Consumption insurance over the business cycle*

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

February 2024

Abstract

In U.S. micro data, consumption smoothing is cyclical: consumption reacts more to idiosyncratic income changes in booms. This matters for average costs of business cycles and aggregate fluctuations in consumption demand. In standard models of self-insurance, where individuals borrow and save to smooth income fluctuations, consumption smoothing is strong when a temporarily low average wage level or low interest rates make current shocks less important for permanent income, and when high savings rates relax future borrowing constraints of low-wealth households. When studying these determinants in the general equilibrium of a standard business-cycle model with incomplete markets and idiosyncratic risk, procyclical wealth accumulation makes consumption smoothing substantially more effective in booms. To solve this "countercyclical consumption smoothing puzzle", we explore alternative market structures, income processes and misperception of idiosyncratic risk. Procyclical bias in the perceived persistence of idiosyncratic shocks can help solve the puzzle.

JEL Classification: E32, G22
Keywords: consumption smoothing, risk sharing, limited enforcement, business cycles, misperceived income risk

*I thank, in particular, Paul Klein. Without his support this paper would not have been written. I also thank Kieran Larkin for help with US micro-date. For comments on previous drafts I thank Árpád Ábrahám, Chris Busch, Dirk Krueger, and Gianluca Violante, as well as Per Krusell, Fabrizio Perri, Morten Ravn, and Kjetil Storesletten. Mattias Hallberg provided outstanding research assistance.

[†]Paris School of Economics, 48 Bd Jourdan, 750014 Paris; Email: `tobias.broer@psemail.eu`.

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 how fluctuations in aggregate economic activity affect this ability of households to insulate consumption from 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.

Motivated by these considerations, our paper makes three contributions: first, we show that according to U.S. micro data from both the Consumer Expenditure Survey (CEX) and the Panel Study of Income Dynamics (PSID), household consumption reacts more strongly to individual income changes in booms. In other words, the sensitivity of individual consumption to income changes is procyclical, making consumption insurance, defined as the degree to which individual consumption is less volatile than individual disposable income, countercyclical on average.

Second, we show how in standard economies with (exogenously) incomplete markets, households can more effectively smooth idiosyncratic shocks in periods of low average wages or interest rates (when the role of current idiosyncratic shocks for permanent income is small); and when previously high incomes have raised average wealth levels through increased savings (which makes borrowing constraints less likely to bind in the future, and thus flattens the consumption policy function). In the general-equilibrium of a standard self-insurance (SI) economy with incomplete markets and idiosyncratic risk, it is precisely this procyclical wealth accumulation that makes consumption smoothing substantially more effective in booms.

Our third contribution is to explore potential solutions to this ”countercyclical consumption insurance puzzle", by relaxing the assumptions of the standard model.

Because our benchmark model abstracts from cyclical fluctuations in earnings risk other than cyclical employment transitions, we first investigate whether countercyclical income risk as in Storesletten et al. (2004) can help explain the puzzle. In particular, they find shocks to the persistent component of individual earnings to be more volatile in recessions. This, however, makes income shocks on average more persistent in recessions, and thus aggravates the puzzle by making consumption smoothing even more procyclical.

This strong role of income persistence for consumption smoothing, and previous evidence that households misperceive their individual income risk, motivates us to study an extension of the SI economy that relaxes rational expectations. Specifically, we build on Balleer et al. (2023), who find that U.S. workers are overoptimistic about their labor-market prospects on average, and that this optimism is significantly larger in times of low unemployment. This motivates us to study a version of the benchmark economy with negative and countercyclical bias in perceived separation probabilities. This suggests that in booms, income shocks are perceived to be more persistent, which increases their effect on consumption. We show that, when strong enough, such misperception of idiosyncratic risk indeed makes consumption smoothing countercyclical, as observed in the data.

Finally, we explore a more fundamental alternative to the SI environment, where restrictions to asset trade are not exogenous, but endogenous, and therefore potentially cyclical. In particular, we consider financial frictions that arise when households have access to a complet set of state-contingent insurance contracts, but cannot commit to honoring contractual payments. We first show a separation result for this "LC" environment 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. Using a new solution method based on the near-analytical solution to the stationary LC economy (Krueger and Perri, 2011, Broer, 2013), we show that with risk aversion greater than 1 and fluctuations in idiosyncratic income risk that capture key features of unemployment risk in U.S. micro-data, consumption smoothing is even more countercyclical than in our benchmark SI economy.

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 characterise the cyclical determinants of the degree of consumption smoothing in standard SI economies. And we extend the stationary LC environment where insurance is limited by the risk of default, studied in Krueger and Uhlig (2006), Krueger and Perri (2011), Krueger and Perri (2006), Broer (2013), or Bold and Broer (2021) to include aggregate stochastic fluctuations and characterize it both quantitatively and analytically. This extension is related to, but substantially pre-dates, Ando et al. (2023), who provide an analytical characterisation of a much more general LC model with capital accumulation under the (more stringent) assumption that one of two idiosyncratic income states equals zero.[1]

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

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), Imrohoroğlu (1989), Aiyagari (1994), Huggett (1997), Krusell and Smith (1998a), and on the relationship between inequality and aggregate economic performance including the effect of policies[2]. The effect of aggregate conditions on the ability of consumers to protect their consumption from income fluctuations has received less attention.[3] Relative to this literature, recently surveyed in Krueger et al. (2016), we make three contributions: first, we study the implications of aggregate shocks and cyclical income processes 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. Second, we explore the role of misperceived income risk in a standard, quantitative SI economy. And third, we 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)). Overall, our study is similar to Krueger and Perri (2005), but focusing on the role of aggregate fluctuations in activity and prices for key moments of risk sharing.

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. A recent literature has documented the cyclicality of income risk including its higher moments, and how it affects consumption smoothing (see Busch and Ludwig (2020) and the references therein). 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. (2014) who document acyclical variance (but procylical skewness) of income

[2] See 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.

[3] An 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.

changes in US administrative income data. While we do consider the role of cyclical income risk for the cyclical nature of consumption insurance, we mainly view our analysis as complementary to this literature, as we highlight the cyclical nature of consumption insurance for any given income process.

Section 2 shows that cyclical consumption smoothing changes average measures of inequality, the welfare costs of business cycles, and the dynamics of aggregate demand. presents our empirical findings based on U.S. CEX and PSID data. Section 3 discusses the effect of aggregate fluctuations on consumption smoothing in a cyclical version of the standard income fluctuations problem. Section 4 studies a quantitative general equilibrium model with both aggregate and idiosyncratic risk and exogenously incomplete markets. Section 6 presents extensions to cyclical and misperceived idiosyncratic risk. Section 6 discusses the LC economy with endogenous financial frictions. Section 7 concludes.

2 Cyclical insurance: Why should we care?

Relative to a benchmark of complete insurance, dispersion and volatility in the consumption of individual households reduces average welfare. The desirability of public insurance and redistributional policies crucially depends on a quantification of this welfare cost, and on the frictions that cause limited insurance and determine the effectiveness of policies in reducing consumption heterogeneity. In addition, consumption heterogeneity may affect the level and dynamics of asset prices and aggregate demand through, potentially time-varying, precautionary savings and financial frictions. The rest of this section briefly illustrates this in a simple, illustrative setting.

Consider an economy in discrete time $t = 1, 2, ...$ with a continuum of households and aggregate output of a single consumption good Y_t, where a reduced-form partial-insurance function links individual consumption and income shares: $\tilde{c} = \theta_t \tilde{y}_t^{\beta_t}$, for $\tilde{x} = \frac{x}{Y_t}$. So only a time-varying fraction β_t of idiosyncratic shocks to log-income shares are passed through to log-consumption. Assume that $log(\tilde{y}_t)$ and $log(Y_t)$ are i.i.d. normally distributed as $N(0, V_y)$ and $N(ln(\bar{Y}), V_Y)$, respectively, and denote the (for simplicity constant) mean and variance of β_t as, respectively $\overline{\beta}$ and Var_β.

Cyclical insurance raises consumption inequality and reduces welfare

Because most inequality measures and typical utility functions are concave, time-fluctuations in consumption caused by cyclical insurance raise inequality and lower welfare on average. For example, the difference in the variance of log consumption between this simple setting and an alternative where $\beta_t = \overline{\beta}$, $\forall t$, equals

$$\Delta Var(log(c)) = E_t[\beta_t^2 Var_y] - \overline{\beta}^2 Var_y = Var_\beta Var_y > 0 \tag{1}$$

Following Lucas (1987), with constant risk aversion equal to σ, the cost of cyclical fluctuations, including in consumption insurance, expressed as a difference in permanent consumption equals

$$\Delta ln\bar{c} \approx \frac{1}{2}\sigma V_Y \left[1 + \Phi_{cY}\left(V_\beta - (\sigma - 1)Cov(\beta, Y)\right)\right] \tag{2}$$

where $\Phi_{cY} = \frac{V_y \bar{\beta}^2}{V_Y}$ is the variance of (idiosyncratic) consumption shares relative to that of aggregate consumption / income.

The utility cost of cyclical fluctuations in aggregate income is thus approximately equal to σV_Y, as in Lucas (1987). With cyclical consumption smoothing there are additional welfare losses from cyclical fluctuations in β_t proportional to the relative variance $\frac{V_y}{V_Y}$. Because the variance of individual incomes V_y is one to two orders of magnitude larger than that of aggregate income V_Y, the loss from variations in consumption insurance may easily dominate that of aggregate income movements.[4]

[4]The 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 percent 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.

An additional welfare loss arises whenever consumption insurance is low in booms ($Cov(\beta, Y) < 0$): because of declining marginal utility, the cost of any consumption fluctuations from limited insurance increases during times of low average output and consumption.

Note that the welfare loss (2) abstracts from a potentially important additional source of welfare costs of business cycles that arises whenever cycles reduce the *average* degree of consumption insurance, which is constant and equal to $\overline{\beta}$ in our illustrative environment.

Cyclical insurance changes the path of aggregate demand

Within the illustrative environment of this section, consider a simple equilibrium where, in addition to the reduced-form insurance, households can smooth consumption by trading a bond in zero net supply and subject to a zero-borrowing limit, as in Werning (2015). The Euler equation for individual bond holdings then defines the interest rate R_t as

$$1 = \delta R_t \max_i \{E_t[\left(\frac{\theta_{t+1}\tilde{y}_{it+1}^{\beta_{t+1}}}{\theta_t \tilde{y}_{it}^{\beta_t}}\right)^{-\sigma} \left(\frac{Y_{t+1}}{Y_t}\right)^{-\sigma}]\} \tag{3}$$

where the maximum is taken across all individuals i and we follow Werning (2015) by imposing the equilibrium condition $C_t = Y_t$.

When the process for income shares $\tilde{y}_t$ is independent of Y_t, Werning (2015)'s "as if" result holds: the equilibrium is the same as that with a representative agent whose discount factor equals $E[\hat{\delta}_t]$, where $\hat{\delta}_{it} = \delta \left(\frac{(\theta_{t+1}\tilde{y}_{it+1}^{\beta_{t+1}})}{(\theta_t \tilde{y}_{it}^{\beta_t})^{-\sigma}}\right)^{-\sigma}$, and $E_t[\hat{\delta}_{it}]$ is increasing in β_{t+1}[5]. In particular, the elasticity of current consumption to (policy-induced) changes in future consumption or the interest rate is unchanged by limited risk sharing. Correlation between β_t and aggregate output Y_t breaks this as if result in the same way as correlation between income risk and aggregate output in Werning (2015). In particular, whenever β_t is a decreasing function of Y_t at given interest

[5]To see this, note that $E_t[\hat{\delta}_{it}] = e^{\sigma(\beta_{t+1}^2(1+\sigma)-\beta_t^2)V_y+\sigma\beta_t ln(\tilde{y}_{it})}$

rate R_t, such that consumption insurance is procyclical, time-varying fluctuations in consumption insurance amplify the response of current aggregate consumption to changes in future consumption. In other words, procyclical consumption insurance amplifies aggregate demand fluctuations as rising future Y_t reduces current precautionary savings. Mutatis mutandis, counteryclical insurance dampens aggregate fluctuations.

3 Evidence from US Micro Data

This section presents evidence that consumption insurance, as measured by the comovement of individual consumption and income, is weaker during times when aggregate income is above its trend level. The ideal data for studying consumption insurance over the cycle would have high-frequency information on consumption expenditure and its determinants such as asset portfolios and returns, earnings and income shocks, as well as non-financial transfers, of a sample of households through time. Existing data sets, in contrast, contain information at medium frequency (lower than the frequency of typical income shocks), on subsets of household consumption, have little or no information on exogenous shocks, no or irregular information on wealth, and follow households over short periods. We therefore focus mostly on a particularly simple reduced-form measure of the sensitivity of individual household consumption to individual income changes, namely the slope of the conditional mean of consumption growth as a function of income growth, equal to the coefficient β in the following regression

$$\Delta c_t \quad = \quad \alpha + \beta \Delta y_t + \varepsilon_t \tag{4}$$

where Δc_t and Δy_t denote the 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 the coefficient β that characterises the pass-through from income to consumption changes as a simple, and classical, measure of consumption smoothing (see Gervais and Klein (2010)'s discussion of the literature) that we compare to the same moments in models of consumption insurance, or to indirectly infer model parameters. Importantly, any measurement error in consumption leaves the regression coefficient unaffected, while error in measured incomes 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 $\mathsf{STD}(\Delta c)$. As we argued in the introduction, it is the cyclicality of $\mathsf{STD}(\Delta c)$ that may strongly affect the cost of business cycles. Although measurement error in consumption strongly affects the level of $\mathsf{STD}(\Delta c)$, it does not affect its cyclical behavior as long as the error is not in itself cyclical.

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 the income measures in the surveys we study. We define as booms, or good times, those of above-trend aggregate activity, and as busts, or bad times, those with activity below trend.[6]

3.1 Evidence from CEX data

We first consider evidence from the US Consumer Expenditure Survey (CEX), the standard source of consumption information in the US. The CEX is a four-quarter rotating panel with detailed information about quarterly household consumption expenditures. The survey 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.[7] Together, these facts point to substantial noise in CEX data on the joint income-consumption distribution. We never-

[6]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 CEX 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.

[7]For 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.

theless choose the CEX as our prefered source of information because of its high frequency, and its broad coverage of consumption items.[8] We interpret our results below bearing in mind the issue of measurement error and the somewhat odd timing of measurements and income in the CEX that may further attenuate measured consumption-income comovement, and discuss the extent to which it may influence our measurement of the comovement between consumption smoothing and the business cycle.

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 alternative measures for income (family earnings) and consumption (nondurable consumption excluding rental payments, *ND*). 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.[9] Appendix I contains details about data construction.

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 quarterly coefficient estimates β_t, 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

[8]See 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.

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

Figure 1: Consumption Insurance in CEX data

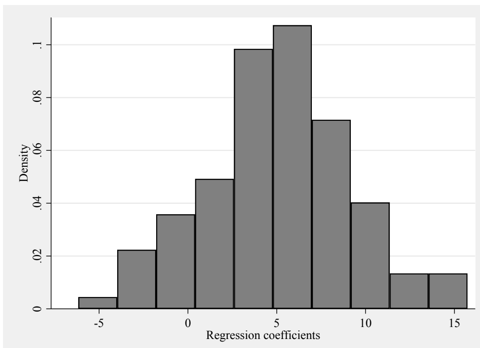

(a) Histogram of β_t in (4) estimated quarter-by-quarter (in percent) for consumption measure ND+ and income measure disposable income.

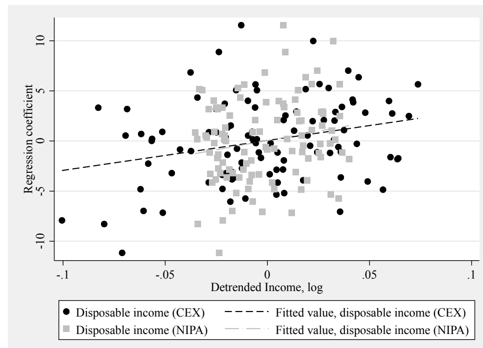

(b) Scatter plot of β_t (in percent, vertical axis) against detrended disposable income (NIPA) / GDP (log, horizontal axis), quarterly frequency. Both variables are calculated as deviations from their mean.

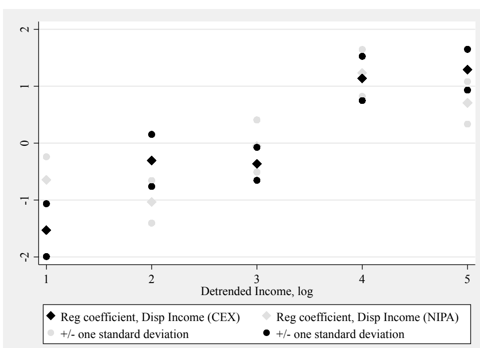

(c) β_t (in percent, vertical axis) estimated quarter-by-quarter and averaged within quintiles of detrended disposable income (CEX and NIPA) (log, horizontal axis). Both variables are calculated as deviations from their mean.

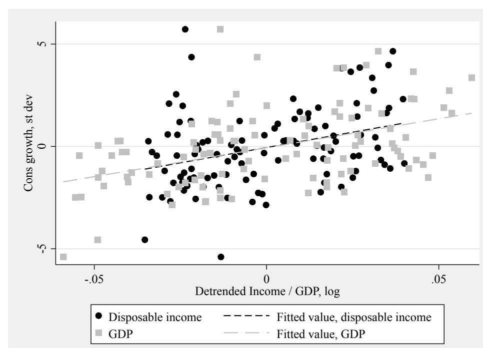

(d) Scatter plot of the standard deviation of ND+ consumption growth (in percent, vertical axis) against detrended disposable income (NIPA) / GDP (log, horizontal axis).

Table 1: Regression coefficient β, in percent (CEX)

	(1)	(2)	(3)	(4)
	ND+	ND	ND+	ND
Disp income growth	5.043***	6.213***		
	(15.17)	(14.53)		
Earnings growth			2.802***	3.033***
			(11.04)	(9.28)
Constant	0.261	0.746***	0.322**	0.830***
	(1.60)	(3.55)	(1.97)	(3.94)
r2	0.00664	0.00609	0.00353	0.00250
N	34444	34444	34443	34443

The table reports the regression coefficient β, in percent, in different versions of (4) when c_t is nondurable consumption plus (imputed) rental services (ND+, columns 1 and 3), nondurable consumption (ND, columns 2 and 4), 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).

with higher-than-average sensitivity of individual consumption to income changes (higher values of β, along the vertical axis). Moreover, this cyclicality is substantial: according to panel $c)$ the regression coefficient β_t is on average three quarters higher when the cyclical component of income is in the top quintile, compared to the bottom quintile. Table 2 depicts the corresponding coefficients γ_0 and γ in a regression of β_t, the coefficients from a quarter-by-quarter estimation of (4), on deviations of aggregate income measures from their log-linear trend $\hat{Y}_t$

$$\beta_t \quad = \quad \gamma_0 + \gamma \hat{Y}_t + u_t. \tag{5}$$

The table 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).

The procyclical sensitivity of consumption to income changes documented in Panel

Table 2: Regression coefficients γ (CEX)

	(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 estimates of γ in (5) using quarterly time series of coefficients β_t (using ND+ consumption growth in column 1 and 2, ND consumption growth in column 3, and Food consumption growth in column 4), in percent, and 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).

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).

3.2 Evidence from PSID data

The PSID is a longitudinal survey of a sample of US households originally designed for studying the dynamic evolution of poverty and income. The original 1968 sample comprised a representative sample of 3000 households, plus 2000 poor families. Since then, the PSID has followed the original families and the families of their offspring. The survey changed to a computer-based interview in 1993, and from annual to biannual frequency in 1997.

Relative to the CEX, the PSID has two disadvantages for studying the dynamics of consumption insurance over the business cycle, as it only contains information about food consumption (at home, and away from home), and at a lower frequency (annual or biannual). The information about income is, however, considered as more accurate, and pertains to exactly the same time period as that on consumption, thus, presumably, reducing income measurement error. We use two measures of household income. First, the PSID aggregate of all household income components, labeled "family income". And second, the sum of labor earnings and transfers of all household members, labeled "household income".

The survey asks for a wide variety of information about the household and its members, in particular its "head" (typically the husband in a married couple) and the spouse, if present. We use sample information from the 1980 to 2017 waves of the survey.[10] We use a sample definition that follows Heathcote et al. (2010). In particular, we focus on households whose head is male, and between 30 and 60 years of age. We also drop single households and eliminate outliers.[11] As for the CEX, we base our analysis on the residuals of consumption and income from a regression on household observables, including dummies for education, race, and a polynomial in the age of the household head.

Table 4 presents the estimated slope coefficients β in (4) when Δc_t and Δy_t correspond to the two-year difference in, respectively the logarithm of PSID food consumption and that of family and household income. The point estimates are about twice as large as for the CEX. Figure 2 shows that the business-cycle pattern of insurance, as summarised by comovement of the time varying coefficients β_t and deviations of aggregate income measures from their log-linear trend, is very similar to that observed in CEX data. Specifically, the regression coefficient β_t is higher in periods when aggregate activity is above trend. Table 4 shows that the magnitude of this correlation, as summarised by the coefficient γ in an estimation of (5) using PSID data, is also similar to that observed in CEX data. The point estimates are, however, less statistically significant. We attribute this to the small number of

[10]We choose the starting year to coincide with that of our CEX sample.

[11]Specifically, we winsorize observations whose level or growth rates of income or consumption lie in the two most extreme percentiles of the distribution in a given year.

Table 3: Regression coefficient β, in percent (PSID)

	(1)	(2)
	1	2
Family income growth	8.772***	
	(10.75)	
Household income growth		8.922***
		(10.82)
Constant	0.629***	0.639***
	(3.50)	(3.56)
r2	0.00512	0.00524
N	23675	23675

The table reports the regression coefficient β, in percent, in different versions of (4) using data from the PSID. c_t denotes food consumption, and y_t is total disposable family income after taxes and transfers (row 1), or household income (row 2). Robust standard errors are used; stars denote conventional significance levels: * (p<.1), ** (p<.05), *** (p<.01).

observations.

4 The cyclical income fluctuations problem

Motivated by the procyclical consumption-income comovement observed in US data, this section studies how temporary fluctuations in aggregate output, the interest rate and the persistence of individual income movements change the consumption response to idiosyncratic income shocks according to standard consumption theory.

For this, we consider a "cyclical" income fluctuations problem, where consumers "self-insure" income shocks by buying or selling one-period bonds at a given interest rate, but where we distinguish between aggregate and idiosyncratic income movements, and let the interest rate, as well as the persistence of idiosyncratic shocks, vary over time. Specifically, we derive the period-0 response of consumption to an additive, one-time idiosyncratic shock ϵ_0 to income share $y_{i0} = 1$ for different exogenous sequences of aggregate income and interest rates $\{Y_t, R_t\}_0^\infty$. ϵ_t converges back to its mean 0 according to $\epsilon_t = \rho_t \epsilon_{t-1}$ where the persistence $\rho_t < 1$ is potentially

Figure 2: β_t and aggregate income (PSID)

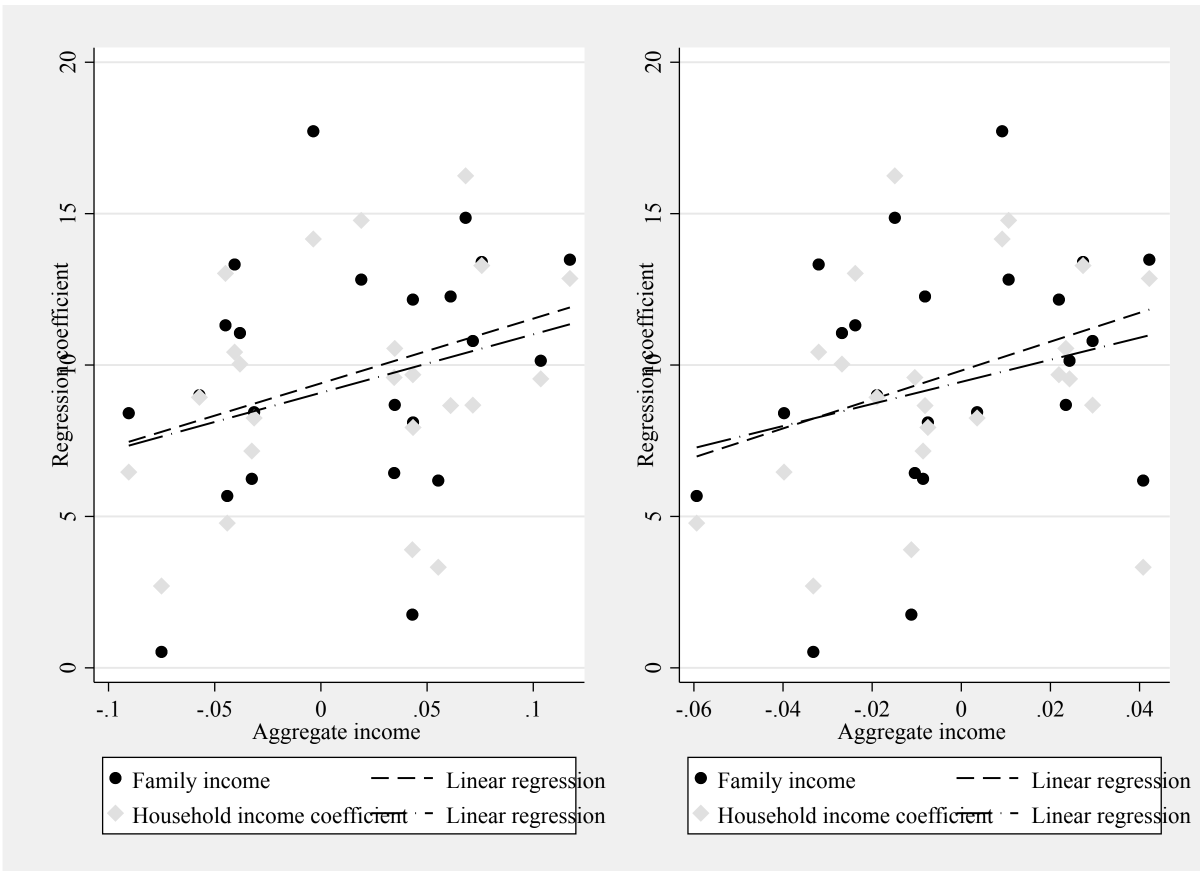

Scatter plot of β_t in (4) estimated period-by-period (in percent, vertical axis) against detrended disposable income (PSID / NIPA, horizontal axis), biannual frequency.

Table 4: Regression coefficient γ (PSID)

	(1)	(2)	(3)	(4)
	1	2	3	4
Disp Y (PSID)	21.36*	19.30*		
	(1.73)	(1.75)		
Disp Y (NIPA)			47.75*	36.46
			(1.73)	(1.23)
Constant	9.397***	9.084***	9.816***	9.445***
	(10.71)	(12.23)	(12.05)	(12.40)
r2_o				
N	23	23	23	23

t statistics in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

The table reports estimates of γ in (5) using time series of coefficients β_t (using PSID Food consumption growth and respectively, family income growth (column 1 and 3) and household income growth (column 2 and 4), in percent, and log-trend deviations of aggregate disposable income (from the PSID, row 1, and from NIPA, row 2). Robust standard errors are used; stars denote conventional significance levels: * (p<.1), ** (p<.05), *** (p<.01).

time-varying.

4.1 Permanent-income consumption

When consumers have log-preferences, discount the future at rate δ, and can save and borrow at a given interest rate without borrowing constraints, standard derivations (see Appendix 12 for details) yield period-0 consumption as

$$C_0 = (1-\delta)\left(\sum_{t=1}^{\infty} Y_t Q_t y_{it} + Y_0 y_{i0} + a_{i0}\right) \tag{6}$$

Here $Q_t = \frac{1}{\Pi_{s=0}^{t} R_s}$ denotes the period-0 price of period-t consumption and a_{i0} consumer i's asset holdings in period 0. Period-0 consumption thus simply equals $(1-\delta)$ times the present discounted value of life-time income plus current wealth. To consider cyclical fluctuations, denote the 'steady-state' values of output, interest rates, and income persistence as, respectively $\bar{Y} = 1, \bar{R} = \frac{1}{\delta}, \bar{\rho}$ and assume that $\{Y_t, R_t\}_0^{\infty}$ converge monotonically from Y_0, R_0 to their steady-state values. This allows to calculate β_0, the response of consumption to an unanticipated income change ϵ_{i0} as

$$\beta_0 = \frac{d\frac{C_0 - C_0^{y=1}}{C_0^{y=1}}}{dy_{i0}} = \frac{1}{dy_{i0}} \frac{\sum_{t=0}^{\infty} \delta^t \omega_t q_t y_{it} + \alpha_{i0} - \sum_{t=0}^{\infty} \delta^t \omega_t q_t - \alpha_{i0}}{\sum_{t=0}^{\infty} \delta^t \omega_t q_t + \alpha_{i0}} = \frac{\sum_{t=0}^{\infty} \delta^t \omega_t q_t \rho_0^t}{\sum_{t=0}^{\infty} \delta^t \omega_t q_t + \alpha_{i0}} \tag{7}$$

where $C_0^{y=1}$ denotes consumption in the absence of a shock, $\omega_t = \frac{Y_t}{\bar{Y}}$ and $q_t = \frac{Q_t}{\bar{Q}_t} = \frac{Q_t}{\delta^t}$ are the ratios of, respectively, total income and period-t consumption prices to their steady-state values, and $\alpha_{i0} = \frac{a_{i0}}{\bar{Y}}$.

According to (7), the percentage change in consumption from an idiosyncratic income shock $dy_{i0} = \epsilon_{i0}$ simply equals the relative change in present discounted lifetime income. Proposition 1 summarises how cyclical factors affect the coefficient β_0.

Proposition 1 *Countercyclical self-insurance of permanent-income consumers*

Consider $\Delta = \beta_0 - \beta$, the deviation of the coefficient β_0 from its steady-state value $\beta = \frac{1-\delta}{1-\delta\rho_0}\frac{1}{1+(1-\delta)a_0}$.

Δ is positive, and the impact of current income shocks on consumption thus larger than in steady state, if any of the following conditions holds:

1. *Current aggregate income is above steady state $Y_0 > \bar{Y}$ and $R_t = \frac{1}{\delta}\ \forall t$.*
2. *The current interest rate is above steady state $R_0 > \bar{R} = \frac{1}{\delta}$, $Y_t = \bar{Y}\ \forall t$, and financial assets are below a strictly positive threshold $\bar{\alpha}_0$.*
3. *The current interest rate is below steady state $R_0 < \bar{R} = \frac{1}{\delta}$, $Y_t = \bar{Y}\ \forall t$, and financial assets are above a strictly positive threshold $\bar{\alpha}'_0$.*
4. *The current persistence ρ_0 is above its steady state $\rho_0 > \bar{\rho}$, $Y_t = \bar{Y} \forall t$, and $R_t = \frac{1}{\delta} \forall t$.*

The converse statements hold mutatis mutandis.

Proof. See Appendix 12. ■

According to proposition 1, idiosyncratic shocks affect consumption of permanent-income consumers more when their effect on total life-time wealth is strong. To understand when this is the case, note that life-time human wealth is more affected by an idiosyncratic shock to income share y_{i0} today when current aggregate income is above average, which front-loads the effect of current incomes on lifetime human wealth. In addition, higher-than-average aggregate incomes increase the share of human wealth in total wealth. Together, this strengthens the impact of current idiosyncratic shocks on consumption when current aggregate income is high.

When current interest rates are high, future incomes are discounted more heavily. This increases the relative effect of current incomes on life-time human wealth. High interest rates, however, decrease the share of human wealth in total wealth. The former effect is more important for individuals whose life-time wealth is dominated by human wealth. For individuals with little financial wealth, idiosyncratic shocks thus affect consumption more when current interest rates are above average, but the opposite is true for the financially wealthy.

Finally, when persistence is above average, any current shock affects life-time human wealth more strongly.

4.2 Introducing borrowing limits

How do occasionally binding borrowing limits affect the results in Proposition 1? A binding borrowing limit today breaks the link between current and future consumption. The possibility of binding constraints in the future thus limits the set of states or periods over which current consumption is smoothed. This increases the effect of current income shocks on current consumption. In the absence of uncertainty, borrowing constraints are more likely to bind the more upward-sloping is the income path relative to the optimal consumption path (implying temporary borrowing to finance front-loaded consumption), and the lower are current assets a_{i0}. With borrowing constraints, there is thus an offsetting force to the procyclical consumption-income comovement implied by permanent-income theory. In particular, a slump in aggregate output, or a fall in the interest rate, raise current consumption above current income, and thus make borrowing limits more likely to bind in the future, raising the effect of idiosyncratic income shocks on individual consumption. This effect is larger at lower asset levels, where borrowing limits are more likely to bind.

Instead of identifying (restrictive) conditions under which this change in the likelihood of binding borrowing constraints dominates the permanent-income effect, we illustrate the overall effect of cyclical fluctuations on self-insurance along the wealth distribution numerically in Figure 3, for a somewhat more general setting with persistent and transitory income shocks with unit mean, relative risk aversion equal to 2 and a borrowing limit equal to per-period-income.[12] We study the responses of individual consumption to, respectively, a persistent income shock (red lines in Figure 3) and a purely transitory income shock (blue line), both of approximately

[12]Specifically, we set $\beta = 0.96$, and use an income process that is the sum of a purely transitory shock ζ_t and an AR(1) process ϵ_t with persistence 0.952, whose innovations have standard deviations of respectively, 0.255 and 0.168 of steady-state income, see Storesletten et al. (2004). The persistence of aggregate shocks to output and interest rates is 0.95.

one standard deviation.

The left panel of Figure 3 plots the difference $\Delta = \beta_0 - \beta$, where β_0 denotes the response of consumption to an idiosyncratic shock ϵ_0 that coincides with a persistent positive shock to aggregate output Y_0 of 2.5 percent and β the steady state response, as a percentage of the latter. In line with Proposition 1, the difference is positive at high asset levels: when borrowing limits are unlikely to bind, consumption equals permanent income, which responds more strongly to idiosyncratic shocks when current aggregate output is high. At low asset levels, in contrast, the difference turns negative: saving during output booms makes borrowing limits less likely to bind in the future, and thus dampens the response of consumption to idiosyncratic shocks. Because the difference between current and permanent income implied by persistent aggregate output booms is smaller, the difference in consumption responses to persistent shocks (the blue line) is dampenend relative to the case of transitory shocks (and would trivially be zero for permanent output increases).

The central panel considers the percentage difference Δ in the consumption response to an idiosyncratic shock ϵ_0 that coincides with a persistent positive shock to the interest rate R_0 of 2.5 percentage points and the response in steady state, again as a percentage of the latter. Again in line with Proposition 1, the difference is positive at medium-to-high asset levels as higher interet rates raise the effect of current income on permanent incomes, an effect that is again smaller for transitory increases in aggregate output. The difference declines above a threshold value of assets, as high interest rates reduce the role of human in total wealth. This effect is stronger for persistent increases in the interest rate, and reduces the consumption response to shocks at high asset values below that in steady state, reversing the sign of Δ in that case. As in the case of a positive aggregate income shock in the left panel, at low values of assets the increase in optimal savings with a higher interest rate reduces the likelihood of binding borrowing limits and thus makes consumption respond less to idiosyncratic shocks, implying a negative Δ.

Finally, the right panel depicts the difference in the individual consumption response with (for simplicity permanently) higher persistence of idiosyncratic shocks relative to its steady state value. Higher persistence increases the effect of persistent idiosyncratic shocks on permanent income, and thus their effect on consumption. But it also makes borrowing limits less likely to bind (by reducing the likelihood of low income states in the future) and thus reduces the effect of transitory shocks on consumption at low asset values.

Figure 3: Individual consumption responses compared to steady state: numerical illustration

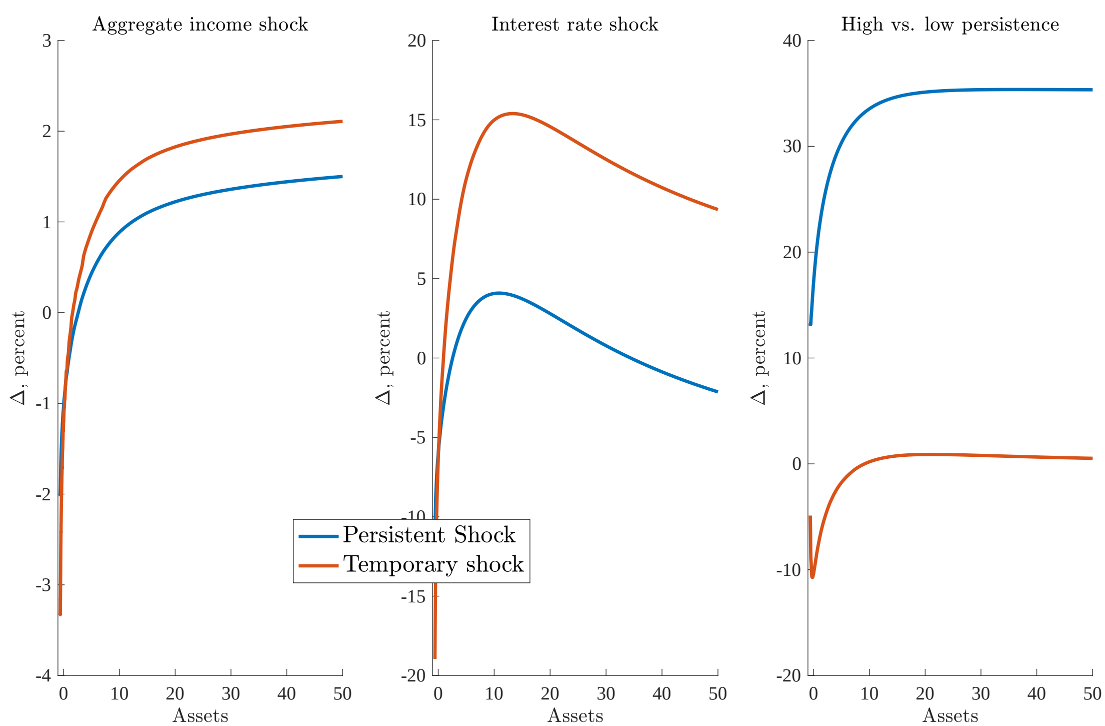

The figure plots the difference $\Delta = \beta_0 - \beta$, as a percentage of β. β_0 denotes the response of consumption to an idiosyncratic shock ϵ_0 that coincides with a persistent positive shock to aggregate output Y_0 (left panel) or the interest rate R_0 (central panel), or a permanent change in the persistence of idiosyncratic shocks (right panel), and β denotes the steady state response. The blue line pertains to a purely temporary idiosyncratic shock, the red line to a persistent shock.

In sum, this analysis of the "cyclical" income fluctuation problem has identified movements in aggregate output, interest rates, and the persistence of income shocks, as important determinants of cyclical fluctuations in the degree of self-insurance. Moreover, it highlighted substantial heterogeneity across the wealth distribution. In particular, at high wealth a temporary increase in aggregate income or interest rates raises consumption-income comovement as in Proposition 1 as borrowing limits can be ignored. At low levels of wealth, in contrast, procyclical savings lengthen the horizon of consumption smoothing by reducing the likelihood of binding borrowing constraints in the future, which dampens consumption-income comovement in booms.

The overall cyclicality of consumption insurance therefore depends on the comovement of interest rates and aggregate output over time, and on the equilibrium wealth distribution. In fact, Werning (2015) provides an example where the general-equilibrium comovement of output and interest rates completely undoes the cyclicality of insurance that results from isolated movements in either of them in partial equilibrium. In other words, to quantify the average cyclicality of consumption insurance requires a general-equilibrium analysis of a quantitative model that is able to capture the wealth distribution observed in the data. We turn to such an analysis next.

5 Quantitative analysis: the countercyclical-self-insurance puzzle

This section quantifies cyclical fluctuations in consumption insurance in a state-of-the-art quantitative general-equilibrium environment with idiosyncratic income risk and incomplete markets (which we call the "SI", or "self-insurance" model). In particular, we consider a version of the model economy in Krueger et al. (2016) that extends the classical environment of Krusell and Smith (1998b) to include a stylised life-cycle with social security transfers for retirees, heterogeneity in the degree of patience across households, and persistent shocks to earnings of employed households.

5.1 The model environment[13]

Time is discrete $t = 0, 1, \ldots.$ The unique consumption good is produced using a standard constant-return-to-scale technology $Y = ZF(K, N)$ where $Z \in \{Z_l, Z_h\}$ is an aggregate productivity shock that follows a Markov process $\pi(Z'|Z)$ and K and N are capital and labor inputs, respectively. Capital depreciates at rate $\tilde{\delta} \in [0, 1]$.

There is a constant measure 1 of potentially infinitely-lived households. Households are born young and turn old (retire) with a constant probability $1-\theta \in [0, 1]$ every period. Old households die with constant $1 - \nu \in [0, 1]$. Households have CRRA preferences with heterogeneous discount factor δ_i. Young households are born without assets, and provide one unit of labor. Income risk comes from two sources. First, there is risk to change unemployment status $s \in \{e, u\}$, where e denotes employment and u unemployment. The conditional distribution of s' $\pi(s'|s, Z', Z)$, that describes unemployment risk and whose typical element we denote as $\pi_{ij}^{kl}, j, i \in \{e, u\}; k, l \in \{h, l\}$, depends on the current employment status s as well as the current and future level of aggregate productivity. So unemployment risk is cyclical, with higher job-finding rates and lower separation rates in good aggregate times ($Z = Z_h$). We assume a law of large numbers, and specify π such that the number of households in each employment state only depends on current (but not past) aggregate productivity. In addition to unemployment risk, there are fluctuations in idiosyncratic labor endowments $e \in \{e_1, \ldots, e_7\}$ that follow a Markov process $\pi(e'|e)$ that is independent of aggregate productivity. To self-insure against income fluctuations, in additional to a perfect annuity market, households have access to the market for physical capital of which they hold $k \geq 0$. We denote the joint cross-sectional distribution of employment status s, individual productivity y and asset holdings k as Φ.

The government runs balanced-budget unemployment insurance and social security systems. Households receive unemployment benefits that replace a constant fraction $\tilde{\rho}$ of their most recent labor income, financed by a proportional labor-income tax ξ. The government charges a constant payroll tax ξ_{SS} that finances social security

[13] This section may be skipped by readers familiar with the Krueger et al. (2016)-environment.

benefits that are the same for all retirees.

Appendix 13 describes the problem of a working household and defines the recursive competitive equilibrium, both of which are standard.

5.2 Calibration

We closely follow the calibration to quarterly data in Krueger et al. (2016) with the exception of our specification of aggregate and unemployment risk. While Krueger et al. (2016) are interested in the interaction between household heterogeneity and output fluctuations in rare, deep recessions (such as the Great Recession of the early 2000s), our focus is on more standard cyclical fluctuations. To transparently capture key features of cyclical fluctations in the post-II war U.S. economy, and in line with our empirical analysis in Section 3, we define “good" times, and “bad" times in the data as those when the U.S. unemployment rate is, respectively, above and below its trend level.[14] We choose transition probabilities $\pi(Z'|Z)$ such that the Markov process matches the average length of bad times, and the average time the economy spends there, in the data. We then choose Z_h and Z_l equal to the average levels of total factor productivity in good and bad times, respectively. We choose a Markov process for s such that the job finding probability in the model equals that in U.S. data (between 1948 and 2012) in good and bad times, respectively, and set the separation rates such that the unemployment rate in the ergodic distribution equals 4 (6) percent during good (bad) times.[15] Good times are more persistent than bad ones, such that the economy spends a little more than 55 percent of periods in good times, where aggregate productivity is 2.4 percent higher than its mean (normalized to 1). This yields job-finding rates close to 80 percent per quarter on average, and 10 percentage points higher in booms, so strongly procyclical. Separation rates average 4 percent, but are 1.2 percentage points higher in bad times, so counter-

[14]To allow for slow-moving fluctuations in the natural rate of unemployment we include a quadratic term in the trend.

[15]The remaining ”off-diagonal" parameters are identified by the requirement that the unemployment rate only be a function of current aggregate productivity Z.

cyclical. For comparison, we also study a version of the model where employment risk is independent of the aggregate economy (and both separation and job-finding rates equal their weighted average values), which we label "acyclical unemployment risk". Appendix 13 describes the calibration of the remaining parameters that follows Krueger et al. (2016) .

5.3 Consumption insurance in the quantitative model

In line with the results in Krueger et al. (2016), the model captures key moments of the joint distribution of consumption, income and wealth, in particular a realistic ratio of capital to annual output of 2.8, and substantial wealth dispersion with a Gini coefficient of 0.77.[16] Here we concentrate on the joint distribution of consumption and income growth, and its comovement with the business cycle.

Table 5 shows the regression coefficient β in (4), in percent, estimated on a long simulated panel of households from the model's dynamic equilibrium. Columns two and three show, respectively, the regression coefficients estimated on subsamples of periods when the economy is in good (β_h) or bad (β_l) times of, respectively, high and low levels of productivity and low / high levels of unemployment.

Row 1 shows the regression coefficient in a version of the model without any aggregate fluctuations, and with acyclical unemployment risk (essentially a version of the Aiyagari (1994) economy). The regression coefficient is substantially higher than that estimated on U.S. microdata. We take this as evidence of measurement error in reported income growth rates.[17]

[16]The features that allow a Gini coefficient for wealth about twice that in Krusell and Smith (1998b) are a higher replacement ratio of unemployment benefits and the life-cycle with social security transfers for retirees (that reduce the incentives to save during or in anticipation of, respectively, unemployment or retirement, widening the left tail of the distribution) and persistent earnings risk of the employed (that increases incentives for precautionary savings in particular for households with high labor income).

[17]Classical measurement error in the log-*level* of incomes whose standard deviation equals 20 percent that of actual income *growth* (or log-differences) would bring the coefficients estimated on model data close to those in PSID data, as it raises the variance of measured income growth to $Var(dlogy) \times (1 + 2 * 0.2)^2 \approx 2Var(dlogy)$. An alternative calibration would target the observed degree of consumption insurance in the data. But since we strongly suspect measurement error

The cyclical fluctuations in output, interest rates, and idiosyncratic risk in the benchmark model (row 2 of Table 5) reduce the average regression coefficient only slightly, by less than 2 percent. Importantly, however, the degree of consumption-income comovement is *countercyclical* in the model: the regression coefficient in bad times β_l exceeds that in good times β_l by 13 percent, implying a γ coefficient that is negative, equal to about minus 11 percent. So the model predicts *procyclical* consumption insurance, in contrast to the data.

Row 3 of Table 5 shows the regression coefficients in a version of the model that abstracts from movements in aggregate productivity, such that pre-tax wages are constant (implying that the coefficient γ is not defined as aggregate income is constant). Even in this case β_l exceeds β_h, by about half the baseline amount. So even in the absence of productivity fluctuations, cyclical unemployment risk in line with U.S. evidence alone implies procyclical consumption insurance, again in contrast to the data. The reason for this is that unemployment shocks are more persistent in busts, when job-finding rates are low, and thus cause larger consumption falls.

Row 4 depicts the regression coefficients in a version of the model with acyclical unemployment risk (where the unemployment rate equals the average of good and bad times in all periods) but fluctuations in aggregate productivity equal to those in the baseline model. The difference between β_h and β_l is reduced by a third (as more persistent unemployment shocks during booms increase consumption-income comovement as measured by β_h), but remains negative. And because aggregate output fluctuates less with constant employment rates, the regression coefficient γ is only slightly lower than in row 2.

5.4 Consumption insurance along the equilibrium wealth distribution

What lies behind the procyclical nature of consumption insurance in the model even without cyclical unemployment risk (row 4 of Table 5)? Figure 4 depicts the differ-

in incomes, but ignore its magnitude, we follow the calibration in Krueger et al. (2016) that yields realistic levels of average wealth (or the capital-output ratio) and its dispersion.

Table 5: Regression coefficient β, in percent, in the quantitative model

	β	β_l	β_h	γ
Stationary	17.8	17.8	17.8	NaN
Benchmark	17.6	18.8	16.7	-11.3
$PE: w, r, u \; at \; SS$	17.7	18.2	17.2	NaN
iid unemp	18.0	18.7	17.3	-9.3

The table reports the regression coefficient β in (4) in a simulation of the quantitative model environment on average (column 1), as well as in times of low and high aggregate productivity (β_{low} and β_{high}, columns 2 and 3, respectively). Column 4 reports the regression coefficient γ in Equation (5), a measure of how procyclical individual consumption-income comovement is.

ence $\Delta = \beta_h - \beta_l$, when both coefficients are estimated separately for households within (average) decile ranges of the equilibrium wealth distribution of the model (depicted along the bottom axes). The upward-sloping concave pattern of Δ is very similar to that in Figure 3: because booms make borrowing constraints less binding by increasing savings buffers, consumption insurance is procyclical ($\beta_h - \beta_l < 0$) for the bottom 40 percent of the wealth distribution. Because booms increase both equilibrium wages and interest rates, they imply a stronger effect of idiosyncratic earnings shocks on permanent incomes. This makes consumption insurance countercyclical for the top half of the wealth distribution ($\beta_h - \beta_l > 0$). The overall pattern in Figure 4 is very similar to that in Figure 3, notwithstanding the fact that the former results from a quantitative equilibrium model, and is drawn along the equilibrium wealth distribution.

The pattern in Figure explains why insurance is procyclical on average in our benchmark model: the coefficients in Table 5 are, essentially, an average of those underlying the figure. Because insurance is strongly procyclical for households in the bottom third of the wealth distribution, but only mildly countercyclical in its upper half, the resulting average degree of comovement is procyclical.

Figure 4: Difference in individual consumption responses

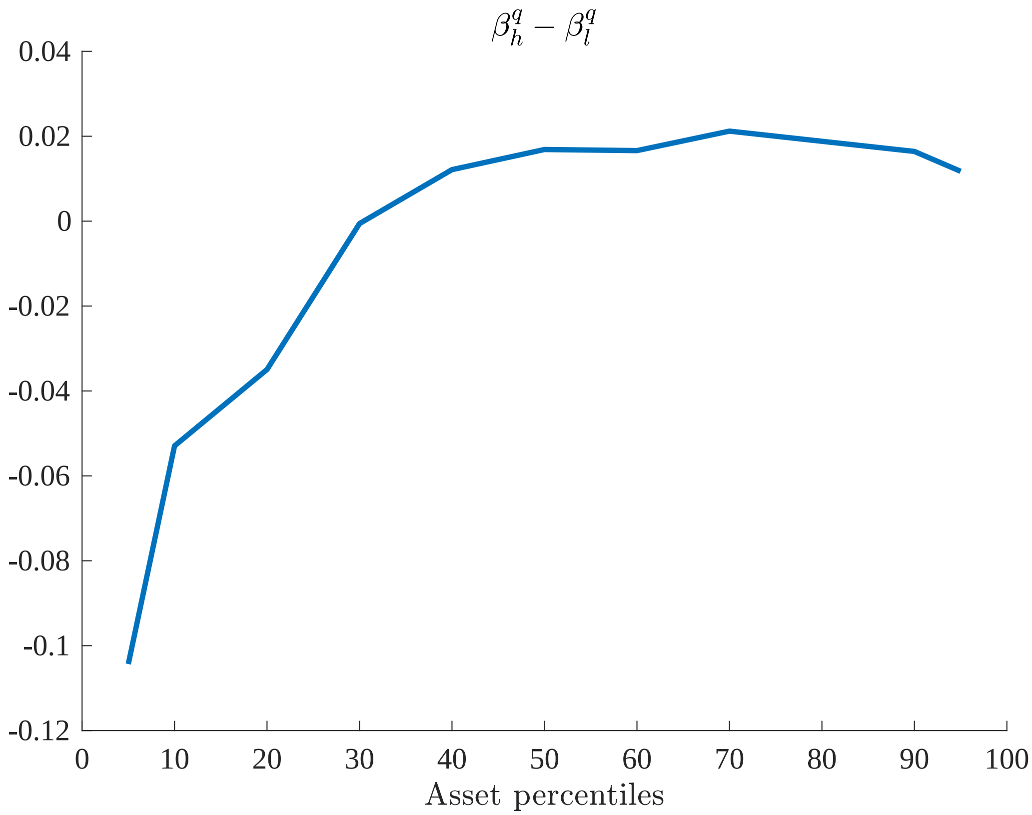

The figure depicts $\Delta = \beta_h - \beta_l$, the difference between the pass-trough coefficient in good and bad times, when both coefficients are estimated separately for households within (average) decile ranges of the equilibrium wealth distribution of the model (depicted along the bottom axes).

5.5 Consumption insurance of low-wealth PSID families

The strong heterogeneity in the cyclicality of consumption insurance along the equilibrium wealth distribution in the quantitative model raises the question of how the corresponding pattern looks in micro-data. Unfortunately the main survey of households' financial situation in the U.S. (the Survey of Consumer Finances, SCF) has no information about consumption. The CEX, in contrast, lacks data on household wealth. Only the PSID contains information for both, but wealth data is only collected at 5 year intervals prior to 2001. This leaves us with 12 observations.

Figure 5: β_t and aggregate income: low-wealth PSID households

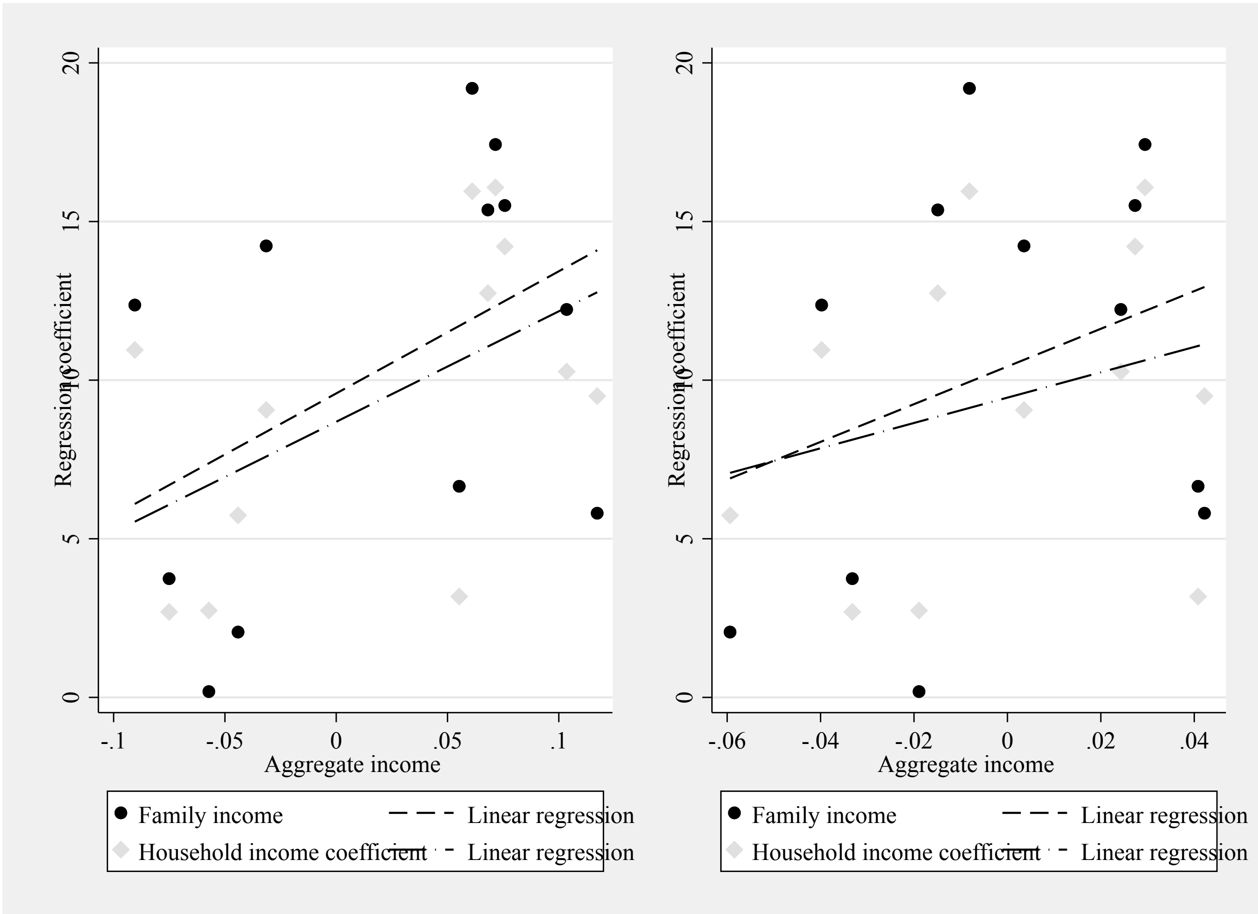

For families in the bottom wealth quartile of the PSID, the figure presents a scatter plot of β_t, in percent, in (4) estimated period-by-period (vertical axis) against detrended disposable income (PSID / NIPA, horizontal axis), biannual frequency.

We nevertheless perform the same analysis as in Section 3.2 for a subsample of the PSID families in the bottom wealth quartile. Figure 5 and Table 6 present the results. Although less precisely estimated, the point estimates of the slope parameters are if anything higher than those for the whole sample. So we find no evidence of the procyclical insurance predicted by the model even when focusing on low-wealth households in the PSID.

Table 6: Regression coefficient γ (PSID, low-wealth households)

	(1)	(2)	(3)	(4)
	1	2	3	4
Disp Y (PSID)	38.49	34.81*		
	(1.52)	(2.05)		
Disp Y (NIPA)			59.41	39.92
			(1.12)	(0.97)
Constant	9.582***	8.689***	10.43***	9.448***
	(5.20)	(6.64)	(5.66)	(6.57)
r2	0.201	0.279	0.0976	0.0748
N	12	12	12	12

t statistics in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

For families in the bottom wealth quartile, the table reports estimates of (5) using quarterly time series of coefficients β_t (using PSID Food consumption growth and respectively, family income growth (column 1 and 3) and household income growth (column 2 and 40), in percent, and log-trend deviations of aggregate disposable income (from the PSID, row 1, and from NIPA, row 2). Robust standard errors are used; stars denote conventional significance levels: * (p<.1), ** (p<.05), *** (p<.01).

6 Extensions of the SI model

From the perspective of standard incomplete-markets theory, the procyclical degree of consumption-income comovement (the countercyclical nature of consumption insurance) in U.S. microdata poses a puzzle. To explain it, this section explores two extensions of the quantitative SI model to cyclical earnings risk and biased expecations of workers' employment risk.

6.1 Cyclical income risk

Following Krueger et al. (2016), the only cyclical element of the idiosyncratic income process in the benchmark model of section 5 are the transition probabilities across employment states $\pi(s'|s, Z', Z)$. In contrast to this cyclical unemployment risk, the persistent and transitory components of earnings risk of the employed, captured by the transition probabilities $\pi(y'|y)$, were assumed to be independent of the aggregate state of the economy, and their variances thus constant over time. This income process does therefore not capture the evidence in Storesletten et al. (2004), who argue that the variance of persistent earnings shocks is countercyclical in the U.S. economy. Because time variation in earnings risk potentially provides an additional source of cyclicality in consumption insurance, we study in this section an extension of the benchmark model to cyclical fluctuations in the dispersion of persistent shocks to earnings.

A priori, we expect the degree of insurance to be lower in times of high persistent earnings risk, for two reasons: first, because optimal savings buffers increase with risk, a cyclical increase in idiosyncratic risk starts with lower-than-optimal savings buffers, which increases the effect of income shocks on consumption at high risk. When risk declines again after some time, in contrast, savings buffers are large relative to income shocks, reducing their effect on consumption during low-risk periods. Second, time variation in persistent earnings risk changes its importance relative to less persistent sources of idiosyncratic risk, in the form of unemployment shocks that are relatively short lived or purely temporary earnings risk. A high variance of persistent shocks therefore increases the correlation of current income with permanent income, and thus its comovement with consumption. This should act to improve consumption insurance in good times and reduce it in bad times, worsening the puzzle.

Nevertheless, we study quantitatively an extension of our benchmark model to countercyclical persistent earnings risk. For this, we follow Storesletten et al. (2004) by introducing time variation in the dispersion of shocks η to the persistent income component p in (38) (see Appendix 13). Specifically, we leave the average variance $\hat{\sigma}^2_{\eta,h}$ unchanged, but let $\hat{\sigma}^2_{\eta,l} = \varphi\hat{\sigma}^2_{\eta,h}$ with $\varphi > 1$, such that the persistent component of earnings risk is more volatile in bad times, and idiosyncratic risk thus countercyclical. We choose φ to target the evidence in Storesletten et al. (2004) who find a ratio of the variances of the persistent income component equal to 1.75 (see their Table 1). This ratio does not equal φ, because of additional persistent income risk from unemployment in our benchmark model. We therefore estimate a standard

Table 7: Cyclical consumption insurance with cyclical idiosyncratic risk

	β	β_{low}	β_{high}	γ
Benchmark	0.178	0.189	0.168	-0.115
Benchmark, with Y, R, W **at SS**	0.177	0.182	0.172	NaN
Weakly countercyclical risk	0.168	0.227	0.120	-0.554
... with Y, R, W **at SS**	0.156	0.214	0.123	NaN
Medium countercyclical risk	0.170	0.233	0.111	-0.645
... with Y, R, W **at SS**	0.170	0.230	0.115	NaN
Highly countercyclical risk	0.133	0.211	0.090	-0.676
... with Y, R, W **at SS**	0.134	0.212	0.092	NaN

The table reports the regression coefficient β in (4) in a simulation of the quantitative model environment extended to countercyclical idiosyncratic risk on average (column 1), as well as in times of low and high aggregate productivity (β_{low} and β_{high}, columns 2 and 3, respectively). Column 4 reports the regression coefficient γ in Equation (5), a measure of how procyclical individual consumption-income comovement is.

persistent-plus-transitory income process (that does not condition on employment status) on income panel data from a long model simulation, and set φ such that the relative variances of the estimated persistent shocks across good and bad times fit the evidence in Storesletten et al. (2004). Moreover, we consider two additional cases with, respectively, stronger and weaker degrees of cyclicality.

Table 7 presents the results. In line with the intuition, countercyclical persistent earnings risk makes insurance even more procyclical than in the benchmark version of our model. The γ coefficient is negative, and increases about six-fold in magnitude for our preferred choice of φ. Focusing on the effect of cyclical risk by eliminating aggregate fluctuations in prices and productivity (in even rows of Table 7) changes the result very little.

6.2 Biased perceptions of idiosyncratic risks

So far, the analysis has maintained the assumption of rational expectations. This implies, among other things, that households correctly perceive the idiosyncratic risks they fact, including how these vary over the business cycle. The recent literature on household expectations, and how they depart from the benchmark of rationality, has, however, established that many households have in fact a biased view of the income shocks they are likely to experience in the future.[18] The specific departure from rational expectations that we consider is motivated by Balleer et al. (2023), who show that U.S. workers have biased perceptions of their labor market prospects. In particular, they document a negative bias in perceived separation probabilities that is countercyclical: employed workers are more over-optimistic, relative to the truth, in booms than in busts.[19] This motivates us to study misperceived idiosyncratic income risk as a potential solution to the puzzle.

6.2.1 Misperceived unemployment risk

We study a particularly simple form of misperception, whereby households have correct and complete knowledge of the structure of the economy including the sequences of past and possible future exogenous states (that consist of histories of aggregate productivity and idiosyncratic shocks). Given the evidence of biased perceptions of unemployment risk (Balleer et al. (2023)), households misperceive the Markov process for employment states $\pi(s'|s, Z', Z)$, such that the perceived markov probabilities include a bias component: $\hat{\pi}_{ij}^{kl} = \pi_{ij}^{kl} + bias_{ij}^{kl}$. We concentrate on the bias of the employed and set the bias in perceived job-keeping probabilities to $bias_{ee}^{kk} = \overline{bias} + (1 - 2 * \mathbb{I}_{ll})bias^{cyc}$, where $\mathbb{I}_{jj}$ is an indicator that equals 1 when the economy remains in aggregate state j, and 0 otherwise.

[18]Rozsypal and Schlafmann (2017) document how U.S. households on average perceive their current income to be more persistent than it actually is. Balleer et al. (2023) show that U.S. households underestimate employment risk. See also Koşar and Van der Klaauw (2023), and Caplin et al. (2023).

[19]Specifically, Balleer et al. (2023) find that the bias in employed workers' perception of job-separation probabilities is more negative in times when the unemployment rate in their state of residence is below trend, see their Table 6.

Our choice of $\overline{bias}$ and $bias^{cyc}$ targets the evidence in Balleer et al. (2023), based on the U.S. *Survey of Consumer Expectations*: first, we set $\overline{bias}$ equal to 1 percentage point, equal to the average over-optimism about their job-keeping probability of employed U.S. workers. Second, to match the finding that this optimistic bias is stronger in times of below-trend unemployment, we set $bias^{cyc} > 0$. Because of the short sample of observations in the SCE, whose only recession period is that of the Covid-19 pandemic, we are somewhat agnostic about the precise magnitude of this cyclical component. While Balleer et al. (2023)'s point estimates imply a difference in bias between periods of below- and above-trend unemployment equal to 0.6 percentage points, we study a range of positive values of $bias^{cyc}$. Conditional on this, we choose the remaining perceived transition rates to be consistent with unchanged unemployment rates in booms and recessions, respectively, and with the requirement that there be no transition period in unemployment (as described in Section 5).

We denote the resulting misperceived expectation operator $\mathbb{E}$. Given this misperception, we study a minimal departure from the standard rational expectations equilibrium, similar to the conditional rational expectation equilibrium in Caines and Winkler (2021) (but noting that households in our environment misperceive the probability distribution of an exogenous, idiosyncratic variable). For this, we assume that households do not observe the cross-sectional distributions of income and assets. In the SI economy, this holds by the definition of a Krusell and Smith (1998a)-equilibrium. Given the change in household expectation operator, the computational solution of the SI economy is equally unchanged (noting that the simulation step uses the true transition matrix).

6.2.2 Results

Table 8 presents the results. Procyclical optimism about job-keeping probabilities reduces the effect of unemployment shocks on consumption in bad times, when both employment and unemployment are perceived to be more transitory. In good times,

Table 8: Cyclical consumption insurance with mis-perceived idiosyncratic risk

	β	β_{low}	β_{high}	γ
$bias^{cyc} = 0$	0.175	0.188	0.167	-0.110
$bias^{cyc} = 0$**, Y,R,W at SS**	0.177	0.182	0.172	NaN
$bias^{cyc} = 0.5pp$	0.184	0.190	0.180	-0.054
$bias^{cyc} = 0.5pp$**, Y,R,W at SS**	0.192	0.186	0.196	NaN
$bias^{cyc} = 1pp$	0.187	0.188	0.186	-0.006
$bias^{cyc} = 1pp$**, Y,R,W at SS**	0.187	0.181	0.191	NaN
$bias^{cyc} = 2pp$	0.199	0.186	0.207	0.113
$bias^{cyc} = 2pp$**, Y,R,W at SS**	0.199	0.179	0.212	NaN

For different values of $bias^{cyc}$, the table reports the regression coefficient β in (4) in a simulation of the quantitative model environment with misperceived idiosyncratic income risk on average (column 1), as well as in times of low and high aggregate productivity (β_{low} and β_{high}, columns 2 and 3, respectively). Column 4 reports the regression coefficient γ in Equation (5), a measure of how procyclical individual consumption-income comovement is.

in contrast, misperception increases the effect of income shocks on consumption, because changes in employment state are perceived to be more persistent. In the absence of aggregate movements in prices and income (rows 4,6 and 8 of Table 8), this effect makes consumption insurance countercyclical. For small values of $bias^{cyc}$, however, this effect is dominated by the pro-cyclicality of consumption insurance due to changes in aggregate income levels and prices discussed in Section 4. With a strongly procyclical bias ($bias^{cyc}$ equal to 2 percentage points) , the degree of countercyclicality in consumption insurance as measured by the coeficient γ is strongly positive, if only a third of that estimated in CEX data.

7 Beyond self-insurance: cyclical frictions with complete markets

This section looks at a more fundamental departure form the benchmark SI model where both the assets available and limits to borrowing were exogenous and independent of the state of the economy. Instead, we study an alternative ("LC") environment with endogenous financial frictions due to participation constraints that arise from limited contract enforcement.[20] Relative to stationary LC economies, whose properties are well-understood (Krueger and Perri (2011), or Broer (2013)), the introduction of aggregate fluctuations makes both sides of these participation constraint cyclical. And because the outside option to the contract is particularly attractive when aggregate output is high, this could help explain countercyclical insurance in U.S. data by tightening financial constraints in good times.

7.1 Market Structure and Competitive Equilibrium

We consider a substantially simplified version of the economy of section 5 but maintain the structure and cyclicality of idiosyncratic risk and its comovement with aggregate output. Specifically, we abstract from capital accumulation and assume a simple linear technology operated by competitive firms, such that in equilibrium wages are equal to the level of total factor productivity $Z_t \in \{Z_l, Z_h\}$. In addition, we assume that agents are infinitely lived workers (abstracting thus from any life-cycle) and that there is no earnings risk of the employed. Given the exogenous nature of labor supply, the economy is thus, essentially, an endowment economy, allowing us to write individual incomes as individual shares y_t of aggregate income Y_t as in Section 2. Independently of this normalisation, fluctuations in aggregate output are governed by the Markov process $\pi(Y'|Y)$ and fluctuations in individual incomes arise from unemployment risk in the same way as in Section 5, described by the same Markov process $\pi(y'|y, Y', Y)$.

In contrast to the SI model of the previous sections, we assume that individuals can trade a complete set of state-contigent contracts to insure against income risk. Contracts are only enforced, however, by the threat of exclusion from financial trade.

[20]Even 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.

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 consumption equals income $c_t = y_t \cdot Y_t$. The discounted utility associated with autarky is denoted by $V(y_t, Y_t)$, i.e.

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

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

$$E\left[\sum_{s=0}^{\infty} \delta^s u(c_{t+s}) \middle| \, y_t, Y_t \right] \geq V(y_t, Y_t) \tag{8}$$

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

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 $\boldsymbol{\tau} = \{\tau_t(y^t, Y^t)\}_{t=0}^{\infty}$. 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 $\boldsymbol{\tau}$ according to the profit function

$$P_{y_0, Y_0}(\boldsymbol{\tau}) = -E\left[\sum_{t=0}^{\infty} q_t(Y^t) \tau_t(y^t, Y^t)\right]. \tag{9}$$

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 $\boldsymbol{\tau}$ that, given prices $\mathbf{q} := \{q_t(Y^t)\}_{t=0}^{\infty}$, maximizes (9) subject to (8) and

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

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 Φ_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 $\mathbf{q}$, an assignment of transfer policies $\boldsymbol{\tau}$ 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 $\mathbf{q}$.
2. Markets clear, i.e.
$$\int_0^1 \tau_t(i)di = 0$$
for all $t = 0, 1, \ldots$ and all possible histories.

7.2 Analytical equilibrium characterization

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

$$\tilde{c}_t = \frac{c_t}{Y_t} = (\delta \widetilde{R}_t)^{-\frac{1}{\sigma}} \tilde{c}_{t-1} \tag{10}$$

where $\widetilde{R}_t = R_t \left(\frac{Y_t}{Y_{t+1}}\right)^{\sigma}$. Denote the cross-sectional distribution of consumption shares in period t as $\Phi_t^c : \mathbb{B}([0, y^n]) \longrightarrow [0,1]$ 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 $Y^1 = Y^2 = ... = Y^N = 1$) 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 $\Phi^c(\tilde{c}) : R^+ \longrightarrow [0,1]$ the discrete stationary distribution of consumption shares in this equilibrium, V_j the autarky value at income y^j, $j = 1, .., n$, 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_{ll} on a consumption support given by (10) starting from c^h, with a lower bound equal to $c^l = y^l$.[21]

7.2.1 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, multipicative 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.

Proposition 2 *Independence of idiosyncratic consumption movements from aggregate income*

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

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

$$R_t = R\frac{Y_t}{Y_{t-1}} \tag{11}$$

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

$$\Phi^c_t(\tilde{c}) = \Phi^c(\tilde{c}).$$

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

$$\begin{aligned} E\sum_{s=0}^{\infty}\delta^s\ln(c_{t+s}) &= E\sum_{s=0}^{\infty}\delta^s\left[\ln(\tilde{c}_{t+s})+\ln(Y_{t+s})\right] \\ &\geq V(y_t)+E\sum_{s=0}^{\infty}\delta^s\ln(Y_{t+s}) = V(y_t\cdot Y_t) \end{aligned} \tag{12}$$

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 (10) of unconstrained agents at the prices R_t given by (11). 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. ■

7.3 Quantitative analysis

Outside the special case of log-preferences, it is difficult to derive general features of consumption insurance in the LC economy. Broer (2020) derives conditions for insurance to be procyclical that do not, however, cover the empirically relevant case where both aggregate endowments and idiosyncratic income risk are cyclical and insurance is partial. The rest of this section therefore studies the LC economy quantitatively.

7.3.1 Calibration

The calibration of the economy is identical to that in Section 5, adjusted for the simplified nature of the LC economy. In particular, we keep unchanged the period-utility function, the Markov processes for aggregate and idiosyncratic risk, and the replacement rate of $\tilde{\rho} = 0.5$. Since the aggregate wealth holdings in the economy are zero, we choose the discount factor δ to target a degree of insurance in the stationary LC economy, without fluctuations in aggregate output or unemployment risk, that is close to that in the SI economy. As is well known, the threat of autarky must be counteracted by substantial impatience to avoid full insurance in equilibrium. Indeed, the value of the discount factor we find is equal to 0.69.

7.3.2 Model solution

In LC economies, neither a state space reduction of the type used to solve the standard incomplete-markets model in the previous section, following Krusell and Smith (1998a), 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, it is easy to construct examples 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 that summarise the time-varying cross-sectional distribution by a small number of moments (typically the mean of capital holdings) 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 (10)) 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 small (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.[22] 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 (10).

7.3.3 Results

For different parameterisations of the LC economy, Table 9 reports the regression coefficient β in Equation (4) for the whole sample and estimated during times of high / low productivity (β_l/β_h). The final column reports the regression coefficient γ in equation (5). The first row depicts the coefficient in the stationary economy (where we set aggregate output to the average $Z_h = Z_l = \overline{Z}$), of similar magnitude as in the SI economy. Row 2 reports results for the benchmark version of the LC economy: fluctuations in aggregate output and idiosyncratic risk together reduce the β coefficient on average by about 30 percent. This is due to a strong increase in consumption insurance in good times of high output: β_l is about three times the

[22]Note that when the lower idiosyncratic income state equals 0, low-income agents are never constrained. Ando et al. (2023) exploit this feature to analytically characterise a more general LC model with capital accumulation.

Table 9: Cyclical consumption insurance in a simple limited-commitment economy

	β	β_{low}	β_{high}	γ
Stationary	18	18	18	NaN
Benchmark	14	22	6	-2.83
$Z_h = Z_l = 1$	14	21	7	NaN
iid unemp	18	18	17	-14

The table reports the regression coefficient β in (4) in a simulation of the LC economy on average (column 1), as well as in times of low and high aggregate productivity (β_{low} and β_{high}, columns 2 and 3, respectively). Column 4 reports the regression coefficient γ in Equation (5), a measure of how procyclical individual consumption-income comovement is.

size of β_h, implying a coefficient γ is substantially larger in magnitude than both the counterpart in the SI model, and the data. So endogenous market incompleteness strengthens the puzzle.

8 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. This seems important because the costs of business cycles may be substantially altered when they affect the average degree of consumption insurance or its volatility. In addition, the cyclicality of consumption-income comovement provides an additional dimension along which to discipline models for designing redistributional or insurance policies.

We first showed that in U.S. micro data from the Consumer Expenditure Survey and the Panel Study of Income Dynamics, household consumption reacts more strongly to individual income changes in booms. We then identified key determinants of fluctuations in consumption-income comovement in standard theory where consumers self-insure against income shocks through borrowing and saving. Temporarily high aggregate income or interest rates increase the effect of individual income shocks on permanent income, and thus on current consumption for high-wealth households whose borrowing constraints don't bind. At the same time, both high income and high interest rates encourage saving, which lengthens wealth-poor consumers' smoothing horizon by reducing the likelihood of binding borrowing constraints in the future. This acts to dampen the effect of shocks on consumption. The average cyclicality of consumption smoothing thus depends on the general-equilibrium comovement between interest rates and aggregate income, as well as the equilibrium wealth distribution. In the general equilibrium of an otherwise data-consistent heterogeneous-agent model with incomplete markets and rich sources of idiosyncratic risk (Krueger et al., 2016), consumption responds less to income shocks in booms, driven by procyclical savings rates. Consumption insurance is thus procyclical in this "SI" environment, in contrast to the data.

The inability of the SI model to explain the data motivated us to explore potential solutions to this "countercyclical consumption insurance puzzle".

We first studied the ability of two extensions of the SI model to explain the puzzle. Countercyclical variance of income shocks makes consumption insurance even more procyclical, reinforcing the puzzle. This is because buffers are smaller-than-optimal at the beginning of high-volatility periods, and because increased volatility of persistent shocks increases the effect of a typical income shock on consumption. Both effects decrease consumption insurance in bad times relative to good times. Because income persistence strongly affects consumption-income comovement, we also explored a cyclical form of persistence bias (Rozsypal and Schlafmann, 2017). Specifically, in line with the evidence presented in Balleer et al. (2023), we studied a version of the benchmark SI environment where workers become optimistic about their employment prospects in booms and pessimistic in recessions. Together with an unchanged observed unemployment rate, this makes insurance less procyclical because employment and unemployment are perceived to be more persistent in booms. We show that this misperception makes consumption insurance indeed countercyclical in the SI model, and may thus help the puzzle.

We then studied a more fundamental alternative to the SI environment without exogenous restrictions to asset trade, but with endogenous, and therefore potentially cyclical, financial frictions arising from limited contract enforcement. In a parameterization where unemployment risk was similarly countercyclical as in US employment data, consumption insurance was even more procyclical than in the SI model. Importantly, while the average sensitivity of consumption to income changes in the SI economy was not much affected by aggregate fluctuations relative to a stationary version of the economy, it increased strongly in the LC environment with cyclical unemployment risk. The cyclical nature of consumption insurance thus argues against limited-commitment frictions as a main source of limited consumption risk sharing.

Because the analysis of extensions and alternatives to the standard model was exploratory in nature, these results ask for more research on the determinants of consumption insurance, in particular with endogenous financial frictions and departures from rational expectations. Our LC economy abstracted from capital investment and abstracted from other frictions. Future research should therefore generalize our analysis to account for additional endogenous state variables[23] and other frictions, such as limited information about household incomes or effort levels as in Broer et al. (2017) that may be less cyclical. It would also be interesting to investigate alternative environments with endogenous borrowing constraints 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 (Chatterjee et al., 2007, Auclert and Mitman, 2018). Similarly, the simple form of misperception we study is but an encouragement to study how departures from the benchmark of full information and rational expectations affect consumption insurance.

References

Ábrahám, Á. and E. Cárceles-Poveda (2010). Endogenous trading constraints with

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

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.

Ando, Y., D. Krueger, and H. Uhlig (2023). One-sided limited commitment and aggregate risk. Technical report, National Bureau of Economic Research.

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.

Balleer, A., G. Duernecker, S. Forstner, and J. Goensch (2023). The effects of biased labor market expectations on consumption, wealth inequality, and welfare.

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.

Bold, T. and T. Broer (2021). Risk sharing in village economies revisited. *Journal of the European Economic Association 19*(6), 3207–3248.

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. (2020). Consumption insurance over the business cycle.

Broer, T., M. Kapička, and P. Klein (2017). Consumption risk sharing with private information and limited enforcement. *Review of Economic Dynamics 23*, 170–190.

Busch, C. and A. Ludwig (2020). Higher-order income risk over the business cycle.

Caines, C. and F. Winkler (2021). Asset price beliefs and optimal monetary policy. *Journal of Monetary Economics 123*, 53–67.

Caplin, A., V. Gregory, E. Lee, S. Leth-Petersen, and J. Sæverud (2023). Subjective earnings risk. Technical report, National Bureau of Economic Research.

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. RÃos-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.

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 (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Heathcote, J., F. Perri, and G. L. Violante (2010, January). Unequal we stand: An empirical analysis of economic inequality in the united states: 1967-2006. *Review of Economic Dynamics 13*(1), 15–51.

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.

Koşar, G. and W. Van der Klaauw (2023). Workersâ perceptions of earnings growth and employment risk. *FRB of New York Staff Report* (1056).

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and Household Heterogeneity. In *Handbook of Macroeconomics*, Volume 2, pp. 843–921. Amsterdam: Elsevier.

Krueger, D. and F. Perri (2005). Understanding consumption smoothing: Evidence from the us consumer expenditure data. *Journal of the European Economic Association 3*(2-3), 340–349.

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 (1998a). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy 106*(5).

Krusell, P. and A. Smith (1998b). 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.

Lucas, R. E. (1987). Models of business cycles.

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.

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.

9 Appendix I: CEX 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 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.

10 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 interest rates are a function only of 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 (10). 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.

10.1 Existence of a 'history-independent' equilibrium

Result 1 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 $\{R_s\}, s = t, t+1, ..., \infty$ (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

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

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 $\{R_{t+1}, ..., R_{t+s}\}$ (corresponding to 2^s possible aggregate histories between $t+1$ and $t+s$) according to (10). $\mathbb{P}^s$ 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 $R_{t+1}, R_{t+2}, ..R_{t+s}$ 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 income 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

$$\begin{aligned}
0 &= \int_{i:y_{it}=y_h}\left(\frac{c_{it}}{Y^I} - \frac{c_{it-1}}{Y^J}\right)di + \int_{i:y_{it}=y_l}\left(\frac{c_{it}}{Y^I} - \frac{c_{it-1}}{Y^J}\right)di & (14)\\
&= \Phi_h\frac{c_t^{hI}}{Y^I} - \int_{i:y_{it}=y_h}\left(\frac{c_{it-1}}{Y^J}\right)di + \left[(\delta\widetilde{R}_t^{IJ})^{\frac{1}{\sigma}} - 1\right]\int_{i:y_{it}=y_l}\frac{c_{it-1}}{Y^J}di \\
&= \Phi_h\left[\frac{c_t^{hI}}{Y^I} - \left(p_{hh}\Phi_h\frac{c_{t-1}^{hJ}}{Y^J} + p_{lh}(1-\Phi_h)(1-\frac{c_{t-1}^{hJ}}{Y^J})\right)\right] & (15)\\
&\quad +(1-\Phi_h)\left[(\delta\widetilde{R}_t^{IJ})^{\frac{1}{\sigma}} - 1\right]\left(p_{hl}\Phi_h\frac{c_{t-1}^{hJ}}{Y^J} + p_{ll}(1-\Phi_h)(1-\frac{c_{t-1}^{hJ}}{Y^J})\right) & (16)
\end{aligned}$$

where $Y_{t-1} = Y^J$ is the aggregate income value in period $t-1$, $\widetilde{R}_t^{JI} = R_t^{JI}\left(\frac{Y^I}{Y^J}\right)^{\sigma}$ 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 inviduals both equal 1. (16) thus reduces to

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

which implies that $\widetilde{R}_t^{IJ}$ is strictly decreasing in c_t^{hI} (as the consumption of the unconstrained must fall by more in (10), implying a lower R_t^I, when the consumption of the constrained is higher). Also, (10) shows how consumption of the unconstrained in aggregate state I in period t is stricly increasing in R_t^I.

Result 1: 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 ($p_{ij} = p_i, \forall j = 1, 2, .., n$). 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

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

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 (10)). 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.

10.2 Computation

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

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

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, $\overline{c}^L$ and $\overline{c}^H$, construct two finite grids for consumption, each appropriate for a distinct aggregate state. Denote them by $\mathbb{c}^1$ and $\mathbb{c}^2$. 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 $\mathbb{c}^I$ is $y^1 Y^i$ and the last element of $\mathbb{c}^I$ 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 $\mathbb{c}^j_m$ in the next period if, in the current period, aggregate output is Y^i and individual consumption is $\mathbb{c}^j_k$.

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 $\mathbb{c}^i_k$. 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 $\Phi(2)$, next period's individual income share is high and $c_{t+1} = \bar{c}^j$. That is the easy case. Alternatively, with probability $\Phi(1)$, 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

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

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

$$\mathbb{c}^j_m \leq c_{t+1} \leq \mathbb{c}^j_{m+1}.$$

So we assign some of the probability mass $\Pi(j|i)\Phi(1)$ to $\mathbb{c}^j_m$ and the rest to $\mathbb{c}^j_{m+1}$. More precisely,

$$P(j,m|i,k) = \Phi(1)\Pi(j|i) \frac{\mathbb{c}^j_{m+1} - c_{t+1}}{\mathbb{c}^j_{m+1} - \mathbb{c}^j_m}$$

and

$$P(j, m+1|i,k) = \Phi(1)\Pi(j|i)\frac{c_{t+1} - \mathbb{c}^j_m}{\mathbb{c}^j_{m+1} - \mathbb{c}^j_m}.$$

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

$$P(j, m+1|i,k) = P(j, m+1|i,k) + \Phi(1)\Pi(j|i)\frac{c_{t+1} - \mathbb{c}^j_m}{\mathbb{c}^j_{m+1} - \mathbb{c}^j_m}$$

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.

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

This section derives the welfare cost approximation in Section 2 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 aggainst aggregate fluctuations but that idiosyncratic income shocks are partially insured: $\tilde{c} = \theta_t y_{it}^{\beta_t}$, such that only a time-varying fraction β_t of idiosyncratic shocks to log-income shares are passed through to log-consumption. We assume $log(y_{it})$ is i.i.d. normally distributed as $N(\mu_y, V_y)$. Also, we assume, for simplicity that β_t is distributed iid across different time periods with mean $\overline{\beta}$ variance Var_β.

Feasibility implies that income and consumption shares sum to 1 ($E[\tilde{c}] = E[y] = 1$), which yields $\mu_y = -\frac{1}{2}V_y$ and $log(\theta_t) = \frac{1}{2}V_y\beta_t(1-\beta_t)$. Note that this implies $E[log(c_{it})] = log(\theta_t) + \beta\mu_y = \frac{1}{2}V_y\beta_t(1-\beta_t) - \beta\frac{1}{2}V_y = -\frac{1}{2}V_y\beta_t^2$. For a given β_t, utility of an individual agent equals

$$
\begin{aligned}
u(c_{it}) &= \frac{(\theta_t y_{it}^{\beta_t} Y_t)^{1-\sigma}}{1-\sigma} && (19)\\
&= \frac{(e^{-\frac{1}{2}V_y\beta_t^2+\beta_t(logy_{it}-\mu_y)}Y_t)^{1-\sigma}}{1-\sigma} && (20)\\
&= \frac{e^{-(1-\sigma)\frac{1}{2}V_y\beta_t^2+(1-\sigma)\beta_t(logy_{it}-\mu_y)}Y_t^{1-\sigma}}{1-\sigma} && (21)
\end{aligned}
$$

Taking expectations over i yields

$$
\begin{aligned}
E[u(c_{it})] &= \frac{e^{-(1-\sigma)\frac{1}{2}V_y\beta_t^2+0+\frac{1}{2}(1-\sigma)^2\beta^2 Var_y]}Y_t^{1-\sigma}}{1-\sigma} && (22)\\
&= \frac{e^{\beta_t^2\frac{V_y}{2}[-(1-\sigma)+(1-\sigma)^2]}Y_t^{1-\sigma}}{1-\sigma} && (23)\\
&= \frac{e^{\sigma(\sigma-1)\frac{V_y}{2}\beta_t^2}Y_t^{1-\sigma}}{1-\sigma} && (24)
\end{aligned}
$$

Expanding in β_t and Y_t to second order yields

$$E[u(c))] \approx U(\bar{c})E\big[1+\bar{\beta}^2\sigma(\sigma-1)V_y dln\beta_t + (1-\sigma)dlnY_t$$
$$+\frac{1}{2}\sigma(\sigma-1)[\sigma(\sigma-1)V_y\bar{\beta}^2+1)]V_y\bar{\beta}^2 dln\beta_t^2 + \frac{1}{2}\sigma(\sigma-1)dlnY_t^2$$
$$-\frac{1}{2}\sigma(\sigma-1)^2\bar{\beta}^2 V_y dln\beta_t dlnY_t\big] \tag{25}$$

where $dlogx_t \approx \frac{x_t-\bar{x}}{\bar{x}}$ denotes the deviation of the $log(x_t)$ from its mean. Taking expectations, this yields

$$E[u(c))] \approx U(\bar{c})\big[1+\frac{1}{2}\sigma(\sigma-1)[\sigma(\sigma-1)V_y\bar{\beta}^2+1)]V_y\bar{\beta}^2 V_\beta \tag{26}$$
$$+\frac{1}{2}\sigma(\sigma-1)V_Y$$
$$-\frac{1}{2}\sigma(\sigma-1)^2\bar{\beta}^2 V_y Cov(\beta,Y)\big] \tag{27}$$

This gives the welfare costs of fluctuations in β_t and Y_t in terms of expected consumption

$$\Delta ln\bar{c} \approx \frac{1}{2}\sigma V_Y\left[1+\Phi_{cY}\left(V_\beta-(\sigma-1)Cov(\beta,Y)\right)\right] \tag{28}$$

where $\Phi_{cY} = \frac{V_y\bar{\beta}^2}{V_Y}$ is the variance of (idiosyncratic) consumption shares relative to that of aggregate consumption / income, and where we set $\sigma(\sigma-1)(V_y\bar{\beta}^2)^2 > 0$ to 0, which holds for realistic consumption volatility and standard levels of risk aversion.

12 Appendix IV: Proofs of results in Section 4

12.1 Derivation of Equation (7)

In the absence of uncertainty, the lifetime budget constaint in period 0 can be written as

$$C_0 + \sum_{t=1}^{\infty}\frac{1}{\Pi_{s=1}^t R_s}C_t = \sum_{t=1}^{\infty}Y_t\frac{1}{\Pi_{s=1}^t R_s}y_{it} + Y_0 y_{i0} + a_{i0} \tag{29}$$

Substituting for C_t from the Euler equation $C_t = \delta^t\Pi_{s=1}^t R_s C_0$ yields

$$C_0\sum_{t=0}^{\infty}\delta^t = \left(\sum_{t=1}^{\infty}Y_t Q_t y_{it} + Y_0 y_{i0} + a_{i0}\right) \tag{30}$$

Simplifying using $Q_t = \frac{1}{\Pi_{s=1}^t R_s}$ yields the result.

12.2 Proof of Proposition 1

Note that $\beta_{it} > \beta_i$ whenever

$$\frac{\sum_{t=0}^{\infty}\delta^t \omega_t q_t \rho_0^t}{\sum_{t=0}^{\infty}\delta^t \omega_t q_t + \alpha_{i0}} > \frac{\sum_{t=0}^{\infty}\delta^t \rho_0^t}{\sum_{t=0}^{\infty}\delta^t + \alpha_{i0}} \tag{31}$$

$$(\sum_{t=0}^{\infty}\delta^t \omega_t q_t \rho_0^t)\left(\sum_{t=0}^{\infty}\delta^t + \alpha_{i0}\right) > (\sum_{t=0}^{\infty}\delta^t \rho_0^t)\left(\sum_{t=0}^{\infty}\delta^t \omega_t q_t + \alpha_{i0}\right) \tag{32}$$

$$\frac{\sum_{t=0}^{\infty}\delta^t \omega_t q_t \rho_0^t}{(1-\delta)^{-1}}\left(\frac{1}{(1-\delta)^2} + \frac{\alpha_{i0}}{(1-\delta)}\right) > \frac{\sum_{t=0}^{\infty}\delta^t \rho_0^t}{(1-\delta)^{-1}}\frac{\sum_{t=0}^{\infty}\delta^t \omega_t q_t}{(1-\delta)^{-1}}\left(\frac{1}{(1-\delta)^2} + \frac{\alpha_{i0}}{\frac{\sum_{t=0}^{\infty}\delta^t\omega_t q_t}{(1-\delta)^{-1}}(1-\delta)}\right) \tag{33}$$

$$E_\delta\left[\omega_t q_t \rho_0^t\right]\left(\frac{1}{(1-\delta)^2} + \frac{\alpha_{i0}}{(1-\delta)}\right) > E_\delta\left[\rho_0^t\right]E_\delta\left[\omega_t q_t\right]\left(\frac{1}{(1-\delta)^2} + \frac{\alpha_{i0}}{\frac{\sum_{t=0}^{\infty}\delta^t\omega_t q_t}{(1-\delta)^{-1}}(1-\delta)}\right) \tag{34}$$

where E_δ denotes the infinite sum with respect to the discrete mass function $p(i) = \frac{\delta^i}{(1-\delta)^{-1}}, i = 1, 2, 3, ...$ Note that $\sum_{i=0}^{\infty} p(i) = 1$, and $q_t = \frac{Q_t}{\bar{Q}_t} = \frac{1}{\Pi_0^t(R_t\delta)} < (>)1$ is strictly monotonically decreasing (increasing) when $R_0 > (<)1/\delta$. Thus ω_i, q_i, ρ^i are strictly monotonic functions of i, and we have $E_\delta\left[\omega_t q_t \rho_0^t\right] > (<) E_\delta\left[\rho_0^t\right]E_\delta\left[\omega_t q_t\right]$ whenever $\omega_0 > (<)1$ and $q_0 = 1$, or $q_0 > (<)1$ and $\omega_0 = 1$, or both $\omega_0 > (<)1$ and $q_0 > (<)1$.

Statement *1.* then follows noting that $\frac{\alpha_{i0}}{E_\delta[\omega_t q_t](1-\delta)^{-2}}$ is declining in ω_0, so the right hand side of (34) is smaller than the left-hand side. Statement *2.* follows because, in that case, the right hand side of (34) is smaller than the left-hand side for $\alpha_{i0} = 0$. Since $\frac{\alpha_{i0}}{E_\delta[\omega_t q_t](1-\delta)^{-2}}$ is increasing in R_0 and thus $\frac{\alpha_{i0}}{(1-\delta)} < \frac{\alpha_{i0}}{E_\delta[\omega_t q_t](1-\delta)^{-2}}$ whenever $R_0 > \frac{1}{\delta}$, there is a threshold value $\bar{\alpha}_0$ around which the inequality flips (statement 4.).

Statement 4. follows because the right hand side is strictly increasing in ρ_0.

13 Appendix V: Details of the quantitative model in Section

For the benchmark SI model in Section 5, this section describes the household problem, defines the competitive equilibrium, and details the calibration of the remaining parameters. The exposition is identical to Krueger et al. (2016), who provide additional detail.

Household problem

The problem of a working household can be written recursively as follows

$$
\begin{aligned}
& v_W(s, e, a, \beta; Z, \Phi) = \max_{c, a' \geq 0} u(c) \\
& +\beta \sum_{(Z', s', e')} \pi\left(Z' \mid Z\right) \pi\left(s' \mid s, Z', Z\right) \pi\left(e' \mid e\right) \\
& \times\left[\theta v_W\left(s', e', a', \beta; Z', \Phi'\right)+(1-\theta) v_R\left(a', \beta; Z', \Phi'\right)\right] \\
& \textit{subject to} \\
& c+a' = \left(1-\xi(Z; \tilde{\rho})-\xi_{SS}\right) w(Z, \Phi) e\left[1-(1-\tilde{\rho}) 1_{s=u}\right] \\
& +(1+r(Z, \Phi)-\delta) a \\
& \Phi' = H\left(Z, \Phi, Z'\right)
\end{aligned}
$$

where $1_{s=u}$ is an indicator function that identifies unemployed individuals, v_W and v_R are, respectively the value functions of working and retired households, and H is a perceived law of motion that describes transitions of aggregate productivity and the cross-sectional distribution of households across individual state variables Φ. Wages w and capital returns r are functions of aggregate state productivity and the aggregate labor input (via the distribution Φ), and the payroll tax ξ depends only on current productivity Z (conditional on which labor inputs are constant) and the replacement rate $\tilde{\rho}$.

Recursive competitive equilibrium

Using this setup of the household problem, we can define a recursive competitive equilibrium following Krusell and Smith (1998b) as Value and policy functions of working and retired households, v_j, c_j, a'_j, pricing functions r, w & an aggregate law of motion H s.t.

1. Given r, w, taxes and H, v solves the household Bellman equation and c, a' are the associated policy functions.
2. Factor prices are given by

$$w(Z,\Phi) \;= ZF_N(K(Z,\Phi), N(Z,\Phi)) \tag{35}$$

$$r(Z,\Phi) \;= ZF_K(K(Z,\Phi), N(Z,\Phi)) \tag{36}$$

3. The current tax rate balances the unemp insurance system
4. Market clearing

$$N(Z,\Phi) \;= (1-\Pi_Z(u)) \textstyle\sum_{e\in Y} e\Pi(e), \quad K(Z,\Phi) \;= \displaystyle\int ad\Phi \tag{37}$$

5. H is induced by $pi(s'|s, Z', Z), \pi(e'|e), \pi(Z'|Z)$ and the optimal policy function a' for assets.

Calibration of remaining parameters

Following Krueger et al. (2016), we choose a standard Cobb-Douglas production technology with a capital share of 36 percent. We set the depreciation rate to 2.5 percent per quarter. The earnings process of the employed $\pi(e'|e)$ is the sum of a persistent component p, which follows an AR(1) process whose shocks η are i.i.d. with variance $\hat{\sigma}^2_\eta$, and a purely transitory shock ϵ with variance $\hat{\sigma}^2_\epsilon$, where p follows

$$\begin{aligned} \log(e') &= p+\epsilon \\ p' &= \phi p+\eta \end{aligned} \tag{38}$$

We set $\hat{\phi} = 0.9695$ and the variance of persistent shocks $\hat{\sigma}^2_\eta = 0.0384$. We approximate the process using a 21-state approximation. We choose a risk-aversion parameter equal to 1 (log-preferences). The discount factor δ is distributed uniformly on a support $[\bar{\delta}-\epsilon, \bar{\delta}+\epsilon]$ with $\bar{\delta} = 0.9864, \epsilon = 0.0053$. The unemployment replacement rate equals 50 percent. Together, these parameters imply a wealth Gini of 77 percent and a capital-output ratio of 10.3.