ and $\log y_{it-1}$, see appendix A for the exact moment condition). We do in fact use this information, and therefore the components of income risk in our model can be identified off income data alone.

On the other extreme, Blundell and Preston (1998) do not use the time series properties of income at all, but use consumption data to identify transitory from permanent shocks. They define a permanent shock as a change in income that does not translate into a change in consumption. Potentially, this method could be more efficient and, given that most of the interest in decomposing income risk into a transitory and a permanent part has to do with consumption insurance, it has the advantage that one takes the relation between consumption and income into account explicitly. The price to pay is that the conclusions can only be as general as the model specification. In particular, Blundell and Preston impose that consumption can be perfectly insured against transitory shocks, and cannot be insured at all against permanent shocks to income.

In this paper we try to take the best of both worlds. We use both the information

from the time series behavior (the autocovariance) of income, and all information available from the evolution of consumption as well as the comovement of consumption with income. Thus our model is over-identified and we are able to test the over-identifying restrictions. (In particular, we can drop all moment conditions involving consumption.) When we are satisfied that the results do not depend on the specific subset of moment conditions used, we can use the full set to get more efficient estimates. This is particularly important because the data are very noisy.

Compared to Blundell and Preston, we not only use a much larger set of moment conditions, but because of the extra degrees of freedom we are also able to specify a more flexible model of consumption behavior. In particular, as discussed in section 2.5, we allow for the possibility that consumer can partially insure against idiosyncratic permanent shocks. This parameter would not be identified using either income data alone, or without using the autocovariance properties of income.

5 Inequality over the business cycle

In table 3 we present scalar parameter estimates of the model for both estimation techniques. The first seven rows of the table are parameters that determine the distributional impact of aggregate shocks; we will return to these estimates below. The second and third block report estimates for the cross-sectional variance of the idiosyncratic trends in income and consumption,[19] and for the measurement error in different equations. The measurement error in the equations for $\Delta Var (\log y)$, $\Delta Var (\log c)$ and $\Delta Cov (\log y, \log c)$ is set to zero, which is a reasonable assumption if the variance of the measurement error in income and consumption is time invariant.[20] As expected, the measurement error in the variances of $\Delta \log c$ is substantial, but there seems to be little cross-sectional correlation in the measurement error in income and consumption since the estimate of the measurement error in $Cov (\Delta \log y, \Delta \log c)$ is small.

5.1 How incomplete are insurance markets?

The last row in table 3 presents an estimate for ϕ, the degree of risk sharing against permanent idiosyncratic shocks. If $\phi = 0$ then insurance markets are non-existing, whereas $\phi = 1$ represents the complete markets solution. The estimate for ϕ is fairly precise, and seems to suggest that around 75% of permanent idiosyncratic risk can be

---

[19] Although we could easily allow for these variances to differ across cohorts, we treat them as constants for efficiency. This does not affect the results since the variances are very similar for different cohorts.

[20] This is an identifying assumption without which we would not be able to estimate the mean level of the variance of the permanent shocks and the mean change in the variance of the transitory shocks. The mean level of the transitory shocks remains unidentified but is of no interest for the evolution of inequality.

insured through risk sharing. This estimate represents a substantial deviation from the standard incomplete markets model that completely shuts down insurance markets.[21]

Allowing for partial risk sharing is important in reconciling the joint evolution of income and consumption inequality. The variance of transitory shocks is identified either from the autocovariance of income, or from the difference in the change in the variance of income and the change in the variance of consumption. Without ϕ, these estimates conflict with each other. If we introduce ϕ, the variance of the transitory shocks in the expression for the change in the variance of consumption is multiplied by a factor $(1 - \phi)^2$, allowing the estimate for the variance of the permanent shocks to increase without violating the moment condition for the change in the variance of income.

It is possible that the parameter ϕ captures more than only partial risk sharing. In particular, if a small fraction of consumers is credit constrained, this will also imply a quantitatively important difference in the change in the variance of consumption and income inequality. In appendix B we make a case why we think it is reasonable to assume that some consumers are credit constrained or behave as if they are because they are buffer-stock consumers with low asset levels on the steep part of the consumption policy rule. We are currently working on an extension to this paper that will take this complication into account. At this point we are reluctant to give any structural interpretation to the parameter estimate of ϕ.

5.2 Sources of risk and the distributional impact of aggregate shocks

The estimates for the variance of the idiosyncratic responses to aggregate shocks γ_i and δ_i in table 3 seem to indicate that the impact of aggregate shocks on inequality is large. The point estimate for the standard deviation of δ_i is about 3, suggesting that δ_i ranges from -5 to $+7$ for 95% of consumers. However, the reason these parameters are estimated to be large, is that they are identified off the correlation between functions of the aggregate shocks v_t^2, $(\Delta e_t)^2$ and $v_t \Delta e_t$ with the variances and covariance of consumption and income. Because this correlation is small and because v_t^2, $(\Delta e_t)^2$ and $v_t \Delta e_t$ are (very) small numbers, the point estimates for $Var(\gamma)$, $Var(\delta)$ and $Cov(\gamma, \delta)$ are high and imprecise.

The effect of the business cycle on inequality depends on the variances of the idiosyncratic responses as well as on the size of the aggregate shocks. These scalar estimates alone then, give little insight into the impact of aggregate shocks on inequality. Therefore, we have plotted the evolution of income inequality, if the aggregate shocks would

---

[21]This is particularly true when one realizes that the parameter has been estimated using household consumption data, and therefore excludes risk sharing through living together in the same household, something that has been shown to be quantitatively important compared to risk sharing through financial markets.

have been zero in every time period. This counterfactual exercise is presented in figure 5. The upper panel in that graph shows the predicted values for income inequality for the full model, as well as for the model without aggregate shocks. The lower panel graphs the difference between the two lines with one standard error bands. These standard errors take into account all uncertainty related to the aggregate shocks (i.e. in the estimates of v_t and e_t as well as in the variances and covariance of γ_i and δ_i). This lower panel can be interpreted as the evolution of inequality if there would have been only aggregate shocks to the economy. Clearly, the effect is negligible, strongly suggesting that virtually all of inequality is driven by idiosyncratic risk.

In figures 6 and 7 we have similarly isolated the effects on inequality of the idiosyncratic transitory and permanent shocks. It is clear that permanent risk is responsible for the within cohort rise in inequality (the age effects) as expected, as well as for most of the rise in inequality during the 1980s. Transitory inequality increased sharply in the early 1980s but then gradually declined again, and by the end of the sample period was back at the 1980 level or even lower. Qualitatively, these conclusions correspond well to Moffitt and Gottschalk (2002). Compared to Blundell and Preston (1998), who use UK data, we find a much less important contribution of transitory shocks to the rise in income inequality in the 1980s. This difference is largely driven by the fact that we allow for partial insurance against idiosyncratic permanent shocks.

5.3   Comovement risk with the business cycle

In the previous subsection we argued that the direct distributional impact of aggregate shocks on inequality is negligible. Nevertheless it is possible that risk is (counter)cyclical, if the variance of the idiosyncratic shocks is correlated with aggregate shocks. There has been much speculation in the literature about such a correlation. Probably best known, is Constantinides and Duffie's (1996) explanation for the equity premium puzzle. Constantinides and Duffie argue that if stocks offer high return in periods when consumption risk is low, they would provide even worse insurance against such consumption risk than is suggested by the standard consumption CAPM model. Thus, a countercyclical variance of permanent shocks could explain why risky assets offer such a high risk premium.

To answer this question, we have plotted our estimates for the aggregate permanent shocks and the variance of the idiosyncratic permanent shocks in figure 8. The shaded areas are recessions according to the NBER business cycle dating committee (1980:1-1980:3, 1981:3-1982:4 and 1990:3-1991:1). Not surprising given the persistence in aggregate fluctuations, it seems that business cycle is primarily driven by shocks to permanent income and our estimates for the shocks correspond fairly well to the NBER dates. During the 1990s, periods of high volatility in the permanent idiosyncratic shocks seem to correspond to periods of low realizations of the aggregate shocks. However, in

the 1980s exactly the opposite is true. Not surprisingly, over the whole sample period, the correlation is very close to zero. We estimate the correlation coefficient to be 0.065 with a standard error of 0.211 and conclude that our data provide no evidence for any comovement of consumption risk with the business cycle.

ADD: compare results to Cogley (2002), Heaton and Lucas (1996), Ramey and Ramey (1995), Brav et.al. (1999) and Storesletten et.al. (2000b)

6   Concluding remarks

In this paper we used CEX data on income and consumption inequality to estimate four components of income risk: idiosyncratic versus aggregate, and transitory versus permanent shocks to income. We extended the methodology set out in Blundell and Preston (1998) to take aggregate uncertainty into account, and to allow for risk sharing against idiosyncratic permanent shocks. We use a set of moment conditions that summarize all available information in the joint evolution of consumption and income, conditional on the constraint that the CEX has a very short panel dimension. This model performs well in capturing the actual evolution of inequality from the noisy data. We also proposed a formal likelihood based estimation technique that is more robust and solves many of the practical difficulties associated with GMM in problems with a high dimensional parameter space.

Our results indicate that inequality is largely driven by idiosyncratic risk, the impact of aggregate shocks on the cross-sectional variance of income being negligibly small. The estimates of the relative contribution of transitory and permanent risk correspond well to the findings of Moffitt and Gottschalk (2002), despite their very different methodology. Compared to Blundell and Preston (1998) we find a much less important contribution of transitory shocks to income inequality. Finally we find no evidence for any correlation between consumption risk and the business cycle, which casts doubt on Constantinides and Duffie's (1996) explanation for the equity premium puzzle.

In future research we plan to extend the methodology used here, to take into account that some consumers are credit constrained or behave as buffer-stock consumers with low asset levels. In addition we plan to explore how the relative importance of the components of income risk interacts with observable consumer characteristics like education. Particularly relevant for an empirical evaluation of the Constantinides and Duffie result would be to re-do the analysis separately for consumers that participate in the market for risky financial assets and those that do not.

A Moment conditions

With our model for the joint movement of income and consumption within a cohort, we can calculate cross-sectional means and variances for each cohort $j \in \{1, 2, ..., J\}$. First, taking averages over individuals within in each cohort, we get (dropping the logs for convenience)

$$\Delta y_{jt} = \alpha_j + \Delta e_t + v_t + \varepsilon_{jt}^y \tag{4}$$

$$\Delta c_{jt} = b_j + v_t + \varepsilon_{jt}^c \tag{5}$$

where the epsilons represent the difference between the sample and population means and are therefore zero mean.[22] These moment conditions identify the aggregate shocks e_t and v_t as well as the cohort specific fixed effects α_j and b_j.

Similarly, we can obtain $3JT$ moment conditions from the variances and the covariance between Δy_{ijt} and Δc_{ijt} for each cohort at each time period

$$\begin{aligned} Var_{jt}(\Delta y) &= Var_j(\alpha) + Var_t(u) + Var_{t-1}(u) + Var_t(n) \\ &\quad + (\Delta e_t)^2 Var(\gamma) + v_t^2 Var(\delta) + 2v_t \Delta e_t Cov(\gamma, \delta) + \eta_{jt}^y \end{aligned} \tag{6}$$

$$Var_{jt}(\Delta c) = Var_j(b) + (1-\phi)^2 Var_t(n) + (1-\phi)^2 v_t^2 Var(\delta) + \eta_{jt}^c \tag{7}$$

$$\begin{aligned} Cov_{jt}(\Delta y, \Delta c) &= Cov_j(\alpha, b) + (1-\phi) Var_t(n) \\ &\quad + (1-\phi)\left[v_t^2 Var(\delta) + v_t \Delta e_t Cov(\gamma, \delta) \right] + \eta_{jt}^{yc} \end{aligned} \tag{8}$$

From these expressions, we can identify the time varying variances $Var_t(u)$ and $Var_t(n)$ and the distributional parameters $Var(\gamma)$, $Var(\delta)$ and $Cov(\gamma, \delta)$, as well as the cohort fixed effects $Var_j(\alpha)$, $Var_j(b)$ and $Cov_j(\alpha, b)$. The easiest way to interpret these moment conditions is to consider the aggregate shocks as known from the moment conditions above, and see these as a SUR regression with time specific fixed effects $Var_t(u)$ and $Var_t(n)$.

Additional moment conditions can be derived from using the differences of the variances and covariance in levels. The changes in the variance of income, the variance of consumption, and their covariance, are given by

$$\begin{aligned} \Delta Var_{jt}(y) &= Var_{jt}(\Delta y) - 2Var_{t-1}(u) + 2Cov(\alpha_i, y_{it-1}) \\ &\quad + 2\Delta e_t Cov(\gamma_i, y_{it-1}) + 2v_t Cov(\delta_i, y_{it-1}) + \xi_{jt}^y \end{aligned} \tag{9}$$

$$\Delta Var_{jt}(c) = Var_{jt}(\Delta c) + 2Cov(b_i, c_{it-1}) + 2(1-\phi) v_t Cov(\delta_i, c_{it-1}) + \xi_{jt}^c \tag{10}$$

$$\begin{aligned} \Delta Cov_{jt}(y, c) &= Cov_{jt}(\Delta y, \Delta c) + \Delta e_t Cov(\gamma_i, c_{it-1}) + v_t Cov(\delta_i, c_{it-1}) \\ &\quad + (1-\phi) v_t Cov(\delta_i, y_{it-1}) + \xi_{jt}^{yc} \end{aligned} \tag{11}$$

---

[22] These error terms also include measurement error. See the discussion in the main text.

where

$$
\begin{aligned}
Cov\left(\alpha_i, y_{it-1}\right) &= Cov\left(\alpha, y_0\right) + (t-1)\,Var\left(\alpha\right) \\
Cov\left(b_i, c_{it-1}\right) &= Cov\left(b, c_0\right) + (t-1)\,Var\left(b\right) \\
Cov\left(\gamma_i, y_{it-1}\right) &= Cov\left(\gamma, y_0\right) + \left(\sum_{s=1}^{t-1}\Delta e_s\right)Var\left(\gamma\right) + \left(\sum_{s=1}^{t-1}v_s\right)Cov\left(\gamma, \delta\right) \\
Cov\left(\delta_i, y_{it-1}\right) &= Cov\left(\delta, y_0\right) + \left(\sum_{s=1}^{t-1}v_s\right)Var\left(\delta\right) + \left(\sum_{s=1}^{t-1}\Delta e_s\right)Cov\left(\gamma, \delta\right) \\
Cov\left(\gamma_i, c_{it-1}\right) &= Cov\left(\gamma, c_0\right) + (1-\phi)\left(\sum_{s=1}^{t-1}v_s\right)Cov\left(\gamma, \delta\right) \\
Cov\left(\delta_i, c_{it-1}\right) &= Cov\left(\delta, c_0\right) + (1-\phi)\left(\sum_{s=1}^{t-1}v_s\right)Var\left(\delta\right).
\end{aligned}
$$

These $3JT$ moment conditions introduce only 5 extra parameters: $Cov\left(\alpha, y_0\right)$, $Cov\left(\gamma, y_0\right)$, $Cov\left(\delta, y_0\right)$, $Cov\left(b, c_0\right)$, $Cov\left(\gamma, c_0\right)$ and $Cov\left(\delta, c_0\right)$, so they substantially increase efficiency. Moreover, these parameters are all interesting. They describe to what extent sensitivity of income and consumption to the business is correlated with the initial distribution of income and consumption.

Finally, one more set of moment conditions can be derived from the change in the autocovariance of income or, equivalently, from the covariance of changes in income with lagged income.[23]

$$
\begin{aligned}
Cov\left(\Delta y_{it}, y_{it-1}\right) &= \tfrac{1}{2}\left[Cov_{jt}\left(y_{it}, y_{it-1}\right) - Var_{jt-1}\left(y\right)\right] \qquad (12) \\
&= -Var_t\left(u\right) + Cov\left(\alpha_i, y_{it-1}\right) + v_t Cov\left(\delta_i, y_{it-1}\right) \\
&\quad + \Delta e_t Cov\left(\gamma_i, y_{it-1}\right) + \zeta_{jt}^{y}
\end{aligned}
$$

Notice that Blundell and Preston's (1998) estimator for the variance of transitory and permanent shocks is a special case of the moment conditions for the change in the variances of income and consumption (or the change in the variance of income and the covariance of income and consumption). Setting the aggregate shocks (or the variances and covariance of the idiosyncratic responses) equal to zero, and with the additional assumptions $\alpha_i = \alpha$ and $b_i = b$ constant across individuals, it is easy to see that

$$
\Delta Var_t\left(y\right) - \Delta Var_t\left(c\right) \simeq \Delta Var_t\left(y\right) - \Delta Cov_t\left(c, y\right) \simeq \Delta Var_t\left(u\right)
$$

which is Blundell and Preston's difference in difference estimator for the change in the variance of the transitory shocks. The variance of the permanent shocks can then be retrieved by substracting these estimates from the change in the variance of income.

---

[23] We cannot calculate the autocovariance in consumption, nor the autocovariances of changes in income and consumption, because we have only one annual change in consumption for each individual. See the description of the data in appendix C.
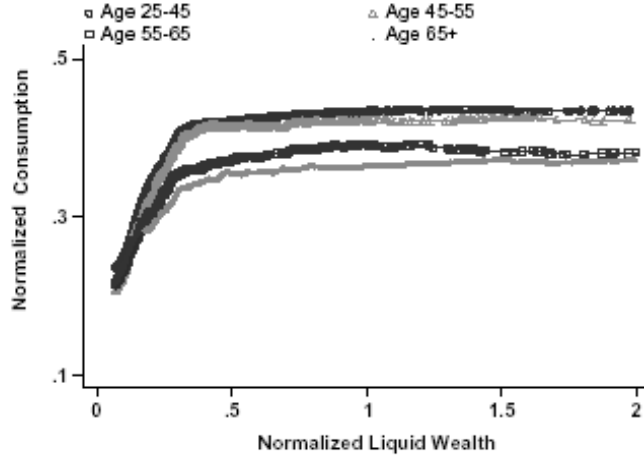
Figure 1: The consumption policy rule, non-parametric estimate using CEX data on consumption and assets, taken from Gourinchas and Parker (2001).

B   Credit constraints and precautionary saving

The model we described so far allows for a time-invariant amount of precautionary saving. However, the model 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 consumes all of current income (In an old paper Campbell and Mankiw 1990 show that a model where about half the consumers is forward looking and the other half 'hand-to-mouth' describes the data better than the PIH). Clearly λ can be interpreted as the fraction of consumers that are (perfectly) credit constrained. We argue however, that it also proxies precautionary savings behavior relatively well.

Precautionary saving is theoretically closely related to liquidity constraints, and empirically virtually indistinguishable (see Carroll 2001a). Figure 1 shows a non-parametric estimate of the consumption policy rule, taken from Gourinchas and Parker (2001). Approximately, the figure shows 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 permanent income hypothesis. If wealth is below $\tilde{A}$, the marginal propensity to consume out of extra cash-on-hand is close to one.

To formalize the role of λ we write the estimation model as follows.[24]

$$\Delta y_{it} = \alpha_i + \Delta u_{it} + \gamma_i \Delta e_t + n_{it} + \delta_i v_t$$

$$\Delta c_{it} = \begin{cases} \Delta y_{it} & \text{for a fraction } \lambda \text{ of the individuals} \\ \Delta c_{it}^* = b_i + (1 - \phi) n_{it} + \tilde{\delta}_i v_t & \text{for the rest} \end{cases}$$

where Δc_{it}^* is the change in consumption for consumers that behave as if they follow the PIH (extended for partial risk-sharing).

With this extension of the model derived in the text, the moment conditions for aggregate changes in consumption change to

$$\Delta c_{jt} = \lambda \Delta y_{jt} + (1 - \lambda) \Delta c_{jt}^* = \lambda \alpha_j + (1 - \lambda) b_j + \lambda \Delta e_t + v_t$$

The moment conditions for the mean and variance of changes in income do not change. The variance of Δc_{ijt} is given by the weighted average of the within group variances plus the between group variance, where the groups are credit constrained consumers (a fraction λ) and PIH consumers (a fraction $1 - \lambda$)[25]

$$Var_{jt}(\Delta c) = \lambda Var_{jt}(\Delta y) + (1 - \lambda) Var_{jt}(\Delta c^*) + \lambda (1 - \lambda) \left(\Delta y_{jt} - \Delta c_{jt}^* \right)^2$$

Similarly, the covariance between Δc and Δy is given by

$$Cov_{jt}(\Delta c, \Delta y) = \lambda Var_{jt}(\Delta y) + (1 - \lambda) Cov(\Delta c^*, \Delta y)$$

The moment conditions for the changes in the variance of consumption, and the change in the covariance between income and consumption are slightly more involved. The change in the variance of c changes to

$$\Delta Var_{jt}(c) - 2Cov(\Delta c_{it}, c_{it-1}) = Var_{jt}(\Delta c)$$

where $Var_{jt}(\Delta c)$ as above. The reason we use $\Delta Var_{jt}(c) - 2Cov(\Delta c_{it}, c_{it-1})$ rather than $\Delta Var(c)$ is that we cannot evaluate $Cov(\Delta c_{it}, c_{it-1})$ because we don't have any expression for c_{it-1} from the model. In particular don't want to assume that for a fraction of the consumers $c_{it} = y_{it}$. That would be a much more restrictive assumption than assuming that for a fraction of the consumers $\Delta c_{it} = \Delta y_{it}$. Using a similar argument we get

$$\Delta Cov_{jt}(y, c) - Cov(\Delta y, c_{it-1}) = Cov_{jt}(\Delta y, \Delta c) + Cov(\Delta c, y_{it-1})$$

where $Cov(\Delta y, \Delta c)$ as above and

$$Cov(\Delta c, y_{it-1}) = \lambda Cov(\Delta y, y_{it-1}) + (1 - \lambda) Cov(\Delta c^*, y_{it-1})$$

---

[24] Notice that we assume λ is the same in every cohort and time-invariant. This is potentially a problem; in particular λ is likely to be correlated with the business cycle.

[25] 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.

C   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 are asked about family characteristics and income, and are given information about how to gather their expenditure information. In the second through fifth quarter 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 categories of expenditures corresponding to the categories in the CEX up to a high level of disaggregation. 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. We can therefore calculate three quarterly changes in consumption, but only one annual change in income. This gives rise to the timing problem discussed in more detail below. Because new households enter every month, we can in principle create a monthly time series of consumption and income (although the reference periods would of course be overlapping). Because of changes in the survey design, households cannot be matched with previous months in the first quarter of 1982 and 1986.[26]

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 expenditures is converted to real expenditures using item specific CPI-U price indices (all urban consumers) for the expenditure categories mentioned as components of consumption. To deflate income we use the CPI-U for total expenditures (including expenditures on durables).

We run a series of checks to identify mismatches, including large changes in family composition.[27] In addition, we drop the following households: non-urban households be-

---

[26]Because the BLS makes five quarters of data available each year, the fifth quarter being the first quarter of the next year, in principle some of the individuals could be matched across the changes in the survey design using the previous data release. Because we can handle missing data in our estimation procedures we have not pursued this.

[27]In particular, we suspect a 'mismatch' if the household changes cohort or because any of the 6 categories of family composition (male 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 daugher 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.

cause 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 (it is common in the literature to remove incomplete income reporters, we should probably check the robustness of our results against not removing households with topcoded income); households of which either the reference person or his or her partner is retired (working zero hours and reported reason is retired); and households of which either the reference person or his or her partner is either a student or living in student housing. The latter two categories are dropped because it seems the standard consumption model is not a good description for these households. 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.

Table 1 reports summary statistics for our estimation sample as well as for the full sample of households between 20 and 65 years old. The variables labeled D_* are dummies for the different categories of drops. The most important are incomplete income reporters (18% of the sample) and retired (4% of the sample). In total we drop slightly over 30,000 observations out of 260,000 corresponding to about 8000 out of 50,000 households for the sample period of 20 years. Notice that the number of households corresponds to about 2500 per year instead of the 5000 mentioned earlier because a lot of households have a reference person over 65 years old.

Comparing the two samples on the basis of their personal characteristics, the households in our sample are slightly younger because we dropped households with a retired reference person, and have somewhat higher income 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 reference category is 'other' which includes two adults living together without being married), the number of adults and kids, 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. Notice that consumption is quarterly whereas income refers to the past year. Notice also that real income and consumption are a factor 100 smaller than their nominal counterparts because the BLS normalizes price indices to 100 for the average in the 36 months period 1982-1984. lrinc and lrcons refer to log real income and consumption, SA stands for seasonally adjusted and SAA means adjusted for seasonality as well as age (see the section on preference shifters).

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. Following the literature, we use 5 year cohorts which are labeled by their average age in 1980, i.e. cohort 43 consists of consumers that were aged 41 through 45 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 or quarterly time series, we use annual time periods. Table 2 reports the cell sizes for the cohort-year cells. We eliminate cells with average age below 21 or above 64 (these cells are shaded in the table). In the remainder of the paper we mainly use the dataset with cohort-year cells.

C.1   Timing problem

This data problem deserves particular attention, because it directly influences our estimation procedure. For each household, we observe income only twice.[28] The most information we can retrieve from these two observations, is one annual first difference. Questions on consumption on the other hand, are asked each quarter for four subsequent quarters. We can add up consumption in four quarters to obtain one observation for annual consumption. However, we can only calculate three quarterly first differences. The changes in consumption therefore, do not refer to the same time period as changes in income. We handle these timing differences by multiplying the difference in consumption over three quarters by four thirds. In order not to loose too many observations, if consumption data are missing in one or more quarters, we use a semi-annual or quarterly first difference, multiplied by 2 or 4 respectively.

This procedure works fine for calculating mean changes in consumption, or the covariance between changes in income and consumption. However, when calculating the variance of changes in consumption, the following problem occurs. The variance we are interested in, is the cross-sectional variance of annual changes in consumption. So, letting t denote quarters, what we want is

$$Var\left(c_{it} - c_{it-4}\right) = Var\left(\Delta_1 c_{it} + ... + \Delta_1 c_{it-3}\right)$$

where Δ_1 denotes the quarterly first difference operator. If, as predicted by the PIH, consumption follows a random walk, changes in consumption are independent, so that

$$Var\left(c_{it} - c_{it-4}\right) = Var\left(\Delta_1 c_{it}\right) + ... + Var\left(\Delta_1 c_{it-3}\right) \simeq 4Var\left(\Delta_1 c_{it}\right)$$

However, the data do not allow us to calculate this variance. Instead, we use

$$Var\left(\tfrac{4}{3}\left(c_{it} - c_{it-3}\right)\right) = \tfrac{16}{9}Var\left(\Delta_1 c_{it} + ... + \Delta_1 c_{it-2}\right) \simeq \underbrace{\tfrac{16}{9}3}_{5.33}Var\left(\Delta_1 c_{it}\right)$$

The bias introduced by this procedure is even much larger if there are missing consumption data in one or more quarters. A similar timing problem occurs for $Var\left(\Delta c_{it}, c_{it-1}\right)$.

---

[28]The exception are households for which the number of adults (over 15 years old) or the number of earners changes, because a household member finds or looses a job. In these cases the interview contains the same questions on income that otherwise are asked in the first and fifth interview only.

Luckily, the changes in the variances and covariance of income and consumption do not suffer from the problem because to calculate these we use the time dimension of the synthetic panel, for which we have annual time periods. Therefore, we are able to avoid the problem by using $\Delta Var\left(c\right) - Cov\left(\Delta c_t, c_{t-1}\right)$ as a proxy for $Var\left(\Delta c\right)$. A more formal way to solve this problem would be to re-write the model in terms of quarters, using three moment conditions for quarterly changes in consumption, and three more conditions for the covariances between the quarterly changes in consumption with the annual change in income.

C.2   Preference shifters

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. The first thing to note is that we would have preferred to enter controls in the estimation procedure rather than purging the data of these effects. This is not possible 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 a bit more subtle and we turn to this issue next. After that we discuss the justification to control for age effects.

The primary reason to control for family composition is that 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.[29] 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.

These estimates come from a simple cross-sectional OLS regression of consumption on the number of adults and kids in the household. However, family composition may

---

[29]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. 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).

change over the sample period. If (possible unobservable) other family characteristics are correlated with preferences over family size, then the change in consumption of a family with three kids that gets a fourth baby does not need to be the same as the average cross-sectional difference between families with three and four children. In fact, comparing the estimates from a cross-section regression with fixed effects estimates, which are identified only off within sample changes in family size, the coefficients turn out to be quite different. Controlling for family size in levels, will therefore introduce spurious changes in consumption for families that change family composition over the sample period. To control for this problem, we first regress log real consumption on the family composition variables, take the residuals and add back up the constant, and then run a FE regression of these residuals on the same variables, again taking the residuals but including the family specific fixed effect.

If we clean the consumption data, clearly we should clean the income data as well. However, there is much less justification to expect income to depend on family composition than consumption. Nevertheless, the coefficient estimates for these effects are quite significant and similar to the coefficients for consumption (although the coefficient on the number of children has the opposite sign). There is an issue whether we want to use these estimated coefficients, or the estimates 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 opt to use the estimated coefficients for income.

Graphing log real consumption and income, controlled for seasonality, attrition bias and family size by cohort over time it is that both income and consumption increase for the younger cohorts, and decrease for the older cohorts. This 'hump shaped' life-cycle profile of consumption and income is clearly non-linear, and will therefore not be captured by the trend α_i in income and b_i in consumption. Therefore, we clean the data for these age-effects by regressing on a fourth order polynomial in age, which captures the shape of the age profiles well, and taking the residuals. This has the additional advantage that the individual specific trends α_i in and b_i now will have mean zero across individuals. Notice that in some sense it is arbitrary to control for age effects. We are interested in the time effects in consumption and income, which represent the aggregate shocks. Therefore, we would like to control for both cohort and age effects, but clearly cannot do both. It seems plausible however, to assume that cohort effects are less important than age effects, and more directly related to the time effects.

To summarize the above: we first regress log real consumption and income on a set of month dummies (seasonality), interview dummies (attrition bias), a fourth order polynomial in age and the number of adults and children in the family. Then we take the residuals of this regression, and add back up the constant. Then we run a fixed

effects regression of these residuals on the number of adults and children, to remove spurious within sample changes that were introduced because of the first regression, and take the residuals including the family specific fixed effects. Notice that this second step will only change consumption and income for families that experience a within sample change in family composition.

We 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, of course these variables are highly correlated with income, so we risk removing exactly the variation we are interested in, which is the comovement between consumption and income.

D   Likelihood based estimation procedure

As noted in the main text, all shocks of the model are assumed to be jointly normal with zero mean and a block diagonal variance-covariance matrix. The only covariance block that is left unconstrained is the covariance between the innovations to the idiosyncratic permanent variance and the business cycle, represented by the aggregate shocks. We use $\omega^T = [\omega'_1, ..., \omega'_T]'$ to denote the history of a generic vector of variables ω_t and D^T to denote all the observable data used to estimate the model.

The objective here is providing a procedure to evaluate the posterior distribution of the parameters. The parameters of interest are the unobservable states, v_t, e_t, $Var_t(u)$ and $Var_t(n)$ and the hyperparameters Σ, θ and ϕ that are defined in the main text. Because of the high-dimensionality of the problem, we use Gibbs sampling to evaluate the posteriors. The Gibbs sampling is carried out in five steps, drawing in turn the aggregate shocks, the hyperparameters θ, the time varying variances of the idiosyncratic shocks, the hyperparameter ϕ and the hyperparameters Σ, conditional on the observed data and the rest of the parameters. We make use of rejection sampling in order to implement some of the restrictions on the parameters suggested by the theory. In the first section of this appendix we discuss each step in turn.

D.1   Markov chain Monte Carlo algorithm

Step 1: Drawing the aggregate shocks   Rewrite equations (4) and (5) in the state-space form

$$
\begin{aligned}
d_t &= H s_t + \varepsilon_t \\
s_t &= F s_{t-1} + \zeta_t,
\end{aligned}
$$

where

$$
d_t \equiv \begin{bmatrix} \Delta y_{1t} \\ \vdots \\ \Delta y_{J-1t} \\ \Delta y_{Jt} \\ \Delta c_{1t} \\ \vdots \\ \Delta c_{J-1t} \\ \Delta c_{Jt} \end{bmatrix} ; \ H = \begin{bmatrix} 1 & 1 & -1 & & & & & & \\ \vdots & \vdots & \vdots & & I_{J-1} & & & 0_{J-1} & \\ \vdots & \vdots & \vdots & & & & & & \\ 1 & 1 & -1 & -1 & \cdots & -1 & 0 & \cdots & 0 \\ 1 & 0 & 0 & & & & & & \\ \vdots & \vdots & \vdots & & 0_{J-1} & & & I_{J-1} & \\ \vdots & \vdots & \vdots & & & & & & \\ 1 & 0 & 0 & 0 & \cdots & 0 & -1 & \cdots & -1 \end{bmatrix} ; \ s_t = \begin{bmatrix} v_t \\ e_t \\ e_{t-1} \\ \alpha_i \\ \vdots \\ \alpha_{J-1} \\ b_1 \\ \vdots \\ b_{J-1} \end{bmatrix} ;
$$

$$
\varepsilon_t = \begin{bmatrix} \varepsilon^y_{1t} \\ \vdots \\ \varepsilon^y_{J-1t} \\ \varepsilon^y_{Jt} \\ \varepsilon^c_{1t} \\ \vdots \\ \varepsilon^c_{J-1t} \\ \varepsilon^c_{Jt} \end{bmatrix} ; \ F = \begin{bmatrix} 0 & 0 & 0 & & & & & \\ 0 & 0 & 0 & & 0_{3,J-1} & & 0_{3,J-1} & \\ 0 & 1 & 0 & & & & & \\ & 0_{J-1} & & & I_{J-1} & & 0_{J-1} & \\ & 0_{J-1} & & & 0_{J-1} & & I_{J-1} & \end{bmatrix} ; \ \zeta_t = \begin{bmatrix} v_t \\ e_t \\ 0 \\ 0 \\ \vdots \\ \vdots \\ \vdots \\ \vdots \\ 0 \end{bmatrix}
$$

Conditional on Σ, the system of observation equations is linear and has Gaussian innovations with known variance. We use the methodology described in Carter and Kohn (1994) to make draws from the posterior of the states. The density $p(s^T|D^T, \Sigma)$ can be factored as[30]

$$
p(s^T|D^T, \Sigma) = p(s_T|D^T, \Sigma) \prod_{t=1}^{T-1} p(s_t|s_{t+1}, D^T, \Sigma),
$$

where

$$
\begin{aligned}
p(s_t|s_{t+1}, D^T, \Sigma) &= N(s_{t|t+1}; V_{t|t+1}), \\
s_{t|t+1} &= E(s_t|s_{t+1}, D^T, \Sigma), \\
V_{t|t+1} &= Var(s_t|s_{t+1}, D^T, \Sigma).
\end{aligned}
$$

This means that the vector of v's and e's can be easily drawn because $s_{t|t+1}$ and $V_{t|t+1}$ can be computed using the forward (Kalman filter) and the backward recursions reported

---

[30]$p(\cdot)$ is used to denote a generic density function.

in appendix D.4. Specifically, the last recursion of the filter provides $s_{T|T}$ and $V_{T|T}$, i.e. mean and variance of the posterior distribution of the states. Drawn a value from this distribution, the draw can be used in the backward recursion to obtain $s_{T-1|T}$ and $V_{T-1|T}$ and so on.[31]

Step 2: Drawing the hyperparameters θ Consider now the system of equations given by the seven moment conditions relative to variances and covariances (7), (6), (8), (10), (9), (11) and (12). Rewrite them in the following linear form:

$$dd_t = GS_t + X_t\theta + \epsilon_t, \tag{13}$$

where

$$
\begin{aligned}
dd_t &= [Var_{1t}(\Delta c), ..., Var_{Jt}(\Delta c), Var_{1t}(\Delta y), ..., Var_{Jt}(\Delta y), cov_{1t}(\Delta y; \Delta c), ..., cov_{Jt}(\Delta y; \Delta c), \\
&\quad \Delta Var_{1t}(c), ..., \Delta Var_{Jt}(c), \Delta Var_{1t}(y), ..., \Delta Var_{Jt}(y), \Delta cov_{1t}(y; c), ..., \Delta cov_{Jt}(y; c), \\
&\quad cov_{1t}(\Delta y; y_{t-1}), ..., cov_{Jt}(\Delta y; y_{t-1})]' \\
S_t &= [Var_t(n), Var_t(u), Var_{t-1}(u)]' \\
\epsilon_t &= [\eta_{1t}^c, ..., \eta_{Jt}^c, \eta_{1t}^y, ..., \eta_{Jt}^y, \eta_{1t}^{yc}, ..., \eta_{Jt}^{yc}, \xi_{1t}^c, ..., \xi_{Jt}^c, \xi_{1t}^y, ..., \xi_{Jt}^y, \xi_{1t}^{yc}, ..., \xi_{Jt}^{yc}, \eta_{1t}^{yy}, ..., \eta_{Jt}^{yy},]',
\end{aligned}
$$

G is the states loading matrix and X_t a set of time varying regressors. For brevity we omit the exact expressions of G and X_t. The system of equations (13) can be rewritten as

$$dd_t^* = (dd_t - GS_t) = X_t\theta + \epsilon_t,$$

where, conditional on $D^T, v^T, e^T, Var(u)^T, Var(n)^T, \phi$ and Σ, dd^{*T} and X^T are observable. We are now in the case of a linear normal system of equations. Therefore θ can be easily drawn from his posterior distribution, which takes the form $p(\theta|D^T, v^T, e^T, Var(u)^T, Var(n)^T, \phi, \Sigma) = N(\widehat{\theta}_{OLS}; Var(\widehat{\theta}_{OLS}))$.

Step 3: Drawing $Var(u)^T$ and $Var(n)^T$ Consider again equation (13) and rewrite it as

$$dd_t^{**} = (dd_t - X_t\theta) = GS_t + \epsilon_t, \tag{14}$$

where, conditional on $D^T, v^T, e^T, \theta, \phi$ and Σ, dd^{**T} is observable. The states transition equation is given by:

$$S_t = BS_{t-1} + w_t,$$

where

$$B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}; \; w_t = \begin{bmatrix} w_{nt} \\ w_{ut} \\ 0 \end{bmatrix}. \tag{15}$$

---

[31] For further details on Gibbs sampling for state space models see Carter and Kohn (1994) or Kim and Nelson (1999).

The same procedure applied in the first step of the sampler can be adopted here to draw the unobservable states $Var(n)^T$ and $Var(u)^T$.

Step 4: Drawing $1 - \phi$ The parameter $1 - \phi$ enters non-linearly some of the moment conditions of A and this complicates the analysis. Therefore, in order to take a draw from the posterior distribution of $1 - \phi$ we use a Metropolis-Hastings step,[32] nested in the Gibbs sampler. The procedure works as follows: we draw a candidate value $1 - \phi^*$ from a proposal distribution, say $\varphi(1 - \phi^* | D^T, v^T, e^T, Var(u)^T, Var(n)^T, \theta, \Sigma)$, which in our case is normal. In order to chose a reasonable value for the mean of the proposal distribution, we run a regression of $\Delta Var(c)$ on $Var(n)$ and we take the square root of the result.[33]   The variance is arbitrarily set to 0.03. We reject the proposal values outside the interval $[0, 1]$. At this point we compute the posterior value associated to the draw, $p\left(1 - \phi^* | D^T, v^T, e^T, Var(u)^T, Var(n)^T, \theta, \Sigma\right)$ that, under flat priors, coincides with the likelihood of the system of equations given by (7), (10), (8) and (11). The new draw is accepted with probability

$$q = \min\left\{ \frac{p(1 - \phi^*) \; / \; \varphi(1 - \phi^*)}{p(1 - \phi^{i-1}) \; / \; \varphi(1 - \phi^{i-1})}, 1 \right\}.$$

where $1 - \phi^{i-1}$ is the previous draw of the chain. If the proposal value is rejected, the next element of the chain is set to be $1 - \phi^{i-1}$.

Step 5: Drawing Σ Conditional on v^T, e^T, $Var(u)^T$, $Var(n)^T$, ϕ and D^T, every element of Σ has an inverse-Gamma posterior distribution, independent of the other elements. Moreover, conditional on v^T, e^T, $Var(u)^T$, $Var(n)^T$, ϕ and D^T, it is easy to draw from these inverse-Gamma posteriors because the innovations are observable.[34]

Summary   Summarizing, the sampler takes the form:

1. Initialize $\Sigma, Var(u)^T, Var(n)^T, \phi$.

2. Sample $v^T, e^T, \alpha_1, ..., \alpha_{J-1}$ and $b_1, ..., b_{J-1}$ from $p(v^T, e^T, \alpha_1, ..., \alpha_{J-1}, b_1, ..., b_{J-1} | D^T, \Sigma)$.

3. Sample θ from $p(\theta | D^T, v^T, e^T, Var(u)^T, Var(n)^T, \phi, \Sigma)$.

4. Sample $Var(u)^T$ and $Var(n)^T$ from $p(Var(u)^T, Var(n)^T | D^T, v^T, e^T, \theta, \phi, \Sigma)$.

5. Sample ϕ from $p(\phi | D^T, v^T, e^T, Var(u)^T, Var(n)^T, \theta, \Sigma)$ through a Metropolis-Hastings step.

---

[32] See Gelman et al. (1995) for further details on Metropolis-Hastings algorithms.

[33] The intuition for this step relies on equation (10).

[34] See Gelman et al. (1995) for a description of the sampling procedure from inverse-gamma or inverse-Wishart distributions, also when combined with the natural conjugate prior described below.

6. Sample Σ, by sampling Σ_i from $p(\Sigma_i | D^T, v^T, e^T, Var(u)^T, Var(n)^T, \theta, \phi)$

7. Rejection of the whole draw if $Var(\delta) < 0$ or $Var(\gamma) < 0$.

8. Go to 2.

D.2   Priors

To make the results comparable to previous studies we use flat and uninformative priors, except for the cases described below, in which a prior is suggested by the theory *and* the use of conjugate prior simplifies the analysis. In the first step, for a_j and b_j, $j = 1, ..., J$, we assumed a normal prior with mean corresponding to the sample mean of $\{\Delta y_{jt}\}_{t=1}^T$ and $\{\Delta c_{jt}\}_{t=1}^T$ and variance arbitrarily fixed to 10. In the third step of the sampler, every element of $S_{0|0}$ is set to be equal to the sample mean of $\{\{Var_{jt}(\Delta y)\}_{t=1}^T\}_{j=1}^J$, while the variance of S_0 is set to the identity matrix. In the fifth step we use conjugate inverse-gamma priors for σ_v^2 and σ_e^2. The scale values are set respectively to the variance of the aggregate consumption and the variance of the aggregate income minus the variance of the aggregate consumption, dividing the result by two. The degrees of freedom are set to 5. The effect of the priors is negligible, except in the last case for σ_v^2 and σ_e^2, where the prior plays the role of slightly increasing the size of the aggregate shocks to bring these more in line with the GMM estimates. Notice that this prior, if anything, strengthens our result that aggregate shocks have a negligible effect on inequality.

D.3   Correlation aggregate shocks and idiosyncratic risk

In order to explicitly take into account the possible correlation between aggregate shocks and innovations to the variance of the idiosyncratic permanent shocks, we slightly modify the procedure described above in the following way. In the fifth step, we draw σ_v^2, σ_e^2, $\sigma_{w_n}^2$, σ_{v,w_n} and σ_{e,w_n} jointly from an inverse-Wishart distribution. Then in the first step we add the equation

$$w_{nt} = \kappa_v v_t + \kappa_e e_t + z_t$$

to the measurement equations, where w_{nt} are observable and κ_v, κ_e and $Var(z_t)$ can be easily computed given σ_v^2, σ_e^2, $\sigma_{w_n}^2$, σ_{v,w_n} and σ_{e,w_n}. After this point the procedure in step 1 is identical to the one described above. In the third step we add the two equations

$$v_t = \varphi_v w_{nt} + h_{vt}$$
$$e_t = \varphi_e w_{nt} + h_{et},$$

to the measurement equations, where v_t and e_t are observable and φ_v, φ_e, $Var(h_{vt})$ and $Var(h_{et})$ can be easily computed given σ_v^2, σ_e^2, $\sigma_{w_n}^2$, σ_{v,w_n} and σ_{e,w_n}. If the state vector is expanded to take $Var_{t-1}(n)$ into account, the procedure in step 3 is identical to the one described above. This adjustment turns out not to make any difference for the results.

D.4   Gibbs sampling for state-space models

These are the forward (Kalman filtering) and the backward recursions for Gibbs sampling for state space models. Measurement equation:

$$d_t = H s_t + \varepsilon_t$$

Transition equation:

$$s_t = F s_{t-1} + \zeta_t$$

where

$$\begin{pmatrix} \varepsilon \\ \zeta \end{pmatrix} \sim N(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} R & 0 \\ 0 & Q \end{bmatrix}).$$

Let

$$\begin{aligned} s_{t|s} &= E(s_t | Y^s, H, R, Q), \\ V_{t|s} &= Var(s_t | Y^s, H, R, Q). \end{aligned}$$

Then, given $s_{0|0}$ and $V_{0|0}$ and following the standard Kalman filter forward recursion:

$$\begin{aligned} s_{t|t-1} &= F s_{t|t-1}, \\ V_{t|t-1} &= F V_{t-1|t-1} F' + Q, \\ K_t &= V_{t|t-1} H_t' (H V_{t|t-1} H' + R)^{-1}, \\ s_{t|t} &= s_{t|t-1} + K_t (y_t - H s_{t|t-1}), \\ V_{t|t} &= V_{t|t-1} - K_t H V_{t|t-1}. \end{aligned}$$

In this way they can be obtained $s_{T|T}$ and $V_{T|T}$, used to draw the first value of the backward recursion. The updating formulas for the following draws are:

$$\begin{aligned} s_{t|t+1} &= s_{t|t} + V_{t|t} V_{t+1|t}^{-1} (s_{t+1} - s_{t|t}), \\ V_{t|t+1} &= V_{t|t} - V_{t|t} V_{t+1|t}^{-1} V_{t|t}. \end{aligned}$$

References

[1] Autor, D., L. Katz and A. Krueger (1997). Handbook article

[2] Richard Blundell and Ian Preston (1998). Consumption Inequality and Income Uncertainty. QJE, May 1998.

[3] Alon Brav, George M. Constantinides, Christopher C. Geczy (1999). Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence. NBER Working Paper No.w7406

[4] John Y. Campbell, N. Gregory Mankiw (1990). Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. NBER Working Paper No.w2924*

[5] George M. Constantinides, Darrell Duffie (1996). Asset Pricing with Heterogeneous Consumers. The Journal of Political Economy, Vol. 104, No. 2. (Apr., 1996), pp. 219-240.

[6] Carroll, Christopher D. (2001a). A Theory of the Consumption Function, With and Without Liquidity Constraints (expanded version). NBER Working Paper 8387.

[7] Carter C. K. and R. Kohn (1994), "On Gibbs Sampling for State Space Models", *Biometrika*, 81, 541-553.

[8] Timothy Cogley (2002). Idiosyncratic risk and the equity premium: evidence from the consumer expenditure survey. *Journal of Monetary Economics,* 49, pp.309-334.

[9] Angus Deaton, Christina Paxson (1994). Intertemporal Choice and Inequality. The Journal of Political Economy, Vol. 102, No. 3. (Jun., 1994), pp. 437-467.

[10] Gelman A., J. B. Carlin, H. S. Stern and D. B. Rubin (1995), "*Bayesian Data Analysis*", London: Chapman and Hall.

[11] Peter Gottschalk, Robert Moffitt (1994). The Growth of Earnings Instability in the U.S. Labor Market. Brookings Papers on Economic Activity, Vol. 1994, No. 2. (1994), pp. 217-254.

[12] Pierre-Olivier Gourinchas and Jonathan A. Parker (2001). The Empirical Importance of Precautionary Saving. AER AEA Papers and Proceedings January 2001.

[13] Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002). Consumption over the Lifecycle. Econometrica, January 2002, 70(1), pp47-91.

[14] John Heaton, Deborah J. Lucas (1996). Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing. The Journal of Political Economy, Vol. 104, No. 3. (Jun., 1996), pp. 443-487.

[15] Peter J. Huber (1981). Robust statistics. New York: Wiley.

[16] Kim C. and C. R. Nelson (1999), *State-Space Models with Regime Switching*, MIT Press: Cambridge, Massachusetts.

[17] Dirk Krueger and Fabrizio Perri (2001). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. mimeo.

[18] Thomas E. MaCurdy (1982). The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis. Journal of Econometrics, vol.18, pp.83-114.

[19] Robert A. Moffitt 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.

[20] Robert A. Moffitt and Peter Gottschalk (2002). Trends in the Transitory Variance of Earnings in the United States. *The Economic Journal,* 112, pp.C68-C73.

[21] Danny Quah (1998). Permanent and Transitory Movements in Labor Income: An Explanation for "Excess Smoothness" in Consumption. The Journal of Political Economy, Vol. 98, No. 3. (Jun., 1990), pp. 449-475.

[22] Danny Quah (1992). The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds. Econometrica, Vol. 60, No. 1. (Jan., 1992), pp. 107-118.

[23] Garey Ramey and Valerie A. Ramey (1995). Cross-Country Evidence on the Link Between Volatility and Growth. American Economic Review, 85(5), pp.1138-1151.

[24] Kjetil Storesletten, Chris I. Telmer, Amir Yaron (2000a). Consumption and Risk Sharing Over the Life Cycle. NBER Working Paper No.w7995

[25] Kjetil Storesletten, Chris I. Telmer, Amir Yaron (2000b). The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk. NBER Working Paper No.w8040

[26] US Department of Commerce, Bureau of the Census and US Department of Labor, Bureau of Labor Statistics (2001). Current Population Survey, Annual Demographic Survey (March CPS Supplement), 1977-2001. http://www.bls.census.gov/cps or http://www.nber.org/data/cps.html

[27] US Department of Labor, Bureau of Labor Statistics (1999). Consumer Expenditure Survey, Interview Survey, 1980-1999. http://www.bls.gov/cex

Table 1.
Summary statistics

Variable	Full sample aged 21-65 (and >5 in 1980)		Estimation sample	
	Mean	Std. Dev.	Mean	Std. Dev.
age	41.27	11.94	40.14	11.42
fam_size	2.84	1.58	2.86	1.59
married	0.58	0.49	0.57	0.50
single	0.31	0.46	0.32	0.46
nrad (16+)	2.04	0.96	2.03	0.96
nrkd (15-)	0.80	1.13	0.83	1.14
no_earnr	1.62	0.94	1.68	0.91
inc_hrs1	42.74	11.15	42.64	11.00
inc_hrs2	37.43	12.09	37.26	12.04
D_cons	0.001	0.036	0	0
D_inc	0.183	0.386	0	0
D_retd	0.043	0.203	0	0
D_student	0.011	0.104	0	0
D_studhous	0.005	0.073	0	0
inc	27907.85	28832.68	31804.09	24189.21
foodbev	1219.14	886.70	1185.21	800.70
cons	2634.86	1696.22	2568.36	1516.42
totexp	7282.63	6240.78	7160.52	5570.45
rinc	218.86	207.99	252.09	177.03
rfoodbev	9.63	6.71	9.47	6.19
rcons	22.07	13.48	21.73	12.30
rtotexp	57.09	45.33	56.86	41.20
lrinc	5.10	1.11	5.24	0.88
lrcons	2.93	0.58	2.93	0.56
lrincSA	5.10	1.13	5.24	0.89
lrconsSA	2.93	0.58	2.93	0.56
lrincSAA	4.92	1.09	5.04	0.85
lrconsSAA	2.83	0.53	2.83	0.51
NrObs	263786		228984	

Table 2.
Cell sizes by cohort and year

year	coh8	coh13	coh18	coh23	coh28	coh33	coh38	coh43	coh48	coh53	coh58	coh63	Total
1980			21	93	112	102	77	72	64	58	68	55	722
1981		0	93	442	511	518	395	306	316	321	289	207	3398
1982		0	56	136	124	105	95	73	70	68	67	44	838
1983		3	148	345	433	373	285	213	226	280	202	115	2623
1984		8	193	354	458	400	331	239	214	254	186	102	2739
1985		51	308	590	578	563	465	354	339	301	267	118	3934
1986	1	31	85	144	97	86	76	63	55	55	30	27	750
1987	0	117	324	481	453	444	337	282	306	230	160	82	3216
1988	1	92	330	385	392	362	275	217	199	192	106	62	2613
1989	8	136	369	398	361	380	273	243	200	148	106	49	2671
1990	32	189	361	421	390	382	296	245	176	157	87	36	2772
1991	47	206	398	405	361	353	287	197	163	135	84	45	2681
1992	83	236	397	437	394	320	254	227	161	115	56	32	2712
1993	104	247	413	420	399	338	243	206	158	99	65	27	2719
1994	140	278	414	437	376	325	279	193	148	84	40	21	2735
1995	193	321	447	487	436	423	297	205	175	76	45	20	3125
1996	117	148	204	186	180	147	101	87	45	23	8	7	1253
1997	212	305	338	376	332	293	190	130	62	43	17	12	2310
1998	247	321	364	364	327	297	211	140	75	46	24	12	2428
1999	292	342	352	349	377	307	192	129	71	47	22	12	2492
2000	209	256	309	290	276	216	169	83	61	33	19	8	1929
Total	1686	3287	5924	7540	7367	6734	5128	3904	3284	2765	1948	1093	50660

Table 3.
Scalar parameter estimates

	GMM		Likelihood based	
	Mean	St. error	Mean	St. error
Impact aggregate shocks				
variance(gamma) (1)	22.03	9.34	10.48	10.62
variance(delta) (1)	22.90	12.50	12.74	13.42
covariance(gamma,delta)	26.82	13.27	5.15	10.34
cov(gamma,y0)	1.07	1.10	0.09	0.68
cov(delta,y0)	0.53	0.76	0.30	0.88
cov(gamma,c0)	1.18	3.18	0.05	0.26
cov(delta,c0)	0.21	0.20	-0.11	0.35
Deterministic trends				
variance(alpha)	0	(2)	0	(2)
variance(b)	0	(3)	0.0002	0.0003
cov(alpha,y0)	0	(2)	0	(2)
cov(b,y0)	0	(3)	-0.0010	0.0035
Measurement error				
in variance(dy) and autocov (4)	-0.041	0.023	-0.028	0.055
in variance(dc)	0.272	0.015	0.286	0.010
in covariance(dy,dc)	0.012	0.004	0.012	0.005
Risk sharing				
phi	0.775	(3)	0.775	0.115
correlation coefficient v and var(n)			0.065	0.211

(1) Restricted to be positive.
(2) Restricted to zero because of identification.
(3) Restricted for reasons of convergence.
(4) Measurement error in both equations not separately identified. Estimate is difference.

Figure 1. Inequality by cohort
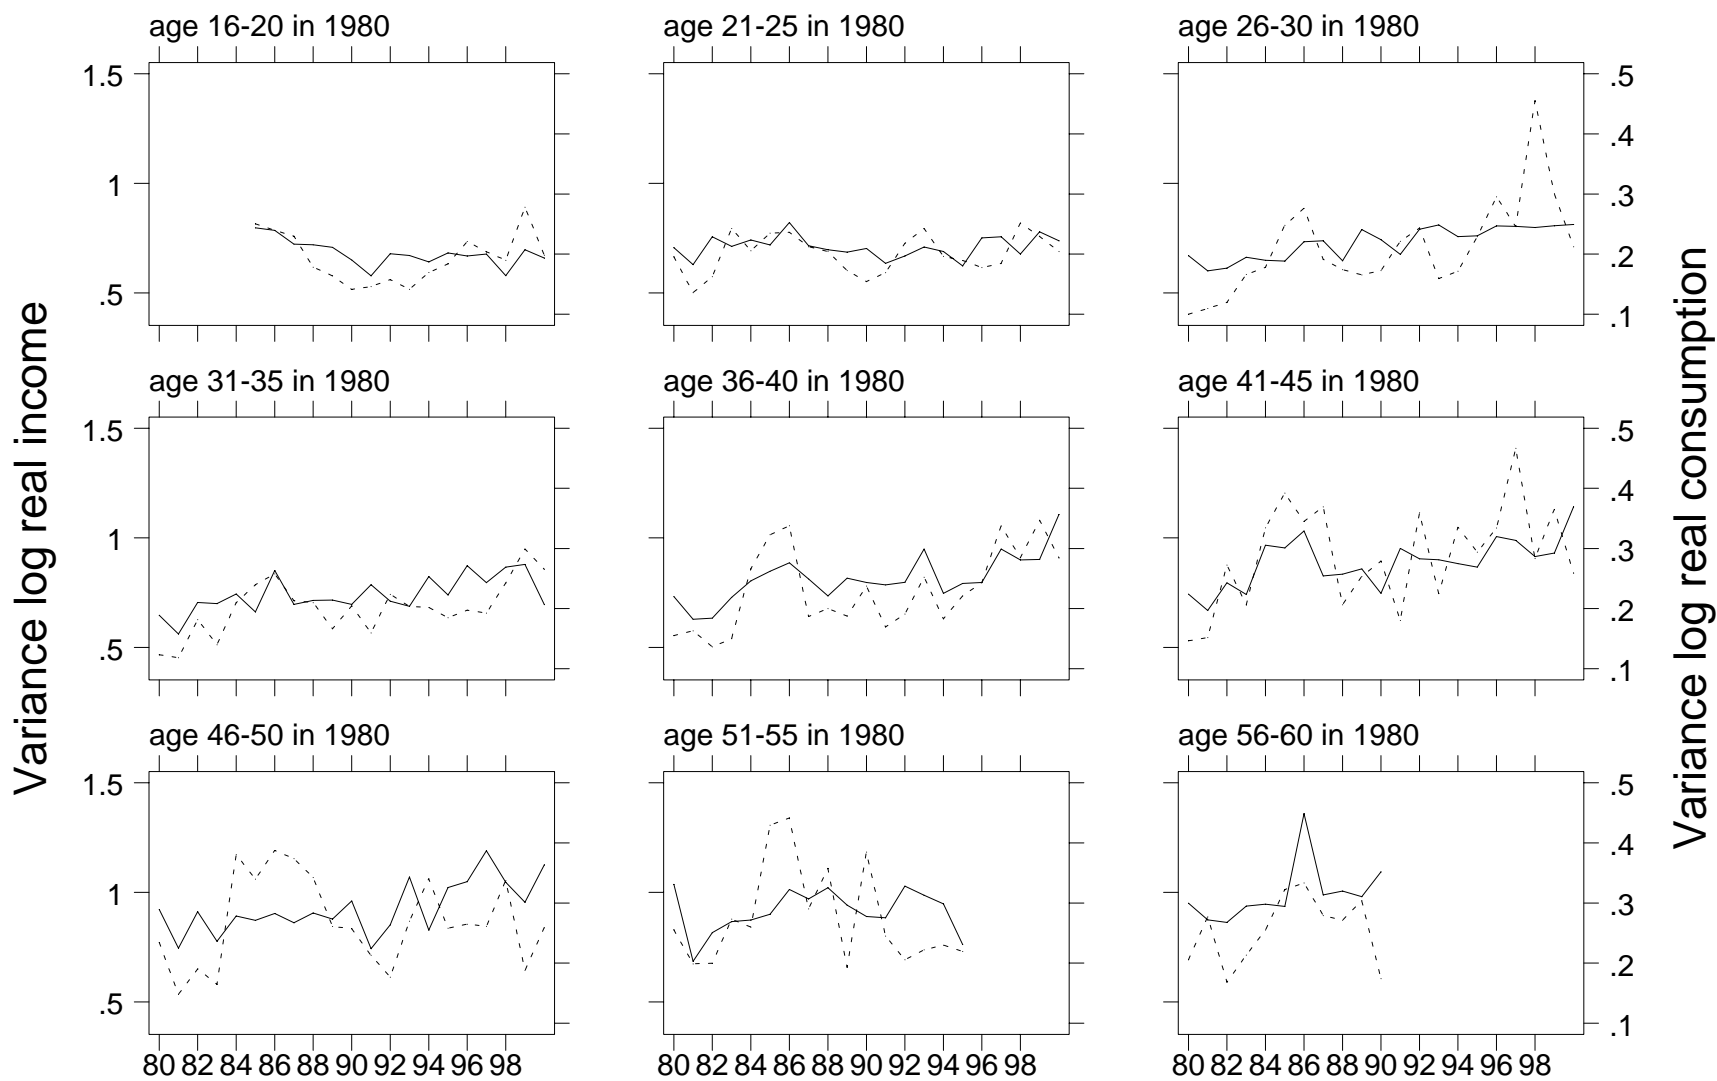
Figure 2. Age profile consumption inequality
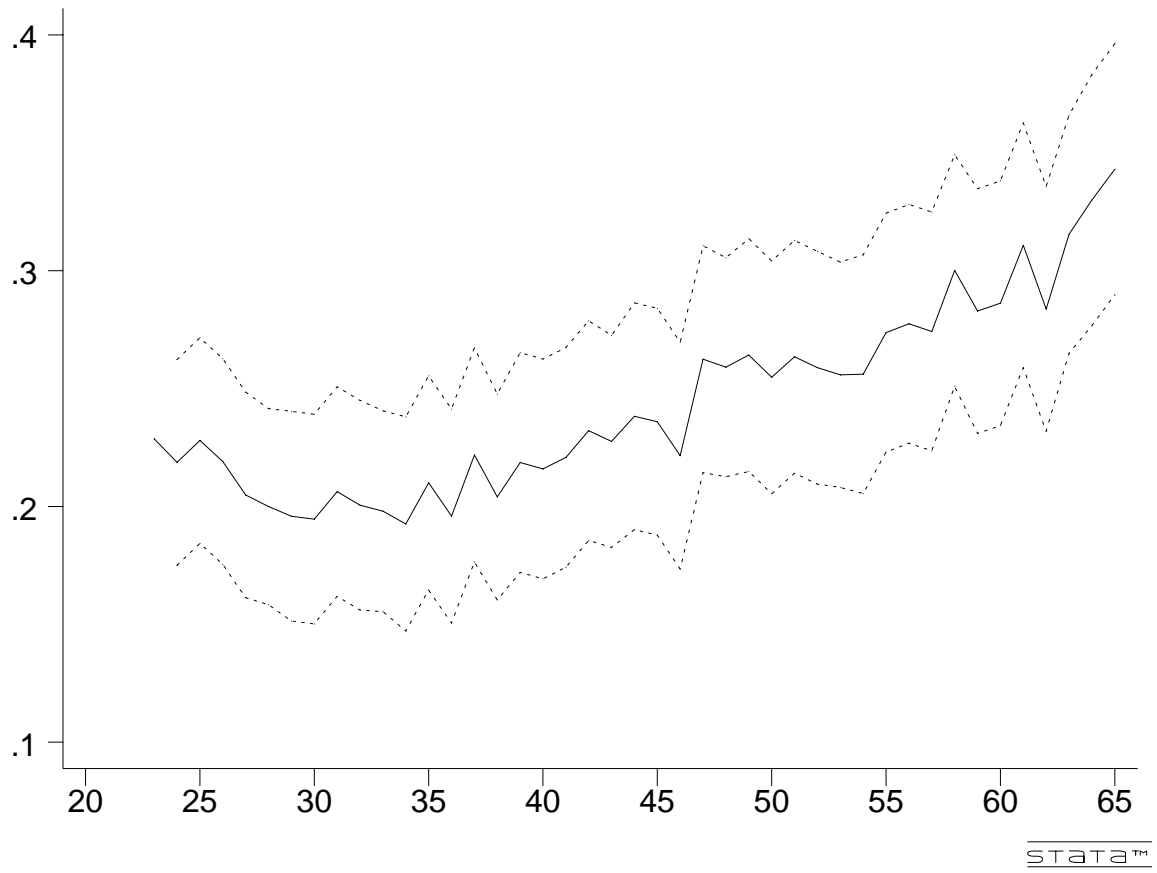
Figure 3.

Actual and model predicted income inequality
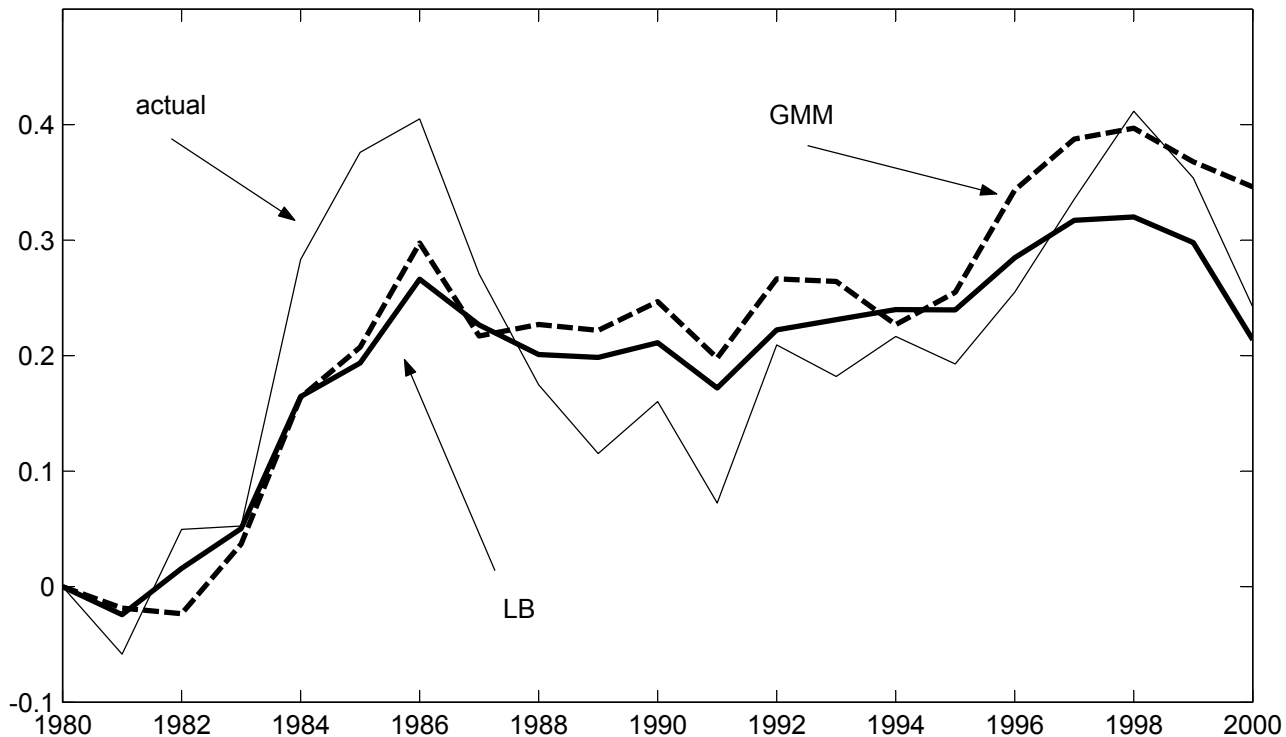
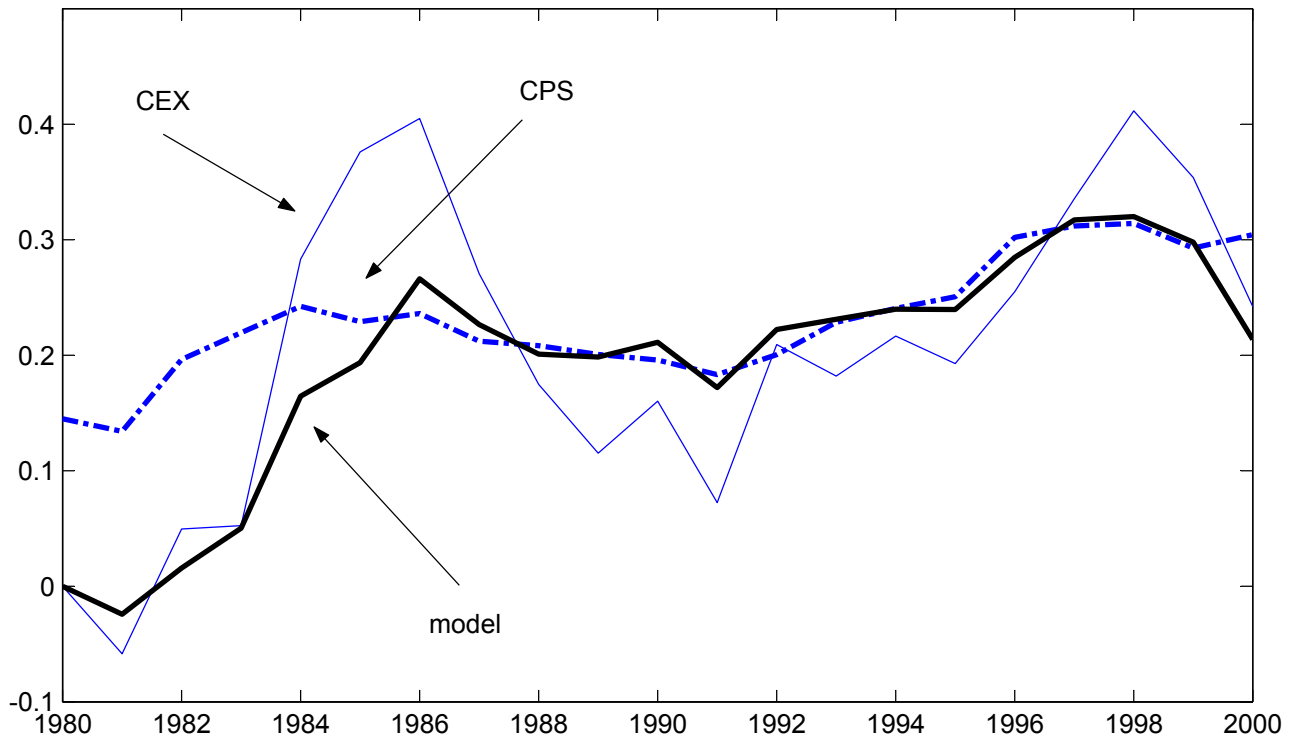## Actual and model predicted consumption inequality
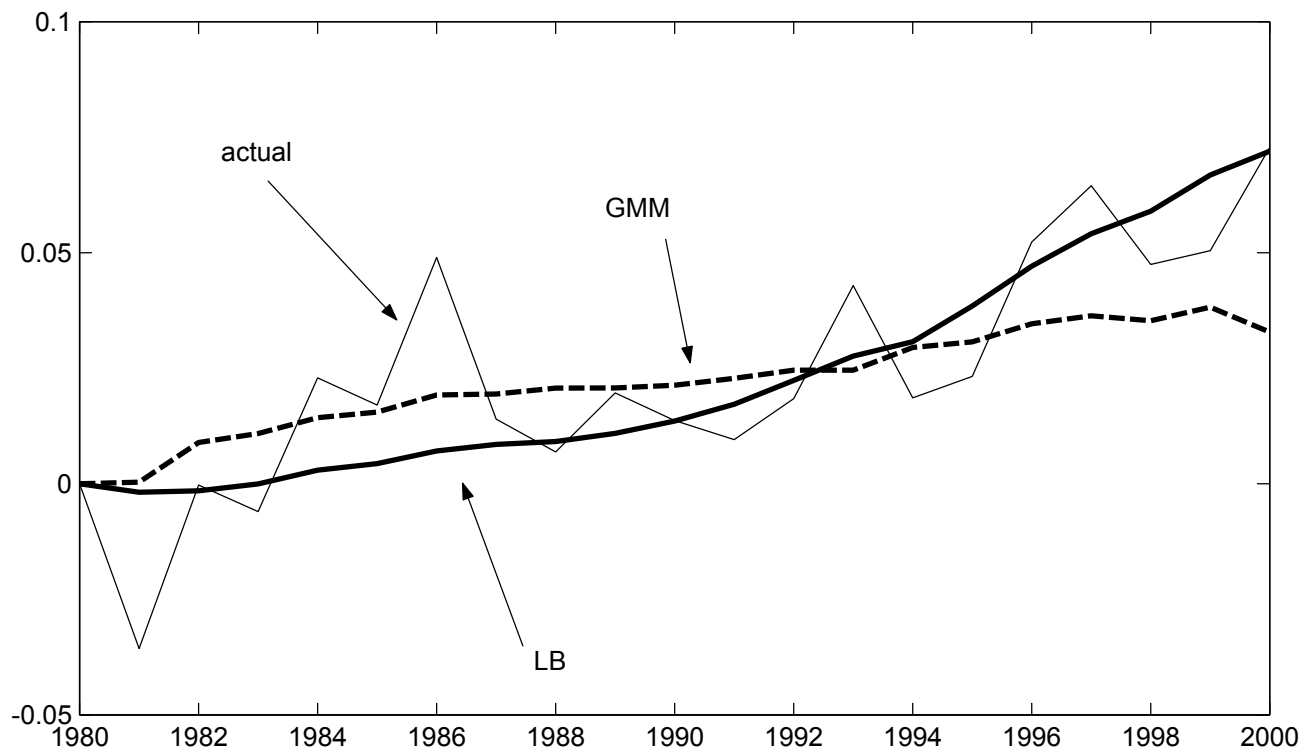
Figure 4.

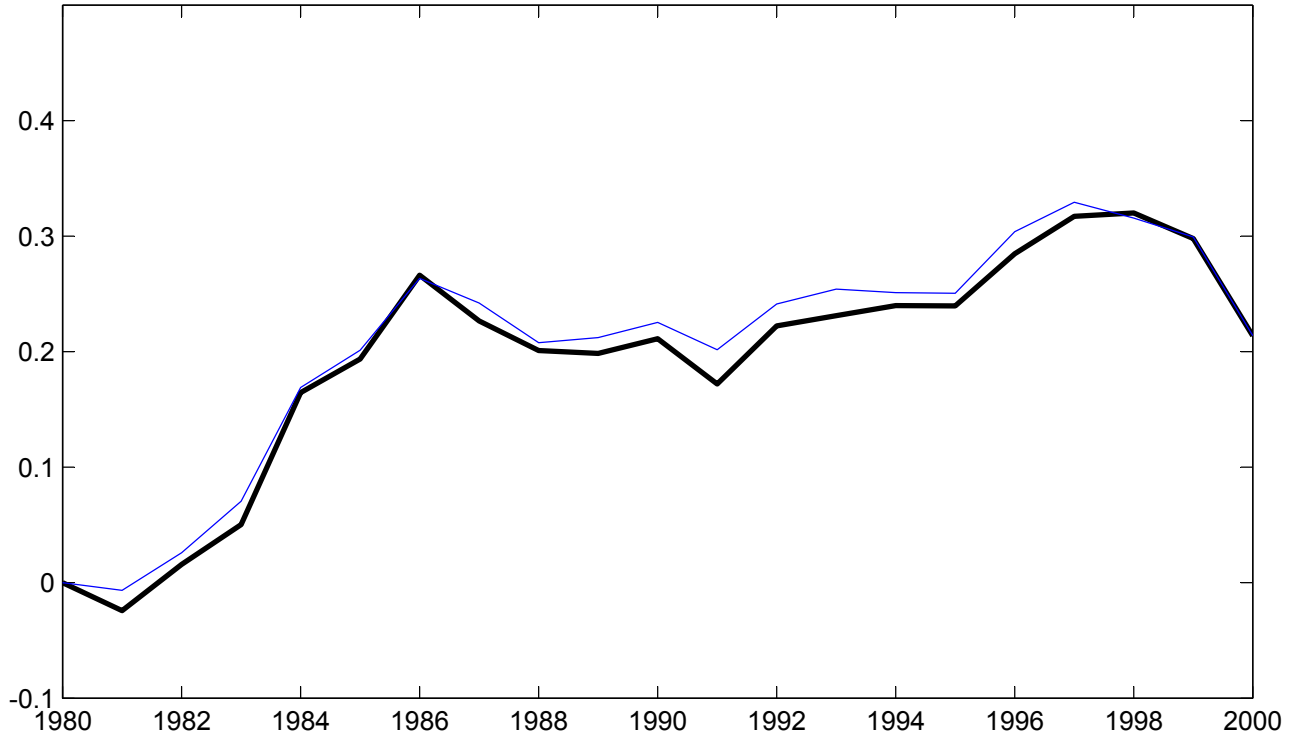### Actual (robust and non robust) and model predicted income inequality
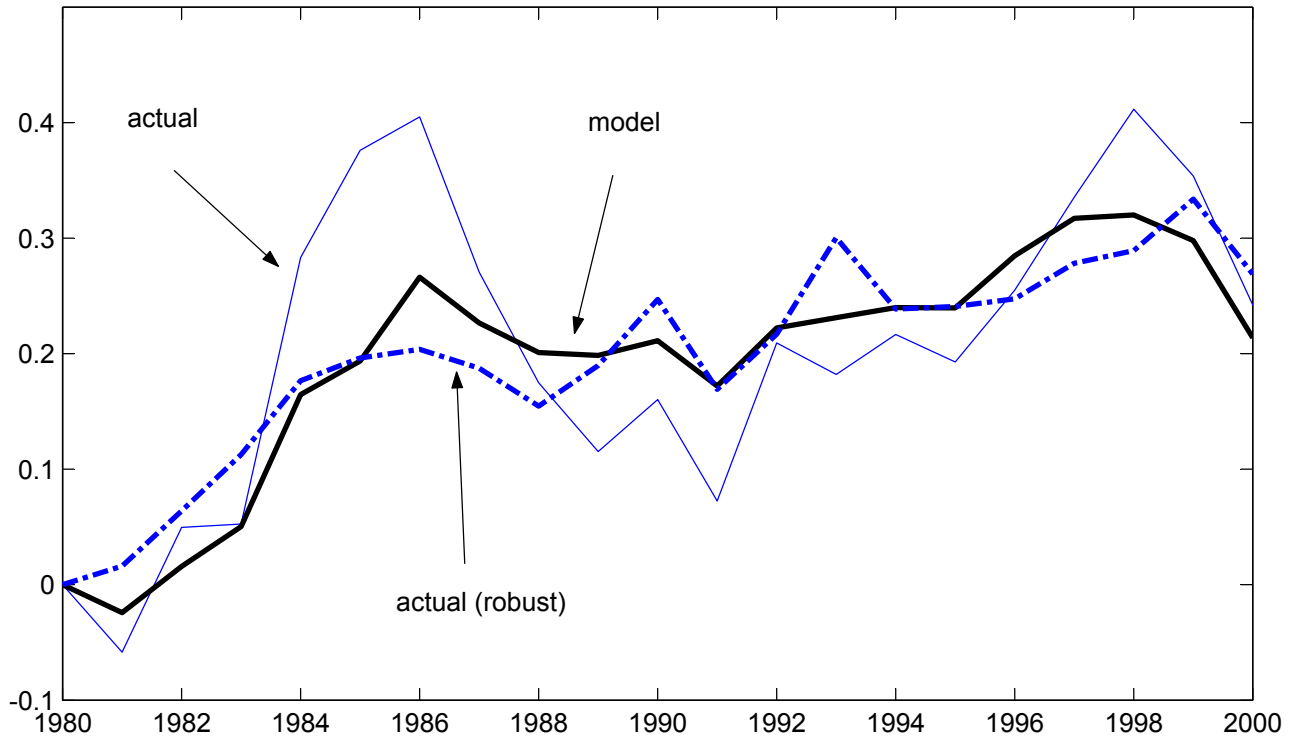
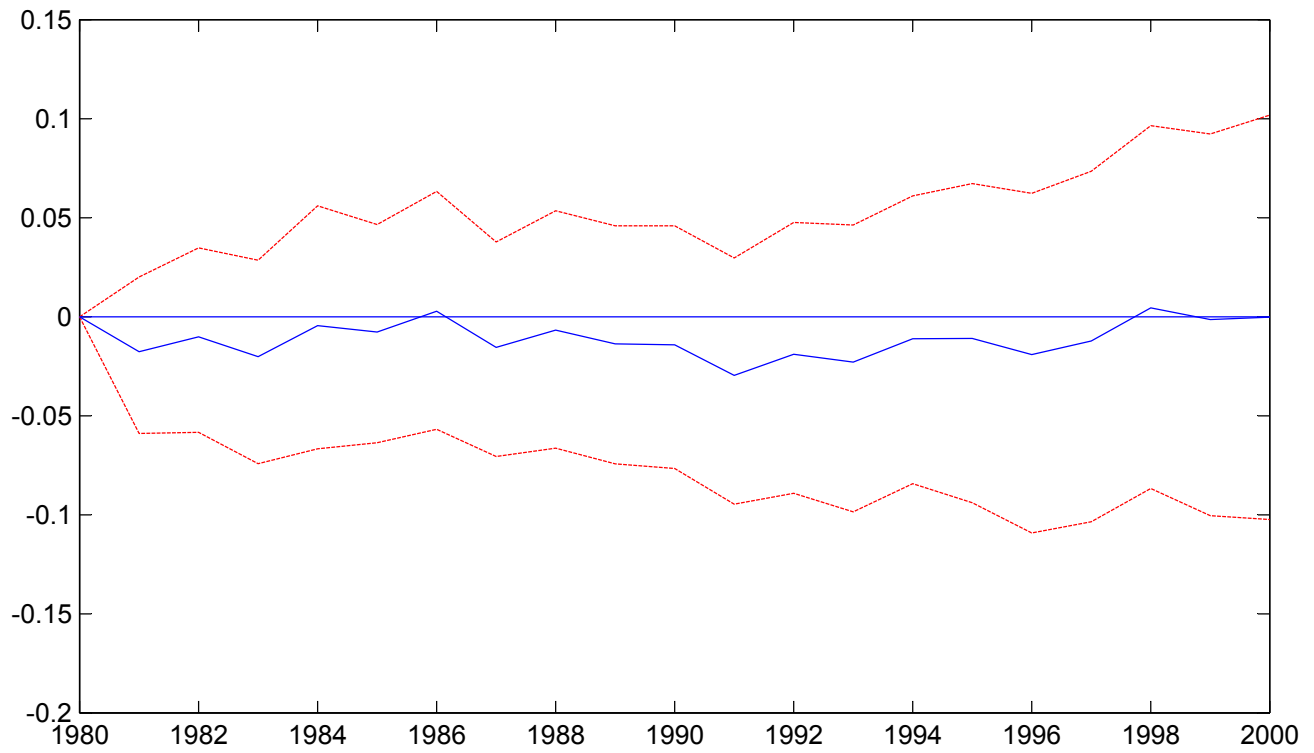### Actual (CEX and CPS) and model predicted income inequality

Figure 5.

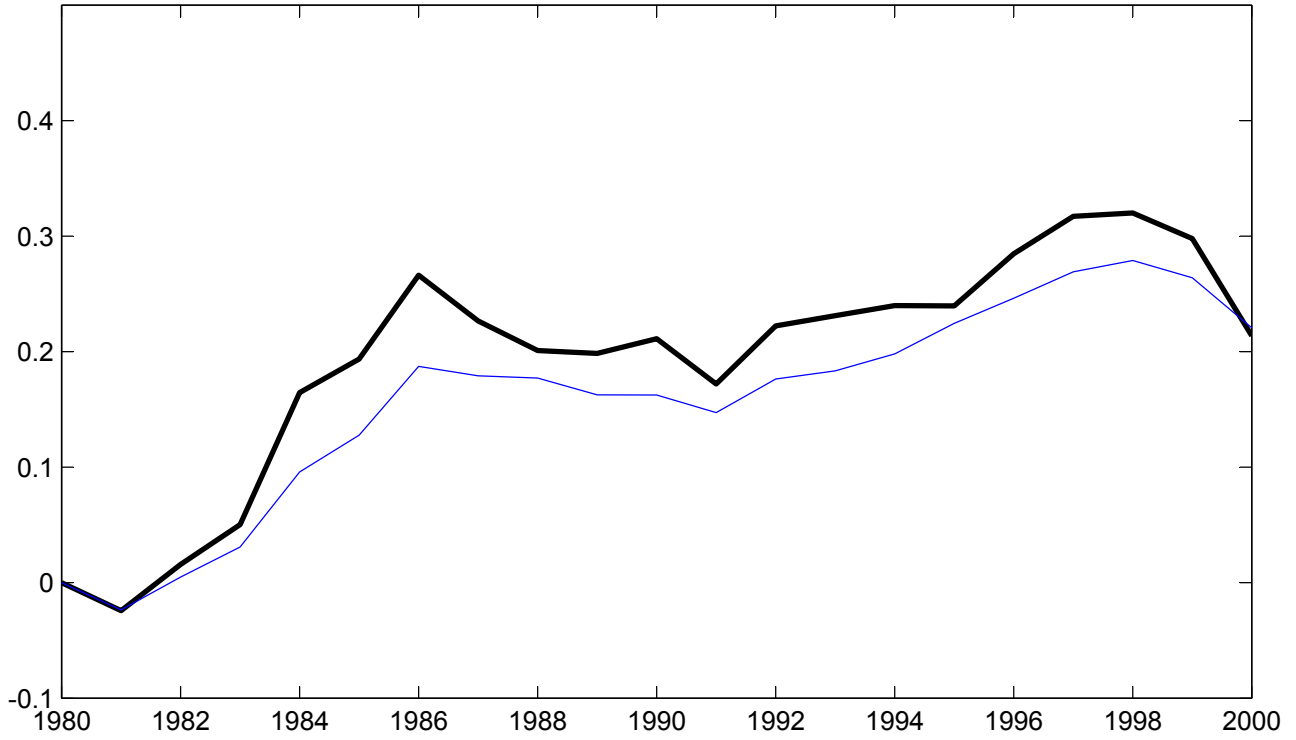## Income inequality without aggregate sources of uncertainty
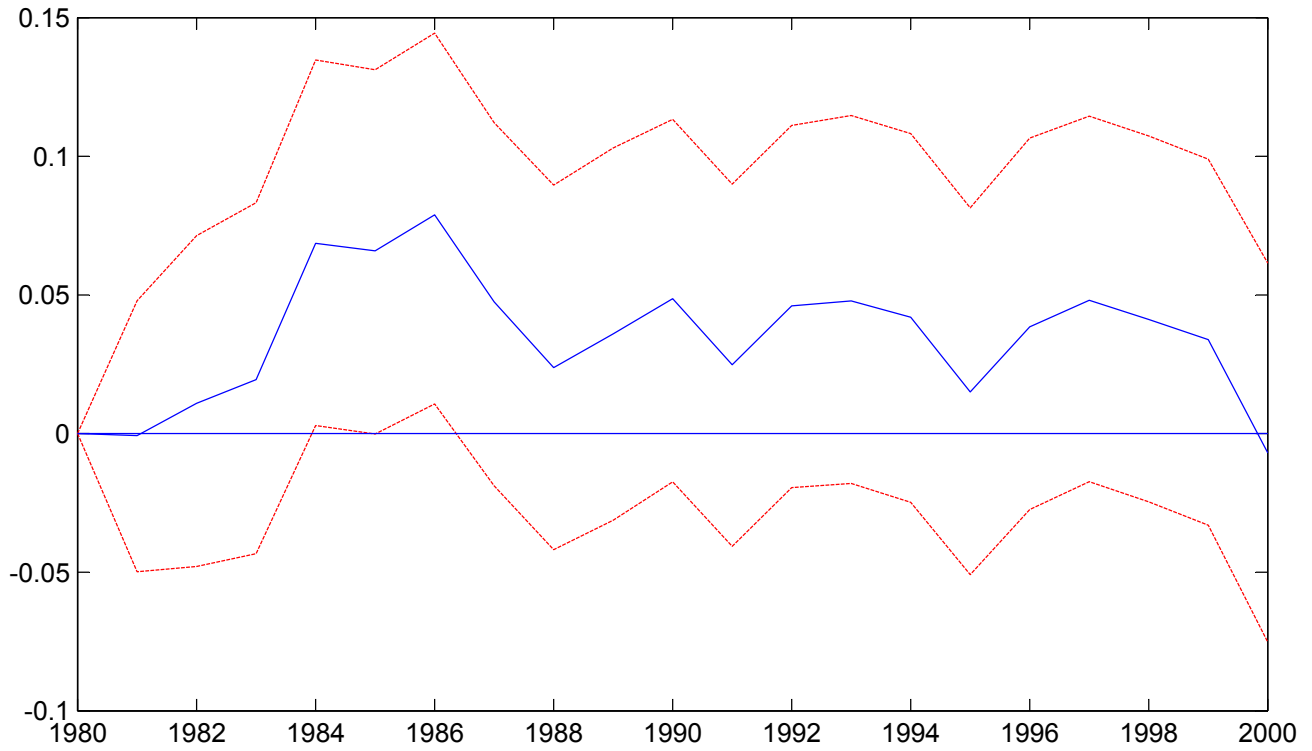

Impact of aggregate sources of uncertainty on income inequality

Figure 6.
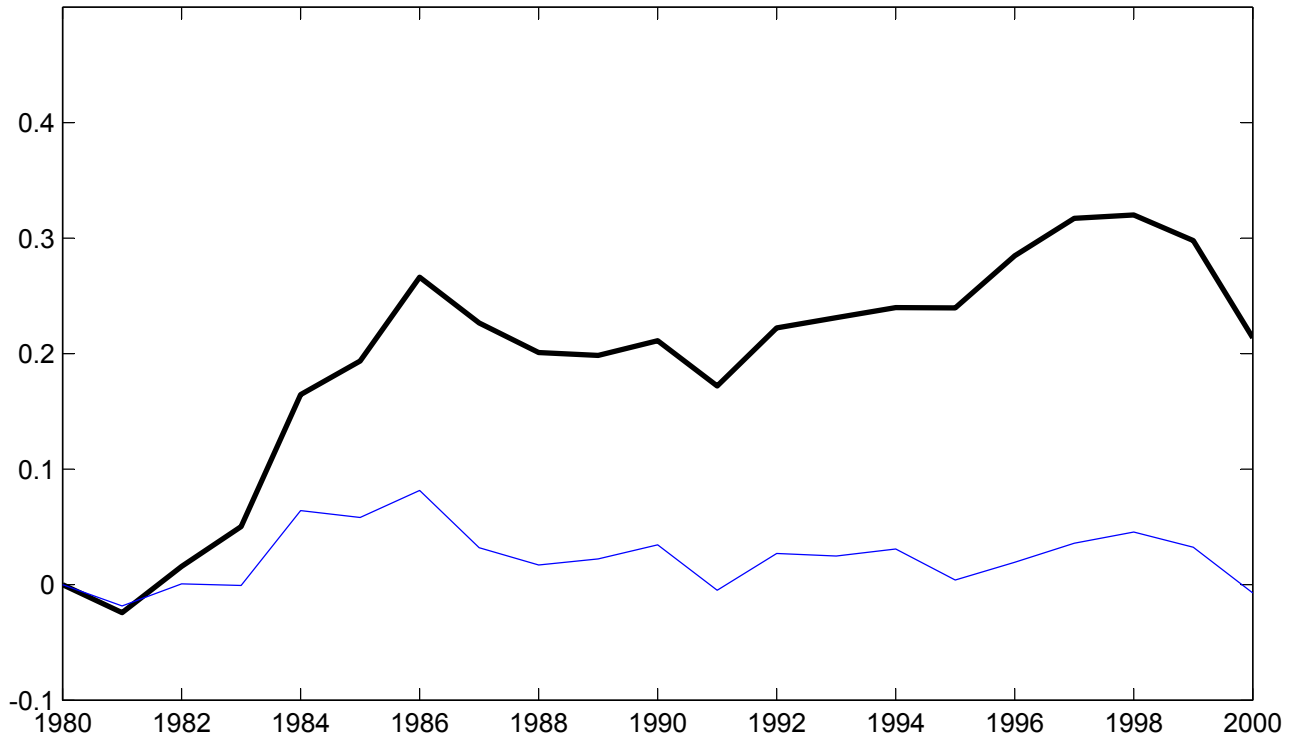
Income inequality without transitory idiosyncratic inequality

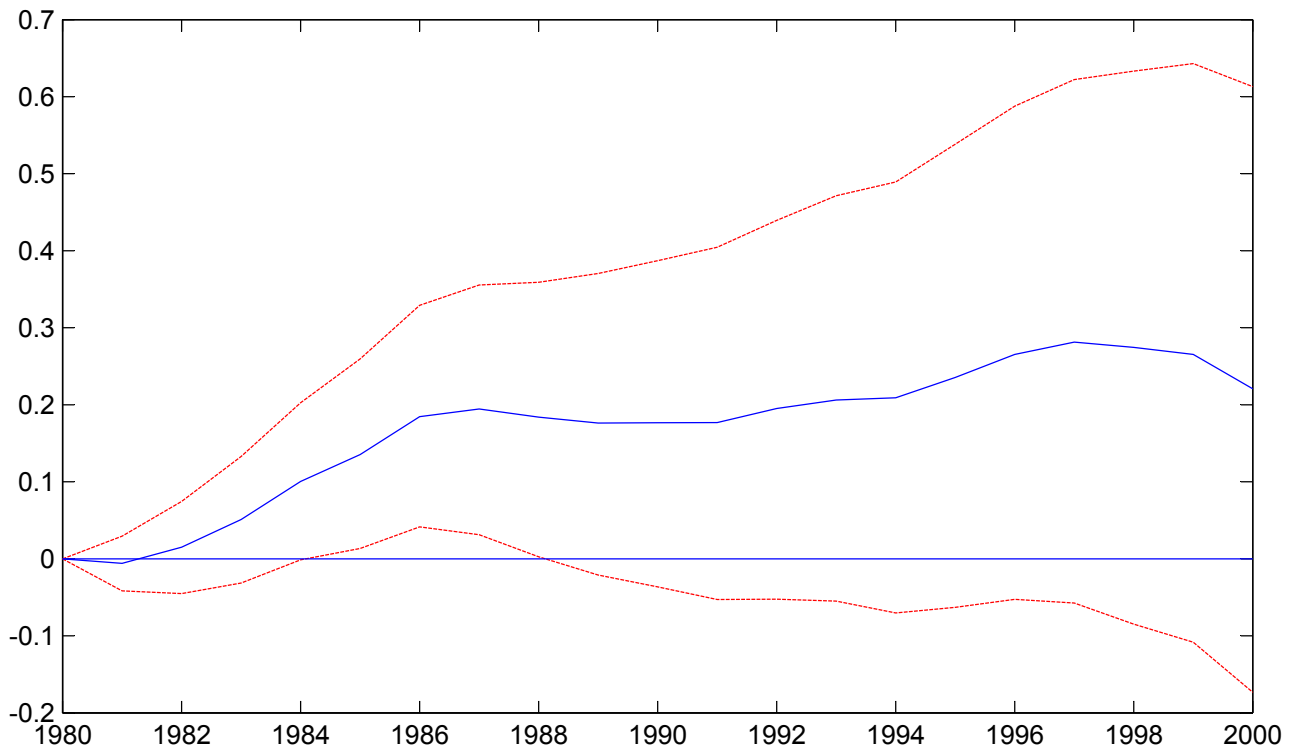### Impact of transitory idiosyncratic inequality on income inequality

Figure 7.

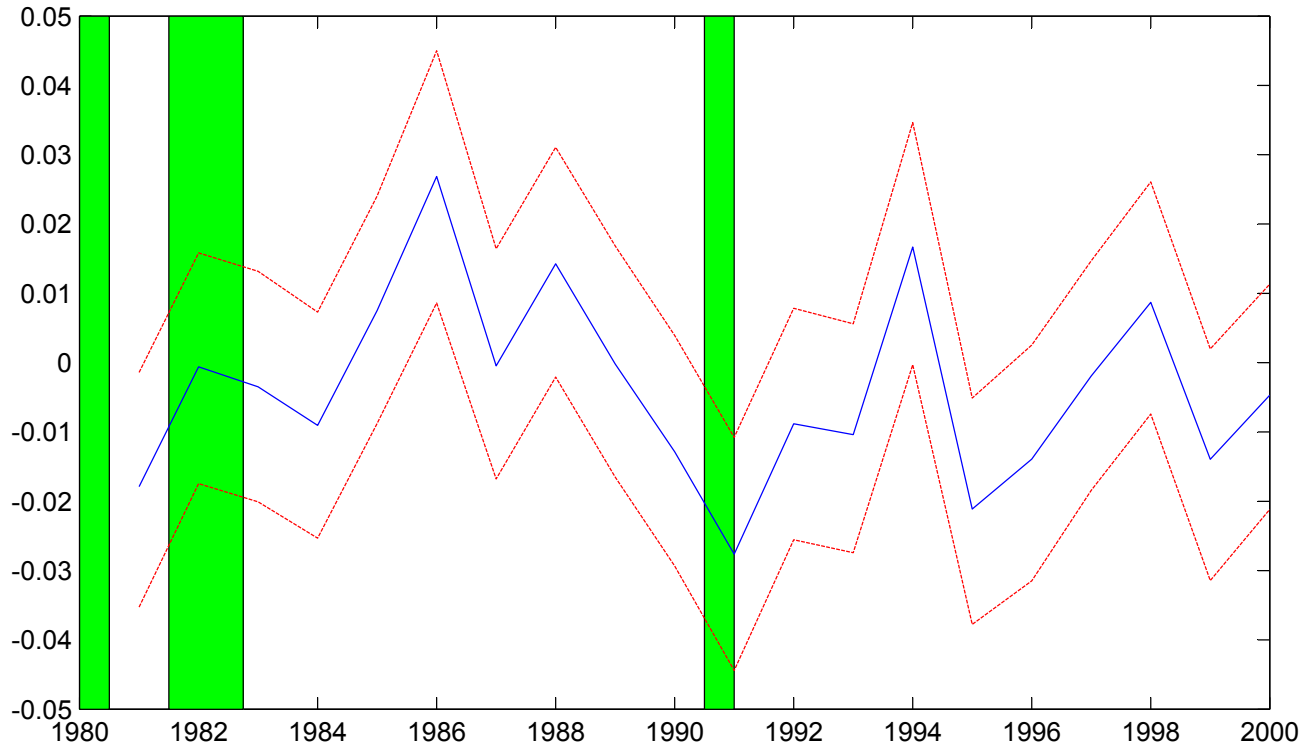## Income inequality without permanent idiosyncratic inequality

Impact of permanent idiosyncratic inequality on income inequality

Figure 8.

permanent aggregate shocks v_t

variance permanent idiosyncratic shocks $Var_t(n)$