DISCUSSION PAPER SERIES

No. 8701

THE WRONG SHAPE OF INSURANCE? WHAT CROSS-SECTIONAL DISTRIBUTIONS TELL US ABOUT MODELS OF CONSUMPTION-SMOOTHING

Tobias Broer

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP8701.asp**

ISSN 0265-8003

THE WRONG SHAPE OF INSURANCE? WHAT CROSS-SECTIONAL DISTRIBUTIONS TELL US ABOUT MODELS OF CONSUMPTION-SMOOTHING

Tobias Broer, IIES, Stockholm University and CEPR

Discussion Paper No. 8701
December 2011

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 8701

December 2011

ABSTRACT

The wrong shape of insurance? What cross-sectional distributions tell us about models of consumption-smoothing*

This paper shows how two standard models of consumption risk-sharing–self-insurance through borrowing and saving and limited commitment to insurance contracts–replicate similarly well the standard, second-moment measures of insurance observed in US micro-data. A non-parametric analysis, however, reveals strongly contrasting and counterfactual joint distributions of consumption, income and wealth. Method of moments estimation shows how measurement error in consumption eliminates excessive skewness and concentration of consumption growth. Moreover, counterfactual non-linearities disappear at high estimated risk-aversion under self-insurance, but are a robust feature of limited commitment. Its "shape of insurance" thus argues strongly in favour of the self-insurance model.

JEL Classification: D31, D52, E21 and E44
Keywords: limited commitment, risk sharing, wealth and consumption distribution

Tobias Broer
Institute for International Economics
Stockholm University
SE-10691 Stockholm
SWEDEN

Email: tobias.broer@iies.su.se

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=162801

* I am indebted to Morten Ravn, Giancarlo Corsetti and Ramon Marimon. I thank Árpád Ábráhám, Jonathan Heathcote, Dirk Krueger, Per Krusell, Sarolta Laczo, Nicola Pavoni, Fabrizio Perri, Gianluca Violante, and especially Kjetil Storesletten, as well as seminar participants at the Federal Reserve Bank of Minneapolis, HECER, IIES, NUS, NORMAC 2010 and the University of Zürich for comments on earlier drafts. The author gratefully acknowledges financial support from the European Research Council under Advanced Grant 230574.

Submitted 28 November 2011

1 Introduction

Understanding the economic determinants of household risk-sharing is important not just for welfare considerations, but also because of the contrasting implications of different models for the behaviour of macro-economic variables, absolute and relative asset prices and the effectiveness of policies to reduce inequality or individual risks. A long literature has tried to identify the nature of household risk-sharing by comparing the predictions of different theories to both macro and micro-data (see Deaton (1992) and Attanasio and Weber (2010) for surveys). Using the sharp and contrasting predictions about the effect of income shocks on consumption that result, respectively, from Keynesian models, the permanent income hypothesis (PIH) and models with perfect insurance, the early empirical literature has focused on the response of household consumption to income changes, and on the reaction of its cross-sectional dispersion as shocks drive household incomes apart over time and the life-cycle. Recent computational advances have allowed researchers to solve for the joint equilibrium distribution of consumption and income in more general incomplete markets models, as well as in micro-founded models where information or contracting frictions in insurance markets imply different degrees of "partial insurance". Although most of these models make less sharp predictions about the consumption effect of income shocks, many studies have continued to assess their performance on the basis of the two standard second order moments, namely the average response of consumption growth to income changes and the cross-sectional variance of consumption relative to that of incomes.[1]

This paper argues that the joint distribution of income, consumption and wealth implied by quantitative equilibrium models contains important information beyond the second-moment statistics that researchers commonly focus on. And it argues that we can use this information to better identify models and discriminate between them. To illustrate this, I look at two standard models of consumption insurance, namely a complete markets model with limited commitment to contracts as in Krueger and Perri (2006), and a simple self-insurance model where agents save and borrow in non-contingent bonds subject to an exogenous borrowing limit (Aiyagari 1993, 1994). I show how calibrated versions of these models with a realistic level of aggregate wealth replicate similarly well the standard second-moment properties of the joint income and consumption distribution in US micro-data, although, as pointed out in Krueger and Perri (2006), the self-insurance model has

[1] See the literature review for further details and references.

somewhat too little, and the limited commitment model somewhat too strong insurance. I then use non-parametric techniques to show first how, despite their relative success in replicating the standard measures, both models have equilibrium distributions with strongly counterfactual features. Particularly, both models predict conditional standard deviations of consumption growth that are an order of magnitude lower than in the data across the whole income distribution, which I take as evidence for measurement error in consumption that tends to affect measured growth rates more strongly than levels. Moreover, apart from positive skewness in the marginal distribution of consumption, conditional means and variances have a non-linear relationship with income in both models, but this non-linearity is significantly stronger in the limited commitment model, whose consumption process combines upward jumps when participation constraints bind with slow downward drift during unconstrained periods. The joint distribution of wealth and income has similar non-linearities, and a counterfactual negative slope in the limited commitment model, where high income comes with high insurance liabilities, rather than high average wealth as under self-insurance and in the data. Building on these non-parametric results I estimate the models using a simulated method of moments procedure to show how the low conditional variances can be used to identify measurement error in consumption, and how a high estimated value of risk-aversion reduces the non-linearity of self-insurance significantly. The limited commitment model, on the other hand, shows strongly counterfactual non-linearities even at very low estimated levels of risk-aversion. Alternative data sources, which show less insurance than the CEX benchmark data, strengthen this finding. For example, Blundell, Pistaferri and Preston's (2008, hf. BPP) PSID dataset shows significantly stronger comovement of income without significant non-linearities, and the volatility of consumption relative to that of incomes is more than twice as high as in the CEX. Both of these facts favour the self-insurance model. I conclude that the shape of stationary model-distributions contains important information beyond second moments, and that it argues strongly in favour of a standard incomplete markets model.

Section II briefly summarises the related literature and the relative contribution of this paper. Section III describes the environment of a simple stationary economy with idiosyncratic income shocks and presents two market structures, with non-contingent assets and complete markets subject to limited enforcement of contracts. Section IV presents the main quantitative findings. Section V looks at alternative data sources and analyses the sensitivity of the results to a different specification of the exogenous income process.

2 Related Literature

Since the early 1980s, a growing number of studies has compared the observable implications of different risk-sharing mechanisms to the empirical evidence.[2] For this, most early studies used the contrasting predictions of different theories about the effect of anticipated and unanticipated income changes on consumption: while in simple Keynesian models both have equal consumption effects, neither affects relative consumption under perfect insurance. The pure PIH with infinitely lived agents and quadratic preferences, on the other hand, predicts that forward-looking consumers do not change their consumption path when they observe anticipated income changes, whereas the consumption response to unanticipated income shocks is approximately 1 if these are permanent in nature, but approximately 0 if they are perfectly transitory. In seminal contributions using data from the US Panel Study of Income Dynamics (PSID), Hall and Mishkin (1982) found smaller but non-zero responses to transitory income shocks, contradicting both Keynesian theory and the PIH, while Altonji and Siow (1987) conclude that the data "generally support the life-cycle model" (p. 293). Early tests of the hypothesis of perfect risk-sharing, that is, for example, implicit in most representative agent models, on the other hand, came in two kinds. A first approach tested the hypothesis that individual income shocks do not cause differences in consumption growth, equivalent to $\widehat{b} = 0$ in the regression

$$dc_{it} = k + \widehat{b} dy_{it} + v_t + \epsilon_{it} \tag{1}$$

where dc_{it} and dy_{it} denote growth rates of individual consumption and income respectively, k is a constant, v_t a vector of time dummies and ϵ_i a residual. The evidence reported in Nelson (1994) (using US Consumer Expenditure Survey (CEX) data and somewhat in contrast to earlier results by Mace (1991)) as well as Cochrane (1991) (on PSID data) reject perfect risk-sharing. Following Deaton and Paxson (1994), an alternative test of perfect insurance is based on its implication that the dispersion of consumption should be constant even when incomes move apart over time in response to persistent income shocks. Thus, Deaton and Paxson (1994) show how the variance of consumption rises with age, while Attanasio and Davis (1996) find that the evolution of relative consumption across education groups is strongly correlated with that of their relative incomes at low frequencies, rejecting perfect insurance.

[2] For a summary of the empirical literature on consumption and savings, see Attanasio and Weber (2010).

In an important contribution Dynarski and Gruber (1997) estimate a $\widehat{b}$ of 10% for non-durables and 17% for durables in CEX data, but find that the variance of food consumption has risen as much as that of earnings between 1970 and 1991 in PSID data. In contrast to many studies, they also look at the non-linearity of income effects on consumption finding somewhat stronger effects of income declines than of increases.

The 1990s and 2000s saw two important developments in the macro-economic literature analysing economies with idiosyncratic income risk. First, computational advances allowed the quantitative analysis of more complex models of consumption-smoothing under incomplete markets (Huggett 1993, Aiyagari 1994, Rios-Rull 1995). And second, evidence that these "self-insurance models", where agents borrow and save exclusively in non-contingent assets, tend to overpredict the cross-sectional dispersion of consumption and the effect of income shocks gave rise to models where insurance through state-contingent contracts is available, but constrained by limits to information (Attanasio and Pavoni 2007) or contract enforcement (Krueger and Perri 2005, 2006, 2011). Although in these quantitative models researchers can easily analyse the whole joint income and consumption distribution, most studies maintained the focus on relative cross-sectional dispersion measures of consumption and the simple regression coefficient $\widehat{b}$ in equations like (1). Thus, Krueger and Perri (2006) find that a partial insurance model with limited commitment can explain a much smaller rise in consumption dispersion relative to that of incomes, as observed in CEX data between 1980 and 2003. Although the limited commitment understates the observed rise in the consumption dispersion, a simple self-insurance model overstates it. In a related paper, Cordoba (2008) shows that both models can replicate reasonably well the level of cross-sectional consumption variance, but that self-insurance models are much better at predicting wealth inequality and particularly its concentration at the top of the distribution. Attanasio and Pavoni (2007) analyse a different model of partial insurance, where missing information about individual effort levels leads to a moral hazard problem for consumption risk-sharing. They show how this can lead to risk-sharing that is not perfect but exceeds that of simple permanent income models. They find evidence for this in data from the UK Family Expenditure Survey, based on both a test for excess smoothness in the consumption response to income shocks and on the response of the consumption variance to changes in the variance of incomes.

In an important contribution, Blundell, Pistaferri and Preston (2008, BPP) have presented important new evidence on the effect of income shocks on consumption, based on a novel US data set they construct by inferring a measure of nondurable consumption for the PSID. Its panel structure allows them to estimate the consumption response to identified permanent and transitory income shocks. They show that the consumption response to transitory income shocks is not significantly different from zero, while that to permanent income shocks is significantly smaller than one. Importantly, Kaplan and Violante (2010) show how a life-cycle version of the self-insurance model has stronger consumption effects than those estimated by BPP. Heathcote, Storesletten and Violante (2010) use a similar model, amended for education choice and a joint labour supply decision within households, to show how it can replicate important moments of US micro-data after 1970, including an important increase in the covariance between hours and wages, and of the cross-sectional variance of consumption.

Relative to this literature, the contribution of this paper is three-fold. First, it is the first to draw attention to the fact that standard models that replicate commonly-used insurance measures similarly well, although not perfectly, have strongly contrasting and counterfactual joint distributions of income, consumption and wealth. Specifically, a non-parametric plot of the joint distributions shows first that both the limited commitment model and self-insurance imply far too little turbulence in consumption, as manifested by a counterfactually low conditional variance of consumption growth. Moreover, the limited commitment model implies a strongly counterfactual negative covariance between wealth and income. And finally, both models have conditional mean and variance functions with counterfactual non-linearities, which are, however, stronger under limited commitment to insurance contracts. Second, I show how we can use moments that summarise these contrasting and counterfactual features in an estimation exercise, rather than through calibration as in previous comparative studies on consumption insurance, to better identify parameters, and to discriminate between models. Particularly, I show how the low turbulence in consumption can be used to identify measurement error, while the asymmetry and non-linearity of consumption-income distributions argues against the limited commitment model at all parameter values, in favour of a self-insurance model with high risk-aversion. Finally, this paper evaluates the theoretical models against both main sources for household data on consumption and income in the US: CEX and PSID, including the BPP imputation of non-durable consumption. This is crucial for its conclusion, because in the PSID, the values of both standard measures of insurance (the variance of consumption relative to that of incomes, and the $\widehat{b}$ coefficient summarising the comovement of consumption and income growth) are about twice those in the CEX, implying significantly less insurance. This reinforces the conclusion in favour the self-insurance model, whose main anomaly, when evaluated on CEX data, was too small a degree of insurance.

Note that this paper is not the first to look at equilibrium distributions of consumption smoothing models or non-linearities in empirical or theoretical distributions.[3] Particularly, the non-linearity in the response of consumption to income rises and declines in the limited commitment model has previously been pointed out by Krueger and Perri (2005). On the basis of a calibrated endowment model with AR(1) shocks, they conclude that a setup which combines characteristics of the limited commitment model with self-insurance, which they show has far too little consumption insurance, would be most promising to explain the facts in US micro-data. Relative to their study, this paper uses a more general model with production and aggregate wealth as well as permanent and transitory shocks, which significantly improves the performance of the self-insurance model in terms of the standard insurance coefficients. Moreover, this paper bases its parametric estimation exercise on a non-parametric analysis of the complete joint distribution of consumption, income and wealth and their growth rates. This identifies the non-linear association of consumption and income growth as but one of the counterfactual and contrasting features of the model distributions. Moreover, the estimation exercise shows how the main anomalies of the limited commitment model are robust to parameter values, while higher risk-aversion and modest estimated measurement error improves further the performance of the self-insurance model, although it does not eliminate the somewhat too strong consumption variance relative to CEX data.

3 The economy

This section presents a simple economy where a continuum of agents faces idiosyncratic income shocks, and the two market structures analysed in this paper: self-insurance through borrowing and saving in non-contingent assets, and insurance contracts without commitment.

[3] For example, Battistin et al (2007) show how Gibrat's law implies log-normality of the marginal consumption distribution under the PIH, for which they find evidence in US micro data. Also, Dynan et al (2006) show that in PSID data, while consumption responds more strongly to negative income shocks, this asymmetry has fallen over time, which they take as evidence of declining liquidity constraints. And Krueger and Perri (2008) show that in the Italian Household Survey the relation between nondurable consumption and income changes unexplained by a first stage regression on household characteristics is largely linear, with a slightly stronger response of consumption to positive income changes.

3.1 The economic environment

The economy consists of a large number of individuals of unit mass. Individuals are indexed by i, located on a unit-interval $i \in I = [0,1]$ with Sigma-Algebra $\mathbb{I}$. Denote as $\Phi_I : \mathbb{I} \to [0,1]$ the (constant) non-atomic measure of individuals. Time is discrete $t \in \{0,1,2,...,\infty\}$ and a unique perishable good is used for consumption.
Idiosyncratic risks arise from fluctuations in individual endowments of effective labour units that agents supply inelastically. The labour endowment of agent i in period t, $z_{i,t}$, takes values in a finite set Z: $z_{i,t} \in Z = \{z^1 > z^2 > ... > z^N\}, N \geq 2$. Let $\mathbb{Z}$ be the power set of Z, and denote as $\Phi_{Z,t} : \mathbb{Z} \to [0,1]$ the measure of agents at all (subsets of) labour endowment realisations in period t. Labour endowments follow a Markov process that is independent of i, and $\mathbb{I}$-measurable (i.e. $\{i : z_{i,t+1} = z^k | z_{i,t} = z^j\} \in \mathbb{I}, \ \forall z^j, z^k$). The process is described by a Markov transition matrix F that has strictly positive entries $\pi_{i,j} > 0, \forall i,j$, is monotone (in the sense that the conditional expectation of an increasing function of tomorrow's income is itself an increasing function of today's income), and has a unique ergodic distribution $\Phi_Z : \mathbb{Z} \to [0,1]$. Thus, in the long-run, the aggregate (or average) labour endowment $L_t = \int z_{it} d\Phi_I$ is constant, while individual labour units fluctuate. Let $Z_0 : I \to Z$ be a measurable function that assigns all individuals an initial labour endowment. For a given wage rate per unit of effective labour w_t, let $y_{it} = w_t z_{i,t}$ denote agent i's labour income. I denote as y^t an individual's history of income until period t, and as $\pi(y'|y)$ transition probabilities in terms of income. Also, I denote as $\Phi_{ay,t}$ the joint distribution of income and assets in period t induced by Z_0, F, and individual decisions as described below.
Agents live forever and order consumption sequences according to the utility function

$$U = E_0 \sum_0^\infty \beta^t u(c_{i,t}) \tag{2}$$

where E_0 is the mathematical expectation conditional on period 0 information, $0 < \beta < 1$ discounts future utility, $c_{i,t}$ is consumption by agent i in period t, and $u : R^+ \to R$ is an increasing, strictly concave, twice-continuously differentiable function that satifies Inada conditions and is identical for all agents in the economy.
A representative competitive firm in the the economy hires labour L_t and capital K_t every period at rental rates r_t and wages w_t, in order to maximise profits from operating a neoclassical production technology $AF(K_t, L_t)$, where A is a productivity parameter. I focus on stationary equilibria, and therefore abstract from aggregate fluctuations. So A

is constant over time. Investment in capital, which depreciates at a constant rate δ every period, is performed by competitive financial intermediaries that live for one period only, and also trade assets with individuals.[4]

In the following, I present two structures of financial markets and associated equilibrium definitions.

3.2 Complete markets with limited commitment to contracts

First, consider an environment where market are complete. So individuals buy or sell, at price $q(\widetilde{y}_t, y_{t+1})$, a quantity $a(\widetilde{y}_t, y_{t+1})$ of a state-contingent asset that specifies delivery of one unit of the consumption good to/from an agent who experiences endowment realisation y_{t+1} after history $\widetilde{y}_t$.However, agents are unable to commit to future payments: after observing their income they can declare default, which cancels all their financial liabilities or assets but also excludes them from insurance markets forever in the future. Following Abreu, Pearce and Stacchetti (1990), it is often assumed that default triggers the worst subgame perfect equilibrium of autarky. In contrast, in this paper I model the consequences of default in line with those that follow individual bankruptcy in the US economy. Particularly, I assume agents can save in capital after default, and can go back to insurance markets with a constant probability $\hat{p}$. Note that the expected value of default $U^{aut}(y^t)$ puts a lower bound on the utility agents derive from participation in state-contingent trade. As in Kehoe and Levine (1993), the problem of a typical household with initial income y_0 and asset holdings a_0 can thus be formulated as a choice of contingent consumption and asset plans, given prices and lower bounds on utility

$$V(a_0, y_0) = max_{c_s(a_s,y^s),\{a_{s+1}(a_s,y^s,y_{s+1})\}} U_0 \tag{3}$$

$$s.t.\ c_t(a_t, y^t) + \sum_{y_{t+1}} q(y^t, y_{t+1}) a_{t+1}(a_t, y^t, y_{t+1}) \leq a_t + y_t \tag{4}$$

$$U_t(a_t, y^t) \geq U^{aut}(y^t) \tag{5}$$

Definition. A competitive equilibrium with participation constraints under limited commitment to state-contingent contracts is a sequence of prices $\{r_t, w_t, \{q(y^t, y_{t+1})\}\}_{t=0}^{\infty}$ and an allocation $\{K_t, c(a_t, y^t), \{a_{t+1}(a_t, y^t, y_{t+1})\}, \Phi_{ay,t}\}_{t=0}^{\infty}$ such that

[4] The short life of intermediaries serves to circumvent the shareholder disagreement problem that arises in economies with heterogeneous individuals.

- $c(a_t, y^t), \{a_{t+1}(a_t, y^t, y_{t+1})\}$ solve the household problem (3) to (5) given prices and initial values of income and assets
- The net marginal returns to capital and labour equal their rental prices

$$AF_K(K_t, 1) + 1 - \delta = r_t \tag{6}$$

$$AF_L(K_t, 1) = w_t \tag{7}$$

- $\Phi_{ay,t}$ is the joint distribution of assets and incomes induced by Z_0, F, and individual decisions $\{a_{t+1}(a_t, y^t, y_{t+1})\}$.
- State-contingent prices charged by intermediaries satisfy the no-arbitrage condition

$$r_{t+1} + 1 - \delta = \frac{\pi(y_{t+1}|y_t)}{q_t(y^t, y_{t+1})} \tag{8}$$

- The sum of net returns to state-contingent assets equals those to capital

$$\sum_{y_{t+1}} \sum_{y^t} \int a_{t+1}(a_t, y^t, y_{t+1}) \pi(y_{t+1}|y_t) d\Phi_{ay}(., y^t) = (r_{t+1} + 1 - \delta) K_{t+1} \tag{9}$$

Given the absence of aggregate risk I concentrate on "stationary" equilibria, where the distribution of individual consumption is stationary through time.[5] Since the price of state-contingent securities in a continuum economy without aggregate uncertainty is constant across states of nature, the price of state-contingent assets is simply a probability-weighted version of the price of uncontingent bonds, or

$$R = \frac{\pi(y_{t+1}|y_t)}{q(y^t, y_{t+1})} \ \forall y_{t+1} \tag{10}$$

where the constant interest rate R is determined by market-clearing. This yields the first-order condition for state-contingent assets as

$$\frac{U'(c_i')}{U'(c_i)} = [R\beta(1 + \gamma_i')]^{-1} \tag{11}$$

where the Lagrange multiplier γ_i' is positive in states where agent i's participation constraint binds. Note how equation (11) has two important implications: first, the marginal

[5]Note that, contrary to an endowment economy, in an environment with production, capital accumulation has externalities on incentives through its impact on wages and thus the outside option to the contract. Hence, the competitive equilibrium allocation is not efficient, as shown by Abraham and Carceles (2006).

rate of substitution between consumption today and in unconstrained states of nature next period, where $\gamma_i' = 0$, is constant across agents, equal to $\frac{1}{R\beta}$. And second, among agents with identical consumption today, those that experience a binding participation constraint tomorrow, with $\gamma_i' > 0$, have strictly higher consumption growth than their unconstrained peers. In a stationary equilibrium without perfect insurance this implies that constrained agents experience different discrete jumps in consumption to a level where, given optimal use of insurance contracts in the future, expected utility is equal to autarky utility $U^{aut}(y^t)$. Since the latter only depends on the current income $y_{i,t}$ of constrained agents, so does their current level of consumption under the insurance scheme. This lack of history dependence in consumption of constrained individuals is well-known as the "amnesia" (Kocherlakota 1996), or "forgiveness" (Thomas and Worrall 1994) property of consumption allocations with limited commitment. In a stationary equilibrium with constant aggregate consumption, the counterpart of these consumption increases for constrained agents are declines in consumption for unconstrained agents. Particularly, from equation (11) with $\gamma_i' = 0$ and $\beta R < 1$, unconstrained individuals experience (in marginal utility terms) equal smooth downward movements in consumption.

Krueger and Uhlig (2006), Thomas and Worrall (2007) and Broer (2009b) (for the economy with $N > 2$) show how the stationary cross-sectional distribution of consumption in this economy is exactly pinned down by the participation constraints and the recursive law of motion of consumption for unconstrained agents (11). This distribution inherits the asymmetry of the insurance mechanism, where downward drift in consumption of unconstrained agents is counteracted by occasional upward jumps to fulfill participation constraints. The marginal distribution of consumption is thus a mixture of geometric distributions, corresponding to the downward-sloping consumption paths of unconstrained agents from any of the N levels of participation-constrained consumption. For the joint distribution of consumption and income, this results in strong nonlinearity and heteroscedasticity. For example, while all individuals at maximum income z^1 are constrained at a common consumption level, individuals at z^N are found throughout the whole consumption distribution, depending on the number of periods since they were last constrained at a higher income level. Similarly, the joint distribution of income and consumption <u>growth</u> is characterised by the fact that individuals with negative income shocks share the same upward drift in marginal utility, and thus have similar downward drift in consumption, while positive income shocks lead to a large variety of positive consumption jumps.

3.3 Self-insurance

I also analyse a version of the economy where state-contingent contracts are assumed to be absent, but agents can exchange non-contingent bonds to smooth consumption over time subject to a borrowing limit. The setup of the economy and the equilibrium are thus essentially those of Aiyagari (1993), and the household problem is

$$V(b_0, y_0) = max_{c_s(b_s,y^s),b_{s+1}(b_s,y^s)} U_0 \tag{12}$$

$$s.t.\ c_t(b_t, y^t) + b_{t+1}(b_t, y^t) \leq R_t b_t + y_t \tag{13}$$

Definition. A competitive equilibrium under self-insurance through non-contingent assets is a sequence of prices $\{r_t, w_t, R_t)\}_{t=0}^{\infty}$ and an allocation $\{K_t, c(b_t, y^t), b_{t+1}(b_t, y^t), \Phi_{by,t}\}_{t=0}^{\infty}$ such that

- $c(a_t, y^t), b_{t+1}(b_t, y^t)$ solve the household problem (12) to (13) given prices
- $\Phi_{by,t}$ is the joint distribution of assets and incomes induced by Z_0, F, and individual decisions $b_{t+1}(b_t, y^t)$.
- the net marginal returns to capital and labour equal their rental prices

$$AF_K(K_t, 1) + 1 - \delta = r_t \tag{14}$$

$$AF_L(K_t, 1) = w_t \tag{15}$$

- The bond market clears

$$\sum_{y^t} \int R_{t+1} b_{t+1}(a_t, y^t) d\Phi_{by}(., y^t) = (r_{t+1} + 1 - \delta) K_{t+1} \tag{16}$$

4 Quantitative Results

4.1 Data

As many previous studies, I look at US micro-data to assess the performance of consumption smoothing models. For this, we would ideally compare the joint stochastic process for individual income and consumption that results from the models with panel data observations for US households or individuals. Unfortunately, these data are not available. While the PSID is a long panel of households with detailed information on income, until recently, it only collected information on food consumption. Comprehensive consumption data is available in the CEX survey, which, however, follows households for at most 5 consecutive quarterly surveys, and only collects information on income in the second and fifth of these. Since the CEX has the best consumption data, and in order to assure comparibility with previous studies such as Nelson (1994), Mace (1991), Dynarski and Gruber (1997), Krueger and Perri (2006, 2005), and many others, the benchmark results of this paper are based on data from the CEX interview survey. This comes at the price that the timing of observations on consumption (quarterly in interviews 2 through 5) and income (retrospective questions on yearly averages in interview 2 and 5 only) is not totally consistent. In order to adress this, a subsequent section compares the benchmark results to those based on two other datasets: first, the new PSID dataset by BPP, who interpolate a comprehensive series for non-durable consumption using the individual PSID data on food consumption and an inverted food demand function estimated on CEX data; and second, a truely quarterly CEX dataset, based on the income interpolation by Gervais and Klein (2008).

In the benchmark results of this section, I use CEX data for the years 1998 to 2003. The length of the time period is a compromise between a sufficient number of observations and sufficiently small changes in the income process, which in the model is constant and exogenous. Moreover, since I specify the consequences of default in accordance with US bankruptcy regulation, which underwent important changes in 2005, the final period is 2003. As in Krueger and Perri (2006), I choose a sample that only includes urban households who are full income and consumption respondents and whose head is older than 20 and younger than 65 years of age. I also exclude households who report 0 or only food consumption, who have negative labour income, and those whose hourly wage is below half the mininum wage. For consumption, I use Krueger and Perri's (2006) measure of non-durable consumption including an estimate of service flows from housing and cars. As the income measure, I use after-tax labour earnings plus transfers, defined as the sum of wages and salaries of all household members, plus a fixed fraction of self-employment farm and nonfarm income, minus reported federal, state, and local taxes (net of refunds) and social security contributions. Moreover, I divide both consumption and income by the number of adult equivalents in the household, using the census equivalence scale as in Krueger and Perri (2006).

Following MaCurdy (1982), Hubbard et al (1995), Storesletten (2004) or Krueger and Perri (2006), identify the income process using a standard model with restricted income

profiles.[6] Specifically, I assume $log(y_{it})$ to be the sum of a group specific component α_{jt} and an idiosyncratic part x_{it}. The latter, in turn, is the sum of a persistent AR(1) process m_{it}, with persistence parameter ρ and variance σ_m^2, plus a completely transitory component ε_{it} which has mean zero and variance σ_ε^2.
The process for LEA+ is thus of the form

$$\begin{aligned} log(y_{it}) &= \alpha_{jt} + x_{it} \\ x_{it} &= m_{it} + \varepsilon_{it} \\ m_{it} &= \rho m_{it-1} + \nu_{it} \\ \varepsilon &\sim N(0, \sigma_\varepsilon^2) \\ \nu_{it} &\sim N(0, \sigma_\nu^2) \end{aligned} \tag{17}$$

From both consumption and income, I first partial out the group-specific component α_{jt} through regression on a cubic function of age and dummies for education (and their interaction with age), gender, race and a dummy for professionals. Then, I identify σ_ε^2 and σ_ν^2 from the stationary variance and autocovariance of x_{it} as

$$cov(x_{it}, x_{it-1}) = \frac{\rho}{1-\rho^2}\sigma_\nu^2 \tag{18}$$

$$var(x_{it}) = \frac{1}{\rho}cov(x_{it}, x_{t-1}) + \sigma_\varepsilon^2 \tag{19}$$

Setting $\rho = 0.9989$, the value estimated by Storesletten et al (2004), this yields $var(m_t) = 0.23$ and $var(\nu_{it}) = 0.11$, very close to the values found, for example, by Krueger and Perri (2006) for CEX data in this period. I then use the standard Tauchen and Hussey (1999) method to approximate the resulting process using a 7-state Markov chain for m_{it}, and a binary process for ν_{it}.

4.2 Functional forms, punishment of default and borrowing limits

To compare the model-implied allocation to the data, we have to specify functional forms and their parameter values, as well as the aggregate wealth in the economy and the exact

[6] This limits the heterogeneity of income profiles to a group specific constant and stochastic shocks with identical (autoco)variances. An alternative is to use heterogeneous income profiles with, for example, variation in slope coeficients across groups. See e.g. Guvenen (2009).

punishment of default. I choose a utility function with constant relative risk aversion of σ

$$u(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma} \tag{20}$$

The punishment of default in the limited commitment environment tries to capture key features of US bankruptcy legislation prior to 2005. Particularly, under the most common chapter 7 bankruptcy procedure, a defaulter's liabilities are erased once financial and other assets such as real estate are seized.[7] Default leads to an entry in an individual's credit history that remains on her record for 10 years, which typically makes access to credit more difficult but not impossible. Moreover, she is not allowed to refile for bankruptcy under Chapter 7 for 6 years, although she can file under chapter 13. While under chapter 7 creditors can not seize any of her income, a judge can order payments out of her income under chapter 13.[8] To capture these stylised facts, I assume that, once an individual declares default, she immediately looses all financial assets and liabilities. Moreover, she is excluded from the insurance mechanism but can save (although not borrow) in non-contingent assets. To include the possibility of regaining access to credit markets, I assume that she is re-admitted with a constant probability $\hat{p}$ to the insurance scheme, but not allowed to declare default again.[9] I calibrate $\hat{p} = 15\%$, to achieve an expected duration of exclusion from insurance of approximately 6 years. In the self-insurance environment, I set the exogenous borrowing limit equal to yearly income.

An important parameter of the calibration is medium wealth, as it determines the distribution of financial returns in the economy. This is crucial, since individuals with positive financial returns will never find it optimal to default - they would rather accept financial payments and default tomorrow. Thus, an endowment economy with zero net wealth is likely to overstate default incentives. Krueger and Perri (2006), Cordoba (2008), or Abraham and Carceles (2008) calibrate their models with capital to aggregate wealth-to-output ratios. Since this study wants to capture liquid wealth used for insurance only, and does not aim at capturing the very high wealth holdings in the right-hand tail of the wealth distribution, I follow Kaplan and Violante (2010) and others and exclude the top of the wealth distribution, but include housing assets in my definition of wealth. Particularly, I ensure that the median wealth-to-income ratio in the model equals the median of the

[7] Exemptions vary across states.

[8] See Livshits et al (2007) for details.

[9] I thus exclude the possibility to default under chapter 13. This is not restrictive, as it would never be optimal to default if creditors can seize debtor income.

net worth-to-income ratio of the bottom 90 percent of US households between 21 and 64 years of age according to the 2004 Survey of Consumer Finances (SCF), equal to 1.41. In the calibrated models, I therefore substitute the asset market clearing conditions (9) and (16) by this condition for median wealth.
The remainder of the paper looks at two methods to determine the value of relative risk-aversion σ and the discount factor β, which are the most important coefficients in the model for risk-sharing. First, I look at a standard calibration as in Krueger and Perri (2006), or Cordoba (2008). And second, I show how the results change when I jointly estimate these coefficients to target key moments of the data.

4.3 Insurance coefficients and equilibrium distributions in calibrated models

This section presents quantitative results for a standard calibration of the economies with log-preferences ($\sigma = 1$), and a discount factor β chosen to yield an annual interest rate R of 1.025 in general equilibrium.[10]

4.3.1 Standard insurance measures

As described in more detail in the literature review, comparative analyses of risk-sharing models most commonly focused on second moments of the joint consumption and income distribution, such as the linear association of consumption and income growth, or the relative dispersion of consumption and income and its evolution over time or the life-cycle. Table 1 compares these standard measures of insurance, namely the regression coefficient of consumption on income growth $\widehat{b}$ in equation (1)[11], and the relative variance of log-

[10] The interest rate is chosen in line with an average ex-ante real interest rate on 6-month US treasury bills between 1998 and 2003 of 2.57%.

[11] Note that I focus on the simple linear association between consumption and income growth as summarised by $\widehat{b}$. An alternative would be to differentiate between responses to persistent and transitory shock, in analogy, for example, to the consumption responses to permanent and transitory shocks analysed in BPP. I focus on the simpler measure for three reasons: first, $\widehat{b}$ is one of the classical moments studied in the consumption risk-sharing literature (see, for example, Gervais and Klein (2008) and the references therein). Second, in the environment of this paper, shocks are not truely permanent, which implies a bias in the BPP identification scheme under limited commitment (Broer 2009a) and self-insurance (Kaplan and Violante 2010). And finally, the focus on $\widehat{b}$ allows a simple parametric analysis of non-linearities in the conditional distribution of consumption and income growth in the following section.

consumption and log-income $\frac{Var(C)}{Var(Y)}$ as a measure of the relative dispersion of consumption and income, calculated from a simulation of the models and the data.[12] In CEX data, the variance of annual log consumption residuals has about half the variance of income residuals, while $\widehat{b}$, calculated on log differences of consumption and income over 4 quarters, is with 9% close to the value found for non-durables by Dynarski and Gruber (1997) on an earlier sample.

The results for the calibrated models confirm those from earlier studies (Krueger and Perri (2005), Cordoba (2008) and particularly Krueger and Perri (2006)): the limited commitment model ("LC", row 3) implies an association of income and consumption growth that is with 5.3% somewhat smaller than in the data. And the relative cross-sectional variance of consumption is only about half of that observed in the CEX. The self-insurance economy ("SI", row 4), on the other hand, exhibits differences between model and data moments of similar size, but in the opposite direction, by overstating both the consumption effect of income growth (12%) and the relative cross-sectional variance of consumption (0.57). The literature has concluded from this evidence that both models perform similarly well in replicating key data moments. Moreover, in order to reconcile the theory with the data, it is often suggested to combine features of both models.[13]

This paper takes a different approach: first, in order to discriminate between the two models, it provides further evidence on the underlying distributions of which the insurance coefficients in table 1 are but two moments. And second, it generalises the models by including measurement error in consumption, and identifies its parameters by a transparent estimation procedure, to see how robust the conclusions based on a calibration exercise are.

4.3.2 The shape of insurance: joint distributions of consumption and income

The insurance measures studied in the previous section, namely the relative variance $\frac{Var(C)}{Var(Y)}$ and the regression coefficient $\widehat{b}$, are summary statistics of the joint distributions of, respectively, the levels of consumption and income and their growth rates. Although quantitative models readily allow an analysis of any moment of these joint distributions,

[12] I exploit the stationary nature of the models and calculate their moments from a long simulation over 200.000 periods

[13] See, for example, Krueger and Perri (2006), p. 187-88, or Krueger and Perri (2005). Other studies, such as BPP or Heathcote et al (2010) also conclude from the study of US micro-data that one should combine the standard incomplete markets model with additional insurance possibilities more generally.

Table 1

		β	σ	$\frac{Var(C)}{Var(Y)}$	$\widehat{b}_{dc,dy}$
1	CEX data			0.42	0.09
2	St Error			0.012	0.008
3	LC Calib	0.954	1.00	0.21	0.05
4	SI Calib	0.966	1.00	0.57	0.12

The table presents moments in CEX data (row 1) together with their standard errors (row 2) and in the calibrated models (rows 3 and 4) with limited commitment ("LC") and self-insurance ("SI") that choose β to target an interest rate $R = 1.025$. The standard errors for the data statistics are calculated using a bootstrap procedure with 400 repetitions.

previous studies on consumption insurance in calibrated economies have typically not looked at the their shape beyond second moments. In this section, I propose a simple alternative to the standard parametric comparison of models and the data: a plot of the joint distributions of consumption, income and wealth, and their conditional mean and variance functions. The aim of this exercise is two-fold: first, it shows how, despite similar performance when judged on the basis of common second-order moments, the models' joint distributions not only differ between each other, but also exhibit strongly counterfactual features when compared to the data. And second, the non-parametric analysis of the distributions is a useful heuristic to identify those moments where the discrepancy between models or with the data is largest. This prepares the estimation of more general models with measurement error in the following section.

[Figures 1 and 2 about here.]

Figure 1 presents histograms of the joint consumption-income distributions in the limited commitment model (top panel) and the self-insurance economy (central panel), together with a bivariate kernel density estimate of the distribution in CEX data (bottom panel), using an optimal bandwith (Botev et al 2008).[14] Bigger circles correspond to higher mass at a particular point in consumption-income space. The top panel of figure 1 shows how, in the limited commitment model, consumption rises on average with

[14] As before, I partial out from both series the effect of a vector of observable individual characteristics, to control for ex-ante differences or predictable changes in life-time wealth.

current income, but is highly heteroscedastic. In particular, the conditional variance of consumption declines as we move to higher income values, where the theory showed individuals to be more likely to have binding participation constraints. Interestingly, the self-insurance economy (central panel) also features some decline in the conditional variance of consumption, although less pronounced than the limited commitment economy. The estimate of the joint distribution of consumption and income in the bottom panel is roughly homoscedastic, with a mean of consumption that increases in income.

Figure 2 shows the corresponding mean and variance functions (top and bottom panels respectively), across 5 income bins of equal mass or frequency.[15] The figure can be interpreted as a decomposition of the cross-sectional consumption variance into the variation in means between income groups and variation within groups. The top panel shows how in CEX data, after controlling for education, age, etc., consumption of individuals in the highest income quintile is on average about 50 percent higher than that in the bottom quintile, with increases that are roughly linear across the income distribution. The limited commitment model replicates these conditional means of the data well, while the self-insurance model overstates the rise across the income distribution by 35 pp. As the bottom panel shows, the small relative variance of consumption in the limited commitment economy (shown in table 1) results from counterfactually low within-group variation: conditional standard deviations are less than a third of those in CEX data. Conditional standard deviations in the self-insurance model are on average three quarters of those in the data, so the high relative variance of consumption levels in the model results from the counterfactually strong variation in means across income groups depicted in the top panel. In line with the evidence in figure 1, CEX data exhibit roughly homoscedastic conditional variances, while both models show a decline across the income distribution whose relative magnitude, however, is stronger in the limited commitment model.

Figures 1 and 2, based on consumption and income levels, do not control for unobserved fixed effects over and above those included in the first-stage regression. These unobserved effects, however, leave the joint distribution of growth rates unaffected. Moreover, the theory section provided strong predictions on the behaviour of consumption growth especially under limited commitment to contracts, where unconstrained agents experience common downward drift in consumption, while consumption jumps upward for individuals with binding participation constraints. Figure 3 therefore shows the joint distribution

[15]Note how the figure shows that the marginal distribution of income in the models, whose calibration target is the cross-sectional variance, is similar, but not identical to that in the data.

of consumption and income growth in the models and the data. Most striking is the asymmetry in the limited commitment model (top panel). There, income declines of different magnitude are associated with the same fall in consumption, while positive income growth implies a variety of consumption responses, including infrequent large jumps, leading to strong positive skewness in the marginal distribution of consumption growth. Figure 4 presents the corresponding means and standard deviations of consumption growth across the income distribution, divided into 5 bins of equal mass or frequency.[16] Again, the figure can be interpreted as a decomposition of the consumption growth variance into within group-dispersion and the variation in means across groups, whose average slope is equal to the regression coefficient $\widehat{b}$ in table 1.

The self-insurance model replicates the positive association of consumption and income growth in the data well, with conditional means that lie within, or very close to, the two-standard-error bands. The limited commitment model does significantly worse: the moderate average association between income and consumption growth of 5% in the limited commitment model masks a strong difference between a zero response to income declines, and a stronger average response to income increases.[17] The bottom panel of figure 4 reveals a second striking feature of the joint distributions in the calibrated models: the conditional standard deviation of consumption growth, or "turbulence" in consumption, is an order of magnitude smaller than in the data for both models. And again, the limited commitment model has strongly heteroscedastic conditional dispersions: since consumption declines are of the same magnitude for all agents, the standard deviation is zero for negative values of income growth while consumption increases are more widely dispersed. Both in the data and the self-insurance model conditional standard deviations are much more homoscedastic.

[Figures 3 and 4 about here.]

This section has shown how, despite similar performance when judged on the basis of commonly used measures of insurance, self-insurance and limited commitment to contracts imply very different joint distributions of consumption and income. Particularly, due

[16] As can be seen, the discretised income process, calibrated to have equal cross-sectional variance in levels as the data, implies a slightly different cross-sectional distributions of income growth, relative to the data.

[17] This non-linear response to income rises and declines has been pointed out previuosly in a parametric context by Krueger and Perri (2005).

to its asymmetric insurance mechanism, the limited commitment model implies strong positive skewness of consumption growth, a strongly non-linear response of consumption to income growth, and important heteroscedasticity of conditional distributions. Moreover, both models predict conditional variances of consumption growth that are an order of magnitude smaller than in the data.

4.3.3 The distribution of wealth and income

Figure 5 and 6 perform an exercise similar to that in the previous section for the joint distribution of wealth and income, using the net worth variable of the 2004 wave of the Survey of Consumer Finances (SCF).[18]
Figure 5 shows how the self-insurance economy generates a stronger dispersion of wealth, and higher wealth levels for some agents, than the limited commitment economy, as already pointed out by Cordoba (2008). More striking, however, are the slopes of the conditional mean and variance functions in figure 6, where the limited commitment model shows a counterfactual negative relationship of mean wealth with income, and a strong decline in conditional variances. This is because in that model, a positive income shock implies a rise in financial liabilities in the forms of higher (expected) payments into the insurance scheme. This implies that individuals with highest income have minimum wealth.[19] Negative income shocks, on the other hand, give rise to insurance claims and thus increase wealth. Since individuals slowly deplete their wealth levels after a negative income shock, the income poor have a variety of positive wealth levels, including the highest in the economy. In the self-insurance economy, on the other hand, the bufferstock nature of wealth leads to a positive relationship between income and wealth levels on average. But there is large variation around the mean, as individuals slowly build up, or draw down, their wealth after income changes. Also, the mass of individuals at the borrowing constraint clearly rises as income falls. Finally, in SCF data, after an initial decline in the variance, we see both an increase in mean wealth, as well as in its variance, as income, measured as salaries plus a proportion of business income, rises.

[Figures 5 and 6 about here.]

[18] Again, I use residuals from a regression on individual characteristics, and eliminate the top and bottom percent of observations to control for outliers and top-coding.

[19] In a model with purely transitory shocks, where autarky values are not necessarily monotonous in income, this only holds approximately, as figure 5 shows.

4.4 Estimated Models with Measurement Error

Although both self-insurance and limited commitment to complete contracts were shown to deliver realistic degrees of insurance, the previous section showed how their implied joint distributions of income, consumption and wealth have contrasting and partly counterfactual features. This section shows how we can use these features, together with standard estimation techniques, to discriminate between models, and to better identify their preference parameters, as well as measurement error in consumption. For this, I concentrate on the joint distribution of consumption and income, and thus abstract from the distribution of wealth, which we saw has a counterfactual negative association with income in the limited commitment model. This is because, arguably, the wealth measure in the SCF does not capture (fully) the present discounted value of future contingent claims, such as unemployment benefits or disability insurance, that give rise to the negative correlation between income and wealth under limited commitment. Rather, in order to further discriminate between the two models, that performed similarly on second moments, and to assess their ability to explain the data, I focus on those features of their implied consumption-income distributions that the non-parametric analysis showed to be either counterfactual, or different between the models. So I include in the parametric analysis, first, the relative variance of consumption and income growth $\frac{Var(dc)}{Var(dy)}$, that was counterfactually low in both models and thus suggested the presence of measurement error. Second, I look at moments that capture the asymmetry of the distributions, more prominent in both models than the data. Specifically, I include the skewness of the marginal distribution of consumption growth[20] $skew(dc)$, and the regression coefficients $\widehat{b}_{dc,dy}|_{dy>0}, \widehat{b}_{dc,dy}|_{dy<0}$ in separate regressions for income increases and income declines

$$dc_{it} = k_1 + \widehat{b}_{dc,dy}|_{dy>0} dy_{it} + v_t + \epsilon_{it} \quad if dy_{it} > 0 \tag{21}$$

$$dc_{it} = k_2 + \widehat{b}_{dc,dy}|_{dy<0} dy_{it} + v_t + \epsilon_{it} \quad if dy_{it} \leq 0 \tag{22}$$

Note that so far, I have assumed that both consumption and income are measured without error. The high variance of consumption growth in the data relative to that in the models, however, could be an indicator of important measurement error in consumption. This is because, first, if consumption levels are measured with error, the error of measured growth rates will have a higher variance than that of levels unless errors are strongly serially correlated. Moreover, if consumption is heterogeneous across households but

[20] I use the skewness of growth, rather than levels, because the borrowing limit in the self-insurance model has a direct effect on the skewness in levels, but not on that of growth rates.

smooth between periods, the signal to noise ratio in growth rates will be further reduced relative to that in levels, potentially explaining the high observed variance of consumption growth even with modest measurement error in levels. Previous empirical studies have indeed found evidence of strong measurement error in the CEX consumption measures and in consumption data more generally.[21] This section thus generalises the model by introducing measurement error in consumption. For this, I make the strong assumption that income is measured without error. One reason for this is that recall error is likely to be smaller for income, which often consists of a single, documented monthly flow from one source, than for consumption, which is the sum of numerous smaller transactions. Moreover, as I argue in the conclusion, the main argument of this paper is likely to be strengthened if income is measured with significant error. I thus assume that observed consumption data c_{it} contain classical uncorrelated measurement error μ_{it}[22]

$$c_{it} = c_{it}^{mod} + \mu_{it} \tag{23}$$

where c_{it} is observed log-consumption of individual i at time t, c_{it}^{mod} is the log-consumption implied by the model and μ_{it} is assumed to be an i.i.d. mean-zero random variable with variance Var_{μ}.

According to (23) measured consumption growth equals $dc_{it} = dc_{it}^{mod} + \mu_{it} - \mu_{it-1}$, implying an error term of $\epsilon_{it} = \mu_{it} - \mu_{it-1}$. Since ϵ_{it} is uncorrelated with dy_{it}, measurement error does not bias the estimates of $\widehat{b}$ in equation (1). But its variance is equal to twice the variance in levels, potentially explaining the strong empirical variance of consumption growth relative to those predicted by the models.[23] This effect is even stronger for CEX

[21] For example, Runkle (1991) estimates 76 percent of consumption growth in the PSID to be measurement error, while Ahmed et al (2006) find 70 to 80 percent of the cross-sectional variance in the recall consumption measure of the Canadian Food Expenditure Survey to be due to measurement error. Finally, Cogley (2002) attributes more than 90 percent of the variation in consumption growth in the CEX to measurement error, and Aguiar and Bils (2011) find that, after accounting for income and product-specific measurement error, the increase in CEX consumption inequality since 1980 is similar to that in income inequality, in contrast to the much lower increase of inequality in measured consumption documented by Krueger and Perri (2006).

[22] Ahmed et al (2006) report a small but significant negative correlation between measurement error and consumption data from diaries, their proxy for true consumption. While this is prima-facie evidence against the assumption of classical measurement error, it could also result from measurement error in the proxy.

[23] Of course the assumption of classical measurement error is crucial here. In a model with persistent errors, the variance of consumption growth errors is smaller than that in levels for strong enough persistence.

data, where it is common to measure consumption levels as annual averages, which further reduces the variance of measurement error in levels relative to growth rates, which are based on quarterly data.

Note that the specification of measurement error implies that, for a given variance of consumption in the model $Var(C^{mod})$, its variance Var_μ can, in principle, be exactly identified using only the information on relative variances in levels $\frac{Var(C)}{Var(Y)}$ presented in table 1. However, since there may be individual-specific effects on consumption levels that are not captured by our first-stage regression, including information on the relative variance of consumption growth disciplines this estimate.

For the data and the simple calibrated models presented above, the first four lines of table 2 report statistics for the four additional moments chosen to capture those features of the model distributions that were either counterfactual or differed between the models: the skewness of consumption growth $skew(dc)$, the relative variances of consumption and income growth $\frac{Var(dc)}{Var(dy)}$ and the regression coefficients $\widehat{b_1}, \widehat{b_2}$ in separate regressions for income increases and income declines. CEX data has relative growth variances not too

Table 2

		β	σ	$\frac{Var(C)}{Var(Y)}$	$\frac{Var(dc)}{Var(dy)}$	$skew(dc)$	$\widehat{b}_{dc,dy}\|_{dy<0}$	$\widehat{b}_{dc,dy}\|_{dy>0}$	$\frac{Var_\mu}{Var(C^{mod})}$
1	CEX data			0.42	0.50	0.18	0.07	0.06	
2	St Error			0.012	0.019	0.120	0.014	0.015	
3	LC Calib	0.954	1.00	0.21	0.01	3.89	0.00	0.09	
4	SI Calib	0.966	1.00	0.57	0.05	0.85	0.07	0.11	
5	LC Estim	0.967	0.40	0.29	0.50	0.00	0.00	0.09	0.17
6	SI Estim	0.899	4.00	0.61	0.52	0.02	0.08	0.09	0.08

The table presents moments of CEX data (row 1, with standard errors in row 2), as well as their counterparts in the calibrated models with limited commitment ("LC", row 3) and self-insurance ("SI", row 4). Rows 5 and 6 show moments from estimated models that choose β, σ and the variance of classical measurement error Var_μ to target an equilibrium interest rate $R = 1.025$ the median wealth to income ration and the 4 distributional moments. The standard errors for the data statistics are calculated using a bootstrap procedure with 400 repetitions.

different from those in levels. There is small positive skewness in the consumption growth distribution, but no significant non-linearities in the regression coefficient of consumption growth on income growth. In line with the non-parametric evidence in the previous section, the simple calibrated insurance models (rows 3 and 4) have much lower consumption growth dispersion. More strikingly, while the consumption response to income declines in the self-insurance model is 2/3 that to income increases, in the limited commitment model the consumption response does not change across income declines of different size, implying a 0 coefficient, while that to income rises has a value similar to that under self-insurance.[24] Moreover, skewness, while positive in both models, is an order of magnitude higher in the limited commitment model relative to both the data and the self-insurance model.

Rows 5 and 6 of table 2 provide results from a simulated method of moments estimation of the preference parameters β and σ, and of the variance of measurement error Var_μ, using a diagonally weighted minimum distance estimator and the 5 moments of table 2 as targets.[25] Interestingly, higher estimated risk aversion and a lower discount factor of 0.9 all but eliminate the non-linearity in the consumption response to income shocks in the self-insurance model. Non-linearities in the limited commitment model, however, remain unchanged even at very low estimated risk-aversion. Importantly, the estimated measurement error, whose variance is moderate relative to that of model-implied consumption levels but more important in growth rates, together with the estimated risk-aversion parameters brings the relative volatility of consumption and income growth in both models very close to the data. The increase in the variance of measured consumption growth also increases strongly the denominator of the skewness measure, and thus reduces its magnitude to close to 0 even in the limited commitment model. To the contrary, the impact of measurement error on the relative variance of consumption and income levels is modest. Thus, the limited commitment model continues to underpredict relative consumption dispersion in levels, while the overprediction of the self-insurance model is, with high risk aversion, little changed relative to the simple calibration even when including measurement error. The reason for this is, as mentioned above, that the effect of uncorrelated measurement error on growth rates, based on quarterly data, is much stronger than that on levels, measured as averages over four quarters. This also explains why, despite its strong impact on the variance and skewness of consumption growth, as a percentage of average annual levels implied by the model, measurement error is relatively small.

[24] Note how the regression coefficients for the self-insurance economy are an order of magnitude smaller than in Krueger and Perri (2005), who interpret their large effects as partial evidence against self-insurance. The smaller values here result from the permanent-plus-transitory structure of income shocks which strongly increases the average insurance properties of self-insurance.

[25] To avoid local minima, I estimate the model by solving it on a three-dimensional grid of $0.8 < \beta < 1$, $0.4 < \sigma < 10$ and Var_μ. I continue to only consider general equilibria where $R = 1.025$ and the median wealth to income ratio equals that in SCF data. I then choose those equilibrium parameter combinations that minimise the sum of squared moment deviations weighted by the inverse of the variance of each data moment.

The self-insurance model with measurement error and estimated preference parameters thus performs significantly better than its calibrated version in replicating the key features of the data identified in the non-parametric analysis, although the model continues to overpredict the relative variance of consumption and income levels. Moreover, while strong non-linearities are a robust feature under limited commitment, they disappear with higher estimated risk aversion in the self-insurance model. The evidence on the "shape of insurance" thus strongly argues in favour of a standard incomplete markets model, rather than a model where insurance markets are available but suffer from limited enforcement of contracts.

5 Sensitivity of the results

This section generalises the results presented above to different data sources and different specifications of the exogenous income process.

5.1 Alternative data sources

The CEX is the most important source of information on household consumption in the US economy. The PSID is generally considered to have more accurate information on incomes, but only provides data on food consumption. Recently, in a seminal contribution, Blundell, Pistaferri and Preston (2008, BPP) have imputed a series of non-durable consumption for the PSID by using its food expenditure information and a consumption demand function estimated on CEX data. The long time dimension of the resulting panel allowed them, under some assumptions, to identify the variances of permanent and transitory income shocks, as well as their impact on current consumption. Importantly, while they cannot reject perfect insurance against transitory shocks to income, they find evidence of excess smoothness, with only 2/3 of permanent income shocks translating into current consumption.

Table 3 compares the benchmark estimates of key moments in CEX data considered so far, to those that result from different variable definitions and alternative data sources. The first row of table 3 reports the benchmark results, which were based on a consumption series that includes a measure of service flows from key durables ("ND+"). Row 3 and 4 compare them to those for non-durables only ("ND") and food consumption. The differences between the different consumption measures are small relative to standard errors

Table 3

		$\frac{Var(C)}{Var(Y)}$	$\frac{Var(dc)}{Var(dy)}$	$skew(dc)$	$\widehat{b}_{dc,dy}\|_{dy<0}$	$\widehat{b}_{dc,dy}\|_{dy>0}$	$\widehat{b}_{dc,dy}$
1	CEXND+	0.418	0.508	0.178	0.074	0.065	0.095
2	St Error	0.012	0.019	0.120	0.014	0.015	0.008
3	CEXND	0.425	0.639	0.067	0.068	0.043	0.085
4	St Error	0.011	0.024	0.112	0.015	0.017	0.008
5	CEXFOOD	0.425	1.105	-0.159	0.078	0.049	0.076
6	St Error	0.011	0.043	0.083	0.019	0.024	0.011
7	CEX GK ND+					0.13	
8	CEX GK ND					0.11	
9	PSIDND	1.065	2.002	-0.089	0.182	0.175	0.182
10	St Error	0.032	0.082	0.193	0.034	0.032	0.017
11	PSIDFOOD	0.806	1.509	0.040	0.117	0.134	0.130
12	St Error	0.020	0.055	0.159	0.024	0.029	0.014

The table shows data moments for CEX consumption data on nondurables plus imputed services (Krueger and Perri 2006, row 1 with standard errors in row 2), nondurables (rows 3 and 4) and food consumption (rows 5 and 6). Rows 9 to 13 show the same moments for PSID data on nondurables as imputed by BPP, and food consumption. Rows 7 and 8 calculate the regression coefficient of consumption on income growth in CEX data adjusting for the different timing of consumption and income observations as in Gervais and Klein (2008). The standard errors for the data statistics are calculated using a bootstrap procedure with 400 repetitions.

for all statistics, apart from the relative variance of consumption growth. Interestingly, this is highest for food consumption. This is in line with difficulties to recall frequent but small purchases as an important source of error in the survey.

Rows 4 and 5 report estimates for $\widehat{b}_{dcdy}$, the regression coefficient of consumption on income changes, using the interpolation procedure by Gervais and Klein (2008) that corrects for the different timing of the consumption and income observations in the CEX.[26] The quarterly regression coefficients $\widehat{b}_{dcdy}$ based on their method are higher than the annual coefficients in rows 2 and 4, and similar to their original results despite the different sample selection. Again, the coefficient for non-durables is slightly smaller than that for a

[26] Gervais and Klein (2008) first use a GMM procedure to estimate a monthly AR(1) process for income that fits the information on annual individual income from interviews 2 and 5 of the CEX survey. They then show how the fitted values from this process can be used to consistently estimate a quarterly version of the regression coefficient $\widehat{b}_{dcdy}$. As in their paper, I use a longer sample for the interpolation, starting in 1980.

more comprehensive consumption measure.
Rows 5 and 6 report moments of the joint consumption-income distribution in a sample of the BPP dataset that, contrary to their paper, also includes households with female heads and those with changing household composition, to be comparable to the CEX benchmark sample. As is known from the original BPP results, the variance of consumption is higher relative to that of incomes than in CEX data. This is true for consumption levels but particularly for growth rates of consumption, which are twice as volatile as those of incomes. While BPP attribute the strong variance of consumption to imputation error in the new non-durable series, the relative variance of the original food consumption measure is also about twice as high as in the CEX. And food consumption growth remains substantially more volatile than income growth. The association between consumption and income growth is also twice as high in the BPP PSID data, especially for non-durable consumption. Similar to the evidence from the CEX, however, both the non-linearity and the skewness are small relative to the standard errors of the statistics.

Table 4 show the calibration and estimation results for the two risk-sharing models when the income process is estimated on the BPP income measure and the target moments are calculated on the basis of their measure of non-durable consumption. To understand the results, the first thing to notice is that, when estimated using PSID income data, the variances of both persistent ($var(m_t)$)and transitory components of income ($var(\nu_{it})$) are approximately two thirds the size of those estimated from the sample of CEX data. As the results for the calibrated models in rows 3 and 4 show, this lower estimated variance of income shocks reduces consumption dispersion in both models, and more than that of incomes under limited commitment, resulting in a lower value of the relative consumption-income variance in levels. Again, both models predict a relative dispersion of consumption growth that is an order of magnitude smaller than that for levels, in contrast to the data. For an unchanged equilibrium interest rate of 2.5 percent, the remaining moments in the calibrated self-insurance economy are largely unchanged. The limited commitment economy has the familiar non-linearity in the relationship of consumption and income growth. But with less volatile incomes, the response of consumption to income increases is smaller, and the skewness stronger, than in the benchmark estimation using CEX data. As table 3 showed, both levels and growth rates of the consumption measures in PSID data are more volatile relative to incomes than in the CEX, and relative consumption growth volatility is about twice as high as that of levels for the non-durable series. Moreover, contrary to the CEX, in the PSID both levels and growth rates are calculated from annual observations. So the variance of the error in measured growth rates is only twice that of the measurement error in annual levels. Taken together, this leads to higher estimates of measurement error in levels relative to the benchmark case. Particularly, as row 5 of table 4 shows, in the limited commitment model measurement error is estimated to have a variance almost four times higher than model-implied consumption in levels. This brings all measured statistics close to the data, apart from the regression coefficients $\widehat{b}_{dc,dy}|_{dy<0}$ and $\widehat{b}_{dc,dy}|_{dy>0}$, which are unaffected by measurement error, and unchanged despite a risk-aversion coefficient estimated to be somewhat below 1. So again, the non-linearity in the conditional distributions is a crucial statistic to evaluate the limited commitment model. In the self-insurance model, on the other hand, the estimated variance of measurement error in consumption is similar to that of consumption predicted by the model, and leads to measured relative variances that are about 20 percent too high in levels and 5 percent too low in growth rates. As in the limited commitment model, this reduces measured skewness to 0. Again, higher estimated risk-aversion of 4.75 all but eliminates the non-linearity in the self-insurance model, although the $\widehat{b}$ coefficients are estimated at only about half the value as in the data.

In summary, with higher relative variances of consumption and stronger consumption-income comovements, the moments from alternative data sources are generally closer to those that arise from simple self-insurance than the benchmark results based on CEX data. This reinforces the conclusion that, overall, the evidence provided in this study strongly favours the self-insurance model relative to one with limited commitment to contracts.

Table 4

		β	σ	$\frac{Var(C)}{Var(Y)}$	$\frac{Var(dc)}{Var(dy)}$	$skew(dc)$	$\widehat{b}_{dc,dy}\|_{dy<0}$	$\widehat{b}_{dc,dy}\|_{dy>0}$	$\frac{Var_\mu}{Var(C^{mod})}$
1	PSID data			1.07	2.02	-0.09	0.18	0.18	
2	St Error			0.032	0.082	0.193	0.034	0.032	
3	LC Calib	0.963	1.00	0.19	0.01	4.63	0.00	0.07	
4	SI Calib	0.968	1.00	0.57	0.05	0.81	0.07	0.11	
5	LC Estim	0.967	0.67	1.01	2.25	0.00	0.00	0.07	3.94
6	SI Estim	0.906	4.75	1.24	1.89	0.00	0.08	0.09	1.21

The table presents moments of PSID data (row 1), in the calibrated models (rows 2 and 3) with limited commitment ("LC") and self-insurance ("SI") as well as estimated versions that choose β, σ and Var_μ to target an equilibrium interest rate $R = 1.025$ the median wealth to income ration and the 4 distributional moments (rows 4 and 5). The standard errors for the data statistics are calculated using a bootstrap procedure with 400 repetitions.

5.2 Alternative Income Processes

Table 5

		β	σ	$\frac{Var(C)}{Var(Y)}$	$\frac{Var(dc)}{Var(dy)}$	$skew(dc)$	$\widehat{b}_{dc,dy}\|_{dy<0}$	$\widehat{b}_{dc,dy}\|_{dy>0}$	$\frac{Var_{\mu}}{Var(C^{mod})}$
1	CEX data			0.42	0.50	0.18	0.07	0.06	
2	St Error			0.012	0.019	0.120	0.014	0.015	
3	LC Estim	0.960	0.47	0.27	0.55	0.00	0.00	0.17	0.19
4	SI Estim	0.871	4.25	0.53	0.57	0.02	0.13	0.20	0.09
5	PSID data			1.07	2.02	-0.09	0.18	0.18	
6	St Error			0.032	0.082	0.193	0.034	0.032	
7	LC Estim	0.971	0.40	0.98	2.53	0.00	0.01	0.08	5.86
8	SI Estim	0.918	3.75	1.09	1.86	0.00	0.12	0.18	1.31

The table presents moments of CEX data (row 1, with standard errors in row 2) and PSID data (row 5, with standard errors in row 6), and in models that are estimated to meet the data moments with limited commitment ("LC", row 3 and 5) and self-insurance ("SI", row 4 and 6) for a lower persistence of the income process equal to $\rho = 0.9$, as opposed to the benchmark value of $\rho = 0.9989$.

Previous studies have found the performance of risk-sharing models to depend on the specification of the income process, in particular the persistence parameter ρ. Table 5 presents the results when the models are estimated on both CEX data (rows 1 to 4) and PSID data (rows 5 to 8) with a reduced value of income persistence equal to $\rho = 0.9$. To understand how this changes the results in the self-insurance model, it is worth pointing out how the relative importance of persistent and transitory shocks changes with ρ for an unchanged variance-covariance of income observed in data. Particular, for a given level of persistence and autocovariance, the variance of shocks to the persistent component equals

$$\sigma_{\nu}^2 = \frac{1-\rho^2}{\rho} cov(x_t, x_{t-1}) \tag{24}$$

So the variance of persistent shocks is decreasing in ρ, while that of transitory shocks σ_{ε}^2 is increasing in ρ. Moreover, for high levels of ρ, income growth is, approximately, simply the sum of persistent and transitory shocks. In other words, conditional on unchanged variance and covariance of x_t, the contribution of persistent shocks to income growth falls as ρ rises. This implies, somewhat counterintuitively, that the response of consumption growth to income changes <u>increases</u> as income becomes less persistent. This effect can be seen very clearly in rows 4 and 8 in table 5. The main difference with respect to earlier results is that, in the self-insurance model with lower persistence, consumption responses to income changes are larger and more non-linear: the response of consumption growth to income increases is now more than 50% larger than that to income falls in both estimations. The higher average value of the $\widehat{b}$ coefficients is further away from CEX data, but significantly closer to the values found in the PSID. Cross-sectional dispersion is smaller when incomes are less persistent, despite slightly greater estimated measurement error, in line with the intuition from permanent income models. The relative variance of consumption nonetheless still exceeds that observed in CEX data, where measurement error is therefore again estimated to be small. As before, with measurement errror of a similar magnitude as that of model-implied consumption levels, the model fits the relative variances in PSID data extremely well.

In the limited commitment model, lower persistence reduces the attractiveness of autarky for high-income individuals who expect their high income to last less long. This makes participation constraints less binding, and thus leads to an increase in insurance in calibrated models. In the results of table 5, this is offset through a rise in estimated measurement error (whose variance is now almost 6 times that of model-implied consumption in row 5) and a fall in risk aversion for the estimation on PSID data (where the estimate for σ is now on the lower bound of the parameter space). The counterfactual non-linearities of the limited commitment model, however, are even stronger at lower persistence, particularly in the model estimated on CEX data.

6 Conclusion

This article has shown how the joint distribution of consumption, income and wealth contains important information about consumption-smoothing models beyond its second-moment properties. Particular, I have analysed two standard models where limited risk-sharing results either from exogenous incompleteness of financial markets, or from the inability of households to commit to insurance contracts. Both models replicate similarly well, although not perfectly, standard insurance measures in US micro-data, namely the relative variances of consumption and income or the average comovement of their growth rates. Nevertheless, their joint distributions are strongly counterfactual. Specifically, both models exhibit counterfactually low variances of consumption growth and asymmetries not present in the data, such as declining conditional variances of consumption along the income distribution, positive skewness of consumption growth, and stronger responses to income rises than to income falls. Estimation of the models via a simulated method of moment procedure shows how measurement error in consumption improves the performance of both models in terms of relative variances of consumption and the skewness of its growth rate. Moreover, the asymmetry all but vanishes in the self-insurance model at higher estimated risk-aversion. The same is not true for the limited commitment model, where the relationship between consumption and income growth remains highly non-linear even at low estimated risk-aversion. The "shape of insurance" thus contains important information to discriminate between models of consumption risk-sharing. Particularly, although the self-insurance model continues to predict cross-sectional variances of consumption that are somewhat higher than in CEX data, its overall performance is better than that of its limited commitment counterpart, although both models perform worse when incomes are less persistent. The comparison with BPP's PSID dataset, where consumption is more volatile than in the CEX and its association with income movements is stronger, reinforces this conclusion.

Future research should generalise these results in at least two directions: first, a similar "non-parametric-to-parametric" approach should be applied to alternative models of consumption-smoothing. For example, it would be interesting to perform a similar analysis for the Attanasio and Pavoni (2007) model, where imperfect risk-sharing arises due to information, rather than contracting, frictions. Second, we need further documentation of the relative structure of different datasources on consumption and income in the US. Particularly, there seem to be important differences between key moments in the two main datasources, CEX and PSID, that call for further analysis. In this context, it would also be interesting to include measurement error in income in the estimation. To the extent that measurement error in income decreases the observed relative variance of consumption and attenuates estimates of $\widehat{b}$, it should, if anything, reinforce the conclusion that the self-insurance model does well in replicating key features of US micro-data, by bringing its degree of insurance closer to that observed in the data. But a more systematic study of this issue would be very useful.

7 References

1. Ábrahám, Árpád and Cárceles-Poveda, Eva (2006). "Endogenous incomplete markets, enforcement constraints, and intermediation", Theoretical Economics Vol. 1 No. 4, pp. 439-459.

2. Abreu, Dilip, David G. Pearce and Ennio Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", Econometrica, Vol. 58 No. 5, pp. 1041-1063.

3. Aguiar, Mark A. and Mark Bils (2011). "Has Consumption Inequality Mirrored Income Inequality?", NBER Working Paper 16807.

4. Ahmed, Naeem, Matthew Brzozowski and Thomas F. Crossley (2006). "Measurement Errors in Recall Food Consumption Data", IFS Working Paper 06/21.

5. Albarran, Pedro and Orazio P. Attanasio (2003). "Limited Commitment and Crowding out of Private Transfers: Evidence from a Randomised Experiment", Economic Journal Vol. 113 No. 486, C77-C85.

6. Aiyagari, S. Rao (1993). "Uninsured idiosyncratic Risk and Aggregate Saving", Federal Reserve Bank of Minneapolis Working Paper 502.

7. Aiyagari, S. Rao (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving," Quarterly Journal of Economics, Vol. 109 No. 3, pp. 659-84.

8. Altonji, Joseph G and Aloysius Siow (1987). "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data", The Quarterly Journal of Economics, Vol. 102 No. 2, pp. 293-328.

9. Alvarez, Fernando and Urban J. Jermann (2000). "Efficiency, Equilibrium, and Asset Pricing with Risk of Default", Econometrica, Vol. 68 No. 4, 775-798.

10. Atkeson, Andrew (1991). "International Lending with Moral Hazard and Risk of Repudiation", Econometrica Vol. 59 No. 4, 1069-1089.

11. Atkeson, Andrew and Robert E. Lucas (1992). "On Efficient Distribution with Private Information", Review of Economic Studies, Vol. 59 No. 3, 427-53.

12. Atkeson, Andrew and Robert E. Lucas (1995). “Efficiency and Equality in a Simple Model of Unemployment Insurance”, Journal of Economic Theory, Vol. 66 No. 1, 64-88.

13. Attanasio, Orazio P. and Steven J. Davis (1996). “Relative Wage Movements and the Distribution of Consumption”, Journal of Political Economy, Vol. 104 No. 6, 1227-62.

14. Attanasio, Orazio P. and Nicola Pavoni (2007). “Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption”, NBER Working Papers 12994.

15. Attanasio, Orazio P. and Guglielmo Weber (2010). “Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy”, Journal of Economic Literature Vol. 48 No. 3, pp. 693-751.

16. Barr, Abigail and Garance Genicot (2008). “Risk Sharing, Commitment, and Information: An Experimental Analysis”, Journal of the European Economic Association, Vol. 6 No. 6, 1151-1185.

17. Battistin, Erich, Richard Blundell and Arthur Lewbel (2007). “Why is consumption more log normal than income? Gibrat’s law revisited”, IFS Working Papers W07/08, Institute for Fiscal Studies.

18. Bloise, Gaetano, Pietro Reichlin, and Marco Tirelli (2009). “Indeterminacy of Competitive Equilibrium with Risk of Default”. CEPR Discussion Paper No. 7477.

19. Blundell, Richard, Luigi Pistaferri and Ian Preston (2008). “Consumption Inequality and Partial Insurance”, American Economic Review, Vol. 98 No. 5, 1887-1921.

20. Botev, Z. I., J. F. Grotowski and D. P. Kroese (2008). “Kernel density estimation via diffusion”, University of Queensland, mimeo.

21. Broer, Tobias (2009a). “Heterogeneous Individuals in the International Economy”, PhD Thesis, European University Institute.

22. Broer, Tobias (2009b). “Stationary equilibrium distributions under limited commitment”, EUI ECO Working Paper 2009/39.

23. Broer, Tobias (2010). "Crowding out and crowding in: When does redistribution improve risk-sharing in limited commitment economies? ". Journal of Economic Theory, Vol 146 No 3, pp. 957-975.

24. Campbell, John Y. and Angus Deaton (1989). "Why Is Consumption So Smooth?", Review of Economic Studies Vol. 56 No. 3, pp. 357-73.

25. Campbell, John Y. and N. Gregory Mankiw (1990). "Permanent Income, Current Income, and Consumption", Journal of Business and Economic Statistics Vol. 8 No. 3, pp. 265-79.

26. Campbell, John Y. and N. Gregory Mankiw (1991). "The Response of Consumption to Income: A Cross-Country Investigation", European Economic Review Vol. 35 No. 4, pp. 723-756.

27. Cogley, Timothy (2002), "Idiosyncratic risk and the equity premium: evidence from the consumer expenditure survey", Journal of Monetary Economics Volume 49, No. 2, pp. 309-334.

28. Cordoba, Juan-Carlos (2008). "U.S. inequality: Debt constraints or incomplete asset markets?", Journal of Monetary Economics, Vol. 55 No.2, 350-364.

29. Dardanoni, Valentino (1995). "Income distribution dynamics: monotone Markov chains make light work", Social Choice and Welfare, Vol. 12, 181-192.

30. Deaton, Angus (1992). *Understanding Consumption*, Oxford University Press, New York.

31. Dynan, Karen E., Douglas W. Elmendorf and Daniel E. Sichel (2006). "Financial innovation and the Great Moderation: what do household data say?", Proceedings, Federal Reserve Bank of San Francisco.

32. Eozenou, Patrick (2009). "Optimal Risk Sharing Under Limited Commitment: Evidence From Rural Vietnam", MPRA Paper 12688, University Library of Munich, Germany.

33. Guvenen, Fatih (2009) "An Empirical Investigation of Labor Income Processes", Review of Economic Dynamics Vol. 12 No.1, 5879.

34. Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1995), "Precautionary Saving and Social Insurance", The Journal of Political Economy, 103 No. 2, 360399.

35. Hall, Robert E. and Mishkin, Frederic S. (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households", Econometrica Vol. 50 No. 2, 461-81.

36. Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante (2009). "Consumption and Labor Supply with Partial Insurance: An Analytical Framework", NBER Working Paper 15257.

37. Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante (2010). "The Macroeconomic Implications of Rising Wage Inequality in the United States", Journal of Political Economy Vol. 118 No. 4, pp. 681-722.

38. Heathcote, Jonathan, Fabrizio Perri and Giovanni L. Violante (2008). "Cross-Sectional Facts for Macroeconomists: United States (1967-2006)", mimeo.

39. Huggett, Mark (1993). "The risk-free rate in heterogeneous-agent incomplete-insurance economies", Journal of Economic Dynamics and Control Vol. 17 No. 5-6, pp. 953-969.

40. Kehoe, Timothy J. and David K. Levine (2001). "Liquidity constrained markets versus debt constrained markets", Econometrica, Vol. 69 No. 3, 575-598.

41. Kehoe, Timothy J. and David K. Levine (1993). "Debt-Constrained Asset Markets", Review of Economic Studies, Vol. 60 No. 4, 865-888.

42. Kocherlakota, Narayana R. (1996). "Implications of Efficient Risk Sharing without Commitment", Review of Economic Studies, Vol. 63 No. 4, 595-609.

43. Krueger, Dirk and Fabrizio Perri (2009). "How do Households Respond to Income Shocks?", mimeo.

44. Krueger, Dirk and Fabrizio Perri (2006). "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory", Review of Economic Studies, Vol. 73 No. 1, 169-99.

45. Krueger, Dirk and Fabrizio Perri (2011). “Public or Private Risk Sharing”, Journal of Economic Theory, Vol. 146 No. 3, 920-956.

46. Krueger, Dirk und Harald Uhlig (2006). “Competitive risk sharing contracts with one-sided commitment”, Journal of Monetary Economics, Vol. 53 No. 7, 1661-1691.

47. Krusell, Per and Anthony A. Smith Jr. (1998). “Income and Wealth Heterogeneity in the Macroeconomy”, Journal of Political Economy, Vol. 106 No. 5, 867-896.

48. Lazcó, Sarolta (2011), “On Risk Sharing in Village Economies: Structural Estimation and Testing”, mimeo.

49. Ligon, Ethan, Jonathan P. Thomas and Tim Worrall (2002). “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies”, Review of Economic Studies, Vol. 69 No. 1, 209-44.

50. Ligon, Ethan, Jonathan P. Thomas and Tim Worrall (1996),“Informal Insurance Arrangements in Village Economies” (University of Warwick, mimeo).

51. Livshits, Igor, James MacGee and Michèle Tertilt (2007). “Consumer Bankruptcy: A Fresh Start”, American Economic Review, Vol. 97 No. 1, pp. 402-418.

52. MaCurdy, Thomas E. (1982), The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis, Journal of Econometrics, January , 18 No.1, 83114.

53. Marcet, Albert and Ramon Marimon (2009). “Recursive Contracts”, mimeo.

54. Phelan, Christopher (1995). “Repeated Moral Hazard and One-Sided Commitment,” Journal of Economic Theory Vol. 66(2), 488-506.

55. Ríos-Rull, José-Victor (1995).“Models with Heterogeneous Agents”, in Cooley , Tom (ed.): *Frontiers of Business Cycles Research*, Princeton University Press.

56. Runkle, David E. (1991). “Liquidity constraints and the permanent-income hypothesis: Evidence from panel data”, Journal of Monetary Economics Vol. 27 No. 1, 73-98.

57. Storesletten, Kjetil, C. Telmer, and A. Yaron (2004). “Consumption and Risk Sharing over the Life-Cycle”, Journal of Monetary Economics, Vol. 51 No. 3, 609633.

58. Tauchen, George and Robert Hussey (1991). “Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models”, Econometrica, Vol. 59 No. 2, 371-396.

59. Thomas, Jonathan and Tim Worrall (1994). “Informal Insurance Arrangements in Village Economies”, Liverpool Research Papers in Economics and Finance No. 9402.

60. Thomas, Jonathan and Tim Worrall (1988). “Self-enforcing Wage Contracts”, Review of Economic Studies, Vol. 55 No. 4, 541-54.

61. Thomas, Jonathan and Tim Worrall (2007). “Unemployment insurance under moral hazard and limited commitment: public versus private provision”, Journal of Public Economic Theory Vol. 9 No. 1, 151-181.

62. Townsend, Robert M (1994). “Risk and Insurance in Village India”, Econometrica, Vol. 62 No. 3, 539-91.

63. Uhlig, Harald (1996). “A Law of Large Numbers for Large Economies”, Economic Theory, Vol. 8 No. 1, 41-50.

64. Zeldes, Stephen P. (1989). “Consumption and liquidity constraints: an empirical investigation”, Journal of Political Economy, Vol. 97 No. 2, 305-346.

8 Appendix: Sample and Data Description

8.1 Sample selection

CEX: As in Krueger and Perri (2006, hf. KP), I use a sample of complete income respondents between 20 and 65 years of age, and exclude: rural households, households with only food consumption, average hourly earnings or the household head below half the minimum, with no labour income or less then 1 week worked in either year before the income observation, with no hours worked but positive earnings for either spouse or the household head, with non-positive food consumption or only food consumption.

PSID: BPP consider only continuously married households with a male head aged between 30 and 65, and neglect the PSID's SEO low-income sample. Using their original dataset, I also include households with heads aged 25 to 29, female heads, and those not continuously married.

8.2 Comparison of sample sizes and data definitions

Table A1: Data comparison

Dataset	Year	Sample size	Frequency	Def: Y	Def: C
CEX KP 06	1999 - 2004	10827 HH	quarterly	LEA+	ND,ND+
PSID BPP 08	1980 - 1992	2329 HH (20484 obs)	annual	LEA+	Food, NDBPP

LEA+: Wages and salary plus a fraction (0.864) of self-employment income plus transfers minus taxes deflated by annual CPI (see KP). LEA: Sum of labor income and transfers, such as welfare payments, minus taxes paid (BPP p. 1894).

ND: Expenditure on nondurables, services, and small durables (such as household equipment), deflated by category-specific CPI indices (see KP, p. 165). ND+: Expenditure on all items plus rental equivalent of housing and quarterly service flow from vehicles (1/32 times the purchase price calculated in line with Cutler and Katz (1991), deflated by category-specific CPI indices (see KP p. 165).

Food: Sum of expenditure on food at home and away from home. NDBPP: Sum of expenditure on food, alcohol, tobacco, other nondurables (services, heating fuel, public and private transport including gasoline), personal care, and semidurables (clothing and footwear) (see BPP p. 1893).

9 Figures

Figure 1: Joint distribution of consumption and income

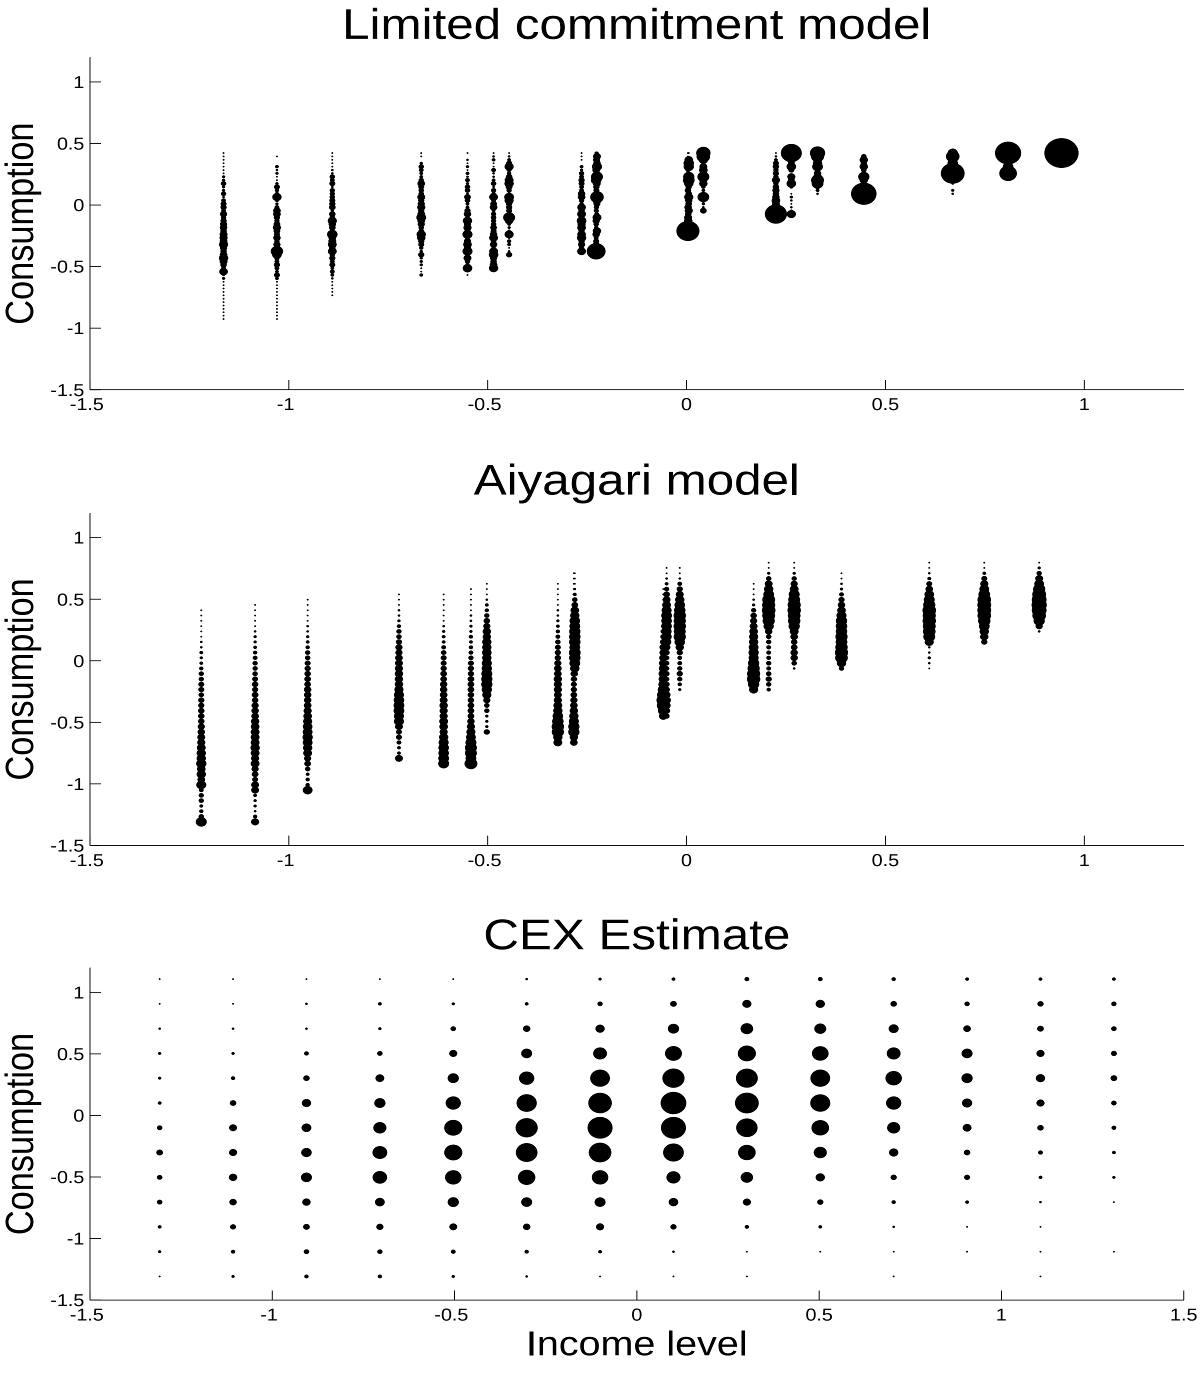

The figure shows histogrammes of the joint distributions of (the logarithms of) consumption and income in the limited commitment economy and a simple Aiyagari economy, together with a kernel density estimate of the empirical distribution in CEX data. The size of dots is proportional to the frequency mass at the corresponding point. The kernel density is estimated using an optimal bandwith (Botev et al 2008), and is based on residuals from a first-stage regression of the variables on observable individual characteristics as described in the main text.

Figure 2: Conditional mean and variance function of consumption levels

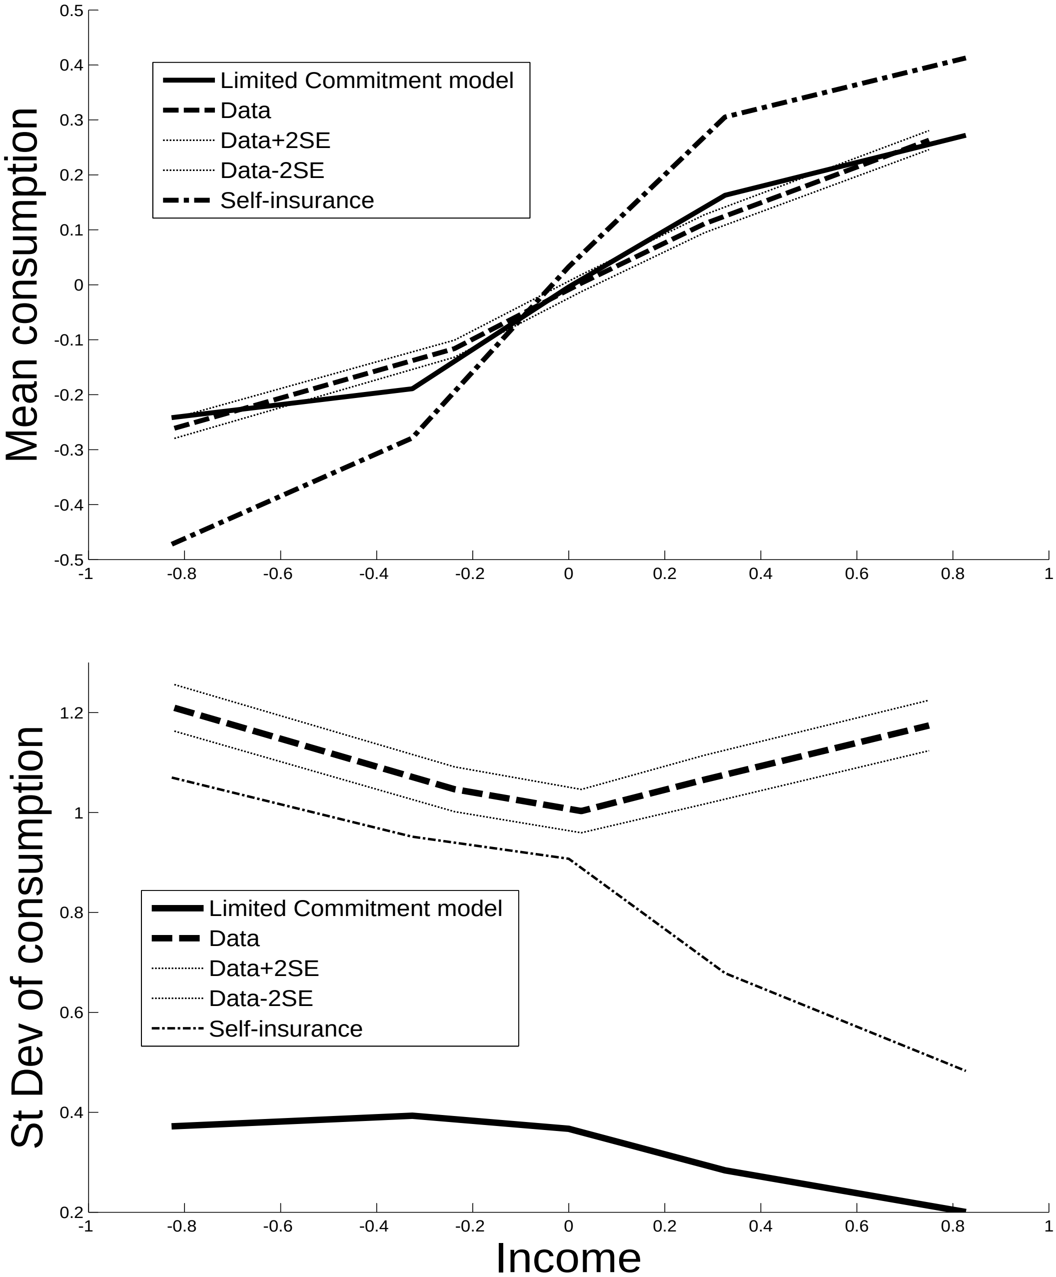

For every decile of the income distribution, the figure shows the mean (top panel) and variance (bottom panel) of (the logarithms of) consumption from the limited commitment model (solid lines), the self-insurance economy (dotted lines) and CEX data (dashed lines). The standard errors for the data statistics are calculated using a bootstrap procedure with 400 repetitions.

Figure 3: Joint distribution of consumption and income growth

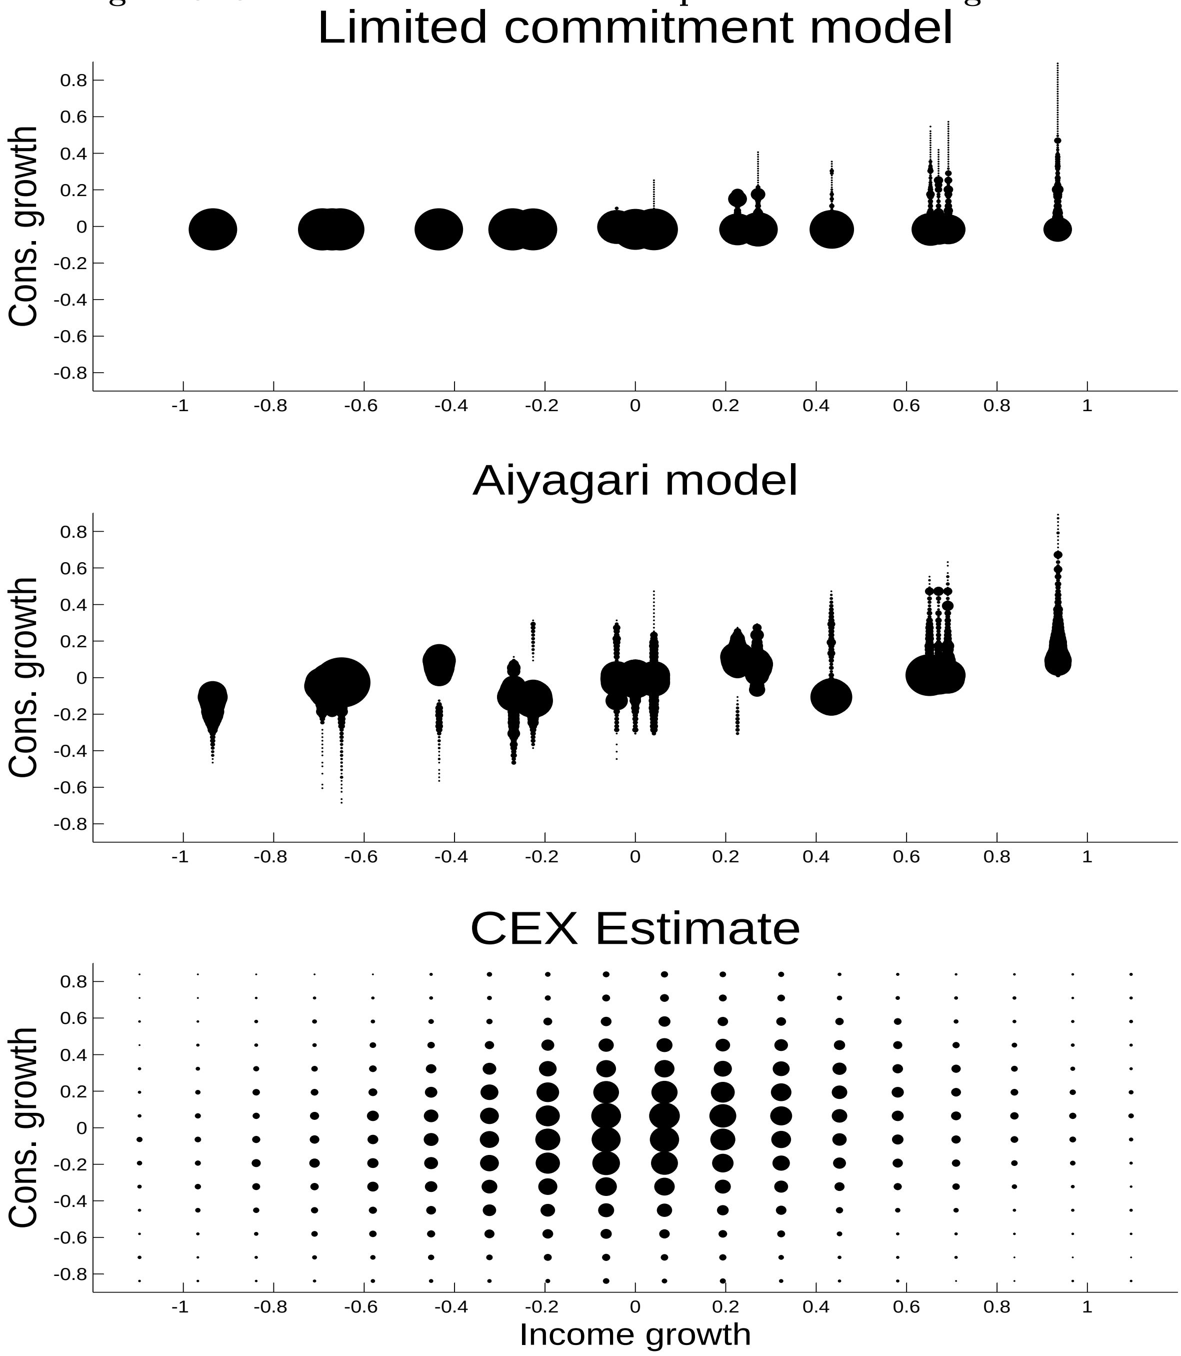

The figure shows histogrammes of the joint distributions of consumption and income growth in the limited commitment economy and a simple Aiyagari economy, together with a kernel density estimate of the empirical distribution in CEX data. The size of dots is proportional to the frequency mass at that point. The kernel density estimate of the empirical distribution uses an optimal bandwith (Botev et al 2008), and is based on differences in the raw data.

Figure 4: Conditional mean and variance function of consumption growth

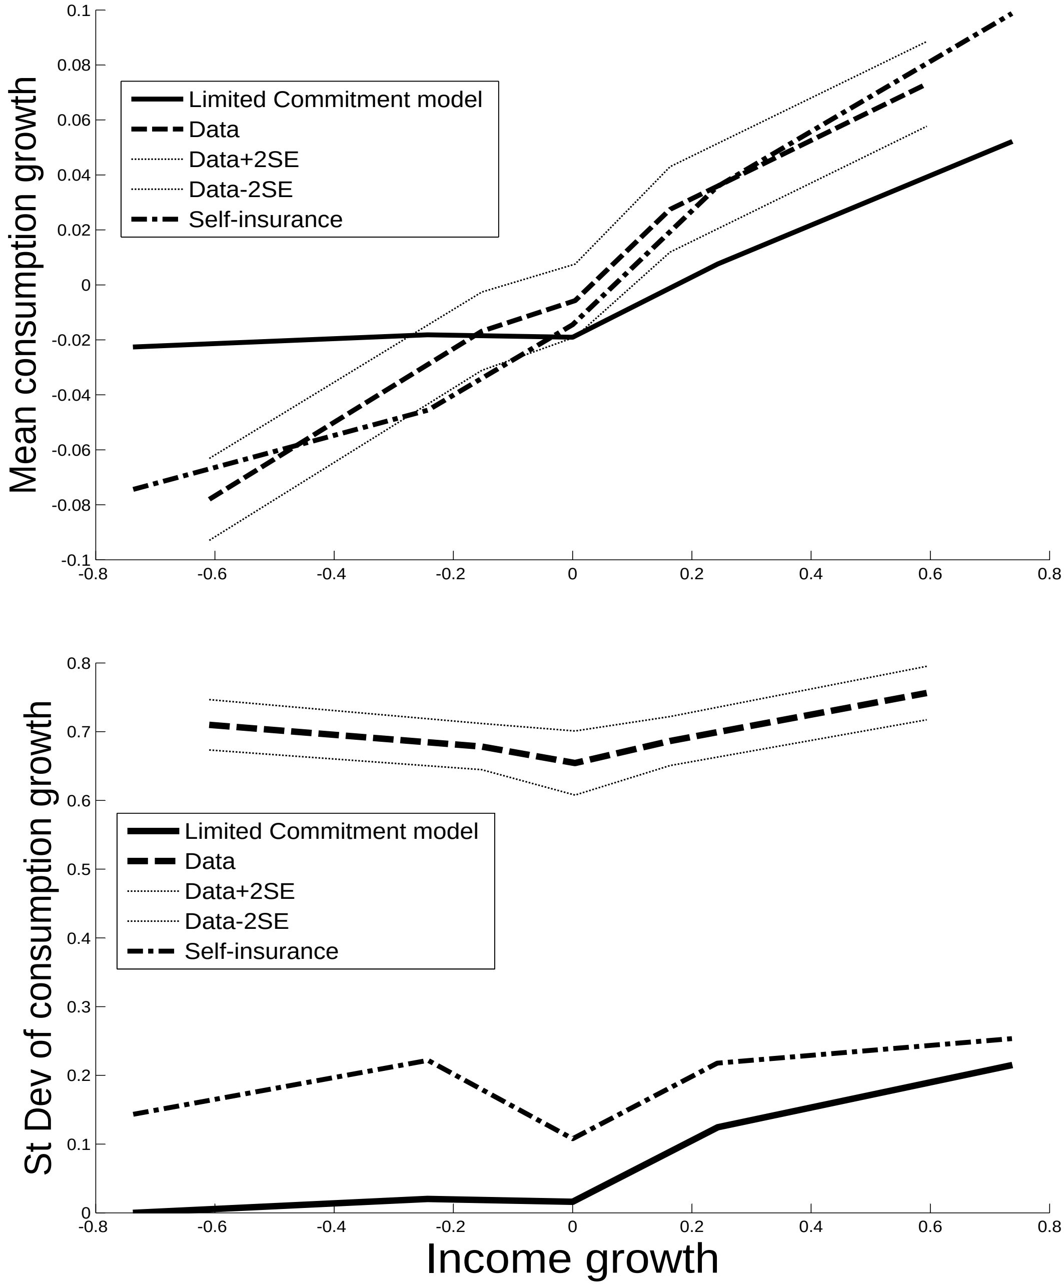

For every decile of the income growth distribution, the figure shows the mean (top panel) and standard devations (bottom panel) of consumption growth from the limited commitment model (solid lines), the self-insurance economy (dotted lines) and CEX data (dashed lines). The standard errors for the data statistics are calculated using a bootstrap procedure with 400 repetitions.

Figure 5: Joint distribution of wealth and income

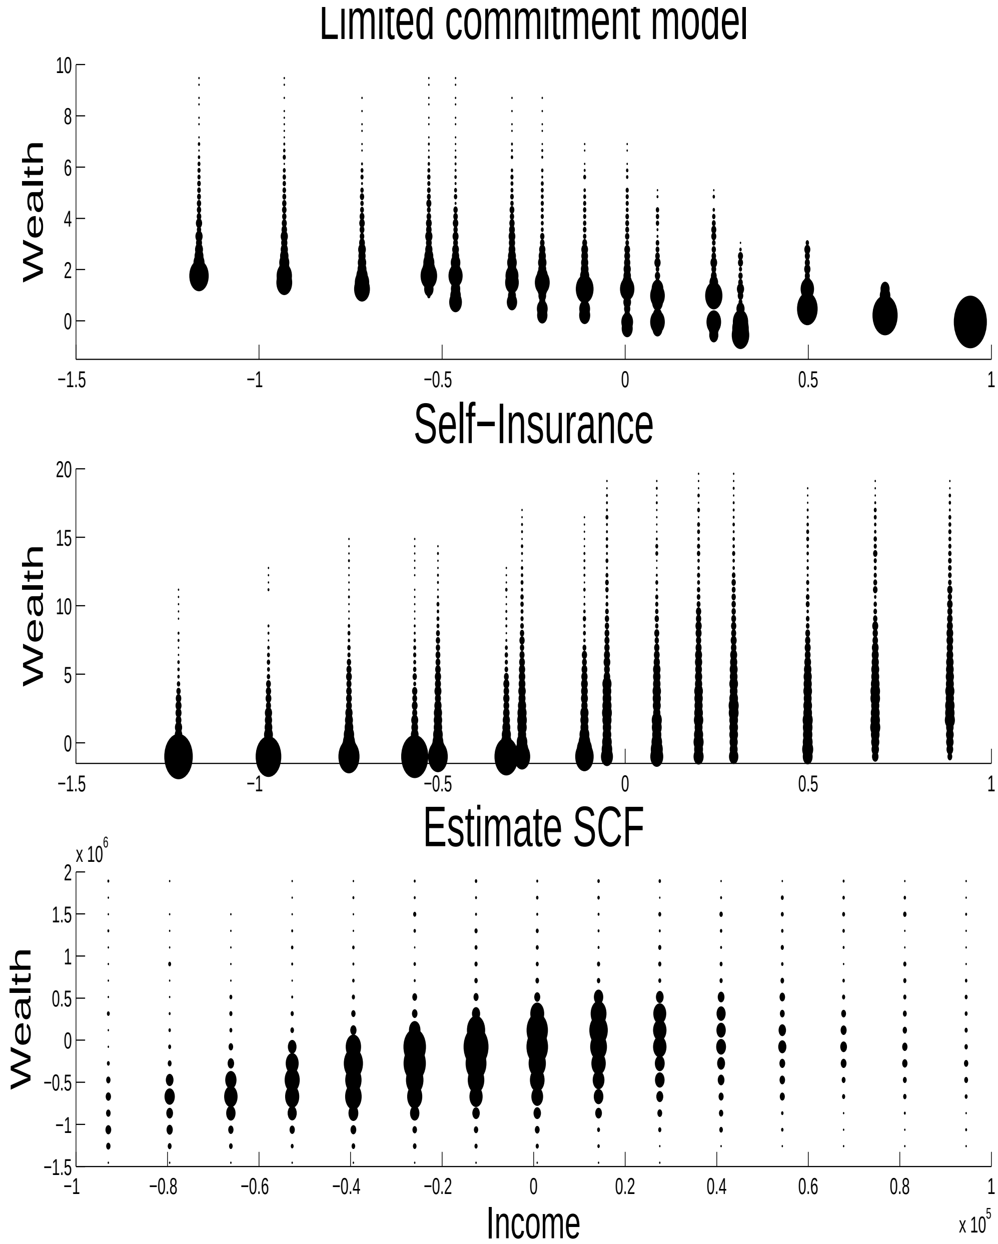

The figure shows histogramm estimates of the joint distributions of (the levels of) wealth and income in the limited commitment and self-insurance economies, together with a kernel density estimate of the empirical distribution in SCF data. The size of dots is proportional to the frequency mass at that point. The kernel density is estimated using an optimal bandwith (Botev et al 2008), and is based on residuals from a first-stage regression of the variables on observable individual characteristics as described in the main text.

Figure 6: Conditional mean and variance function of wealth

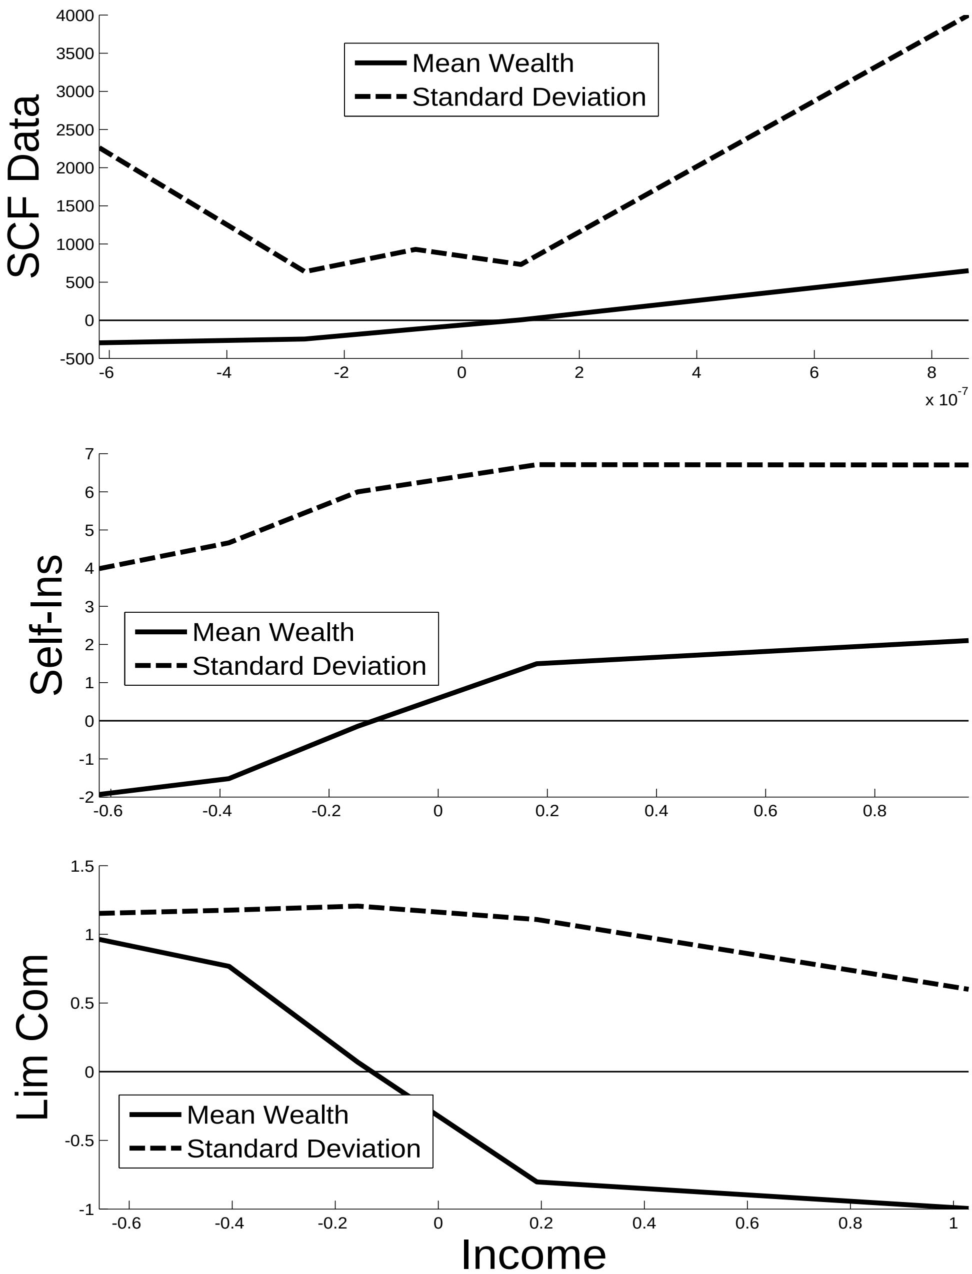

For every decile of the income distribution, the top panel shows the mean and standard deviation of net worth in SCF data. The mean and standard deviation from the limited commitment model (solid lines) and the self-insurance economy (dotted lines) are depicted in the central and bottom panels, respectively. Both income and wealth distributions are demeaned before plotting. The figure shows no standard errors for the data statistics because of the highly variant sampling probabilities in the SCF that preclude a simple bootstrap procedure.