

DISCUSSION PAPER SERIES

DP14579

CONSUMPTION INSURANCE OVER THE BUSINESS CYCLE

Tobias Broer

MONETARY ECONOMICS AND FLUCTUATIONS



CONSUMPTION INSURANCE OVER THE BUSINESS CYCLE

Tobias Broer

Discussion Paper DP14579

Published 07 April 2020

Submitted 06 April 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Tobias Broer

CONSUMPTION INSURANCE OVER THE BUSINESS CYCLE

Abstract

How do business cycle fluctuations affect the ability of households to smooth consumption against idiosyncratic shocks? To answer this question, we first document that, in U.S. micro-data, individual consumption reacts more to income changes in booms. Standard incomplete markets models, in contrast, where individuals borrow and save to smooth consumption, predict a lower sensitivity of consumption to individual income changes during times of high output. This motivates us to consider an alternative environment where financial frictions are endogenous and arise from lack of contract enforcement, whose business cycle properties have so far not been studied. We show analytically that this model is consistent with a wide variety of cyclical patterns of insurance. In a quantitative application with unemployment risk, we show that the response of individual consumption to job losses differs strongly between times of high and low output, and identify the conditions under which it is procyclical, as in the data.

JEL Classification: E32, G22

Keywords: Consumption Smoothing, Risk Sharing, Limited Enforcement, business cycles

Tobias Broer - tobias.broer@iies.su.se
IIES Stockholm University and CEPR

Acknowledgements

I thank, in particular, Paul Klein. Without his support this paper would not have been written. I also thank Per Krusell, Fabrizio Perri, Morten Ravn, and Kjetil Storesletten, as well as participants in several seminars and conferences, for comments on previous drafts.

Consumption insurance over the business cycle*

Tobias Broer[†]
Paris School of Economics
and IIES, Stockholm University

January 2020

Abstract

How do business cycle fluctuations affect the ability of households to smooth consumption against idiosyncratic shocks? To answer this question, we first document that, in U.S. micro-data, individual consumption reacts more to income changes in booms. Standard incomplete markets models, in contrast, where individuals borrow and save to smooth consumption, predict a lower sensitivity of consumption to individual income changes during times of high output. This motivates us to consider an alternative environment where financial frictions are endogenous and arise from lack of contract enforcement, whose business cycle properties have so far not been studied. We show analytically that this model is consistent with a wide variety of cyclical patterns of insurance. In a quantitative application with unemployment risk, we show that the response of individual consumption to job losses differs strongly between times of high and low output, and identify the conditions under which it is procyclical, as in the data.

JEL Classification: E32, G22

Keywords: consumption smoothing, risk sharing, limited enforcement, business cycles

1 Introduction

The degree to which households can smooth consumption in the face of unexpected shocks is an important determinant of their average well-being. This paper studies

*I thank, in particular, Paul Klein. Without his support this paper would not have been written. I also thank Per Krusell, Fabrizio Perri, Morten Ravn, and Kjetil Storesletten, as well as participants in several seminars and conferences, for comments on previous drafts.

[†]IIES, Stockholm University, SE-10691 Stockholm; Email: tobias.broer@iies.su.se.

how fluctuations in aggregate economic activity affect this ability of households to smooth consumption against idiosyncratic shocks. We think this is interesting for at least two reasons: first, fluctuations in the degree of consumption smoothing, and their comovement with aggregate economic conditions, may inform us about the economic frictions that make risk sharing imperfect in many economic contexts. Such information is useful because the effect of policies aimed at reducing consumption volatility depends on those frictions. Second, business cycles have additional welfare costs when they decrease the degree of risk sharing, or consumption smoothing, on average or make it more variable over time. In fact, the utility cost of cyclical fluctuations in aggregate income is approximately proportional to its variance V_Y .¹ That of cyclical fluctuations in consumption smoothing, in contrast, is proportional to $E(\beta^2 - \bar{\beta}^2)V_y$, where V_y is the variance of idiosyncratic incomes while β and $\bar{\beta}$ denote the fractions of income shocks that pass through to consumption in, respectively, the cyclical economy and an economy without cyclical fluctuations. Business cycles are thus costly not only whenever they lower insurance against idiosyncratic shocks (or raise β) on average, but also when they increase the variance of insurance. Since idiosyncratic incomes are substantially more volatile than their aggregate, cyclical fluctuations in consumption insurance are likely to have a substantial additional welfare cost, relative to standard measures of the cost of business cycles.²

Motivated by these considerations, our paper makes three contributions: first, we show that in data from the US consumer expenditure survey (CEX), household consumption reacts more strongly to individual income changes in booms. In other words, the sensitivity of individual consumption to income changes is procyclical.

¹See Appendix III for a derivation.

²The standard deviation of the log-difference in family disposable post-tax income in data from the US consumer expenditure survey (CEX) is, after accounting for time-fixed effects, 50 percent (compared to 65 for family earnings). The standard deviation of log-differences in aggregate US disposable household incomes (when log-differences are calculated as the same year-on-year overlapping averages as in the CEX, see the data section for details), is 1.2 percent. The variance of β , measured using a yearly regression of total consumption growth on income growth, is 6.4 percent for disposable income (4.0 percent for family earnings). This implies a ratio $\frac{V_\beta V_y}{V_Y}$ of about 110 for both measures of individual incomes. While measurement error in individual incomes inflates this ratio, it attenuates estimates of β , whose mean and variance are around 10 and 6 percent, respectively, in US CEX data.

Second, we show how in standard economies with (exogenously) incomplete markets, the cyclical behavior of consumption smoothing is determined by the cyclicity of interest rates and income persistence, and (when preferences exhibit prudence or borrowing limits are imposed) by that of the wealth and income distribution. A quantitative analysis shows that, in a bond economy similar to Huggett (1997), consumption reacts *less* to income changes in booms, in contrast to the data. Moreover, aggregate risk strongly increases the average sensitivity of consumption to income changes. In an economy with capital (Krusell and Smith, 1998), where the average distance of household assets from their lower bound is larger and interest rates are less cyclical, this sensitivity is still mildly lower in booms (but on average unaffected by aggregate fluctuations), again in contrast to the data.

One reason for this counterfactual cyclicity of consumption insurance in simple models of consumption smoothing is the exogenous character of financial frictions, which do not respond to business cycle conditions: both the set of assets available, and the (borrowing) constraints under which they are traded are exogenous. This motivates us to study the cyclicity of consumption insurance in an alternative environment without exogenous restrictions to asset trade, but with endogenous, and therefore potentially cyclical, financial frictions arising from limited enforcement of co-insurance contracts. Since insurance is more valuable in times of high unemployment risk and low average income we would expect participation constraints to loosen, and insurance to improve, in bad times. This could explain the finding of countercyclical insurance in CEX data.

Our third contribution is to characterise the cyclical properties of risk sharing in a standard continuum ("LC") economy with complete markets and limited commitment (Krueger and Perri, 2011, Krueger and Uhlig, 2006, Krueger and Perri, 2006) analytically and quantitatively. We first show a separation result similar to Werning (2015)'s for incomplete markets economies (but not limited to deterministic aggregate fluctuations): with relative risk aversion equal to 1 (log-preferences) and idiosyncratic risk that is independent of aggregate conditions, individual consumption shares are independent of the history of aggregate shocks in the LC economy,

and equal to those in the stationary environment without aggregate fluctuations. So consumption insurance is acyclical. With risk aversion greater than 1 or cyclical income risk, in contrast, the sensitivity of consumption to income changes is cyclical, and in line with the data whenever participation constraints are more binding in booms. We show how standard techniques to solve heterogeneous agent economies (Krusell and Smith, 1998, Boppart et al., 2018, Reiter, 2009) are problematic in the LC economy, and propose a new solution method, based on the near-analytical solution to the stationary LC economy (Krueger and Perri, 2011, Broer, 2013). In a quantitative application to unemployment risk, we find that cyclical fluctuations in consumption risk sharing in the LC economy are substantially larger than in the standard incomplete markets economy, while the average degree of risk sharing is similar to that in the stationary economy. With procyclical income persistence, the LC economy predicts a stronger response of consumption to idiosyncratic income shocks in booms, as in the data.

Relation to the literature

Our analysis links two literatures that have so far remained largely separate. First, the extensive theoretical and empirical literature on consumption risk sharing, surveyed e.g. in Attanasio and Weber (2010), has largely abstracted from cyclical fluctuations, concentrating on the average degree of consumption risk sharing or its trend (Krueger and Perri, 2006, Blundell et al., 2008), and on stationary, or deterministic, model environments. We contribute to this literature in two ways: empirically we document substantial and significant comovement between cyclical components of aggregate GDP or aggregate disposable household income and standard indicators of consumption risk sharing in US data. Theoretically we extend the stationary environment where insurance is limited by the risk of default, studied in Krueger and Uhlig (2006), Krueger and Perri (2011), Krueger and Perri (2006), or Broer (2013), to include aggregate stochastic fluctuations and characterize it both quantitatively and analytically.³

³See Lepetyuk and Stoltenberg (2013) for an environment with limited commitment and one-time uncertainty about a future aggregate state where transfers are constrained to only depend on current income, in contrast to our setting with aggregate risk where transfers to unconstrained

Our results link this literature on risk sharing to that on aggregate economic fluctuations in economies with heterogeneous agents, which has concentrated mainly on exogenously incomplete markets (Bewley, 1977, Imrohoroglu, 1989, Aiyagari, 1994, Huggett, 1997, Krusell and Smith, 1998), and on the relationship between inequality and aggregate economic performance including the effect of policies⁴. The effect of aggregate conditions on the ability of consumers to protect their consumption from income fluctuations has received less attention.⁵ Relative to this literature, recently surveyed in Krueger et al. (2016), we make two contributions: first, we study the implications of aggregate fluctuations not for the marginal cross-sectional distribution of consumption or wealth, but for standard measures of consumption insurance related to the joint distribution of consumption and income growth. Our second contribution to this literature is to study, analytically and quantitatively, an alternative market structure of complete markets with participation constraints (as in Alvarez and Jermann (2000) or Kehoe and Levine (1993)).

Previous studies of cyclical movements in consumption inequality studied the dynamics of the cross-sectional distribution of consumption over the cycle (De Giorgi and Gambetti, 2017), or in response to monetary policy shocks (Coibion et al., 2017), but not the joint distribution of individual consumption and income. Because current income shocks matter more for permanent incomes when their persistence is high, the cyclicity of income risk is relevant for that of consumption smoothing. While Storesletten et al. (2004) found income risk to be countercyclical in US survey data, our evidence from CEX data is in line with Guvenen et al. (2014a) who document acyclical variance (but procyclical skewness) of income changes in US administrative income data.

agents have (potentially long) history dependence. Chien and Lustig (2009) study a related setting with aggregate fluctuations and financial portfolio constraints.

⁴See Kaplan and Violante (2018) for a survey of the recent literature on heterogeneous-agent new Keynesian (HANK) models. This literature has shown how the heterogeneity in marginal propensities to consume implied by idiosyncratic risk and incomplete markets changes the transmission of shocks, and particularly the effect of policies that have distributional implications.

⁵An exception to this is a recent paper by Acharya and Dogra (2018), who derive a closed form expression for consumption as a function of asset holdings, individual income, and aggregate conditions in an endowment economy without aggregate risk under the assumption of exponential utility.

Section 2 presents our empirical findings based on CEX data. Section 3 discusses the effect of aggregate risk on consumption smoothing in environments with exogenously incomplete markets. Section 4 discusses the LC economy with endogenous financial frictions, and Section 6 concludes.

2 Consumption Risk and Insurance Over the Business Cycle - Evidence from CEX Data

In this section, we document how consumption risk and standard measures of consumption insurance fluctuate over the U.S. business cycle. For this we use data from the Consumer Expenditure Survey (CEX), a 4-quarter rotating panel with detailed information about quarterly household consumption expenditures. The CEX only collects information about *annual* household labor earnings and disposable income in the first and fourth survey round. Moreover, the gap between aggregate CEX consumption and that in the U.S. National Income and Product Accounts has widened over time, which has been interpreted as a decline in the quality of CEX consumption data.⁶ Together, these facts point to substantial noise in CEX data on the joint income-consumption distribution. We nevertheless choose the CEX as a source of information over alternative datasets, such as the PSID, because of its high frequency, and its broad coverage of consumption items.⁷ We interpret our results below bearing in mind the issue of measurement error, and discuss the extent to which it may influence our measurement of the comovement between consumption smoothing and the business cycle. In addition, we are aware of the somewhat odd timing of measurements of consumption and income in the CEX. We deal with it in the following way: rather than follow Gervais and Klein (2010) and adjust our econometric approach to measuring consumption smoothing, we compute moments

⁶For a discussion, see Davis (2003), and Battistin (2003) who argues that in particular the quality of CEX interview survey data on frequently purchased small items has declined.

⁷See Gervais and Klein (2010) for details on the timing in the CEX. In addition, information about disposable income is missing in the years 2004 and 2005. The PSID, in contrast, only contains annual information about food consumption until 1996, when it broadens its coverage of consumption categories but changes from annual to biannual frequency.

in all model simulations in such a way as to be consistent with the timing in the CEX. Specifically, we use 4-quarter differences in log quarterly consumption, and changes in annual log income observations which partially overlap with the consumption observations, just as in the CEX.

Since we are interested in private smoothing, or risk sharing, of overall consumption conditional on public insurance through taxes and transfers, our benchmark results focus on the joint distribution of the growth rates of the CEX definition of family disposable income, on the one hand, and a broad nondurable consumption aggregate (including rental payments and imputed rental services for house owners), denoted $ND+$, on the other. We also consider an alternative income measure, family earnings, and two alternative consumption measures, the CEX aggregate of nondurable consumption excluding rental payments (ND), and food consumption ($Food$). Our CEX sample starts in 1983 and ends in 2012. We focus on households whose head is of working age (between 21 and 64 years of age), and who are labeled as complete income respondents, and whose income is not top-coded.⁸ Appendix *I* contains details about data construction.

As measures of the business cycle we consider deviations from a log-linear trend of three aggregate output or income measures: real GDP, real household disposable income from the National Income and Product Accounts (NIPA), and the mean of disposable income in the CEX sample. We define as booms, or good times, those of above-trend aggregate activity, and as bad times those with activity below trend. We prefer this definition to the natural alternative of using recessions as determined by the NBER Business Cycle Dating Committee. Apart from the fact that our sample period only comprises three such recessions, our definition is more in line with our model predictions, relating to times of below-trend activity, while NBER recessions aim to identify the period between peak and trough of the cycle, and thus periods of *declining* activity.

A particularly simple measure of the sensitivity of individual household consump-

⁸In addition, we exclude households whose composition changes, and those that have not completed all four surveys.

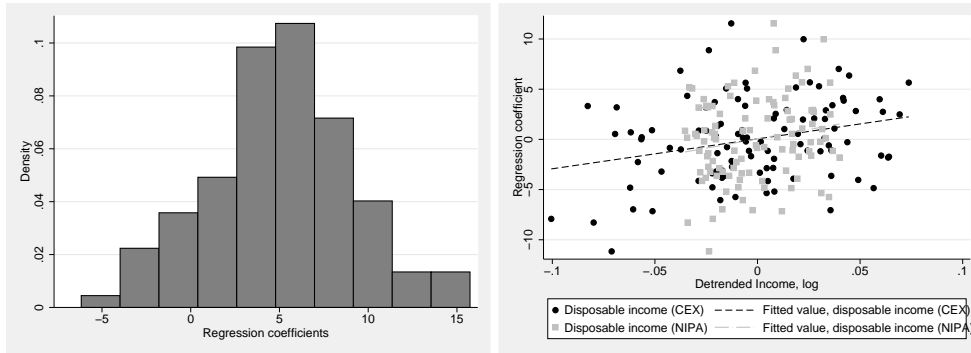
tion to individual income changes is the slope of the conditional mean of consumption growth as a function of income growth, equal to the coefficient β in the following regression

$$\Delta c_t = \alpha + \beta \Delta y_t + \varepsilon_t \quad (1)$$

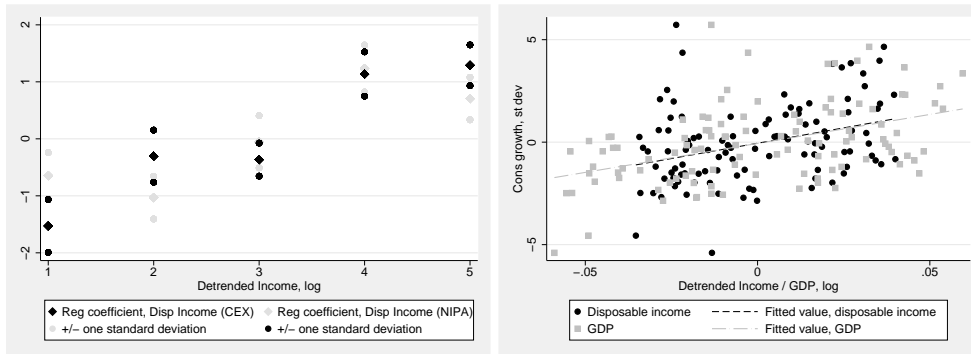
where Δc_t and Δy_t denote the (4-quarter) log-difference of individual consumption and income, respectively, α is a constant, and ε_t is an error term. Although not a structural parameter per se, we use β as a simple, and classical, measure of consumption insurance (see Gervais and Klein (2010)'s discussion of the literature) that allows us to indirectly infer model parameters. Importantly, any measurement error in consumption leaves the regression coefficient unaffected, while error in measured incomes over and above that introduced by the asynchronicity with consumption (which we deal with explicitly) attenuates it. In addition to β , we also look at the cyclical behavior of a simple measure of consumption risk, namely the cross-sectional standard deviation of consumption growth $\text{STD}(\Delta c)$. As we argued in the introduction, it is the cyclicity of $\text{STD}(\Delta c)$ that may strongly affect the cost of business cycles. Although measurement error in consumption strongly affects the level of $\text{STD}(\Delta c)$, it does not affect its cyclical behavior as long as the error is not in itself cyclical.

According to Table 1 the regression coefficient β equals between about 3 and 6 percent on average in our sample, in line with values found in previous studies, indicative of strong average insurance and / or measurement error in income. Panel *a*) of Figure 1 shows how, when we estimate (1) quarter by quarter, this average masks a wide dispersion of coefficient values, whose standard deviation equals more than 80 percent of their mean. This substantial dispersion is not, however, simply due to noise: According to Panel *b*) of Figure 1 periods when aggregate disposable income or GDP are above trend (along the bottom axis) are associated with higher-than-average sensitivity of individual consumption to income changes (higher values of β , along the vertical axis). Moreover, this cyclicity is substantial: according to panel *c*) the regression coefficient β is on average three quarters higher when the cyclical component of income is in the top quintile, compared to the bottom quintile.

Figure 1: Consumption Insurance in CEX data



(a) Histogram of β , in percent, in (1), consumption measure ND+, income measure disposable income. (b) Scatter plot of β , in percent, in (1) estimated quarter-by-quarter (vertical axis) against detrended disposable income (NIPA) / GDP (log, horizontal axis), quarterly frequency.



(c) β in (1) (in percent, vertical axis) estimated quarter-by-quarter and averaged within quintiles of detrended disposable income (CEX and NIPA) (log, horizontal axis). (d) Scatter plot of the standard deviation of ND+ consumption growth (in percent, vertical axis) against detrended disposable income (NIPA) / GDP (log, horizontal axis), quarterly frequency.

Table 1: Regression coefficient β , in percent

	(1)	(2)	(3)	(4)	(5)	(6)
	ND+	ND	Food	ND+	ND	Food
Disp income growth	5.043*** (15.17)	6.213*** (14.53)	5.010*** (9.47)			
Earnings growth				2.802*** (11.04)	3.033*** (9.28)	2.692*** (6.67)
Constant	0.261 (1.60)	0.746*** (3.55)	-1.129*** (-4.34)	0.322** (1.97)	0.830*** (3.94)	-1.063*** (-4.09)
r2	0.00664	0.00609	0.00260	0.00353	0.00250	0.00129
N	34444	34444	34444	34443	34443	34443

The table reports the regression coefficient β , in percent, in different versions of (1) when c_t is nondurable consumption plus (imputed) rental services (ND+, columns 1 and 4), nondurable consumption (ND, columns 2 and 5), or food consumption (Food, columns 3 and 6), and when y_t is total disposable family income after taxes and transfers (row 1), or family earnings (row 2). Robust standard errors are used; stars denote conventional significance levels: * ($p < .1$), ** ($p < .05$), *** ($p < .01$).

Table 2 depicts the corresponding coefficients in a regression of quarterly β coefficients on deviations of aggregate income measures from their log-linear trend. It shows that the procyclical sensitivity of individual consumption to income changes is highly significant for both measures of nondurable consumption (ND+ and ND, in columns 1 and 2, respectively), and for both the CEX and the NIPA measure of aggregate disposable income as indicators of business cycles (columns 1 and 3). It is insignificant, however, for food consumption (column 4).

The procyclical sensitivity of consumption to income changes documented in Panel *b*) and *c*) of Figure 1 results in strongly procyclical cross-sectional dispersion of consumption growth, as shown by the strongly increasing standard deviation in Panel *d*). This relationship holds not only when taking disposable income (CEX or NIPA) as an indicator of the cycle, but also for deviations of log US GDP from its trend (in gray). Moreover, it is, again, economically important, and not due to extreme observations in the micro-data: the 90-10 percentile ratio of ND+ consumption growth in the cross-section is 10 percent higher when detrended disposable income is in

Table 2: β and aggregate income

	(1)	(2)	(3)	(4)
	ND+	ND+	ND	Food
Disp income (CEX)	29.78***		31.57**	-7.626
	(2.66)		(2.18)	(-0.45)
Disp income (NIPA)		37.08**		
		(2.00)		
Constant	0.0510	0.0531	0.0540	-0.0131
	(0.13)	(0.13)	(0.11)	(-0.02)
r2	0.0760	0.0376	0.0565	0.00253
N	102	102	102	102

The table reports results from a regression of the quarterly time series of coefficients β (using ND+ consumption growth in column 1 and 2, ND consumption growth in column 3, and Food consumption growth in column 4), in percent, on log-trend deviations of aggregate disposable income (from the CEX, columns 1, 3 and 4, and from NIPA, column 2). Robust standard errors are used, stars denote conventional significance levels: * (p<.1), ** (p<.05), *** (p<.01).

the top quintile, relative to the bottom quintile. Given that classical measurement in (log-) consumption levels introduces substantial acyclical measurement error in consumption growth (as the measurement error in differences is twice that in levels and we expect the fundamental dispersion of consumption growth rates to be small relative to levels), this value is likely to underestimate the true cyclical increase in consumption risk. At the same time, the cross-sectional dispersion of disposable income growth is constant across the cycle in our data, in line with Guvenen et al. (2014b)'s finding that income risk is acyclical. Together, this implies that the relative dispersion of consumption and income growth increases significantly in booms, independently of the particular indicator of business cycles, as Table 3 shows.⁹

⁹Table 3 looks at differences in percentile ratios, which are more robust to outliers in the quarterly micro-data. Results for differences in standard deviations are similar, and significant at the 1 percent level for the NIPA measure of disposable income and GDP as measures of the business cycle, but not for CEX aggregate disposable income.

Table 3: Relative dispersion in consumption and income growth and aggregate income

	(1)	(2)	(3)
	Relative dispersion	Relative dispersion	Relative dispersion
Disp Y (CEX)	0.487** (2.52)		
Disp Y (NIPA)		0.919** (2.59)	
GDP			0.809*** (3.16)
Constant	0.000833 (0.11)	0.00132 (0.18)	0.00123 (0.17)
r2	0.0548	0.0624	0.0953
N	102	102	102

The table reports results from a regression of the quarterly time series of the log-percentile ratio difference (the 90-10 log percentile ratio of consumption growth (ND+ column 1 and 2, ND column 3) minus that of disposable income growth) on deviations of logarithms of aggregate disposable income (from the CEX, column 1, and from NIPA, column 2) and of GDP (column 3) from a linear trend. Robust standard errors are used, stars denote conventional significance levels: * (p<.1), ** (p<.05), *** (p<.01).

3 Cyclical consumption smoothing with exogenously incomplete markets

The previous section documented substantial cyclical consumption smoothing and risk around their average values in CEX data. This gives rise to two questions: what are the sources of these cyclical fluctuations? And how, if at all, do they affect the average level of insurance? To study these questions, this section looks at a standard “self-insurance” (SI) environment where asset trade is exogenously restricted to a non-contingent bond, as in Huggett (1997), or to non-contingent claims to capital, as in Krusell and Smith (1998).

3.1 The model environment

Consider an economy in discrete time ($t = 0, 1, \dots$) populated by a continuum of agents of total mass one indexed by $i \in [0, 1]$ who receive a stochastic amount of a perishable consumption good as endowment income e_{it} every period. Aggregate income $Y_t = \int_0^1 e_{it} di$ is confined to a finite set $\mathbb{Y} := \{Y^1, Y^2, \dots, Y^N\}$ where $Y^1 < Y^2 < \dots < Y^N$ with a time-homogeneous probability transition function P so that

$$\text{Prob}[Y_{t+1} = Y^k | Y_t = Y^j] = P_{jk}$$

We assume that this defines an irreducible Markov chain. Let $\Pi_t^k(j)$ denote the probability that $Y_t = Y^j$ given that $Y_0 = Y^k$. Let $Y^t = (Y_0, Y_1, \dots, Y_t)$ denote a history of aggregate income shocks from 0 to t and let \mathbb{Y}^t denote the set of possible such sequences.

Suppressing the subscript i from now on, we denote individual income shares as $y_t = \frac{e_t}{Y_t}$. They too are confined to a finite set $\mathbb{y} := \{y^1, y^2, \dots, y^n\}$ with $y^1 < y^2 < \dots < y^n$. We assume aggregate income transitions to be independent of realisations of y_t and thus of their distribution. Individual transition probabilities, in contrast, may depend on aggregate income transitions, and are described by a time-homogeneous

probability transition function p so that

$$\text{Prob}[y_{t+1} = y^l | Y_{t+1} = Y^k, Y_t = Y^j, y_t = y^i] = p_{il}^{jk}$$

Again, we assume that p defines an irreducible Markov chain.

Together, p and P define the joint probability transition function $Q_{lk}^{jk} = \text{Prob}[Y_{t+1} = Y^k, y_{t+1} = y^l | Y_t = Y^j, y_t = y^i] = p_{il}^{jk} P_{jk}$, with an associated unique stationary distribution. We confine our attention to transition functions p_{il}^{jk} with the property that, once the economy has reached that ergodic distribution, in period t , say, the distribution of individual income shares y^j conditional on aggregate income Y^k , denoted $\pi(j|k)$, is constant over time, such that $\pi(j|k)_{t+s} = \pi(j|k)$; $s = 0, 1, 2, \dots$. In other words there is no transition in the conditional distribution of income shares, which "jumps" between its stationary conditional distributions (that may vary, however, across aggregate income states). Within this class of individual transitions, we look at two special cases: we call "independent" individual transitions that are independent of aggregate income, implying $p_{il}^{jk} = p_{il}^{mn} \equiv p_{il}$ for any $j, k, m, n \in \{1, \dots, N\}$, implying that $\pi(j|k) = \pi(j|l) \equiv \pi(j)$, $\forall k, l \in \{1, \dots, N\}, j \in \{1, \dots, n\}$. Second, we call "iid" individual transitions that in addition are independent of individual income, such that $p_{il} = p_{kl} \equiv p_l = \pi(l)$, $\forall j \in \{1, \dots, n\}$.

Let $y^t = (y_0, y_1, \dots, y_t)$ denote a history of individual income shares from 0 to t and let \mathbb{y}^t denote the set of possible such sequences.

Individuals rank random consumption sequences $\{c_t\}_{t=0}^{\infty}$ according to

$$E_0 \left[\sum_{t=0}^{\infty} \delta^t u(c_t) \right].$$

where δ is a common discount factor, E_t denotes the mathematical expectation conditional on time t information, and u is a strictly increasing, strictly concave and twice continuously differentiable utility function.

3.2 Market structure and competitive equilibrium

Individuals trade, in period t , non-contingent bonds at price q_t . Individual bond holdings b_t are constrained by a borrowing limit $B > 0$, such that $b_t \geq -B, \forall t$. Denote as $\pi(y, b)_t$ the cross sectional joint distribution of bond holdings and income shares in period t and as Γ the transition function for this distribution conditional on Y_t and Y_{t+1} . For V the individual's value function, and denoting next period's variables with primes, the individual problem is thus

$$\begin{aligned}
 V(b, y, \pi(y, b), Y) &= \max_{c, b'} u(c) + \delta E [V(b', y', \pi(y, b)', Y')] \\
 &\quad s.t. \quad c + qb' = b + y \\
 &\quad \pi(y, b)' = \Gamma(\pi(y, b), Y', Y) \\
 &\quad b \geq -B
 \end{aligned} \tag{2}$$

where individuals use the law of motion Γ to predict next period's distribution $\pi(y, b)'$, which determines next period's demand for bonds and thus their price q' .

We define a competitive equilibrium following Krusell and Smith (1998): A recursive competitive equilibrium consists of an equilibrium law of motion Γ , a bond price function q , individual decision rules c, b' , such that the latter solve the individual's problem (2) given q and Γ , the bond market clears, and Γ is induced by the exogenous transition function p and optimal savings behavior.

3.3 Analytical characterization of equilibrium

This section shows how in SI economies periods of low interest rates, low income persistence, and high average resources are associated with a low sensitivity of consumption to idiosyncratic income shocks. With acyclical income risk, this implies a countercyclical sensitivity of consumption to income changes in endowment economies, where interest rates are countercyclical, in contrast to the procyclical sensitivity in CEX data.

3.3.1 Quadratic preferences: the role of interest rates and cyclical income risk

Consider a simplified version of the SI environment without borrowing constraints and without aggregate risk, populated by consumers with quadratic preferences $u(c) = c - \frac{1}{\theta}c^2$. Aggregate output equals $Y_{t+s} = \bar{Y}$, $s = 1, 2, \dots$ in all but the first period, and individual incomes are iid, such that $E[y_{it}] = 1 \forall i, t > 0$. We study how an aggregate, transitory shock in period t changes the response of individual consumption to individual income shocks in that period. To capture the fact that both market prices and idiosyncratic risk vary over the business cycle, the shock has two effects: first, it raises current output by ϵ such that $Y_t = \bar{Y} + \epsilon$. And second, it increases the probability that individual income remains unchanged next period by $p(y)$, not necessarily positive and potentially dependent on current income share y , such that the expected income share is $E[y_{it+1}] = 1 + p(y)(y_{it} - 1)$. As Appendix IV shows, period t consumption shares are then characterized by

$$\tilde{c}_t - 1 = \frac{r \left(\rho_t^2 + \frac{\rho_t p}{1+r} \frac{\bar{Y}}{Y_t} \right)}{1 + \rho_t^2 r} \cdot (y_{it} - 1) + \frac{r}{1 + \rho_t^2 r} \cdot \tilde{a}_{it} \quad (3)$$

where $\tilde{x}_t = \frac{x_t}{Y_t}$, $r = \frac{1}{\delta} - 1$ is the stationary net interest rate (the inverse of the bond price) and $\rho_t = \frac{1+r_t}{1+r}$ is the relative current gross interest rate. The sensitivity of consumption shares to individual income shocks ($y_{it} - 1$) is thus increasing in the current interest rate (through its impact on ρ_t), as higher interest rates increase the weight of current income in permanent income. The sensitivity also increases when aggregate shocks make individual incomes more persistent by raising $p(y)$, thus, again, increasing the impact of current income shocks on permanent income.¹⁰

This simple example highlights two sources of cyclical fluctuations in consumption smoothing: the cyclicity of income risk, and that of interest rates. Equilibrium (relative) interest rates equal $\rho_t = 1 - \frac{\epsilon}{\theta + \bar{y}}$ and are thus countercyclical (as mean

¹⁰Acharya and Dogra (2018) show a similar result in a heterogeneous-agent New Keynesian economy with CARA preferences.

reversion implies higher expected output growth when current output is low). The cyclicity of interest rates thus implies a countercyclical sensitivity of individual consumption shares to income shares, contrary to the data. Sufficiently procyclical income persistence, however, may still lead to procyclical sensitivity. The sensitivity of consumption shares to individual income shocks ($y_{it} - 1$) is independent, in contrast, of the current level (and thus the cross-sectional distribution) of assets a_{it} , and of the dispersion of future incomes (current or permanent). The next section shows how more general preferences, and constraints to intertemporal trade in the form of borrowing limits, break this irrelevance-of-distributions feature of quadratic utility.

3.3.2 General preferences and borrowing constraints

With convex marginal utility, or “prudence”, income uncertainty reduces consumption through precautionary savings, and more so at lower levels of wealth (where risky incomes are a higher share of total wealth, and current asset holdings closer to any borrowing limit, Kimball (1990)). In stationary SI environments with purely transitory income shocks, a constant borrowing limit and constant relative risk aversion, precautionary savings thus typically make consumption a concave function of current resources.¹¹ An unanticipated change in income therefore affects consumption more at lower levels of current resources, making the wealth and income distribution a crucial determinant of average consumption insurance. In particular, at a given interest rate, an increase in the dispersion of current wealth or incomes deteriorates average insurance (as, by concavity, the higher sensitivity of consumption to current resources of the poor is less than offset by a lower sensitivity of the rich), while an increase in average wealth or incomes improves it. In general equilibrium, Werning (2015) shows how the sensitivity of consumption to individual

¹¹Carroll and Kimball (1996) derive conditions on preferences (satisfied e.g. by constant relative (CRRA) or absolute (CARA) risk aversion) for the consumption function to be concave when borrowing constraints never bind in finite horizon problems. With occasionally binding liquidity constraints Carroll and Kimball (2001) show this for CARA and CRRA preferences, Nishiyama and Kato (2011) for quadratic utility.

income may nevertheless be independent of (deterministic movements in) aggregate income (and the coefficient β in (1) thus constant) when preferences have the log-form and borrowing limits and idiosyncratic risk are exactly proportional to aggregate income. With constant borrowing limits, or less-than-proportional income risk, in contrast, concave consumption functions are an additional reason why in SI economies the sensitivity of individual consumption to income changes is lower in booms. In endowment economies, this reinforces the effect of countercyclical interest rates, implying a (counterfactually) countercyclical sensitivity whenever income risk is acyclical.

3.4 Quantitative Analysis

To see how the forces highlighted in the previous combine to determine business cycle fluctuations in the degree of consumption smoothing, and how aggregate uncertainty affects average consumption smoothing in equilibrium, this section studies the SI economy quantitatively.

3.4.1 Functional forms and benchmark parameter values

We consider utility that exhibits constant relative risk aversion (CRRA) with risk-aversion parameter $\sigma \geq 1$, such that

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

As a benchmark, we consider risk aversion slightly higher than implied by log preferences, and set $\sigma = 1.5$. We choose a value of the discount factor δ such that our preferred measure of risk sharing, the regression coefficient β in (1), approximately equals 6 percent, in line with the values observed in CEX data.

To transparently capture key features of the post-II war U.S. economy, we consider an economy similar to Krusell and Smith (1998) that goes through times of high and low aggregate output (which we label booms or “good” times, and “bad” times,

respectively), and where individual uncertainty arises from unemployment risk. We thus have $N = n = 2$, and $Y \in \{Y^H, Y^L\}$, $y \in \{y^h, y^l\}$ where superscripts h, H and l, L denote high and low individual endowment shares and aggregate endowments, respectively. We define good and bad times in the data as those when U.S. GDP is, respectively, above and below its log-linear trend. We choose transition probabilities P_{jk} such that the Markov process matches the average length of bad times, and the average time the economy spends there, in the data. Finally, we choose the ratio of Y^H and Y^L to equal that of average levels of total factor productivity in good and bad times. We specify idiosyncratic states y^h, y^l as corresponding to employment and unemployment, respectively, with $y^l = 0.6 * y^h$, to mimick a replacement rate of 60 percent, approximately equal to the net replacement rates reported by the OECD for the U.S. in recent times.¹² We choose a Markov process for y such that the job finding probability p_{lh}^{JK} in the model equals that in U.S. data (between 1948 and 2012) in good and bad times, respectively, and set the separation rate p_{hl}^{JK} such that the unemployment rate in the ergodic distribution equals those observed on average in good and bad times.

Row 1 of Table 4 reports the transition probabilities we identify for aggregate endowments from US GDP, and for individual endowment shares from US data and job-finding rates calculated as in Shimer (2005). Good times are more persistent than bad ones, such that the economy spends a little more than 70 percent of periods in good times, where the aggregate endowment is 1.5 percent higher than its mean (normalized to 1), implying aggregate endowments of 0.959 in bad times. Job-finding rates are high, close to 80 percent per quarter, and procyclical. Separation rates are countercyclical. This finding based on observed transitions in unemployment status is somewhat in contrast to the acyclical income risk we observe in CEX data. We therefore also consider two alternative parameterizations: first, we denote as “independent” (row 2) a case where transition probabilities are identical in good and bad times such that unemployment is constant at its average, observed value (equal to 5.8 percent). Second, we also consider a parameterisation that raises the potential

¹²<https://stats.oecd.org/Index.aspx?DataSetCode=NRR>.

Table 4: Benchmark parameter values

	P_{HH}	P_{LL}	p_{hl}^{HH}	p_{hl}^{LL}	p_{lh}^{HH}	p_{lh}^{LL}
Countercyclical risk	96	88	4.77	5.10	79.3	77.4
Independent	96	88	4.85	4.85	78.7	78.7
Procyclical persistence	96	88	4.75	4.95	77.2	80.3

The table reports job-finding and separation rates conditional on (staying in) good and bad times, in percent, as used in the quantitative analysis below. See Appendix VII for the full transition matrix.

of SI economies to explain the procyclical sensitivity of individual consumption to income changes in the data by strengthening the persistence of incomes, and thus the link between current and permanent incomes, in booms. We label this case “procyclical persistence”, where the persistence parameters p_{ii}^I , $i = h, l$ are two percent higher (lower) in good (bad) times than their value with independent transitions, such that $p_{hh}^{HH} = p_{ll}^{LL} > p_{hh}^{LL} = p_{ll}^{HH}$.¹³

3.4.2 Model solution

Solving for equilibrium allocations and prices is difficult in SI economies because the current joint distribution of individual income and bond holdings matters for “average” decisions that determine future prices and thus also matters for current individual decisions. Since that endogenous distribution is a highly dimensional object we use an approximate solution concept, whereby agents only consider a small number of moments of the distribution to forecast prices, as in Krusell and Smith, 1998. Specifically, agents forecast bond prices in both aggregate states in the future using a linear law of motion conditional on the current cross-sectional variance of bond holdings, the current aggregate state and the current bond price. These forecasts are, as typical in these economies, extremely accurate.

¹³Note that the remaining transition probabilities are pinned down by the assumption that unemployment is constant within good and bad times, as in Krusell and Smith (1998).

Table 5: Sensitivity parameters β (in percent) in the bond economy

	β	β^{YL}	β^{YH}
Stationary	5.964		
Constant r	5.119	5.609	4.676
Independent	9.288	10.114	8.945
Procyclical persistence	9.341	9.893	9.101
Countercyclical risk	9.346	10.439	8.899

The table reports, for the bond economy, the regression coefficient β in (1) (in percent, column 1), as well as in times of low and high aggregate endowments (β^{YL} and β^{YH} , columns 2 and 3 respectively). The coefficients are calculated in the stationary economy without aggregate fluctuations (row 1), an economy with aggregate fluctuations in endowments and acyclical unemployment risk at an exogenous interest rate equal to the equilibrium rate in the stationary economy (row 2); and general equilibrium economies with aggregate fluctuations in endowments and acyclical unemployment risk (row 3), procyclical income persistence (row 4), and countercyclical income risk (row 5). To be consistent with CEX data, consumption and income growth are calculated as 4-quarter changes in quarterly consumption and in (backward-looking) annual incomes.

3.4.3 Results for the bond economy

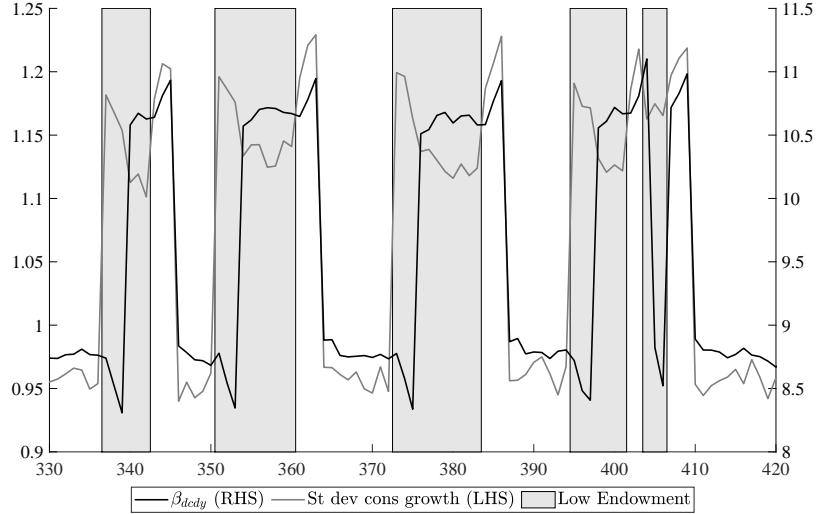
Table 5 reports key features of the benchmark bond economy, when we set σ to its benchmark value of 1.5, and choose a borrowing limit B equal to twice average quarterly income. With a discount factor $\delta = 0.9945$ this implies an insurance coefficient β in the stationary economy (without fluctuations in aggregate endowments or income risk) approximately equal to 6 percent, in line with values observed in CEX data. Table 5 reports the regression coefficient β on average (column 1), as well as those observed in good and bad times (β^{YH} and β^{YL} in columns 3 and 2, respectively), for 5 versions of the bond economy, keeping the discount factor δ unchanged. When aggregate endowments fluctuate, but income risk is acyclical and the interest rate kept constant at its equilibrium level in the stationary economy (row 2), the sensitivity of consumption to income shocks β falls by about 15 percent on average, as the addition of aggregate risk makes consumers accumulate a small positive stock of bonds on average that serves as an additional buffer also against idiosyncratic shocks. Importantly, even in this partial equilibrium of the bond economy, insurance is procyclical ($\beta^{YH} < \beta^{YL}$), due to the concavity of the consumption function at given interest rates: as higher incomes in good times move consumers away from

their borrowing constraint, consumption becomes less sensitive to income shocks.

Insurance is substantially lower in the general equilibrium bond economy with aggregate fluctuations (rows 3 and 4 of Table 5) than with constant interest rates: the sensitivity of consumption to income changes β is about a third higher. Importantly, the procyclicality of insurance is strengthened further. This is because mean reversion in aggregate endowments raises expected income growth in bad times, making equilibrium interest rates countercyclical and raising the sensitivity of consumption to income shocks in bad times in line with (3). With cyclical income risk (rows 4 and 5), average consumption smoothing deteriorates further slightly. Even with substantial procyclical persistence (row 4), implying a stronger effect of current idiosyncratic shocks on permanent incomes in booms, consumption remains less sensitive to income when output is high. And when we account for the countercyclical unemployment risk (row 5) observed in US employment data, the sensitivity is even more countercyclical (as individuals who lose their job in bad times expect unemployment to last longer and thus reduce consumption by more).

Figure 2 presents a time series of the coefficient β and the standard deviation of consumption growth from a simulation of the general equilibrium bond economy with independent transitions. The variance of both moments is dominated by differences between good and bad times (noting the time-lag implied by the CEX-consistent 4-quarter growth rates and the annual nature of output in (1)), with little history dependence within aggregate states. Both moments show spikes when the aggregate state changes, however, as consumption levels adjust discretely to the new aggregate state (and thus increase the standard deviation of consumption growth), and asset-poor unemployed moving to employment dominate the covariance of income and consumption growth, which boosts (reduces) the regression coefficient when aggregate income growth is positive (negative).

Figure 2: Dynamics of consumption insurance in the bond economy



The figure shows time series of the regression coefficient β (in percent, right-hand axis) in (1) and the cross-sectional standard deviation of consumption growth (in percent, left-hand axis) in simulations of the bond economy of Section 3.4.3. To be consistent with CEX data, consumption and income growth are calculated as 4-quarter changes in quarterly consumption and in (backward-looking) annual incomes.

3.4.4 A Krusell and Smith (1998) economy

Both the countercyclical nature of interest rates and the strong effect of additional risk on average insurance in the bond economy are likely to hinge on the inability of individuals to accumulate assets on average. The remainder of this section thus looks at a version of our economy where consumers can build a buffer to income fluctuations by accumulating (a non-negative amount of) productive capital that is an input to a standard neoclassical production function and earns a competitive rate of return r_t , as in Krusell and Smith (1998). Specifically, we choose an aggregate production function of the Cobb-Douglas form, with capital share $\alpha = 0.36$, and a linear depreciation rate of 2.5 percent per quarter. Aggregate fluctuations in total factor productivity are such that (in the case of independent idiosyncratic risk, implying a constant unemployment rate) output fluctuations equal the endowment fluctuations in the bond economy (in turn chosen to capture fluctuations in total fac-

Table 6: Sensitivity parameters β (in percent) in a Krusell and Smith (1998) economy

	β	β^{YL}	β^{YH}
Stationary	6.008		
Independent	6.004	6.093	5.963
Proc persistence	6.028	6.019	6.032
Counterc risk	5.999	6.201	5.905

For the economy with capital and Cobb-Douglas production, the table reports the cross-sectional regression coefficient β , in percent, on average (column 1), as well as in times of lower- (β^{YL} , column 2) and higher-than-average output (β^{YH} , column 3) without fluctuations (row 1), with idiosyncratic risk independent of aggregate fluctuations (row 2), with procyclical persistence (row 3), and with countercyclical unemployment risk (row 4). To be consistent with CEX data, consumption and income growth are calculated as 4-quarter changes in quarterly consumption and in (backward-looking) annual incomes.

tor productivity in the US economy across good and bad times). Employed agents earn the equilibrium wage paid by competitive firms w_t , and pay taxes to finance a balanced-budget unemployment insurance scheme that finances benefits for the unemployed with an unchanged replacement rate of 60 percent. Transition probabilities are unchanged with respect to the bond economy. Appendix VII describes the environment in more detail.

Table 6 reports the regression coefficients β for a discount factor δ equal to 0.9656, which, again, implies a value of β in the stationary economy approximately equal to 6 percent. Note that this implies a quarterly capital-output ratio of approximately 6, lower than typical values for advanced economies, but implying that individuals hold substantial assets on average. The message of the table is simple: relative to the bond economy, the effect of aggregate fluctuations on consumption smoothing is strongly reduced by the possibility to accumulate productive assets as a buffer against income fluctuations. The sensitivity of individual consumption to idiosyncratic income changes β is mildly countercyclical ($\beta^{YH} < \beta^{YL}$, in contrast to CEX data), apart from the case of procyclical persistence (row 3), where the stronger effect of current income on permanent income in good times counteracts the effect of stronger insurance due to higher resources there, leaving the sensitivity β acyclical.

Importantly, the negative effect of aggregate risk on average consumption smoothing in the bond economy is completely absent in this economy with capital.

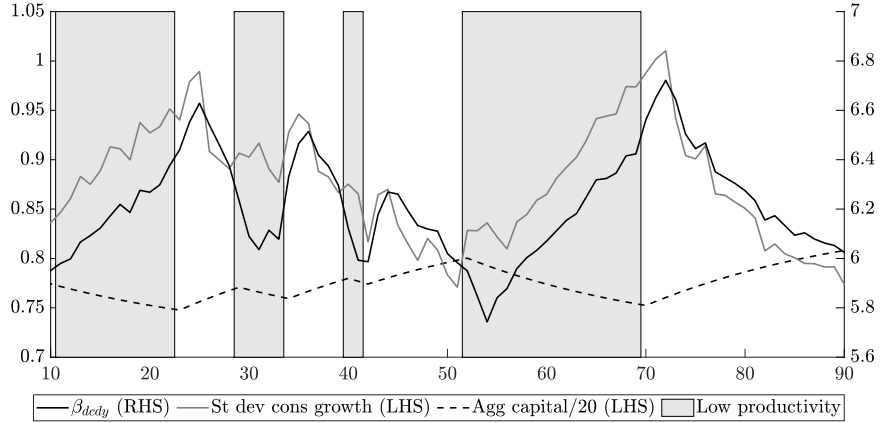
The smaller effect of aggregate fluctuations in the economy with capital has two origins: first, the availability of aggregate assets allows consumers to avoid, most of the time, the low asset levels where consumption responds strongly to income changes and the consumption function has substantial curvature. When most consumers are located in the range of asset holdings where the consumption function is (approximately) linear, however, movements in the mean and dispersion of the asset distribution caused by aggregate fluctuations in productivity have much less effect on average insurance. Second, although high rates of return on capital implied by a low discount factor increase the average effect of current income shocks on permanent income, this effect is approximately constant across the business cycle. This is because the two separate sources of fluctuations in the marginal return to capital (aggregate TFP shocks and fluctuations in the stock of aggregate capital) imply, approximately, acyclical rates of return.

Figure 3 shows that the degree of insurance in the economy with capital is, in fact, mildly more variable than suggested by Table 6. More than with output and productivity, however, variations in insurance are correlated with fluctuations in the capital stock. In particular, the sensitivity coefficient β and the dispersion of growth rates fall when the capital stock, and thus the the average buffer against idiosyncratic shocks, rises.

One caveat to the results in this subsection is that our parameterization, without any other sources of heterogeneity than that implied by heterogeneous histories of idiosyncratic income shocks, implies wealth holdings that are more concentrated than in US data. In particular, the wealth distribution that underlies the results in Table 6 has too few individuals at low asset levels, where consumption reacts strongly to income changes.¹⁴ The results may thus understate the cyclical fluctuations in insurance. This is partly intentional: we think that the two parameterisa-

¹⁴See, e.g., Krueger et al. (2016) for a discussion of this feature of Krusell and Smith (1998)-type economies and extensions that imply a more realistic wealth distribution

Figure 3: Dynamics of consumption insurance in a Krusell and Smith (1998) economy



The figure shows time series of the regression coefficient β (in percent, right-hand axis) in (1) and the cross-sectional standard deviation of consumption growth (in percent, left-hand axis) as well as the capital stock (divided by 20 to fit on the left-hand axis) in simulations of the economy with capital of Section 3.4.4, with independent individual transitions. To be consistent with CEX data, consumption and income growth are calculated as 4-quarter changes in quarterly consumption and in (backward-looking) annual incomes.

tions, of a bond economy without net assets on the one hand, and an economy where all agents use productive assets as a buffer against income shocks on the other, span the possible range of outcomes in incomplete markets models, and thus provide a robust prediction of procyclical consumption smoothing. In fact, we would expect additional sources of heterogeneity, such as in discount factors or along the life-cycle, to increase the counterfactual cyclical features of the incomplete markets model by raising the mass of individuals at low assets. Table 8 in Appendix V provides support for this conjecture, by depicting the sensitivity parameter β in an economy with capital and production that is identical to our benchmark in this section, but where 10 percent of individuals have a lower discount factor $\delta' = \delta - 0.002$. With $\delta = 0.98$, this yields an average coefficient β similar to the benchmark parameterisation in this section, but increases the mean capital (by about 50 percent in the stationary

economy) as well as its variance (almost exactly three-fold), and strongly widens the left-hand-tail of capital holdings (the 5th percentile, equal to about 1.5 times average wages in the stationary benchmark economy, falls by approximately 40 percent). In line with our intuition, the sensitivity of individual consumption to idiosyncratic income changes in Table 8 is substantially more countercyclical ($\beta^{YH} < \beta_{YL}$) in all versions of that model than that with homogeneous discount factors in Table 6.

4 Cyclical consumption insurance with endogenous financial frictions

According to the previous section, standard incomplete markets models of consumption smoothing predict consumption to be less sensitive to income changes in booms, in contrast to the procyclical sensitivity in CEX data documented in Section 2. In the analysis, both the degree of market incompleteness and borrowing limits were exogenous and thus acyclical. The strong effect of a positive asset supply on the degree and cyclicity of consumption-sensitivity, however, as well as Werning (2015)’s result of a constant sensitivity with proportional borrowing limits, suggest that fluctuations in financial constraints might be an important determinant of cyclical fluctuations in consumption smoothing, and thus a potential explanation of our stylised facts in CEX data. In this section, we therefore look at an (“LC”) environment without any exogenous restrictions to financial contracts, but with endogenous financial frictions due to participation constraints that arise from limited contract enforcement.¹⁵

Relative to stationary LC economies, whose properties are well-understood (Krueger and Perri (2011), or Broer (2013)), the introduction of aggregate risk affects equilibrium insurance contracts in two ways: first, the relative costs and bene-

¹⁵Even with exogenous market incompleteness financial frictions can be made responsive to business cycles by allowing individuals to default on their debt, thus endogenising a riskless borrowing limit B (as the maximum that consumers would always pay back) or individual interest rates r_t (to account for probabilities of default). Although both approaches tend to imply a small amount of borrowing, thus limiting any effect of cyclical fluctuations in borrowing conditions on the equilibrium, they are interesting in their own right.

fits of any given transfer policy may differ between periods of high and low aggregate income. And second, because idiosyncratic risk is itself cyclical, the expected likelihood of making or receiving such transfers in the future also changes. Since both effects change the tightness of participation constraints, the effect of aggregate economic fluctuations on consumption smoothing in LC economies is fundamentally different from that in SI economies. In particular, since insurance is more valuable in times of low output, and high unemployment risk, and the outside option of financial autarky less attractive then, we would expect participation constraints to loosen in bad times. This could explain the finding of countercyclical insurance in CEX data.

4.1 Market Structure and Competitive Equilibrium

Consider an economy that shares the physical environment of the SI economy of section 3, including the structure of uncertainty and preferences, but where individuals can trade a complete set of state-contingent contracts that, however, are only enforced by the threat of exclusion from financial trade.

We choose a formulation of insurance where each individual signs a contract with an insurance provider. At the beginning of each period, individual income shares and aggregate income are revealed and individuals decide whether to honour their insurance contract or whether to move permanently into autarky, where $c_t = y_t \cdot Y_t$. The discounted utility associated with autarky is denoted by $V(y_t, Y_t)$, i.e.

$$V(y_t, Y_t) = \mathbb{E} \left[\sum_{s=0}^{\infty} \delta^s u(y_{t+s} Y_{t+s}) \middle| y_t, Y_t \right].$$

To ensure that individuals stay in their insurance contract, we must have

$$\mathbb{E} \left[\sum_{s=0}^{\infty} \delta^s u(c_{t+s}) \middle| y_t, Y_t \right] \geq V(y_t, Y_t) \quad (4)$$

for all $t = 0, 1, \dots$

There is a large number of insurance providers who offer, at time $t = 0$, mutually agreeable insurance contracts to individuals. An insurance contract is a transfer

program $\tau = \{\tau_t(y^t, Y^t)\}_{t=0}^\infty$, where

$$\tau_t : \mathbb{Y}^t \times \mathbb{y}^t \rightarrow \mathbb{R}.$$

This sequence of transfer functions defines individual consumption according to

$$c_t = y_t \cdot Y_t + \tau_t.$$

Insurance providers evaluate a transfer policy τ according to the profit function

$$P_{y_0, Y_0}(\tau) = -\mathbb{E} \left[\sum_{t=0}^{\infty} q_t(Y^t) \tau_t(y^t, Y^t) \right]. \quad (5)$$

where $q_t(Y^t)$ are the intertemporal/state prices of consumption, noting that the law of large numbers ensures that there is no uncertainty about the distribution of individual income shares in any period t . Insurance providers are constrained to deliver a given expected utility $V_0(i)$ to any individual i in period 0, where from now on we suppress the dependence on i .

The profit maximization problem of an insurance provider is to choose, for each individual it offers a contract to, a transfer policy τ that, given prices $\mathbf{q} := \{q_t(Y^t)\}_{t=0}^\infty$, maximizes (5) subject to (4) and

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \delta^s u(c_t^i) \middle| y_0, Y_0 \right] = V_0.$$

Notice that insurance providers are irrevocably committed to a contract once it is signed, but that consumers can renege at any time and move into autarky, as described above.

Finally, define a sequence of “interest rates” via

$$R_t = \frac{q_t}{q_{t-1}}.$$

We assume that asset trade starts in $t = 0$ after endowments are realized. A simple arbitrage argument implies that R_t may depend on the history of aggregate shocks Y^t and the initial income distribution of agents in period 0 π_0^j , but not on idiosyncratic income histories between 0 and t .

We define a competitive equilibrium for a given assignment of initial promised utilities V_0 as a stochastic process q , an assignment of transfer policies τ for all individuals i as a function of idiosyncratic and aggregate shocks such that

1. The transfer policy solves the insurance provider's problem given V_0 and q
2. Markets clear, i.e.

$$\int_0^1 \tau_t(i) di = 0$$

for all $t = 0, 1, \dots$ and all possible histories.

4.2 Analytical equilibrium characterization

Because insurance is more valuable in utility terms at low average levels of consumption, and the outside option of high income agents deteriorates when separation rates rise, we would expect participation constraints to loosen, and consumption to be less sensitive to income shocks, in bad times, in line with our findings in CEX data. The analytical results in this section qualify this intuition: with acyclical idiosyncratic risk, consumption insurance is acyclical with log-preferences, and, under some additional conditions, procyclical with relative risk aversion larger than 1. But we identify conditions such that, with cyclical individual risk, individual consumption is more sensitive to idiosyncratic income shocks in booms.

We concentrate on equilibria with imperfect risk sharing, i.e. where at least one participation constraint binds every period. Consumption is characterized by an income-dependent (and generally time-varying) value $c_t = c_t^j$ for constrained individuals that is independent of their individual histories, and by a standard Euler Equation for unconstrained individuals, which can be written in terms of consumption shares as

$$\tilde{c}_t = \frac{c_t}{Y_t} = (\delta \tilde{R}_t)^{-\frac{1}{\sigma}} \tilde{c}_{t-1} \tag{6}$$

where $\tilde{R}_t = R_t \left(\frac{Y_t}{Y_{t+1}} \right)^\sigma$. Denote the measure of consumption shares $\pi_t^c : \mathbb{B}([0, y^n]) \rightarrow$

$[0, 1]$ where \mathbb{B} denotes the Borel algebra (and where we exploit that $\tilde{c}_t \leq y^n$ in any equilibrium).

Call “stationary” an equilibrium in the economy without aggregate fluctuations (such that $Y^1 = Y^2 = \dots = Y^N = 1$) where the distribution of consumption shares (equal to consumption levels by the normalization of aggregate income) is constant over time, and denote the consumption share of individuals who are constrained at income y^j as \tilde{c}^j . Denote $\pi^c(\tilde{c}) : R^+ \rightarrow [0, 1]$ the discrete stationary distribution of consumption shares in this equilibrium, V_j the autarky value at income y^j , $j = 1, \dots, n$, and R the constant equilibrium interest rate. It is easy to see that with partial risk sharing and two income states ($n = 2$), such that $y \in \{y^h, y^l\}$ with $y_h > y_l$, all h, or “high”, types are constrained at a constant level c^h . The mass of low-income agents declines geometrically at rate p_l on a consumption support given by (6) starting from c^h , with a lower bound equal to $c^l = y^l$.¹⁶

4.2.1 Independence with logarithmic preferences

This section provides a benchmark “separability” result for the LC environment with any number of income states N, n : with log-preference and acyclical, multiplicative income risk, aggregate fluctuations and idiosyncratic risk are independent. This is similar to Werning (2015)’s result for SI economies, but also holds for stochastic aggregate fluctuations.

Result 1: With logarithmic preferences ($u(c_t) = \ln(c_t)$), there exists an equilibrium where the state-contingent interest rate equals

$$R_t = R \frac{Y_t}{Y_{t-1}} \quad (7)$$

and the distribution of consumption shares is the same as in the stationary distribution, such that (with a slight abuse of notation)

$$\pi_t^c(\tilde{c}) = \pi^c(\tilde{c}).$$

¹⁶See Krueger and Perri (2011) or Broer (2013) for a detailed characterization of the stationary joint distribution of consumption and income.

Proof

Let \tilde{c}_{t+s} denote the state-contingent sequence of consumption shares implied by R starting from any consumption share in the stationary distribution. To see how this sequence solves participation constraints in the economy with aggregate fluctuations with equality, note that

$$\begin{aligned} \mathbb{E} \sum_{s=0}^{\infty} \delta^s \ln(c_{t+s}) &= \mathbb{E} \sum_{s=0}^{\infty} \delta^s [\ln(\tilde{c}_{t+s}) + \ln(Y_{t+s})] \\ &\geq V(y_t) + \mathbb{E} \sum_{s=0}^{\infty} \delta^s \ln(Y_{t+s}) = V(y_t \cdot Y_t) \end{aligned} \quad (8)$$

where the inequality follows from the fact that \tilde{c}_{t+s} satisfies the participation constraints in the economy without aggregate fluctuations. Moreover, \tilde{c}_{t+s} also solves the optimality condition (6) of unconstrained agents at the prices R_t given by (7). Finally, by the definition of consumption shares in the stationary equilibrium, they sum to 1 and thus imply market-clearing in the economy with aggregate fluctuations. ■

4.2.2 CRRA preferences

In the empirically relevant case of risk aversion σ larger than 1 and cyclical income risk, it is typically impossible to characterize the time-varying joint distribution of incomes (aggregate and idiosyncratic) and consumption analytically. The remainder of this section, however, provides some benchmark results for the case of 2 aggregate and idiosyncratic states ($N = n = 2$). We denote aggregate endowments as $Y^H = 1 + \epsilon$, $Y^L = 1 - \epsilon$, and endowment shares as y^h, y^l . The probabilities of remaining in the same aggregate state are equal, denoted $P_{HH} = P_{LL} = P$, implying that average aggregate income is 1 and ϵ a measure of aggregate income fluctuations. We assume non-negative (but potentially state-dependent) serial correlation in idiosyncratic incomes ($p_{hh}^{IJ}, p; IJ_{ll} \geq \frac{1}{2}$, for $I, J \in \{H, L\}$). An intuitive, tractable measure of state-dependent insurance is the average transfer made by h agents in aggregate state I as a share of total aggregate income, denoted by $\theta^I = \mathbb{E}[y^h - \tilde{c}_t^h]$, for $I = H, L$. Finally, we define a measure of insurance in good relative to bad times by

$\Delta\theta = \theta^H - \theta^L$. We talk of “procyclical” insurance whenever this measure is positive. Note that the procyclical sensitivity of individual consumption to income changes observed in CEX data implies countercyclical insurance.

Procyclical insurance with independent transitions

This section shows how, when idiosyncratic risk is acyclical and aggregate risk iid, the introduction of small aggregate endowment fluctuations tightens participation constraints more in bad times of low average income and consumption, when current insurance transfers paid by high income agents are more costly in utility terms. Under the, somewhat restrictive, condition of independent aggregate and idiosyncratic transitions, the LC economy thus cannot replicate the countercyclical insurance / procyclical sensitivity of consumption to income changes observed in CEX data.

Assume that aggregate and individual transitions are independent across aggregate states, implying $p_{ii}^I = p_{jj}^J$, $i, j = h, l$, and $P = \frac{1}{2}$. Consider an allocation that simply scales the stationary allocation by aggregate income, such that individual consumption shares take the same values, and follow the same transitions, as in the stationary equilibrium. Denote the difference between the resulting autarky and contract values at high individual income as $\Delta^I = V^I(y_h) - U^I(\tilde{c}_h)$, $I = H, L$, where \tilde{c}^h is the constant consumption share of high-income agents in the stationary distribution. Define $\Delta = \Delta^H - \Delta^L$. Finally, note that, without aggregate fluctuations ($\epsilon = 0$), the proposed allocation is trivially an equilibrium at the stationary interest rate R , implying $\Delta^I = 0$, $I = H, L$.

Lemma 1: With independent transitions, as ϵ rises from 0, participation constraints of high types tighten more at low aggregate income.

Proof:

$$\frac{\delta\Delta}{\delta\epsilon}\Big|_{\epsilon=0} = \frac{\delta}{\delta\epsilon}\Big|_{\epsilon=0} \left[\frac{(y_h^{1-\sigma} - c_h^{1-\sigma}) [(1+\epsilon)^{1-\sigma} - (1-\epsilon)^{1-\sigma}]}{1-\sigma} \right] \tag{9}$$

$$= 2(y_h^{1-\sigma} - c_h^{1-\sigma}) < 0 \tag{10}$$

$$\Leftrightarrow \frac{\delta\Delta^L}{\delta\epsilon}\Big|_{\epsilon=0} > \frac{\delta\Delta^H}{\delta\epsilon}\Big|_{\epsilon=0} \tag{11}$$

where the first equality exploits the assumption that aggregate incomes are iid, so expected continuation utilities in autarky and with stationary consumption shares are independent of current aggregate income (and therefore cancel in the difference). The inequality then follows from the declining power function. ■

Result 2: With independent transitions, insurance is procyclical at low levels of ϵ .

Proof: Write $\tilde{c}^{hI} = 1 + e^I$, $I = H, L$ and totally differentiate Δ^I to get

$$\frac{de^I}{d\epsilon} \Big|_{\epsilon, e=0} = \frac{\frac{\delta \Delta^I}{\delta \epsilon} \Big|_{\epsilon, e=0}}{\frac{\delta \Delta^I}{\delta e^I} \Big|_{\epsilon, e=0}} \quad (12)$$

Note that $\frac{\delta \Delta^H}{\delta e^I} \Big|_{\epsilon, e=0} = \frac{\delta \Delta^L}{\delta e^I} \Big|_{\epsilon, e=0}$. Together with (11), this implies that the difference in insurance transfers between H and L states increases with ϵ : $\frac{\delta \Delta^{\theta}}{\delta \epsilon} \Big|_{\epsilon, e=0} = \left[\frac{de^L}{d\epsilon} - \frac{de^H}{d\epsilon} \right] \Big|_{\epsilon, e=0} > 0$. So insurance transfers turn procyclical as aggregate fluctuations rise from 0. By continuity they are also procyclical for sufficiently small $\epsilon > 0$.

■

Countercyclical insurance with procyclical outside options

This section shows how, when individual income risk is not independent of aggregate fluctuations but such that incentive constraints tighten in good times, the model predicts countercyclical insurance, in line with the data, as long as aggregate endowments are sufficiently persistent.

Result 3: When individual transition probabilities are such that participation constraints are more binding at high aggregate income with full persistence ($P = 1$), there is a value ν such that insurance is countercyclical whenever $P > 1 - \nu$.

Proof: Consider a version of the economy where aggregate incomes are fully persistent ($P = 1$). Consider two pairs of values p_{hh}^{HH}, p_{ll}^{HH} and p_{hh}^{LL}, p_{ll}^{LL} such that the consumption share of high types is higher in booms. By the assumption that an equilibrium with partial insurance exists, and by the fact that transition probabilities change autarky values, such pairs exist. Now consider a small $d\nu > 0$ that introduces transitions between H and L states. Since contract values and autarky values are continuous in ν at $\nu = 0$, the result follows. ■

One example where the result holds is when income risk is symmetric and persistence procyclical, such that $p_{hh}^{HH} = p_{ll}^{LL} > p_{hh}^{LL} = p_{ll}^{HH}$. Intuition suggests it may also hold when income risk is countercyclical: lower risk in booms could then reduce the value of insurance to tighten participation constraints and lower insurance. We discuss these possibilities in the context of the quantitative analysis below.

Aggregate risk and the degree of insurance

Two forces act to decrease average insurance in LC economies with aggregate income fluctuations. The first follows from Result 2: transfers are smaller at low aggregate income as participation constraints tighten in bad times. The second force simply stems from history dependence in consumption shares: tighter participation constraints in bad times raise high-type consumption shares immediately. Looser participation constraints in good times, in contrast, increase transfers only slowly, as inherited relative consumption shares are unchanged among unconstrained individuals according to (6). In the general case of a time-varying joint distribution of (aggregate and individual) consumption and income, the effect of these forces is difficult to show analytically. The following two examples, however, consider benchmark cases where participation constraints of high types hold with equality, respectively, at perfect insurance and in autarky in the stationary economy.

Example 1: Consider a case of independent aggregate transitions, ($P = \frac{1}{2}$), and preference parameters and individual transition probabilities such that the participation constraint of high types holds with equality at equal consumption shares in the stationary economy ($\tilde{c}^i = 1, \forall i$). In this case, increasing aggregate uncertainty from 0 ($d\epsilon > 0$) reduces insurance.

Proof: Since participation constraints initially hold with equality at perfect insurance, but tighten in the low aggregate state when ϵ rises (Lemma 1), perfect insurance violates participation constraints as ϵ rises above 0. ■

Example 2: Consider a case of independent aggregate transitions, ($P = \frac{1}{2}$), and preference parameters and individual transition probabilities such that the participation constraint of high types holds with equality in autarky ($\tilde{c}^i = y^i, \forall i$). In this

case, increasing aggregate uncertainty from 0 ($d\epsilon > 0$) increases insurance. ■

Proof: Since participation constraints initially hold with equality in autarky, but loosen in the high aggregate state when ϵ rises (Lemma 1), intermediaries can offer insurance to high types as ϵ rises above 0. ■

4.3 Quantitative analysis

The results in the previous section highlight the different forces that affect the cyclicity, and the average level, of consumption insurance in LC economies. In the empirically relevant case where both aggregate endowments and idiosyncratic income risk are cyclical and insurance is partial, however, it is impossible to characterise the resulting cyclicity analytically. The rest of this section therefore studies the LC economy quantitatively.

4.3.1 Model solution

In LC economies, neither a state space reduction of the type we used in the previous section, following Krusell and Smith (1998), nor linearization techniques (Boppart et al., 2018, Reiter, 2009) are feasible. To see this, note that linearization is based on a stationary consumption distribution associated with a version of the economy without aggregate fluctuations. In our limited commitment economy, however, the equilibrium degree of consumption risk sharing is a highly non-linear function of aggregate shocks. In fact, in line with Result 2, it is easy to construct examples (of the type of example 1) where insurance is perfect in the absence of aggregate fluctuations, but limited when the latter are sufficiently large. In this case, there is no unique distribution that a linearization approach could be based on.

Similarly, state space reduction techniques are difficult to apply by the fact that, even in the two-type economy, market-clearing state prices of consumption depend on the time-varying mass of constrained agents (whose consumption is, for given future prices, unaffected by a change in current prices, such that the higher their

mass, the more need current state prices vary to clear markets). Forecasting future interest rates (that determine future consumption and thus current contract values via (6)) thus requires forecasts of the mass of constrained agents. That mass, however, depends on the mass of agents that become constrained every period, and thus on the entire current consumption distribution of unconstrained agents.

In Appendix *II* we show how the quasi-analytical characterization of the consumption distribution in the stationary version of the model (Krueger and Perri, 2011, Broer, 2013) can be used to solve its version with aggregate fluctuations. Specifically, we show that, in the economy with $N = n = 2$, as long as the mass of agents constrained at low income is negligible (which we check in all our computations), there exists an equilibrium that is characterised by values of interest rates and high-type consumption that are history-independent, in the sense that they only depend on the current and, in the case of interest rates, previous, aggregate state. The consumption distribution of unconstrained agents, in contrast, has full history-dependence, as it depends on the sequence of realised interest rates, and thus aggregate shocks, in preceding periods through (6).

4.3.2 Results

This section reports results from simulations of the LC economy, assuming, as in Section 3.4.3, a constant relative risk aversion equal to 1.5, and using the process for aggregate endowments, as well as the three processes for individual endowment shares with acyclical or independent transitions, with procyclical persistence, and countercyclical risk from Table 4. Table 7 reports the regression coefficient β for different parameterisations of the LC economy. Again, we choose the discount factor δ such that the average sensitivity of consumption to income changes in the stationary economy (without aggregate fluctuations, row 1) approximately equals 6 percent, in line with the values observed in CEX data. As known from the analysis of stationary LC environments, the resulting value of about 0.72 is substantially lower than those in the SI economy. Row 2 to 4 illustrate the theoretical results

Table 7: Cyclical consumption insurance in a simple limited-commitment economy

	β	β^L	β^H
Stationary	5.97		
Proc persistence, $Y^H = Y^L$	6.55	4.95	7.18
Counterc risk, $Y^H = Y^L$	6.05	8.62	5.04
Independent	5.97	6.91	5.60
Proc persistence	6.53	5.86	6.80
Countercyclical risk	6.05	9.60	4.65

The table reports the regression coefficient β in (1) in simulations of different versions of the LC economy on average (column 1), as well as in times of low and high aggregate productivity (β_{Y^L} and β_{Y^H} , columns 2 and 3, respectively). To be consistent with CEX data, consumption and income growth are calculated as 4-quarter changes in quarterly consumption and in (backward-looking) annual incomes.

in Section 4.2.2 by introducing cyclical income risk in the stationary LC economy without aggregate endowment fluctuations (thus maintaining $Y^H = Y^L = 1$). In line with Result 3, with procyclical income persistence ($p_{ll}^H > p_{ll}^L$, $p_{hh}^H > p_{hh}^L$, in row 2), participation constraints bind more in good times, thus implying a procyclical sensitivity of consumption to income shocks ($\beta^{Y^H} > \beta^{Y^L}$), as in CEX data.¹⁷ At the high level of consumption risk sharing found in CEX data, however, our preferred parameterisation of unemployment risk, with countercyclical probabilities of unemployment ($p_{ll}^L > p_{ll}^H$, $p_{hl}^L > p_{hl}^H$, row 3), implies countercyclical sensitivity ($\beta^{Y^H} < \beta^{Y^L}$). This result is, however, not general: with lower or less cyclical job-finding rates, the LC economy with countercyclical risk is consistent also with a procyclical sensitivity of consumption to income changes (although the differences between $\beta^{Y^H} > \beta^{Y^L}$ are typically small in magnitude).¹⁸ Importantly, as suggested by example 1, at the high level of insurance implied by CEX data, cyclical transition probabilities reduce the average level, particularly with procyclical persistence.

Row 4 of Table 7 adds the benchmark structure of aggregate endowment fluctua-

¹⁷The magnitude of the implied coefficient in a regression of β on log-deviations of aggregate endowments from their mean is with 39.5 close to those found in Table 2. But since the procyclical persistence parameters are only illustrative, we do not emphasize this result.

¹⁸Quantitative examples are available from the authors.

tions to the stationary economy, maintaining acyclical individual transitions. As proven in Result 2 for small aggregate fluctuations, this results in a countercyclical sensitivity of consumption to individual income shocks. Although this effect is strong, it does not suffice to overturn the procyclical sensitivity of consumption to income changes implied by procyclical income persistence. The cyclical variation with procyclical persistence and aggregate endowment fluctuations (row 6) is thus again in line with CEX data, if only about half that with constant aggregate endowments. Finally, row 7 of Table 7 reports results for the economy with aggregate fluctuation and the benchmark structure of cyclical income risk: both effects add up to make consumption respond substantially more to income shocks in bad times, in contrast to CEX data.

An important feature of LC economies is typically a pronounced asymmetry between consumption responses to income increases (which tighten participation constraints and thus increase consumption) and decreases (which are largely insured, see Broer (2013)). In our quantitative version of the LC economy, this asymmetry is not actually very pronounced, with an average consumption response to income declines of 5.5 percent, compared to 7 percent for income rises (in our preferred parameterisation with countercyclical risk). This is because high job-finding rates make unemployment duration of more than one quarter unlikely, thus limiting the asymmetry implied by large responses of consumption of the long-term unemployed that find jobs.

Compared to the SI models in section 3, cyclical fluctuations in aggregate income and idiosyncratic risks thus cause substantially stronger fluctuations in the sensitivity of consumption to idiosyncratic shocks when financial frictions are endogenous, and thus cyclical themselves. In fact, with procyclical income persistence, consumption is more sensitive to income changes in booms, qualitatively in line with the key facts we document for CEX data. Because transfers are more costly in marginal utility terms during bad times of low aggregate incomes, the LC model tends to predict a countercyclical sensitivity, however. This is particularly true when accounting for the slightly countercyclical unemployment risk we find in U.S.

data. Overall, accounting for the cyclical nature of aggregate endowments and idiosyncratic risks reduces the average level of risk sharing (by slightly less than 10 percent with procyclical persistence), while risk sharing is unchanged with independent transition.

5 Conclusion

In this paper, we have studied how cycles in aggregate economic activity determine the ability of households to smooth consumption in the face of idiosyncratic shocks to their incomes. We first showed that in data from the US consumer expenditure survey (CEX), household consumption reacts more strongly to individual income changes in booms. This contrasts with the predictions from standard incomplete markets models of consumption smoothing where consumption reacts *less* to income changes in booms. We also showed how the magnitude of this effect strongly depends on the cyclical nature of interest rates and income risk, and particularly on the net asset supply that determines the average distance of household assets from their borrowing limit. The counterfactual prediction by the standard model motivated us to study the cyclical nature of consumption insurance in an alternative environment without exogenous restrictions to asset trade, but with endogenous, and therefore potentially cyclical, financial frictions arising from limited contract enforcement. We showed that fluctuations in consumption risk sharing there tend to be larger than with exogenously incomplete markets, and that the sensitivity of consumption to income changes may indeed be higher in booms, although we did not find this to be true in a parameterization where unemployment risk was similarly countercyclical as in US employment data. Importantly, while the average sensitivity of consumption to income changes in the SI economy with capital was not much affected by aggregate fluctuations, relative to stationary economies, it increased strongly in the bond economy, and by a small amount in the LC environment with cyclical unemployment risk.

These results point towards the importance of endogenous financial frictions for con-

sumption smoothing. Our analysis of an environment with endogenous constraints to risk sharing was intentionally simple: first, we abstracted from capital investment. Future research should therefore generalize our solution method to account for additional endogenous state variables.¹⁹ Second, the simple model we studied attributed financial constraints to only one friction: limited contract enforcement. This was because, a priori, we would expect other frictions, such as limited information about household incomes or effort levels as in Broer et al. (2017), to be less cyclical. It would nevertheless be interesting to investigate alternative environments with endogenous frictions. For example, it would be interesting to contrast the predictions from our model to one where borrowing conditions are endogenous under the maintained assumption of incomplete markets, for example because consumers can default on their debt and aggregate conditions change their incentives to do so, and thus the interest rates at which they can borrow. Although debt levels are often found to be small in such environments, the results in Chatterjee et al. (2007) and Auclert and Mitman (2018) suggest that the interaction between aggregate economic conditions and individual default decisions may be important. More generally, it would be interesting to compare the cyclical properties of the joint distribution of consumption and income growth, of which the sensitivity parameter β is just one moment, across countries with different institutions governing consumer finances, such as bankruptcy regulation, the possibility to collateralize consumer borrowing through mortgage equity withdrawal, etc.

Our results showed how, in the SI economy, the cyclical behavior of interest rates, and the cross-sectional distribution of assets, were critical for the response of consumption to income shocks in booms vs recessions. Our analysis was simple in both dimensions. In future research it would be interesting to consider other aggregate shocks (to demand, or policies), consider nominal rigidities (that make interest rates countercyclical in response to demand or fiscal shocks if they follow standard Taylor rules), and add model elements that have been found important to explain the observed wealth distribution in the US economy (such as life-cycle motives for sav-

¹⁹See Krueger and Perri (2006) or Ábrahám and Cárceles-Poveda (2010) for the analysis of stationary LC economies with capital.

ing and heterogeneity in returns, as discussed, e.g., in Krueger et al. (2016)). An additional element that may effect the relative ability of households to smooth consumption in booms vs recessions could be cyclical fiscal policy. To the extent that our measure of insurance was calculated with respect to post-tax-post-transfer income, our data analysis takes account of this. Nevertheless, it would be interesting to investigate the relationship between different fiscal policies and the cyclicity of consumption smoothing more in detail.

Both models we considered predicted the sensitivity of consumption to income changes to be more procyclical, and thus more in line with the data, when individual income shocks are more persistent in booms. In our simple application to unemployment risk, procyclical persistence, however, contrasts with the countercyclical risk we found in US employment data. More generally, it also contrasts with the findings of Storesletten et al. (2004) that persistent shocks are more important relative to transitory ones in recessions. One possibility to reconcile our results with this evidence would be misperception of consumers in the persistence of their income shocks, as in Rozsypal and Schlafmann (2017). In fact, intuition suggests that if individuals perceive shocks to be more persistent in booms, they react stronger to income changes, as observed in the data.

References

- Ábrahám, Á. and E. Cárceles-Poveda (2010). Endogenous trading constraints with incomplete asset markets. *Journal of Economic Theory* 145(3), 974–1004.
- Acharya, S. and K. Dogra (2018, February). Understanding HANK: insights from a PRANK. Staff Reports 835, Federal Reserve Bank of New York.
- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109(3), 659–84.
- Alvarez, F. and U. J. Jermann (2000). Efficiency, equilibrium, and asset pricing with risk of default. *Econometrica* 68(4), 775–797.
- Attanasio, O. P. and G. Weber (2010). Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature* 48(3), 693–751.

- Auclert, A. and K. Mitman (2018). Bankruptcy and Aggregate Demand. Technical report.
- Battistin, E. (2003). Errors in survey reports of consumption expenditures. Technical report, IFS Working Papers, Institute for Fiscal Studies (IFS).
- Bewley, T. (1977, December). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory* 16(2), 252–292.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control* 89(C), 68–92.
- Broer, T. (2013). The wrong shape of insurance? What cross-sectional distributions tell us about models of consumption smoothing. *American Economic Journal: Macroeconomics* 5(4), 107–40.
- Broer, T., M. Kapička, and P. Klein (2017). Consumption risk sharing with private information and limited enforcement. *Review of Economic Dynamics* 23, 170–190.
- Carroll, C. D. and M. S. Kimball (1996). On the concavity of the consumption function. *Econometrica: Journal of the Econometric Society*, 981–992.
- Carroll, C. D. and M. S. Kimball (2001). Liquidity constraints and precautionary saving. *NBER Working Paper* (w8496).
- Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Ríos-Rull (2007, November). A Quantitative Theory of Unsecured Consumer Credit with Risk of Default. *Econometrica* 75(6), 1525–1589.
- Chien, Y. and H. Lustig (2009). The market price of aggregate risk and the wealth distribution. *The Review of Financial Studies* 23(4), 1596–1650.
- Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics* 88, 70–89.
- Davis, S. J. (2003). [on the welfare consequences of the increase in inequality in the united states]: Comment. *NBER Macroeconomics Annual* 18, 121–132.
- De Giorgi, G. and L. Gambetti (2017). Business cycle fluctuations and the distribution of consumption. *Review of Economic Dynamics* 23, 19 – 41.

- Den Haan, W. J., K. L. Judd, and M. Juillard (2010). Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty. *Journal of Economic Dynamics and Control* 34(1), 1 – 3. Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.
- Gervais, M. and P. Klein (2010). Measuring consumption smoothing in CEX data. *Journal of Monetary Economics* 57, 988–999.
- Guvenen, F., S. Ozkan, and J. Song (2014a). The nature of countercyclical income risk. *Journal of Political Economy* 122(3), 621–660.
- Guvenen, F., S. Ozkan, and J. Song (2014b). The nature of countercyclical income risk. *Journal of Political Economy* 122(3), 621–660.
- Huggett, M. (1997). The one-sector growth model with idiosyncratic shocks: Steady states and dynamics. *Journal of Monetary Economics* 39(3), 385–403.
- Imrohoroğlu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy* 97(6), 1364–1383.
- Kaplan, G. and G. L. Violante (2018, August). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives* 32(3), 167–94.
- Kehoe, T. J. and D. K. Levine (1993). Debt-constrained asset markets. *Review of Economic Studies* 60, 865–888.
- Kimball, M. S. (1990, July). Precautionary Saving and the Marginal Propensity to Consume. NBER Working Papers 3403, National Bureau of Economic Research, Inc.
- Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and household heterogeneity. In *Handbook of Macroeconomics*, Volume 2, pp. 843–921. Elsevier.
- Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? Evidence and theory. *Review of Economic Studies* 73, 163–193.
- Krueger, D. and F. Perri (2011). Public versus private risk sharing. *Journal of Economic Theory* 146(3), 920–956.
- Krueger, D. and H. Uhlig (2006). Competitive risk sharing contracts with one-sided commitment. *Journal of Monetary Economics* 53(7), 1661–1691.
- Krusell, P. and A. Smith (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106(5).
- Lepetyuk, V. and C. A. Stoltenberg (2013). Policy announcements and welfare. *The Economic Journal* 123(571), 962–997.

- Nishiyama, S.-I. and R. Kato (2011, September). On the Concavity of the Consumption Function with a Quadratic Utility under Liquidity Constraints. TERG Discussion Papers 274, Graduate School of Economics and Management, Tohoku University.
- Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control* 33(3), 649–665.
- Rozsypal, F. and K. Schlafmann (2017). Overpersistence bias in individual income expectations and its aggregate implications.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review* 95(1), 25–49.
- Storesletten, K., C. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor-market risk. *Journal of Political Economy* 112(3), 695–717.
- Werning, I. (2015, August). Incomplete markets and aggregate demand. Working Paper 21448, National Bureau of Economic Research.

6 Appendix I: Data construction

We use a CEX sample that covers the years 1983 to 2012. We use the following main variables in the analysis: a broad nondurable consumption aggregate (including rental payments and imputed rental services for house owners), denoted $ND+$; the CEX aggregates of total nondurable consumption excluding rental payments (ND), and food consumption ($Food$); family earnings (the sum of labor earnings of the head and spouse of a household); family disposable income (including labor earnings, business or farm income, professional income, financial income and income from social security, unemployment and other benefits including food stamps) minus federal, state and local taxes paid, as reported by the household. To construct log-differences of income and consumption, we use residuals from a regression on year-quarter dummies and a number of household characteristics: the number of dependent children, the number of adults in the household, a cubic function of the household’s head’s age, a dummy that equals 1 if a spouse is present, a set of dummies capturing the head’s and spouse’s education status, and their interaction, and a full set of race dummies. We also perform our analysis using raw data,

which yields very similar results (since most of the household characteristics are unchanged and our regressors are log-differences).

We restrict our sample to households whose head is of working age (between 21 and 64 years of age), labeled as a complete income respondent (meaning the respondent provided values for major sources of income, such as wages and salaries, self-employment income, and Social Security income), who is present in all 4 quarterly interviews, and whose income is not top-coded. We also exclude households whose composition changes. For the time-varying standard deviation (Panel d) of Figure 1), we trim the log-differences in ND+ consumption at percentiles 2 and 98. All time-series that enter Figure 1 or the regressions enter as deviations from a log-linear trend.

As measures of the business cycle we consider deviations from a log-linear trend of three aggregate output or income measures: real GDP, real household disposable income from the National Income and Product Accounts (NIPA), and the sum of disposable income in the CEX sample.

7 Appendix II: Solution and computation of equilibrium in the LC economy

This section shows the existence of a ‘history-independent’ equilibrium in the economy with $N = 2$ and $n = 2$, where $y \in \{y^l, y^h\}$, $Y \in \{Y^L, Y^H\}$. The equilibrium is history-independent in the sense that the consumption of h agents, all of them constrained, only depends on the current aggregate state, and that of interest rates only on the current and preceding aggregate states, while consumption of agents that remain unconstrained at low income depends on a (potentially long) finite history of aggregate states. The intuition behind this result is that participation constraints are purely forward looking: the contract value only depends on current consumption and future interest rates that determine consumption in unconstrained future periods via (6). Moreover, when only h agents are constrained, the market clearing

condition can be written as a function of h agent consumption in the current and the preceding period and the current interest rate. We can thus write market-clearing conditions and participation constraints independently of history, which implies the existence of a history-independent equilibrium.

7.1 Existence of a ‘history-independent’ equilibrium

Result 4 below considers the economy with $N = 2$ and $n = 2$, $y \in \{y^l, y^h\}$, $Y \in \{Y^L, Y^H\}$ and independent individual transitions ($p_{ij}^{kl} = p_{ij}^{mn}$, $k, l, m, n \in \{H, L\}$), for parameters such that all high-income agents are constrained but risk sharing is strong enough for the mass of agents that are constrained at $y = y^l$ to be negligible. In this case, the participation constraint of high-income agents allows us to write a condition for the constrained level of consumption at y^h , Y^I , denoted c_t^{hI} only in terms of future prices of consumption $\{R_s\}$, $s = t, t + 1, \dots, \infty$ (which may depend on the history of aggregate shocks until s).

$$V^{hI} = u(c_t^{hI}) + \delta[p_{hh}\mathbb{P}^1 V^1 + p_{hl}\mathbb{P}^1 u^1(c_t^{hI})] + p_{hl} \sum_{s=2}^{\infty} \delta^s p_{ll}^{s-2} [p_{lh}\mathbb{P}^s V^s + p_{ll}\mathbb{P}^s u^s(c_t^{hI})] \quad (13)$$

where we neglected the possibility to become constrained at low income in line with the condition of the result, and

$$u^s(c_t^{hI}) = u((\delta^s \Pi_{k=1}^s R_{t+k})^{\frac{1}{\sigma}} c_t^{hI}), s = 1, 2, \dots$$

is the $2^s \times 1$ vector of current utilities from consumption after s periods of low income, and thus s unconstrained transitions of consumption at 2^s possible interest rate sequences $\{R_{t+1}, \dots, R_{t+s}\}$ (corresponding to 2^s possible aggregate histories between $t + 1$ and $t + s$) according to (6). \mathbb{P}^s is the corresponding vector of probabilities of each sequence. And V^s with typical element V^{hI} is the vector of autarky values at high income in the last period of each sequence. Note that c_t^{hI} is history-dependent only insofar future interest rates $R_{t+1}, R_{t+2}, \dots, R_{t+s}$ are.

The market-clearing condition can be expressed in terms of the change in the consumption share of individuals whose current income is high vs of those whose in-

come is low. Take first-differences of individual consumption shares and integrate over all agents whose income is currently high and low respectively, and impose that no low-income agents are constrained, to get

$$0 = \int_{i:y_{it}=y_h} \left(\frac{c_{it}}{Y^I} - \frac{c_{it-1}}{Y^J} \right) di + \int_{i:y_{it}=y_l} \left(\frac{c_{it}}{Y^I} - \frac{c_{it-1}}{Y^J} \right) di \quad (14)$$

$$= \pi_h \frac{c_t^{hI}}{Y^I} - \int_{i:y_{it}=y_h} \left(\frac{c_{it-1}}{Y^J} \right) di + \left[(\delta \tilde{R}_t^{IJ})^{\frac{1}{\sigma}} - 1 \right] \int_{i:y_{it}=y_l} \frac{c_{it-1}}{Y^J} di$$

$$= \pi_h \left[\frac{c_t^{hI}}{Y^I} - \left(p_{hh} \pi_h \frac{c_{t-1}^{hJ}}{Y^J} + p_{lh} (1 - \pi_h) \left(1 - \frac{c_{t-1}^{hJ}}{Y^J} \right) \right) \right] \quad (15)$$

$$+ (1 - \pi_h) \left[(\delta \tilde{R}_t^{IJ})^{\frac{1}{\sigma}} - 1 \right] \left(p_{hl} \pi_h \frac{c_{t-1}^{hJ}}{Y^J} + p_{ll} (1 - \pi_h) \left(1 - \frac{c_{t-1}^{hJ}}{Y^J} \right) \right) \quad (16)$$

where $Y_{t-1} = Y^J$ is the aggregate income value in period $t - 1$, $\tilde{R}_t^{JI} = R_t^{JI} \left(\frac{Y^I}{Y^J} \right)^\sigma$ and we exploit the fact that the consumption share of unconstrained agents in $t - 1$ equals $(1 - \frac{c_{t-1}^{hJ}}{Y^J})$. Given c_{t-1}^{hJ} and c_t^{hI} , equation (16) defines a unique value R_t^{IJ} that is consistent with market clearing when the economy transits from Y^J to Y^I between periods $t - 1$ and t .

Note that with iid individual incomes, average consumption shares in the *previous* period of *current* high and low income individuals both equal 1. Equation (16) thus reduces to

$$\pi_h \left[\frac{c_t^{hI}}{Y^I} - 1 \right] + (1 - \pi_h) \left[(\delta \tilde{R}_t^{IJ})^{\frac{1}{\sigma}} - 1 \right] = 0 \quad (17)$$

which implies that \tilde{R}_t^{IJ} is strictly decreasing in c_t^{hI} (as the consumption of the unconstrained must fall by more in (6), implying a lower R_t^I , when the consumption of the constrained is higher). Also, Equation (6) shows how consumption of the unconstrained in aggregate state I in period t is strictly increasing in R_t^I .

Result 4: Consider $N = 2$ and $n = 2$, such that $y \in \{y^l, y^h\}$, $Y \in \{Y^L, Y^H\}$ and assume that an equilibrium with partial risk sharing exists and that full risk sharing is not incentive compatible. Assume, for simplicity, that transitions of individual income shares are iid over time ($p_{ij} = p_i, \forall j = 1, 2, \dots, n$). When u exhibits CRRA with RRA σ and parameters are such that all high-income agents are constrained

but the mass of agents that are constrained at $y = y^l$ is negligible, there exists a ‘history-independent’ equilibrium, where the interest rate between periods t and $t + 1$ depends only on the aggregate income in those periods so that

$$R_{t+1} = R^{I,J} = R^I \left(\frac{Y_I}{Y_J} \right)^\sigma \quad (18)$$

when $Y_{t+1} = Y^I$ and $Y_t = Y^J$. Moreover, individuals with high income experience only two consumption levels c_h^L, c_h^H that only depend on aggregate income $Y^I, I \in \{L, H\}$.

Proof: We show existence of an equilibrium that does not depend on history by construction. Consider the autarky equilibrium with $c^{hI} = y^h Y^I, I \in \{H, L\}$, and a pair of small changes $dc^{hI} < 0$, implying $dR^I > 0$ according to (17). By the assumption that an equilibrium with partial risk-sharing exists, this makes participation constraints in both aggregate states slack. It also satisfies resource constraints. Now reduce c^{hI} further, adjusting interest rates in states where aggregate income equals Y^I in the same fashion, until the participation constraint binds in I . Do the same for state $J \neq I$, noting that by increasing interest rates in state J , this makes participation constraints slack in state I (as consumption in future unconstrained states J rises for those currently at c^{hI} according to (6)). Iterate until both participation constraints bind, which implies partial risk sharing by the assumption that full risk sharing is not incentive compatible. ■

In practice, we solve for $c^{hI}, R^I, I \in \{H, L\}$ using a numerical solver. Note that this ‘history-independent’ equilibrium is history-independent only insofar as interest rates and constrained levels of consumption of high types are concerned. Consumption of low-income agents depends on the history of aggregate states (that defines interest rate sequences). Equilibria where individual transition probabilities are not iid ($p_{ij} \neq p_{i,k}, k \neq j$) or where individual transition probabilities depend on aggregate states can be constructed in the same fashion (although in the latter case, (18) does not hold, as interest rates depend on aggregate transitions not just through output growth). Similarly for equilibria with more aggregate states $N > 2$. Equilibria with more idiosyncratic states $n > 2$ can only be constructed in

this fashion when the mass of constrained agents at all but one income values is negligible.

7.2 Computation

Suppose $n = 2$, $N = 2$ and that

$$u(c) = \lim_{\tilde{\sigma} \rightarrow \sigma} \frac{c^{1-\tilde{\sigma}} - 1}{1 - \tilde{\sigma}}.$$

Since (17) and (13) evaluated at $I = L, H$ provide four equations in four unknowns c^{hI}, R^I , $I = L, H$, one can solve for the history-independent equilibrium easily using a computer. In this equilibrium, the assumption that no agent is constrained at low income has to be verified (in the sense that the probability mass function of consumption shares declines to negligible values at its lower bound y^l). Since consumption at low income is history-dependent, one could simulate the economy using the history-independent values of c^{hI}, R^I , $I = L, H$. In practice, we proceed as follows: Given R^1, R^2, \bar{c}^L and \bar{c}^H , construct two finite grids for consumption, each appropriate for a distinct aggregate state. Denote them by \mathfrak{c}^1 and \mathfrak{c}^2 . Sometimes we will consider these as column vectors whose elements are in ascending order.

We cannot be sure that, in fact, consumption is confined to a finite set, so this is an approximation, but one that can be made arbitrarily accurate by making the grids fine enough. The first element of \mathfrak{c}^I is $y^1 Y^i$ and the last element of \mathfrak{c}^I is c^{hI} .

Now we construct a probability transition function $P(j, m|i, k)$ for all the grid points. To be explicit, $P(j, m|i, k)$ is the probability that aggregate output will be Y^j and individual consumption will be \mathfrak{c}_m^j in the next period if, in the current period, aggregate output is Y^i and individual consumption is \mathfrak{c}_k^i .

Suppose we start in an arbitrary period t at grid point k in grid i so that this period's aggregate output is Y^i and this period's individual consumption is \mathfrak{c}_k^i . Then there are several possibilities for the next period.

With probability $\Pi(j|i)$ we transit to aggregate state j . For each such j , there are

two possibilities. With probability $\pi(2)$, next period's individual income share is high and $c_{t+1} = \bar{c}^j$. That is the easy case. Alternatively, with probability $\pi(1)$, the individual income share is low. Then it is possible that the agent is unconstrained in the next period. If so, next period's consumption satisfies

$$c_{t+1} = (\delta R^{i,j})^{1/\sigma} \mathbb{C}_k^i$$

If this c_{t+1} satisfies $c_{t+1} \geq y^1 Y^j$, then we have the correct value for c_{t+1} . Otherwise $c_{t+1} = y^1 Y^j$.

At this point, we may find that c_{t+1} is nowhere to be found on the grid \mathbb{C}^j . However, by construction, we can find an integer m so that

$$\mathbb{C}_m^j \leq c_{t+1} \leq \mathbb{C}_{m+1}^j.$$

So we assign some of the probability mass $\Pi(j|i)\pi(1)$ to \mathbb{C}_m^j and the rest to \mathbb{C}_{m+1}^j . More precisely,

$$P(j, m|i, k) = \pi(1)\Pi(j|i) \frac{\mathbb{C}_{m+1}^j - c_{t+1}}{\mathbb{C}_{m+1}^j - \mathbb{C}_m^j}$$

and

$$P(j, m+1|i, k) = \pi(1)\Pi(j|i) \frac{c_{t+1} - \mathbb{C}_m^j}{\mathbb{C}_{m+1}^j - \mathbb{C}_m^j}.$$

Subtlety: the step down from \bar{c}^j may be less than the distance between gridpoints. Then the total probability mass associated with the destination \bar{c}^j may come from distinct contingencies. The above formula then needs to be modified so that

$$P(j, m+1|i, k) = P(j, m+1|i, k) + \pi(1)\Pi(j|i) \frac{c_{t+1} - \mathbb{C}_m^j}{\mathbb{C}_{m+1}^j - \mathbb{C}_m^j}$$

where it is understood that P is initially set to all zeros and the formula is applied for each distinct contingency (high or low individual income share).

Now organize all these transition probabilities into a matrix P and stack the two grids according to

$$\mathbb{C} := \begin{bmatrix} \mathbb{C}^1 \\ \mathbb{C}^2 \end{bmatrix}.$$

Then the value of staying in the prescribed risk sharing arrangement can be written as

$$V_0 = \sum_{t=0}^{\infty} \delta^t \mu^0 P^t u(\mathbb{c}) = \mu^0 [I - \delta P]^{-1} u(\mathbb{c})$$

where the row vector μ^0 is the initial distribution over \mathbb{c} , typically a unit vector with all mass concentrated at a candidate value of constrained consumption of h agents c^{hI} . Given guesses for $c^{hI}, R^{IJ}, I, J \in \{H, L\}$ we can construct P and check participation constraints and market clearing. This allows us to solve for the history-independent equilibrium using a computer. The advantage to simulation is that the equilibrium P immediately defines transitions.

8 Appendix III: Welfare costs of cyclical fluctuations in aggregate consumption and the degree of insurance

This section derives the welfare cost approximation mentioned in the introduction using a simple static example. Write individual consumption as the product of an individual consumption share \tilde{c} and aggregate income Y , so $c = \tilde{c}Y$. Assume that there is no insurance against aggregate fluctuations, so $c = \tilde{c}Y$, but that idiosyncratic income shocks are partially insured: $\tilde{c} = 1 + \beta y$, such that only a time-varying fraction β of idiosyncratic shocks to income share y passes through to consumption. Assume that shocks to idiosyncratic and aggregate income are independent of each other, with constant means equal to $\bar{Y} = 1$ and 0 and variances V_Y and V_y , respectively. The cost of aggregate fluctuations is then approximately

$$E[u(c) - u(\tilde{c} \bar{Y})] \approx -\frac{\sigma}{2} \left[V_\beta + \frac{V_Y}{V_y} \right] V_y \quad (19)$$

where V_β is the variance of β . To see this, approximate utility using a second order Taylor expansion of consumption around $E[\tilde{c}] = 1$ and $E[Y] = \bar{Y}$

$$E[u(c)] \approx \frac{\bar{Y}^{1-\sigma}}{1-\sigma} + \bar{Y}^{1-\sigma} \left(\delta c + \frac{\delta Y}{\bar{Y}} \right) - \frac{\sigma}{2} \cdot \bar{Y}^{1-\sigma} \left(\delta c^2 + \left(\frac{\delta Y}{\bar{Y}} \right)^2 + \delta c \frac{\delta Y}{\bar{Y}} \right) \quad (20)$$

Taking expectations after deriving a similar second-order approximation for $u(\tilde{c}\bar{Y})$ and substituting $\delta c = \beta\delta y$ yields

$$\begin{aligned} \mathbb{E}[u(c) - u(\tilde{c} \cdot \bar{Y})] &\approx -\frac{\sigma}{2}\bar{Y}^{-(\sigma+1)}\mathbb{E}\left[(\beta^2 - \bar{\beta}^2)V_y + V_Y\right] \\ &= -\frac{\sigma}{2}\left[V_\beta + [\hat{\beta}^2 - \bar{\beta}^2] + \frac{V_Y}{V_y}\right]V_y \end{aligned} \quad (21)$$

where $\hat{\beta}$ denotes the average of β in the cyclical economy. Costs of business cycles are thus higher than when aggregate fluctuations increase β on average ($\hat{\beta} > \bar{\beta}$), or when β varies with the cycle ($V_\beta > 0$).

9 Appendix IV: A simple SI economy with quadratic preferences

Derivation of equation (3)

The individual Euler equation reads

$$1 - \frac{1}{\theta}c_{it} = \delta R_t \mathbb{E}\left[1 - \frac{1}{\theta}c_{it+1}\right] \quad (22)$$

Market clearing implies that $1 + r_t = \frac{1}{\delta} \frac{\theta - Y_t}{\theta - Y_{t+1}}$. From period $t + 1$, in the absence of aggregate shocks, we therefore have $1 + r = \frac{1}{\delta}$ and consumption is characterized by the standard permanent-income formula for consumption as

$$c_{it+1} = \frac{r}{1+r} [a_{t+1} + (y_{t+1} - 1)\bar{Y}] + \bar{Y} \quad (23)$$

Using this expression, and $\mathbb{E}[(y_{t+1} - 1)] = p(y_{it} - 1)$ in (22) yields

$$\begin{aligned} c_{it} &= \theta(1 - \delta(1 + r_t)) + \delta(1 + r_t) \left[\frac{r}{1+r} (a_{t+1} + p(y_{it} - 1) + \bar{Y}) + \bar{Y} \right] \\ &= \theta(1 - \delta(1 + r_t)) + \delta(1 + r_t) \left\{ \frac{r}{1+r} [(1 + r_t)(y_{it} + a_{it} - c_{it}) + p(y_{it} - 1)] + \bar{Y} \right\} \\ &= \frac{\theta(1 - \delta(1 + r_t)) + \delta(1 + r_t) \left\{ \frac{r}{1+r} [(1 + r_t)(y_{it} + a_{it}) + p(y_{it} - 1)] + \bar{Y} \right\}}{1 + \delta(1 + r_t)^2 \frac{r}{1+r}} \end{aligned} \quad (24)$$

where the second equality follows from the individual budget constraint, and the third from rearranging for c_{it} . Subtracting $C_t = Y_t$ from both sides, yields

$$c_{it} - C_t = r \cdot \frac{\left(\frac{1+r_t}{1+r}\right)^2 \cdot \left[\left(1 + \frac{p}{1+r_t} \cdot \frac{\bar{Y}}{Y_t}\right)(y_{it} - 1)Y_t\right] + a_{it}}{1 + \left(\frac{1+r_t}{1+r}\right)^2 r} \quad (25)$$

Note that the interest rate in period t , paid in period $t + 1$, is

$$1 + r_t = \frac{1}{\delta} \left[1 - \frac{Y_t - \bar{Y}}{\theta - \bar{Y}} \right] \quad (26)$$

$$= \frac{1}{\delta} \left[1 - \frac{\epsilon}{\theta - \bar{Y}} \right] \quad (27)$$

$$= \frac{1}{\delta} [1 - \hat{\epsilon}] \quad (28)$$

Using this in 3 yields

$$c_{it} - C_t = \frac{\left\{ r[\rho_t^2 (1 + \frac{\rho_t p}{\rho_t(1+r)} \frac{\bar{Y}}{Y_t} (y_{it} - 1) Y_t + a_{it}] \right\}}{1 + (1 - \hat{\epsilon})^2 r} \quad (29)$$

10 Appendix V: Sensitivity parameter β in a Krusell and Smith (1998) economy with heterogeneous δ

Table 8: Sensitivity parameters β (in percent) in a Krusell and Smith (1998) economy

	β	β^{YL}	β^{YH}
Stationary	6.390		
Independent	6.320	6.760	6.119
Proc persistence	6.285	6.849	6.026
Counterc risk	6.348	6.640	6.214

For the economy with capital and Cobb-Douglas production and heterogeneous discount factors, such that 90 (10) percent of agents have $\delta = 0.98(0.978)$, the table reports the cross-sectional regression coefficient β , in percent, on average (column 1), as well as in times of lower- (column 2) and higher-than-average output (column 3) without fluctuations (row 1), with idiosyncratic risk independent of aggregate fluctuations (row 2), with procyclical persistence (row 3), and with countercyclical unemployment risk (row 4). To be consistent with CEX data, consumption and income growth are calculated as 4-quarter changes in quarterly consumption and in (backward-looking) annual incomes.

11 Appendix VI: Transition probabilities

Table 9: Independent individual transitions

	y_l, Y^L	y_h, Y^L	y_l, Y^H	y_h, Y^H
y_l, Y^L	0.1876	0.6947	0.0250	0.0926
y_h, Y^L	0.0428	0.8395	0.0057	0.1119
y_l, Y^H	0.0090	0.0335	0.2036	0.7538
y_h, Y^H	0.0021	0.0405	0.0464	0.9110

Table 10: Procyclical persistence

	y_l, Y^L	y_h, Y^L	y_l, Y^H	y_h, Y^H
y_l, Y^L	0.1737	0.7086	0.0232	0.0945
y_h, Y^L	0.0437	0.8387	0.0058	0.1118
y_l, Y^H	0.0097	0.0328	0.2187	0.7388
y_h, Y^H	0.0020	0.0405	0.0455	0.9119

Table 11: Countercyclical risk

	y_l, Y^L	y_h, Y^L	y_l, Y^H	y_h, Y^H
y_l, Y^L	0.1996	0.6828	0.0280	0.0896
y_h, Y^L	0.0450	0.8373	0.0053	0.1124
y_l, Y^H	0.0093	0.0332	0.1980	0.7594
y_h, Y^H	0.0022	0.0403	0.0456	0.9118

12 Appendix VII: An SI economy with capital and production (Krusell and Smith, 1998)

The economic environment is a version of the Krusell and Smith (1998) economy with unemployment benefits, as studied in Den Haan et al. (2010) (and other ar-

ticles in the same issue). Employed individuals earn wage $y_t^h = w_t$, while the unemployed receive unemployment benefits $y_t^l = \mu w_t$. The only asset in the economy is physical capital, whose net return equals $r_t - \delta$, the rental rate net of depreciation, and is equal for all households. Financial markets are thus incomplete, but households can smooth consumption through their choice of capital k_{t+1}^i by saving and dissaving, subject to a no-borrowing limit ($k_t^i > 0$). The household thus solves the problem

$$\max_{\{c_t^i, k_{t+1}^i\}_{t=0}^{\infty}} E_t \sum_{t=0}^{\infty} \beta^t \frac{(c_t^i)^{1-\sigma} - 1}{1-\sigma} \quad (30)$$

$$s.t. \quad c_t^i + k_{t+1}^i = r_t k_t^i + (1 - \tau_t) \bar{l} y_t^h \mathbb{I} + y_t^l (1 - \mathbb{I}) + (1 - \delta) k_t^i \quad (31)$$

$$k_{t+1}^i \geq 0 \quad (32)$$

where \mathbb{I} equals 1 for the employed and 0 for the unemployed, E_t is the mathematical expectation conditional on information in period t , \bar{l} is the time endowment (normalised to have labor supply in bad times equal 1) and τ a proportional tax on labor income.

Competitive firms rent capital and hire labor to produce the single output good in the economy using a standard Cobb-Douglas production technology

$$Y_t = z_t K_t^\alpha (\bar{l} L_t)^{1-\alpha} \quad (33)$$

where z_t is an exogenous process for aggregate productivity that takes on two values ($z_t \in \{Y^L, Y^H\}$).

Markets for labor, capital, and consumption goods are competitive, so factor prices are given by

$$w_t = z_t (1 - \alpha) \left(\frac{K_t}{\bar{l} L_t} \right)^\alpha \quad (34)$$

$$r_t = z_t \alpha \left(\frac{K_t}{\bar{l} L_t} \right)^{\alpha-1} \quad (35)$$

Taxes τ_t are set as to balance the budget of an unemployment insurance scheme, such that

$$\tau_t = \frac{\mu u_t}{\bar{l} L_t} \quad (36)$$

where L_t and $u_t = 1 - L_t$ are, respectively, the employment and unemployment rates in the economy.

We set $\alpha = 0.36$, and $\delta = 0.025$, standard values also used in Den Haan et al. (2010).