

August 2003

Does Income Inequality Lead to Consumption Inequality? Evidence and Theory¹

Dirk Krueger

University of Pennsylvania and NBER

Fabrizio Perri

New York University, NBER and CEPR

ABSTRACT

This paper first documents the evolution of the cross-sectional income and consumption distribution in the US in the past 25 years. Using data from the Consumer Expenditure Survey we find that rising income inequality has not been accompanied by a corresponding rise in consumption inequality. Over the period from 1972 to 1998 the standard deviation of the log of after-tax labor income has increased by 20% while the standard deviation of log consumption has increased by less than 2%. Furthermore, income inequality has increased both between and within education groups while consumption inequality has increased between education groups but mildly declined within groups. We then argue that these empirical findings are consistent with the hypothesis that an increase in income volatility has been an important factor in the increase in income inequality, but at the same time has, endogenously, led to a better development of credit markets, allowing households to more effectively smooth their consumption against idiosyncratic income fluctuations. We develop a consumption model in which the sharing of income risk is limited by imperfect enforcement of credit contracts and in which the development of financial markets depends on the volatility of the individual income process. This model is shown to be quantitatively consistent with the joint evolution of income and consumption inequality in US, while other commonly used consumption models are not.

KEYWORDS: Limited Enforcement, Risk Sharing, Consumption Inequality

JEL CLASSIFICATION CODES : E21, D91, D63, D31, G22

¹We thank seminar participants at the SED meetings, ET Conference, NBER Summer Institute, NBER EFG Meeting, Econometric Society Summer Meetings, Minnesota Macro Week, CEPR ESSET, Arizona State, Board of Governors, Bocconi, Bureau of Labor Statistics, Carnegie Mellon, Chicago GSB, Columbia, Di Tella, Duke, the FEDs at Atlanta, Dallas, Minneapolis and Philadelphia, Iowa, Humboldt, Maryland, MIT, Northwestern, NYU, Princeton, Roma La Sapienza, Rutgers, Stockholm School of Economics, TU Berlin, UCLA, UC San Diego and UPenn for many and very helpful comments. Thanks to Harris Dellas and Pierre-Olivier Gourinchas for their thoughtful discussions and to Cristobal Huneus for excellent research assistance. Krueger acknowledges financial support from the NSF under grant SES-0004376. This paper was in part written while the first author was visiting the University of Pennsylvania and the second author was visiting Princeton University. We thank both institutions for their hospitality.

1. Introduction

The sharp increase in earnings and income inequality for the US in the last 25 years is a well-documented fact. Many authors have found that the dispersion of US 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.² 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 measure well how economic well-being is allocated among households in the US.³ Moreover the same change in current income inequality might have a very different impact on the welfare distribution, depending on the structure of credit 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 focused on measures of individual *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) and Blundell and Preston (1998).⁴

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 US developed over the last 25 years. Using data from the Consumer Expenditure Survey, the paper extends and complements the studies mentioned

²See, e.g., Gottschalk and Moffitt (1994), Gottschalk and Smeeding (1997) or Katz and Autor (1998) for recent surveys of these empirical findings.

³Blundell 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.

⁴Even the popular press has been occupied with the cross-sectional consumption distribution. The bestseller by Cox and Alm (1999) argues that the last 25 years were a dazzling economic success story for (almost) all Americans when judged from the consumption experience of households.

in the last paragraph. Our main finding is that despite the surge in income inequality in US in the last quarter of the century, (the Gini index has increased about 10 percentage points while the standard deviation of the logs displays an increase of around 20%) consumption inequality has increased much less (around 2 percentage points in the Gini and less than 2% in the standard deviation of logs). We also document another important difference between the two distributions: income inequality has increased both between and within education groups while consumption inequality has increased between groups but has actually slightly declined within groups.

Second, we go on to propose a theoretical explanation for these stylized facts. It is our hypothesis that an increase in the volatility of idiosyncratic labor income has been an important cause of the increase in income inequality, but that it has also caused a change in the development of financial markets, allowing individual households to better insure against idiosyncratic income fluctuations. We present a simple model of endogenously incomplete markets, building on earlier work by Alvarez and Jermann (2000) and Kehoe and Levine (1993, 2001) that allows us to *analytically* characterize the relationship between income and consumption inequality. In the model agents enter risk sharing contracts, but at any point of time have the option to renege on their obligations, at the cost of being excluded from future risk sharing. Our main result is that whenever there is some sharing of idiosyncratic income risk in the economy an increase in the volatility of income, keeping the persistence of the income process constant, *always* leads to a *reduction* in consumption inequality within the group that shares income risk. Intuitively, higher income volatility increases the value of risk sharing opportunities, therefore reducing the incentives to default. As a consequence, more risk sharing is possible and the consumption distribution becomes less dispersed. We view this model as a simple and analytically tractable way of capturing 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.

Finally, we assess whether an extension of the simple model is *quantitatively* consistent with the stylized facts established in the empirical section of the paper. We develop a model with a large number of agents that face a stochastic labor income process. We choose this income process to match the level and trend of income inequality, both between and within different gender-education 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 consumption distributions. We demonstrate that this model, consistent with the data, can reproduce an increasing between-group consumption inequality and a moderately decreasing within-group consumption inequality. A standard incomplete markets model along the lines of Huggett (1993) and Aiyagari (1994), on the other hand, predicts a significant increase in between- *and* within-group consumption inequality in response to increasing income inequality.

Our quantitative results are consistent with a recent empirical study by Blundell et al. (2002) that rejects full consumption insurance, but documents that households are able to insure income shocks to a larger extent than the permanent income hypothesis (whose general equilibrium extension is the standard incomplete markets model we consider) predicts, pointing to risk sharing mechanisms that we explicitly attempt to model with our endogenous incomplete markets model. Similarly, Storesletten et al. (2000) document that both empirically and in a calibrated life-cycle version of the standard incomplete markets model the cross-sectional dispersion of consumption increases with cohort age, following a similar pattern as for income. The increase with age is less pronounced for consumption than for income, both in the data as well as in their model, but the model overpredicts

the life-cycle increase in consumption dispersion. Again, their study suggests that households appear to have more powerful consumption smoothing abilities than simple self-insurance as the standard incomplete markets model hypothesizes.

The paper is organized as follows: in Section 2 we document our main stylized facts. Section 3 develops a simple two-agent model that can be solved analytically and aims at providing intuition for the quantitative results presented for the models with a large number of agents which is presented in section 4. In Section 5 we lay out our quantitative thought experiment and in Section 6 we discuss the calibration of both models. Section 7 presents our numerical results and assesses the success of both models in explaining the stylized facts documented in Section 2. Section 8 concludes. The recursive formulation of the models as well as computational details can be found in Appendix A1 and details about the data used in the paper are contained in Appendix A2.

2. Trends in Income and Consumption Inequality

In this section we report our main empirical findings. In particular, we document how US income and consumption inequality has evolved over the last quarter of the century. For this purpose our main object of analysis is the Consumer Expenditure (CE) Survey, which is currently the only micro-level data set for the US that reports comprehensive measures of consumption expenditures and income measures for repeated large cross-sections of households.⁵

Our sample is composed of all households in the CE who are complete income respondents, with the reference person between the age of 25 and 64 and who report positive income and positive total consumption expenditure for the interview year (1972-73 samples) or interview quarter (post 1980 samples). This selection generates a sample of around 6300 households per year for the income and consumption distribution in the years 1972-1973, while for the post-1980 period it leaves an

⁵The Panel Study of Income Dynamics (PSID) reports both income and consumption data. The consumption data, however, contains only food consumption and therefore is of limited use for our analysis.

average of 3000 households per quarter in the consumption distribution and 1500 households per quarter in the income distribution.

For these distributions we compute two common measures of income and consumption inequality: the Gini coefficient and the standard deviation of the logarithm. The evolution of these measures is reported in figure 1. The top panels report the actual values of the inequality measures while the bottom panels report, to facilitate the comparison of the trends, the deviation of the indexes from their 1972 value. In each panel the solid line represents inequality of after tax labor income including transfers while the dashed line represents inequality of nondurable consumption expenditures plus expenditures on household equipment plus imputed services from houses and cars (henceforth ND+ consumption expenditures). All variables are measured in constant 1982-84 dollars, deflated by expenditure component-specific CPI's. Income and consumption for each household is divided by the number of adult equivalents in the household using the Census equivalence scale.⁶ The standard deviations are computed on the residuals from regressing income and consumption for each cross section on a quartic in age and on a dummy for the race of the reference person in the household. We treated the data in this way to control for compositional effects stemming from a potential change in the age/race/family structure of the US population over time. Finally the thin dash-dotted lines are standard errors of the inequality measures, computed by performing a bootstrap procedure with hundred repetitions.

Figure 1 confirms the well-known fact that income inequality in the US has increased significantly in the last quarter of the century: the Gini index has risen by about 10 percentage points while the standard deviation of the logs displays an increase of around 20%.⁷ The figure also presents

⁶See Dalaker and Naifeh (1998). We 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.

⁷Krueger and Perri (2003) compare the increase in labor income inequality using CE data with the increase in labor income inequality obtained by using PSID data (from Heathcote, Storesletten and Violante, 2003) and the increase measured by using CPS (from Katz and Autor, 1999) and find that, for the same sample selection, the magnitude of the increase is very similar.

our main empirical finding, namely that the increase in *consumption* inequality has been much less marked; the increase has been around 2 percentage points for the Gini and less than 2% in the standard deviation of logs.⁸ Note that the impact of redistributive public policies (such as progressive income taxation or unemployment insurance) is already included in our income definition so it cannot be responsible for the divergence between the two series. Although the evolution of consumption inequality has been studied much less than 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 80's* has been accompanied by an increase in consumption inequality. Our measures also display an increase in consumption inequality in the early 80's, but it is less marked than the increase in income inequality;⁹ moreover in the 1990s income inequality has continued to rise (although at a slower pace) while consumption inequality has actually slightly declined. This last fact has also been also reported by FED chairman Greenspan (1998) in his introductory remarks to a symposium dedicated to income inequality. In the next subsection we check the robustness of these findings to alternative measures of income and consumption inequality.

A. Alternative Measures of Income and Consumption Inequality

In table 1 we summarize various robustness checks of our main empirical finding. In the first three columns we report the change in consumption inequality obtained using different definitions of consumption expenditures; the first column uses ND+ consumption expenditures (the same definition used in figure 1), the second column uses nondurable consumption expenditures (this is the definition of consumption used by Attanasio and Davis, 1996), while the third column reports the change in inequality for total consumption expenditures. Both alternative measures confirm

⁸Pendakur (1998) finds similar results for Canada between 1978 and 1992 for his preferred measure of consumption.

⁹The increase in consumption inequality from 1972 to 1984 reported by Cutler and Katz (1991) is higher than the increase we find. In a separate appendix (available upon request) we investigate the causes of this discrepancy and find that a significant part of the difference is due to the fact that they select a smaller sample, relative to our study. This strategy increases the time variation of inequality measures. Slesnick (2001) and Attanasio (2002) also computed consumption inequality measures from the CE and find increases that are comparable to the one we document.

that consumption inequality has been quite stable relative to income inequality, with inequality in nondurable consumption expenditures actually decreasing and inequality in total consumption expenditures increasing by less than 1/3 of the increase in income inequality.

TABLE 1. ADDITIONAL MEASURES OF INCOME AND CONSUMPTION INEQUALITY

Period	%Δ Std. Dev. Cons			% Share of Btm Quint.		90/10 Ratio		Gini Income*	
	ND+	ND	TE	LYA+	ND+	LYA+	ND+	CE	CPS
72-73	0	0	0	5.91	9.20	4.96	3.09	0.381	0.397
80-81	NA	-4.6	0.4	5.22	9.57	6.59	3.15	0.410	0.405
85-86	1.7	-2.6	7.9	3.91	9.45	10.24	3.13	0.444	0.422
90-91	0.6	-4.5	5.6	4.41	9.67	8.47	3.14	0.428	0.428
97-98	1.5	-4.0	7.4	4.01	9.24	9.13	3.35	0.456	0.458

*The Gini Income refers to the income concept of household money income before taxes

It is important to keep in mind that total consumption expenditures include cash payments for homes, purchases of cars and even cash contributions toward retirement, and therefore contains a significant part of households' savings, which biases measured consumption inequality towards measured income inequality. For this reason we think of the latter definition as an upper bound for the true change in consumption inequality rather than as an accurate measure for its trend.

The next two columns report the fraction of total after tax labor income plus transfers (labeled LYA+) and ND+ consumption expenditures that accrue to the lowest 20% of the population (where the quintiles are defined with respect to the corresponding cross-sectional income or consumption distribution) in a given period. We view this statistic as an important indicator of how the poorest group in the population has fared in terms of income and consumption. The numbers in the table reveal patterns very similar to those emerging from the data plotted in Figure 1. In particular, we observe a decline in the income share earned by the poorest 20% of the population, from almost 6%

in 1972/73 to 4% in 1997/98. The share of ND+ consumption expenditures of the poorest 20% of the population, however, has remained stable. These 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 in the same period declined from 9.9% to 9.5%.

Columns 6 and 7 address the potential concern about the presence of top-coding in the CE data set. In our empirical analysis we set top-coded data entries equal to the top-coding thresholds; since these thresholds change over time our inequality measures may be affected by these changes. To partially control for these effects we report a statistic that, although less informative about overall inequality, is much less sensitive to the change in top coding thresholds. The 90/10 ratio is the ratio between the income (or consumption) of the household at the 90-th percentile and the income of the household at the 10-th percentile of the distribution. The 90/10 ratio, again, reveals a similar pattern, displaying a large increase in income inequality and a much less marked increase in consumption inequality.

Finally, the last two columns address the issue of the quality of CE income data by comparing the Gini coefficient for income computed from the Current Population Survey, which draws a much larger cross section of US households, to the Gini coefficient computed from the CE. Note that, although the two measures do not perfectly track each other, they both reveal an increase in income inequality of similar magnitude.¹⁰

B. Between and Within-Group Income and Consumption Inequality

Before turning to the theoretical explanation for the empirical findings it is helpful to further investigate the differences between income and consumption inequality by decomposing them in

¹⁰The CPS does not report taxes paid by households so we could not construct an income measure comparable to LYA+. Therefore we report the Gini for total money income in the CPS (see Jones and Weinberg, 2000) and we computed the Gini for the same income measure in the CE. In addition, for the years after 1980 we used NBER's TAXSIM to evaluate taxes paid by households in the CPS and constructed a Gini time series for LYA+ based on the CPS. We found again that this time series displays an increase in inequality very similar to the one obtained from the CE.

between- and within-group inequality. Our theoretical explanation for the lack of consumption inequality to increase with income inequality relies on better insurance of idiosyncratic shocks over time, induced by higher volatility of the idiosyncratic part of income. If our theory is correct, we would expect higher variability in the idiosyncratic, within-group component of income, but constant or even declining variability of the idiosyncratic part of consumption (because of better insurance). We thus decompose the inequality measures reported in Figure 1 into their between- and within-group component.

The empirical decomposition we employ is simple and widely used (see Katz and Autor, 1999). For each labor income and consumption expenditure cross-section (after controlling for age and race effects) we regress income and consumption on education and sex of the head of the household. We choose education and sex to define groups since the increase in the wage skill premium and the decline of the wage gender gap are the two most important determinants of the changes in between-group income inequality in the last 25 years. We then denote the cross sectional variance explained by education and sex 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- and within-group income (panel a) and consumption (panel b) inequality, measured by the log-standard deviation.¹¹ Note that for income both the between- and within-group components display an increase. For consumption, on the other hand, the between-group component displays an increase, not very different in magnitude from that of income.¹² But, most importantly for our purposes, for consumption the within-group idiosyncratic part is actually slightly declining over time, partially offsetting the increase in between-group inequality.

¹¹We show changes in standard deviations since these are easier to interpret than changes in variances.

¹²This finding is highly consistent with the results by 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.

To summarize, the data presented in this section document the well-known increase in income inequality in the last 25 years and the surprising lack of any substantial increase in consumption inequality. These findings are robust to different definitions of consumption and income and to different measures of inequality. We have also demonstrated that consumption inequality has diverged from income inequality mainly because within-group income inequality has increased significantly while within-group consumption inequality has actually slightly decreased. The remaining part of the paper first develops a simple analytical and then a richer dynamic computable general equilibrium model designed to help us understand these facts. In the next section we present a simple model in which we can *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 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. There are two (types of) agents $i = 1, 2$ and a single, nonstorable consumption good in each period. In each period one consumer has income $1 + \varepsilon$ and the other has income $1 - \varepsilon$, so that the aggregate endowment is constant at 2 in each period. Let $s_t \in S = \{1, 2\}$ denote the consumer that has endowment $1 + \varepsilon$. We assume that $\{s_t\}_{t=0}^{\infty}$ follows a Markov process with transition matrix

$$\pi = \begin{bmatrix} \delta & 1 - \delta \\ 1 - \delta & \delta \end{bmatrix}$$

Note that $\delta \in (0, 1)$ governs the persistence of the endowment process while $\varepsilon \in [0, 1)$ measures the variability of the income process.

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 . We assume that $\pi(s_0) = \frac{1}{2}$ for all $s_0 \in S$, so that both agents are ex-ante identical. 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 $\lim_{c \rightarrow 0} u'(c) = \infty$. 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 , from event history s^t onwards and denote by $e = (e^1, e^2)$ the autarkic allocation of consuming the 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 bear the associated costs, which we specify as exclusion from intertemporal trade. This implies that any risk sharing mechanism must yield allocations that deliver to each consumer a continuation utility at least as high as from the autarkic allocation, for all event histories s^t . This is formalized by imposing the following individual rationality constraints on allocations:

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

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

$$c^1 + c^2 = e^1 + e^2$$

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 inequality, as measured by ε .

A. The Constrained Efficient Consumption Distribution

We focus on symmetric allocations.¹³ In order to analyze how the constrained efficient consumption allocations vary with ε it is convenient to solve analytically for the value of autarky.

In this simple economy the continuation value from the autarkic allocation is given by

$$U(1 + \varepsilon) = \frac{1}{D} \{(1 - \beta) u(1 + \varepsilon) + \beta(1 - \delta) [u(1 + \varepsilon) + u(1 - \varepsilon)]\}$$

$$U(1 - \varepsilon) = \frac{1}{D} \{(1 - \beta) u(1 - \varepsilon) + \beta(1 - \delta) [u(1 + \varepsilon) + u(1 - \varepsilon)]\}$$

where $D = [(1 - \beta\delta)^2 - (\beta - \beta\delta)^2] / (1 - \beta) > 0$. Here $U(1 + \varepsilon)$ denotes the continuation utility of autarky for the agent with the currently high income and $U(1 - \varepsilon)$ denotes the continuation utility of the agent with the currently low income. 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 onwards.

Notice that the value of autarky for the agent with high income, $U(1 + \varepsilon)$, 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 $\varepsilon_1 = \arg \max_{\varepsilon} U(1 + \varepsilon) \in (0, 1)$. 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 consumption from tomorrow onwards more and 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 (see Figure 3), since an increase in ε reduces consumption today for this agent and makes it more risky from tomorrow onwards.

Using these properties of the continuation utilities from autarky and the results by Alvarez

¹³A consumption allocation is symmetric if $c_t^1(s^t) = c_t^2(\tilde{s}^t)$, for all t and all s^t, \tilde{s}^t such that $s_\tau = 1$ implies $\tilde{s}_\tau = 2$ for all $\tau \leq t$.

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

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

where $U^{FB} = u(1)$ is the lifetime utility of the first best allocation in which there is complete risk sharing and consumption of both agents is constant at 1, and $U(1 + \varepsilon_c(\varepsilon))$ is the continuation utility of the consumption allocation characterized by $\varepsilon_c(\varepsilon)$.

Note that if $U^{FB} \geq U(1 + \varepsilon)$ the only solution to the above equation is $\varepsilon_c(\varepsilon) = 0$ and the constrained efficient allocation implies full risk sharing. If $U^{FB} < U(1 + \varepsilon)$ the equation above has in general two solutions, with $\varepsilon_c(\varepsilon) = \varepsilon$ (autarky) always being a solution, but not necessarily the smallest one.

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 the rich agent to walk away from the risk-sharing arrangement, with positive time discounting 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.

B. Income Variability and Consumption Inequality

The following proposition characterizes how the constrained efficient symmetric consumption distribution varies with the variability of income, ε .

PROPOSITION 2. *For fixed $\delta \in (0,1)$, starting from a given income dispersion $\varepsilon = \varepsilon_0$ a marginal increase in ε leads to a strict decrease in consumption inequality in the constrained efficient symmetric consumption distribution if and only if $0 < \varepsilon_c(\varepsilon_0) < \varepsilon_0$ (that is, in the initial distribution there is positive, but not complete risk sharing).*

The proof of this proposition follows immediately from proposition 1 and the properties of $U(1 \pm \varepsilon)$. We aim at providing some intuition for the proposition above in Figure 3, in which we plot the value of autarky in the two states and the value of full risk sharing, as a function of the dispersion of income ε . We can divide all possible values for ε into three regions. If $\varepsilon \geq \varepsilon_2$ then the value of autarky in both states is below the value of full risk sharing, hence the full risk sharing allocation ($\varepsilon_c = 0$) satisfies the individual rationality constraints (1) and thus is the constrained efficient consumption allocation. Obviously in this range an increase in ε has no effect on the consumption distribution.

Suppose now that $\varepsilon_1 < \varepsilon < \varepsilon_2$. Consider for example the point $\varepsilon = \varepsilon_h$. From proposition 1 the constrained efficient consumption allocation is given by the smallest solution to $U(1 + \varepsilon_c(\varepsilon_h)) = U(1 + \varepsilon_h)$; from the figure it is immediate that the solution is $\varepsilon_c(\varepsilon_h) = \varepsilon_l$. In this allocation the agent with high income will receive a continuation utility equal to the value of autarky, while the agent with low income receives a continuation utility strictly higher than the value of autarky. Notice from the figure that in this range there is partial, but positive risk sharing ($0 < \varepsilon_c(\varepsilon_h) = \varepsilon_l < \varepsilon_h$). Now a marginal increase in ε from ε_h (an increase in income inequality) reduces the value of autarky for the rich agent and she has less of an incentive to walk away from the risk sharing arrangement.

A smaller current level of consumption is required to make her not default (ε_l moves to the left), thereby reducing the amount of consumption inequality in this economy.

Finally, if $\varepsilon < \varepsilon_1$ (consider for example $\varepsilon = \varepsilon_l$) then autarky is the constrained efficient allocation and $\varepsilon_c(\varepsilon_l) = \varepsilon_l$. Note that in this case there is no risk sharing and a marginal increase in income inequality leads to a one-to-one increase in consumption inequality.¹⁴

To summarize, in this environment with limited commitment an increase of income dispersion has ambiguous effects on consumption inequality, but in general, if the amount of risk sharing in the economy is positive (full), an increase in income inequality will reduce (not increase) consumption inequality. The intuition behind the result 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 it is the fear of being poor again in the future that makes a currently rich agent transfer resources to his currently poor brethren.

This analysis suggests that the endogenous evolution of (formal, market-based or informal) risk-sharing mechanisms can indeed generate a declining within-group consumption inequality despite an increasing within-group income inequality. In the next section we will consider the same mechanism in an economy with a continuum of agents which face a more realistic income process that also allows for changes in between-group inequality, with the goal of evaluating its quantitative relevance.

¹⁴It is also straightforward to show that an increase in persistence δ leads to an increase in consumption inequality 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 and Levine (2001).

The intuition for this result 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 less transfers to the poor agent being sustainable. Graphically, in Figure 3, the graph for $U(1 + \varepsilon)$ tilts around the origin, upward for an increase in δ . For a given $\varepsilon = \varepsilon_h$ with partial risk sharing, the corresponding consumption allocation $\varepsilon_c(\varepsilon, \delta) = \varepsilon_l$ shifts to the right due to this increase in δ .

4. The Model with Large Number of Agents

A. The Environment

There is a continuum of consumers of measure 1. Individuals are of types $i \in \{1, \dots, M\}$, with p_i denoting the fraction of the population being of type i . We interpret these different types or “groups” of agents as capturing heterogeneity in the population with respect to sex and education, fixed characteristics that affect an individuals’ wage and therefore income. Since relative wages for educated individuals and women have increased substantially over the last 25 years, and thus are partially responsible for the recent trends in income inequality, an incorporation of this type of heterogeneity appears to be critical for any quantitative study on income and consumption inequality.

There is a single, nonstorable consumption good. An individual of type i has a stochastic endowment process $\{\alpha_{it}y_t\}$ where α_{it} is the deterministic type-specific, possibly time-varying mean endowment and $\{y_t\}$ follows a Markov process with finite support Y_t , a set with cardinality N . 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 will 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 denote by y_t the current period 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 endowment shock history y^t . We furthermore assume that at date 0 the measure over current endowments is given by $\Pi_0(\cdot)$. At date 0 agents are distinguished by their type i , their initial asset holdings (claims to period zero consumption) a_0 and by the their initial income shock y_0 . Let Φ_0 be the initial distribution over types (i, a_0, y_0) . Finally, agents’ preferences are exactly

as described in the simple model of the previous section.

B. Market Structures

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

Endogenous Incomplete Markets

An individual of type (i, a_0, y_0) starts with initial assets a_0 and trades Arrow securities subject to pre-specified credit lines $A_t^i(y^t, y_{t+1})$ that are contingent on observable endowment histories and an agents' type,¹⁵ 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 endowment shock history and time, in order to reflect deterministic changes in the income process and hence in the magnitude of endowments $\alpha_{it}y_t$.

Consider the problem of an agent of type i with initial conditions (i, a_0, y_0) (we suppress the dependence of functions on i whenever there is no room for confusion). The agent chooses, conditional on his 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)

$$(2) \quad (1 - \beta) \left(u(c_0(a_0, y_0)) + \sum_{t=1}^{\infty} \sum_{y^t|y_0} \beta^t \pi(y^t|y_0) u(c_t(a_0, y^t)) \right)$$

$$(3) \quad \text{s.t.} \quad c_t(a_0, y^t) + \sum_{y_{t+1}} q_t(y^t, y_{t+1}) a_{t+1}(a_0, y^t, y_{t+1}) = \alpha_{it}y_t + a_t(a_0, y^t) \quad \forall y^t$$

$$(4) \quad a_{t+1}(a_0, y^t, y_{t+1}) \geq A_{t+1}^i(y^t, y_{t+1}) \quad \forall y^t, y_{t+1}$$

Now we will specify the short-sale constraints $A_t^i(y^t, y_{t+1})$ in more detail. Following Alvarez

¹⁵Note that we rule out any insurance against being of a particular "type" i . We will further comment on this assumption and its implications in the calibration section.

and Jermann (2000) we will define “solvency constraints” that are not too tight. As before let by $U_t^{Aut}(i, y_t)$ denote the continuation utility from consuming the endowment from period t onwards, given current endowment realization $\alpha_{it}y_t$. Given a sequence of prices $\{q_t\}_{t=0}^{\infty}$ and short-sale constraints $\{A_t^i(y^t, y_{t+1})\}_{t=0}^{\infty}$, 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 (3) and (4). Short-sale constraints $\{A_t^i(y^t, y_{t+1})\}_{t=0}^{\infty}$ are not “too tight” if they satisfy

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

That is, the constraints are such that an agent of type i , having borrowed up to the borrowing constraint, $a_{t+1}(a, y^t, y_{t+1}) = A_{t+1}^i(y^t, y_{t+1})$ for state (y^t, y_{t+1}) , is indifferent between repaying his debt and defaulting, with the consequence of default being specified as exclusion from future access to financial markets (i.e. being expelled into autarky), as in the simple model of the previous section.

DEFINITION 1. *Given Φ_0 , a competitive equilibrium with solvency constraints $\{A_t^i(y^t, y_{t+1})\}_{t=0}^{\infty}$ that are not too tight is allocations $\{c_t^i(a_0, y^t), a_{t+1}^i(a_0, y^t, y_{t+1})\}_{t=0, i \in M}^{\infty}$, prices $\{q_t\}_{t=0}^{\infty}$ and measures $\{\Phi_t\}_{t=1}^{\infty}$ such that*

1. (Optimization) *Given prices, allocations $\{c_t^i(a_0, y^t), a_{t+1}^i(a_0, y^t, y_{t+1})\}_{t=0}^{\infty}$ maximize (2) subject to (3) and (4) and the solvency constraints are not “too tight”.*
2. (Market clearing)

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

3. (Equilibrium Laws of Motion) $\Phi_{t+1} = H_t(\Phi_t)$.

In this definition the equilibrium laws of motion for measures $\{H_t\}$ are induced by the transition probabilities π_t 's and the functions $a_t^i(a_0, y^t, y_{t+1})$. Now suppose that the deterministic part of income is constant across time for all types, $\alpha_{it} = \alpha_i$. Then we define a stationary equilibrium (or steady state) as an equilibrium for which, for all $t \geq 0$ we have $\Phi_t = \Phi$ and $q_t = q$.

Notice that the dispersion of the income process affects the solvency constraints and thus the extent to which individual agents can borrow in exactly the same way as it affected the extent of risk sharing in the simple model of Section 3. In particular, an increase in the dispersion of the income process not only increases the necessity of extended borrowing to smooth consumption, but also the possibility of extended borrowing, since the default option becomes 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 the cross-sectional dispersion of consumption.

Standard Incomplete Markets

We will compare our results to those obtained in a standard incomplete markets model, as in Huggett (1993) or Aiyagari (1994). That model is a special case of the economy described above, with $A_t^i(y^t, y_{t+1}) = -\alpha_{it}\bar{B}$ and the absence of a full set of contingent claims. Let q_t^{in} denote the price, at period t , of a safe claim to one unit of the consumption good in period $t+1$. The sequential budget constraints the agent faces are (again suppressing type indexation for the allocations)

$$(5) \quad c_t(a_0, y^t) + q_t^{in} a_{t+1}(a_0, y^t) = \alpha_{it} y_t + a_t(a_0, y^{t-1})$$

and the short-sale constraints become

$$(6) \quad a_{t+1}(a_0, y^t) \geq -\alpha_{it}\bar{B}.$$

We let $R_t^{in} = \frac{1}{q_t^{in}}$ denote the risk free gross real interest rate in the standard incomplete markets economy. The definition of equilibrium and stationary equilibrium for this economy is similar to the one discussed above and hence omitted. Notice that the only difference between the two economies

is the set of financial assets that are traded (a full set of contingent claims in our economy, only a single uncontingent bond in the standard incomplete markets economy) 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. For details and for the computational algorithm employed please refer to Appendix A1. Note that the computation of the equilibrium in the endogenous incomplete markets model is nonstandard as one has to solve for both prices and borrowing constraints simultaneously.

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 for both economies so that the stationary equilibrium in both economies matches key observations of the US economy in the early 70's. This applies, in particular, to the deterministic and stochastic part of the income process, the key quantitative ingredient of our models.
2. We then introduce a finite path of changes in the dispersion of the income process to mimic the increase in income inequality observed in US data as documented in Section 2. We assume that this change in the income process is unforeseen by agents, but that all future changes in the income process are fully learned once the first change has occurred.
3. The change in the income process for a finite number of periods induces a transition in both models from the initial to a final stationary equilibrium corresponding to the income process that prevails once the path of income dispersion changes 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 and other macroeconomic statistics of interest for both models and compare them to the main stylized

empirical facts established in Section 2. In order to carry out these steps we first have to specify the parameters of both models in our calibration section.

6. Calibration

The following parameters need to be chosen: a) preference parameters: the time discount factor β and the coefficient of relative risk aversion σ (as we will assume a constant relative risk aversion utility function) b) the individual agents' endowment process $\{\alpha_{it}y_t\}_{t=0}^{\infty}$ with $y_t \in Y_t = \{y_{1t}, y_{2t}, \dots, y_{Nt}\}$ and $\alpha_{it} \in A_t = \{a_{1t}, a_{2t} \dots a_{Mt}\}$ and the transition matrices π_t for the stochastic part of the endowment process c) the fractions of the population p_i being of a particular type i and d) the borrowing constraint \bar{B} for the standard incomplete markets model.

A. Income Process

We take the length of a model period to be one year. An individuals' income $\alpha_{it}y_t$ consists of a type-specific, possibly time-dependent deterministic part α_{it} and an idiosyncratic, type-independent part y_t . In order to map income inequality in our model to that in the data we have to give empirical content to the notion of a "type" or "group" i and measure how the inequality of income between types changed over the time period of interest. The residual cross-sectional variability of income (and its changes over time) will then be attributed to the idiosyncratic part of income y_t .

In the empirical section our main measure of inequality was the standard deviation of the logarithm of income (and consumption) after age and race effects were removed. The main remaining, and therefore modeled, elements of observable heterogeneity are education and sex of the household. We interpret the types i as standing in for this heterogeneity.

As described in the empirical section, in order to decompose the cross-sectional variance of household income into between-type and within-type variance we follow Katz and Author (1999)

and write the logarithm of income $\ln(e_{it})$ as

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

and obtain $\sigma_{et}^2 = \sigma_{\alpha t}^2 + \sigma_{yt}^2$ where $\sigma_{et}^2 = \text{Var}[\ln(e_{it})]$, $\sigma_{\alpha t}^2 = \text{Var}[\ln(\alpha_{it})]$ and $\sigma_{yt}^2 = \text{Var}[\ln(y_t)]$. For each year we then regress the logarithm of income on household head education and sex dummies and identify $\sigma_{\alpha t}^2$ with the part of the variance in log-income that is explained by the regression and attribute the remainder of the variance to σ_{yt}^2 , the variance of the idiosyncratic part.

In doing so we end up with three time series from the data, $\{\sigma_{et}^2, \sigma_{\alpha t}^2, \sigma_{yt}^2\}_{t=1972}^{1998}$ (where $\{\sigma_{\alpha t}, \sigma_{yt}\}_{t=1972}^{1998}$ were plotted in Figure 2). Our calibration strategy is to choose parameters governing the model income process so that a) in the initial stationary equilibrium both the between- and within-type income variance of the model matches the data for the early 70's and b) along the transition trends in between- and within-type income variances are reproduced by the model income process. The results from this procedure, to be described in detail below, along with the empirically estimated standard deviations, are summarized in Figure 4.

Between-Group Income Inequality

We pick the number of types to be 2 with equal mass $p_i = 0.5$ in the population. For the initial stationary equilibrium we choose the type-specific mean for type 1 as $\alpha_1 = e^{-\sigma_{\alpha}1972}$ and for type 2 as $\alpha_2 = e^{\sigma_{\alpha}1972}$. Similarly, using $\sigma_{\alpha}1998$ we obtain average group incomes for the final steady state, persisting from 1998 into the indefinite future. For the transition path we then select $\{\alpha_{1t}, \alpha_{2t}\}_{t=1973}^{1997}$ so that the trend of between-group income inequality follows that in the data.¹⁶ See Figure 4 for the results.

Our specification of between-group income inequality deserves further discussion. First, remember that $\{\alpha_{1t}, \alpha_{2t}\}_{t=1972}^{1998}$ is a deterministic sequence; second, notice that after 1998 between-

¹⁶We do not attempt to explain the high frequency movements in consumption inequality with our model and thus do not fit the high frequency movements in income inequality in our calibration. For that reason also we matched the average variance between 1972 and 1973 in the initial steady state.

group inequality is assumed to permanently remain at its higher, 1998 level. These facts imply, in the context of both models considered in this paper, that the increase in between-group income inequality unambiguously translates 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 theoretically in Section 3.

We chose this specification for two reasons. First, in an influential paper Attanasio and Davis (1996) show that between-group consumption insurance spectacularly 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. Second, we will be able to quantify exactly to what extent the endogenous evolution of credit lines predicted by the endogenous incomplete markets model is able to offset the increase in between-group consumption inequality via reducing within-group consumption inequality. The magnitude of this effect depends on the exact calibration of the idiosyncratic part of the income process, to which we turn next.

Within Group Income Variability

We model the idiosyncratic part of an individuals’ income process, $\ln(y_t)$, as a simple AR(1) process, as, for example, in Storesletten et al. (1998, 2000), Aiyagari (1994), Heaton and Lucas (1996) and many others. In particular, we assume that

$$\ln(y_t) = \rho \ln(y_{t-1}) + \varepsilon_t$$

where ε_t is a serially uncorrelated and normally distributed random variable with zero mean and variance $\sigma_{\varepsilon_t}^2$. Following Storesletten et al.’ (1998) estimates from PSID data we choose $\rho = 0.98$,

that is, idiosyncratic income shocks are quite persistent.¹⁷ We will report below how sensitive our results are to changes in the persistence of the income process.

The choice of $\sigma_{\varepsilon t}^2$ is governed by the same principle as the specification of between-group income variability above. For the initial steady state we pick $\sigma_{\varepsilon 1972}^2$ so that the associated stationary distribution of the stochastic process, discretized with the Tauchen method¹⁸, has variance equal to $\sigma_{y 1972}^2$ (in logs) as in the data. A similar procedure is followed for the final steady state. We then pick $\{\sigma_{\varepsilon t}^2\}_{t=1973}^{1997}$ so that the variance of $\ln(y_t)$ implied by the model stochastic process follows the trend of the within-group variance of log-income as measured in our data, $\{\sigma_{y t}^2\}_{t=1973}^{1997}$. Again Figure 4 shows the result.¹⁹

Note that after 27 model periods (1998 in real time) the change in the dispersion of the income process is completed and the income process in the model does not change anymore. However, due to the endogenous wealth dynamics in both models it may take substantially longer than these 27 years for both economies to complete the transition to the new stationary consumption and wealth distribution.

B. Preference Parameters and the Borrowing Limit

We assume that the period utility is logarithmic, $u(c) = \log(c)$. We then chose β to match a real risk-free interest rate of 2.5%²⁰ for the initial steady state of the endogenous incomplete

¹⁷After taking out individual fixed effects, the process Storesletten et al. (1998) estimate can be written as

$$\begin{aligned}\ln(y_t) &= \eta_t + \nu_t \\ \eta_t &= \hat{\rho}\eta_{t-1} + \xi_t\end{aligned}$$

In the simple AR(1) process we use, the correlation between $\ln(y_t)$ and $\ln(y_{t-1})$ equals ρ , whereas in theirs it equals $\frac{\hat{\rho}\sigma_{\xi}^2}{\sigma_{\xi}^2 + (1-\hat{\rho}^2)\sigma_{\nu}^2} = 0.97$ with their empirical parameter estimates (see their Table 3, row 3). We choose our ρ so that the implied correlation between $\ln(y_t)$ and $\ln(y_{t-1})$ of the discretized Markov chain equals 0.97, which requires roughly $\rho = 0.98$ as input into the Tauchen procedure (see Tauchen and Hussey, 1991), as reported in the text.

¹⁸We choose, as compromise between computational feasibility and realism, the number of states of the discretized process to be $N = 9$.

¹⁹The calibrated income process not only tracks the income standard deviations in the data, but is also capable to reproduce the trends of the income Gini coefficient from the data.

²⁰This is the average real return of AAA municipal bonds (which are tax-exempt) for the sample period.

markets economy, which yields a value of $\beta = 0.971$. For the standard incomplete markets economy we will report results for various combinations of the discount factor β and the exogenous borrowing constraint \bar{B} to obtain the same real risk free interest rate of 2.5%. As a benchmark we set $\bar{B} = 2$; we normalize endowment in such a way that this borrowing limit corresponds to a generous 2 times the average annual income for each type i . The time discount factor required to obtain an interest rate of 2.5% with this borrowing constraint is $\beta^{in} = 0.9495$.

7. Quantitative Results

A. Income and Consumption Inequality

Figure 5 summarizes the first main quantitative result of our paper. It shows the dynamics of the standard deviation of log-consumption for US data and contrasts it with those generated by the endogenous incomplete markets model as well as the standard incomplete markets model. Quantitatively, the endogenous incomplete markets is quite closer to the data as it slightly overpredicts an increase in the standard deviation of log consumption of 7 percentage points while the standard incomplete markets predicts an increase in the same measure of almost 20 percentage points, very close to the increase in the log-income standard deviation (23 percentage points, by construction, in both models). Our findings for the standard incomplete markets model mirror those of Blundell et al. (2002) and Storesletten et al. (2000) discussed in the introduction and suggest that households have ways of insuring against consumption fluctuations beyond self-insurance. The quantitative results from both models are robust to different measures of inequality. A picture almost identical to Figure 5 would emerge if, instead of the standard deviation of log-consumption we would plot the Gini coefficient. Also note that the standard incomplete markets substantially overpredicts the increase in consumption inequality even if consumption is measured as total consumption expenditures in the data (20 percentage points in the model v/s 7 points in the data).

B. Between- and Within-Group Consumption Inequality

One important feature of the endogenous incomplete markets model, motivated by the empirical findings of Attanasio and Davis (1996), is the differential response of consumption inequality to increases in between- and within-group income inequality. The model implies that in response to increases in within-group income inequality, within-group consumption inequality should *decline*, due to improved intra-group risk sharing, whereas increases in between-group income inequality should translate into increasing between-group consumption inequality. In the empirical section above we documented that, in fact, the data shows exactly this differential response of between- and within-group consumption inequality to increased income inequality.

To what extent does our model predict this response correctly, in a quantitative sense? Figure 6 shows the decomposition, for the data (as in Figure 2) and for both models. Panel (a) shows that *quantitatively* the endogenous incomplete markets model reproduces the trends of consumption inequality within and between groups fairly closely; the mechanism of endogenously expanding credit leads to a decline of within-group consumption inequality of 1.7%, slightly smaller than in the data, but partially offsetting the increase in between-group consumption inequality of 9.1%. Combining these two observations the model predicts a moderate increase of consumption inequality over the last 27 years. One discrepancy between the model and the data is the timing of the increase in between group inequality. In the data the increase is gradual while the model predicts a very sharp increase at the beginning of the transition. We conjecture that this discrepancy has to do with our simplifying assumption that in 1972 agents in the model learn the entire future evolution of the between-group part of their income, and start to act on this information immediately.

Panel (b) shows why the standard incomplete markets model overstates the increase in consumption inequality relative to the data: that model predicts a substantial increase in within-group consumption inequality while in the data it actually slightly declined.

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

C. Sensitivity Analysis

In this section we document how sensitive our main findings are to changes in the parameterization of the model. In particular, we want to investigate whether the relative failure of the standard incomplete markets model is due to the fact that borrowing constraints are set tight or that income shocks are assumed to be very persistent.

In Table 2 we report results for different (β, \bar{B}) combinations that lead to an interest rate of 2.5% in the initial steady state of the standard incomplete markets model. For comparison we also repeat the corresponding results from the endogenous incomplete markets model and the empirical statistics. The first column of the table (labeled ΔStd) reports the change in total consumption inequality between 1972 to 1998 in percentage points, the second (labeled ΔStd_p) and the third (labeled ΔStd_t) report the change in between- and within-group consumption inequality.

TABLE 2. CHANGE IN CONSUMPTION INEQUALITY

Economy	ΔStd	ΔStd_p	ΔStd_t
$(\beta, \bar{B}) = (0.9405, 1)$	20.6	11.6	17.4
$(\beta, \bar{B}) = (0.9495, 2)$	19.4	11.0	16.4
$(\beta, \bar{B}) = (0.957, 4)$	19.0	10.7	16.0
Endo. Inc. Markets	7.0	9.1	-1.7
Data	1.5	7.6	-2.2

We observe that in the standard incomplete markets model, as the exogenous borrowing

limit is relaxed, the increase in consumption inequality induced by the increase in income inequality declines. This is to be expected as looser borrowing constraints allow agents to better smooth consumption by borrowing against future income. Increasing the borrowing constraint further below $\bar{B} = 4$ does not result in a substantial further decline in the consumption inequality increase: in the light of persistent income shocks an individuals' willingness to borrow against future income and to pay interest on the corresponding loans is limited. Most importantly, independent of the borrowing constraint the standard incomplete markets model tends to overpredict the increase in consumption inequality quite substantially.

At least since Friedman (1957) it is well-understood that more persistent income shocks are harder to self-insure against than income shocks that are transitory in nature. Storesletten et al. (1998) measure income shocks as being very persistent; we now document how the standard incomplete markets model performs if income shocks weren't as persistent as they appear to be in the data. Table 3 (in which we report the same model statistics as in table 2) documents our findings; all experiments take as given a borrowing constraint of $\bar{B} = 2$, and adjust the time discount factor to obtain an interest rate of 2.5% in the initial steady state of the corresponding model. Apart from the benchmark value of $\rho = 0.98$ we experiment with a $\rho = 0.925$ (which is an alternative estimate of the persistence parameter reported in Storesletten et al. 1998) and with a $\rho = 0.53$, which is the value estimated by Heaton and Lucas (1996).

TABLE 3. CHANGE IN CONSUMPTION INEQUALITY

Economy	ΔStd	ΔStd_p	ΔStd_t
Stan. Inc. Markets, $\rho = 0.98$	19.4	11.0	16.4
Stan. Inc. Markets, $\rho = 0.925$	19.1	11.3	15.7
Stan. Inc. Markets, $\rho = 0.53$	21.3	12.7	17.2
Endo. Inc. Markets, $\rho = 0.98$	7.0	9.1	-1.7
Endo. Inc. Markets, $\rho = 0.925$	8.4	9.2	-1.4
Endo. Inc. Markets, $\rho = 0.53$	9.3	9.3	-0.0
Data	1.5	7.6	-2.2

We see that the increase in consumption inequality does *not* decrease with the persistence of the income shocks. This is true for both models, albeit for different reasons. Note that, as the persistence of the income shocks is reduced one has to increase their variability $\sigma_{\varepsilon t}^2$ in order to obtain the same level of income inequality. In the endogenous incomplete markets model a reduction of ρ and an increase in $\sigma_{\varepsilon t}^2$ reduces the value of default, therefore making more consumption insurance of idiosyncratic shocks enforceable. Our mechanism for reducing within-group consumption variability is weaker if there is less within-group consumption variability to start with. At the extreme, for a $\rho = 0.53$ or lower all *idiosyncratic* income shocks (although not between-group income shocks) can be perfectly insured against in the initial steady state already, and no further reduction in within-group consumption dispersion is possible.

On the other hand, for the standard incomplete markets model, while Friedman's intuition is correct and less persistent shocks *of the same magnitude* can be smoothed better, in the presence of borrowing constraints shocks of *larger magnitude* may not be, as Table 4 shows. Thus, if one adheres to the principle that the model income process has to reproduce the empirical observations about trends in income inequality, lowering the persistence of the income process does not help the

standard incomplete markets model in explaining the recent trends in consumption inequality.

Therefore, even with persistence as low as $\rho = 0.53$, the standard incomplete markets model seems to predict too strong an increase in consumption inequality, compared to what is observed in US data.

D. Market Completeness or Endogenous Borrowing Constraints?²¹

There are two main differences between our endogenous incomplete markets model and the standard incomplete markets model. First, our model features a full set of Arrow securities and second, borrowing constraints adjust endogenously to changes in the income process. It is therefore instructive to analyze whether the result of roughly constant consumption inequality in our model is mainly due to the fact that our set of assets completely spans the underlying uncertainty or from flexible short-sale constraints.

Table 4 (in which we report the same model statistics as in table 2) summarizes the predictions for consumption inequality implied by the endogenous incomplete markets model, the standard incomplete markets model and two additional models. The first is identical to our endogenous incomplete markets model, with the exception that we freeze the short-sale constraints for Arrow securities at their initial levels, i.e. don't let them respond endogenously to changes in the income process over time. The other model, named after Zhang (1997), is similar to the standard incomplete markets model in that agents can only trade a risk-free, uncontingent bond. The borrowing constraint, however, is now allowed to vary over time. In particular, 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 financial autarky from thereon.

²¹We thank Pierre Olivier Gourinchas for helpful discussions leading to this subsection.

TABLE 4. CHANGE IN CONSUMPTION INEQUALITY

Economy	ΔStd	ΔStd_p	ΔStd_t
Endogenous. Inc. Markets	7.0	9.1	-1.7
Standard. Inc. Markets	19.4	11.0	16.4
Zhang Economy	23.7	13.5	19.5
Fixed Bor. Constraints	11.1	8.7	6.9
Data:	1.5	7.6	-2.2

Table 4 suggests that the success of our model to generate fairly flat consumption inequality requires the combination of *both* full spanning and endogenously evolving debt constraints. If these constraints are not allowed to adjust, even with a full set of Arrow securities the increase in income inequality is accompanied by a fairly substantial increase in consumption inequality, inconsistent with the data. Not surprisingly, this difference to the endogenous incomplete markets model is due to the fact that now within-group consumption inequality increases drastically, since the mechanism by which better insurance against (higher) idiosyncratic income fluctuations emerges is shut down.

On the other hand, in the Zhang economy consumption inequality follows income inequality almost one to one. With highly persistent idiosyncratic income shocks agents are hesitant to borrow to smooth consumption. Therefore the value agents place on access to credit markets is not very high (compared to autarky). Thus, credit lines are very tight and increase only minimally with the rise in income volatility. Consequently the equilibrium allocation in this economy is very close to autarky and consumption inequality follows income inequality very closely.

We conclude that, in order to explain the diverging trends of income and consumption inequality requires, within the context of the models explored, both a large scope of insurance markets (full spanning) as well as an increase in the scale of these markets (an expansion of credit lines).

8. Conclusions

In this paper we use CE survey data to document that the increase in income inequality for the US 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, at least partially, driven by an increase in idiosyncratic labor income risk, then the value households place on access to formal market-based and informal credit and insurance mechanisms rises. If these mechanisms, as endogenous response to increasing income risk, become more developed then agents can make more use of them. Individual consumption is better insulated against (higher) income risk and the cross-sectional consumption distribution fails to fan out with the cross-sectional income distribution. 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 rise in consumption inequality.

The mechanism through which agents, in the endogenous incomplete markets model of the last section, keep their consumption profiles stable in the light of more volatile income, is an expansion in the amount of non-collateralized credit available to consumers. Did this expansion take place in the data? One simple (but of course only partial) measure of the credit available and used by US consumers is the ratio of unsecured consumer credit to disposable income. In Figure 7 we plot the trends for the ratio of aggregate consumer credit to disposable income from US *data* for the last 40 years, and again the income Gini.²² Both are quite flat until the mid 1970s and then show a similar upward 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

²²The series for consumer credit is from the 2002 Economic Report of the President, table B77. The inequality index is the Gini Index for income of families that is available starting in 1959 in the US Census Historical Income Inequality Tables, Table F4. We eliminate the cyclical components from each series using a Hodrick-Prescott filter with a smoothing parameter equal to 100.

to (starting in the mid 1970's) in order to insulate consumption from bigger income fluctuations. We want to stress that we view the expansion of credit lines in the endogenous incomplete markets 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 of these mechanisms.

Conditional on our empirical and quantitative findings a logical next step for future research is to identify and decompose the importance of these exact mechanisms that enable better insurance against income risk over the last decades. A more detailed analysis of cross-section micro-level data sets, with particular emphasis on variables that measure in-kind transfers, market- and non-market intermediated credit as well as other explicit or implicit income insurance mechanisms appears to be called for, given the results documented in this paper. We defer this to ongoing and future research.

References

- [1] Aiyagari, R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109, 659-684.
- [2] Alvarez, F. and U. Jermann (2000), “Efficiency, Equilibrium, and Asset Pricing with Risk of Default,” *Econometrica*, 68, 775-798.
- [3] Attanasio O. (2002), “Consumption and Income Inequality: What Do We Know and What Can We Learn from it?”, *mimeo*.
- [4] Attanasio, O. and S. Davis (1996), “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104, 1227-1262.
- [5] Blundell, R. and I. Preston (1998), “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, 113, 603-640.
- [6] Blundell, R., L. Pistaferri and I. Preston (2002), “Partial Insurance, Information and Consumption Dynamics,” *mimeo*.
- [7] Bureau of Labor Statistics, (1978) *Consumer Expenditure Survey: Integrated Diary and Interview Survey, 1972-73, Bulletin No. 1992*, Washington, D.C.
- [8] Cox, M. and R. Alm (1999): *Myths of Rich & Poor: Why We're Better Off Than We Think*, Basic Books, New York.
- [9] Cutler, D. and L. Katz (1991a), “Macroeconomic Performance and the Disadvantaged,” *Brookings Papers on Economic Activity*, 1-74.
- [10] Cutler, D. and L. Katz (1991b), “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's,” *American Economic Review*, 82, 546-551.
- [11] Dalaker, J. and M. Naifeh (1998), “Poverty in the United States: 1997,” *Current Population Report P60-201*, Bureau of the Census, Washington.
- [12] Deaton, A. and C. Paxson (1994), “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 102, 437-467.
- [13] 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.
- [14] Dynarski, S. and J. Gruber (1997), “Can Families Smooth Variable Earnings?,” *Brookings Papers on Economic Activity*, 229-303.
- [15] *Economic Report of the President* (2002), Washington, D.C.
- [16] Friedman, M. (1957), *A Theory of the Consumption Function*, Princeton University Press, Princeton.
- [17] Greenspan, A. (1998), “Opening Remarks to the Symposium on Income Inequality: Issues and Policy Options,” Federal Reserve Bank of Kansas City.

- [18] Gottschalk, P. and R. Moffitt (1994), "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 217-272.
- [19] Gottschalk, P. and T. Smeeding (1997), "Cross-National Comparisons of Earnings and Income Inequality," *Journal of Economic Literature*, 35, 633-687.
- [20] Heathcote J. , Storsletten K. and G. Violante (2003), "The Macroeconomic Implications of Rising Wage Inequality in the US ", *mimeo*
- [21] 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.
- [22] Huggett, M. (1993), "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953-969.
- [23] 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.
- [24] Johnson, D. and T. Smeeding (1998), "Measuring the Trends in Inequality of Individuals and Families: Income and Consumption," *mimeo*.
- [25] Jones A. and D. Weinberg (2000), "The Changing Shape of the Nation's Income Distribution," *Current Population Reports P60-204*, U.S. Census Bureau.
- [26] Katz, L. and D. Autor (1999), "Changes in the Wage Structure and Earnings Inequality," in: Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics*.
- [27] Kehoe, T. and D. Levine (1993), "Debt Constrained Asset Markets," *Review of Economic Studies*, 60, 865-888.
- [28] Kehoe, T. and D. Levine (2001), "Liquidity Constrained Markets versus Debt Constrained Markets," *Econometrica*, 69, 575-598.
- [29] Krueger, D. and F. Perri (1999), "Risk Sharing: Private Insurance Markets or Redistributive Taxes?," *Federal Reserve Bank of Minneapolis Staff Report 262*.
- [30] Krueger, D. and F. Perri (2003), "On the Welfare consequences of the increase in Inequality in the United States", forthcoming, *18th NBER Macro Annual*.
- [31] Mayer, S. and C. Jencks (1993), "Recent Trends in Economic Inequality in the United States: Income versus Expenditures versus Well-Being," in Papadimitriou, D and E. Wolff (eds.), *Poverty and Prosperity in the USA in the Late Twentieth Century*, St. Martin's Press, New York.
- [32] Nelson, J. (1994), "On Testing for Full Insurance Using Consumer Expenditure Survey Data," *Journal of Political Economy*, 102, 384-394.
- [33] Pendakur, K. (1998), "Changes in Canadian Family Income and Family Consumption Inequality between 1978 and 1992," *Review of Income and Wealth*, 44, 259-83.
- [34] Slesnick, D. (1993), "Gaining Ground: Poverty in the Postwar United States," *Journal of Political Economy*, 101, 1-38.

- [35] Slesnick, D. (2001), *Consumption and Social Welfare: Living Standards and Their Distribution in the United States*, Cambridge University Press.
- [36] Storesletten, K., C. Telmer and A. Yaron (1998), "The Risk Sharing Implications of Alternative Social Security Arrangements," *Carnegie-Rochester Conference Series on Public Policy*, 50, 213-59.
- [37] Storesletten, K., C. Telmer and A. Yaron (2000), "Consumption and Risk Sharing over the Life Cycle," *mimeo*.
- [38] Tauchen, G. and R. Hussey (1991), "Quadrature-Based Methods for Obtaining Approximate Solutions to Non Linear Asset Pricing Models," *Econometrica*, 59, 371-396.
- [39] Zhang, H. (1997), "Endogenous Borrowing Constraints with Incomplete Markets," *Journal of Finance*, 52, 2187-2209.

Appendix

A1. Recursive Formulation and Computational Algorithm

Now we formulate the consumer problem for the endogenous incomplete markets economy 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 will also omit the distinction by types. The equilibrium problem is nonstandard as 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

$$U^{Aut}(y) = (1 - \beta)u(y) + \beta \sum_{y' \in Y} \pi(y'|y)U^{Aut}(y')$$

We 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 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:

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

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

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

for all y' and i . 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 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 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 condition. 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 piece-wise linear functions over the state space. For more details on the algorithm and on the theoretical characterization of the stationary equilibrium, see Krueger and Perri (1999).

A2. Data Description

Our statistics are based on repeated cross sections constructed from the Consumer Expenditure Survey (CE) for the years 1972-1973, 1980-81 and 1984 to 1998, as provided by the Bureau of Labor Statistics. The 1972-1973 samples were conducted quarterly, but only annual totals were released; thus for these years we have only two cross sections, each reporting consumption and income for the year of the interview. The surveys from 1980 on were conducted and reported on a quarterly basis; therefore we have four cross sections for each year. A fraction of households in the

survey is interviewed for four consecutive quarters. Households report consumption expenditures for the quarter preceding the interview month and income data for the year preceding the interview month. Income questions are asked only in the first and fourth quarter. Following the suggestions by Nelson (1994) we exclude incomplete income respondents from our sample. We also only include households with the reference person between the age of 25 and 64 who report positive income and positive total consumption expenditure for the interview year (1972-73 samples) or interview quarter (post 1980 samples).

Income Data

Definition The income definition we use is total household labor income after taxes plus transfers. We construct labor income as total wages and salaries plus a fixed fraction of self-employment farm- and non-farm income (The exact fraction is 0.864 and is taken from Diaz Jimenez, Quadrini and Rios Rull 1996). From labor income we subtract reported federal, state and local taxes (net of refunds) and social security contributions paid by the household. We then add reported government transfers: in particular we add unemployment insurance, food stamps and welfare.

Top-Coding In the 1972 and 1973 samples income is top-coded and bottom-coded; if the total annual income before taxes of a household is below 2,000\$ or above 35,000\$ no component of income is reported in dollar values, and only information about whether the household had received a positive amount of income in that component is available. In this case we proceed as follows. If the household does not report positive income in any of the components of after tax labor income plus transfers (as defined above) we set its after-tax labor income plus transfers to 0 (and thus exclude the household from our sample); if it reports positive income in at least one of the components we set its after-tax labor income plus transfers to match average after tax labor income plus transfers for households with a total income below 3,000\$ (in the case of bottom-coded individuals) or with a total income above 25,000\$ (in the case of top-coded individuals). The latter two figures are obtained from table 1 in the Bureau of Labor Statistics Bulletin (1978). For the samples starting in the first quarter of 1980, whenever income components are top-coded we set them to their top-coding thresholds. We have experimented with changing the values of top-coded income components; inequality measures are robust to these changes, as in general the number of top-coded households is very small.

Aggregation Once we constructed an income measure for all households we compute all inequality statistics, weighting the households by their sample weight provided by the CE. For the years 1972-1973 we have only one cross section per year; hence there are no time aggregation issues. For calendar years after 1980 in each year we have four observations for each inequality statistic. The annual statistic is then computed by taking a weighted average of the quarterly statistics. The weights are proportional to the overlap between the calendar year and the year for which the income is reported by the household. For example, the Gini for the first quarter of 1981 enters with a weight of 1/4 in the 1981 Gini and with a weight of 3/4 in the 1980 Gini.

Consumption Data

Definitions In the paper we use three definitions of consumption expenditures. The first definition, labeled ND+, includes expenditures on nondurable goods and services, expenditures on household equipment and imputed service flows from houses and cars. Expenditures on nondurable goods and services include consumption expenditures for food, alcoholic beverages, tobacco, utilities, personal care, household operations, public transportation, gasoline and motor oil, apparel, education, reading, health services and miscellaneous expenditures. Each component of consumption is

deflated by its corresponding monthly CPI from the Bureau of Labor Statistics. Expenditures on household equipment include items such as furniture, appliances and floor coverings, such as rugs. The reason for why we use expenditures and not imputed services is that in the CE no information is available for the value or the inventory of the stock of this type of equipment, and the panel dimension of the CE is too short to carry out perpetual inventory techniques. With respect to vehicles, we impute services from cars in the following manner, following closely the procedure outlined by Cutler and Katz (1991a). From the CE data we have expenditures for purchases of new and used vehicles. We also have data on the number of cars a consumer unit possesses. For each year we first select all households that report positive expenditures for vehicle purchases, and run a regression of vehicle expenditures on a constant, age, sex and education of the reference person of the consumer unit, total consumption expenditures, excluding vehicle expenditures of the consumer, the same variable squared, total income before taxes, family size and quarter dummies. We use the estimated regression coefficients to predict expenditures for vehicles for all households in that quarter (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 vehicles the consumer unit owns, times $\frac{1}{32}$ (reflecting the assumption of average complete depreciation of a vehicle after 32 quarters) plus other expenditures for cars, such as insurance, maintenance and finance charges. With respect to housing services the CE provides information on rent paid for the residence of the consumer unit, including insurance and other out-of-pocket expenses paid by the renter. To impute housing services for those consumer units that own their residence 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.²³ This variable is not available for all years of the sample, in particular not for the years 1980-81 and 1993-94; for those years we do not compute inequality measures for ND+ consumption expenditures.²⁴ As with nondurable consumption, all imputed services from consumer durables and housing are deflated with the corresponding CPI. The second definition is nondurable consumption expenditure (as defined above). The third definition is total consumption expenditure, which includes all direct out-of-pocket expenditures made by the consumer units and is a variable reported in the CE.

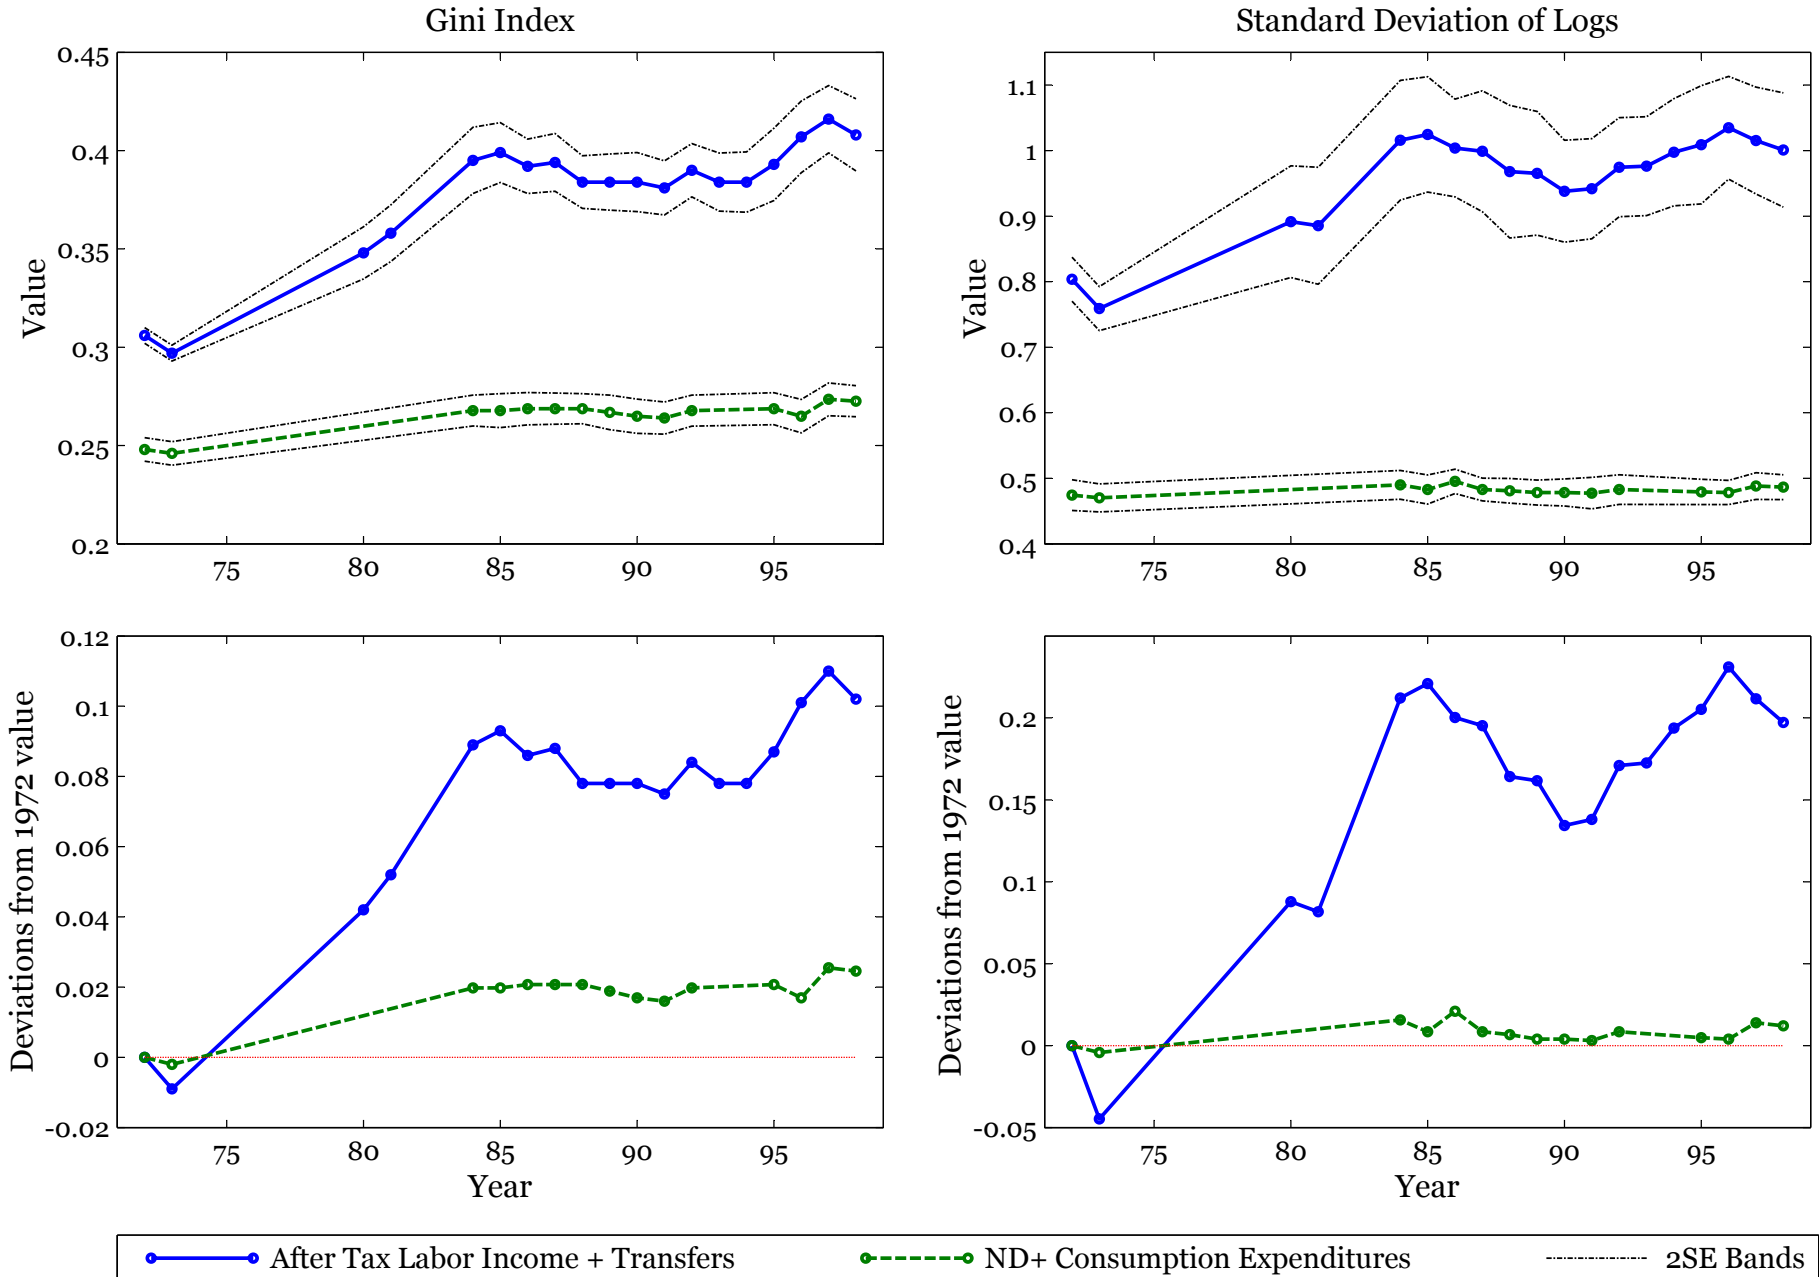
Top-Coding Top-coded consumption expenditures are set equal to the top-coding threshold.

Aggregation Weighting and time aggregation for consumption is dealt with similarly to income. The one additional issue is that inequality measures for 72-73 are based on annual consumption expenditures, while post-1980 inequality measures are based on quarterly expenditures. In order to make these measures comparable we use the following procedure: for the post-1980 sample we select only households that are interviewed for 4 consecutive quarters (this procedure reduces the sample by a half). For these households we aggregate consumption over the year and compute annual inequality measures. We then use the average ratio between the annual and quarterly inequality measures to re-scale the quarterly consumption inequality measures and make them comparable to the annual measures of 1972-1973. We have also experimented with directly using the annual consumption inequality measures for the post-1980 sample; the trends of consumption inequality were unaffected. The time series, however, displays more volatility and higher standard errors due to the smaller sample size.

²³The exact question that the reference person of the CU is asked is “If you were to rent your home today, how much you think it would rent for monthly, without furnishings and utilities?”

²⁴We experimented with using an imputation procedure similar to the one used for vehicles in order to obtain housing services for the four missing years. Results were very similar and are available upon request.

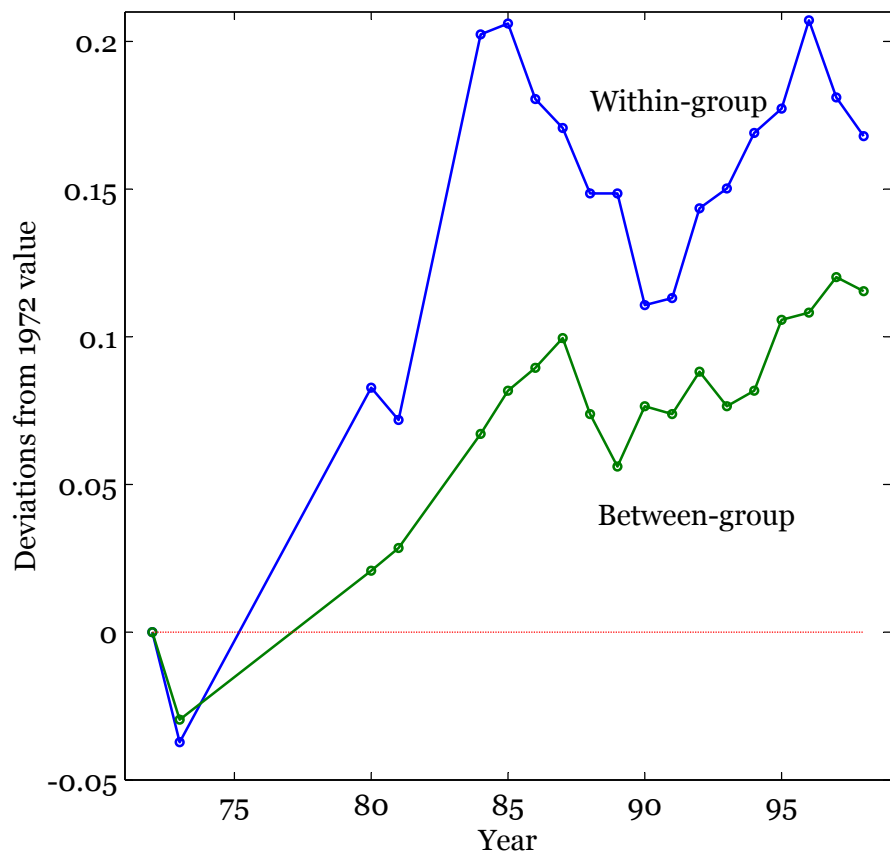
Figure 1. The Evolution of Income and Consumption Inequality in the US



Note: Circles represent actual data points. All indexes are computed on cross sections of households in the CE survey. Income and Consumption are per adult equivalent. Standard deviations are based on the residuals from regressing, in each cross section, household consumption or income on controls for age and race of the reference person

Figure 2. Decomposition of Income and Consumption Inequality: Data

(a) Decomposition of Std. Dev. of Log-Income



Note: Circles represent actual data points

(b) Decomposition of Std. Dev. of Log-Consumption

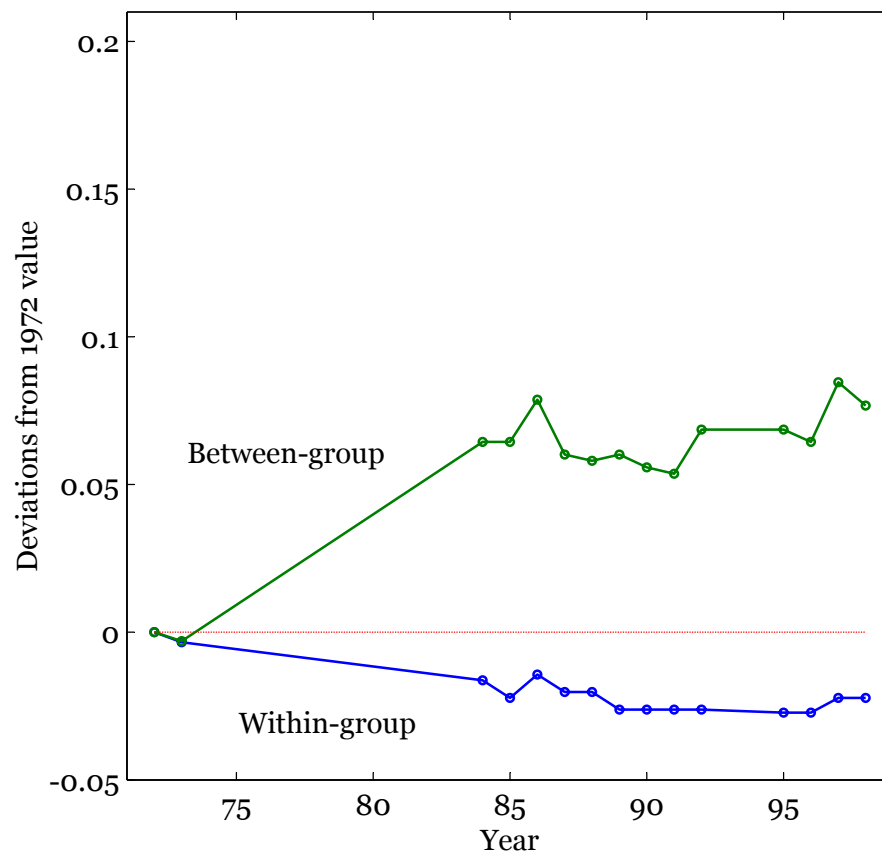


Figure 3. Characterizing the Efficient Allocation

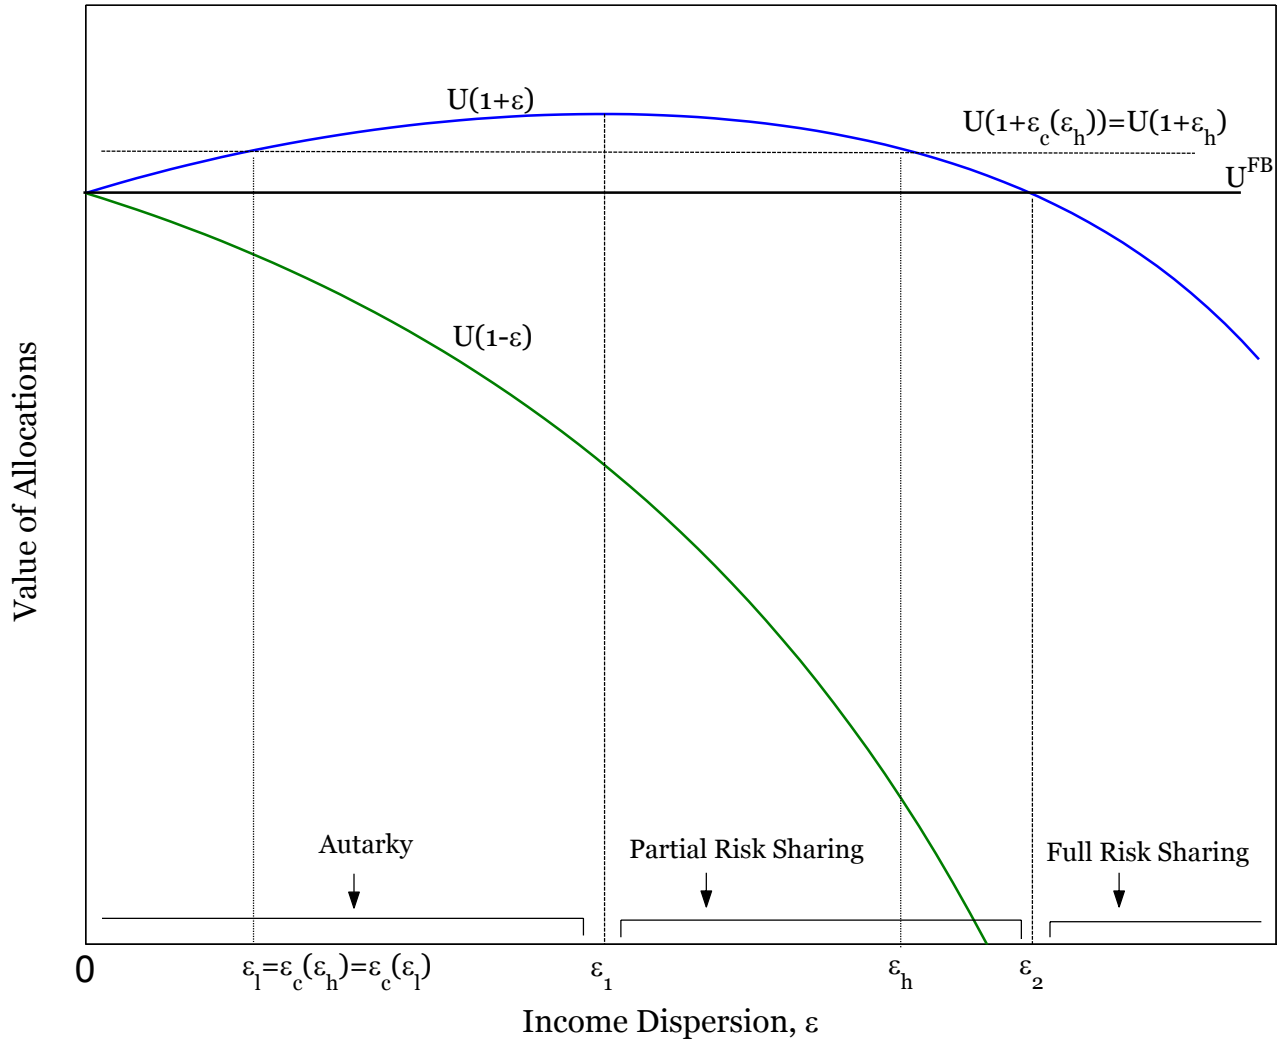


Figure 4. Standard Deviation of Log-Income for Data and Models

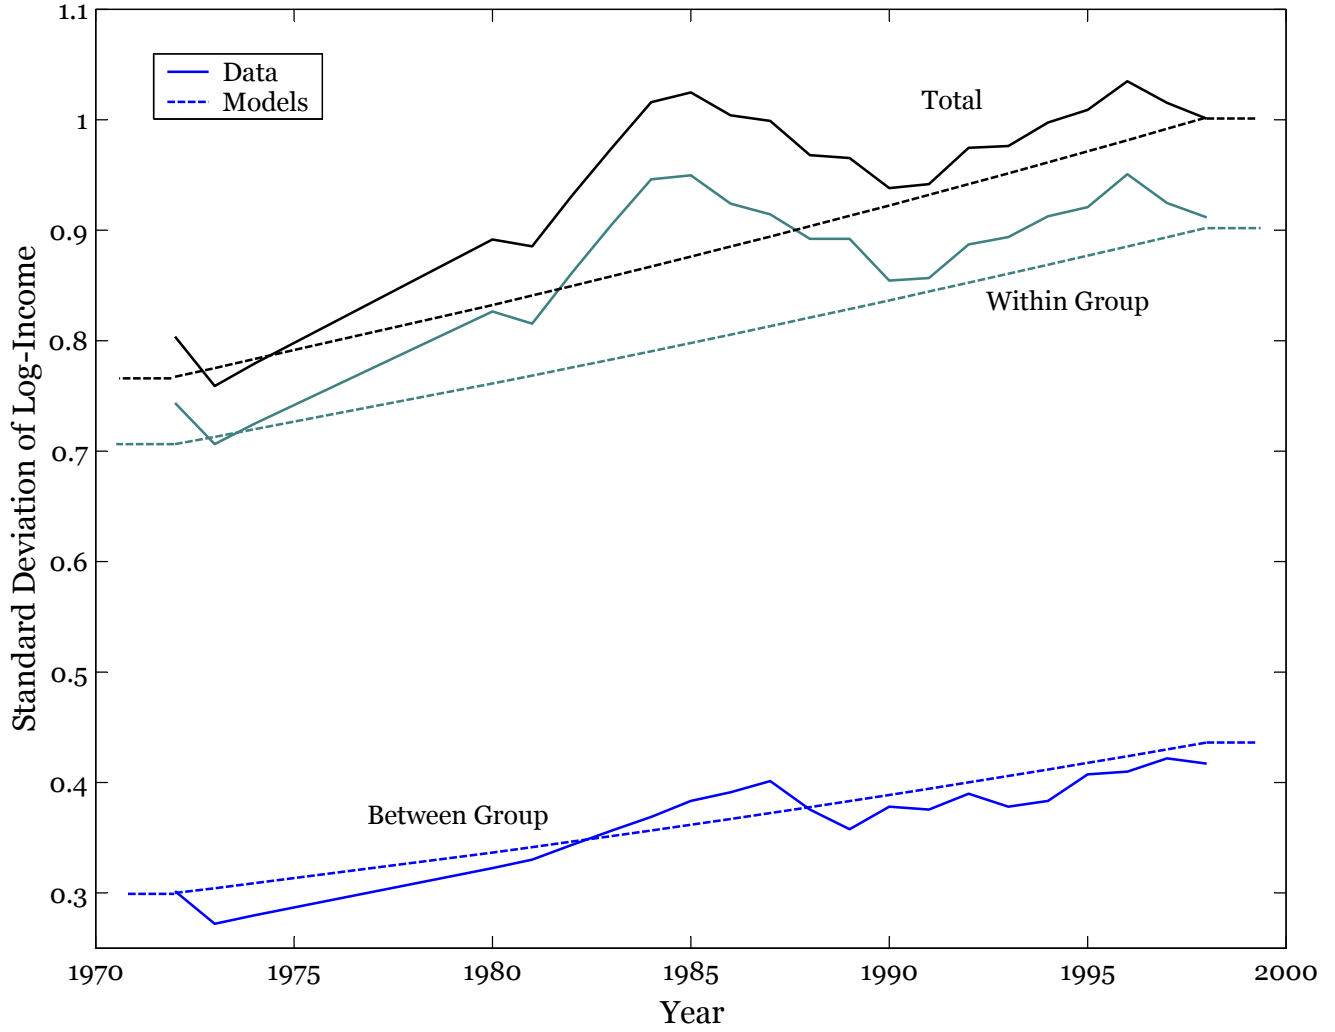


Figure 5. Standard Deviation of Log-Consumption: Data and Models

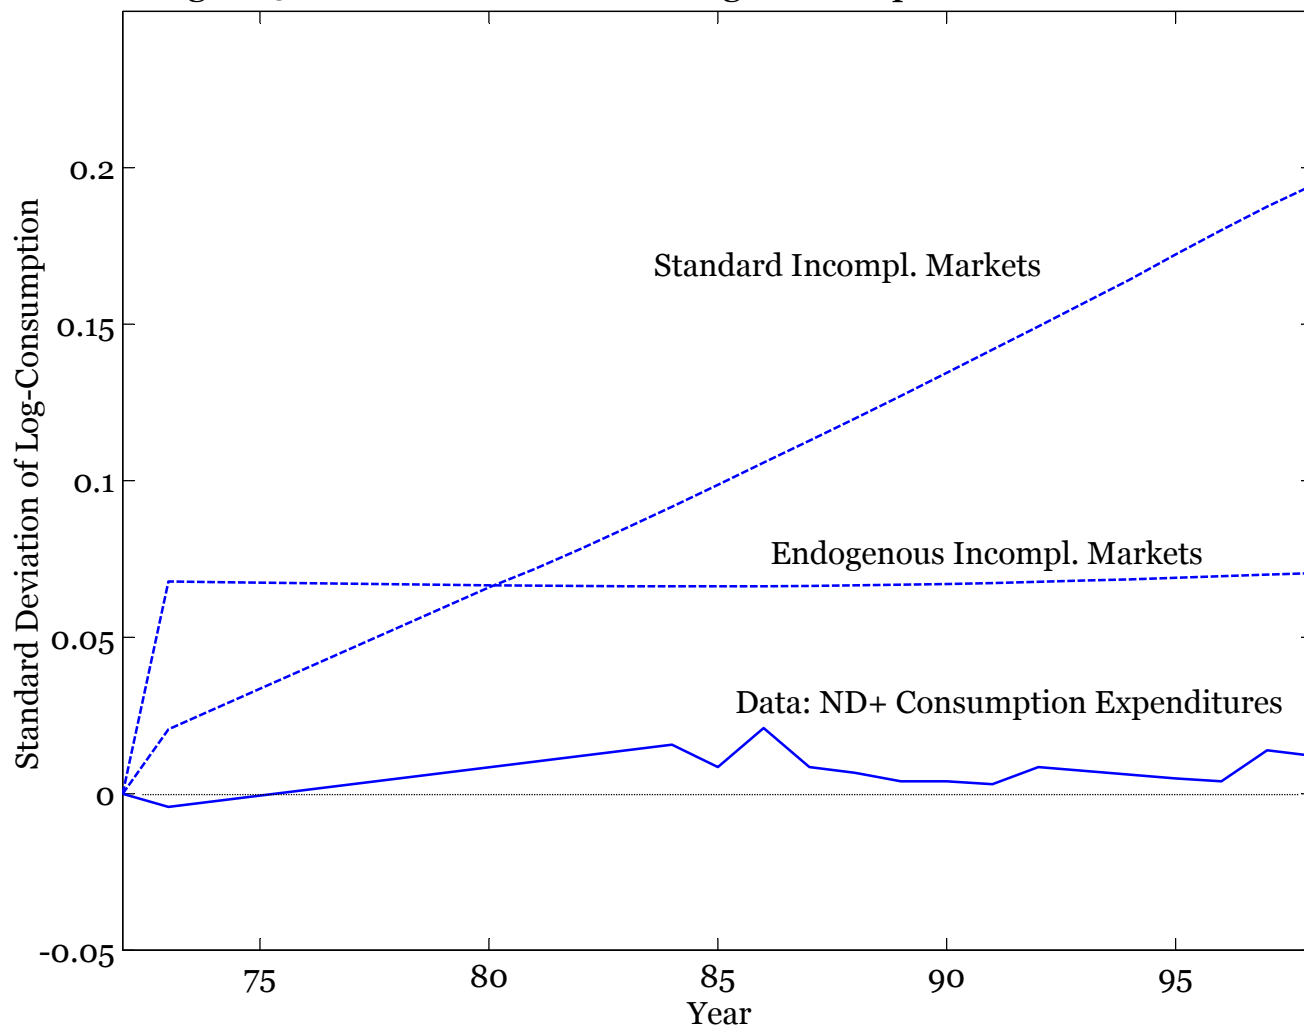


Figure 6. Decomposition of Standard Deviation of Log Consumption: Data and Models

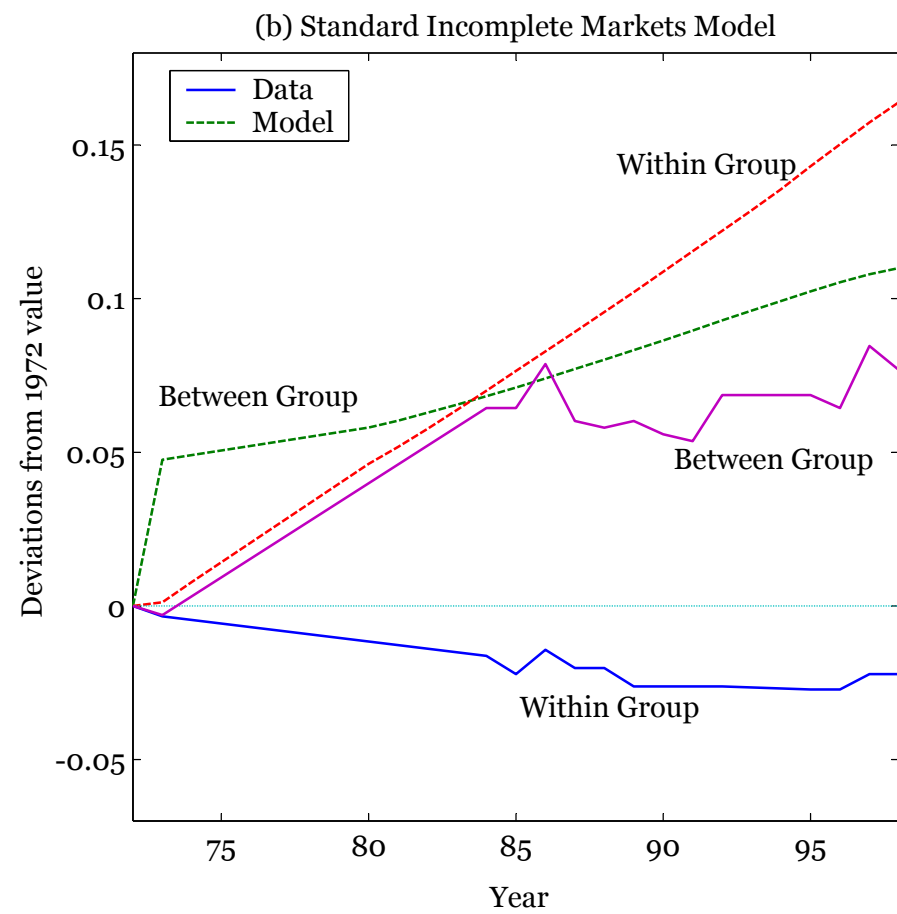
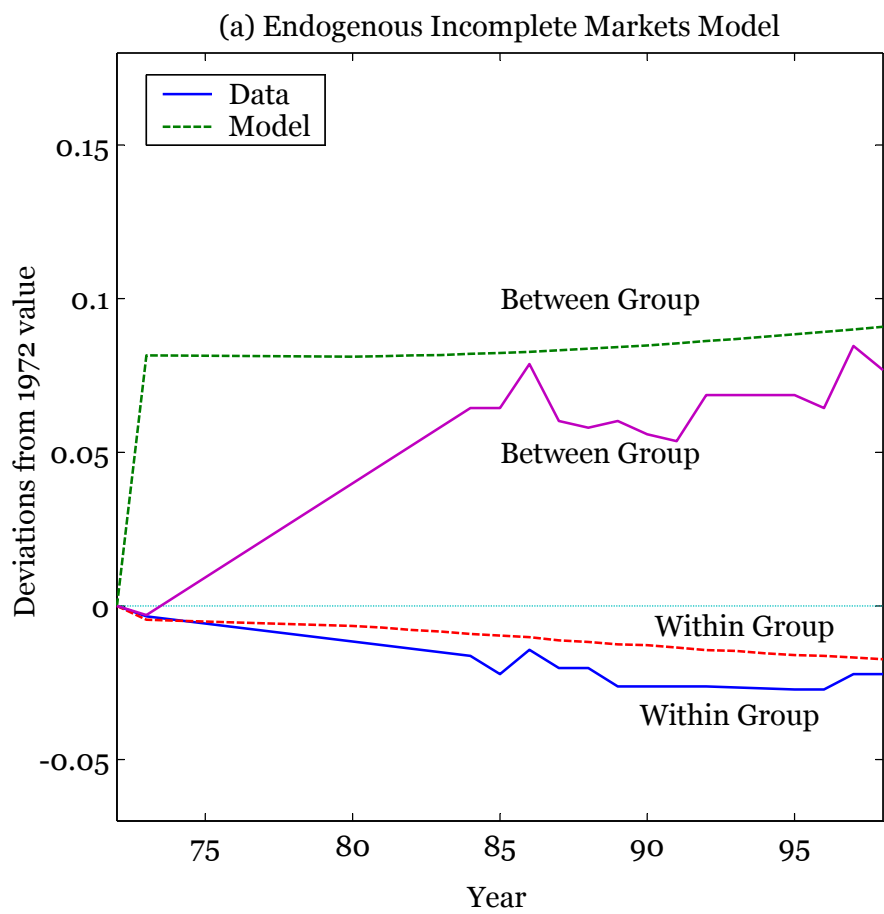


Figure 7. Income Inequality and Consumer Credit

