

June 2005

Does Income Inequality Lead to Consumption Equality? Evidence and Theory

Dirk Krueger*

Goethe University Frankfurt,
University of Pennsylvania,
National Bureau of Economic Research,
and Centre for Economic Policy Research

Fabrizio Perri*

Federal Reserve Bank of Minneapolis,
New York University,
National Bureau of Economic Research,
and Centre for Economic Policy Research

ABSTRACT

Using data from the Consumer Expenditure Survey, we first document that the recent increase in income inequality in the United States has not been accompanied by a corresponding rise in consumption inequality. Much of this divergence is due to different trends in within-group inequality, which has increased significantly for income but little for consumption. We then develop a simple framework that allows us to analytically characterize how within-group income inequality affects consumption inequality in a world in which agents can trade a full set of contingent consumption claims, subject to endogenous constraints emanating from the limited enforcement of intertemporal contracts (as in Kehoe and Levine, 1993). Finally, we quantitatively evaluate, in the context of a calibrated general equilibrium production economy, whether this setup, or alternatively a standard incomplete markets model (as in Aiyagari, 1994), can account for the documented stylized consumption inequality facts from the U.S. data.

*We thank the editor, Fabrizio Zilibotti, three anonymous referees, and seminar participants at numerous institutions and conferences for their many and very helpful comments. Many thanks also to Harris Dellas and Pierre-Olivier Gourinchas for their thoughtful discussions and to Carlos Serrano and Cristobal Huneus for outstanding research assistance. Krueger gratefully acknowledges financial support from the National Science Foundation under grant SES-0004376. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

The sharp increase in earnings and income inequality for the United States in the last 25 years is a well-documented fact. Many authors have found that the dispersions of U.S. household earnings and incomes have a strong upward trend, attributable to increases in the dispersion of the permanent component of income as well as to an increase in the volatility of the transitory component of income.¹ If one is interested in the welfare impact of these changes, however, the distribution of current income might not be a sufficient statistic. Since a significant fraction of variations of income appear to be due to variations in its transitory component, current income may not be the appropriate measure of lifetime resources available to agents; and thus the distribution of current income might not be a good measure of how economic welfare is allocated among households in the United States.² Moreover, the same change in current or permanent income inequality might have a very different impact on the welfare distribution, depending on the structure of credit and insurance markets available to agents for smoothing income fluctuations. For these reasons, several authors have moved beyond income and earnings as indicators of well-being and have studied the distribution of individual or household *consumption*. Contributors include Cutler and Katz (1991a,b), Johnson and Shipp (1991), Johnson and Smeeding (1998), Mayer and Jencks (1993), Slesnick (1993, 2001), Deaton and Paxson (1994), Dynarski and Gruber (1997), Blundell and Preston (1998), and Krueger and Perri (2004).

Our paper follows this line of research and aims at making three contributions, one empirical, one theoretical, and one quantitative in nature. On the empirical side it investigates how the cross-sectional income and consumption distribution in the United States developed over the period 1980–2003. Using data from the Consumer Expenditure Survey, the paper extends and complements the studies mentioned above. Our main finding is that despite the surge in income inequality in the United States in the last quarter century, consumption inequality has increased only moderately. In particular, income inequality has increased substantially both between and within groups of households with the same characteristics (such as education, sex, and race), but even though between-group consumption inequality has tracked between-group

¹See, e.g., Gottschalk and Moffitt (1994), Gottschalk and Smeeding (1997), or Katz and Autor (1999) for recent surveys of these empirical findings.

²Blundell and Preston (1998) provide theoretical conditions under which the cross-sectional distribution of current consumption is a sufficient statistic for the cross-sectional distribution of welfare.

income inequality quite closely, within-group *consumption* inequality has increased much less than within-group income inequality.

Second, we propose a theoretical explanation for these stylized facts. Our hypothesis is that an increase in the volatility of idiosyncratic labor income (which we identify as the increase in within-group inequality) not only has been an important factor in the increase in income inequality, but also has caused a change in the development of financial markets, allowing individual households to better insure against these (now bigger) idiosyncratic income fluctuations. We present a simple model with endogenous debt constraints (henceforth referred to as the debt constraint markets (DCM) model), building on earlier work by Alvarez and Jermann (2000), Kehoe and Levine (1993, 2001), and Kocherlakota (1996), that allows us to *analytically* characterize the relationship between within-group income and consumption inequality. In the model agents enter risk-sharing contracts, but at any point in time have the option to renege on their obligations, at the cost of losing their assets and being excluded from future risk sharing. Our main result is that an increase in the volatility of income, keeping the persistence of the income process constant, *always* leads to a *smaller* increase in consumption inequality within the group that shares income risk, unless there is no capital income in the economy. The nondegenerate range of income dispersion is such that an increase in income volatility leads to a decline in consumption inequality. Intuitively, we know that higher income volatility increases the value of risk-sharing opportunities and therefore reduces the incentives to default. As a consequence, more risk sharing is possible and the consumption distribution fans out less than the income distribution (and may even “fan in”). This model captures, in a simple and analytically tractable way, the idea that the structure of the credit markets in an economy is endogenous and that, in response to higher income volatility, credit markets have more value and thus will tend to deepen.³

Finally, we assess whether an extension of this simple model is *quantitatively* consistent with the stylized facts established in the empirical section of the paper. We develop a production economy with capital and a large number of agents that face a stochastic labor income process.

³The endogenous response of credit markets to income risk has interesting policy implications. In Krueger and Perri (1999) we show that in the DCM model public insurance (unemployment insurance, progressive taxes, etc.) may crowd out private insurance, possibly more than one-for-one. Empirical studies by Cutler and Gruber (1996) and Albarran and Attanasio (2003) find a sizeable crowding-out effect of public insurance programs.

We choose this income process to match the level and trend of income inequality, both between and within different groups. In particular, we also allow for changes in income inequality that are *not* due to changes in income volatility. The extent to which agents can borrow to insulate consumption from idiosyncratic income fluctuations is derived endogenously. It is a function of the volatility of the stochastic income process, which, as before in the simple model, affects the incentives to repay loans by determining how valuable future access to credit markets is. Our model, for a given time series of cross-sectional income distributions, produces a time series of cross-sectional between- and within-group consumption distributions. We also evaluate the quantitative implications for within- and between-group consumption inequality of a standard incomplete markets model (henceforth referred to as the SIM model) along the lines of Aiyagari (1994). We find that the DCM model slightly understates the increase in consumption inequality and the SIM model somewhat overstates that increase, relative to the data.

The paper is organized as follows. In Section 2 we document our main stylized facts. Section 3 presents the simple model, and Section 4 lays out the economy we use for the quantitative analysis. Section 5 describes our quantitative experiment and parameter choices, and Section 6 presents and discusses the results. Section 7 concludes. Appendix A contains more details about the data and Appendix B more details about computational issues.

2 Trends in Income and Consumption Inequality

In this section we document how income and consumption inequality have evolved in the United States during the last 25 years. For this purpose we use the Consumer Expenditure (CE) Interview Survey, which is the only micro-level data set for the United States that reports comprehensive measures of consumption expenditures and earnings for a repeated large cross section of households.⁴

2.1 The Consumer Expenditure Survey

The CE Interview Survey is a rotating panel of households that are selected to be representative of the U.S. population. It started in 1960, but continuous data are available only from the

⁴The Panel Study of Income Dynamics (PSID) reports both income and consumption data. The consumption data, however, contain only food consumption and therefore are of limited use for our analysis.

first quarter of 1980 until the first quarter of 2004. Each quarter the survey contains detailed information on quarterly consumption expenditures for all households interviewed during that quarter. After a first preliminary interview, each household is interviewed for a maximum of four consecutive times. In the second and fifth interviews, household members are asked questions about earnings, other sources of income, hours worked, and taxes paid for the past year.

2.2 Income and Consumption

Our measure of income is meant to capture all sources of household revenues that are exogenous (to a large extent) to the consumption and saving decisions of households (which are endogenous in our models). Therefore we define income as labor earnings after taxes plus transfers (from now on, LEA+ income). We measure after-tax labor earnings as the sum of wages and salaries of all household members, plus a fixed fraction of self-employment farm and nonfarm income,⁵ minus reported federal, state, and local taxes (net of refunds) and Social Security contributions. We then add reported government transfers (unemployment insurance, food stamps, and welfare).

Our measure of consumption is meant to capture the flow of consumption services that accrue to a household in a given period. For nondurable or small semidurable goods as well as services, expenditures are a good approximation for that flow. For large durable goods such as cars and houses, the relation between current expenditures and consumption service flows is less direct. Thus we impute service flows from the (value of the) stock of durables of a household. Our measure of service flows from housing is the rent paid by the households who indeed rent their home and the self-reported (by the CE respondent) hypothetical rent by the households who own. Our measure of service flows of cars is a fixed fraction (1/32) of the value of the stock of vehicles owned by the household. Since we do not have direct information on the value of the stock of cars, we follow closely the procedure used by Cutler and Katz (1991a) and use information from households who currently purchase vehicles (and for which we therefore observe the value of the purchase) to impute the value of the stock of vehicles for all households. Our benchmark measure of a household's consumption is then the sum of expenditures on nondurables, services, and small durables (such as household equipment), plus imputed services from housing and vehicles. Each expenditure component is deflated by expenditure-specific,

⁵The exact fraction is 0.864 and is taken from Díaz-Giménez, Quadrini, and Ríos-Rull (1996).

quarter-specific consumer price indexes (CPIs). From now on we label this benchmark measure ND+ consumption.⁶

As we are interested in the distribution of resources per capita, before computing inequality measures we divide household income and consumption by the number of adult equivalents in the household using the census equivalence scale (see Dalaker and Naifeh, 1998).

2.3 Sample Selection

Our objective is to characterize the link between labor income and consumption inequality. Therefore we want to select a benchmark sample of households for which we have reliable data for both labor income and consumption for the same time interval. To do so we restrict our sample to households that are complete income respondents and interviewed four times.⁷ For these households, income measured in the fifth interview and the sum of consumption figures reported in the second through the fifth interviews are our measures of their yearly income and consumption. For comparability with previous inequality studies, we performed additional sample selections, such as excluding elderly households, rural households, and households whose reference person reports an implausibly low real wage.⁸

In Table A2 in the appendix we report the benchmark sample sizes for every year in the period 1980–2003, along with weighted averages for income and consumption. Note that the data display no growth of expenditures on nondurables over time, as Slesnick (2001) already highlights. This is puzzling since aggregate nondurable consumption expenditures from the National Income and Product Accounts (NIPA) show significant growth (see again Slesnick, 2001). This might be a signal for growing underreporting of nondurable consumption expenditures in the CE. Also note, however, that ND+ consumption includes services from durables which, on average, are almost as large as expenditures on nondurables (see again Table A2) and display a growth trend over time that matches up better with NIPA data. So if underreporting of consumption exists, it is likely to be less severe for our benchmark ND+ consumption measure than for the more commonly used nondurable consumption expenditures.

⁶In Appendix A we provide a detailed description of all the items included in our consumption measures and of our imputation and deflation procedures.

⁷The CE classifies as incomplete income respondents those households who report zero income for all the major income categories, suggesting nonreliability of their earning figures (see also Nelson, 1994).

⁸See Appendix A for a precise list of our sample restrictions.

2.4 Inequality Trends

Figure 1 displays the trend for four commonly used measures of cross-sectional inequality, computed on the benchmark sample and on the measures of income and consumption described above. All measures are computed using CE population weights. We report the Gini coefficient, the variance of the logs, and the 90/10 and 50/10 ratios. In each panel the dashed line represents inequality in LEA+ income and the solid line represents inequality in ND+ consumption. Finally the thin dash-dotted lines are standard errors of the inequality measures, which we computed by performing a bootstrap procedure with 100 repetitions.

Figure 1 confirms the well-known fact that labor income inequality in the United States has increased significantly in the last quarter century: the Gini index has risen from around 0.3 to around 0.37, and the variance of the logs displays an increase of more than 20%. The 90/10 ratio for income surges from around 4.2 to over 6, suggesting a large divergence between the two tails of the income distribution over time. Finally the 50/10 ratio displays an increase from 2.2 to about 2.7, revealing that households in the bottom tail of the income distribution have lost ground relative to the median.⁹

The figure also presents our main empirical finding, namely that the increase in *consumption* inequality has been much less marked;¹⁰ the increase has been from 0.23 to 0.26 for the Gini and about 5% in the variance of logs. The 90/10 ratio has increased from 2.9 to around 3.4, suggesting a much more moderate fanning out of the consumption distribution.¹¹ Finally, the 50/10 ratio increases only from about 1.7 to 1.9, suggesting that in terms of consumption, households in the bottom part of the distribution have lost less ground relative to the median.¹²

Note that our income definition includes government taxes and transfers, so that changes in government income redistribution policies cannot be responsible for the divergence between the two series. Although the evolution of consumption inequality has been studied less than

⁹Increases of similar magnitude are found in other cross-sectional data sets. Krueger and Perri (2004) compare the increase in wage inequality using CE data with that obtained by using PSID data (from Heathcote et al., 2004) and the increase measured by using the Current Population Survey (CPS) data (from Katz and Autor, 1999) and find that, for the same sample selection, the magnitude of the increase is very similar. This suggests that the quality of the CE earnings/wage data is comparable to those of other cross-sectional data sets.

¹⁰Pendakur (1998) finds similar results for Canada for 1978–1992 and his preferred measure of consumption.

¹¹One nice property of the 90/10 ratios is that they are not sensitive to changes in top-coding thresholds. The divergence of the 90/10 ratios in income and consumption thus suggests that changes in top-coding thresholds play no important role in explaining the measured divergence in inequality.

¹²These findings are consistent with those of Slesnick (2001), who found that poverty rates for income increased from 11.1% in 1973 to 13.8% in 1995, while poverty rates for consumption declined from 9.9% to 9.5%.

the evolution of income inequality, some authors (Cutler and Katz, 1991a,b, and Johnson and Shipp, 1991) have noted that the sharp increase in income inequality of the *early 1980s* has been accompanied by an increase in consumption inequality. Our measures also display an increase in consumption inequality in the early 1980s, but as noted by Slesnick (2001), it is less marked than the increase in income inequality; moreover, in the 1990s income inequality has continued to rise (although at a slower pace) while consumption inequality has remained substantially flat. This last fact has also been reported by Federal Reserve chairman Alan Greenspan (1998) in his introductory remarks to a symposium on income inequality. Attanasio et al. (2005) have recently looked at consumption distributions using the CE Diary Survey, which surveys a different group of households and collects information on consumption of small items frequently purchased, such as food and personal care items. They find that for comparable categories in the Diary and Interview Surveys, mean per capita consumption is very similar but consumption inequality grows more in the Diary Survey. Based on this finding, they construct a measure of the variance of logs of nondurable consumption that uses information from both the Diary and the Interview Surveys and find that it increases about 4.6% over the period 1986–2000. By contrast our measure, based solely on the Interview Survey, displays an increase of about 2.5%. These increases in inequality are different, but they are both significantly smaller than the increase in variance of log income (which over the same period was over 12%). Moreover, we conjecture that, due to its limited consumption coverage, the impact of using the Diary Survey is likely to be even smaller if one focuses on a broader definition of consumption such as our benchmark ND+ consumption. In the next subsection we check the robustness of the trends just described to alternative definitions of consumption and to alternative sample selection choices.

2.5 Alternative Definitions and Samples

Panel (a) of Table 1 reports the increase in consumption inequality (measured as the variance of logs) from the first two years of our sample (1980–1981) to the last two (2002–2003) obtained for definitions of consumption expenditures. As a reference point, in the first two columns we again report the increase in inequality for our benchmark measures of income (LEA+) and

consumption (ND+).¹³ The third and fourth columns report the increase in inequality in food consumption and nondurable consumption (the definition of consumption used by Attanasio and Davis, 1996, among others). For these definitions of consumption, the increase in inequality is slightly smaller than for our benchmark consumption (ND+) measure. Finally, the last column of panel (a) reports the change in inequality for total consumption expenditures (TCE). For this definition the increase in consumption inequality is larger, although still less than half of the increase in income inequality. One should keep in mind, however, that this consumption measure includes cash payments for homes and vehicles, and therefore contains a significant part of households' savings, which biases measured consumption inequality toward measured income inequality. In addition, this measure, being based on expenditures on durables rather than on service flows from these durables, is affected by changes in the frequency of durables purchases over time. For these reasons we think of the latter statistic as an upper bound for the true change in consumption inequality rather than as the best estimate of it.

TABLE 1. CHANGES IN INEQUALITY

(a) Alternative definitions					
	Income	Consumption			
	LEA+	ND+	Food	ND	TCE
Change in var. log	21.4%	5.3%	2.3%	3.3%	10.4%
(b) Alternative samples					
	Benchmark	Quarterly	Inc. Rural	Inc. Low Wages	
Change in var. log ND+	5.3%	6.3%	4.5%	5.5%	
Change in var. log LEA+	21.4%	22.0%	20.5%	27.5%	

Panel (b) of Table 1 reports the increase in LEA+ income and ND+ consumption inequality as computed using different sample selections. The first column contains the results for our benchmark sample. In the second column we report inequality measures computed by using quarterly (as opposed to yearly) consumption measures. These numbers also use information from households that are interviewed in the CE *less* than four times, thus increasing the sample

¹³In the remainder of the paper we focus on the variance of logs as our main measure of inequality. Therefore, in Table 1 we restrict attention to this measure. Using other measures of inequality yields similar results.

size significantly.¹⁴ It has the disadvantage, however, that for any given household, income and consumption information do not refer to the same period. The third and fourth columns report measures for samples that include rural households and households with a reference person reporting a very low wage. We observe that with all sample selection criteria employed, the increase in consumption inequality is quite similar and much smaller than the increase in income inequality.¹⁵

2.6 Between- and Within-Group Income and Consumption Inequality Trends

Before turning to the theoretical explanation for the empirical findings, we further investigate the nature of the change in income and consumption inequality by decomposing them into between- and within-group inequality. Between-group inequality is attributable to fixed (and observable) characteristics of the household (for example, education, experience, and sex). Although between-group inequality changes over time (returns to these characteristics can change over time, as in the case of the increase in the college premium), it is unlikely that households can insure against these changes; thus, increases in between-group inequality should translate into similar increases in between-group consumption inequality.

Within-group income inequality is a residual measure that also includes inequality caused by idiosyncratic income shocks. Therefore, increases in within-group income inequality are (at least partly) attributable to an increase in the volatility of idiosyncratic income shocks. In the models discussed in the next sections, the main question is how well households can insulate their consumption from an increase in the volatility of these idiosyncratic income shocks. The better households can insure against these shocks the less we expect within-group consumption inequality to increase in response to an increase in within-group income inequality. Therefore, we now empirically measure the changes in both within-group income and consumption inequality.

The empirical decomposition we employ is simple and commonly used. Following Katz and Autor (1999), for each labor income and consumption expenditure cross section (after controlling for age effects), we regress income and consumption on the following characteristics of the

¹⁴Our benchmark sample has an average of 6,660 quarter/household data points per year, and the quarterly sample has an average of 11,300 quarter/household data points per year.

¹⁵We also experimented with per-household (as opposed to per-adult-equivalent) income and consumption measures and with different equivalence scales. These changes affect the level of inequality measures but have very little effect on the trends.

reference person and the spouse (if present): sex, race, years of education, experience, interaction terms between experience and education, dummies for managerial/professional occupation, and region of residence. These characteristics explain about 25% of the cross-sectional variation of income and consumption in 1980. We denote the cross-sectional variance explained by these characteristics as “between-group” inequality and the residual variance as “within-group” inequality. By construction the two variances sum to the total variance.

Figure 2 shows the evolution of between-group (panel a) and within-group (panel b) inequality for income and consumption. As documented by many studies, for income both the between- and within-group components display an increase. Panel (a) shows that for consumption, the between-group component displays an increase similar in magnitude to that of income.¹⁶ Panel (b) reveals a very different picture for the within-group component: the increase in consumption inequality is an order of magnitude smaller than the increase in income inequality. Consequently, understanding the trends in panel (b) is crucial for understanding the patterns of income and consumption inequality in the United States.

In the next section we present a simple model in which we *analytically* characterize the relation between income and consumption inequality within a group of ex ante identical agents and show how the endogenous expansion of risk sharing may lead to a small increase (or even a decline) in within-group consumption inequality in the wake of increasing income inequality.

3 A Simple Model

We analyze a pure exchange economy similar to Kocherlakota (1996), Alvarez and Jermann (2000), and Kehoe and Levine (2001). Time is discrete and the number of time periods is infinite. Each period has two (types of) agents $i = 1, 2$ and a single, nonstorable consumption good. Agents obtain endowments of the consumption good from two sources: labor and capital “income.” This distinction, although not necessary in the simple model, helps to explain our quantitative results in the next section. First, an agent receives endowments in the form of stochastic labor income. If one consumer has labor income $1 + \varepsilon$, the other has $1 - \varepsilon$, so that the

¹⁶This finding is highly consistent with the results of Attanasio and Davis (1996), which suggest that changes in relative wages between education groups are fully reflected in consumption changes of these groups. We will revisit this point below in our model-based quantitative exercise.

aggregate endowment from labor is constant at 2 in each period. Let $s_t \in S = \{1, 2\}$ denote the consumer that has labor income $1 + \varepsilon$. We assume that $\{s_t\}_{t=0}^\infty$ is a sequence of i.i.d. random variables with $\pi(s_t = 1) = \pi(s_t = 2) = \frac{1}{2}$, so that households are ex ante identical. Note that the parameter $\varepsilon \in [0, 1)$ measures the variability of the income process.

In addition, two trees, one initially owned by each agent, each yield a constant endowment of r per period. Thus, aggregate endowment from capital equals $2r$, total endowment per period is constant at $2(1 + r)$, and the capital share is given by $\frac{r}{1+r}$.

Let $s^t = (s_0, \dots, s_t)$ denote an event history and $\pi(s^t)$ the time 0 probability of event history s^t . An allocation $c = (c^1, c^2)$ maps event histories s^t into consumption. Agents have preferences representable by

$$U(c^i) = (1 - \beta) \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \pi(s^t) u(c_t^i(s^t)),$$

where $\beta < 1$ and u is continuous, twice differentiable, strictly increasing and strictly concave on $(0, \infty)$, and satisfies the Inada condition $\lim_{c \rightarrow 0} u'(c) = \infty$. Let us define as

$$U(c^i, s^t) = (1 - \beta) \sum_{\tau=t}^{\infty} \sum_{s^\tau | s^t} \beta^{\tau-t} \pi(s^\tau | s^t) u(c_\tau^i(s^\tau))$$

the continuation utility of agent i from allocation c^i , from event history s^t onward, and denote by $e = (e^1, e^2)$ the autarkic allocation of consuming the *labor* endowment in each event history.

In this economy both agents have an incentive to share their endowment risk. We assume, however, that at any point in time both agents have the option of renegeing on the risk-sharing arrangement obligations and bearing the associated costs, specified as exclusion from intertemporal trade and loss of any tree in their possession. This implies that any risk-sharing mechanism must yield allocations that deliver to each consumer a continuation utility at least as high as the autarkic allocation, for all histories s^t . Formally, we impose the individual rationality constraints

$$U(c^i, s^t) \geq U(e^i) = (1 - \beta) \sum_{\tau=t}^{\infty} \sum_{s^\tau | s^t} \beta^{\tau-t} \pi(s^\tau | s^t) u(e_\tau^i(s^\tau)) \quad \forall i, s^t. \quad (1)$$

We say that an allocation (c^1, c^2) is constrained efficient if it satisfies the resource constraint

$$c^1 + c^2 = 2(1 + r) \quad (2)$$

and the individual rationality constraints (1). Alvarez and Jermann (2000) show how constrained efficient allocations can be decentralized as competitive equilibria with state-dependent borrowing constraints. Now we study the cross-sectional consumption distribution associated with a constrained efficient allocation; we are particularly interested in how this distribution changes in response to an increase in income volatility, as measured by ε .

3.1 The Constrained Efficient Consumption Distribution

We focus on symmetric allocations.¹⁷ In order to analyze how constrained efficient consumption allocations vary with ε , we now solve for the continuation value of autarky, which is given by

$$\begin{aligned} U(1 + \varepsilon) &= (1 - \beta) u(1 + \varepsilon) + \frac{\beta}{2} [u(1 + \varepsilon) + u(1 - \varepsilon)] \\ U(1 - \varepsilon) &= (1 - \beta) u(1 - \varepsilon) + \frac{\beta}{2} [u(1 + \varepsilon) + u(1 - \varepsilon)]. \end{aligned}$$

Here $U(1 + \varepsilon)$ and $U(1 - \varepsilon)$ denote the continuation utility of autarky for the agents with the currently high income and the currently low income, respectively. The continuation utility from autarky is a convex combination of utility obtained from consumption today, $(1 - \beta) u(1 + \varepsilon)$ or $(1 - \beta) u(1 - \varepsilon)$, and the expected utility from tomorrow onward. The next lemma, whose proof is straightforward and hence omitted, states properties of $U(1 + \varepsilon)$ as a function of income variability ε . Define $U^{FB}(r) = u(1 + r)$ as the lifetime utility of the first best allocation in which there is complete risk sharing and consumption of both agents is constant at $1 + r$.

Lemma 1 *$U(1 + \varepsilon)$ is strictly increasing in ε at $\varepsilon = 0$, is strictly decreasing in ε as $\varepsilon \rightarrow 1$, and is strictly concave in ε , with a unique maximum $\varepsilon_1 = \arg \max_{\varepsilon} U(1 + \varepsilon) \in (0, 1)$. Furthermore:*

1. *Either $U(1 + \varepsilon_1) \leq U^{FB}(r)$ or*
2. *$U(1 + \varepsilon_1) > U^{FB}(r)$. In this case there exist $\underline{\varepsilon}(r) < 1$ such that $U(1 + \underline{\varepsilon}(r)) = U^{FB}(r)$. The number $\underline{\varepsilon}(r)$ is strictly increasing in r . Furthermore, either there exists a number $\bar{\varepsilon}(r) \in (\underline{\varepsilon}(r), 1)$ such that $U(1 + \bar{\varepsilon}(r)) = U^{FB}(r)$ and the number $\bar{\varepsilon}(r)$ is strictly decreasing in r , or $U(2) > U^{FB}(r)$, in which case we take $\bar{\varepsilon}(r) = 1$.*

¹⁷Consider two histories s^t, \tilde{s}^t that satisfy, for all $\tau \leq t$,

$$s_\tau = 1 \text{ if and only if } \tilde{s}_\tau = 2.$$

A consumption allocation is symmetric if $c_t^1(s^t) = c_t^2(\tilde{s}^t)$ for all such histories s^t, \tilde{s}^t .

The nonmonotonicity of $U(1 + \varepsilon)$, shown in Figure 3, stems from two opposing effects. For small ε the direct effect of higher consumption today outweighs the higher risk faced by the agent from tomorrow onward and $U(1 + \varepsilon)$ increases with ε . As ε becomes larger and future consumption more risky, $U(1 + \varepsilon)$ declines with ε as the risk effect dominates the direct effect. On the other hand, the value of autarky for the agent with currently low income, $U(1 - \varepsilon)$, is strictly decreasing (and concave) in ε (also plotted in Figure 3), since an increase in ε reduces consumption today for this agent and makes it more risky from tomorrow onward.

By using these properties of the continuation utilities from autarky and the results by Alvarez and Jermann (2000) and Kehoe and Levine (2001) (in particular their Proposition 5), one immediately obtains the following characterization of the consumption distribution for this economy.

Proposition 2 *The constrained efficient symmetric consumption distribution is completely characterized by a number $\varepsilon_c(\varepsilon) \geq 0$. Agents with labor income $1 + \varepsilon$ consume $1 + r + \varepsilon_c(\varepsilon)$, and agents with labor income $1 - \varepsilon$ consume $1 + r - \varepsilon_c(\varepsilon)$. The number $\varepsilon_c(\varepsilon)$ is the smallest nonnegative solution of the following equation:*

$$U(1 + r + \varepsilon_c(\varepsilon)) = \max(U^{FB}(r), U(1 + \varepsilon)), \quad (3)$$

and $U(1 + r + \varepsilon_c(\varepsilon))$ is the lifetime utility of the consumption allocation characterized by $\varepsilon_c(\varepsilon)$.

The intuition for this result is simple: in any efficient risk-sharing arrangement, the currently rich agent has to transfer resources to the currently poor agent. To prevent this agent from defaulting, she needs to be awarded sufficiently high current consumption in order to be made at least indifferent between the risk-sharing arrangement and the autarkic allocation. The proposition simply states that the efficient consumption allocation features maximal risk sharing, subject to providing the currently rich agent with sufficient incentives not to walk away.

Note that if $U^{FB}(r) \geq U(1 + \varepsilon)$, the smallest solution to equation (3) is $\varepsilon_c(\varepsilon) = 0$ and the constrained efficient allocation implies full risk sharing. This also implies that, since the value of complete markets is strictly increasing in capital income r , a higher r expands the region of income dispersions ε for which perfect risk sharing is feasible and thus constrained efficient. Also

note that unless $r = 0$, autarky is never constrained efficient, since the equation

$$U(1 + r + \varepsilon_c(\varepsilon)) = U(1 + \varepsilon)$$

is never solved by $\varepsilon_c(\varepsilon) = \varepsilon$, unless $r = 0$.

3.2 Income Variability and Consumption Inequality

We now characterize how the constrained efficient consumption distribution varies with the variability of income, ε . Remember that ε_1 was defined as a unique maximizer of $U(1 + \varepsilon)$.

Proposition 3 *Fix $\beta \in (0, 1)$ and $r \geq 0$.*

1. *If $U(1 + \varepsilon_1) \leq U^{FB}(r)$, then perfect consumption insurance is feasible for all ε and a change in ε has no effect on consumption inequality.*
2. *If $U(1 + \varepsilon_1) > U^{FB}(r)$, then for $\varepsilon \in [0, \underline{\varepsilon}(r))$ and $\varepsilon \in [\bar{\varepsilon}(r), 1)$ perfect consumption insurance is feasible and a marginal increase in ε has no effect on consumption inequality. If $\varepsilon \in [\varepsilon_1, \bar{\varepsilon}(r))$ a marginal increase in ε leads to a reduction in consumption inequality, whereas for $\varepsilon \in [\underline{\varepsilon}(r), \varepsilon_1)$ a marginal increase in ε increases consumption inequality. If $r > 0$, the increase in consumption inequality is strictly smaller than the increase in income inequality.*

The proof of this proposition follows immediately from Proposition 1 and the properties of $U(1 + \varepsilon)$ stated in Lemma 1, apart from the very last part. The fact that consumption inequality always increases less than income inequality is obvious for the regions $\varepsilon \in [0, \underline{\varepsilon}(r))$, $\varepsilon \in [\bar{\varepsilon}(r), 1)$, and $\varepsilon \in [\varepsilon_1, \bar{\varepsilon}(r))$, since in these regions consumption inequality does not change or is even declining in income volatility. For the region $\varepsilon \in [\varepsilon_1, \bar{\varepsilon}(r))$ we have

$$U(1 + r + \varepsilon_c(\varepsilon)) = U(1 + \varepsilon)$$

and thus by the implicit function theorem

$$\frac{d\varepsilon_c(\varepsilon)}{d\varepsilon} = \frac{U'(1 + \varepsilon)}{U'(1 + r + \varepsilon_c(\varepsilon))} \in (0, 1)$$

since $U(\cdot)$ is strictly concave, and efficient risk sharing implies $r + \varepsilon_c(\varepsilon) < \varepsilon$ for $r > 0$.

Figure 3 provides some intuition for the proposition above. In the top panel we plot the value of autarky in the two states, the value of full risk sharing and $U(1 + r + \varepsilon)$, and in the bottom panel we plot income and consumption dispersion as a function of income dispersion ε .

As shown in the top panel, we see that for $\varepsilon \in [0, \underline{\varepsilon}(r))$ and $\varepsilon \in [\bar{\varepsilon}(r), 1)$, $U^{FB}(r) > U(1 + \varepsilon)$, and thus the first best allocation can be implemented. In this case, as shown in the bottom panel, consumption inequality does not vary with income inequality.

Suppose now that $\varepsilon_1 < \varepsilon < \bar{\varepsilon}(r)$. For example, consider the point $\varepsilon = \varepsilon_a$ on the x-axis; from Proposition 1 the constrained efficient consumption allocation is given by the smallest solution to $U(1 + r + \varepsilon_c(\varepsilon_a)) = U(1 + \varepsilon_a)$. The top panel of Figure 3 displays the solution $\varepsilon_c(\varepsilon_a)$. In this allocation, which involves partial risk sharing as $\varepsilon_c(\varepsilon_a) < \varepsilon_a$, the agent with high income receives a continuation utility equal to the value of autarky, whereas the agent with low income receives $U(1 + r - \varepsilon_c(\varepsilon_a))$, strictly higher than its value of autarky ($U(1 - \varepsilon_a)$). In this range a marginal increase in income inequality reduces the value of autarky for the high-income agent and less current consumption is required to make her not default ($\varepsilon_c(\varepsilon_a)$ moves to the left). This reduces consumption dispersion in the economy, as shown in the bottom panel of the figure.

Finally, in the range $\underline{\varepsilon}(r) < \varepsilon < \varepsilon_1$ (consider, for example, the point $\varepsilon = \varepsilon_b$ in the figure), the constrained efficient allocation is characterized by $\varepsilon_c(\varepsilon_b)$. In this case a marginal increase in ε increases the value of autarky for the constrained agent, and so her current consumption has to increase to prevent her from defaulting: consumption inequality increases.

To summarize, in this environment with limited commitment an increase of income dispersion always leads to a smaller increase in consumption dispersion as long as there is some capital income. It may even lead to a reduction in consumption dispersion. The intuition behind these results is that an increase in income inequality, by making exclusion from future risk sharing more costly, renders the individual rationality constraint less binding. It thereby allows individuals to share risk to a larger extent and thus reduces fluctuations in their consumption profiles. It is crucial for this result that income shocks are not perfectly permanent (although they may be highly persistent) because the fear of being poor again in the future is what makes a currently rich agent transfer resources to his currently poor brethren.¹⁸ This analysis suggests

¹⁸It is straightforward to generalize our results to a serially correlated endowment process. An increase in persistence leads to an increase in consumption dispersion in the constrained efficient consumption distribution. This increase is strict if initially there is some, but not complete, risk sharing. For a proof of this result, see Kehoe

that the endogenous evolution of (formal, market-based, and informal) risk-sharing mechanisms can indeed generate a modestly increasing or even declining within-group consumption inequality despite a substantially increasing within-group income inequality.

3.3 Capital Income and the Extent of Risk Sharing

Finally, we show how the extent of risk sharing depends on how abundant the capital income r . Since we will study a production economy with capital in our quantitative exercise, it is instructive to provide some intuition for how the presence (and magnitude) of capital income affects the extent to which households can share risk. We find that risk sharing is increasing in r , strictly so if risk sharing is not perfect (we already argued above that the region of ε for which perfect risk sharing obtains is strictly larger the larger the capital income r).

Proposition 4 *Let $\varepsilon_c(\varepsilon; r)$ characterize the constrained efficient consumption allocation, as a function of capital income r . Then if $\hat{r} > r$, we have*

$$\varepsilon_c(\varepsilon; \hat{r}) \leq \varepsilon_c(\varepsilon; r)$$

for all $\varepsilon \in (0, 1)$, with inequality the strict if and only if $\varepsilon_c(\varepsilon; r) > 0$. That is, more risk sharing is possible with capital income \hat{r} than with r .

Proof. *The only-if part is obvious, since $0 \leq \varepsilon_c(\varepsilon; \hat{r}) < \varepsilon_c(\varepsilon; r)$. For the if part, if $\varepsilon_c(\varepsilon; \hat{r}) = 0$ the result follows. So suppose $\varepsilon_c(\varepsilon; \hat{r}) > 0$. Then*

$$U(1 + \hat{r} + \varepsilon_c(\varepsilon, \hat{r})) = U(1 + \varepsilon). \quad (4)$$

But if $\varepsilon_c(\varepsilon; \hat{r}) \geq \varepsilon_c(\varepsilon; r) > 0$ (perfect risk sharing at r is impossible since the perfect risk sharing region is smaller at r than at \hat{r}), then we obtain a contradiction since

$$U(1 + \varepsilon) = U(1 + r + \varepsilon_c(\varepsilon, r)) < U(1 + \hat{r} + \varepsilon_c(\varepsilon, r)) \leq U(1 + \hat{r} + \varepsilon_c(\varepsilon, \hat{r})) = U(1 + \varepsilon). \quad (5)$$

■

and Levine (2001). The intuition is again simple: the value of autarky for the agent with high current income increases (as the agent is more likely to have high income in the future with higher persistence), which makes the individual rationality constraint more stringent and leads to fewer transfers to the poor agent being sustainable.

In the next section we evaluate the quantitative importance of the mechanism of extended consumption insurance due to a relaxation of default constraints just described. For this we employ a production economy with a continuum of agents that face a more realistic income process than in the simple model; in particular, we will also allow for changes in between-group inequality.

4 The Model with a Large Number of Agents

4.1 The Environment

A single good being produced in a given period can be used for consumption or investment in the physical capital stock K . The representative firm produces output according to a Cobb-Douglas production technology $Y = AK^\alpha L^{1-\alpha}$, where L denotes the homogeneous labor input and A is a technology parameter. The aggregate resource constraint reads as

$$C_t + K_{t+1} - (1 - \delta)K_t = AK_t^\alpha L_t^{1-\alpha}, \quad (6)$$

with C_t denoting aggregate consumption and δ denoting the depreciation rate of physical capital.

Labor is inelastically supplied by a continuum of consumers of measure 1. Individuals belong to different groups $i \in \{1, \dots, M\}$, with p_i denoting the fraction of the population in group i . We interpret these different groups of agents as capturing heterogeneity in the population with respect to fixed characteristics that affect an individual's wage and therefore income, such as education or sex. Since we documented above that an important part of the rise in income inequality is due to increased between-group inequality, an incorporation of this type of heterogeneity appears to be critical for any quantitative study of income and consumption inequality.

An individual in group i has a stochastic labor endowment process $\{\alpha_{it}y_t\}$, where α_{it} is the deterministic, group-specific, and possibly time-varying mean labor endowment, and the idiosyncratic component $\{y_t\}$ follows a Markov process with finite support Y_t , a set with cardinality N . Since labor income will be the product of individual labor endowment and an economy-wide wage per efficiency unit of labor, we use the terms *labor endowment* and *labor income* interchangeably. Let $\pi_t(y'|y)$ denote the transition probabilities of the Markov chain, assumed to be

identical for all agents. The set Y_t and the matrix π_t are indexed by t since we allow for the idiosyncratic part of the income process to change over time. Furthermore, we assume a law of large numbers, so that the fraction of agents facing shock y' tomorrow with shock y today in the population is equal to $\pi_t(y'|y)$. Finally, we assume that $\pi_0(y'|y)$ has a unique invariant measure $\Pi(\cdot)$. Let us denote by y_t the current period labor endowment and by $y^t = (y_0, \dots, y_t)$ the history of realizations of endowment shocks; also $\pi(y^t|y_0) = \pi_{t-1}(y_t|y_{t-1}) \cdots \pi_0(y_1|y_0)$. We intend the notation $y^s|y^t$ to mean that y^s is a possible continuation of labor endowment shock history y^t . We furthermore assume that at date zero the measure over current labor endowments is given by $\Pi_0(\cdot)$. At date zero agents are distinguished by their group i , their initial asset holdings (claims to period zero consumption) a_0 , and by their initial labor endowment shock y_0 . Let Φ_0 be the initial distribution over types (i, a_0, y_0) . Thus, total labor supply is given by

$$L_t = \int \sum_{y^t} \alpha_{it} y_t \pi(y^t|y_0) d\Phi_0, \quad (7)$$

that is, by the sum of all labor endowments in the population. Finally, agents' preferences are exactly as described in the simple model of the previous section.

4.2 Market Structures

We now describe the market structure of two incomplete markets economies whose quantitative properties we will contrast with the stylized empirical facts established in Section 2.

4.2.1 Debt Constraint Markets (DCM)

An individual of type (i, a_0, y_0) starts with initial assets a_0 and trades Arrow securities subject to prespecified credit lines $A_t^i(y^t, y_{t+1})$ that are contingent on observable labor endowment histories and an agents' group, and whose exact form is specified below. The prices for these Arrow securities are denoted by $q_t(y^t, y_{t+1})$ and depend only on an agent's own labor endowment history and time, in order to reflect deterministic changes in the income process and hence in the magnitude of labor endowments $\alpha_{it} y_t$.

Consider the problem of an agent of type i with initial conditions (i, a_0, y_0) . The agent chooses, conditional on his labor endowment history, consumption $\{c_t(a_0, y^t)\}$, and one-period

Arrow securities $\{a_{t+1}(a_0, y^t, y_{t+1})\}$ whose payoff is conditional on his own endowment realization y_{t+1} tomorrow, to maximize for given (a_0, y_0)

$$(1 - \beta) \left(u(c_0(a_0, y_0)) + \sum_{t=1}^{\infty} \sum_{y^t|y_0} \beta^t \pi(y^t|y_0) u(c_t(a_0, y^t)) \right) \quad (8)$$

$$\text{s.t. } c_t(a_0, y^t) + \sum_{y_{t+1}} q_t(y^t, y_{t+1}) a_{t+1}(a_0, y^t, y_{t+1}) = w_t \alpha_{it} y_t + a_t(a_0, y^t) \quad \forall y^t \quad (9)$$

$$a_{t+1}(a_0, y^t, y_{t+1}) \geq A_{t+1}^i(y^t, y_{t+1}) \quad \forall y^t, y_{t+1}. \quad (10)$$

Here w_t is the wage rate per effective unit of labor.

Following Alvarez and Jermann (2000), we specify the short-sale constraints $A_t^i(y^t, y_{t+1})$ as “solvency constraints” that are not too tight. As before, let $U_t^{Aut}(i, y_t)$ denote the continuation utility from autarky, given current labor endowment realization $\alpha_{it} y_t$. Given a sequence of prices $\{q_t\}_{t=0}^{\infty}$ and short-sale constraints $\{A_t^i(y^t, y_{t+1})\}_{t=0}^{\infty}$, let us define the continuation utility $V_t(i, a, y^t)$ of an agent of type i with endowment shock history y^t and current asset holdings a at time t as

$$V_t(i, a, y^t) = \max_{\{c_s(a, y^s), a_{s+1}(a, y^s, y_{s+1})\}} (1 - \beta) \left(u(c_t(a, y^t)) + \sum_{s=t+1}^{\infty} \sum_{y^s|y^t} \beta^t \pi(y^s|y^t) u(c_s(a, y^s)) \right)$$

subject to (9) and (10). Short-sale constraints $\{A_t^i(y^t, y_{t+1})\}_{t=0}^{\infty}$ are not “too tight” if they satisfy

$$V_{t+1}(i, A_{t+1}^i(y^t, y_{t+1}), y^{t+1}) = U_{t+1}^{Aut}(i, y_{t+1}) \text{ for all } (y^t, y_{t+1}). \quad (11)$$

That is, the constraints are such that an agent of type i , having borrowed up to maximum, $a_{t+1}(a, y^t, y_{t+1}) = A_{t+1}^i(y^t, y_{t+1})$, is indifferent between repaying his debt and defaulting, with the default consequence being specified as limited future access to financial markets. In contrast to the simple model, we now allow households to at least save after default, at a state-uncontingent interest rate r_d (a parameter of the model). The value of autarky is then given

by

$$U_t^{Aut}(i, y_t) = \max_{\{c_s(a_0, y^s), b_{s+1}(a_0, y^s)\}} (1 - \beta) \left(u(c_t(a_0, y^t)) + \sum_{s=t+1}^{\infty} \sum_{y^s|y^t} \beta^t \pi(y^s|y^t) u(c_s(a_0, y^s)) \right)$$

$$\text{s.t. } c_s(a_0, y^s) + \frac{b_{s+1}(a_0, y^s)}{1 + r_d} = w_s \alpha_{is} y_s + b_s(a_0, y^{s-1}) \quad \forall y^s$$

$$b_{s+1}(a_0, y^s) \geq 0$$

and subject to $b_t(a_0, y^{t-1}) = 0$. That is, the defaulting agent starts her life in financial autarky with no assets and can save (but not borrow) at the interest rate r_d . For $r_d = -1$, the household would optimally never save after default, and thus in this case the value of autarky is determined exactly as in the simple model above.

Definition 1 Given Φ_0, K_0 , a competitive equilibrium with solvency constraints $\{A_t^i(y^t, y_{t+1})\}$ that are not too tight is allocations $\{c_t(i, a_0, y^t), a_{t+1}(i, a_0, y^t, y_{t+1})\}$ for households, allocations $\{K_t, L_t\}$ for the firm, and prices $\{w_t, r_t, q_t(y^t, y_{t+1})\}$ such that

1. (Household Optimization) Given prices, household allocations maximize (8) subject to (9) and (10), and the solvency constraints are not “too tight” (in the sense of 11).

2. (Firm Optimization)

$$w_t = (1 - \alpha) A \left(\frac{K_t}{L_t} \right)^\alpha$$

$$r_t = \alpha A \left(\frac{K_t}{L_t} \right)^{\alpha-1} - \delta$$

3. (Market Clearing) L_t is given by (7), the goods market clearing condition (6) holds, with

$$C_t = \int \sum_{y^t} c_t(i, a_0, y^t) \pi(y^t|y_0) d\Phi_0,$$

and the asset market clearing condition holds, with

$$K_{t+1} = \frac{1}{1 + r_{t+1}} \int \sum_{y^t} a_{t+1}(i, a_0, y^t, y_{t+1}) \pi(y^{t+1}|y_0) d\Phi_0 \equiv \frac{A_{t+1}}{1 + r_{t+1}}.$$

A stationary equilibrium is one in which $\{r_t, w_t\}$ as well as the cross-sectional asset and consumption distributions are constant over time.

Two comments on the equilibrium definition: First, in the asset market clearing condition we divide the right-hand side by the interest rate since the Arrow securities are state-contingent *zero coupon* bonds. Second, no arbitrage implies that

$$q_t(y^t, y_{t+1}) = \frac{\pi(y_{t+1}|y_t)}{1 + r_{t+1}}$$

because households have to be indifferent between saving with risk-free capital or reconstructing a risk-free asset with the full set of Arrow securities. (Note that with the full set of Arrow securities, risk-free capital is a redundant asset for households, so we explicitly abstained from introducing purchases of risk-free capital in the household problem.) Physical capital in our economy is not important as an additional asset, but is important in the sense that it provides the economy as a whole with an asset in positive net supply and therefore a positive wealth-to-income ratio. The simple model above demonstrated that the more abundant the capital (income), the better the extent of consumption insurance achievable in the DCM model.

Notice that the dispersion of the income process affects the debt constraints, and thus the extent to which individual agents can borrow, in exactly the same way it affected the extent of risk sharing in the simple model of Section 3. An increase in the dispersion of the income process increases not only the necessity of extended borrowing to smooth consumption, but also the possibility of extended borrowing since the default option may become less attractive. This effect is the driving force behind our main quantitative result that an increase in the cross-sectional dispersion of income may not lead to a significant increase in cross-sectional consumption inequality.

4.2.2 Standard Incomplete Markets (SIM)

We now compare our results to those obtained in a standard incomplete markets model, as in Huggett (1993) or Aiyagari (1994). Let q_t^{in} denote the price, at period t , of an uncontingent claim to one unit of the consumption good in period $t + 1$. The sequential budget constraints

the agent faces are

$$c_t(a_0, y^t) + q_t^{in} a_{t+1}(a_0, y^t) = w_t \alpha_{it} y_t + a_t(a_0, y^{t-1}) \quad (12)$$

and the short-sale constraints become

$$a_{t+1}(a_0, y^t) \geq -\alpha_{it} \bar{B}. \quad (13)$$

The definitions of equilibrium and stationary equilibrium for this economy are similar to the ones discussed above and are hence omitted. Notice that the only differences between the two economies are the set of financial assets that are traded (a full set of contingent claims in the DCM model and only a single uncontingent bond in the SIM model) and how the short-sale constraints that limit these asset trades are specified.

In order to compute calibrated versions of both economies, we reformulate them recursively. Note that the computation of equilibrium in the DCM model requires us to solve for both prices and borrowing constraints simultaneously (see Appendix B for details).

5 The Quantitative Exercise

We now explain the quantitative exercise we carry out below. It involves the following steps:

1. We first choose parameter values so that the stationary equilibrium in both economies matches key observations of the U.S. economy in the 1980s. This applies, in particular, to the deterministic and stochastic part of the labor productivity and thus income process.
2. We then introduce a finite path of changes in the dispersion of the labor productivity process to mimic the increase in income inequality observed in the U.S. data. We assume that this change in the labor productivity process is unforeseen by agents, but that all future changes in the process are fully learned once the first change occurs.
3. The change in the labor productivity process for a finite number of periods induces a transition in both models from the initial to a final stationary equilibrium associated with the process that prevails once the change in that process has been completed.

4. Both models endogenously generate consumption distributions along the transition from the old to the new steady state. We compute measures of consumption inequality for both models and compare them to the empirical facts established in Section 2. In order to carry out these steps, we now specify the parameters of both models.

5.1 Calibration

We need to choose (a) technology parameters A, δ, α , (b) preference parameters β and σ (the coefficient of relative risk aversion in the constant relative risk aversion (CRRA) utility function), (c) the endowment process $\{\alpha_{it}y_t\}_{t=0}^{\infty}$ with $y_t \in Y_t$ and $\alpha_{it} \in A_t$, (d) the transition matrices π_t , (e) the group sizes p_i , (f) the borrowing constraint \bar{B} for the SIM model, and (g) the autarkic interest rate r_d for the DCM model.

5.1.1 Income Process

We take the length of a model period to be one year. An individual's labor income $e_{it} = w_t\alpha_{it}y_t$ consists of a common wage; a group-specific, time-dependent deterministic part, α_{it} ; and an idiosyncratic stochastic component, y_t . In the empirical section we decomposed individual income and consumption data into a group-specific component and an idiosyncratic component. Our calibration strategy follows the same approach. The logarithm of labor income is given by

$$\ln(e_{it}) = \ln(w_t) + \ln(\alpha_{it}) + \ln(y_t),$$

and thus the cross-sectional variance of log-labor income is equal to

$$\sigma_{et}^2 = \sigma_{\alpha t}^2 + \sigma_{yt}^2,$$

where $\sigma_{et}^2 = Var[\ln(e_{it})]$, $\sigma_{\alpha t}^2 = Var[\ln(\alpha_{it})]$, and $\sigma_{yt}^2 = Var[\ln(y_t)]$. We identify $\{\sigma_{\alpha t}^2, \sigma_{yt}^2\}_{t=1980}^{2003}$ with the between- and within-group income variances plotted in Figure 2. We first HP-filter (with a smoothing parameter of 400) the time series $\{\sigma_{\alpha t}^2, \sigma_{yt}^2\}_{t=1980}^{2003}$, in order to remove high-frequency variation in our empirical variances. We then choose parameters governing the model income process so that (a) in the initial stationary equilibrium, both the between- and within-group income variances of the model match the filtered data for the early 1980s and (b) along

the transition, trends in between- and within-group income variances are reproduced by the model.

5.1.2 Between-Group Income Inequality

We pick the number of groups to be two with equal mass $p_i = 0.5$. For the initial stationary equilibrium we choose the group-specific means as $\alpha_1 = e^{-\sigma_{1980}}$ and $\alpha_2 = e^{\sigma_{1980}}$. Similarly, using $\sigma_{\alpha_{2003}}$ we obtain average group incomes for the final steady state, persisting from 2003 into the indefinite future. For the transition path we then select $\{\alpha_{1t}, \alpha_{2t}\}_{t=1981}^{2003}$ so that the trend of between-group income inequality follows that in the data.

Since $\{\alpha_{1t}, \alpha_{2t}\}_{t=1980}^{\infty}$ is a deterministic sequence, in the context of both models considered in this paper, the increase in between-group income inequality translates fully into an increase in between-group consumption inequality. Furthermore, by construction, the change in between-group inequality does not affect the quantitative importance of the risk-sharing mechanism at work for within-group stochastic income variability described in Section 3.

We choose this specification for two reasons. First, in an influential paper Attanasio and Davis (1996) show that between-group consumption insurance fails, and they conclude that “the evidence is highly favorable to an extreme alternative hypothesis under which relative consumption growth equals relative wage growth” (p. 1247). With our specification of average group income, changes in this income component are not (self-)insurable, consistent with their findings. Second, we will be able to quantify exactly to what extent the (self-)insurance mechanisms of both models can offset the increase in idiosyncratic income volatility. The potency of these mechanisms depends on the properties of the idiosyncratic income process, discussed next.

5.1.3 Within-Group Income Variability

We model the idiosyncratic part of the income process, $\ln(y_t)$, as the sum of a persistent and a transitory component, as in Storesletten et al. (1998, 2004) or Heathcote et al. (2004):

$$\begin{aligned} \ln(y_t) &= z_t + \varepsilon_t \\ z_t &= \rho z_{t-1} + \eta_t. \end{aligned} \tag{14}$$

Here ε_t, η_t are independent, serially uncorrelated, and normally distributed random variables with zero mean and variances $\sigma_{\varepsilon t}^2, \sigma_{\eta t}^2$, respectively. We explicitly allow these variances to change over time, whereas we treat ρ as a time-invariant parameter. Note, however, that if $\sigma_{\varepsilon t}^2$ and $\sigma_{\eta t}^2$ increase at different rates over time, the implied persistence of the idiosyncratic component of income, $\ln(y_t)$, changes. Thus, the process we use is flexible enough to allow for time-varying income persistence even if ρ is constant over time. As a benchmark value for ρ we choose $\rho = 0.9989$, the value Storesletten et al. (2004) find when estimating the process in (14). As a sensitivity analysis we also report results for lower persistence parameters.

We now describe how, from our data on $\{\sigma_{yt}^2\}_{t=1980}^{2003}$ and conditional on the value of ρ , we identify the unobserved variance of the transitory part, $\sigma_{\varepsilon t}^2$, and the variance of the persistent part, σ_{zt}^2 (or equivalently $\sigma_{\eta t}^2$). The key statistics that allow us to identify σ , and σ_{zt}^2 are the observed cross-sectional within-group income variance σ_{yt}^2 and the cross-sectional within-group income auto-covariance $Cov(y_t, y_{t+1})$, which, thanks to the short panel dimension of the CE data, we can measure in our sample. The two identifying equations for $(\sigma_{zt}^2, \sigma_{\varepsilon t}^2)$ are easily derived for our income process:

$$Cov(y_t, y_{t+1}) = E((z_t + \varepsilon_t)(\rho z_t + \eta_{t+1} + \varepsilon_{t+1})) = \rho \sigma_{zt}^2 \quad (15)$$

and

$$\sigma_{yt}^2 = \sigma_{zt}^2 + \sigma_{\varepsilon t}^2. \quad (16)$$

Given ρ , $Cov(y_t, y_{t+1})$, and σ_{yt}^2 , the values for σ_{zt}^2 and $\sigma_{\varepsilon t}^2$ are uniquely determined.

Equipped with time series for $\{\sigma_{\varepsilon t}^2, \sigma_{zt}^2\}_{t=1980}^{2003}$ from the data, we now specify a discretized version of the process in (14). The purely transitory component, at period t with equal probability, takes one of the two values $\{\varepsilon_{1t}, \varepsilon_{2t}\}$. If we choose $\varepsilon_{1t} = -\sigma_{\varepsilon t}$ and $\varepsilon_{2t} = \sigma_{\varepsilon t}$, the transitory shock fed into the model has a variance exactly as big as identified in the data, for all time periods. Finally, for the persistent part of the process we use a seven-state Markov chain with time-varying states and transition probabilities such that the variance of this process, in each period, equals σ_{zt}^2 as identified from the data.¹⁹ Note that after 24 model periods (2003 in real

¹⁹We employ the Tauchen and Hussey (1991) procedure to discretize the $AR(1)$ process. As inputs this procedure requires that $\rho = 0.9989$ (in the benchmark) and a (time-varying) variance $\sigma_{\eta t}^2$. The variance $\sigma_{\eta t}^2$ was chosen in such a way that the implied variance σ_{zt}^2 from the Markov chain matches the σ_{zt}^2 in the data.

time), the change in the dispersion of the income process is completed. However, due to the endogenous wealth dynamics in both models, it may take substantially longer than 24 years for both economies to complete the transition to the new stationary consumption distribution.

To summarize, the income process fed into the model perfectly reproduces the empirically identified time series of between-group income variance $\sigma_{\alpha t}^2$, as well as the within-group income variance due to the transitory shock, $\sigma_{\varepsilon t}^2$, and the persistent shock, $\sigma_{z t}^2$. Figure 4 displays the original and HP-filtered time series of these variances, identical for data and both models. We observe that all three components increase substantially over time, contributing to the overall increase in income inequality $\sigma_{y t}^2$. Of the overall increase of 18 percentage points in the variance of the filtered data, 36% is due to the change in the variance of between-group income, 40% is due to the persistent part, and 24% is due to the transitory part. Thus we confirm Violante’s (2002) findings that a significant part of the increase in wage or earnings inequality is due to bigger transitory shocks. The implied overall persistence of the idiosyncratic income process $\ln(y_t) = z_t + \varepsilon_t$ is roughly constant over time, with a very slight decline.

5.1.4 Exogenous Borrowing Limit and Autarkic Interest Rate

As a benchmark borrowing limit in the SIM model, we set $\bar{B} = 1$. Note that we normalize endowment in such a way that this borrowing limit corresponds to a generous one times the average annual income for each group i . As a benchmark in the DCM model, we allow households to save at the initial equilibrium interest rate, $r_d = r$ in autarky. We then report how sensitive our quantitative results are to the choice of the borrowing limit and the autarkic interest rate.

5.1.5 Technology and Preference Parameters

We assume that the period utility is logarithmic, $u(c) = \log(c)$. We then choose the technology parameters (A, α, δ) so that in both models the initial steady state has a wage rate of 1 (a normalization), a capital income share of 30%, and a return on physical capital of 4% per annum, as suggested in McGrattan and Prescott (2003). The time discount factor β is then set in both models such that the initial steady state in both models has a capital (wealth) -to-output ratio of 2.6. This value is equal to the average wealth (including financial wealth and housing wealth) for CE households in the benchmark sample in 1980–1981, and it is also close

to the value estimated in NIPA data by Fernández-Villaverde and Krueger (2002). Note that, conditional on the capital-output ratio target, the three technology parameters solve the three equations $\alpha = 0.3$ and

$$\begin{aligned} r &= 0.04 = \alpha \frac{Y}{K} - \delta \\ w &= 1 = (1 - \alpha)A \left(\frac{K}{N} \right)^\alpha = (1 - \alpha)A \left(\frac{r + \delta}{\alpha A} \right)^{\frac{\alpha}{\alpha-1}}, \end{aligned}$$

which yields $\delta = 7.54\%$ and $A = 0.9637$ in both models, since the production side of the economy is identical in both. The appropriate choice of the time discount factor β then ensures that households indeed have the incentive to save exactly the amount required to make the capital-output ratio equal to 2.6. This requires a $\beta = 0.959$ for the DCM model and a $\beta = 0.954$ for the SIM model. In our sensitivity analyses we always recalibrate β to maintain the same equilibrium capital-output ratio in the initial steady state.²⁰

6 Quantitative Results

6.1 Benchmark Calibration

We now document the quantitative predictions of both models with respect to the evolution of consumption inequality and compare them to the data. Since in our calibration we used filtered income variances as targets, we also filter the consumption inequality statistics to remove high-frequency fluctuations. Figure 5 summarizes the main quantitative results of our paper. The left panel displays the change in the *between-group* variance of log-consumption implied by both models as well as the data. The right panel does the same for the *within-group* variance of log-consumption, our main focus of interest. Since the sum of the between- and within-group variance of consumption equals the overall variance, the change in the overall variance from the data and both models can be readily deduced from the two panels.

Since the change in between-group income inequality is modeled by a deterministic process, by construction there is no (self-)insurance possible against the increase in between-group income variability. Thus, in the long run all of the increase in between-group income inequality is

²⁰Since the other statistics are determined exclusively from the production side of the economy, changes in \bar{B} or r_b do not require recalibration of A, α, δ .

reflected in a one-for-one increase in between-group consumption inequality in both models. It is therefore not surprising that both models have nearly identical predictions for that part of consumption inequality. As shown in panel (a) of Figure 2, the fact that in the data the increase in between-group income inequality is similar to the increase in between-group consumption inequality implies that the increase in between-group inequality predicted by both models is similar to the one observed in the data.

The crucial differences between the two models are the financial instruments that can be used to (self-)insure idiosyncratic income risk, the borrowing limits that constrain the use of these instruments, and the change of these limits over time. Thus, the crucial quantitative question is how well both models can capture the trend in within-group consumption inequality. The right panel of Figure 5 answers this question. It shows that, for our benchmark parameterization, the DCM model understates and the SIM model overstates the increase in within-group consumption inequality, compared to the data. Although the data display an increase in the variance of about 2.0%, the DCM model shows an increase of only 0.5%, whereas the standard incomplete markets model predicts an increase of 4.5% compared to the data.²¹ To put these numbers in perspective, note that the increase in within-group income variance from the data and fed into the models is 11%. Thus, both models are successful in generating an increase in within-group consumption inequality that is substantially lower than that for income inequality, in line with the data.

The empirical evidence of increasing both between-group and within-group consumption inequality also speaks against the standard *complete* markets model. That model, by allowing perfect consumption insurance between and within groups, counterfactually predicts that between-group, within-group, and total consumption inequality should remain completely unchanged over time.

The quantitative difference in the change of within-group consumption inequality in the two models is due to the differential response of consumer credit to increased income volatility. In the SIM model the increase in the variance of income leads to higher precautionary savings. In addition, households facing larger shocks become more hesitant to borrow, plus their ability

²¹Note that if both panels of Figure 5 are combined, the overall increase in consumption inequality in the data is, even in a quantitative sense, almost perfectly matched by the DCM model. But since this is due to the fact that this model overstates the increase in between-group consumption inequality and understates the increase in within-group consumption inequality, we do not want to stress this finding. The SIM model, however, overstates both components of inequality and thus the extent of the overall increase in consumption inequality.

to borrow remains unchanged. Thus, outstanding unsecured consumer credit (as a fraction of output) *declines* by 0.6%, equilibrium asset holdings and thus the physical capital stock increase by 2.1%, and the real return on capital declines by 17 basis points. In contrast, in the DCM model credit limits expand for the purchase of all Arrow securities, and households can and do borrow more, at least against the contingency of having higher income tomorrow. This is reflected in an increase in outstanding unsecured consumer credit (again, as a fraction of output) by 2.1%. But keeping consumption as smooth as it was before the change in the income process may require a bigger expansion in borrowing than is feasible with the new, wider constraints, so some of the increase in income volatility is reflected in consumption: the within-group consumption variance increases, albeit only very mildly. The expansion of credit in the DCM model is met by an increase in purchases of Arrow securities, as households have a stronger need to save for the contingency of being income-poor tomorrow. On net, aggregate savings and thus the capital stock increases by 1.8% and the return on capital falls by 15 basis points. The increase in the capital stock and the decline in the interest rate are smaller than those in the SIM model because the increase in asset accumulation in the debt constraint model is partially offset by a higher demand for credit, an effect that is absent in the SIM model. A precise quantitative evaluation of both models with respect to the CE data along the credit dimension is not possible because data on unsecured consumer credit are not available in our CE sample. At least qualitatively, however, the DCM model seems more consistent with recent developments in U.S. credit markets (see also our Figure 7 in the conclusion).

6.2 Sensitivity Analysis

In this section we document how sensitive our main findings are to changes in the tightness of the borrowing constraints and the high persistence of the income shock. Finally we assess the performance of a hybrid model that inherits elements of both the SIM and DCM models.

6.2.1 Borrowing Constraints

In Table 2 we report results for different values of the borrowing constraint for the SIM model. We normalized the production function in such a way that, in the initial steady state, average wages are equal to 1. Thus, a borrowing constraint of $B = 2$ implies that a household can take

out (noncollateralized) loans up to twice her annual average labor income. We also document how our results for the DCM model change if we reduce the net real interest rate at which households can save in autarky to zero, making the default option less attractive. The statistics we report are the change, between 1980 and 2003, in the *within-group* consumption (in logs) variance, the change in the outstanding credit-to-GDP ratio, and the change in the real interest rate.

TABLE 2. CHANGE IN BORROWING CONSTRAINTS

Economy	ΔVar	$\Delta Credit$ (%)	Δr (%)
SIM, $B = 0$	0.0456	0	-0.20
SIM, $B = 1$	0.0454	-0.62	-0.17
SIM, $B = 2$	0.0463	-0.88	-0.15
DCM, $r_d = 4\%$	0.0049	2.05	-0.15
DCM, $r_d = 0\%$	0.0012	12.34	-0.03

The most important observation from Table 2 is that more generous credit lines per se, if they *do not change over time*, do not help to keep the increase in within-group consumption inequality low in the SIM model. For this result, the existence of capital as an asset in positive net supply is crucial. Since the wealth-to-output ratio is calibrated to 2.6 in the model, most households are far to the right from the borrowing constraint in the wealth distribution. In addition, with highly persistent income shocks, they are very reluctant to obtain credit in response to bad income shocks. In fact, as these shocks become bigger over time, households become even more timid in using credit and try to stay away from high debt positions. We observe from the table that the credit-to-GDP ratio declines with the increase in income volatility, the more so the looser the borrowing constraint and thus the higher the debt position of those at the constraint.

For the DCM model, reducing the interest rate at which people can save after default leads to a worse autarkic option, and thus to more borrowing being enforceable. In fact, with an $r_d = -2\%$ (i.e., households can save only at a negative real interest rate of 2%), the default option is so unattractive that perfect risk sharing is possible and the predictions of the DCM model collapse to those of the standard complete markets model. Reducing the autarkic savings interest rate gives rise not only to better risk allocation in the initial steady state, but also to a stronger relaxation of borrowing constraints over time. As a result, the use of credit

expands substantially and there is almost no increase in within-group consumption inequality over time, reducing the ability of the DCM model to match the data relative to the benchmark parameterization. Also note that with $r_d = 0\%$, the increase in credit demand almost matches the increased savings demand, so that the real interest rate and the capital stock remain virtually unchanged as income variability increases.

6.2.2 Persistence of Income Shocks

At least since Friedman (1957) it has been well understood that very persistent income shocks are harder to self-insure against than income shocks that are transitory in nature. Our idiosyncratic income process is the sum of a highly persistent and a purely transitory component. Although authors that estimate this process from wage data consistently find the persistence parameter ρ close to 1, some disagreement exists about its exact magnitude. We have repeated our analysis for various other choices of ρ and report, in Table 3, results for $\rho = 0.8$, the value estimated by Guvenen (2005), which is the lowest estimate for ρ we are aware of, for the exact income process we use.²² In each case we recalibrate the income process and time discount factor such that the cross-sectional income dispersion of the process fed into the model matches the empirical facts from Figure 4, and the initial real return on capital remains at 4%. This procedure keeps the volatility of the income process unchanged, but reduces its persistence.

TABLE 3. CHANGE IN PERSISTENCE: $\rho = 0.8$

Economy	ΔVar	$\Delta Credit$ (%)	Δr (%)
SIM, $B = 0$	0.0355	0	-0.28
SIM, $B = 1$	0.0334	-0.01	-0.22
SIM, $B = 2$	0.0341	0.0	-0.19
DCM, $r_d = 4\%$	0.0025	1.33	-0.10
DCM, $r_d = 2.5\%$	0.0013	3.93	-0.04

Comparing the results of Table 3 with those in Table 2, we see that a lower persistence of income shocks in the SIM model indeed reduces the rise in within-group consumption inequality.

²²The $\rho = 0.8$ is the lowest estimate reported in his Table 1. Heaton and Lucas (1996) estimate a simple $AR(1)$ process; that is, they do not have an independent purely transitory shock. Thus, if the true process is the one we use, their estimated $\rho = 0.53$ is a downward-biased estimate of the true autoregressive coefficient.

We also repeated our exercises with $\rho = 0.95$, the value reported by Storesletten et al. (1998). The results, available upon request, are quite similar to those for $\rho = 0.9989$.

Whereas for $\rho = 0.9989$ this increase rise was about 4.5%, with $\rho = 0.8$ it drops to about 3.4% (compared to 2% in the data). Again, the results are fairly independent of the borrowing constraint. With lower persistence, households in the SIM model find it easier to self-insure by accumulating capital and using it to smooth income shocks. The increase in the capital stock (and corresponding decline in the real return) is more pronounced for $\rho = 0.8$ than for $\rho = 0.9989$. Also, households now are not as timid as before to use credit to smooth income shocks; instead of a decline of credit as a fraction of GDP, we now observe this statistic to be virtually unchanged. We conclude that setting persistence of the income shocks to the lower bound from the empirical literature leads to a lesser increase of within-group consumption inequality in the SIM model. Still, that model somewhat overstates the increase observed in the data.

In the DCM model, a lower ρ reduces the value of autarky for households with currently high income whose constraints are binding because it is now less likely that they will remain income-rich. Thus, *ceteris paribus*, the implications of the DCM model are closer to those of the complete markets model. In fact, for $\rho = 0.8$ any autarkic interest rate below 2% leads to complete consumption insurance. Thus, although the results for this persistence are not identical to the complete markets model, they are quantitatively close, making the model's understatement of the increase in within-group consumption inequality more severe.

6.2.3 Market Completeness or Endogenous Borrowing Constraints?

There are two main differences between our DCM model and the SIM model.²³ First, our model features a full set of Arrow securities, and second, borrowing constraints adjust endogenously to changes in the income process. To evaluate whether introducing time-varying borrowing constraints in the SIM model helps to bring that model's predictions closer to the data, we now present results for a hybrid model. This model, named after Zhang (1997), retains the market structure of the SIM model in that agents can trade only a risk-free, uncontingent bond. The borrowing constraint, however, is now allowed to vary over time: agents can borrow up to the maximum amount such that in *all possible* states tomorrow, they are at least weakly better off repaying their debt rather than defaulting and living in autarky from there on.

Table 4 displays a crucial lesson: introducing time-dependent borrowing constraints into the

²³We thank Pierre-Olivier Gourinchas for helpful discussions leading to this subsection.

SIM model, at least in the way proposed by Zhang, does not alter the quantitative implications of the model with respect to the change in within-group consumption inequality. It does, however, enable the model to reproduce the empirically observed increase in the use of credit.

TABLE 4. A HYBRID MODEL

Economy	ΔVar	ΔVar_0	$\Delta Credit$ (%)	Δr (%)
SIM, $\rho = 0.9989, B = 1$	0.0454	0.0095	-0.62	-0.17
Zhang, $\rho = 0.9989$	0.0469	-0.0112	1.7	-0.16
DCM, $\rho = 0.9989, r_d = 4\%$	0.0049	0.0015	2.05	-0.15
SIM, $\rho = 0.8, B = 1$	0.0334	0.0049	-0.01	-0.22
Zhang, $\rho = 0.8$	0.0348	-0.0154	1.6	-0.22
DCM, $\rho = 0.8, r_d = 4\%$	0.0025	0.0007	1.33	-0.10

To explain the driving forces behind these results, Table 4 reports not only the change in within-consumption inequality between the first and last periods of the income transition (the column labeled ΔVar), but also the change between the initial steady state and the first period of the transition (the column labeled ΔVar_0). Households, on impact, respond very differently to the unexpected change in the income process in the Zhang economy, compared to the SIM economy. In the Zhang economy, credit lines expand on impact and households make use of these expanded credit lines. As a result, consumption inequality initially falls. But this comes at a price: households that have borrowed more now face higher debt levels and income processes with more extreme realizations. Over time, the debt needs to be serviced, consumption has to respond, and consumption inequality eventually goes up (about five periods after impact). After 25 years the change in consumption inequality is the same as in the benchmark SIM model.²⁴ Thus, borrowing constraints that are relaxed over time help households in the short run to better smooth more volatile income fluctuations, but in the long run their debt burden catches up with them. As a result, consumption inequality does not increase in the short run, but eventually it does increase to the same extent as in the standard model.

Why doesn't the same logic apply in the DCM model? Here is where the second feature of

²⁴Another way to explain this is to note that with highly persistent shocks, an increase in the variance of these shocks leads to an increase in the cross-sectional variance of *permanent* income of comparable magnitude, something that Bowlus and Robin (2004) argue occurred empirically in the United States. Not surprisingly, a higher dispersion of permanent income eventually leads to a larger dispersion of consumption in the SIM model.

that model, state-contingent borrowing, comes in. In the DCM model households with currently high income enter the period with high debt, exactly because they have borrowed against the contingency of being income-rich. Thus, the high debt is not such a high burden, and households can and do make use of the higher credit lines without the consequence of particularly low consumption in the future. In sharp contrast, in the SIM model income-poor people start the period with outstanding credit (because of the strong positive correlation of income and the fact that they went into debt because of bad income realizations). Thus, in that model highly indebted households eventually have to accept (persistently) low consumption. Put another way, in the SIM model assets and income are highly positively correlated, whereas in the DCM model they are negatively correlated. Therefore, in one model being in debt has fairly persistent negative consequences for consumption; in the other model it does not.

7 Conclusions

In this paper we use CE survey data to document that the increase in income inequality for the United States in the last 25 years has not been accompanied by a substantial increase in consumption inequality. We propose a theory that provides a simple explanation for this observation. If the increase in income inequality has been driven, at least partially, by an increase in idiosyncratic labor income risk, then the value households place on access to formal and informal credit and insurance mechanisms rises, and the scope of these mechanisms may endogenously broaden. Individual consumption may then be better insulated against (higher) income risk, and cross-sectional consumption inequality may increase only mildly. If, however, the structure of private financial markets and informal insurance arrangements does not respond to changes in the underlying stochastic income process of individuals, then no further hedging against the increasing risk is possible, and the increase in income inequality leads to a more pronounced rise in consumption inequality.

The mechanism through which agents in the DCM model of the last section keep their consumption profiles stable in the light of more volatile income is an expansion in the use of noncollateralized credit. Did this expansion take place in the data? One simple (but of course only partial) measure of credit available that is used by U.S. consumers is the ratio of aggregate

unsecured consumer credit to disposable income. In Figure 6 we plot this ratio from U.S. *data* for the last 40 years, as well as the Gini coefficient for U.S. household income.²⁵ Despite some idiosyncratic cyclical variations, the two series display a remarkably similar trend. Combining this figure with our consumption inequality observations may suggest that consumers could and in fact did make stronger use of credit markets exactly when they needed to (starting in the mid-1970s), in order to insulate consumption from bigger income fluctuations. We want to stress that we view the expansion of credit lines in the DCM model as a metaphor for the expansion of a variety of formal and informal risk-sharing mechanisms, with formal credit being an important, but by no means the only, component.

Both models evaluated in this paper have their shortcomings. The DCM model hypothesizes very well-developed—maybe too developed—direct insurance markets. Although the complete set of Arrow securities can be interpreted to stand in for the host of assets traded in financial markets and for informal insurance mechanisms working on the level of the extended family, some may question the empirical realism of this assumption, which is important for our results, as demonstrated in the last section. This model is consistent with the empirically observed increase in credit and the only slight increase in consumption inequality over time. On the one hand, it overstates the degree to which individual consumption is insulated from income shocks and understates the response of within-group consumption inequality to increased income volatility. On the other hand, the SIM model may have a more realistic market structure, but seems to allow somewhat too little explicit insurance against income shocks, over and above simple self-insurance. In our application it therefore overstates the level and the increase in within-group consumption inequality. It also does not display an increase in the use of uncollateralized credit over time, unless more flexible borrowing constraints are introduced. In their work, Blundell et al. (2004) and Storesletten et al. (2004) come to the same qualitative conclusion.

Conditional on these findings, we conjecture that a model that, like the DCM, has an endogenous evolution of credit markets but that, unlike the DCM, restricts the insurance possibilities available to agents (possibly restricting the spanning of the insurance contracts) might be even

²⁵The series for consumer credit is from the 2005 *Economic Report of the President*, Table B-77. It only includes revolving consumer credit, which is entirely unsecured. The series is available starting from 1968. Personal disposable income is obtained from the same source, Table B-30. The Gini index for household income is available, starting from 1967, from the U.S. Census Bureau, Historical Income Tables, Table H-4.

more empirically successful in matching the data. Further empirical work using micro data may inform us how to more precisely model the mechanisms that households can use to smooth out idiosyncratic income shocks. We defer this work to ongoing and future research.

References

- Aiyagari, S. R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109, 659–684.
- Albarran, P. and O. Attanasio (2003), “Limited Commitment and Crowding Out of Private Transfers: Evidence from a Randomised Experiment,” *Economic Journal*, 113, C77–C85.
- Alvarez, F. and U. Jermann (2000), “Efficiency, Equilibrium, and Asset Pricing with Risk of Default,” *Econometrica*, 68, 775–798.
- Attanasio O., E. Battistin, and H. Ichimura (2005), “What Really Happened to Consumption Inequality in the US?” in E. Berndt and C. Hulten (eds.), *Measurement Issues in Economics—The Paths Ahead: Essays in Honor of Zvi Griliches*, University of Chicago Press, Chicago, forthcoming.
- Attanasio, O. and S. Davis (1996), “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104, 1227–1262.
- Battistin, E. (2003), “Errors in Survey Reports of Consumption Expenditures,” IFS Working Papers, W03/07.
- Blundell, R., L. Pistaferri, and I. Preston (2004), “Consumption Inequality and Partial Insurance,” ISF Working Papers, W04/28.
- Blundell, R. and I. Preston (1998), “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, 113, 603–640.
- Bowlus, A. and J.-M. Robin (2004), “Twenty Years of Rising Inequality in U.S. Lifetime Labour Income Values,” *Review of Economic Studies*, 71, 709–742.
- Cutler, D. and J. Gruber (1996), “Does Public Insurance Crowd Out Private Insurance?” *Quarterly Journal of Economics*, 111, 391–430.
- Cutler, D. and L. Katz (1991a), “Macroeconomic Performance and the Disadvantaged,” *Brookings Papers on Economic Activity*, 1–74.
- _____ (1991b), “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s,” *American Economic Review*, 82, 546–551.
- Dalaker, J. and M. Naifeh (1998), *Poverty in the United States: 1997*, U.S. Bureau of the Census, Current Population Reports Series P60-201, U.S. Government Printing Office, Washington, D.C.

Deaton, A. and C. Paxson (1994), “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 102, 437–467.

Díaz-Giménez, J., V. Quadrini, and J.-V. Ríos-Rull (1997), “Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 21, 3–21.

Dynarski, S. and J. Gruber (1997), “Can Families Smooth Variable Earnings?” *Brookings Papers on Economic Activity*, 229–303.

Economic Report of the President, 2005 (2005), Executive Office of the President, Council of Economic Advisers, Washington, D.C., Table B-77.

Fernández-Villaverde, J. and D. Krueger (2002), “Consumption and Saving over the Life Cycle: How Important Are Consumer Durables?” mimeo, University of Pennsylvania.

Friedman, M. (1957), *A Theory of the Consumption Function*, Princeton University Press, Princeton, N.J.

Gottschalk, P. and R. Moffitt (1994), “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 217–272.

Gottschalk, P. and T. Smeeding (1997), “Cross-National Comparisons of Earnings and Income Inequality,” *Journal of Economic Literature*, 35, 633–687.

Greenspan, A. (1998), “Opening Remarks to the Symposium on Income Inequality: Issues and Policy Options,” Federal Reserve Bank of Kansas City.

Guvenen, F. (2005), “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?” mimeo, University of Rochester.

Heathcote J., K. Storesletten, and G. Violante (2004), “The Macroeconomic Implications of Rising Wage Inequality in the US,” mimeo, Georgetown University.

Heaton, J. and D. Lucas (1996), “Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing,” *Journal of Political Economy*, 104, 443–487.

Huggett, M. (1993), “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 17, 953–969.

Johnson, D. and S. Shipp (1991), “Trends in Inequality Using Consumer Expenditures: 1960 to 1993,” *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 1–7.

Johnson, D. and T. Smeeding (1998), “Measuring the Trends in Inequality of Individuals and Families: Income and Consumption,” mimeo, Bureau of Labor Statistics.

Katz, L. and D. Autor (1999), “Changes in the Wage Structure and Earnings Inequality,” in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 3A, North-Holland, Amsterdam.

Kehoe, T. and D. Levine (1993), “Debt Constrained Asset Markets,” *Review of Economic Studies*, 60, 865–888.

_____ (2001), “Liquidity Constrained Markets versus Debt Constrained Markets,” *Econometrica*, 69, 575–598.

Kocherlakota, N. (1996), “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 63(4), 595–609.

Krueger, D. and F. Perri (1999), “Risk Sharing: Private Insurance Markets or Redistributive Taxes?” Federal Reserve Bank of Minneapolis Staff Report 262.

_____ (2004), “On the Welfare Consequences of the Increase in Inequality in the United States,” *NBER Macroeconomics Annual 2003*, 83–121.

Mayer, S. and C. Jencks (1993), “Recent Trends in Economic Inequality in the United States: Income versus Expenditures versus Well-Being,” in D. Papadimitriou and E. Wolff (eds.), *Poverty and Prosperity in the USA in the Late Twentieth Century*, St. Martin’s Press, New York.

McGrattan, E. and E. Prescott (2003), “Average Debt and Equity Returns: Puzzling?” *American Economic Review*, 93, 392–397.

Nelson, J. (1994), “On Testing for Full Insurance Using Consumer Expenditure Survey Data: Comment,” *Journal of Political Economy*, 102, 384–394.

Pendakur, K. (1998), “Changes in Canadian Family Income and Family Consumption Inequality between 1978 and 1992,” *Review of Income and Wealth*, 44, 259–283.

Slesnick, D. (1993), “Gaining Ground: Poverty in the Postwar United States,” *Journal of Political Economy*, 101, 1–38.

_____ (2001), *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*, Cambridge University Press, Cambridge.

Storesletten, K., C. Telmer, and A. Yaron (1998), “The Risk Sharing Implications of Alter-

native Social Security Arrangements,” *Carnegie-Rochester Conference Series on Public Policy*, 50, 213–59.

_____ 2004, “Consumption and Risk Sharing over the Life Cycle,” *Journal of Monetary Economics*, 51, 609–633.

Tauchen, G. and R. Hussey (1991), “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica*, 59, 371–396.

U.S. Census Bureau (2004), “Gini Ratios for Households, by Race and Hispanic Origin of Householder: 1967 to 2000,” <http://www.census.gov/hhes/income/histinc/h04.html>.

Violante, G. (2002), “Technological Acceleration, Skill Transferability and the Rise in Residual Inequality,” *Quarterly Journal of Economics*, 117, 297–338.

Zhang, H. (1997), “Endogenous Borrowing Constraints with Incomplete Markets,” *Journal of Finance*, 52, 2187–2209.

Appendix A Data Description

In this appendix we provide a detailed description of the income and consumption data we use in the paper.

A.1 CE Files and Observations

Our data come from the CE Interview Surveys 1980 through 2003 (which also contain data for the first quarter of 2004) provided by the Bureau of Labor Statistics (BLS). Consumption expenditure data are from the Family Characteristics and Income (FAMILY) files except for the years 1982 and 1983, for which the FAMILY files do not contain consumption information. For those years we obtain consumption data from the Detailed Expenditures (MTAB) files. Note that the consumption information in the FAMILY files is just an aggregation of the information in the MTAB files. Income data are from the FAMILY files and hours worked by household members (used to construct wages) are from the Member Characteristics and Income (MEMBER) files. An observation in our data set consists of all the information collected in a given interview for a given household (identified by a unique ID number). Each household is present in no more than four observations.

A.2 Consumption Categories

In Table A1 we report all the categories of consumption expenditures we use, together with the BLS price index we use to deflate them. Categories 1 through 20 are reported directly in the CE for every observation, and we derived categories 21 and 22 as described below. Our definition of ND consumption (used in Table 1) includes categories 1 through 13. Our measure of ND+ consumption includes ND consumption, categories 20 and 21, plus categories 14 through 18. These latter categories contain expenditures on durable goods, but since we do not have enough information to compute imputed services from those durables, we simply include expenditures. Finally, our definition of total consumption expenditures (used in Table 1) includes categories 1 through 20. We now describe in more detail some of the categories.

A.2.1 Food Expenditures Correction

A change in survey methodology (see Battistin, 2003, for details) causes a sizeable (about 15%) systematic downward bias in reported food expenditures for all the observations in the years 1982–1987. In order to correct for this bias, we regress the log of food expenditures for all years on a quadratic time trend, on quadratics in income and total nonfood consumption expenditures, on weeks worked, on a complete set of household characteristics (including age, education, region of residence, and family composition), on a dummy for the period 1982–1987, and on the interactions term of the dummy with all other independent variables. We then use the regression coefficients to scale up food expenditures for every observation in the period 1982–1987.

A.2.2 Services from Vehicles

Each CE observation contains reports of expenditures for purchases of new and used vehicles. The CE also reports the number of cars owned by the household in that quarter. For each year we first select all observations that report positive expenditures for vehicle purchases, and run

a regression of these expenditures on quadratics in income and total nonvehicle consumption expenditures, weeks worked by household members, expenditures on gasoline, expenditures on public transportation, vehicle maintenance expenditures, the number of cars owned, a complete set of household characteristics (including age, education, region of residence, and family composition), plus quarter dummies. These regressions have an R^2 that ranges from 74% to 94% in our sample years. On average, in every year a little more than 10% of households report positive expenditures on vehicles. We use the estimated regression coefficients to predict expenditures for vehicles for all households in that year (i.e., for those who did and for those who did not report positive vehicle expenditures). Our measure of consumption services from vehicles, then, is the predicted expenditures on vehicles, times the number of cars the consumer unit owns, times $\frac{1}{32}$ (reflecting the assumption of average complete depreciation of a vehicle after 32 quarters).

A.2.3 Services from Primary Residence

Each observation in the CE provides information on whether the household rents or owns its primary residence. If the household rents, we measure housing services as the rent paid, including insurance and other out-of-pocket expenses paid by the renter. To impute housing services for those households that own, we use a variable from the CE that measures the market rent (as estimated by the reference person of the consumer unit) the residence would command if rented out. This variable is not available for all years of the sample, in particular not for the years 1980–1981 and 1993–1994. Thus, in order to compute this variable in a uniform way across our sample, we use an imputation procedure similar to the one used for vehicles. For the year for which we have the reported market rent of the unit, we regressed it on self-reported property values, quadratics in income and total nonhousing consumption expenditures, a complete set of household characteristics (including age, education, region of residence, and family composition), plus quarter dummies. Since property values are reported by only a subset of the home owners, we allow the coefficient of the regression to be different for those who reported the property value from those who did not. These regressions have an R^2 that ranges from 30% to 55% in our sample years. We then use estimated regression coefficients to predict the rental value of owned properties for all the home owners. For the years 1980–1981 we use the coefficient estimated in 1982, and for the years 1993–1994 we use the coefficients estimated in 1995. For the years 1982–1992 and 1995–2003, we have both the actual and the imputed rental equivalent of the owned home, so we computed the trends in consumption inequality using housing services computed in both ways. The resulting trends are extremely similar.

A.3 Sample Selection

We first exclude observations for which there is clear evidence of measurement error. In particular we exclude observations classified as incomplete income respondents, observations that report zero food expenditures for the quarter, those who report only food expenditures for the quarter, and those who report positive labor income but no hours worked. We then exclude all observations with an age of the household head below 21 or above 64, and those with negative or zero LEA+ earnings. In our benchmark sample we also exclude observations with weekly

wages of the reference person below half of the minimum wage, households classified as rural, and those households that have not completed the full set of four interviews.

A.4 Aggregation and Top-Coding

Inequality measures are computed on annual cross sections. We assign an observation to a given year if the last interview of that household is completed between April of that year and March of the following year. Whenever income or consumption expenditures are top-coded, we set them to their top-coding thresholds. We have experimented with increasing the values of top-coded income/consumption components (multiplying the threshold by 1.5). Inequality measures are robust to these changes because in general, the number of observations with top-coded income or consumption figures never exceeds 2% in a given quarter.

A.5 Data Availability

Our data set in Stata format, including brief documentation, is available at http://pages.stern.nyu.edu/~fperri/research_data.htm.

Appendix B Recursive Formulation and Computational Algorithm

In this appendix we formulate the consumer problem for the DCM model recursively and provide a sketch of the algorithm used to compute a stationary equilibrium. In the nonstationary case (that is, along the transition), the logic remains the same but all functions have to be indexed by t . For simplicity here we omit the distinction by types and lump into y the transitory and persistent income shocks. The equilibrium problem is nonstandard because one needs to solve not only for prices but also for endogenous borrowing constraints. We first compute the value of autarky as the fixed point to the functional equation

$$U^{Aut}(y, b) = \max_{c, b' \geq 0} \left((1 - \beta)u(b + y - \frac{b'}{1 + r^d}) + \beta \sum_{y' \in Y} \pi(y'|y)U^{Aut}(y', b') \right).$$

We then guess the risk-free rate $R = 1/q$. No arbitrage implies that the prices of the Arrow securities $q(y'|y)$ are a function of our guess and are given by $q\pi(y'|y)$. We guess borrowing constraints $A^i(y')$ and solve the consumer problem, taking these borrowing constraints $A^i(y')$ and prices for Arrow securities $q\pi(y'|y)$ as given:

$$\begin{aligned} V(y, a) &= \max_{c, \{a'(y')\}_{y' \in Y}} \left\{ (1 - \beta)u(c) + \beta \sum_{y' \in Y} \pi(y'|y)V(i, a'(y'), y') \right\} \\ &\text{s.t.} \\ c + \sum_{y' \in Y} q(y'|y)a'(y') &= \alpha_i y + a \quad a' \geq A^i(y'). \end{aligned}$$

We finally check to make sure the borrowing constraints are not too tight by asking whether

$$V(y', A^i(y')) = U^{Aut}(y', 0)$$

for all y' . If the equalities hold, then we have solved for the borrowing constraints associated with the guessed interest rate; if not, we update the guesses for $A^i(y')$ until all equalities hold. Once we have found the borrowing constraints that are not too tight, we use the associated optimal asset policies $a'(y, a; y')$ together with the transition probabilities π to define the operator H that maps current measures over wealth and income shocks into tomorrow's measures. We then compute the (unique) fixed point of the operator H and denote it by Φ . Given Φ and the optimal consumption policies, we can check the market clearing conditions. If market clearing holds, we have found a stationary equilibrium; if not, we update our guess of the interest rate $R = \frac{1}{q}$. We implement this procedure numerically by approximating value and policy functions with piecewise linear functions over the state space. For more details on the basic algorithm and on the theoretical characterization of the stationary equilibrium, see Krueger and Perri (1999).

Table A1. Consumption Categories and Deflators

Number	Category	CPI Used to Deflate ^a	BLS CPI Code
Nondurable Expenditures			
1	Food	Food	SAF1
2	Alcoholic beverages	Alcoholic beverages	SAF116
3	Tobacco	Tobacco and smoking products	SEGA
4	Personal care	Personal care	SAG1
5	Fuels, utilities, and public services	Fuels and utilities	SAH2
6	Household operations	Household furnishings and operations	SAH3
7	Public transportation	Public transportation	SETG
8	Gasoline and motor oil	Motor fuels	SETB
9	Apparel	Apparel	SAA
10	Education	Tuition expenditures	SEEB
11	Reading	Recreational reading material	SERG
12	Health services	Medical care	SAM
13	Miscellaneous expenditures ^b	Miscellaneous personal services	SEGD
Other Expenditures			
14	Entertainment	Entertainment ^c	SA6
15	Household equipment	Household furnishings and operations	SAH3
16	Other lodging expenses ^d	Shelter	SAH1
17	Other vehicle expenses ^e	Car maintenance and repair	SETD
18	Rented dwellings	Rent of primary residence	SEHA
19	Owned dwellings	Shelter	SAH1
20	Purchases of vehicles	Purchase of new vehicles	SETA01
Imputed Services			
21	Services from owned primary residence	Rent of primary residence	SEHA
22	Services from vehicles	Purchase of new vehicles	SETA01

^a CPI are monthly average city data for all urban consumers, not seasonally adjusted, and are in base 1982–1984 = 100.

^b Mostly fees for services such as banking or legal assistance.

^c BLS CPI for entertainment ends in 1998. We extend it to 2004 using the CPI for recreation (SAR).

^d Includes mostly expenditures on vacation homes.

^e Includes expenditures on maintenance, repairs, insurance, and finance charges.

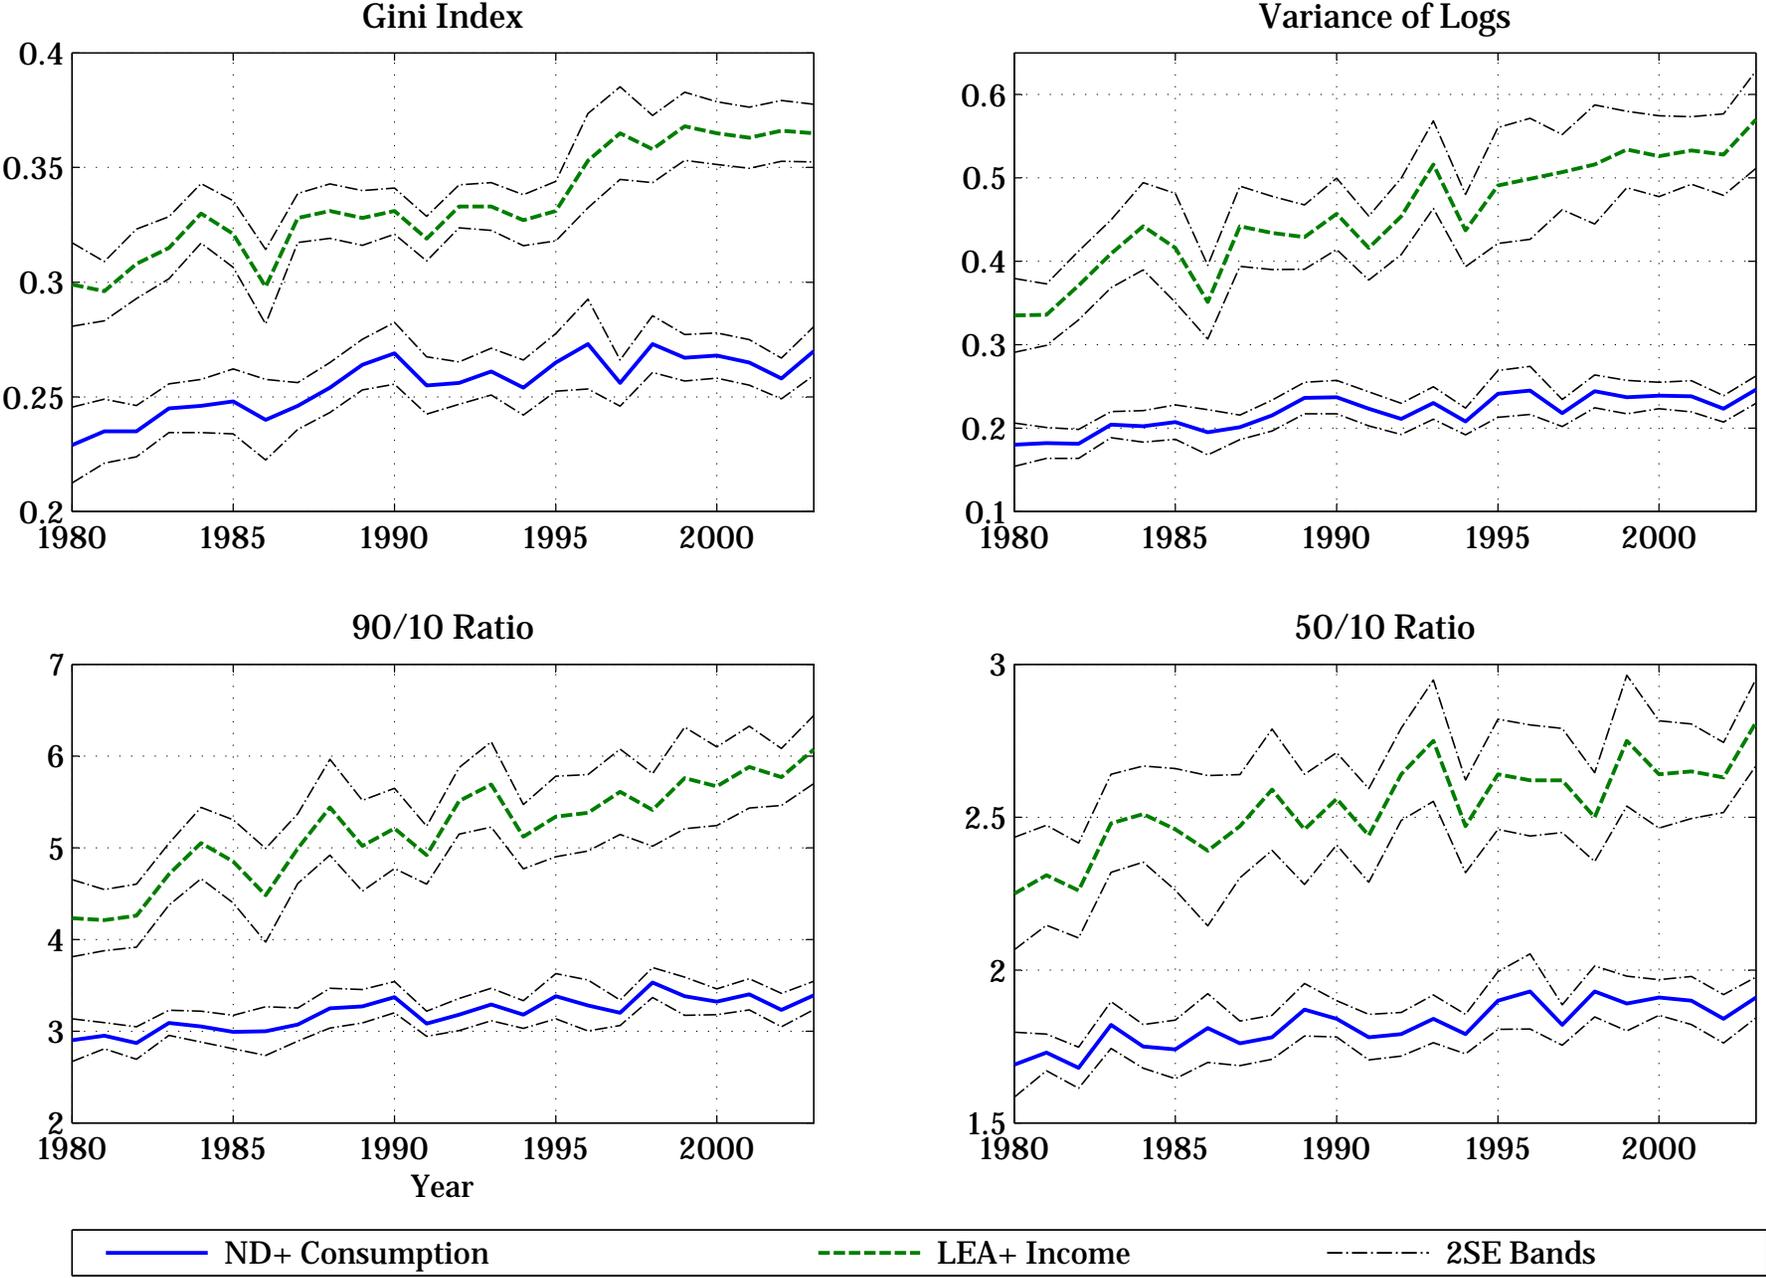
Table A2. Summary Statistics for the Benchmark Sample

Year	Households	Average Income	Average Consumption	
		LEA+	ND	ND+
1980	638	15736	7447	12940
1981	1439	15419	7348	12875
1982	1313	16453	6949	12605
1983	1700	16746	7134	12873
1984	1771	18017	7398	13495
1985	1267	17862	7323	13463
1986	748	18678	7308	13725
1987	1840	19318	7558	13992
1988	1621	18881	7391	13752
1989	1762	19585	7517	14183
1990	1789	18893	7455	14171
1991	1749	19170	7376	13858
1992	1725	19046	7126	13737
1993	1786	18841	7115	13686
1994	1744	18657	7114	13498
1995	1325	19269	7127	13865
1996	881	20203	7433	14674
1997	1609	20557	7203	14204
1998	1578	20952	7323	14689
1999	1763	22060	7263	14532
2000	2164	21508	7191	14508
2001	2227	21468	6967	14060
2002	2387	23427	7280	14766
2003	2593	22544	7053	14359

Note: Income and consumption measures are in 1982–1984 constant dollars per adult equivalent.

Averages are weighted using CE population weights. A household belongs to year x if its fifth interview is between the second quarter of year x and the first quarter of year $x+1$.

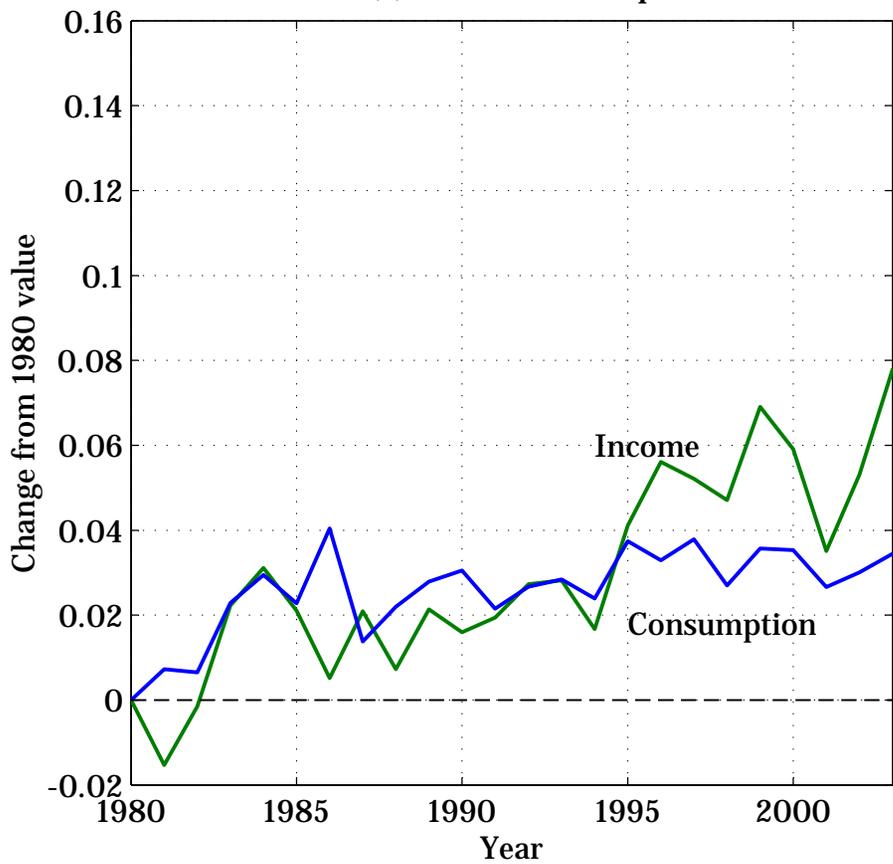
Figure 1. The Evolution of Income and Consumption Inequality in the United States, 1980-2003



Note: The standard errors are computed using a bootstrap procedure with 100 repetitions

Figure 2. Changes in Between- and Within-Group Income and Consumption Inequality

(a) Between-Group



(b) Within-Group

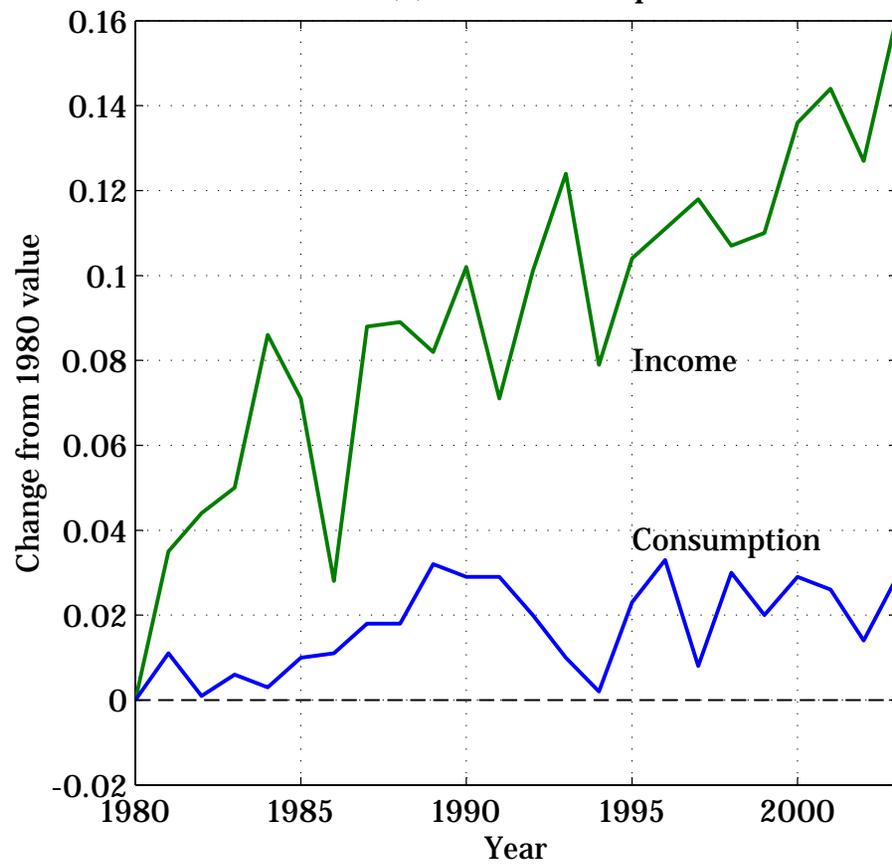


Figure 3. Characterizing the Link Between Income and Consumption Dispersion

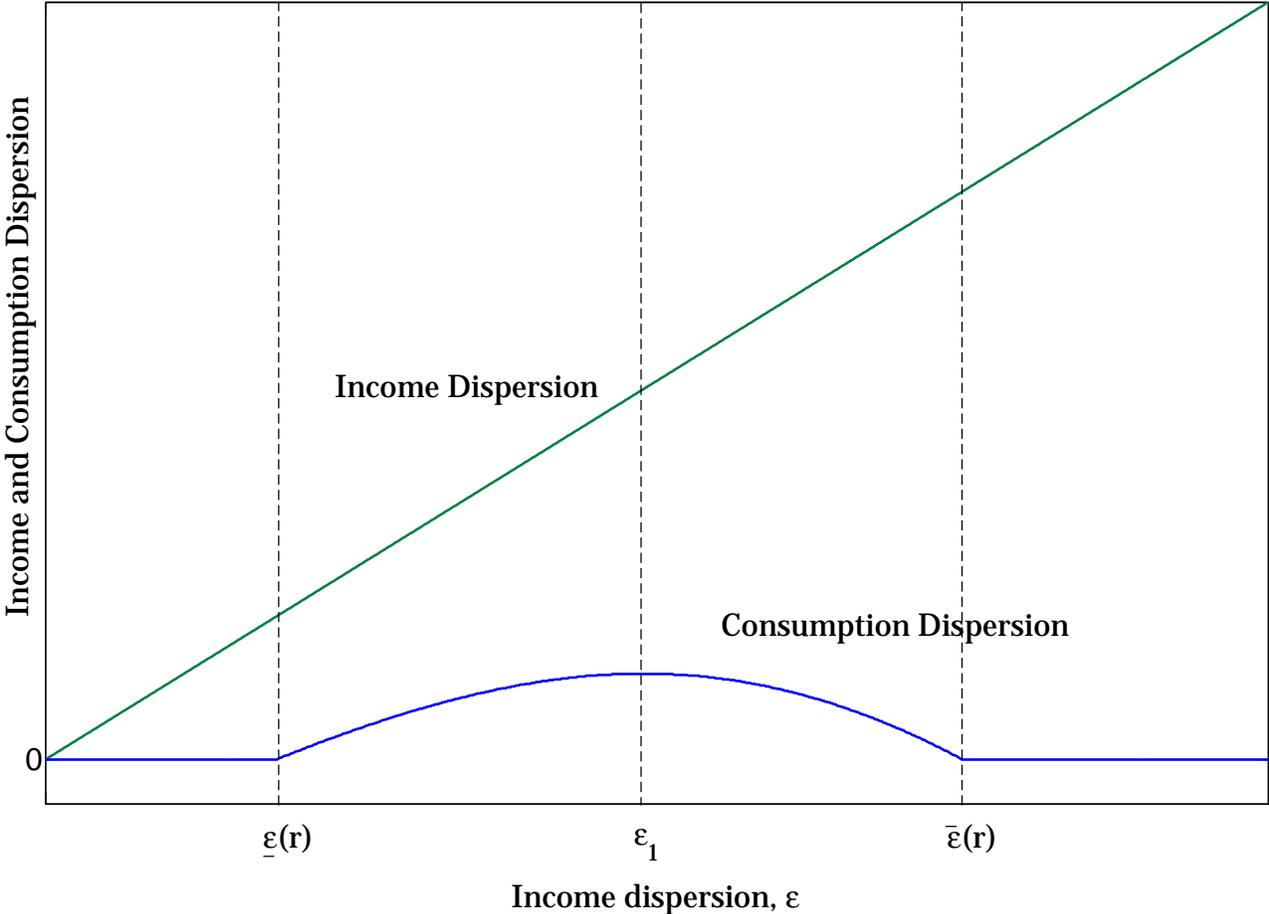
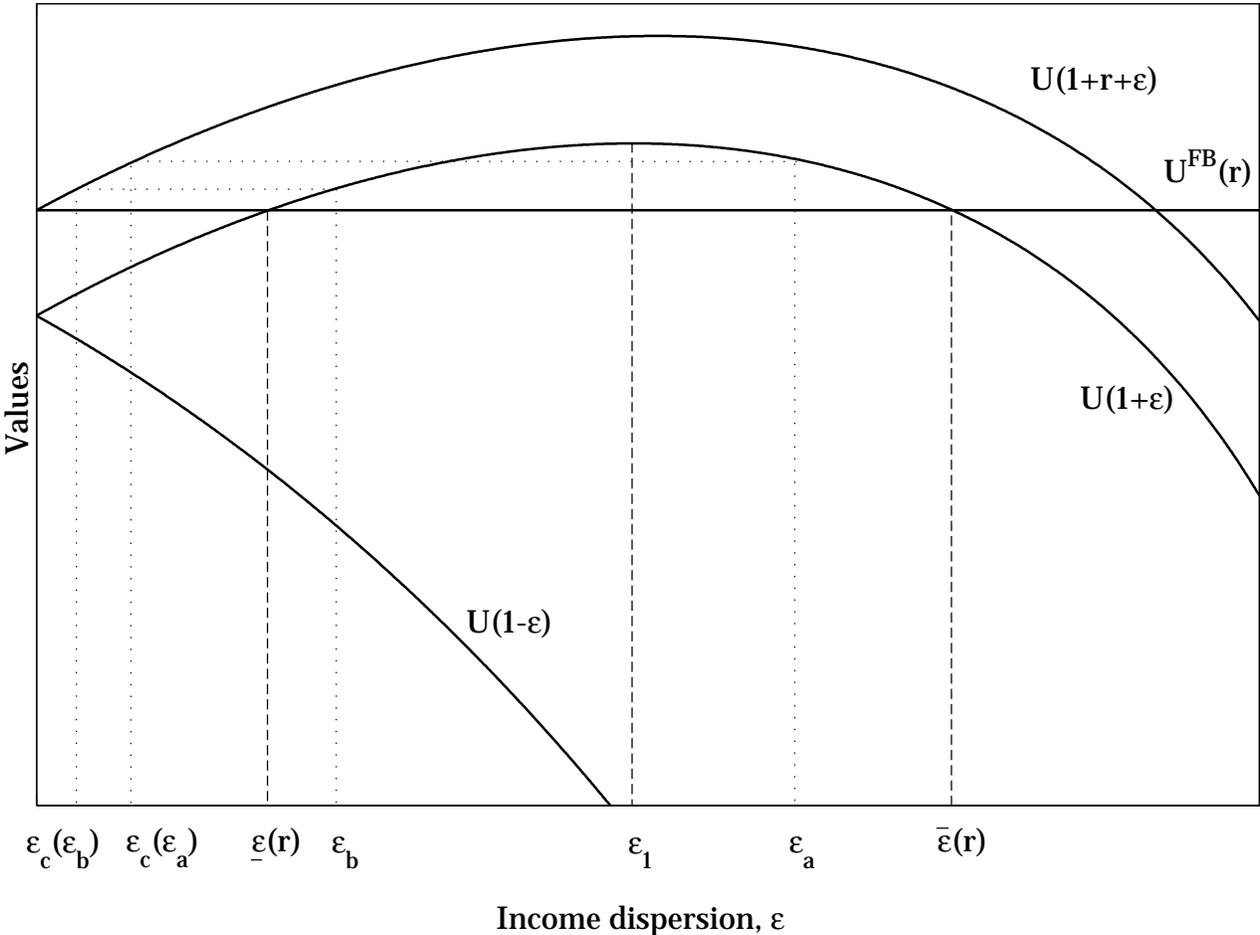
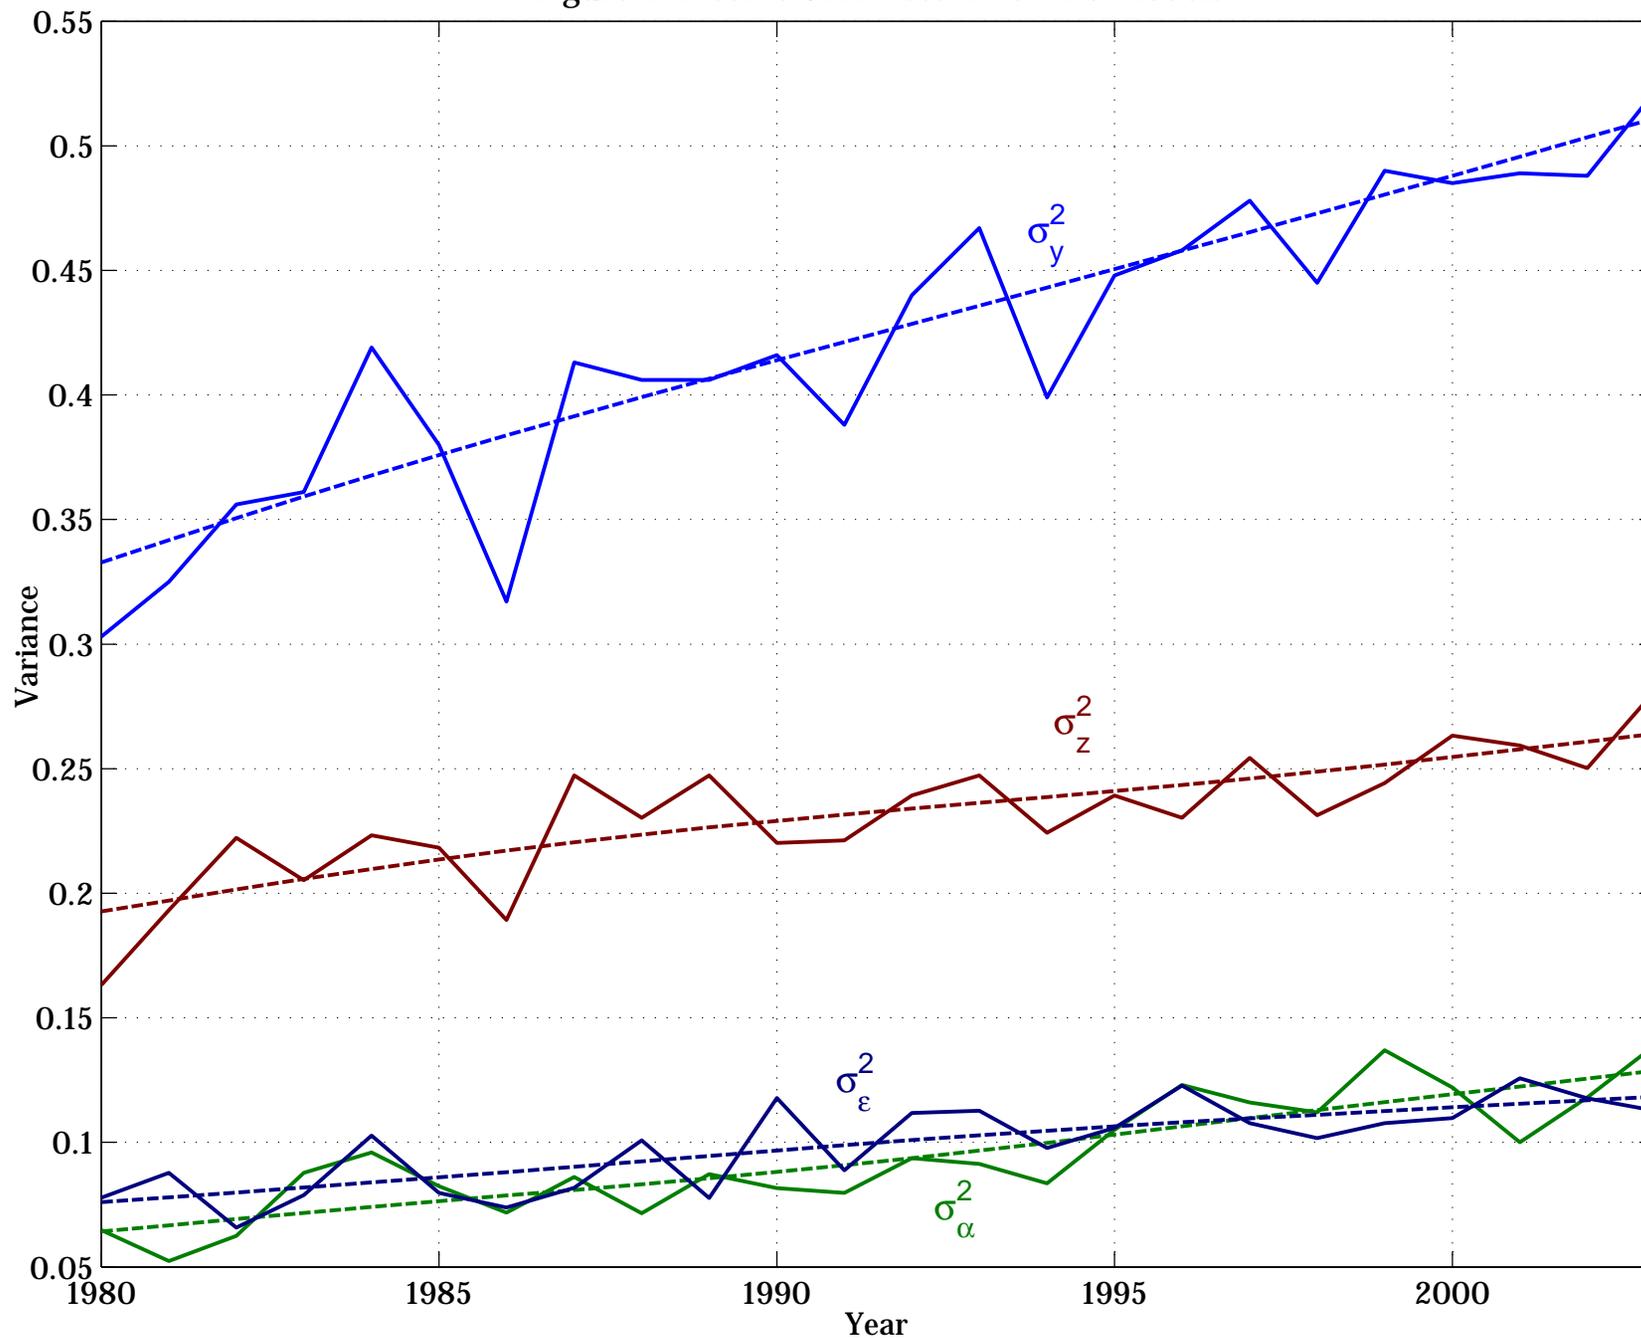


Figure 4. Income Variances: Data and Models



Note: The solid lines are the time series of the cross-sectional variances estimated in the data. The dashed lines (which are the solid lines HP filtered) are the cross-sectional variances used in the models.

Figure 5. Changes in Between- and Within-Group Consumption Inequality: Data and Models

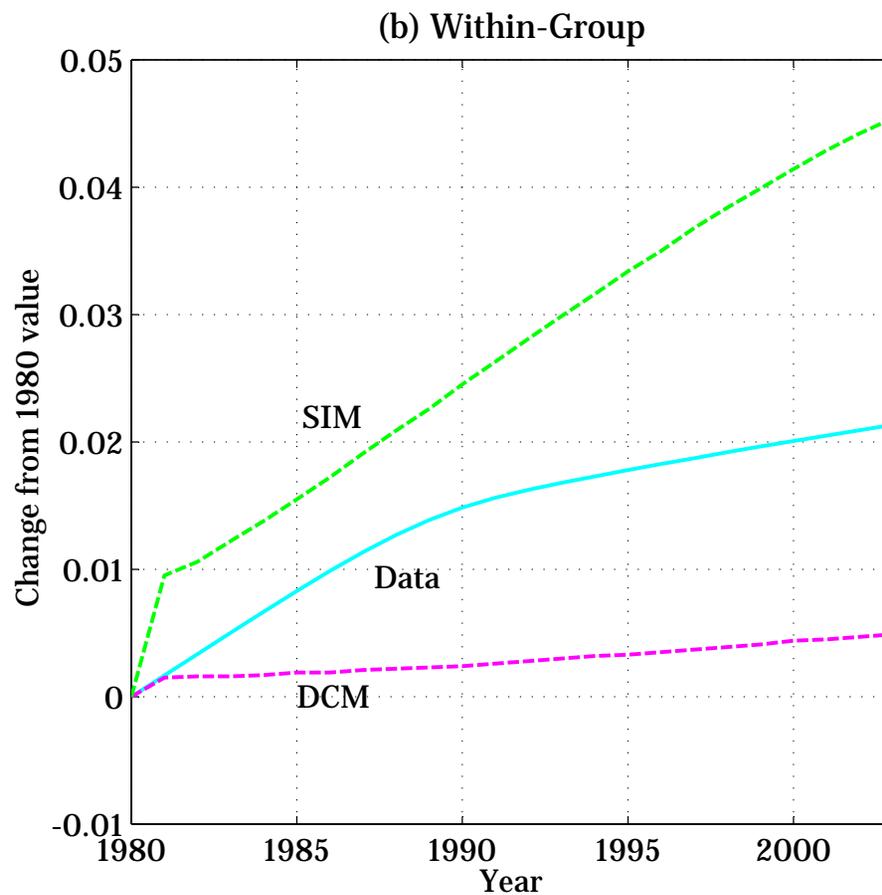
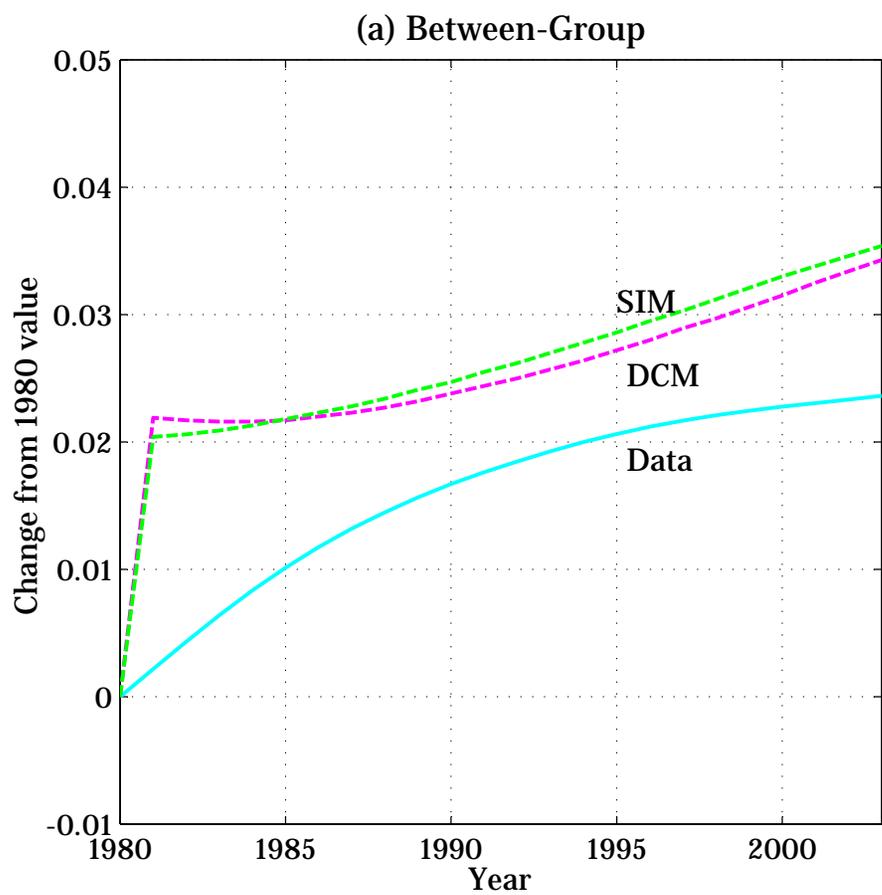


Figure 6. Income Inequality and Consumer Credit

