Does Income Inequality Lead to Consumption Inequality? Evidence and Theory[1]

DIRK KRUEGER

Goethe University Frankfurt, University of Pennsylvania, NBER, and CEPR

and

FABRIZIO PERRI

New York University, Federal Reserve Bank of Minneapolis, NBER, and CEPR

First version received August 2003; *final version accepted August* 2005 (*Eds.*)

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

1. INTRODUCTION

The sharp increase in earnings and income inequality for the U.S. in the last 25 years is a well-documented fact. Many authors have found that this trend is attributable both to increases in the dispersion of the permanent component of income and 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, considering the distribution of current income might not be sufficient. Since a significant fraction of variations of income is due to variations in its transitory component, current income may not be the appropriate measure of lifetime resources available to agents; thus, its distribution might not be a good measure of how economic welfare is allocated among households.[3] Moreover, the same change in current or permanent income inequality might have a very different impact on the welfare distribution, depending on the structure of credit and insurance markets available to agents for smoothing income fluctuations. For these reasons, several authors have moved beyond income and earnings as indicators of well-being and have studied the distribution of individual or household *consumption*. Contributors include Cutler and Katz (1991*a,b*), Johnson and Shipp (1991), Mayer and Jencks (1993), Slesnick (1993, 2001),

---

1. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.
2. See, for example, Gottschalk and Moffitt (1994), Gottschalk and Smeeding (1997), or Katz and Autor (1999).
3. Blundell 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.

Deaton and Paxson (1994), Dynarski and Gruber (1997), Blundell and Preston (1998), Johnson and Smeeding (1998), and Krueger and Perri (2004).

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

Second, we explore a theoretical explanation for these stylized facts. Our hypothesis is that an increase in the volatility of idiosyncratic labour income (which we identify as the increase in within-group inequality) has not only been an important factor in the increase in income inequality, but has also caused a change in the development of financial markets, allowing individual households to better insure against these (now bigger) idiosyncratic income fluctuations. We present a simple model with endogenous debt constraints (henceforth referred to as the debt constraint markets (DCM) model), building on earlier work by Kehoe and Levine (1993, 2001), Kocherlakota (1996), and Alvarez and Jermann (2000), that allows us to characterize *analytically*, the relationship between within-group income and consumption inequality. In the model, agents enter risk-sharing contracts, but at any point in time have the option to renege on their obligations at the cost of losing their assets and being excluded from future risk sharing. Our main result is that an increase in the volatility of income, keeping the persistence of the income process constant, *always* leads to a *smaller* increase in consumption inequality within the group that shares income risk. Intuitively, higher income volatility lowers the value of being excluded from credit markets and therefore reduces the incentives to default.[4] As a consequence, more risk sharing is possible, and the consumption distribution fans out less than the income distribution (and may even "fan in"). This model captures, in a simple and analytically tractable way, the idea that the structure of the credit markets in an economy is endogenous and may evolve in response to higher income volatility.[5]

Finally, we assess whether an extension of this simple model is *quantitatively* consistent with the trends for within- and between-group consumption inequality in the U.S. We develop a production economy with capital and a large number of agents that face a stochastic labour income process. We choose this income process to match the level and trend of income inequality in the U.S., both between and within different groups. The extent to which agents can borrow is derived endogenously, as in the simple model. We also evaluate the quantitative implications of a standard incomplete markets (henceforth SIM) model along the lines of Aiyagari (1994). We find that the DCM model slightly understates the increase in consumption inequality, and the SIM model somewhat overstates it, relative to the U.S. data. Finally, we investigate the reason for this difference.

The paper is organized as follows. Section 2 documents the main stylized facts. Section 3 presents the simple model, and Section 4 lays out the model used for quantitative analysis.

---

4. This mechanism would not work in an environment in which default does not involve any exclusion from credit markets. See, for example, Krueger and Uhlig (2005).
5. The endogenous response of credit markets to income risk has interesting policy implications. In Krueger and Perri (1999) we show that in the DCM model public insurance (unemployment insurance, progressive taxes, etc.) may crowd out private insurance, possibly more than one-for-one. Empirical studies by Cutler and Gruber (1996) and Albarran and Attanasio (2003) find a sizeable crowding-out effect of public insurance programmes.

Section 5 describes our experiment and parameter choices, and Section 6 presents and discusses the results. Section 7 concludes. Appendices A and B contain more details about the data and computational issues.

2. TRENDS IN INCOME AND CONSUMPTION INEQUALITY

This section documents how income and consumption inequality have evolved in the U.S. during the last 25 years. For this purpose, we use the CE Interview Survey, which is the only micro-level data-set for the U.S. that reports comprehensive measures of consumption expenditures and earnings for a repeated large cross section of households.[6]

2.1. *The Consumer Expenditure Survey*

The CE Interview Survey is a rotating panel of households who are selected to be representatives of the U.S. population. It started in 1960, but continuous data are available only from the first quarter of 1980 until the first quarter of 2004. For each quarter, the survey contains detailed information on quarterly consumption expenditures for all households interviewed during that quarter. After a first preliminary interview, each household is interviewed for a maximum of four consecutive times. In the second and fifth interviews, household members are asked questions about earnings, other sources of income, hours worked, and taxes paid for the past year.

2.2. *Income and consumption*

Our measure of income is meant to capture all sources of household revenues that are exogenous to the consumption and saving decisions of households (which are the object of our analysis). Therefore, we define income as after-tax labour earnings plus transfers (henceforth LEA+ income). We measure after-tax labour earnings as the sum of wages and salaries of all household members, plus a fixed fraction of self-employment farm and nonfarm income,[7] minus reported federal, state, and local taxes (net of refunds) and Social Security contributions. We then add reported government transfers (unemployment insurance, food stamps, and welfare).

Our measure of consumption is meant to capture the flow of consumption services that accrue to a household in a given period. For nondurable or small semidurable goods as well as services, expenditures are a good approximation for that flow. For large durable goods such as cars and houses, the relation between current expenditures and service flows is less direct. Thus, we impute service flows from the (value of the) stock of durables of a household. Our measure of service flows from housing is the rent paid by the households who, indeed, rent their homes and the self-reported hypothetical rent by households who own them. Our measure of quarterly service flows of cars is a fixed fraction (1/32) of the value of the stock of vehicles owned by the household. Since we do not have direct information on the value of the stock of cars, we follow the procedure used by Cutler and Katz (1991*a*) and use information from households who currently purchase vehicles (and for which we therefore observe the value of the purchase) to impute the value of the stock of vehicles for all households. Our benchmark measure of a household's consumption is then the sum of expenditures on nondurables, services, and small durables (such as household equipment), plus imputed services from housing and vehicles. Each expenditure component is deflated by expenditure-specific, quarter-specific consumer price indexes (CPIs). We label this benchmark measure ND+ consumption.[8] As we are interested in the

---

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

7. The exact fraction is 0·864 and is taken from Díaz-Giménez, Quadrini and Ríos-Rull (1997).

8. In Appendix A we provide more detailed description of our imputation and deflation procedures.

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

2.3. *Sample selection*

We want to select a sample of households for which we have reliable data for both labour income and consumption for the same time interval. For this reason, we only include households who are complete income respondents and have been interviewed five times.[9] For these households, income measured in the fifth interview and the sum of consumption reported in the second through the fifth interviews are our measures of their yearly income and consumption. For comparability with previous studies, we perform additional sample selections, such as excluding elderly and rural households and households whose reference person reports an implausibly low real wage.[10]

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

2.4. *Inequality trends*

Figure 1 displays the trend for four common measures of cross-sectional inequality of LEA+ income and ND+ consumption computed on the benchmark sample.[11] The figure confirms the fact that labour income inequality in the U.S. has increased significantly in the last quarter century: the Gini index has risen from 0·3 to around 0·37, and the variance of the logs displays an increase of more than 20%. The 90/10 ratio for income surges from 4·2 to over 6, suggesting a large divergence between the two tails of the income distribution over time. Finally, the 50/10 ratio displays an increase from 2·2 to 2·7, revealing that households in the bottom tail of the income distribution have lost ground relative to the median.[12]

The figure also presents our main empirical finding, namely that the increase in *consumption* inequality has been much less marked;[13] the increase has been from 0·23 to 0·26 for the Gini index and about 5% in the variance of logs. The 90/10 ratio has increased from 2·9 to around 3·4, suggesting a much more moderate fanning out of the consumption distribution.[14] Finally,
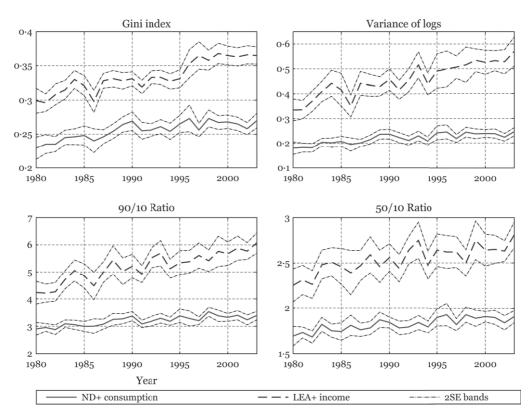
---

9. The CE classifies as incomplete income respondents in those households that report 0 income for all the major income categories, suggesting nonreliability of their earning figures (see also Nelson, 1994).

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

11. All measures are computed using CE population weights.

12. Increases of similar magnitude are found in other cross-sectional data-sets. Krueger and Perri (2004) compare the increase in wage inequality using CE data with that obtained by using PSID data (from Heathcote, Storesletten and Violante, 2004) and the increase measured by using the Current Population Survey data (from Katz and Autor, 1999) and find that, for the same sample selection, the magnitude of the increase is very similar.

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

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

Gini index · Variance of logs · 90/10 Ratio · 50/10 Ratio

— ND+ consumption   — — - LEA+ income   —·—·— 2SE bands

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

FIGURE 1

The evolution of income and consumption inequality in the U.S., 1980–2003

the 50/10 ratio increases only from about 1·7 to 1·9, implying that in terms of consumption (compared to income), households in the bottom part of the distribution have lost less ground relative to the median.[15]

Note that since our income definition includes government taxes and transfers, changes in government income redistribution policies cannot be responsible for the divergence between the two series. Although the evolution of consumption inequality has been studied less than the evolution of income inequality, some authors (Cutler and Katz, 1991*a*,*b*; Johnson and Shipp, 1991) have noted that the sharp increase in income inequality of the *early 1980's* has been accompanied by an increase in consumption inequality. Our measures also display an increase in consumption inequality in the early 1980's, but as noted by Slesnick (2001), it is less marked than the increase in income inequality; moreover, in the 1990's, income inequality has continued to rise (although at a slower pace) while consumption inequality has remained flat.[16] We now discuss the robustness of the trends just described to alternative definitions of consumption and to alternative sample selection choices.

15. 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 declined from 9·9% to 9·5%.

16. This last fact has also been reported by Federal Reserve chairman Alan Greenspan (1998) in his introductory remarks to a symposium on income inequality.

TABLE 1

Changes in inequality

(a) Alternative definitions

	Income	Consumption			
	LEA+	ND+	Food	ND	TCE
Change in variance of log	21·4%	5·3%	2·3%	3·3%	10·4%

(b) Alternative samples

	Benchmark		Quarterly	Income rural	Income low wages
Change in variance of log LEA+	5·3%		6·3%	4·5%	5·5%
Change in variance of log ND+	21·4%		22·0%	20·5%	27·5%

2.5. *Alternative definitions and samples*

Panel (a) of Table 1 reports the increase in consumption inequality, measured as the variance of logs, from the first two years of our sample (1980–1981) to the last two (2002–2003) obtained for alternative definitions of consumption. The first two columns, as reference, again report the increase in inequality for our benchmark measures of income and consumption.[17] The third and fourth columns report the increase in inequality in food consumption and nondurable consumption. For these measures, the increase in inequality is smaller than for our benchmark consumption (ND+) measure. Finally, the last column of Panel (a) reports the change in inequality for total consumption expenditures (TCE). For this definition, the increase in consumption inequality is larger, although still less than half of the increase in income inequality. One should keep in mind, however, that this consumption measure includes cash payments for homes and vehicles, and therefore contains a significant part of households' savings which biases measured consumption inequality towards measured income inequality. In addition, this measure, being based on expenditures on durables rather than on service flows from them, is affected by changes in the frequency of durables purchases over time. We therefore think of the latter statistic as an upper bound for the true change in consumption inequality rather than as the best estimate of it.

Panel (b) of Table 1 reports the increase in LEA+ income and ND+ consumption inequality as computed using different sample selections. The first column contains results for our benchmark sample. In the second column, we report inequality measures computed by using quarterly (as opposed to yearly) consumption measures. This sample contains households that are interviewed in the CE *less* than five times and therefore is significantly larger.[18] It has the disadvantage, however, that for any given household, income and consumption information do not refer to the same period. The third and fourth columns report measures for samples that include rural households and households with a reference person reporting a very low wage. We conclude that in all samples the increase in consumption inequality is quite similar and much smaller than the increase in income inequality.[19] Finally, it is worth mentioning that Attanasio, Battistin and Ichimura (2005) recently analysed the CE Diary Survey, which covers a different group of households and collects information on consumption of small items frequently purchased, such as food and personal care items. They find that for comparable categories in the Diary and

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

18. Our benchmark sample has an average of 6660 quarter/household data points per year, and the quarterly sample has an average of 11,300 quarter/household data points per year.

19. We also experimented with per-household (rather than per-adult-equivalent) income and consumption measures and different equivalence scales, with little effect on the results.

Interview Surveys, consumption inequality grows more in the Diary Survey. They then construct a measure of the variance of logs of nondurable consumption that uses information from both the Diary and the Interview Surveys and find that it increases by about 4·6% over the period 1986–2000. By contrast our ND measure, based solely on the Interview Survey, displays an increase of about 2·5% over the same years. These increases in inequality are different, but they are both significantly smaller than the increase in variance of log income (which over the same period was over 12%). Moreover, we conjecture that due to its limited consumption coverage, the impact of using the Diary Survey is likely to be even smaller if one focuses on a broader definition of consumption such as our benchmark ND+ consumption.

2.6. *Between- and within-group income and consumption inequality trends*

In this section, we decompose the change in income and consumption inequality into changes in between- and within-group inequality. Between-group inequality is attributable to fixed (and observable) characteristics of the household (*e.g.* education, experience, and sex). Although between-group inequality changes over time (returns to these characteristics can change over time, as in the case of the increase in the college premium), it is unlikely that households can insure against these changes. Therefore, increases in between-group inequality should translate into similar increases in between-group consumption inequality.

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

Following Katz and Autor (1999), for each labour income and consumption expenditure cross section (after controlling for age effects), we regress income and consumption on the following characteristics of the reference person and the spouse (if present): sex, race, years of education, experience, interaction terms between experience and education, dummies for managerial/professional occupation, and region of residence. These characteristics explain about 25% of the cross-sectional variation of income and consumption in 1980. We denote the cross-sectional variance explained by these characteristics as "between-group" inequality and the residual variance as "within-group" inequality. By construction the two variances sum to the total variance.

Figure 2 shows the evolution of between-group (Figure 2(a)) and within-group (Figure 2(b)) income and consumption inequality. For income, both the between- and within-group components display an increase. Figure 2(a) shows that for consumption, the between-group component displays an increase similar in magnitude to that of income.[20] Figure 2(b) reveals a very different picture for the within-group component: the increase in consumption inequality is an order of magnitude smaller than the increase in income inequality. Consequently, understanding the trends in Figure 2(b) is crucial for understanding the patterns of income and consumption inequality in the U.S. Note that increasing consumption inequality speaks against a model with *complete* insurance markets against income risk, as such a model would counterfactually predict constant consumption inequality over time. For this reason, in the remainder of the paper, we will focus on models in which insurance against income shocks and thus risk sharing is imperfect.

---

20. This finding is highly consistent with the results of Attanasio and Davis (1996), which suggest that changes in relative wages between education groups are fully reflected in consumption changes of these groups.
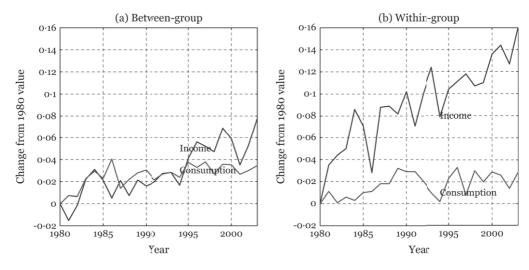
FIGURE 2

Changes in between- and within-group income and consumption inequality

We start by presenting a simple model in which risk sharing is endogenously limited by the limited enforcement of contracts and in which we can *analytically* characterize the relation between income and consumption inequality within a group of *ex ante* identical agents.

3. A SIMPLE MODEL

We analyse a pure exchange economy similar to that of Kehoe and Levine (2001). Time is discrete, $t = 0, \ldots, \infty$. There are two *ex ante* identical households $i = 1, 2$ and a single, nonstorable consumption good per period. Households obtain endowments of the consumption good from two sources: labour and capital "income". First, an agent receives endowments in the form of stochastic labour income. If one consumer has labour income $1 + \varepsilon$, the other has $1 - \varepsilon$. Let $s_t \in S = \{1, 2\}$ denote the consumer that has labour income $1 + \varepsilon$. We assume that $\{s_t\}_{t=0}^{\infty}$ is a sequence of i.i.d. random variables with $\pi(s_t = 1) = \pi(s_2 = 2) = \frac{1}{2}$. The parameter $\varepsilon \in [0, 1)$ measures the variability of the income process. In addition, two trees, one initially owned by each agent, each yield a constant endowment of capital income of r per period.

Let $s^t = (s_0, \ldots, s_t)$ denote an event history and $\pi(s^t)$ the time 0 probability of s^t. An allocation $c = (c^1, c^2)$ maps event histories s^t into consumption which agents value according to

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

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

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

the continuation utility of agent i from allocation c^i, from event history s^t onward, and denote by $e = (e^1, e^2)$ the autarkic allocation of consuming the *labour* 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 reneging on the risk-sharing arrangement obligations and bearing the associated costs, specified as exclusion from intertemporal trade and loss of any tree in their possession. This implies that any risk-sharing mechanism must yield allocations that deliver to each consumer, a continuation utility at least as high as the autarkic allocation, for all histories s^t. Formally, we impose the individual rationality constraints

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

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

$$c^1 + c^2 = 2(1+r) \tag{2}$$

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

3.1. *The constrained efficient consumption distribution*

We focus on symmetric allocations.[21] In order to analyse how constrained efficient consumption allocations vary with ε, we now solve for the continuation value of autarky, which is given by
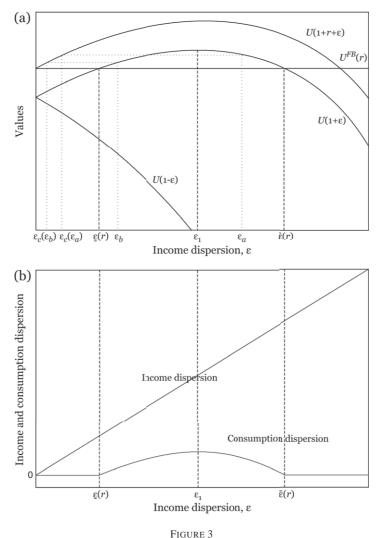
$$U(1+\varepsilon) = (1-\beta)u(1+\varepsilon) + \frac{\beta}{2}[u(1+\varepsilon) + u(1-\varepsilon)]$$

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

$U(1+\varepsilon)$ and $U(1-\varepsilon)$ denote the lifetime utility of always consuming $1+\varepsilon$ when rich and $1-\varepsilon$ when poor (*i.e.* the autarkic allocation), for the currently rich and poor agent, respectively. The continuation utility from autarky is a convex combination of utility obtained from consumption today, $(1-\beta)u(1+\varepsilon)$ or $(1-\beta)u(1-\varepsilon)$, and the expected utility from tomorrow onwards. The next lemma, whose proof is straightforward and hence omitted, states properties of $U(1+\varepsilon)$ as a function of income variability ε. Define $U^{FB}(r) = u(1+r)$ as the lifetime utility of the first best, perfect risk-sharing allocation in which consumption of both agents is constant at $1+r$.

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

  (i) *Either $U(1+\varepsilon_1) \leq U^{FB}(r)$ or*
  (ii) *$U(1+\varepsilon_1) > U^{FB}(r)$. In this case, there exist $\underline{\varepsilon}(r) < 1$ such that $U(1+\underline{\varepsilon}(r)) = U^{FB}(r)$. The number $\underline{\varepsilon}(r)$ is strictly increasing in r. Furthermore, either there exists a number $\bar{\varepsilon}(r) \in (\underline{\varepsilon}(r), 1)$ such that $U(1+\bar{\varepsilon}(r)) = U^{FB}(r)$ and the number $\bar{\varepsilon}(r)$ is strictly decreasing in r, or $U(2) > U^{FB}(r)$, in which case we take $\bar{\varepsilon}(r) = 1$. Consequently, $U(1+\varepsilon) > U^{FB}(r)$ for $\varepsilon \in (\underline{\varepsilon}(r), \bar{\varepsilon}(r))$, and thus for these ε complete risk sharing is worse than autarky for the currently rich agent.*

21. Consider two histories $s^t, \tilde{s}^t$ that satisfy, for all $\tau \leq t$, $s_\tau = 1$ if and only if $\tilde{s}_\tau = 2$. A consumption allocation is symmetric if $c_t^1(s^t) = c_t^2(\tilde{s}^t)$ for all such histories $s^t, \tilde{s}^t$.

FIGURE 3

Characterizing the link between income and consumption dispersion

The nonmonotonicity of $U(1+\varepsilon)$, shown in Figure 3(a), stems from two opposing effects. For small ε the direct effect of higher consumption today outweighs the higher risk faced by the agent from tomorrow onwards: $U(1+\varepsilon)$ increases with ε. As ε becomes larger and future consumption more risky, $U(1+\varepsilon)$ declines with ε as the risk effect dominates. 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, since an increase in ε reduces consumption today for this agent and makes it more risky from tomorrow onwards. In addition, since with complete risk sharing agents share capital income r, whereas in autarky they only consume their labour income, higher r makes risk sharing more attractive relative to autarky. Thus, the region $[\underline{\varepsilon}(r), \bar{\varepsilon}(r)]$ shrinks with increasing r.

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

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

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

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

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

Note that if $U^{FB}(r) \geq U(1 + \varepsilon)$, the smallest solution to equation (3) is $\varepsilon_c(\varepsilon; r) = 0$ and the constrained efficient allocation implies full risk sharing. Also note that unless $r = 0$, autarky is never constrained efficient, since the equation

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

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

3.2. *Income variability and consumption inequality*

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

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

(i) *If $U(1 + \varepsilon_1) \leq U^{FB}(r)$, then perfect consumption insurance is feasible for all $\varepsilon \in [0, 1]$ and a change in ε has no effect on consumption inequality.*

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

The proof of this proposition follows immediately from Proposition 2 and the properties of $U(1 + \varepsilon)$ stated in Lemma 1, apart from the very last part. The fact that consumption inequality always increases less than income inequality is obvious for the regions $\varepsilon \in [0, \underline{\varepsilon}(r))$, $\varepsilon \in [\bar{\varepsilon}(r), 1)$, and $\varepsilon \in [\varepsilon_1, \bar{\varepsilon}(r))$, since in these regions consumption inequality does not change or is even declining in income volatility. For the region $\varepsilon \in [\varepsilon_1, \bar{\varepsilon}(r))$ we have

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

and thus by the implicit function theorem

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

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

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

As shown in Figure 3(a), we see that for $\varepsilon \in [0, \underline{\varepsilon}(r))$ and $\varepsilon \in [\bar{\varepsilon}(r), 1)$, $U^{FB}(r) > U(1+\varepsilon)$, and thus the first best allocation can be implemented. In this case, as shown in Figure 3(b), consumption inequality does not vary with income inequality.

Suppose now that $\varepsilon_1 < \varepsilon < \bar{\varepsilon}(r)$. For example, consider the point $\varepsilon = \varepsilon_a$ on the x-axis. From Proposition 2, the constrained efficient consumption allocation is given by the smallest solution to $U(1+r+\varepsilon_c(\varepsilon_a)) = U(1+\varepsilon_a)$. Figure 3(a) displays the solution $\varepsilon_c(\varepsilon_a)$. In this allocation, which involves partial risk sharing as $\varepsilon_c(\varepsilon_a) < \varepsilon_a$, the household with high income receives a continuation utility equal to the value of autarky, whereas the low-income household receives $U(1+r-\varepsilon_c(\varepsilon_a))$, strictly higher than its value of autarky ($U(1-\varepsilon_a)$). In this range a marginal increase in income inequality reduces the value of autarky for the high-income agent, and less current consumption is required to make it not default ($\varepsilon_c(\varepsilon_a)$ moves to the left). This reduces consumption dispersion in the economy, as shown in Figure 3(b).

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

To summarize, in this environment with limited commitment, an increase of income dispersion always leads to a smaller increase in consumption dispersion as long as there is some capital income. It may even lead to a reduction in consumption dispersion. The intuition behind these results is that an increase in income inequality, by making exclusion from future risk sharing more costly, renders the individual rationality constraint less binding. It thereby allows individuals to share risk to a larger extent and thus reduces fluctuations in their consumption profiles. It is crucial for this result that income shocks are not perfectly permanent (although they may be highly persistent), because the fear of being poor again in the future is what makes a currently rich agent transfer resources to his currently poor brethren.[22]

3.3. *Capital income and the extent of risk sharing*

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

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

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

22. It is straightforward to generalize our results to a serially correlated endowment process. An increase in persistence leads to an increase in consumption dispersion in the constrained efficient consumption distribution. This increase is strict if initially there is some, but not complete, risk sharing. For a proof of this result, see Kehoe and Levine (2001). The intuition is again simple: the value of autarky for the agent with high current income increases (as the agent is more likely to have high income in the future with higher persistence), which makes the individual rationality constraint more stringent and leads to fewer transfers to the poor agent being sustainable.

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

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

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

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

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

$\parallel$

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

4. THE MODEL WITH A LARGE NUMBER OF AGENTS

4.1. *The environment*

A single good being produced in a given period can be used for consumption or investment in the physical capital stock K. The representative firm produces output according to a Cobb–Douglas production technology. The aggregate resource constraint reads as

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

where L_t denotes labour input, K_t the aggregate capital stock, C_t aggregate consumption, A is a technology parameter, and δ the depreciation rate.

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

A household in group i has a stochastic labour endowment process $\{\alpha_{it} y_t\}$, where α_{it} is the deterministic, group-specific, and possibly time-varying mean labour endowment, and the idiosyncratic component $\{y_t\}$ follows a Markov process with finite support Y_t, a set with cardinality N. Since labour income will be the product of individual labour endowment and an economy-wide wage per efficiency unit of labour, we use the terms *labour endowment* and *labour income* interchangeably. Let $\pi_t(y'|y)$ denote the transition probabilities of the Markov chain, assumed to be identical for all agents. The set Y_t and the matrix π_t are indexed by t, since we allow the idiosyncratic part of the income process to change over time. Furthermore, we assume that a law of large numbers applies, so that the fraction of households facing shock y' tomorrow with shock y today in the population equals $\pi_t(y'|y)$. Finally, assume that $\pi_0(y'|y)$ has a unique invariant measure $\Pi_0(\cdot)$. Let y_t denote the current-period labour endowment and $y^t = (y_0, \ldots, y_t)$, the history of endowment shocks; also, $\pi(y^t|y_0) = \pi_{t-1}(y_t|y_{t-1}) \cdots \pi_0(y_1|y_0)$. The notation

$y^s | y^t$ means that y^s is a possible continuation of history y^t. We further assume that at date 0 the measure over current labour endowments is given by $\Pi_0(\cdot)$. At date 0, households are distinguished by their group i, their initial asset holdings a_0, and by their initial labour endowment shock y_0. Let Φ_0 be the initial distribution over types (i, a_0, y_0). Total labour supply is given by

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

Finally, households' preferences are exactly as described in the simple model of the previous section. We now describe the market structure of the economies whose quantitative properties we will contrast with the stylized empirical facts established in Section 2.

4.2. *Market structures*

4.2.1. Debt constraint markets. A household of type (i, a_0, y_0) trades Arrow securities subject to pre-specified credit lines $A_t^i(y^t, y_{t+1})$ that are contingent on observable labour endowment histories and a households' group. Its exact form is specified below. The prices for Arrow securities are denoted by $q_t(y^t, y_{t+1})$ and depend on a household's labour endowment history and time, to reflect deterministic changes in the income process and hence in the magnitude of labour endowments $\alpha_{it} y_t$.

Consider the problem of a household of type (i, a_0, y_0). It chooses consumption $\{c_t(a_0, y^t)\}$, and one-period Arrow securities $\{a_{t+1}(a_0, y^t, y_{t+1})\}$ whose pay-off is conditional on his own endowment realization y_{t+1} tomorrow, to maximize, for given (a_0, y_0) and wage rates w_t

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

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

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

Following Alvarez and Jermann (2000), we specify the constraints $A_t^i(y^t, y_{t+1})$ as "solvency constraints" that are not too tight. As before, let $U_t^{Aut}(i, y_t)$ denote the continuation utility from autarky, given current labour endowment $\alpha_{it} y_t$. Define the continuation utility $V_t(i, a, y^t)$ of a household of type i with history y^t and current asset holdings a at time t as

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

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

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

That is, the constraints are such that a type i household, having borrowed up to maximum, $a_{t+1}(a, y^t, y_{t+1}) = A_{t+1}^i(y^t, y_{t+1})$, is indifferent between repaying his or her debt and defaulting, with the default consequence being specified as limited future access to financial markets. The defaulting household starts with neither assets nor liabilities but, in contrast to the simple model,

is allowed to save (but not borrow), at a state-uncontingent interest rate r_d (a parameter of the model). The value of autarky is given by

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

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

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

subject to $b_t(a_0, y^{t-1}) = 0$. Note that for $r_d = -1$, the household optimally never saves after default; in this case, the value of autarky coincides with that in the simple model above.

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

1. *Household optimization.*   Given prices, household allocations maximize (8) subject to (9) and (10), and the solvency constraints are not "too tight" in the sense of (11).
2. *Firm optimization.*

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

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

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

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

and the asset market-clearing condition holds, with

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

As stationary equilibrium, we define an equilibrium in which $\{r_t, w_t\}$ and the cross-sectional asset and consumption distributions are constant over time. Three comments on the equilibrium definition are in order. First, in the asset market-clearing condition we divide the R.H.S. by the interest rate since the Arrow securities are state-contingent *zero coupon* bonds. Second, no arbitrage implies that

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

because households have to be indifferent between saving with risk-free capital or reconstructing a risk-free asset with the full set of Arrow securities. Third, with the full set of Arrow securities, risk-free capital is a redundant asset for households, so we have abstained from introducing purchases of capital in the household problem explicitly. Physical capital in our economy is not important as an additional asset, but is important because it provides the economy with an

asset in positive net supply and therefore generates a positive wealth-to-income ratio. The simple model above demonstrated that the more abundant is capital (income), the better the extent of consumption insurance that is achievable in the DCM model.

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

4.2.2. Incomplete markets. We will first consider a SIM model, as in Huggett (1993) or Aiyagari (1994). Let q_t^{in} denote the price, at period t, of an uncontingent claim to one unit of the consumption good in period $t+1$. The sequential budget constraints the agent faces are

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

and the borrowing constraints become

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

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

We will also consider, as in Zhang (1997), a single-bond incomplete markets economy with *endogenous* borrowing constraint. In this version of the model (from now on ZIM), households can sell the uncontingent bond, subject to a constraint

$$a_{t+1}(a_0, y^t) \geq A_t^i(y^t),$$

where $A_t^i(y^t)$ is the smallest number satisfying

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

That is, in the ZIM economy, households can short-sell the bond today up to the point where they are at most indifferent between repaying and defaulting tomorrow in all possible states. As in the DCM model, a change in the income process, therefore, changes the borrowing constraint.[23]

5. THE QUANTITATIVE EXERCISE

Our quantitative exercise involves the following steps: (i) We choose parameter values so that the stationary equilibrium in all economies matches key observations of the U.S. economy in the 1980's. This applies, in particular, to the deterministic and stochastic part of the labour productivity process. (ii) We introduce changes in the dispersion of the labour productivity process to mimic the increase in income inequality observed in the U.S. data. It is assumed that this change

23. In order to compute calibrated versions of all models, we reformulate them recursively and then solve them numerically. The computation of equilibrium in the DCM model and the Zhang model requires us to solve for both prices and borrowing constraints simultaneously. See Appendix B for details.

is unforeseen by households, but that all future changes in the process are fully learned once the first change occurs. The change in the labour productivity process for a finite number of periods induces a transition from the initial to final stationary equilibrium associated with the process that prevails, once the change in that process is completed. (iii) The models endogenously generate consumption distributions along the transition from the old to the new steady state. We compute measures of consumption inequality in the models and compare them to the facts in Section 2. We now proceed with the first step.

5.1. *Calibration*

5.1.1. Income process.   We take the length of a model period to be 1 year. A household's labour income $e_{it} = w_t \alpha_{it} y_t$ consists of a common wage rate; a group-specific, time-dependent deterministic part, α_{it}; and an idiosyncratic stochastic component, y_t. In the empirical section, we decomposed household income and consumption data into a group-specific component and an idiosyncratic component. Our calibration strategy follows the same approach. The logarithm of labour income is given by

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

and thus $\sigma_{et}^2 = \sigma_{\alpha t}^2 + \sigma_{yt}^2$, where $\sigma_{et}^2 = \text{Var}[\ln(e_{it})]$ is the cross-sectional variance of log-labour income and $\sigma_{\alpha t}^2 = \text{Var}[\ln(\alpha_{it})]$, $\sigma_{yt}^2 = \text{Var}[\ln(y_t)]$. We identify $\{\sigma_{\alpha t}^2, \sigma_{yt}^2\}_{t=1980}^{2003}$ with the between- and within-group income variances plotted in Figure 2. We first HP-filter (with smoothing parameter of 400) the time series $\{\sigma_{\alpha t}^2, \sigma_{yt}^2\}_{t=1980}^{2003}$, in order to remove high-frequency variation. We then choose parameters governing the model income process so that (a) in the initial stationary equilibrium, between- and within-group income variances of the model match the data for the early 1980's and (b) along the transition, trends in between- and within-group income variances are reproduced by the model.

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

Since $\{\alpha_{1t}, \alpha_{2t}\}_{t=1980}^{\infty}$ is a deterministic sequence, in both models the increase in between-group income inequality translates fully into an increase in between-group consumption inequality. Furthermore, by construction, the change in between-group inequality does not affect the quantitative importance of the risk-sharing mechanism at work for within-group stochastic income variability described in Section 3. We choose this specification for two reasons. First, Attanasio and Davis (1996) show that between-group consumption insurance fails, and they conclude that "the evidence is highly favorable to an extreme alternative hypothesis under which relative consumption growth equals relative wage growth" (p. 1247). With our specification of average group income, changes in this income component are not (self-)insurable, consistent with their findings. Second, we can quantify exactly to what extent the (self-)insurance mechanisms of both models can offset the increase in idiosyncratic income volatility. The potency of these mechanisms depends on the properties of the idiosyncratic income process discussed next.

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

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

Here, ε_t, η_t are independent, serially uncorrelated, and normally distributed random variables with 0 mean and variances $\sigma_{\varepsilon t}^2, \sigma_{\eta t}^2$, respectively. We allow these variances to change over time, whereas we treat ρ as time invariant. Note, however, that if $\sigma_{\varepsilon t}^2$ and $\sigma_{\eta t}^2$ increase at different rates over time, the implied persistence of the idiosyncratic component of income, $\ln(y_t)$, changes. Thus, the process we use is flexible enough to allow for time-varying income persistence even if ρ is constant over time. As a benchmark value for ρ we choose $\rho = 0.9989$, the value Storesletten *et al.* (2004) find when estimating the process in (14). As sensitivity analysis we also report results for lower values of ρ.

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

$$\text{Cov}(y_t, y_{t+1}) = E\left((z_t + \varepsilon_t)(\rho z_t + \eta_{t+1} + \varepsilon_{t+1})\right) = \rho \sigma_{zt}^2 \tag{15}$$
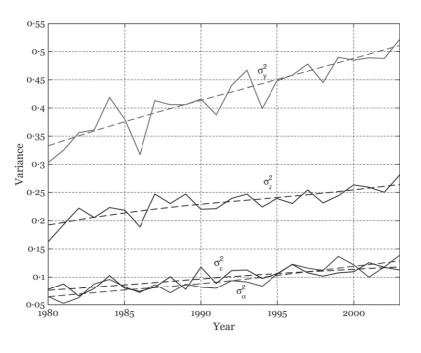
$$\sigma_{yt}^2 = \sigma_{zt}^2 + \sigma_{\varepsilon t}^2. \tag{16}$$

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

Equipped with empirical time series for $\{\sigma_{\varepsilon t}^2, \sigma_{zt}^2\}_{t=1980}^{2003}$, we now specify a discretized version of the process in (14). The purely transitory component takes one of the two values $\varepsilon_{1t} = -\sigma_{\varepsilon t}$ and $\varepsilon_{2t} = \sigma_{\varepsilon t}$ with equal probability. Consequently, the transitory shock fed into the model has a variance exactly as big as identified in the data for all time periods. For the persistent part of the process, we use a seven-state Markov chain with time-varying states such that the variance of this process, in each period, equals σ_{zt}^2 as identified from the data.[24] Note that after 24 model periods (2003 in real time), the change in the dispersion of the income process is completed. However, due to the endogenous wealth dynamics in both models, it may take substantially longer than 24 years for both economies to reach the new stationary consumption distribution.

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

5.1.4. Exogenous borrowing limit and autarkic interest rate.
As a benchmark borrowing limit in the SIM model, we set $\bar{B} = 1$. Note that we normalize endowment in such a way

---

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

Note: The solid lines represent cross-sectional variances estimated in the data. The dashed
lines are the HP-filered series which are also exactly matched by the cross-sectional
variances in the model.

FIGURE 4

Income variances: data and models.

that this borrowing limit corresponds to a generous one times the average annual income for each
group i. As a benchmark in the DCM model, we allow households to save at the initial equilib-
rium interest rate, $r_d = r$ in autarky. We then report how sensitive our quantitative results are to
the choice of the borrowing limit and the autarkic interest rate.

5.1.5. Technology and preference parameters.   We assume that the period utility is log-
arithmic, $u(c) = \log(c)$. We then choose the technology parameters (A, α, δ) so that in both
models the initial steady state has a wage rate of 1 (a normalization), a capital income share of
30%, and a return on physical capital of 4% per annum, as suggested in McGrattan and Prescott
(2003).[25] The time discount factor β is then set in both models such that the initial steady state in
both models has a capital (wealth)-to-output ratio of 2·6. This value is equal to the average wealth
(including financial wealth and housing wealth) for CE households in the benchmark sample in
1980–1981, and it is also close to the value estimated in NIPA data by Fernández-Villaverde and
Krueger (2002).

The resulting technology parameters are $\delta = 7\cdot54\%$ and $A = 0\cdot9637$ in both models, since
the production side of the economy is identical in both. The appropriate choice of the time dis-
count factor β then ensures that households indeed have the incentive to save exactly the amount
required to make the capital-output ratio equal to 2·6. This requires a $\beta = 0\cdot959$ for the DCM

25. Capital is a risk-free asset in the model, so the calibration target for its return should be the empirical return to
the U.S. aggregate capital stock, adjusted by a risk premium and economic growth. Choices in the range of 1–7% seem
defendable and are used in the literature, which places our target in the middle of that range.

model and a $\beta = 0.954$ for the SIM model. In our sensitivity analyses we always re-calibrate β to maintain the same equilibrium capital-output ratio in the initial steady state.

6. QUANTITATIVE RESULTS

6.1. *Benchmark calibration*

Figure 5 summarizes our main quantitative results.[26] Figure 5(a) displays the change in the *between-group* variance of log consumption implied by the DCM and SIM models and the data. Figure 5(b) does the same for the *within-group* variance of log consumption, our main focus of interest.[27]

Since the change in between-group income inequality is modelled by a deterministic process, by construction there is no (self-)insurance possible against the increase in between-group income variability. Thus, in the long run all of the increase in between-group income inequality is reflected in a one-for-one increase in between-group consumption inequality in both models. As shown in Figure 2(a), the fact that in the data the increase in between-group income inequality is similar to the increase in between-group consumption inequality therefore implies that the increase in between-group consumption inequality predicted by both models is similar to the one observed in the data.

The crucial quantitative question is how well both models can capture the trend in within-group consumption inequality. Figure 5(b) answers this question. It shows that for our benchmark parameterization, the DCM model understates and the SIM model overstates the increase in within-group consumption inequality, compared to the data. The data display an increase in the variance of about 2·0%, the DCM model shows an increase of only 0·5%, whereas the SIM model predicts an increase of 4·5%. Note that the increase in within-group income variance from the data is 11%, so both models generate an increase in within-group consumption inequality substantially lower than the increase in income inequality, in line with the data.

The quantitative difference in the change of within-group consumption inequality in the two models is due to the differential response of financial markets to increased income volatility. In the SIM model, the increase in the variance of income leads to higher precautionary savings. In addition, households facing larger shocks become more hesitant to borrow, but their ability to borrow remains unchanged. Thus, outstanding unsecured consumer credit (as a fraction of output) *declines* by 0·6%, equilibrium asset holdings and so the physical capital stock increase by 2·1%, and the real return on capital declines by 17 basis points. In contrast, in the DCM model credit limits expand for the purchase of all Arrow securities, and households can and do borrow more, at least against the contingency of having higher income tomorrow. This is reflected in an increase in outstanding unsecured consumer credit (again, as a fraction of output) by 2·1%. But keeping consumption as smooth as it was before the change in the income process may require a bigger expansion in state-contingent borrowing than is feasible with the wider constraints, so that some of the increase in income volatility is reflected in consumption. Within-group consumption variance increases, albeit very mildly. The expansion of borrowing in the DCM model is met by an increase in purchases of Arrow securities, as households have a stronger need to save for the contingency of being income poor tomorrow. On net, aggregate savings and thus the capital stock increases by 1·8%, and the return on capital falls by 15 basis points. The increase in the capital stock and the decline in the interest rate are smaller than those in the SIM model because

26. Since in our calibration, we used filtered income variances, we also filter the data consumption inequality statistics.

27. Since the sum of the between- and within-group variance of consumption equals the overall variance, the change in the overall variance from the data and both models can be readily deduced from Figure 5(a) and (b).
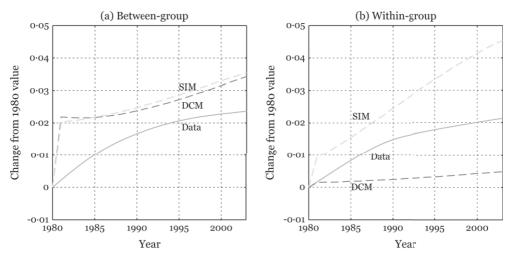
FIGURE 5

Changes in between- and within-group consumption inequality: data and models

the increase in asset accumulation in the DCM model is partially offset by a higher demand for credit, an effect that is absent in the SIM model. A precise quantitative evaluation of both models with respect to the CE data along the credit dimension is not possible because data on unsecured consumer credit are not available in our CE sample. At least qualitatively, however, the DCM model seems more consistent with recent developments in the U.S. markets for uncollateralized credit (see also our Figure 7 in the Conclusions).

6.2. *Sensitivity analysis*

The goal of this section is to better understand the quantitative results presented above. We do so by first illustrating the impact of changing some key parameters and then by evaluating the ZIM model, which helps to assess the importance of endogenous borrowing constraints vs. full spanning of idiosyncratic uncertainty in financial markets.

     6.2.1. Borrowing constraints.   The first three rows of Table 2 report results for different values of the borrowing constraint for the SIM model. A borrowing constraint of $B = 2$ implies that a household can take out (noncollateralized) loans up to twice his or her annual average labour income. The last two lines document how our results for the DCM model change if we reduce the net real interest rate at which households can save in autarky to zero, making the default option less attractive. The second column (labelled ΔVar) reports the change, between 1980 and 2003, in the *within-group* consumption variance, the third (labelled ΔCredit) the change in the outstanding credit-to-GDP ratio, and the last (labelled Δr) the change in the real interest rate.

     The most important observation from Table 2 is that more generous credit lines, if they *do not change over time*, do not help to keep the increase in within-group consumption inequality low in the SIM model. With a wealth-to-output ratio of 2·6 in the model, most households are far to the right from the borrowing constraint in the wealth distribution. In addition, with highly persistent income shocks, they are very reluctant to obtain credit in response to bad income shocks. In fact, as these shocks become bigger over time, households become even more timid

TABLE 2

Change in borrowing constraints

Economy	Δ Var	ΔCredit (%)	Δr (%)
SIM, $B=0$	0·0456	0	−0·20
SIM, $B=1$	0·0454	−0·62	−0·17
SIM, $B=2$	0·0463	−0·88	−0·15
DCM, $r_d=4\%$	0·0049	2·05	−0·15
DCM, $r_d=0\%$	0·0012	12·34	−0·03

in using credit and try to stay away from high-debt positions. We observe from Table 2 that the credit-to-gross domestic product (GDP) ratio declines with the increase in income volatility, the more so for looser the borrowing constraint and thus the higher the debt position of those at the constraint.

For the DCM model, reducing the interest rate at which people can save after default leads to a worse autarkic option, and thus to more borrowing being enforceable. In fact, with an $r_d = -2\%$ (*i.e.* households can save only at a negative real interest rate of 2%), the default option is so unattractive that perfect risk sharing is possible, and the predictions of the DCM model collapse to those of the standard complete markets model. Reducing the autarkic savings interest rate gives rise not only to better risk allocation in the initial steady state, but also to a stronger relaxation of borrowing constraints over time. As a result, the use of credit expands substantially, and there is almost no increase in within-group consumption inequality over time, reducing the ability of the DCM model to match the data relative to the benchmark parameterization. Also note that with $r_d = 0\%$, the increase in credit demand almost matches the increased savings demand, so that the real interest rate and the capital stock remain virtually unchanged as income variability increases.

6.2.2. Persistence of income shocks.   At least since Friedman (1957), it is well understood that very persistent income shocks are harder to self-insure against than transitory ones. Our idiosyncratic income process is the sum of a highly persistent and a purely transitory component. Although authors who estimate this process from wage data consistently find the persistence parameter ρ close to 1, some disagreement exists about its exact magnitude. We have repeated our analysis for various other choices of ρ and report, in Table 3, results for $\rho = 0.8$, the value estimated by Guvenen (2005), which is the lowest estimate for ρ we are aware of, for the exact income process we use.[28] In each case we re-calibrate the income process and time discount factor such that the cross-sectional income dispersion of the process fed into the model matches the empirical facts from Figure 4, and the initial real return on capital remains at 4%. This procedure keeps the volatility of the income process unchanged, but reduces its persistence.

Comparing the results of Table 3 with those in Table 2, we see that lower persistence of income shocks in the SIM model indeed reduces the rise in within-group consumption inequality. Whereas for $\rho = 0.9989$, this rise was about 4·5%, with $\rho = 0.8$ it drops to about 3·4% (compared with 2% in the data). Again, the results are fairly independent of the borrowing constraint. With lower persistence, households in the SIM model find it easier to self-insure by accumulating capital and using it to smooth income shocks. The increase in the capital stock (and corresponding decline in the real return) is more pronounced for $\rho = 0.8$ than for $\rho = 0.9989$.

---

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

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

TABLE 3

Change in persistence: $\rho = 0 \cdot 8$

Economy	ΔVar	ΔCredit (%)	Δr (%)
SIM, $B = 0$	0·0355	0	−0·28
SIM, $B = 1$	0·0334	−0·01	−0·22
SIM, $B = 2$	0·0341	0·0	−0·19
DCM, $r_d = 4\%$	0·0025	1·33	−0·10
DCM, $r_d = 2 \cdot 5\%$	0·0013	3·93	−0·04

Also, households now are not as timid as before to use credit to smooth income shocks; instead of a decline of credit as a fraction of GDP, we now observe this statistic to be virtually unchanged.

In the DCM model, a lower ρ reduces the value of autarky for households with currently high income whose constraints are binding as it is now less likely that they will remain income rich. Thus, insurance possibilities, *ceteris paribus*, are better, and the DCM model implies a smaller increase in within-group consumption inequality. We conclude that setting persistence of the income shocks to the lower bound from the empirical literature leads to a lesser increase of within-group consumption inequality both in the SIM and in the DCM model, but it does not change our basic finding that the SIM slightly overstates the increase in within-group consumption inequality and the DCM slightly understates it.[29]

6.2.3. Market completeness or endogenous borrowing constraints?   In this section we would like to examine the exact cause of the different quantitative predictions of the two models discussed so far, with respect to the increase in within-group consumption inequality. To evaluate the role of differences in the span of assets vs. differences in the structure of borrowing constraints, we repeat the quantitative experiment performed above with the ZIM model (as described in Section 4.2.2), which has a single financial asset (as the SIM model) and endogenous borrowing constraints (as the DCM model).

Table 4 displays the results. The first column of the table reports the change in the *within-group* consumption variance between 1980 and 2003 and shows that the ZIM model predicts an increase that is very similar to the one predicted by the SIM model, and hence larger than the one predicted by the DCM model. This suggests that the degree of spanning (which is the only difference between the ZIM and DCM models) is the key determinant of the differences in the quantitative predictions between the SIM and DCM models.

Note however that, differently from the SIM model, the ZIM model reproduces the empirically observed increase in the use of credit. To better understand this, Table 4 also reports the change between the initial steady state and the first period of the transition (the column labelled ΔVar0). Households, on impact, respond very differently to the unexpected change in the income process in the ZIM economy, compared to the SIM economy. In the ZIM economy, credit lines expand on impact and households make use of these expanded credit lines. As a result, consumption inequality initially falls. But households that have borrowed more now face higher debt levels and income processes with more extreme realizations. Over time, the debt needs to be

29. An issue worth mentioning is that in order to precisely approximate the persistent part of the income process with $\rho = 0 \cdot 9989$ requires a Markov chain with a high number of states (see Tauchen and Hussey, 1991). Due to computational time constraints, our Markov chain uses only seven states. We conducted experiments with nine instead of seven states; doing so raises the persistence of the income process. By using nine states we found a higher predicted increase in within-group consumption inequality for *both* models. The changes, however, were quantitatively small (the increase in consumption inequality rises by around 0·4% in both models). Therefore, this experiment again does not overturn our main finding that the DCM slightly understates the increase in consumption inequality and the SIM slightly overstates it.

TABLE 4

A hybrid model

Economy	Δ Var	Δ Var0	Δ Credit (%)	Δr (%)
SIM	0·0454	0·0095	−0·62	−0·17
ZIM	0·0469	−0·0112	1·7	−0·16
DCM	0·0049	0·0015	2·05	−0·15

serviced, consumption has to respond, and consumption inequality eventually goes up (about five periods after impact). After 25 years the change in consumption inequality is the same as in the benchmark SIM model.[30] Thus, borrowing constraints that are relaxed over time help households in the short run to better smoothen more volatile income fluctuations, but in the long run their debt burden catches up with them.
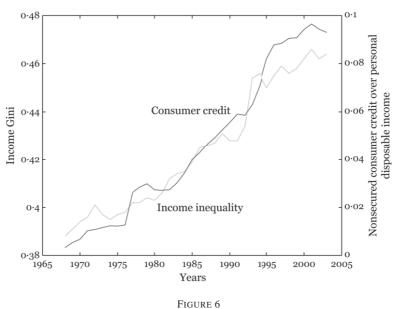
Why does the same logic not apply in the DCM model? Here is where the second crucial feature of that model, state-contingent borrowing, comes in. In the DCM model, households with currently high income enter the period with high debt, exactly because they have borrowed against the contingency of being income rich. Thus, the high debt is not such a high burden, and households can and do make use of the higher credit lines without the consequence of particularly low consumption in the future. In sharp contrast, in the SIM model, income-poor people start the period with outstanding credit (because of the strong positive correlation of income and the fact that they went into debt because of bad income realizations). Thus, in that model, highly indebted households eventually have to accept (persistently) low consumption. Put another way, in the SIM model assets and income are highly positively correlated, whereas in the DCM model, they are negatively correlated. Therefore, in one model being in debt has fairly persistent negative consequences for consumption; in the other model it does not.

From this, we conclude that the span of assets plays a crucial role for determining how much consumption inequality respond to increase in income inequality. Since the DCM model, which features complete spanning, predicts too little increase in income inequality and the SIM and ZIM models with their single bond predict too much increase, we conjecture that a model with an intermediate set of assets may be most successful in quantitatively matching the data.

7. CONCLUSIONS

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

---

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

FIGURE 6

Income inequality and consumer credit

    The mechanism through which agents in the DCM model of the last section keep their consumption profiles stable in the light of more volatile income is an expansion in the use of noncollateralized, state-contingent credit. Did this expansion take place in the data? One simple (but of course only partial) measure of such credit used by the U.S. consumers is the ratio of aggregate unsecured consumer credit to disposable income. Consumer credit is mostly uncollateralized and, since it can be defaulted on, has some element of state contingency. In Figure 6, we plot this ratio from the U.S. *data* for the last 40 years, as well as the Gini coefficient for the U.S. household income.[31] Despite some idiosyncratic cyclical variations, the two series display a remarkably similar trend. Combining this figure with our consumption inequality observations may suggest that consumers could, and in fact did, make stronger use of credit markets exactly when they needed to (starting in the mid-1970's), in order to insulate consumption from bigger income fluctuations.

    Both the models we evaluated have their shortcomings. The DCM model assumes the existence of a complete set of Arrow securities which stand in for extremely well-developed direct insurance markets. Although we see a host of assets traded in financial markets and informal insurance mechanisms working on the level of the extended family, some may question the empirical realism of this assumption. Also, full spanning in asset markets is the key reason why the DCM model produces too small a response of within-group consumption inequality to increased income volatility. On the other hand, the SIM model seems to allow too little explicit insurance against income shocks, over and above simple self-insurance.[32] In this model, it is the absence of spanning that generates an increase in within-group consumption inequality that is too high

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

    32. In their work, Blundell, Pistaferri and Preston (2004) and Storesletten *et al.* (2004) come to the same qualitative conclusion.

relative to the data. The model also does not display an increase in the use of uncollateralized credit over time, unless more flexible borrowing constraints are introduced.

On the basis of these findings, we conjecture that a model with an endogenous evolution of credit markets (like the DCM), but with a span of asset markets in between the one in the DCM and the SIM model, might be even more empirically successful in matching the data. Further empirical work using micro-data may inform us how to model more precisely the mechanisms that households can use to smooth out idiosyncratic income shocks. We defer this work to ongoing and future research.

APPENDIX A. DATA DESCRIPTION

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

A.1. *CE files and observations*

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

A.2. *Consumption categories*

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

A.2.1. **Food expenditures correction.** A change in survey methodology (see Battistin, 2003, for details) causes a sizeable (about 15%) systematic downward bias in reported food expenditures for all the observations in the years 1982–1987. In order to correct for this bias, we regress the log of food expenditures for all years on a quadratic time trend, on quadratics in income and total nonfood consumption expenditures, on weeks worked, on a complete set of household characteristics (including age, education, region of residence, and family composition), on a dummy for the period 1982–1987, and on the interactions term of the dummy with all other independent variables. We then use the regression coefficients to scale up food expenditures for every observation in the period 1982–1987.

A.2.2. **Services from vehicles.** Each CE observation contains reports of expenditures for purchases of new and used vehicles. The CE also reports the number of cars owned by the household in that quarter. For each year, we first select all observations that report positive expenditures for vehicle purchases and run a regression of these expenditures on quadratics in income and total nonvehicle consumption expenditures, weeks worked by household members, expenditures on gasoline, expenditures on public transportation, vehicle maintenance expenditures, the number of cars owned, a complete set of household characteristics (including age, education, region of residence, and family composition), plus quarter dummies. These regressions have an R^2 that ranges from 74% to 94% in our sample years. On average, in every year, a little more than 10% of households report positive expenditures on vehicles. We use the estimated regression coefficients to predict expenditures for vehicles for all households in that year (*i.e.* for those who did and for those who did not report positive vehicle expenditures). Our measure of consumption services from vehicles, then, is the predicted expenditures on vehicles, times the number of cars the consumer unit owns, times $\frac{1}{32}$ (reflecting the assumption of average complete depreciation of a vehicle after 32 quarters).

A.2.3. **Services from primary residence.** Each observation in the CE provides information on whether the household rents or owns its primary residence. If the household rents, we measure housing services as

TABLE A1

Consumption categories and deflators

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

[a]The CPI are monthly average city data for all urban consumers, not seasonally adjusted and they are in base $1982 - 1984 = 100$.
[b]These are mostly fee for services such us banking or legal assistance.
[c]The BLS CPI for entertainment ends in 1998. We extend it to 2004 using the CPI for recreation (SAR).
[d]It includes mostly expenditures on vacation homes.
[e]It includes expenditures on maintenance, repairs, insurance, and finance charges.

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

A.3. *Sample selection*

We first exclude observations for which there is clear evidence of measurement error. In particular, we exclude observations classified as incomplete income respondents, observations that report 0 food expenditures for the quarter, those who report only food expenditures for the quarter, and those who report positive labour income but no hours worked. We then exclude all observations with an age of the household head below 21 or above 64, and those with negative or 0 LEA+

earnings. In our benchmark sample, we also exclude observations with weekly wages of the reference person below half of the minimum wage, households classified as rural, and those households that have not completed the full set of four interviews.

A.4. *Aggregation and top coding*

Inequality measures are computed on annual cross sections. We assign an observation to a given year if the last interview of that household is completed between April of that year and March of the following year. Whenever income or consumption expenditures are top coded, we set them to their top-coding thresholds. We have experimented with increasing the values of top-coded income/consumption components (multiplying the threshold by 1·5). Inequality measures are robust to these changes because, in general, the number of observations with top-coded income or consumption figures never exceeds 2% in a given quarter.

A.5. *Data availability*

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

APPENDIX B. RECURSIVE FORMULATION AND COMPUTATIONAL ALGORITHM

In this appendix, we formulate the consumer problem for the DCM model recursively and provide a sketch of the algorithm used to compute a stationary equilibrium. In the nonstationary case (*i.e.* along the transition), the logic remains the same but all functions have to be indexed by t. For simplicity, here we omit the distinction by types and lump into y the transitory and persistent income shocks. The equilibrium problem is nonstandard because one needs to solve not only

TABLE A2

Summary statistics for the benchmark sample

Year	Households	Average income LEA+	Average consumption	
			ND	ND+
1980	638	15,736	7447	12,940
1981	1439	15,419	7348	12,875
1982	1313	16,453	6949	12,605
1983	1700	16,746	7134	12,873
1984	1771	18,017	7398	13,495
1985	1267	17,862	7323	13,463
1986	748	18,678	7308	13,725
1987	1840	19,318	7558	13,992
1988	1621	18,881	7391	13,752
1989	1762	19,585	7517	14,183
1990	1789	18,893	7455	14,171
1991	1749	19,170	7376	13,858
1992	1725	19,046	7126	13,737
1993	1786	18,841	7115	13,686
1994	1744	18,657	7114	13,498
1995	1325	19,269	7127	13,865
1996	881	20,203	7433	14,674
1997	1609	20,557	7203	14,204
1998	1578	20,952	7323	14,689
1999	1763	22,060	7263	14,532
2000	2164	21,508	7191	14,508
2001	2227	21,468	6967	14,060
2002	2387	23,427	7280	14,766
2003	2593	22,544	7053	14,359

Note: Income and consumption measures are in 1982–1984 constant dollars per adult equivalent. Averages are weighted using CE population weights. A household belongs to year x if its fifth interview is between the second quarter of year x and the first quarter of year $x + 1$.

for prices, but also for endogenous borrowing constraints. We first compute the value of autarky as the fixed point to the functional equation

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

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

$$V(y,a) = \max_{c,\{a'(y')\}_{y' \in Y}} \left\{ (1-\beta)u(c) + \beta \sum_{y' \in Y} \pi(y'|y)V(i,a'(y'),y') \right\}$$

$$\text{s.t.}$$

$$c + \sum_{y' \in Y} q(y'|y)a'(y') = \alpha_i y + a \qquad a' \geq A^i(y').$$

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

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

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

REFERENCES

AIYAGARI, S. R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics*, **109**, 659–684.

ALBARRAN, P. and ATTANASIO, O. (2003), "Limited Commitment and Crowding Out of Private Transfers: Evidence from a Randomised Experiment", *Economic Journal*, **113**, C77–C85.

ALVAREZ, F. and JERMANN, U. (2000), "Efficiency, Equilibrium, and Asset Pricing with Risk of Default", *Econometrica*, **68**, 775–798.

ATTANASIO, O., BATTISTIN, E. and ICHIMURA, H. (2005), "What Really Happened to Consumption Inequality in the US?", in E. Berndt and C. Hulten (eds.) *Measurement Issues in Economics—The Paths Ahead: Essays in Honor of Zvi Griliches* (Chicago: University of Chicago Press) (forthcoming).

ATTANASIO, O. and DAVIS, S. (1996), "Relative Wage Movements and the Distribution of Consumption", *Journal of Political Economy*, **104**, 1227–1262.

BATTISTIN, E. (2003), "Errors in Survey Reports of Consumption Expenditures" (IFS Working Papers W03/07).

BLUNDELL, R., PISTAFERRI, L. and PRESTON, I. (2004), "Consumption Inequality and Partial Insurance" (ISF Working Papers W04/28).

BLUNDELL, R. and PRESTON, I. (1998), "Consumption Inequality and Income Uncertainty", *Quarterly Journal of Economics*, **113**, 603–640.

BOWLUS, A. and ROBIN, J.-M. (2004), "Twenty Years of Rising Inequality in U.S. Lifetime Labour Income Values", *Review of Economic Studies*, **71**, 709–742.

COUNCIL OF ECONOMIC ADVISERS (2005) *Economic Report of the President, 2005*, Table B-77 (Washington, DC: Executive Office of the President).

CUTLER, D. and GRUBER, J. (1996), "Does Public Insurance Crowd Out Private Insurance?", *Quarterly Journal of Economics*, **111**, 391–430.

CUTLER, D. and KATZ, L. (1991*a*), "Macroeconomic Performance and the Disadvantaged", *Brookings Papers on Economic Activity*, 1–74.

CUTLER, D. and KATZ, L. (1991*b*), "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s", *American Economic Review*, **82**, 546–551.

DALAKER, J. and NAIFEH, M. (1998) *Poverty in the United States: 1997, U.S. Bureau of the Census, Current Population Reports, Series P60-201* (Washington, DC: U.S. Government Printing Office).

DEATON, A. and PAXSON, C. (1994), "Intertemporal Choice and Inequality", *Journal of Political Economy*, **102**, 437–467.

DÍAZ-GIMÉNEZ, J., QUADRINI, V. and RÍOS-RULL, J.-V. (1997), "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth", *Federal Reserve Bank of Minneapolis Quarterly Review*, **21**, 3–21.

DYNARSKI, S. and GRUBER, J. (1997), "Can Families Smooth Variable Earnings?", *Brookings Papers on Economic Activity*, 229–303.

FERNÁNDEZ-VILLAVERDE, J. and KRUEGER, D. (2002), "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?" (Mimeo, University of Pennsylvania).

FRIEDMAN, M. (1957) *A Theory of the Consumption Function* (Princeton, NJ: Princeton University Press).

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

GOTTSCHALK, P. and SMEEDING, T. (1997), "Cross-National Comparisons of Earnings and Income Inequality", *Journal of Economic Literature*, **35**, 633–687.

GREENSPAN, A. (1998), "Opening Remarks", in Federal Reserve Bank of Kansas City (ed.) *Income Inequality: Issues and Policy Options*. A Symposium Sponsored by the Federal Reserve Bank of Kansas City.

GUVENEN, F. (2005), "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" (Mimeo, University of Rochester).

HEATHCOTE, J., STORESLETTEN, K. and VIOLANTE, G. (2004), "The Macroeconomic Implications of Rising Wage Inequality in the US" (Mimeo, Georgetown University).

HEATON, J. and LUCAS, D. (1996), "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing", *Journal of Political Economy*, **104**, 443–487.

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

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

JOHNSON, D. and SMEEDING, T. (1998), "Measuring the Trends in Inequality of Individuals and Families: Income and Consumption" (Mimeo, Bureau of Labor Statistics).

KATZ, L. and AUTOR, D. (1999), "Changes in the Wage Structure and Earnings Inequality", in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, Vol. 3A (Amsterdam: North-Holland).

KEHOE, T. and LEVINE, D. (1993), "Debt Constrained Asset Markets", *Review of Economic Studies*, **60**, 865–888.

KEHOE, T. and LEVINE, D. (2001), "Liquidity Constrained Markets versus Debt Constrained Markets", *Econometrica*, **69**, 575–598.

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

KRUEGER, D. and PERRI, F. (1999), "Risk Sharing: Private Insurance Markets or Redistributive Taxes?" (Federal Reserve Bank of Minneapolis Staff Report 262).

KRUEGER, D. and PERRI, F. (2004), "On the Welfare Consequences of the Increase in Inequality in the United States", *18th NBER Macroeconomics Annual 2003*, 83–121.

KRUEGER, D. and UHLIG, H. (2005), "Competitive Risk Sharing Contracts with One-Sided Limited Commitment", *Journal of Monetary Economics* (forthcoming).

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

McGRATTAN, E. and PRESCOTT, E. (2003), "Average Debt and Equity Returns: Puzzling?", *American Economic Review*, **93**, 392–397.

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

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

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

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

STORESLETTEN, K., TELMER, C. and YARON, A. (1998), "The Risk Sharing Implications of Alternative Social Security Arrangements", *Carnegie-Rochester Conference Series on Public Policy*, **50**, 213–259.

STORESLETTEN, K., TELMER, C. and YARON, A. (2004), "Consumption and Risk Sharing over the Life Cycle", *Journal of Monetary Economics*, **51**, 609–633.

TAUCHEN, G. and HUSSEY, R. (1991), "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models", *Econometrica*, **59**, 371–396.

U.S. CENSUS BUREAU (2004), "Gini Ratios for Households, by Race and Hispanic Origin of Householder: 1967 to 2000" (http://www.census.gov/hhes/income/histinc/h04.html).

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

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