

Understanding Consumption Smoothing: Evidence from the US Consumer Expenditure Data*

Dirk Krueger

Goethe University Frankfurt,

University of Pennsylvania,

CEPR and NBER

dirk.krueger@wiwi.uni-frankfurt.de

Fabrizio Perri

New York University,

Federal Reserve Bank of Minneapolis,

CEPR and NBER

fperri@stern.nyu.edu

November 2004

Abstract

Consumption models with endogenous debt constraints differ from standard incomplete markets models in their predictions about an individual household's ability to smooth consumption across time and states of the world. In this paper we develop these differences, both theoretically and quantitatively. We then use data from the US Consumer Expenditure Survey (CE) to assess along which dimensions the predictions of these models are consistent with the empirical evidence. We find that both type models fail to fully account for the data and argue that a model that combines aspects of both might be more successful.

KEYWORDS: Risk Sharing, Insurance, Incomplete Markets, Limited Enforcement, Earning Shocks, Asymmetries

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

*An earlier and substantially different version of this paper circulated under the title "Risk Sharing in Economies with Incomplete Markets". We thank an anonymous referee and seminar participants at the University of Chicago, Stanford University, the 2004 SED and EEA meetings for helpful comments. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Krueger acknowledges financial support from the NSF under grant SES-0004376.

1 Introduction

If complete insurance markets were available to households they could insulate their consumption profile from income risk. This prediction is commonly rejected empirically. In response to this finding a number incomplete risk sharing models have been developed. In many models imperfect risk sharing arises since agents can only trade a single, uncontingent bond and face borrowing constraints.¹ We refer to this setup as the standard incomplete markets model (SIM). Another class of models assume that a full set of state contingent contracts is available to all agents, but that intertemporal contracts can only be legally enforced by exclusion from future intertemporal trade.² Since exclusion from credit markets is not infinitely costly in some states of the world agents might find optimal not to repay their debts and go into autarky. This possibility endogenously restrict the extent to which each contingent asset can be traded and thus limits risk sharing. We refer to this model as the debt constrained markets model (DCM).

In this paper we argue that both types of incomplete risk-sharing models have different, empirically testable implications for the response of consumption growth to idiosyncratic income fluctuations. We demonstrate that a reasonably parameterized version of the SIM model predicts a substantial deviation from perfect consumption insurance, whereas the DCM model predicts a modest deviation. We compare these implications with the United States Consumer Expenditure Survey (CE) data and find that the DCM model does a better job in explaining the data.

A second dimension along which both models differ in their empirical implications is the asymmetric response of consumption to income shocks. Both models predict that consumption growth responds stronger to positive than to negative income growth rate shocks of similar magnitude, but these responses are quantitatively and qualitatively different in both models. Here we discuss why both models predict these asymmetries and then investigate whether the data bear out the asymmetries as predicted by both models; we find that in the data the response is quite symmetric so neither model seems to be able to account for the data.

In sections 2, 3 and 4 we briefly review the two models, characterize some features of the allocations and discuss the choice of parameters. Our results are contained in section 5 and section 6 concludes.

¹See Deaton (1991) or Aiyagari (1994), among others.

²See for example Kehoe and Levine (1993) or Alvarez and Jermann (2000).

2 The Environment

There is a continuum of identical infinitely lived consumers who get period utility $u(\cdot)$ from consuming a single non storable good. Let β be the discount factor. Each individual has a stochastic endowment process and we denote by y_t the current period endowment and by $y^t = (y_0, \dots, y_t)$ the history of realizations of endowment shocks. y_t follows a stationary Markov process with transition probabilities $\pi(y'|y)$ and unique invariant measure $\Pi(\cdot)$. We use the notation $y^s|y^t$ to mean that y^s is a possible continuation of endowment shock realization y^t . At date 0 agents are distinguished by their initial asset holdings a_0 and by their initial shock y_0 . A consumption allocation $\{c_t(a_0, y^t)\}$ specifies how much to an agent of type (a_0, y_0) consumes, conditional on having experienced history y^t .

2.1 Market Structures

The two incomplete markets models whose predictions we will test differ along two dimensions: the set of assets that can be traded and the short-sale constraints these trades are subject to. Within the SIM model agents face budget constraints of the form

$$c_t(a_0, y^t) + \frac{a_{t+1}(a_0, y^t)}{R_t^{in}} = a_t(a_0, y^{t-1}) + y_t$$

where R_t^{in} is the gross interest rate on one-period risk-free bonds. Agents face borrowing constraints of the form $a_{t+1}(a_0, y^t) \geq -\bar{A}$, where \bar{A} is a fixed parameter.

In the DCM model consumers can trade a full set of state contingent commodities. Their budget constraints is

$$c_t(a_0, y^t) + \sum_{y_{t+1}} q_t(a_0, y^t, y_{t+1}) a_{t+1}(a_0, y^t) = a_t(a_0, y^t) + y_t$$

where $q_t(a_0, y^t, y_{t+1})$ is the price of a claim to consumption in period $t + 1$ for a household with characteristics (a_0, y^t) , conditional on shock y_{t+1} being realized.

The feature that separates the DCM model from a complete markets model is the presence of individual rationality (or “debt constraints” as Kehoe and Levine (1993) called them): consumers have the option of opting out of intertemporal trade and renege on their debt. The only punishment for doing so, and hence the only enforcement mechanism for contracts, is that agents that choose to default on their contracts are banned from future intertemporal trade. This assumption together with the existence of the full set of securities impose the following set of enforcement constraints

on equilibrium allocations

$$\begin{aligned}
& u(c(a_0, y^t)) + \sum_{s>t} \sum_{y^s|y^t} \beta^{s-t} \pi(y^s|y^t) u(c(a_0, y^s)) \\
\geq & u(y_t) + \sum_{s>t} \sum_{y^s|y^t} \beta^{s-t} \pi(y^s|y^t) u(y_s) \equiv U^{Aut}(y_t) \quad \forall y^t
\end{aligned} \tag{1}$$

2.2 Recursive Formulation

For the DCM model one can exploit duality theory to characterize allocations as solutions to a cost minimization problem of a social planner that faces a constant interest rate R and has to deliver lifetime utility w to the agent. The planner solves

$$\begin{aligned}
V(w, y) = & \min_{c \geq 0, \{g(y')\}} \left\{ c + \frac{1}{R} \sum_{y' \in Y} \pi(y'|y) V(g(y'), y') \right\} \\
\text{s.t.} &
\end{aligned} \tag{2}$$

$$w = (1 - \beta)u(c) + \beta \sum_{y' \in Y} \pi(y'|y) g(y') \tag{3}$$

$$g(y') \geq U^{Aut}(y') \quad \forall y' \in Y \tag{4}$$

where $V(w, y)$ is the resource cost for the planner to provide an individual with expected utility w when the household's endowment is y and $\frac{1}{R}$ is the relative price of future resources in terms of current resources. The cost consists of the cost of consumption c today and expected cost from tomorrow on, $\sum_{y'} \pi(y'|y) V(g(y'), y')$, discounted to today, where $g(y')$ represents the utility promises received by the household next period if its future income is y' . The promise-keeping constraint (3) ensures that an individual entitled to w in fact receives utility w through the allocation $c, \{g(y')\}_{y' \in Y}$. The participation constraints (4) state that the social planner for each state tomorrow has to guarantee individuals an expected utility promise at least as high as obtained with the autarkic allocation $U^{Aut}(y_t)$.

The definition of a recursive competitive equilibrium and a stationary equilibrium is standard in both models and hence omitted.³

³For the standard incomplete markets model, see Huggett (1993) or Aiyagari (1994), for the endogenous incomplete markets model, see Krueger and Perri (1999).

3 Qualitative Features of the Models

3.1 Standard Incomplete Markets

It is in general hard to obtain sharp characterizations of consumption allocations in the SIM model. If households have quadratic period utility functions, the borrowing constraints are loose enough to never be binding and $\beta R^{in} = 1$, then we are in the special case of the permanent income hypothesis (PIH) and consumption equals permanent income. In this case, if current income declines and increases affect permanent income in a symmetric fashion, consumption responds to both income declines and to income increases and it does so in a symmetric way.

For a more general specification of the utility function and/or binding borrowing constraints is possible to show that individual consumption responds to income fluctuations (because of partial insurance) and that, on average, it responds more to income increases than to income declines. The two features that generate this asymmetry are that the household's optimal consumption function is concave in cash at hand⁴ and that the income process is mean reverting. Note that households with low income have, on average, low cash at hand and, because mean reversion in income, high income growth. Similarly high income households tend to have high cash at hand and negative income growth. The concavity of the consumption function then insures that the consumption increases associated with the income increases of the low income households are larger than the consumption declines associated with the income declines of the high income households

3.2 Debt Constrained Markets

The first order and envelope conditions of problem (2) imply that⁵

$$\begin{aligned} u'(c) &\geq \beta R u'(c'(y')) \\ &= \text{if } g(y') > U^{Aut}(y') \end{aligned} \tag{5}$$

where $c'(y')$ represents optimal consumption next period. That is, in states in which the limited enforcement constraint is not binding the planner allocates consumption to the agent according to a full insurance (complete markets) Euler equation.

⁴See Carroll and Kimball (1996)

⁵Krueger and Perri (1999) prove, for the case that income shocks are *iid*, that the value function is differentiable. Here we go ahead and assume it for the general discussion.

The crucial difficulty in characterizing the optimal consumption allocations in this model is then to identify the states in which the limited enforcement constraints are not binding, and states in which they are binding. Krueger and Perri (1999) obtained the following theoretical characterization of the optimal policy function $g(w, y; y')$ for the *i.i.d.* case, in which case $g(w, y; y') = g(w; y')$, under the assumption that $\beta R < 1$.⁶

- Proposition 1**
1. For each y' tomorrow, the participation constraint is binding or utility promises are decreasing over time: $g(w, y') < w$ for all $w > U^{Aut}(y')$ and $g(w, y') = U^{Aut}(y')$ for all $w \leq U^{Aut}(y')$.
 2. Promises are equalized when possible and increased when the participation constraints bind: If $y' > \hat{y}'$, then $g(w, y') \geq g(w, \hat{y}')$. If $g(w, y') > U^{Aut}(y')$ and $g(w, \hat{y}') > U^{Aut}(\hat{y}')$, then $g(w, \hat{y}') = g(w, y') > U^{Aut}(y')$.
 3. $c(w)$ is strictly increasing in w .

Notice first that, if the constraints for low income states are binding, they are also binding for high income states. Second, for the states for which the constraints are not binding, consumption drifts down at a common rate. For example, if the utility function is of CRRA form, then (5) becomes

$$\log(c') - \log(c) \geq \frac{1}{\sigma} \log(\beta R)$$

with equality if $g(y') > U^{Aut}(y')$. This equation is the basis for the asymmetric response of consumption to income shocks: if income falls, consumption drifts down at rate $\frac{1}{\sigma} \log(\beta R) < 0$ and individual growth rate of consumption is actually independent from the actual income decline, if income increases the participation constraint is binding, consumption growth is positive and the size of consumption growth depends on the size of the income growth.

3.3 Testable implications

To describe the extent of risk allocation of income shocks Mace (1991), and Dynarski and Gruber (1997) run the following regression on the universe of CE data (which contains a short rotating panel for both households' income y_{it} and consumption c_{it})

$$\Delta \log c_{it} = \alpha_1 + \alpha_2 \Delta \log y_{it} + \alpha_3 \text{time dummies} + \beta X_{it} + \varepsilon_{it} \quad (6)$$

⁶We also show that if full insurance is not enforceable, an equilibrium satisfies $\beta R < 1$.

where X_{it} are controls for changes in household composition, seasonal and age effects and ε_{it} represents measurement error. The size of the coefficient α_2 can be interpreted as a measure of the extent of risk allocation for household income shocks. Similarly, in order to test for the asymmetries predicted by both models one can run the regression above restricting the sample to only households who experience positive (negative) income growth. We denote by α_2^+ (α_2^-) the coefficient in this type of regression. In table 1 we summarize the model predictions for the regression coefficients above, which we will compare to the empirical results from CE data. For reference we include also the predictions in the case of absence of credit markets (autarky) and in the case of perfect risk sharing (complete markets)..

Table 1: Consumption Response to Income Shocks

Model	α_2	α_2^-	α_2^+
Autarky	1	1	1
Standard Incomplete Markets	> 0	> 0	$\gg 0$
Debt Constrained Markets	> 0	0	> 0
Complete Markets	0	0	0

4 Calibration

The stochastic endowment process is specified as finite state discretization of an $AR(1)$ process as in Aiyagari (1994) and Heaton and Lucas (1996).

$$\log(y') = \rho \log(y) + \varepsilon$$

where $\varepsilon \sim iidN(0, \sigma_\varepsilon^2)$. Our strategy is to choose ρ from independent evidence and then to pick σ_ε^2 so that the cross-sectional variance of household earnings in the model corresponds to that in the data.

Not all cross-sectional income differences in the data are due to stochastic idiosyncratic income shocks however, but rather to cross-household differences in composition, education and other observable factors. In order to account for this fact we first regress per capita household earnings on age, race, sex and education dummies. We then interpret the residuals from this regression as idiosyncratic household income and choose σ_ε^2 in such a way as to match the cross-sectional variance of idiosyncratic household income in the data in 1991, the midpoint of our sample from 1980 to 2002.

This cross-sectional variance is 0.719.⁷ As a benchmark for the persistence parameter we choose $\rho = 0.98$, motivated by Storesletten et al. (2004), and set σ_ε^2 to match the residual cross-sectional variance of earnings, as discussed above. In the SIM model For the SIM model we need to specify the borrowing limit \bar{A} , which corresponds to the upper limit on uncollateralized loans. A reasonable benchmark borrowing limit may be taken as one year’s average income, that is $\bar{A} = 1$ but we also experiment with $\bar{A} = 4$. Finally with respect to preferences, we choose a log-utility function and pick the time discount factor β so that the stationary equilibrium in each model produces a real interest of 2.5%.

5 Results

Table 2 reports our estimates of α_2 , α_2^+ and α_2^- on CE data and simulated data from both models.⁸ We include all the households in the CE that are classified as complete income respondents, whose reference person is of age between 21 and 64, and which report positive labor earnings and consumption in their 2nd and 5th interview (which are the first and last time each household is interviewed). Our sample starts in the first quarter of 1980 and ends in the last quarter of 2003. For the data we report both OLS estimates as well as instrumental variables estimates. For the latter, the CE contains two independent observations of household income, one that directly relies on income responses, and one that can be constructed based on information for last pay-check and hours worked. To tackle the problem of measurement error in income growth which may bias the estimates for the income coefficient downward we instrument the first income growth measure by the second.⁹ The first three columns in the first row suggest that the DCM model predicts a substantially better risk allocation than the SIM model. From the CE data, we find that perfect insurance is rejected on statistical grounds: in response to a 1% change in income consumption changes range from 0.04% (using OLS estimates) to 0.14% (using IV estimates).¹⁰ Finally, comparing data to

⁷Due to space limitations we have to refer the reader to our detailed description of this procedure in Krueger and Perri (2002).

⁸Our consumption definition is expenditures on nondurables plus imputed service flow from consumer durables and our income definition is labor earnings minus taxes, plus transfers. For more details please see Krueger and Perri (2002).

⁹Our procedure follows Dynarsky and Gruber (1997); see their detailed discussion.

¹⁰In recent work Klein and Gervais (2004) argue that, due timing differences in the reporting of income, the IV estimator is likely to be biased upward and thus that the true value for α_2 might lie between the OLS and the IV

model results we conclude that while the DCM model slightly under-predicts the consumption response to income shocks, the SIM model grossly over predict it, regardless of whether one trusts the OLS estimates or the IV estimates. The last two rows of the table suggest that in the data the response of consumption to positive income shocks is not statistically different from the response to negative income shocks, thus contradicting the asymmetries predicted by both models. The PIH model, described in section 3 above, would perhaps be able to do a better job in accounting for this symmetric pattern of the data. In the next section we assess two simple extensions of the models discussed so far.

Table 2: Estimates of $\alpha_2, \alpha_2^-, \alpha_2^+$ (35157 observations, S.E. in parentheses)						
# of cases	Coef.	DCM	SIM, $\bar{A} = 1$	SIM, $\bar{A} = 4$	CEX (OLS)	CEX (IV)
35157	$\hat{\alpha}_2$	0.02	0.87	0.69	0.04 (0.003)	0.14 (0.011)
16056	$\hat{\alpha}_2^-$	0.00	0.78	0.56	0.04 (0.004)	0.15 (0.018)
19101	$\hat{\alpha}_2^+$	0.06	0.92	0.80	0.03 (0.004)	0.14 (0.017)

5.1 Sensitivity analysis

We first consider the case in which individuals who default are banned from trading risk sharing contracts and from *borrowing* forever, but are allowed to *save* at an interest rate $1 + r$, starting out with zero savings. Allowing agents to save helps them to self-insure against income fluctuations in autarky. Therefore the utility from being in autarky increases and less risk sharing is enforceable in equilibrium. The extent to which risk sharing declines with the introduction of saving depends on r . If we set $r = -1$ then the model coincides with no-saving case, but if we allow agents to save at the market interest rate of 2.5% after default, the estimates of α_2 in the model quadruples from 0.022 to 0.088 (see the first row of table 4) which is quite close to the coefficient estimated in the data. Still in this case the model predicts, contrary to the data, a significant asymmetry in the consumption response to income changes.

Another crucial parameter determining the extent of risk allocation is the persistence of the income shocks. Holding the variance of income shocks fixed, higher persistence of income shocks

estimate.

affect both models. In the DCM model the value of autarky increases for high income realizations (that is, those constraints that tend to be binding), leading to reduced risk sharing and thus a higher coefficient α_2 . On the other hand in the SIM model more persistent income shocks are harder to self-insure, again suggesting α_2 to increase with persistence ρ .

Table 3: Estimates of α_2 as Function of ρ			
Pers. ρ	DCM, $r = -1$	DCM, $r = 0.025$	SIM, $\bar{A} = 1$
0.98	0.02	0.09	0.87
0.93	0.02	0.07	0.81
0.53	0.00	0.03	0.43

Table 3 shows $\hat{\alpha}_2$ as a function of persistence ρ for both models.¹¹ As conjectured above in both models less persistent income shocks are easier to (self-) insure. Even for low persistence, the SIM model still predicts a consumption response to income shocks much higher than the one observed in the data.¹² The table thus suggests an important difference between the two models: the DCM model gets closer to the data with high persistence income shocks while the SIM model does a better job with low persistence income shocks.

6 Conclusions

We studied quantitative versions of two popular partial risk-sharing models and found that neither model can capture the actual consumption response to income fluctuations of US households. The DCM model, due to the presence of insurance markets predicts that households' consumption should be perfectly insulated from income declines and this does not seem true in the data. On the other hand the SIM model, due to the lack of insurance markets, predicts that consumption of US households should react to income shocks much more than it actually does, even when an income process with relatively low persistence is used. One lesson we learn, which echoes recent findings of Blundell et al (2002) and Storesletten et al (2004), is that a set-up which combines aspects of both models could be most helpful in understanding consumption smoothing in the US.

¹¹We always re-calibrate β and σ_ε^2 to match the interest rate and the cross-sectional earnings variance in the data.

¹²The variance σ_ε^2 for a $\rho = 0.53$ required to match the cross-sectional income variability from the data is so substantial that the borrowing constraint of $\bar{A} = 1$ constitutes the natural borrowing constraint, as defined in Aiyagari (1994).

References

- [1] Aiyagari, R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109, 659-684.
- [2] Alvarez, F. and U. Jermann (2000), “Efficiency, Equilibrium, and Asset Pricing with Risk of Default,” *Econometrica*, 68, 775-798.
- [3] Blundell R., L. Pistaferri and I. Preston (2002), “Information, partial insurance and consumption dynamics”, University College of London
- [4] Carroll, C. and M. Kimball (1996), “On the Concavity of the Consumption Function,” *Econometrica*, 64, 981-992.
- [5] Deaton, A. (1991), “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221-48.
- [6] Dynarski, S. and J. Gruber (1997), “Can Families Smooth Variable Earnings?,” *Brookings Papers on Economic Activity*, 229-284.
- [7] Heaton, J. and D. J. Lucas (1996), “Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing,” *Journal of Political Economy*, 104, 443-487.
- [8] Kehoe, T. and D. Levine (1993), “Debt Constrained Asset Markets,” *Review of Economic Studies*, 60, 865-888.
- [9] Klein P. and M. Gervais (2004), “Risk Sharing”, University of Western Ontario
- [10] Krueger, D. and F. Perri (1999), “Risk Sharing: Private Insurance Markets or Redistributive Taxes?,” *Federal Reserve Bank of Minneapolis Staff Report 262*.
- [11] Krueger D. and F. Perri (2002), “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” NBER working paper 9292
- [12] Mace, B. (1991), “Full Insurance in the Presence of Aggregate Uncertainty,” *Journal of Political Economy*, 99, 928-56..
- [13] Storesletten, K., C. Telmer and A. Yaron (2004), “Consumption and Risk Sharing over the Life Cycle”, *Journal of Monetary Economics*, 2004, vol. 51, 609-633,