

How Much Consumption Insurance in Bewley Models with Endogenous Family Labor Supply?*

Chunzan Wu

University of Miami

Dirk Krueger

University of Pennsylvania,
CEPR, CFS, NBER and Netspar

March 26, 2018

Abstract

We show that a calibrated life-cycle two-earner household model with endogenous labor supply can rationalize the extent of consumption insurance against shocks to male and female wages, as estimated empirically by Blundell, Pistaferri, and Saporta-Eksten (2016) in U.S. data. With additively separable preferences, 43% of male and 23% of female permanent wage shocks pass through to consumption, compared to the empirical estimates of 34% and 20%. With non-separable preferences the model predicts more consumption insurance, with pass-through rates of 29% and 16%. Most of the consumption insurance against permanent male wage shocks is provided through the labor supply response of the female earner.

*We thank Greg Kaplan and Luigi Pistaferri for useful conversations at an early stage of this project. Krueger gratefully acknowledges financial support from the NSF under grant SES-0820494.

1 Introduction

How does household consumption respond to shocks to wages of the primary earner? The baseline version of the permanent income hypothesis in which a household has only one bread winner and labor supply of the primary earner is exogenous provides a sharp answer: household consumption responds to permanent wage shocks one for one, and essentially not at all to purely transitory shocks. In a sequence of important papers, Blundell, Pistaferri, and Preston (2008) and Kaplan and Violante (2010) propose to measure the magnitude of the household consumption response to earnings shocks with various persistence by consumption insurance coefficients. This measure is defined as the fraction of the variance of a given shock to log-earnings that does *not* translate into a corresponding change in log-consumption.¹ That is, if the consumption insurance coefficient for a given earnings shock is 1, household consumption growth is completely insulated from the earnings shock, and if it is zero, the earnings shock translates one for one into consumption. Blundell, Pistaferri, and Preston (2008) empirically estimate these consumption responses to transitory and permanent earnings shocks on U.S. data and find close to perfect insurance against purely transitory shocks (except for poor households), as well as substantial insurance against permanent shocks, with a consumption insurance coefficient of 35%.² Kaplan and Violante (2010) evaluate whether a calibrated incomplete-markets life cycle model in which a single-earner household faces transitory and permanent earnings shocks is consistent with the empirical evidence provided by Blundell, Pistaferri, and Preston (2008). They find that in the model households are close to fully insured against transitory earnings shocks, but that there is too little consumption insurance against permanent shocks: the model-

¹Formally, denoting by c_{it} the log of consumption of household i at time or age t , the consumption insurance coefficient for earnings shock x_{it}^n of type (persistence) n is defined as

$$\phi_t^n = 1 - \frac{Cov_i(\Delta \log(c_{it}), x_{it}^n)}{Var_i(x_{it}^n)},$$

where Cov_i and Var_i are taken with respect to the cross-section of household observations indexed by i observed at time (or age) t .

²Also see Santaeuàlia-Llopis and Zheng (2017) for a recent application of the same method to Chinese data.

implied consumption insurance coefficient ranges between 7% and 22%, depending on how tight borrowing constraints are assumed to be, and thus is significantly smaller than the estimates of Blundell, Pistaferri, and Preston (2008).

In both Blundell, Pistaferri, and Preston (2008) and Kaplan and Violante (2010), as in much of the empirical and model-based consumption literature household earnings are treated as exogenous, and the key mechanism through which consumption insurance is achieved is saving (and borrowing, if permitted) at a state non-contingent interest rate. By construction this literature is therefore silent about the underlying shocks behind earnings fluctuations, as well as the alternative mechanisms by which households can and do respond to these underlying shocks. The current paper aims at contributing to the literature by modeling the fundamental sources of consumption risk in an otherwise standard incomplete-markets model with idiosyncratic shocks (of various persistence) to wages of the male and female earner of a two-member household. Our primary *model-based* contribution is to quantify the extent to which *wage* shocks translate into consumption movements, and to evaluate the relative importance of *alternative mechanisms* (adjustment of labor hours of both members of the household and participation of the female earner, as well as precautionary savings) by which consumption insurance occurs in the model.

The key *applied* contribution of this paper is the assessment whether the standard Bewley model with endogenous labor supply of two-earner households can match well the empirically estimated labor supply and consumption responses to transitory and permanent *wage* shocks, as derived in the recent paper by Blundell, Pistaferri, and Saporta-Eksten (2016), henceforth BPS. In this work, which is the natural extension of Blundell, Pistaferri, and Preston (2008) to endogenous household labor supply choices, the authors empirically estimate the transmission coefficients from transitory and permanent wage shocks to labor earnings and consumption in two-earner (male and female) households.³ We treat their estimates as the empirical benchmark against which the transmission coefficients estimated

³For consumption, their transmission coefficients have exactly the same interpretation as the consumption insurance coefficients discussed above, but are now understood as measuring the degree consumption insurance against *wage* rather than earnings shocks. With single earner households and exogenous labor supply the two coincide exactly.

from model-simulated data should be compared, very much in the same spirit Kaplan and Violante (2010) used the empirical estimates of Blundell, Pistaferri, and Preston (2008).

Our findings suggest that the model fits the data quite well. That is, a standard Bewley model with two-earner households and endogenous labor supply can explain most of the responsiveness of household consumption and labor income to wage shocks in the data, in contrast to Kaplan and Violante (2010) who concluded that, treating *income as exogenous*, the standard Bewley model predicts significantly too little consumption insurance against permanent income shocks. In our calibrated model with additively separable preferences between consumption and labor supply, about 43% of male permanent wage shocks and 23% of female permanent wage shocks pass through to household consumption. The corresponding empirical estimates of BPS are 34% and 20%, respectively. Hence the model can account for virtually all of the empirically observed consumption insurance against female wage shocks and most (86%) of the insurance against permanent shocks to male wages. The insurance against transitory wage shocks is almost perfect in the model with close-to-zero pass-through rates for both male and female shocks, while the BPS counterparts are slightly negative but not far away from zero. In addition, both the model and the empirical BPS results indicate that consumption insurance against permanent wage shocks improves over the life cycle, in the model caused by an improved asset position as households age. A decomposition of consumption insurance against *male* permanent wage shocks in the model shows that labor supply responses of the *female* member of the household account for most of this insurance, with a contribution that is almost constant over the life cycle. In contrast, the contribution of male labor supply is negative (that is, male hours fall in response to a permanent negative male wage shock), increasingly so over the life cycle.

The empirical analysis of BPS also finds empirical evidence for complementarity between consumption and leisure, and between leisure of both spouses. Hence we also consider a version of the model with non-separable utility between consumption and labor supply of constant elasticity of substitution functional form and calibrate it to match the complementarities estimated by BPS. Non-separability between consumption and labor supply in preferences increases the amount of con-

sumption insurance against permanent wage shocks in the model: the pass-through rates of male and female permanent wage shocks in the model are now 29% and 16%, slightly lower than the BPS estimates. Non-separable preferences also help the model match the negative pass-through rates of transitory wage shocks to consumption that BPS document empirically.

In order to make sure that the comparison between our model results and the empirical evidence, as represented by the BPS estimates, is not contaminated by potential bias of the estimation method BPS employ, we also evaluate their empirical approach using model-simulated data. What are the potential sources of the bias? For this, note that the derivation of the BPS estimators from a theoretical model with endogenous labor supply and incomplete asset markets requires interior solutions of the household maximization problem, which is not always assured in a model with potentially binding borrowing constraints or an operative extensive margin of female labor supply. However, we show that the performance of the BPS estimator is not affected too strongly by the violation of these assumptions. This is partly due to the fact that BPS restrict attention to households aged 30 to 57, and when we follow this sample selection criterion in model-simulated data most of the households in this age are no longer impacted by the tight borrowing constraints (as they have accumulated away from it). The extensive labor supply margin induces relatively larger biases in the estimates related to female labor supply, but the impact of this model feature is limited, because the female non-participation rate is only moderate, both in our simulated data as well as in the original BPS data set. Hence a potentially operative extensive margin of female labor supply is not a significant source of bias either.

Overall, our results suggest their estimation method performs well on simulated data, and hence the results discussed in the previous paragraphs concerning the comparison between the model-based and the empirical BPS results are not significantly affected by potential biases in the estimation method. The BPS method tends to underestimate consumption insurance against male permanent wage shocks, whereas the direction of the bias on the consumption insurance coefficient against female permanent wage shocks depends on the form of preferences (separable vs. nonseparable preferences). In each case, however, the magnitude of these biases is small,

around 2% in terms of the pass-through rates.⁴

Abstracting from the papers by Blundell, Pistaferri, and Preston (2008), Kaplan and Violante (2010) and Blundell, Pistaferri, and Saporta-Eksten (2016) that directly motivate this study, our work is related to the broader literature that has studied heterogeneous household models with idiosyncratic risks, as pioneered in Bewley (1986), Imrohoroğlu (1989), Huggett (1993) and Aiyagari (1994). The structural life cycle model we employ is most closely related to the models analyzed by Heathcote, Storesletten, and Violante (2010) and Holter, Krueger, and Stepanchuk (2018). However, their applications mainly focus on inequality and tax policy rather than the consumption insurance question we address here, and they do not consider the possibility of non-separable preferences. The paper is also related to the literature on within-household risk-sharing and the role female labor supply plays in this context. For example, Attanasio, Low, and Sánchez-Marcos (2005) study the importance of female labor supply as an insurance mechanism against idiosyncratic income risk within the family, but in their model the labor supply decision is discrete and the intensive margin of labor supply is absent. Ortigueira and Siassi (2013) investigate the impact of within-household risk-sharing on household labor supply and savings. However, only idiosyncratic unemployment risk is considered and there is no life cycle in their model. To our knowledge, no previous work has examined the issue of consumption insurance against wage shocks in a structural life cycle model of two-earner households with endogenous labor supply of both household members.⁵

The rest of this paper proceeds as follows: Section 2 and Section 3 describe the model and its calibration; Section 4 introduces the transmission coefficients and offers a brief review of the BPS method; Section 5 reports and discusses the

⁴BPS add an “*outside insurance*” coefficient to some of their empirical specifications, in order to capture potential insurance obtained through channels other than endogenous labor supply and self-insurance through savings. We find that the estimate of this outside insurance is more significantly downward biased when implemented on model-simulated data, in the order of about 10%, which might in turn bias the estimates of consumption insurance, which depend on the “outside insurance” coefficient, more noticeably.

⁵Perhaps closest to our work in this regard is Heathcote, Storesletten, and Violante (2014) who study consumption insurance against wage shocks in an economy populated by single households making endogenous labor supply decisions and having access to within-group risk-sharing arrangements.

main results; Section 6 evaluates the performance of the BPS method and Section 7 concludes. The appendix contains supplementary results as well as details of the model calibration and implementation of the BPS methodology.

2 Model

2.1 Environment

We study a partial equilibrium life cycle model with idiosyncratic wage risk and endogenous labor supply. We follow a cohort of a continuum of measure 1 households over their life cycle. Households live for T periods, from age $t = 1$ to T , work in the first R periods of life and then are retired from age $R + 1$ onward. Each household has two members: a male and a female. As in BPS we use the notation $X_{j,t}$ to denote the variable X of earner j at age t , with $j = 1$ and 2 corresponding to the male and female members of the household.

In each period, households receive utility from joint household consumption, C_t . A working household's utility is also affected by the levels of their labor supply, $H_{1,t}$ and $H_{2,t}$. Hence the period utility function is assumed to be $u(C_t, H_{1,t}, H_{2,t})$ for a working household and $u^R(C_t)$ for a retired household. The two members of each household are assumed to make joint decisions on consumption and labor supply. Both members of a working household can work and receive wages determined by their own labor productivity. The wages of two members in each household, $W_{j,t}$, are stochastic and can be represented by

$$\log W_{j,t} = g_{j,t} + F_{j,t} + u_{j,t},$$

$$F_{j,t} = F_{j,t-1} + v_{j,t},$$

$$\begin{bmatrix} v_{1,t} \\ v_{2,t} \end{bmatrix} \sim iid N \left(0, \begin{bmatrix} \sigma_{v_1}^2 & \sigma_{v_1,v_2} \\ \sigma_{v_1,v_2} & \sigma_{v_2}^2 \end{bmatrix} \right),$$

$$\begin{bmatrix} u_{1,t} \\ u_{2,t} \end{bmatrix} \sim iid N \left(0, \begin{bmatrix} \sigma_{u_1}^2 & \sigma_{u_1,u_2} \\ \sigma_{u_1,u_2} & \sigma_{u_2}^2 \end{bmatrix} \right),$$

where $g_{j,t}$ is the deterministic mean life-cycle trend of $\log W_{j,t}$; $F_{j,t}$ is the permanent

stochastic component and $u_{j,t}$ is the transitory component of log-wages. Hence $v_{j,t}$, is the permanent shock to earner j 's wage, and $u_{j,t}$ is the transitory shock. Both $v_{j,t}$ and $u_{j,t}$ can be correlated across the two members of each household, but are assumed to be independent between each other and over time. After retirement, labor productivity is assumed to fall to zero, and hence households optimally do not work in retirement. A retired household receives a fixed amount of social security benefits, b , in each period she is alive.⁶

As is common in standard incomplete-markets models, households cannot trade fully state contingent Arrow securities, but they can save, and potentially borrow, at the risk-free interest rate r . They are, however, subject to age-dependent and potentially binding borrowing constraints \underline{A}_t .

2.2 Household Optimization Problem

A working household's problem is then in recursive form

$$\begin{aligned}
 V(A, F_1, F_2, u_1, u_2, t) &= \max_{C, A', H_1, H_2} u(C, H_1, H_2) \\
 &+ \frac{1}{1 + \delta} \sum_{(F'_1, F'_2)} \pi(F'_1, F'_2 | F_1, F_2) \sum_{(u'_1, u'_2)} \pi(u'_1, u'_2) V(A', F'_1, F'_2, u'_1, u'_2, t + 1) \\
 \text{s.t. } & C + A' = W_{1,t}H_1 + W_{2,t}H_2 + (1 + r)A, \\
 & C, H_1, H_2 \geq 0, A' \geq \underline{A}_{t+1},
 \end{aligned}$$

where $\frac{1}{1+\delta}$ is the time discount factor and $\pi(\cdot|\cdot)$ governs the transition probabilities of the wage shocks.⁷ The dynamic programming problem of a retired household is

⁶We abstract from a progressive labor income tax in the model because BPS find that adding this feature to their analysis does not significantly change the results.

⁷Since we will discretize the support of the wage shocks when computing the model, we represent the conditional expectation in the dynamic programming problem as a sum, rather than an integral.

given by:

$$\begin{aligned}
V^R(A, t) &= \max_{C, A'} u^R(C) + \frac{1}{1 + \delta} V^R(A', t + 1) \\
\text{s.t. } C + A' &= b + (1 + r)A, \\
C &\geq 0, A' \geq \underline{A}_{t+1},
\end{aligned}$$

where $u^R(C)$ is the utility function for retired households.

2.3 Choice of the Utility Function

As a benchmark, we assume that the period utility function is additively separable in household consumption and labor of both household members, and has the form:

$$u(C, H_1, H_2) = \frac{C^{1-\sigma}}{1-\sigma} - \psi_1 \frac{H_1^{1+\eta_1^{-1}}}{1+\eta_1^{-1}} - \psi_2 \frac{H_2^{1+\eta_2^{-1}}}{1+\eta_2^{-1}},$$

where the parameter σ governs the intertemporal elasticity of substitution for consumption, and its reciprocal is the Frisch elasticity of consumption with respect to its own price. The parameters η_1 and η_2 are the Frisch elasticities of male and female labor supply with respect to their own wages. The advantage of using this preference structure, apart from simplifying the computation of the model,⁸ is that the intertemporal and Frisch labor supply elasticities are exclusively determined by exogenous parameters which are therefore directly interpretable.

However, this separable preference structure is restrictive in that it does not permit hours worked to affect the marginal utility of consumption (and reversely, does not allow consumption to impact the marginal disutility of hours worked). Furthermore, BPS argue forcefully that non-separabilities between consumption and labor supply are useful for fitting the micro data. Therefore we also study non-separable

⁸With additively separable preferences between consumption and labor supply, the Euler equation of consumption is identical to that in a model with the constant relative risk aversion (CRRA) utility function and exogenous labor supply. Optimal levels of labor supply are easily pinned down by the intratemporal optimality conditions, and are simple functions of optimal consumption in the current period.

preferences of constant elasticity of substitution form:

$$u(C, H_1, H_2) = \frac{\{\alpha C^\gamma + (1 - \alpha)[\xi H_1^\theta + (1 - \xi)H_2^\theta]^{-\frac{\gamma}{\theta}}\}^{\frac{1-\sigma}{\gamma}} - 1}{1 - \sigma}. \quad (2.1)$$

Here γ governs the substitution pattern between consumption and labor supply, and θ governs the substitution pattern between male and female labor supply. The main advantage of this utility function is that it is flexible enough to accommodate different substitution patterns while keeping the computational burden of the problem manageable.⁹ However, now the simple mapping between the preference parameters and Frisch elasticities is lost, in that the Frisch elasticities are no longer deep parameters, but rather depend on the endogenous choices by households.

Given the fact that a significant proportion of females do not participate in the labor market, an operative extensive margin of female labor supply is included in the model by adding a fixed utility cost to the period utility when female hours worked are strictly positive. That is, we add a term $-\mathbf{I}(H_2 > 0)f$ to the period utility where $\mathbf{I}(\cdot)$ is the indicator function and f is the fixed cost.

The period utility function for a retired household is chosen to be of CRRA form:

$$u^R(C) = \Psi \frac{C^{1-\sigma} - 1}{1 - \sigma},$$

where Ψ is a constant that scales marginal utility of consumption after retirement.¹⁰

3 Calibration Strategy

In this section we describe how we parameterize the model, using empirical targets derived from U.S. household data, as measured in the PSID.

⁹Optimal labor supply can still be represented in closed form as a function of optimal consumption.

¹⁰For additively separable preferences, Ψ is simply set to 1. For non-separable preferences, Ψ is chosen such that the age profile of consumption is smooth at retirement. Consumption typically falls upon retirement in the data, but the literature on the retirement consumption puzzle shows that it is mainly due to work-related consumption expenditures not modelled here. See Hurst (2008) for a survey of this literature.

3.1 Data

In BPS, the data used are from the 1999 to 2009 waves of the Panel Study of Income and Dynamics (PSID). The PSID collects data from two groups of households: one group representative of the U.S. population, the other from low income households in the Survey of Economic Opportunity (SEO). BPS' estimation uses only data from non-SEO households with male household heads aged between 30 and 57 that are married and participate in the labor market. Because we want to compare the results of our model with the empirical BPS estimates, the model is calibrated to match the statistics from precisely this group of households whenever possible. This calibration strategy gives the best chance to the model to fit the BPS estimates. If we still find significant differences between our model and the BPS empirical results, they are likely caused by mis-specification of the model rather than inappropriate parameter values. We now briefly describe the calibration of the model, with details relegated to Appendix C.

3.2 Demographic and Initial Wealth

Households are assumed to be born at age 21, retire at age 65 and die at age 80. So age 1 in the model corresponds to age 21 in the data, and T and R are 60 and 45, respectively. Initial household asset levels are set to be zero.¹¹

¹¹The median age at first marriage in the U.S. between 2000 and 2010 is 27.5 for males and 26 for females, according to U.S. Census data. The starting age of 21 of households in the model is therefore younger than in the data. There are mainly two reasons why we made this choice. First, the starting age of a couple cohabitating and thus sharing risk could be much younger than that of eventual marriage. Second, at the time of official marriage couples likely have already accumulated some assets, and have permanent wage components determined by previously realized sequences of shocks. It is difficult to empirically identify the permanent components of wages at individual level, and thus problematic to measure an exact empirical *joint distribution* of the permanent wage components and asset level (the initial states of simulated households). Thus, as an imperfect compromise, we assume that the life cycle of a household begins at an earlier age with zero assets. When calculating the relevant model statistics, only simulated data from households with ages 30 to 57 are used (the same age group employed by BPS), and therefore the choice of initial conditions is not critical for our results.

3.3 Wage Process

The deterministic life cycle profile of wages is taken from Rupert and Zanella (2015). They estimate this wage profile from PSID 1967-2008. The original wage profile starts from age 23 and only has values biennially after age 52, so we interpolate it to an annual profile and extend the age range to 21-65. Because Rupert and Zanella (2015) report only a pooled wage profile for males and females, and the estimation of female wage profile often suffers from the selection bias from females' participation decisions, in the model we assume that deterministic life cycle wage profiles of male and female earners have the same shape, but at different levels.¹² The life cycle average of the male wages is normalized to be 1. In the BPS data, average annual earnings of males is 2.031 times that of working females; we calibrate the level of female wages to match this earnings ratio.

Turning to the stochastic component of wages, the covariance matrices of transitory and permanent wage shocks are taken directly from the BPS estimates. As in their work, both the permanent and transitory wage shocks are assumed to be *iid* across time but potentially positively correlated between the two earners of a household:

$$\begin{bmatrix} \sigma_{u_1}^2 & \sigma_{u_1, u_2} \\ \sigma_{u_1, u_2} & \sigma_{u_2}^2 \end{bmatrix} = \begin{bmatrix} 0.0275 & 0.0058 \\ 0.0058 & 0.0125 \end{bmatrix},$$

$$\begin{bmatrix} \sigma_{v_1}^2 & \sigma_{v_1, v_2} \\ \sigma_{v_1, v_2} & \sigma_{v_2}^2 \end{bmatrix} = \begin{bmatrix} 0.0303 & 0.0027 \\ 0.0027 & 0.0382 \end{bmatrix}.$$

3.4 Borrowing Limit

Since the tightness of the borrowing limit is an important determinant of consumption smoothing opportunities, we consider two extreme cases, and present calibrations and results for both. The first alternative is to rule out borrowing altogether and impose a borrowing constraint of zero (ZBC) on households. The second extreme is to permit the household to borrow freely, subject to certain repayment of

¹²Since the deterministic wage trends are perfectly predictable by households in the model, the behavioral response of households with respect to *wage shocks*, the main focus of this paper, is not significantly affected by the precise life cycle profile of the deterministic wage component.

debt at the end of life. We call the associated constraints non-binding borrowing constraints (NBBC). Reality likely lies somewhere between these two extremes.¹³

3.5 Discount Factor and Interest Rate

We adopt BPS' choice of the real risk-free interest rate of $r = 2\%$ per year. The discount factor δ , a key determinant of household precautionary (and life cycle) saving is calibrated such that the aggregate wealth-income ratio for households aged 30 to 57 in the model is approximately equal to 3, the value in the BPS data set.

3.6 Social Security Benefits

In the U.S., social security benefits are piecewise linear functions of average monthly earnings over the working life. Additional rules govern benefits for spouses. A full representation of the U.S. social security system is costly in terms of computation and the complexity of the model. Hence we model the progressivity of the U.S. social security benefit formula starkly, by assuming that benefits b per household are independent of past contributions. We calibrate b to match the ratio of average social security benefits per household to average household income. In 2012, the *monthly* average social security benefit of a retired worker is \$1262 (Social Security Administration report 2013), and household annual mean income is \$71274 dollars (Current Population Survey 2013). Hence the ratio of the retirement benefit per household to average household income in the model is calibrated to $\frac{1262 \times 2 \times 12}{71274} = 42.5\%$.

¹³Note that with endogenous labor supply the specification of the NBBC case is not obvious. The natural borrowing limits employed by Kaplan and Violante (2010) are the loosest constraints that guarantee, in their environment, that households can pay back their debts with probability one. However, with unbounded endogenous labor supply, households can pay back any amount of debt by working sufficiently long hours. Therefore, we need to use a different set of non-binding borrowing constraints. In particular, we first solve the model with very large negative constraints, and check that in simulations no household is ever at the constraints. Then we trace the lowest asset positions reached by different age groups in any simulation, and use these as sequence of age-dependent borrowing constraints to solve the model again. This procedure is repeated until the implied borrowing constraints do not change any further.

3.7 Fixed Utility Cost of Positive Female Labor Supply

In the data, both males and females have non-trivial non-participation rates in the labor market. However, since the BPS results are based on a sample of households with a working male member, we do not include an extensive margin of male labor supply decision in the model. To take into account female labor supply decisions along the extensive margin in the model, the fixed utility cost of positive female labor supply f is chosen such that the average female non-participation rate of age 30 to 57 households is the same 20% as in the BPS data.

3.8 Additively Separable Preference Parameters

With additively separable preferences, the parameters $1/\sigma$, η_1 , and η_2 are the consumption, male and female labor supply Frisch elasticities with respect to their own prices, i.e., $\eta_{c,p}$, η_{h_1,w_1} and η_{h_2,w_2} .¹⁴ Therefore, we directly adopt the values BPS estimate under the assumption of separability: $\sigma = 1/0.578$, $\eta_1 = 0.528$, and $\eta_2 = 0.850$. The values of the parameters ψ_1 and ψ_2 scaling the disutility of labor are calibrated to match average male and female hours worked in the data. For males, mean hours are normalized to 1, which means that one unit of labor supply in the model corresponds to average hours worked by males in the BPS data, that is, 2302 hours per year. Average hours worked by females, conditional on working, in BPS data set is 1688 hours, and hence the target value for average female hours, conditional on working, is 0.733 in the model.

Table 1 summarizes the calibration of the model with separable preferences.

3.9 Parameters for Non-separable Preferences

With non-separable preferences, Frisch elasticities are no longer deep parameters. Given the functional form in this paper, we are able to derive the Frisch elasticities as functions of preference parameters and allocations. The formulas

¹⁴Following the original BPS paper, Frisch elasticities are denoted by η . The meaning of subscripts are c for consumption, h_j for earner j 's labor supply, p for the price of consumption and w_j for earner j 's wage. For example, $\eta_{c,p}$ is the Frisch elasticity of consumption with respect to its own price.

Table 1: Parameter Values for Additively Separable Preferences

Parameter	Governing	ZBC	NBBC
δ	discount rate of utility	0.0111	0.0093
ψ_1	disutility of male labor supply	0.499	0.523
ψ_2	disutility of female labor supply	0.524	0.550
b	retirement benefits	0.910	0.900
σ	inverse of consumption Frisch elasticity	1/0.578	1/0.578
η_1	male labor supply Frisch elasticity	0.528	0.528
η_2	female labor supply Frisch elasticity	0.850	0.850
f	fixed utility cost of female labor participation	0.0168	0.0175
$e^{g_{1,t}-g_{2,t}}$	male/female wage trend ratio	2.07	2.06

are shown in Appendix D. The preference parameters are then calibrated jointly to match the Frisch elasticities estimated by BPS in the absence of the separability assumption, and average hours worked by males and females in the BPS data. The parameters γ , θ and σ mainly affect the Frisch elasticities, whereas the parameters α and ξ mainly affect labor supply. In particular, for a given set of parameters, the model is first solved and a panel of household data are simulated. Then, based on the simulated data and the formulas for Frisch elasticities, the sample averages of Frisch elasticities and labor supply are calculated and compared with the calibration targets.¹⁵

The values of calibrated parameters of the model with non-separable preferences and the implied Frisch elasticities are summarized in Tables 12 and 13 of Appendix C.

4 Transmission Coefficients and the BPS Method

The main applied purpose of our paper is to quantify how household labor supply, income, and ultimately consumption respond to permanent and transitory wage

¹⁵Note that, given the functional form of the utility function, there are not enough degrees of freedom to match all Frisch elasticities perfectly, and we focus on matching the upper triangular part of the matrix of Frisch elasticities, in particular, $\eta_{c,p}$, η_{c,w_1} , η_{c,w_2} , η_{h_1,w_1} , η_{h_1,w_2} and η_{h_2,w_2} . In fact, there might be no utility function that can match exactly all Frisch elasticities estimated by BPS, due to the theoretical restrictions between Frisch elasticities imposed by their definitions. In the calibration of the non-separable utility function we then need to choose between a good fit of own-price elasticities ($\eta_{c,p}$, η_{h_1,w_1} and η_{h_2,w_2}) and cross-elasticities (η_{c,w_1} , η_{c,w_2} and η_{h_1,w_2}). The chosen parameterization is a compromise between both.

shocks of its two members. In this section we briefly review the concept of the transmission coefficients that BPS use to measure these responses empirically, and provide a description of the BPS method used to estimate them from PSID data.¹⁶ We also discuss possible biases and concerns associated with their method, and describe how we can use our model to assess the potential magnitude of these biases.

Let lower case letters denote logarithms of a variable, so that $\Delta c_t, \Delta y_{j,t}$ are the growth rates of household consumption and labor income of household member j .¹⁷ The response of these household variables to permanent and transitory wage shocks are captured by the transmission coefficients κ in the following equation:

$$\begin{bmatrix} \Delta c_t \\ \Delta y_{1,t} \\ \Delta y_{2,t} \end{bmatrix} = \begin{bmatrix} \kappa_{c,u_1} & \kappa_{c,u_2} & \kappa_{c,v_1} & \kappa_{c,v_2} \\ \kappa_{y_1,u_1} & \kappa_{y_1,u_2} & \kappa_{y_1,v_1} & \kappa_{y_1,v_2} \\ \kappa_{y_2,u_1} & \kappa_{y_2,u_2} & \kappa_{y_2,v_1} & \kappa_{y_2,v_2} \end{bmatrix} \begin{bmatrix} \Delta u_{1,t} \\ \Delta u_{2,t} \\ v_{1,t} \\ v_{2,t} \end{bmatrix}. \quad (4.1)$$

The transmission coefficients measure directly how consumption and labor income respond to wage shocks of different persistence. For example, κ_{c,v_j} measures how consumption responds to earner j 's permanent wage shocks. A value of $\kappa_{c,v_j} = 0.4$ means that 40% of earner j 's permanent wage shocks pass through to household consumption, and hence 60% are insured.

If transitory and permanent wage shocks were separately observable, Equation (4.1), and thus the transmission coefficients, could be estimated directly. However, in practice, only the sum of transitory and permanent wage shocks are observed directly in the data:

$$\Delta w_{j,t} = \Delta u_{j,t} + v_{j,t}.$$

A key contribution of BPS is to provide an empirically applicable method to estimate these transmission coefficients. The authors show that if one log-linearizes the first-order conditions and the intertemporal budget constraint of a two-earner household life cycle model very similar to the one described in Section 2, and as-

¹⁶Detailed derivations of their key formulas relevant for this paper, and the details of how we implement the BPS method are provided in Appendix E.

¹⁷Empirically, c_t and $y_{j,t}$ are the residuals of log consumption and log labor income of earner j at age t after controlling for the effects of household observable characteristics.

sumes interior solutions (thus abstracting from binding borrowing constraints and extensive margin labor supply decisions), then one can derive the transmission coefficients as functions of Frisch elasticities $\{\eta\}$, (human) wealth shares $(\pi_t, s_{j,t})$ and an “outside insurance” coefficient β without fully solving the structural model, as we do in this paper. For example, the transmission coefficient from permanent male wage shocks to consumption κ_{c,v_1} is given as

$$\kappa_{c,v_1} = \frac{(-\eta_{c,p} + \eta_{c,w_1} + \eta_{c,w_2})[\eta_{c,w_1} - (1 - \beta)(1 - \pi_t)(s_{1,t} + \overline{\eta_{h,w_1}})]}{\eta_{c,p} - \eta_{c,w_1} - \eta_{c,w_2} + (1 - \beta)(1 - \pi_t)(\overline{\eta_{h,p}} + \overline{\eta_{h,w_1}} + \overline{\eta_{h,w_2}})} + \eta_{c,w_1}, \quad (4.2)$$

where the η 's are individual Frisch elasticities and $\overline{\eta}$'s are weighted average Frisch elasticities in the household.¹⁸ As the authors argue, π_t is well approximated by the share of assets in total discounted wealth of a household at age t ,

$$\pi_t \approx \frac{Asset_t}{Asset_t + Human\ Wealth_t},$$

where human wealth is the expected value of the discounted future labor income stream of the household and $s_{j,t}$ is approximately the share of earner j 's human wealth in the total human wealth of the household,

$$s_{j,t} \approx \frac{Human\ Wealth_{j,t}}{\sum_{i=1}^2 Human\ Wealth_{i,t}}.$$

BPS introduce the parameter β to capture sources of household insurance that are not explicitly present in their model (neither in ours), such as the insurance provided by networks of relatives and friends. The higher the values of π_t and β are, the less consumption responds to wage shocks, and therefore the better these shocks are insured. The baseline results of BPS are estimated with the restriction of no outside insurance, $\beta = 0$.

Thus BPS' method for estimating the transmission coefficients κ encompasses four steps: (1) Estimate the variance-covariance matrices of the permanent and transitory shocks directly from wage data (with results that were documented in Section 3). (2) Measure the wealth shares π_t and $s_{j,t}$ directly from the asset and labor in-

¹⁸Specifically, $\overline{\eta_{h,p}} \equiv \sum_{j=1}^2 s_j \eta_{h_j,p}$, $\overline{\eta_{h,w_1}} \equiv \sum_{j=1}^2 s_j \eta_{h_j,w_1}$ and, as $\overline{\eta_{h,w_2}} \equiv \sum_{j=1}^2 s_j \eta_{h_j,w_2}$.

come data. (3) Conditional on the results obtained from the first two steps, and using the empirical second-order moments of Δc_t and $\Delta y_{j,t}$, employ a generalized method of moments (GMM) strategy to jointly estimate the Frisch elasticities η 's (and the “outside insurance” coefficient β unless restricted to zero).¹⁹ (4) Calculate the estimates of the transmission coefficients based on the formulas (such as Equation (4.2)) for each household and report as results sample averages across households.

Several assumptions are imposed in the BPS approach which delivers a transparent and empirically operational, and thus very powerful methodology. Violations of these assumptions of course may result in biased estimation results for the Frisch elasticities in step (3) of their procedure. For example, the log-linearization of the household optimality conditions requires interiority of saving and labor supply decisions. This assumption is systematically violated if household borrowing constraints are frequently binding or if nonparticipation of at least one household member is ubiquitous.²⁰ As part of the contribution of the paper, we therefore evaluate, in Section 6, the quality of the BPS methodology for *simulated* data generated from our model with an extensive labor supply margin and potentially binding borrowing constraints for which we know the true parameters, Frisch elasticities and transmission coefficients. But first we study, in Section 5, after briefly documenting basic life cycle properties of the model, its implied wage-shock transmission coefficients. The twofold question we answer is how well is consumption (and labor income) insured against wage shocks in the model, and how does this degree of insurance compare with the empirical evidence as represented by the BPS estimates.

¹⁹For the general non-separable utility function no prior restrictions are imposed in estimation on the Frisch elasticities. The assumption of separability in the utility function translates into restrictions for the cross Frisch elasticities to be zero in the GMM estimation.

²⁰BPS also assume that transitory wage shocks have no wealth effect on consumption or labor, an assumption that is not exactly correct in life cycle models, especially not towards the end of the working life where permanent and transitory shocks become indistinguishable. Finally, a first-order approximation might not capture fully the nonlinear relationship between the wage shocks and endogenous consumption and labor income growth.

5 Results from the Life Cycle Model

As a benchmark we choose the additively separable preference specification, as it permits an intuitive interpretation of the estimation results and transparent comparison to the empirical findings of BPS. We present results for non-separable preferences in Section 5.5. Since the tightness of the borrowing constraints determines the degree of consumption insurance, but also the extent to which the assumption of interior allocations in the BPS approximations is violated, we explore both the case of tight borrowing constraints (the ZBC case) and non-binding borrowing constraints (the NBBC case). We recalibrate all economies to the same empirical targets to make them comparable to each other.

To generate results for each economy, optimal household policy functions are solved numerically using a policy function iteration algorithm combined with the endogenous grid method proposed by Carroll (2006). The policy functions are then used to simulate a panel of 50,000 households from age 21 to age 80 (although we only use observations from age 30 to 57 when comparing model implications to the BPS estimates, consistent with the data set they use). The results we report below are based on these simulated data.²¹

5.1 Life Cycles in the Benchmark Economy

Prior to turning to the main interest of the paper, the magnitude of insurance against wage shocks implied by the model, as measured by the transmission coefficients, we briefly want to document its basic life cycle properties. Figure 1 displays average consumption, asset, labor supply and income profiles, for both specifications of the borrowing constraints. With tight borrowing constraints (In the ZBC economy), average assets are initially zero, but rise fairly quickly as households accumulate away from the borrowing constraint due to their precautionary and life cycle savings motives. By age 30, only a very small share of households remain at or close to zero assets, an important observation for our assessment of the potential bias of the BPS estimates in Section 6. As is common in life cycle models, assets peak at the age of retirement (65 in the model) and are then decumulated to fund

²¹Details about the numerical method are provided in Appendix F.

retirement consumption. The consumption profile initially follows the rising labor income profile, and continues to grow since the product of gross interest rate and discount factor, $\frac{1+r}{1+\delta}$, is greater than one. Declining mean labor supply and rising mean labor income over the life cycle for both the male and female reflect the income effect of the rising wages (and thus incomes and consumption) over the life cycle.

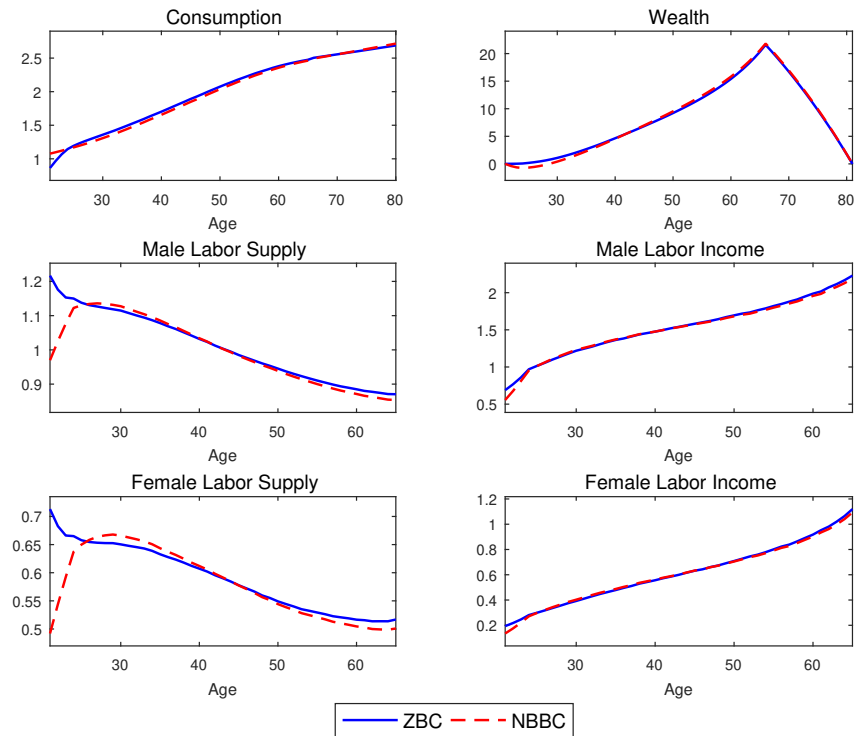


Figure 1: Life Cycles of Cross-sectional Means

Relaxing borrowing constraints (i.e. moving to the NBBC economy) affects average life cycle profiles only in the early stages of economic life. In the NBBC economy households borrow to attain higher consumption and be able to afford lower labor supply during their 20s, thereby attaining smoother consumption, hours and thus utility profiles. By age 30 average life cycle profiles in both economies essentially coincide.²²

²²In Appendix A we document that income and consumption *inequality* is rising over the life cycle as wage distributions fan out, that initially income inequality is smaller in the economy with

5.2 Transmission Coefficients in the Model and the Data

The main question we address in this paper is how household consumption and labor income respond to permanent and transitory wage shocks, as measured by the transmission coefficients κ in Equation (4.1). In the model simulation, realizations of transitory and permanent wage shocks are known separately, in addition to household consumption and labor income. Thus, in the model, the transmission coefficients can be obtained directly from OLS regressions of Equation (4.1). BPS must estimate these transmission coefficients from the data *without separate knowledge* of the different wage shocks; their method was developed precisely to tackle this problem. From now on, results from model-simulated data are labeled as “Model True”,²³ whereas results estimated from PSID data by BPS are labeled as “Data BPS”.

Before discussing the results of this comparison, we want to clarify why the model-based transmission coefficients may differ from the BPS empirical counterparts. First, the BPS methodology might deliver biased estimates of the true transmission coefficients in the data (for example, if potentially binding borrowing constraints are shaping the data). In Section 6 we verify that this is not a major concern, at least not when applied to model-simulated data. Second, our model might be a poor representation of the true data generating process. This is possible despite the fact that we calibrate the model to match the BPS estimates of the Frisch elasticities that then enter the formulas for the transmission coefficients such as Equation (4.2). The transmission coefficients also depend on the joint distribution of the wealth shares π_t and $s_{j,t}$ which BPS derive directly from the data, whereas they are the direct outcome of household saving and labor supply behavior in our model. Only if the model represents well the data along this dimension, will the transmission coefficients estimated by BPS and from our model line up closely. Our results below will therefore be informative about the question whether our model is a good summary of the true data generating process, at least for those aspects related to the

tight borrowing constraints (since wage-poor workers supply high hours to smooth consumption in the absence of borrowing capacities), and that life cycle inequality behaves very similar in both economies after age 30.

²³Due to the large sample size of simulated data, statistic errors are essentially zero. Hence the results from simulated data can be seen as the true values implied by the model.

consumption and labor income dynamics.

Table 2 reports the “Model True” transmission coefficients (both with tight and loose borrowing constraints), together with the “Data BPS” results.²⁴ Since BPS only use data of households aged 30 to 57, the “Model True” transmission coefficients are based on model-simulated data for this age group as well.

Table 2: Transmission Coefficients in the Data and the Model

	Data BPS	Model True	
		ZBC	NBBC
k_{c,u_1}	-0.15(0.06)	0.02	0.02
κ_{c,u_2}	-0.03(0.06)	0.01	0.01
κ_{c,v_1}	0.34(0.05)	0.43	0.43
κ_{c,v_2}	0.20(0.03)	0.23	0.24
κ_{y_1,u_1}	1.59(0.15)	1.51	1.51
κ_{y_1,u_2}	0.10(0.05)	-0.01	-0.01
κ_{y_1,v_1}	0.90(0.08)	1.13	1.13
κ_{y_1,v_2}	-0.23(0.03)	-0.21	-0.21
κ_{y_2,u_1}	0.21(0.11)	-0.02	-0.02
κ_{y_2,u_2}	1.87(0.22)	1.84	1.84
κ_{y_2,v_1}	-0.78(0.14)	-0.56	-0.56
κ_{y_2,v_2}	1.40(0.08)	1.41	1.42

Note: The numbers inside parentheses are standard errors from BPS.

Comparing the consumption transmission coefficients in the ZBC economy with the BPS estimates, we observe that the model implies almost perfect consumption insurance against transitory wage shocks. Only 2% of male and 1% of female temporary wage shocks pass through to household consumption. This result is common in life cycle models with self-insurance through saving and additively separable preferences, and is in line with BPS’ estimates of statistically insignificant or marginally significant but economically small consumption responses to temporary wage shocks of both household members. More importantly, our model predicts that only about 43% of male and 23% of female *permanent* wage shocks pass through to household consumption. The corresponding estimates from BPS are 34% for shocks to the male and 20% for shocks to the female wage. Thus the

²⁴As comparison we use the BPS results without taxes and without the separability assumption.

model-implied consumption insurance against permanent wage shocks is quantitatively close to the empirical estimates. Taking BPS' point estimates, the model can explain about 86% and 96% of the consumption insurance against permanent wage shocks.²⁵

The table also shows that the model not only fits well the consumption insurance patterns in the data, but also gives an accurate account of the empirical transmissions of wage shocks to labor income (and thus labor supply). In the ZBC economy with tight borrowing constraints, the responses of male and female labor income to their own transitory wage shocks, κ_{y_j, u_j} , are larger than one, indicating that labor supply increases when wages are temporarily high.²⁶ In the model, transitory wage shocks have only a small wealth effect, and thus the substitution effect dominates the labor supply response. The small wealth effect (and an absent substitution effect) also explains the slightly negative transmission coefficient $\kappa_{y_j, u_{-j}}$ to labor income from a wage shock of the spouse.

The labor income transmission coefficients to own permanent wage shocks, κ_{y_j, v_j} , are smaller than their transitory counterparts since labor supply responds less to own permanent wage shocks, on account of the stronger wealth effects. In contrast, when the spouse receives a permanent wage shock, labor supply of the other earner responds more strongly. The value of $\kappa_{y_1, v_2} = -0.21$ represents an 0.21% increase in male labor supply in response to a permanent 1% decline in the female wage, and $\kappa_{y_2, v_1} = -0.56$ implies a strong positive response of female hours to a 1% permanent reduction in male wages. These results suggest that the labor supply adjustment of spouses, and especially of females, is a crucial coping mechanism to reduce the adverse effect on household income of permanent wage shocks. Overall, female labor supply is more responsive to wage shocks than male labor supply. Crucially, comparing the transmission coefficients in the ZBC economy with their empirical counterparts, "Data BPS," the model is overall able to reproduce the main patterns in the data quite well, with more pronounced quantitative deviations observed only in the magnitude of the cross income response of one

²⁵Recall that one key finding of Kaplan and Violante (2010) was that there is substantially too little consumption insurance against permanent income shocks in Bewley-type models.

²⁶The labor supply response to wage shocks can be deduced by subtracting the percentage change of wages due to a specific shock from the associated transmission coefficients to labor income.

household member to wage shocks of the other member.

Finally, Table 2 also demonstrates that the transmission coefficients are not substantially affected by the tightness of the borrowing constraints, comparing the results in the ZBC and NBBC economies. This quantitative result is primarily due to the fact that, to facilitate the comparison with BPS' empirical results, only observations for households aged 30 to 57 are used in the calculation of the transmission coefficients, thus excluding young households mostly affected by binding borrowing constraints.²⁷

5.3 Age Profiles of Transmission Coefficients

The transmission coefficients in Table 2 are sample averages for all households aged 30 to 57. Figure 2 documents that there is very substantial age heterogeneity in the response of labor supply, income and consumption by household age. The transmission coefficients to consumption and labor income from transitory shocks, κ_{c,u_j} , κ_{y_j,u_j} , are relatively stable over the life cycle, due to the small wealth effects transitory shocks have. The transmission of permanent wage shocks displays much more significant variation over the life cycle. The transmission coefficients to consumption, κ_{c,v_j} , decreases with age, i.e., the amount of consumption insurance against permanent shocks rises over the life cycle. This is the result of increased asset accumulation and declining human wealth with age, so that permanent wage shocks become less important the older the household turns. In their data, BPS also find increasing consumption insurance with age against permanent male wage shocks, another dimension along which the model fits the data. The transmission coefficient to male labor income from his own permanent wage shocks, κ_{y_1,v_1} , increases with age since the substitution effect on labor from higher wages is offset less by a declining wealth effect over the life cycle. That is, young households increase their male labor supply less than old households in response to a positive male permanent wage shock. Most notably, whereas in young ages female labor

²⁷An additional important factor is that both economies are calibrated separately to fit the same empirical targets. If one would keep all parameters of the ZBC economy intact for the NBBC economy, the amount of consumption insurance would be affected, but so would be all other model-implied statistics, rendering the model counterfactual along many dimensions, most notably the aggregate wealth-income ratio.




Figure 2: Age Profiles of Transmission Coefficients

supply and thus female earnings respond quite strongly to an adverse permanent male wage shock ($\kappa_{y_2,v_1} = -0.7$), with age this adjustment becomes less important (κ_{y_2,v_1} declines in absolute value) since older households smooth consumption primarily using assets, and rely less on adjustments of labor supply.

5.4 How Do Households Insure Against Wage Shocks?

From the results in the previous subsection we know that the degree of consumption insurance especially against permanent wage shocks varies strongly by age, and that importance of labor supply adjustments in providing this consumption insurance display a clear life cycle pattern as well. In this section we provide a systematic decomposition of the importance of the consumption insurance mechanisms present in the model. In our economy, households can smooth wage shocks through four basic mechanisms, two of which are exogenous to the household, and two involve active decisions. First, taking labor supply as given, the fact that wage

shocks are imperfectly correlated among the two household members provides income and thus consumption insurance against individual wage fluctuations. Second, the social security system guarantees some income after retirement, which is independent of past wages and thus mitigates the impact of wage shocks on household lifetime income. The two key margins along which households can adjust behavior in response to wage shocks are the labor supply of both members,²⁸ as well as the accumulation (and decumulation) of assets.

Table 3 displays the transmission coefficient to consumption of permanent and transitory male wage shocks in a sequence of economies that differ in their availability of insurance mechanisms.²⁹ The table also breaks down, for each economy, how much of the consumption insurance is provided by labor supply and thus income adjustments of the male, the female earner as well as through assets and social security (combined). For example, the first line shows that in an economy with exogenous labor supply where the female household member does not work and the male member works full-time, a permanent male wage shock by construction transmits into male income (and thus household income) one for one (i.e. the transmission coefficient to male and total household income is 1). The transmission coefficient to consumption is 70%, and thus 30% of the permanent male wage shock is insured. By construction all of this consumption insurance stems from precautionary saving and from the redistributive social security system. The rows in the table then successively introduce endogenous male labor supply, positive (but exogenous) female labor supply and finally endogenous female labor supply. The bottom panel does the same, but for transitory male wage shocks. Note that in order to maximize comparability across economies we retain the same calibration across all models, implying that they differ in terms of asset accumulation over the life cycle.

²⁸This is true for the intensive and extensive margin. However, our experiments show that for the provision of consumption insurance in the model the extensive margin is of secondary importance. Furthermore, our decompositions below involve changes in log-earnings which is not well-defined for a household member that switches between participation and non-participation. Therefore, the results in this subsection stem from a version of the model that abstracts from the extensive margin of female labor supply.

²⁹We focus on the male wage shocks since they are quantitatively by far the most important one for household consumption, and because female wage shocks display the same qualitative results.

Table 3: Transmission of a 1% Male Wage Shock

Economy	% Change of				Insurance Provided by			Total Insurance
	Male Income	Female Income	Household Income	Household Consumption	Male Earner	Female Earner	Savings+ Social Security	
<i>A. Permanent Shock</i>								
(i) $H_{1,t} = 1, H_{2,t} = 0$	1	–	1	0.70	0	–	30%	30%
(ii) $H_{1,t}$ endogenous, $H_{2,t} = 0$	0.97	–	0.97	0.61	3%	–	36%	39%
(iii) $H_{1,t}$ endogenous, $H_{2,t} = \bar{H}_2$	1.07	0	0.81	0.50	–7%	26%	31%	50%
(iv) $H_{1,t}, H_{2,t}$ both endogenous	1.13	–0.64	0.62	0.43	–13%	51%	19%	57%
<i>B. Transitory Shock</i>								
(i) $H_{1,t} = 1, H_{2,t} = 0$	1	–	1	0.03	0	–	97%	97%
(ii) $H_{1,t}$ endogenous, $H_{2,t} = 0$	1.51	–	1.51	0.02	–51%	–	149%	98%
(iii) $H_{1,t}$ endogenous, $H_{2,t} = \bar{H}_2$	1.51	0	1.14	0.02	–51%	37%	112%	98%
(iv) $H_{1,t}, H_{2,t}$ both endogenous	1.51	–0.02	1.07	0.02	–51%	44%	105%	98%

Note: The results are based on sample averages of transmission coefficients from different economies with the same parameters as in the ZBC benchmark economy, but with different labor supply choice sets. \bar{H}_2 is set equal to the average female labor supply in economy (iv).

Our discussion will focus mostly on the permanent wage shocks, since these are more important determinants of household welfare, are harder to insure, and it is with respect to these shocks that Kaplan and Violante (2010) found the most significant deviations between theory (i.e. a standard Bewley model with exogenous earnings) and data (i.e. the estimates of Blundell, Pistaferri, and Preston (2008)). The key result is that while households obtain some consumption insurance through precautionary asset accumulation, the presence of a secondary earner that can (and optimally does) adjust labor supply in response to permanent wage shocks of the primary earner is the main operative insurance mechanism, and implies that in the benchmark model only 43% of the wage shock transmits to household consumption (i.e. 57% are insured), see the forth row in Table 3, which collects the results for the benchmark economy.

It is instructive to arrive at this result in a sequence of steps. Economy (i) in Table 3 corresponds to the standard incomplete-markets model with exogenous income shocks similar Kaplan and Violante (2010), with all insurance provided by savings and social security. As stated above, this insurance mechanism alone leads to 30% of the permanent wage shock being insured, and 70% being transmitted into consumption.³⁰

³⁰This extent of consumption insurance is slightly higher than documented in Kaplan and Violante (2010) (30% v.s. their 23%). Since we do not recalibrate, economy (i) has a higher asset to income ratio than the one in Kaplan and Violante (2010) (and a higher one than our economy (iv) as well as the data), implying better consumption insurance against permanent income shocks. In addition,

The second row demonstrates that the adjustment of labor supply by the primary earner in response to *permanent* shocks to his own wages is not an effective tool for providing income and thus consumption insurance: 97% of wage fluctuations transmit into earnings fluctuations. Consequently the overall extent of consumption insurance is only moderately affected (the transmission falls from 70% to 61%), and asset accumulation (and social security) remain the predominant mechanisms providing insurance. The third row, economy (iii), shows that the presence of a secondary earner with imperfectly correlated wages, even if this secondary earner does not adjust labor supply in response to shocks to the primary earners' wage, is a significant source of household earnings insurance, as the transmission of the wage shock to *household earnings* falls from 97% to 81%. This is true despite the fact that now, in response to a permanent negative wage shock the primary earner reduces hours, so that the impact of the shock on male earnings exceeds 100%. Only 50% of the wage shock now spills over into household consumption, and the decomposition analysis makes clear that female income insurance of the male wage shock is quantitatively almost as important as asset accumulation for this result.

Finally, in the benchmark economy (iv), when the labor supply of the female earner can adjust to a permanent male wage shock, household income insurance and consumption insurance rises further. Only 62% of the male wage shock transmits to household earnings (despite the fact that male earnings move more than one for one with the wage shock), and only 43% of the permanent wage shock is reflected in consumption, for a consumption insurance coefficient of 57%. Of this consumption insurance, $\frac{51\%}{57\%} \approx 89\%$ is provided by the female earner, both through exogenously imperfectly correlated wages as well as through the endogenous adjustment of hours worked. In the presence of spousal earnings insurance the importance of precautionary saving for consumption insurance diminishes noticeably (providing only 19% consumption insurance), but remains non-negligible. These results demonstrate the quantitatively crucial role that earnings of the second worker in the household plays for the insurance of consumption against permanent

since our empirical comparison is Blundell, Pistaferri, and Saporta-Eksten (2016), we only use households aged 30 to 57, whereas the results in Kaplan and Violante (2010) also include very young households with low assets whose consumption responds especially strongly to permanent earnings shocks.

wage shocks.

Table 3 also shows that, perhaps not surprisingly, household consumption is almost perfectly insured against transitory wage shocks, fairly independently of whether there is a secondary earner in the household and whether this earner can adjust hours worked or not. Furthermore, now the endogenous hours response of the male labor supply amplifies the transitory wage shocks, and asset accumulation and social security are the crucial factors in providing consumption insurance against these temporary wage shocks.

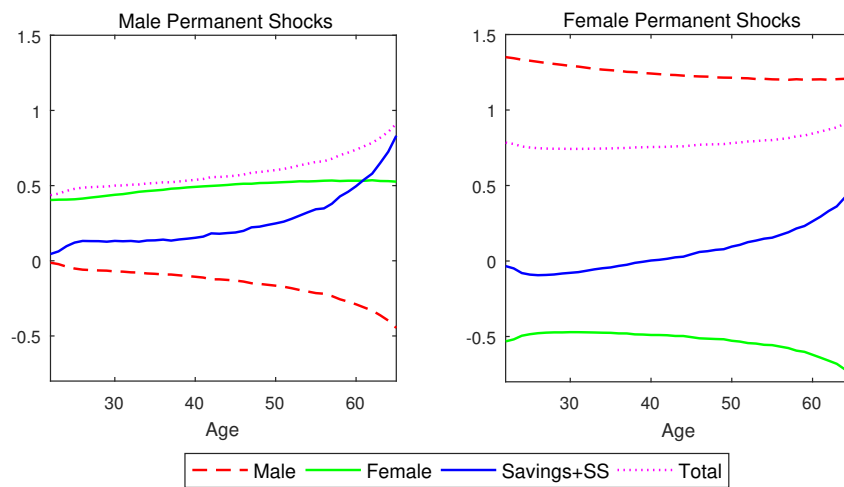


Figure 3: Age Profiles of Consumption Insurance Decomposition

To briefly investigate whether the contributions of the insurance mechanisms we have highlighted in the previous table vary over the life cycle, in Figure 3 we plot the consumption insurance contributions for a permanent wage shock against household age. The left panel pertains to a permanent shock to male wages, whereas the right panel displays the results for a permanent female wage shock. We observe that, for all ages, the labor supply response of the person whose wage is hit by a permanent shock exacerbates the shock (i.e. the insurance contribution is negative, increasingly so as the household ages and the substitution effect increasingly dominates the wealth effect). Consistent with the main theme of Table 3, the labor supply response of the other spouse is the most potent insurance mechanism, and is roughly constant over time, whereas the insurance provided by savings and social

security keeps rising in importance over the life cycle, indicating that older households rely more on self-insurance through savings. This is also the main reason for the increasing age profile of total consumption insurance.

5.5 Non-separable Preferences

The advantage of additively separable preferences between consumption and labor is that there is a direct mapping between preference parameters and Frisch elasticities, so that the BPS parameter estimates are readily interpretable, and can be used in our structural model directly. In addition, Table 2 demonstrates that this specification generates consumption and labor income responses to wage shocks in the model close to their empirical counterparts. However, when BPS relax the assumption of separability in their estimation, they find important Frisch complementarity between consumption and leisure ($\eta_{c,w_j} < 0$, $\eta_{h_j,p} > 0$), and also between the leisure of both spouses ($\eta_{h_j,w_{-j}} > 0$). In addition, introducing non-separable preferences into the model could, in principle, help it to better capture additional aspects of the data the benchmark model thus far struggles with, for example, the negative consumption response to positive transitory wage shocks ($\kappa_{c,u_j} < 0$) in the BPS results.

We therefore now adopt a non-separable preference specification as outlined in Equation (2.1), and calibrate it as described in Section 3.9. This specification is computationally tractable and at the same time can capture the Frisch complementarity between consumption and leisure, and between leisure of both spouses, implied by the BPS estimation results. The purpose of our analysis here is to learn how such non-separability affects the amount of model-implied consumption insurance (and the other transmission coefficients), and to what extent it helps the model match the data better. Table 4 summarizes the most relevant results; the complete set of findings complementing the separable results from the previous sections are available in Appendix B.

Relative to separable utility, now consumption and leisure are Frisch complements, and thus higher labor supply reduces the marginal utility of consumption, and higher consumption increases the marginal disutility of labor supply. The main

Table 4: Transmission Coefficients in the Data and the Model
(Non-separable Preferences)

	Data BPS	Model True	
		ZBC	NBBC
k_{c,u_1}	-0.15(0.06)	-0.20	-0.21
κ_{c,u_2}	-0.03(0.06)	-0.09	-0.09
κ_{c,v_1}	0.34(0.05)	0.29	0.29
κ_{c,v_2}	0.20(0.03)	0.16	0.16
κ_{y_1,u_1}	1.59(0.15)	1.75	1.77
κ_{y_1,u_2}	0.10(0.05)	0.05	0.07
κ_{y_1,v_1}	0.90(0.08)	0.94	0.94
κ_{y_1,v_2}	-0.23(0.03)	-0.35	-0.34
κ_{y_2,u_1}	0.21(0.11)	0.11	0.15
κ_{y_2,u_2}	1.87(0.22)	1.69	1.70
κ_{y_2,v_1}	-0.78(0.14)	-0.60	-0.59
κ_{y_2,v_2}	1.40(0.08)	1.19	1.18

Note: The numbers inside parentheses are standard errors from BPS.

implied difference for consumption insurance is that now the consumption response to transitory wage shocks (of both spouses) is negative in the model, as in the data. Transitory wage shocks only have a small wealth effect, and thus the substitution effect induces labor supply to increase with the transitory wage increase, thus reduces the marginal utility of consumption and thus the household optimally cuts consumption in response to a temporary decline in wages. The same logic applies to a permanent wage shock, and hence the consumption response to such a shock is muted (relative to the separable case) by the consumption leisure complementarity, but of course remains positive. Compared to the “Data BPS” results, the model now matches better the more than 100% consumption insurance against transitory shocks in the data, and generates slightly *higher* consumption insurance against permanent wage shocks than the data. The pass-through rates of permanent wage shocks to consumption in the model are only 29% and 16% for male and female permanent wage shocks, respectively, relative to 34% and 20% as estimated by BPS. Our main conclusion from the previous section thus remains intact, in fact is strengthened: the Bewley model with endogenous dual earner labor supply implies roughly as much insurance against wage shocks as the data appear to exhibit.

Finally, nonseparable preferences also alter, qualitatively, the labor supply and thus earnings response to wage shocks. Specifically, the labor income responses to transitory wage shocks of the *other* earner are now positive, in contrast to the additively separable case, and consistent with the BPS estimates. Overall, we conclude that the transmission coefficients to labor income (and indeed consumption) from wage shocks in the model replicate well the empirical evidence.³¹

6 Performance of the BPS Method

Section 5 argues that the Bewley type model with two-earner households and endogenous labor supply is an empirically successful model of household consumption and labor supply behavior in the presence of wage shocks, in the sense that responses to these shocks of household consumption and labor income in the model match the empirical results from BPS well. This conclusion would not be valid if the BPS estimates were significantly biased. To assess this possibility, in this section we apply the BPS methodology to data simulated from the model. Since we control the data generating process, we know the true Frisch elasticities and transmission coefficients, which we label as “Model True” in the tables below. The estimates based on the BPS methodology are denoted as “Model BPS”.³² Following the BPS baseline results, we do not estimate the “outside insurance” coefficient β , and set it to the true value of zero, unless specified otherwise.³³

To summarize our findings, on model-simulated data the BPS method works well in the estimation of the Frisch elasticities and transmission coefficients, with biases that are in general modest, both for additively separable and non-separable

³¹The transmission coefficients in the NBBC economy are again very close to those of the ZBC economy, for the same reasons already discussed in the additively separable preference case.

³²Because the sample size of simulated data is very large (and can be made arbitrarily large), the standard errors of the “Model BPS” results are practically zero, and hence the differences between the “Model True” and “Model BPS” results must be due to the potential bias of the BPS method, and tell us how well the BPS method identifies these parameters.

³³Although the model includes a social security system with retirement benefit b , the effect of this public insurance is accounted for explicitly in the modified BPS formulas of this paper by q_t . Hence the true value of the “outside insurance” coefficient β in the model is zero. The relevant details are contained in Appendix E, where we also supply the specifics how we implement the BPS methodology on our simulated data.

preferences. These results give us some confidence that the key finding from Section 5 that our Bewley model with dual labor supply captures well the *empirically observed* responses to wage shocks, is not contaminated by the fact that BPS estimate these transmission coefficients with bias.³⁴

6.1 With Additively Separable Preferences

Tables 5 to Table 7 summarize the results from applying the BPS estimation methodology to the benchmark economy with the additively separable preferences.

Table 5: Estimation of Wage Process

	Model True	Model BPS
<i>A. Transitory Shocks</i>		
$\sigma_{u_1}^2$	0.0275	0.0272
$\sigma_{u_2}^2$	0.0125	0.0124
σ_{u_1, u_2}	0.0058	0.0055
<i>B. Permanent Shocks</i>		
$\sigma_{v_1}^2$	0.0303	0.0306
$\sigma_{v_2}^2$	0.0382	0.0365
σ_{v_1, v_2}	0.0027	0.0037

Table 5 reports estimates for the parameters governing the exogenous wage process. We observe that the variances of male shocks are estimated very accurately; the variance of the permanent female wage shocks and the covariance with the male wage shock display slightly larger (but still moderate) biases as a result of the sample selection due to the extensive margin of female labor supply in the model.

The main interest of the paper lies with the transmission coefficients, but as an intermediate step, the Frisch elasticities need to be estimated first. Table 6 contains the relevant results. It shows that the BPS method identifies the Frisch elasticities very precisely, with the possible exception of the female labor supply elasticity to her own wages, η_{h_2, w_2} , which displays a somewhat larger bias due to the extensive margin of female labor supply.³⁵

³⁴Of course, the fact that the BPS method performs well on simulated data does not guarantee that it performs equally well when applied to actual data.

³⁵Correctly imposing the separability assumption on the preferences does not impact the estima-

Table 6: Estimation of Frisch Elasticities and Outside Insurance

	Model True	ZBC			NBBC
		Model BPS			Model
		Baseline ^a	Separable ^b	Outside ^c	BPS ^a
<i>A. Frisch Elasticities</i>					
$\eta_{c,p}$	0.578	0.579	0.615	0.593	0.578
η_{c,w_1}	0	0.004	0	-0.002	0.004
η_{c,w_2}	0	0.021	0	0.022	0.020
$\eta_{h_1,p}$	0	-0.005	0	0.003	-0.005
η_{h_1,w_1}	0.528	0.528	0.520	0.535	0.527
η_{h_1,w_2}	0	-0.002	0	0.002	-0.001
$\eta_{h_2,p}$	0	-0.049	0	-0.052	-0.047
η_{h_2,w_1}	0	-0.004	0	0.003	-0.002
η_{h_2,w_2}	0.850	0.802	0.812	0.805	0.817
<i>B. Outside Insurance</i>					
β	0	-	-	-0.107	-

Note: (a) the baseline BPS method; (b) BPS method assuming separable preferences; (c) BPS method allowing for outside insurance.

Finally, Table 7 reports the results for the transmission coefficients.³⁶ For transitory shocks, the BPS method slightly overestimates the consumption insurance against male shocks, and underestimates it against female shocks. For permanent shocks, it tends to underestimate both the insurance against male and female shocks. However, the main message of the table is that the *magnitude* of the bias in all cases is quite small.

The bias for the wage transmission coefficients to labor income are roughly on the same magnitude as for the consumption insurance coefficients, with the exception of the transmission coefficients to female labor income of male and female permanent wage shocks, κ_{y_2,v_j} . The BPS method overestimates these female labor income responses, likely due to the selection problem caused by the extensive

tion results much. The estimation of the “outside insurance” coefficient is slightly negatively biased, with a magnitude of 10.7%, and the tightness of borrowing constraints does not affect the bias noticeable, since the sample only includes simulated households aged 30 to 57, and young households who are mostly affected by the tight borrowing constraints are not included in the estimation.

³⁶Figure 10 in Appendix E.3 summarizes the true and estimated transmission coefficients to household consumption over the life cycle, and confirms that the BPS method works well fairly uniformly by age of the household.

Table 7: Estimation of Transmission Coefficients

	ZBC				NBBC	
	Model	Model BPS			Model	Model
	True	Baseline ^a	Separable ^b	Outside ^c	True	BPS ^a
κ_{c,u_1}	0.016	0.004	0	-0.002	0.016	0.004
κ_{c,u_2}	0.008	0.021	0	0.022	0.008	0.020
κ_{c,v_1}	0.432	0.445	0.456	0.472	0.432	0.442
κ_{c,v_2}	0.235	0.237	0.238	0.254	0.235	0.238
κ_{y_1,u_1}	1.513	1.528	1.520	1.535	1.513	1.527
κ_{y_1,u_2}	-0.007	-0.002	0	0.002	-0.007	-0.001
κ_{y_1,v_1}	1.133	1.114	1.135	1.089	1.134	1.115
κ_{y_1,v_2}	-0.215	-0.206	-0.201	-0.217	-0.215	-0.206
κ_{y_2,u_1}	-0.020	-0.004	0	0.003	-0.019	-0.002
κ_{y_2,u_2}	1.836	1.802	1.812	1.805	1.840	1.817
κ_{y_2,v_1}	-0.556	-0.599	-0.601	-0.622	-0.559	-0.610
κ_{y_2,v_2}	1.413	1.509	1.498	1.499	1.416	1.515

Note: (a) the baseline BPS method; (b) BPS method assuming separable preferences; (c) BPS method allowing for outside insurance.

margin of female labor supply.

6.2 With Non-separable Preferences

Finally, we assess whether the good performance of the BPS method (with respect to simulated data) extends to the model with preferences that are non-separable between consumption and hours worked. Table 8 presents the results for the estimation of the Frisch elasticities.³⁷ The patterns and magnitudes of biases are similar to those with additively separable preferences, and the “outside insurance” coefficient is still negatively biased by about 10%.

Table 9 reports the results for transmission coefficients. The major difference from the additively separable case is that now the BPS method tends to more significantly overestimate consumption insurance against female permanent wage shocks. However, overall, we conclude that the BPS methodology, when applied to simulated data, measures well the degree of true consumption insurance. To the extent

³⁷The “Model True” Frisch elasticities are no longer deep parameters under non-separable preferences, but rather depend on allocations. The reported results are sample averages of household level elasticities calculated using formulas in Appendix D.

Table 8: Estimation of Frisch Elasticities and Outside Insurance
(Non-separable Preferences)

	ZBC			NBBC	
	Model	Model BPS		Model	Model
	True	Baseline ^a	Outside ^b	True	BPS ^a
<i>A. Frisch Elasticities</i>					
$\eta_{c,p}$	0.371	0.351	0.366	0.370	0.351
η_{c,w_1}	-0.193	-0.210	-0.213	-0.198	-0.217
η_{c,w_2}	-0.114	-0.093	-0.090	-0.117	-0.098
$\eta_{h_1,p}$	0.285	0.256	0.260	0.280	0.254
η_{h_1,w_1}	0.764	0.783	0.786	0.779	0.801
η_{h_1,w_2}	0.080	0.061	0.062	0.100	0.082
$\eta_{h_2,p}$	0.285	0.229	0.224	0.280	0.232
η_{h_2,w_1}	0.139	0.124	0.126	0.173	0.166
η_{h_2,w_2}	0.705	0.645	0.645	0.706	0.663
<i>B. Outside Insurance</i>					
β	0	-	-0.100	0	-

Note: (a) the baseline BPS method; (b) BPS method allowing for outside insurance.

that the same is true for real data, the fact that the Bewley model with dual labor supply replicates well the *empirically observed* transmission coefficients then follows from our results in Section 5.

Table 9: Estimation of Transmission Coefficients
(Non-separable Preferences)

	ZBC			NBBC	
	Model	Model BPS		Model	Model
	True	Baseline ^a	Outside ^b	True	BPS ^a
k_{c,u_1}	-0.199	-0.210	-0.213	-0.207	-0.217
κ_{c,u_2}	-0.086	-0.093	-0.090	-0.090	-0.098
κ_{c,v_1}	0.288	0.307	0.326	0.287	0.305
κ_{c,v_2}	0.156	0.133	0.142	0.155	0.133
κ_{y_1,u_1}	1.753	1.783	1.786	1.773	1.801
κ_{y_1,u_2}	0.052	0.061	0.062	0.070	0.082
κ_{y_1,v_1}	0.943	0.912	0.894	0.941	0.909
κ_{y_1,v_2}	-0.349	-0.318	-0.323	-0.343	-0.312
κ_{y_2,u_1}	0.109	0.124	0.126	0.148	0.166
κ_{y_2,u_2}	1.692	1.645	1.645	1.700	1.663
κ_{y_2,v_1}	-0.604	-0.667	-0.675	-0.593	-0.666
κ_{y_2,v_2}	1.188	1.301	1.299	1.179	1.296

Note: (a) the baseline BPS method; (b) BPS method allowing for outside insurance.

7 Conclusions

In this paper we have argued that a Bewley type model with two-earner households facing idiosyncratic wage shocks and making endogenous labor supply decisions can generate the amount of consumption insurance consistent with BPS' estimates from US data. The successful performance of BPS' estimation methodology on data simulated from such a model suggests that the model fits the actual data well, at least for the U.S.³⁸

Overall, the results in this paper suggest that life cycle models of the form employed in this work can be used in applications where the extent of consumption insurance, and the mechanisms through which it is achieved, are important. This includes the evaluation of social insurance and tax policies for which the adjustment of labor supply of both household members as well as savings responses can be expected to be important. We view the use of the model exhibited in this paper for these questions as important avenues for future work.

³⁸Verifying this statement for data from other countries would of course be an important step for future research.

References

- Aiyagari, S. Rao. 1994. “Uninsured Idiosyncratic Risk and Aggregate Saving”. *The Quarterly Journal of Economics* 109 (3): 659–684.
- Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2005. “Female Labor Supply as Insurance Against Idiosyncratic Risk”. *Journal of the European Economic Association* 3 (2/3): 755–764.
- Bewley, Truman. 1986. “Stationary monetary equilibrium with a continuum of independently fluctuating consumers”. *Contributions to Mathematical Economics in Honor of Gerard Debreu*: 79–102.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. “Consumption Inequality and Partial Insurance”. *American Economic Review* 98 (5): 1887–1921.
- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. “Consumption Inequality and Family Labor Supply”. *American Economic Review* 106 (2): 387–435.
- Carroll, Christopher D. 2006. “The method of endogenous gridpoints for solving dynamic stochastic optimization problems”. *Economics Letters* 91 (3): 312 – 320.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. “Consumption and Labor Supply with Partial Insurance: An Analytical Framework”. *American Economic Review* 104 (7): 2075–2126.
- . 2010. “The Macroeconomic Implications of Rising Wage Inequality in the United States”. *Journal of Political Economy* 118 (4): 681–722.
- Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk. 2018. *How Does Tax Progressivity and Household Heterogeneity Affect Laffer Curves?* Working Paper.
- Huggett, Mark. 1993. “The risk-free rate in heterogeneous-agent incomplete-insurance economies”. *Journal of Economic Dynamics and Control* 17 (5-6): 953 –969.
- Hurst, Erik. 2008. “Understanding Consumption in Retirement: Recent Developments.” In *Recalibrating Retirement Spending and Saving*, ed. by John Ameriks and Olivia S. Mitchell. Oxford University Press.

- Imrohoroğlu, Ayşe. 1989. “Cost of Business Cycles with Indivisibilities and Liquidity Constraints”. *Journal of Political Economy* 97 (6): 1364–1383.
- Kaplan, Greg, and Giovanni L. Violante. 2010. “How Much Consumption Insurance beyond Self-Insurance?” *American Economic Journal: Macroeconomics* 2 (4): 53–87.
- Ortigueira, Salvador, and Nawid Siassi. 2013. “How important is intra-household risk sharing for savings and labor supply?” *Journal of Monetary Economics* 60 (6): 650–666.
- Rupert, Peter, and Giulio Zanella. 2015. “Revisiting wage, earnings, and hours profiles”. *Journal of Monetary Economics* 72:114–130.
- Santaeulàlia-Llopis, Raül, and Yu Zheng. 2017. *The Price of Growth: Consumption Insurance in China 1989-2009*. Working Paper.
- Tauchen, George. 1986. “Finite state markov-chain approximations to univariate and vector autoregressions”. *Economics Letters* 20 (2): 177–181.

Online Appendix: For Review Only

A Additional Model Life Cycle Findings

In this appendix we display life cycle profiles of economic inequality. Specifically, Figure 4 shows the life cycle of income and consumption inequality in both the economy with tight and with loose borrowing constraints, where inequality is measured by the cross-sectional variance of log household income and log household consumption. As expected, both income and consumption inequality rise with age due to the accumulation of the permanent wage shocks during working life. The rise in consumption inequality slows down for working households near retirement, whereas the rise of household income inequality accelerates during the same period. Consumption inequality is constant after retirement because there is no labor income risk after that time in the model (and we omit it from the graph), and consumption levels are simply determined by the total discounted wealth of households at retirement. Consistent with the data, the consumption inequality in the model is smaller than income inequality, reflecting the consumption insurance channels in the model. Income inequality in the NBBC economy is notably higher than that of the ZBC economy in the early years of the life cycle. This is due to the fact that in the ZBC economy, young households hit by transitory negative wage shocks have to work longer hours to smooth consumption, whereas in the NBBC economy these households can simply borrow and avoid to increase their labor supply, which results in lower income for them than in the ZBC economy. This difference vanishes quickly as households accumulate enough assets through saving, and hence the income inequality patterns are almost identical for older households in both economies.

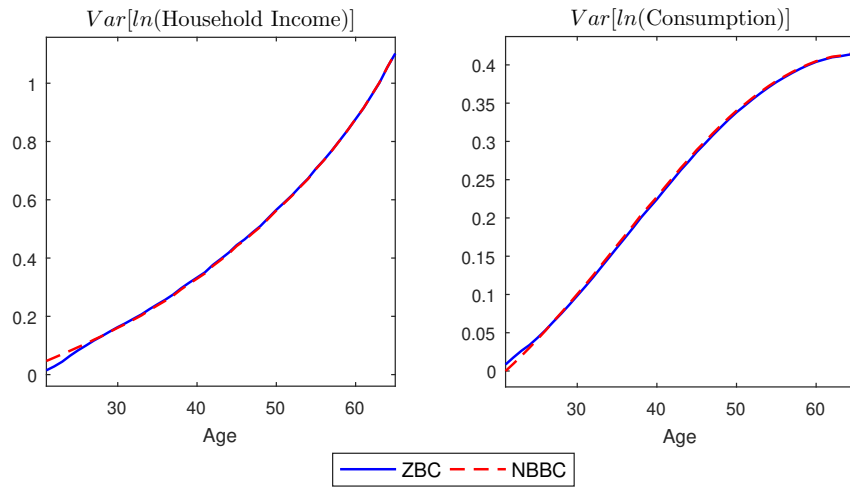


Figure 4: Life Cycles of Income and Consumption Inequality

B Supplementary Results For Non-separable Preferences

This section presents all the figures and tables for the ZBC and NBBC economies with the non-separable preferences that are not included in the main text. Each of them corresponds to one figure or table for the additively separable preferences.

B.1 Life Cycles

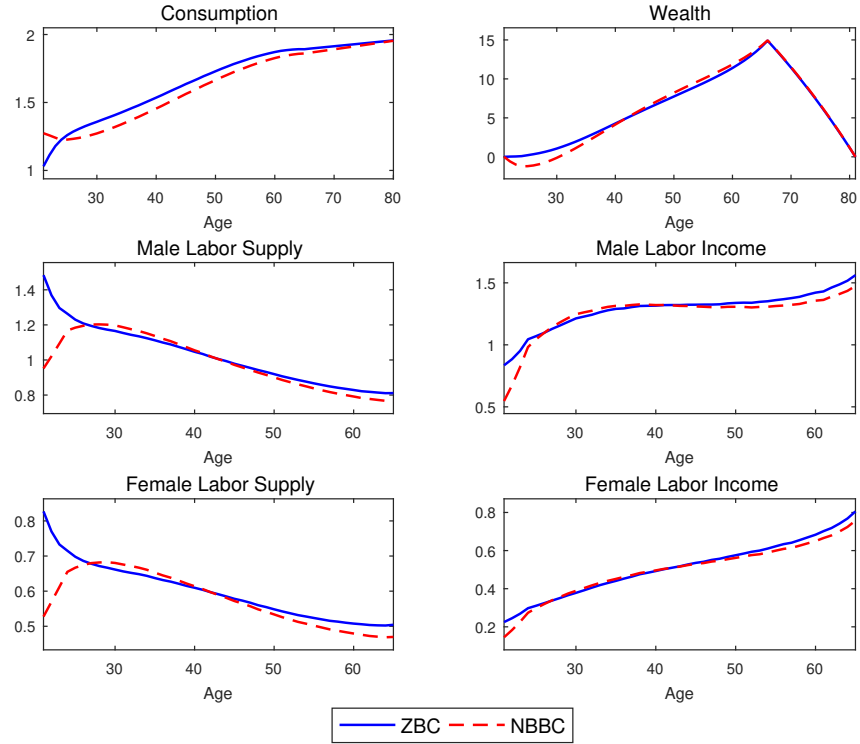


Figure 5: Life Cycles of Cross-sectional Means
(Non-separable Preferences)

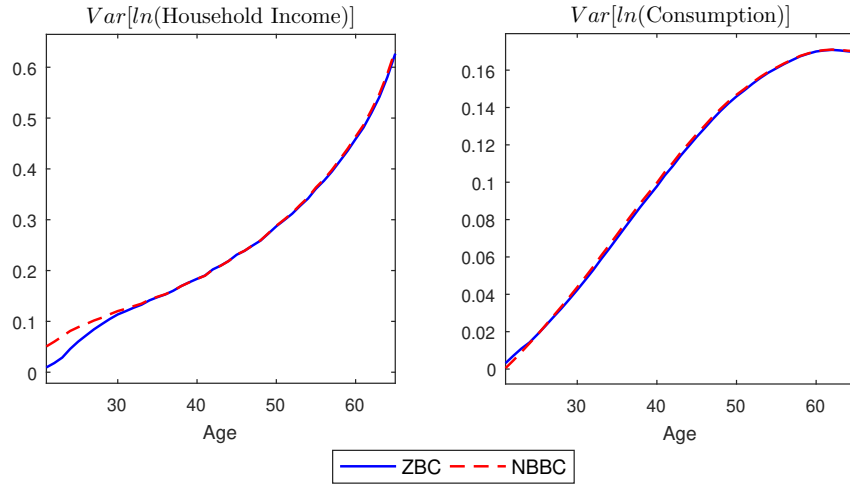


Figure 6: Life Cycles of Income and Consumption Inequality
(Non-separable Preferences)




Figure 7: Age Profiles of Transmission Coefficients
(Non-separable Preferences)

B.2 Consumption Insurance Decomposition

Table 10: Transmission of a 1% Male Wage Shock
(Non-separable Preferences)

Economy	% Change of				Insurance Provided by			Total Insurance
	Male Income	Female Income	Household Income	Household Consumption	Male Earner	Female Earner	Savings+ Social Security	
<i>A. Permanent Shock</i>								
(i) $H_{1,t} = 1, H_{2,t} = 0$	1	–	1	0.63	0	–	37%	37%
(ii) $H_{1,t}$ endogenous, $H_{2,t} = 0$	0.69	–	0.69	0.38	31%	–	31%	62%
(iii) $H_{1,t}$ endogenous, $H_{2,t} = \bar{H}_2$	0.79	0	0.56	0.31	21%	23%	25%	69%
(iv) $H_{1,t}, H_{2,t}$ both endogenous	0.94	–0.68	0.47	0.29	6%	47%	18%	71%
<i>B. Transitory Shock</i>								
(i) $H_{1,t} = 1, H_{2,t} = 0$	1	–	1	0.03	0	–	97%	97%
(ii) $H_{1,t}$ endogenous, $H_{2,t} = 0$	1.78	–	1.78	–0.27	–78%	–	205%	127%
(iii) $H_{1,t}$ endogenous, $H_{2,t} = \bar{H}_2$	1.72	0	1.23	–0.17	–72%	49%	140%	117%
(iv) $H_{1,t}, H_{2,t}$ both endogenous	1.75	0.13	1.28	–0.20	–75%	47%	148%	120%

Note: The results are based on sample averages of transmission coefficients from different economies with the same parameters as the ZBC economy with non-separable preferences but different household structures. \bar{H}_2 is set equal to the average female labor supply in economy (iv).



Figure 8: Age Profiles of Consumption Insurance Decomposition
(Non-separable Preferences)

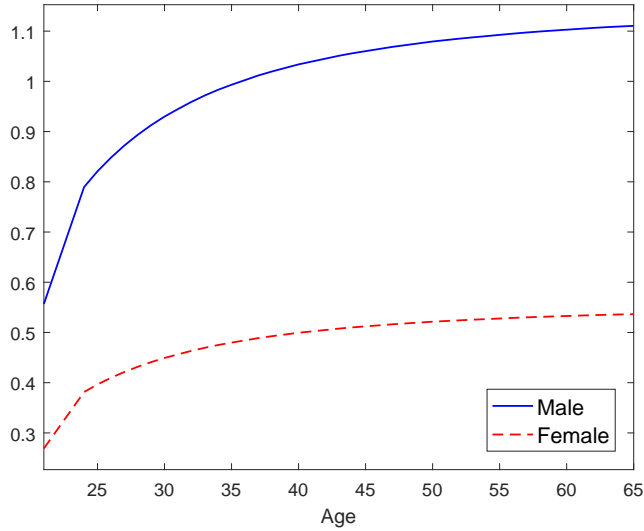


Figure 9: Male and Female Wage Trends

C Calibration Details

C.1 Wage Profiles

The life-cycle male wage trend is interpolated and extrapolated from Rupert and Zanella (2015) and plotted in Figure 9. The scale of it is normalized such that the average male trend wage is 1. The female wage trend is rescaled from the male wage trend to match the ratio of the average earnings between working males and females in the BPS data set.

C.2 Parameter Values and Moments Matched

The parameter values which are the same for all the economies are reported in Table 11. The preference-specific parameter values are reported in Table 1 and 12 for the additively separable and non-separable preferences, respectively. The moments matched in calibration and their values from data are in Table 14. The calibrated Frisch elasticities and their BPS counterparts for the non-separable preferences are in Table 13.

Table 11: Values of Common Parameters

Parameter	Governing	Value
r	interest rate	0.02
<i>Transitory Wage Shocks</i>		
$\sigma_{u_1}^2$	variance of male shocks	0.0275
$\sigma_{u_2}^2$	variance of female shocks	0.0125
σ_{u_1, u_2}	covariance of male and female shocks	0.0058
<i>Permanent Wage Shocks</i>		
$\sigma_{v_1}^2$	variance of male shocks	0.0303
$\sigma_{v_2}^2$	variance of female shocks	0.0382
σ_{v_1, v_2}	covariance of male and female shocks	0.0027

Table 12: Parameter Values for Non-separable Preferences

Parameter	Governing	ZBC	NBBC
δ	discount rate of utility	0.0146	0.0121
α	weight of consumption	0.689	0.625
ξ	weight of male labor supply	0.445	0.438
b	retirement benefits	0.765	0.745
σ	consumption Frisch elasticity	2.380	2.380
γ	substitution between consumption and leisure	-4.00	-3.90
θ	substitution between male and female labor supply	2.60	2.65
Ψ	consumption level after retirement	0.597	0.610
f	fixed utility cost of female labor participation	0.0096	0.0107
$e^{g_{1,t} - g_{2,t}}$	male/female wage trend ratio	2.06	2.05

Table 13: Calibrated Frisch Elasticities with Non-separable Preferences

	Data BPS	ZBC	NBBC
		Model True	Model True
$\eta_{c,p}$	0.372(0.086)	0.371	0.370
η_{c,w_1}	-0.148(0.060)	-0.193	-0.198
η_{c,w_2}	-0.030(0.059)	-0.114	-0.117
$\eta_{h_1,p}$	0.085(0.035)	0.285	0.280
η_{h_1,w_1}	0.594(0.155)	0.764	0.779
η_{h_1,w_2}	0.104(0.053)	0.080	0.100
$\eta_{h_2,p}$	0.035(0.069)	0.285	0.280
η_{h_2,w_1}	0.212(0.108)	0.139	0.173
η_{h_2,w_2}	0.871(0.221)	0.705	0.706

Table 14: Values of Matched Moments

Moment Matched	Target Value
wealth/income ratio	2.999
average male labor supply	1
average female labor supply work	0.733
average benefits/average income ratio	0.425
average female nonparticipation rate	20%
average male earnings/average female earnings work	2.031

D Frisch Elasticities for Non-separable Preferences

In this section, we derive the formulas for the Frisch elasticities with the non-separable preferences. They are generally functions of household allocations, i.e., they are not deep parameters. So we use the sample averages of them as the approximated true values of the Frisch elasticities, i.e., the “Model True” results.

The utility function for the non-separable preferences is

$$u(C, H_1, H_2) = \frac{\{\alpha C^\gamma + (1 - \alpha)[\xi H_1^\theta + (1 - \xi)H_2^\theta]^{-\frac{\gamma}{\theta}}\}^{\frac{1-\sigma}{\gamma}} - 1}{1 - \sigma}.$$

The intertemporal budget constraint is

$$PC + PA' = P(1 + r)A + W_1H_1 + W_2H_2,$$

where P , W_1 , and W_2 are the price of the consumption good and the wages for male and female earners. From the recursive formulation of the household problem, the first-order conditions are

$$\begin{aligned} u_C &= \Delta^{\frac{1-\sigma}{\gamma}-1} \alpha C^{\gamma-1} = \lambda P, \\ u_{H_1} &= -\Delta^{\frac{1-\sigma}{\gamma}-1} (1 - \alpha) \Gamma^{-\frac{\gamma}{\theta}-1} \xi H_1^{\theta-1} = -\lambda W_1, \\ u_{H_2} &= -\Delta^{\frac{1-\sigma}{\gamma}-1} (1 - \alpha) \Gamma^{-\frac{\gamma}{\theta}-1} (1 - \xi) H_2^{\theta-1} = -\lambda W_2, \end{aligned}$$

where λ is the lagrangian multiplier on the budget constraint, $\Delta \equiv \alpha C^\gamma + (1 - \alpha)[\xi H_1^\theta + (1 - \xi)H_2^\theta]^{-\frac{\gamma}{\theta}}$, and $\Gamma \equiv \xi H_1^\theta + (1 - \xi)H_2^\theta$. Taking log difference for both sides of these equations, we get

$$\begin{aligned} \left(\frac{1-\sigma}{\gamma} - 1\right)d \ln \Delta + (\gamma - 1)d \ln C &= d \ln \lambda + d \ln P, \\ \left(\frac{1-\sigma}{\gamma} - 1\right)d \ln \Delta + \left(-\frac{\gamma}{\theta} - 1\right)d \ln \Gamma + (\theta - 1)d \ln H_1 &= d \ln \lambda + d \ln W_1, \\ \left(\frac{1-\sigma}{\gamma} - 1\right)d \ln \Delta + \left(-\frac{\gamma}{\theta} - 1\right)d \ln \Gamma + (\theta - 1)d \ln H_2 &= d \ln \lambda + d \ln W_2, \end{aligned}$$

and

$$\begin{aligned}
d \ln \Gamma &= \theta \mathbb{B} d \ln H_1 + \theta(1 - \mathbb{B}) d \ln H_2, \\
d \ln \Delta &= \gamma \mathbb{A} d \ln C - \frac{\gamma}{\theta} (1 - \mathbb{A}) d \ln \Gamma \\
&= \gamma \mathbb{A} d \ln C - \gamma(1 - \mathbb{A}) \mathbb{B} d \ln H_1 - \gamma(1 - \mathbb{A})(1 - \mathbb{B}) d \ln H_2,
\end{aligned}$$

where $\mathbb{A} \equiv \frac{\alpha C^\gamma}{\Delta}$ and $\mathbb{B} \equiv \frac{\xi H_1^\theta}{\Gamma}$. Substitute $d \ln \Delta$ and $d \ln \Gamma$ with the formulas above, the system of equations becomes

$$G \times \begin{bmatrix} d \ln C \\ d \ln H_1 \\ d \ln H_2 \end{bmatrix} = \begin{bmatrix} d \ln \lambda + d \ln P \\ d \ln \lambda + d \ln W_1 \\ d \ln \lambda + d \ln W_2 \end{bmatrix},$$

where

$$G = \begin{bmatrix} (\gamma - 1)(1 - \mathbb{A}) - \sigma \mathbb{A} & (\gamma - 1 + \sigma)(1 - \mathbb{A}) \mathbb{B} & (\gamma - 1 + \sigma)(1 - \mathbb{A})(1 - \mathbb{B}) \\ (1 - \gamma - \sigma) \mathbb{A} & [(\gamma - 1 + \sigma)(1 - \mathbb{A}) - (\gamma + \theta)] \mathbb{B} + (\theta - 1) & [(\gamma - 1 + \sigma)(1 - \mathbb{A}) - (\gamma + \theta)](1 - \mathbb{B}) \\ (1 - \gamma - \sigma) \mathbb{A} & [(\gamma - 1 + \sigma)(1 - \mathbb{A}) - (\gamma + \theta)] \mathbb{B} & [(\gamma - 1 + \sigma)(1 - \mathbb{A}) - (\gamma + \theta)](1 - \mathbb{B}) + (\theta - 1) \end{bmatrix}.$$

By the definition of Frisch elasticities, we have

$$G^{-1} = \begin{bmatrix} -\eta_{c,p} & \eta_{c,w_1} & \eta_{c,w_2} \\ \eta_{h_1,p} & \eta_{h_1,w_1} & \eta_{h_1,w_2} \\ \eta_{h_2,p} & \eta_{h_2,w_1} & \eta_{h_2,w_2} \end{bmatrix}.$$

Note that because the values of \mathbb{A} and \mathbb{B} depend on the allocations chosen by households, G and the Frisch elasticities all depend on the allocations and are not deep parameters.

If we want the Frisch elasticities to be deep parameters with such utility function, we must have \mathbb{A} and \mathbb{B} as constants. From the FOC's, this would require

$$\begin{aligned}
\Delta^{\frac{1-\sigma}{\gamma}} \mathbb{A} &= \lambda P C, \\
\Delta^{\frac{1-\sigma}{\gamma}} (1 - \mathbb{A}) \mathbb{B} &= \lambda W_1 H_1, \\
\Delta^{\frac{1-\sigma}{\gamma}} (1 - \mathbb{A})(1 - \mathbb{B}) &= \lambda W_2 H_2.
\end{aligned}$$

⇒

$$\frac{\mathbb{A}}{(1 - \mathbb{A})\mathbb{B}} = \frac{PC}{W_1 H_1} = \text{Constant},$$

$$\frac{\mathbb{A}}{(1 - \mathbb{A})(1 - \mathbb{B})} = \frac{PC}{W_2 H_2} = \text{Constant}.$$

This implies the utility function needs to take the Cobb-Douglas form, i.e., $\gamma = 0$ and $\theta = 0$. In that case, the utility function becomes

$$U(C, H_1, H_2) = \frac{[C^\alpha (H_1^\xi H_2^{1-\xi})^{-(1-\alpha)}]^{1-\sigma} - 1}{1 - \sigma}.$$

Follow the same method, we can derive that

$$G = \begin{bmatrix} \alpha - 1 - \alpha\sigma & (1 - \sigma)(\alpha - 1)\xi & (1 - \sigma)(\alpha - 1)(1 - \xi) \\ (1 - \sigma)\alpha & -\sigma(\alpha - 1)\xi + (\alpha - 2)\xi + (\xi - 1) & -\sigma(\alpha - 1)(1 - \xi) + (\alpha - 2)(1 - \xi) + (1 - \xi) \\ (1 - \sigma)\alpha & -\sigma(\alpha - 1)\xi + (\alpha - 2)\xi + \xi & -\sigma(\alpha - 1)(1 - \xi) + (\alpha - 2)(1 - \xi) - \xi \end{bmatrix},$$

and the Frisch elasticity matrix is just G^{-1} . However, the Cobb-Douglas form is not a good choice because it implies that the ratios between male labor income, female labor income and consumption expenditures are all constants independent of the price of consumption and wages, which is counterfactual.

E The BPS Method in This Paper

E.1 Formulas for the Transmission Coefficients

We follow closely the approximation method proposed by Blundell, Pistaferri, and Saporta-Eksten (2016) (BPS) and try to use the notations consistent with the original paper. However, we do make one slight modification to their formulas. In particular, we add the channel of social security explicitly. We define the human wealth as the sum of the discounted labor income and the discounted retirement benefits, and the share of the retirement benefits in the total discounted human wealth is represented by q_t . By doing so, the true value of the “outside insurance” coefficient β in the model should be zero.

As in BPS, the wage of earner j in household i at age t is determined by

$$\ln W_{i,j,t} = Z_t^{W_j'} \beta^{W_j} + F_{i,j,t} + u_{i,j,t},$$

where $Z_t^{W_j}$ is a group of observable characteristics affecting wages such as age, and

$$F_{i,j,t} = F_{i,j,t-1} + v_{i,j,t}.$$

This implies

$$\Delta \ln W_{i,j,t} - \Delta Z_t^{W_j'} \beta^{W_j} = \Delta u_{i,j,t} + v_{i,j,t}.$$

Define $\Delta w_{i,j,t}$ as the unexpected growth of wage that is not explained by the observables, i.e.,

$$\Delta w_{i,j,t} = \Delta u_{i,j,t} + v_{i,j,t}.$$

Log-linearize the first-order conditions of the household problem, and apply the definition of the Frisch elasticities, we have

$$\begin{bmatrix} d \ln C \\ d \ln H_1 \\ d \ln H_2 \end{bmatrix} = \begin{bmatrix} -\eta_{c,p} & \eta_{c,w_1} & \eta_{c,w_2} \\ \eta_{h_1,p} & \eta_{h_1,w_1} & \eta_{h_1,w_2} \\ \eta_{h_2,p} & \eta_{h_2,w_1} & \eta_{h_2,w_2} \end{bmatrix} \times \begin{bmatrix} d \ln \lambda + d \ln P \\ d \ln \lambda + d \ln W_1 \\ d \ln \lambda + d \ln W_2 \end{bmatrix}.$$

BPS derive that

$$\Delta \ln \lambda_t = \omega_t + \varepsilon_t,$$

where ω_t only depends on age, and hence it will be absorbed by the age dummies in regressions. Then we have

$$\begin{aligned} \Delta c_t &= (-\eta_{c,p} + \eta_{c,w_1} + \eta_{c,w_2})\varepsilon_t + \eta_{c,w_1}(\Delta u_{1,t} + v_{1,t}) + \eta_{c,w_2}(\Delta u_{2,t} + v_{2,t}), \\ \Delta h_{1,t} &= (\eta_{h_1,p} + \eta_{h_1,w_1} + \eta_{h_1,w_2})\varepsilon_t + \eta_{h_1,w_1}(\Delta u_{1,t} + v_{1,t}) + \eta_{h_1,w_2}(\Delta u_{2,t} + v_{2,t}), \\ \Delta h_{2,t} &= (\eta_{h_2,p} + \eta_{h_2,w_1} + \eta_{h_2,w_2})\varepsilon_t + \eta_{h_2,w_1}(\Delta u_{1,t} + v_{1,t}) + \eta_{h_2,w_2}(\Delta u_{2,t} + v_{2,t}), \end{aligned}$$

where Δc_t and $\Delta h_{j,t}$ represent the unexpected growth of consumption and labor supply that are not explained by the observables. Transform labor supply into labor

income, and we have

$$\begin{aligned}
\Delta c_t &= (-\eta_{c,p} + \eta_{c,w_1} + \eta_{c,w_2})\varepsilon_t + \eta_{c,w_1}(\Delta u_{1,t} + v_{1,t}) + \eta_{c,w_2}(\Delta u_{2,t} + v_{2,t}), \\
\Delta y_{1,t} &= (\eta_{h_1,p} + \eta_{h_1,w_1} + \eta_{h_1,w_2})\varepsilon_t + (1 + \eta_{h_1,w_1})(\Delta u_{1,t} + v_{1,t}) + \eta_{h_1,w_2}(\Delta u_{2,t} + v_{2,t}), \\
\Delta y_{2,t} &= (\eta_{h_2,p} + \eta_{h_2,w_1} + \eta_{h_2,w_2})\varepsilon_t + \eta_{h_2,w_1}(\Delta u_{1,t} + v_{1,t}) + (1 + \eta_{h_2,w_2})(\Delta u_{2,t} + v_{2,t}),
\end{aligned} \tag{E.1}$$

where $\Delta y_{j,t}$ is the unexpected growth of earner j 's labor income at age t .

From the log-linearization of the intertemporal budget constraint, assuming the transitory shocks have negligible wealth effects, BPS derive that

$$\begin{aligned}
&(-\eta_{c,p} + \eta_{c,w_1} + \eta_{c,w_2})\varepsilon_t + \eta_{c,w_1}v_{1,t} + \eta_{c,w_2}v_{2,t} \\
&= (1 - \pi_t)q_t \sum_{j=1}^2 s_j [(1 + \eta_{h_j,w_j})v_{j,t} + \eta_{h_j,w_{-j}}v_{-j,t} + (\eta_{h_j,p} + \eta_{h_j,w_j} + \eta_{h_j,w_{-j}})\varepsilon_t],
\end{aligned}$$

where π_t is approximately the share of asset in the total discounted wealth for the household at age t ; $s_{j,t}$ is approximately the share of earner j 's discounted labor income in the total discounted labor income of the household; $1 - q_t$ is the share of the retirement benefits in the total discounted wealth of the household at age t .

From this, we can get the formula for ε_t as

$$\begin{aligned}
\varepsilon_t &= \frac{\eta_{c,w_1} - (1 - \pi_t)q_t[s_1 + \overline{\eta_{h,w_1}}]}{\eta_{c,p} - \eta_{c,w_1} - \eta_{c,w_2} + (1 - \pi_t)q_t(\overline{\eta_{h,p}} + \overline{\eta_{h,w_1}} + \overline{\eta_{h,w_2}})}v_{1,t} \\
&+ \frac{\eta_{c,w_2} - (1 - \pi_t)q_t[s_2 + \overline{\eta_{h,w_2}}]}{\eta_{c,p} - \eta_{c,w_1} - \eta_{c,w_2} + (1 - \pi_t)q_t(\overline{\eta_{h,p}} + \overline{\eta_{h,w_1}} + \overline{\eta_{h,w_2}})}v_{2,t} \\
&\equiv \kappa_{\varepsilon,v_1}v_{1,t} + \kappa_{\varepsilon,v_2}v_{2,t},
\end{aligned}$$

where $\overline{\eta_{h,p}} \equiv \sum_{j=1}^2 s_j \eta_{h_j,p}$, $\overline{\eta_{h,w_1}} \equiv \sum_{j=1}^2 s_j \eta_{h_j,w_1}$, and $\overline{\eta_{h,w_2}} \equiv \sum_{j=1}^2 s_j \eta_{h_j,w_2}$.

Substitute ε_t in Equation (E.1) with the formula above, and we get the formulas

for the transmission coefficients:

$$\begin{bmatrix} \Delta c_t \\ \Delta y_{1,t} \\ \Delta y_{2,t} \end{bmatrix} = \begin{bmatrix} \kappa_{c,u_1} & \kappa_{c,u_2} & \kappa_{c,v_1} & \kappa_{c,v_2} \\ \kappa_{y_1,u_1} & \kappa_{y_1,u_2} & \kappa_{y_1,v_1} & \kappa_{y_1,v_2} \\ \kappa_{y_2,u_1} & \kappa_{y_2,u_2} & \kappa_{y_2,v_1} & \kappa_{y_2,v_2} \end{bmatrix} \begin{bmatrix} \Delta u_{1,t} \\ \Delta u_{2,t} \\ v_{1,t} \\ v_{2,t} \end{bmatrix},$$

where

$$\begin{aligned} \kappa_{c,u_j} &= \eta_{c,w_j}, \\ \kappa_{c,v_j} &= (-\eta_{c,p} + \eta_{c,w_1} + \eta_{c,w_2})\kappa_{\varepsilon,v_j} + \eta_{c,w_j}, \\ \kappa_{y_j,u_j} &= 1 + \eta_{h_j,w_j}, \\ \kappa_{y_j,u_{-j}} &= \eta_{h_j,w_{-j}}, \\ \kappa_{y_j,v_j} &= (\eta_{h_j,p} + \eta_{h_j,w_j} + \eta_{h_j,w_{-j}})\kappa_{\varepsilon,v_j} + (1 + \eta_{h_j,w_j}), \\ \kappa_{y_j,v_{-j}} &= (\eta_{h_j,p} + \eta_{h_j,w_j} + \eta_{h_j,w_{-j}})\kappa_{\varepsilon,v_{-j}} + \eta_{h_j,w_{-j}}, \end{aligned}$$

and

$$\kappa_{\varepsilon,v_j} = \frac{\eta_{c,w_j} - (1 - \pi_t)q_t[s_j + \overline{\eta_{h,w_j}}]}{\eta_{c,p} - \eta_{c,w_1} - \eta_{c,w_2} + (1 - \pi_t)q_t(\overline{\eta_{h,p}} + \overline{\eta_{h,w_1}} + \overline{\eta_{h,w_2}})}.$$

The formulas when the separable preference assumption is imposed can be obtained by assuming the values of all the cross Frisch elasticities to be zero.

E.2 Estimation

The estimation method in this paper follows the empirical strategy in the original BPS paper. To apply the method, we first need the data on the unexpected wage growth $\Delta w_{j,t}$, unexpected consumption growth Δc_t , and unexpected labor income growth $\Delta y_{j,t}$ at household level. These can be obtained by regressing the log-differences of the corresponding variables on observable characteristics and constructing the residuals, i.e.,

$$\Delta \log X_t = Z\hat{\beta} + \Delta x_t,$$

where Z represents the observable characteristics, and $\hat{\beta}$ is the estimated coefficients. For the simulated data, because all the households are ex ante identical, Z contains only a group of age dummies.

E.2.1 Wage Covariances

From the wage process, we know

$$\Delta w_{j,t} = \Delta u_{j,t} + v_{j,t}.$$

Following BPS, we estimate the six wage process parameters, $\sigma_{u_1}^2$, $\sigma_{u_2}^2$, σ_{u_1, u_2} , $\sigma_{v_1}^2$, $\sigma_{v_2}^2$, and σ_{v_1, v_2} , by GMM with an identity weighting matrix using 7 moment conditions: $E[\Delta w_{1,t}^2]$, $E[\Delta w_{2,t}^2]$, $E[\Delta w_{1,t} \Delta w_{2,t}]$, $E[\Delta w_{1,t} \Delta w_{1,t-1}]$, $E[\Delta w_{2,t} \Delta w_{2,t-1}]$, $E[\Delta w_{1,t} \Delta w_{2,t-1}]$, and $E[\Delta w_{2,t} \Delta w_{1,t-1}]$. This step requires only wage data.

E.2.2 Smoothing Parameters

The smoothing parameters π_t , $s_{j,t}$ and q_t are calculated directly from the data. The labor wealth of earner j at age t is calculated as

$$Labor\ Wealth_{j,t} = Y_{j,t} + E_t \sum_{k=1}^{R-t} \frac{Y_{j,t+k}}{(1+r)^k}.$$

Note that the expected future income levels should technically depend on the current states of the households. However, in practice, it is difficult to calculate such conditional expectations exactly, so following BPS, the expected future income levels are only conditional on characteristics that either do not change over time (e.g. education) or change in a perfectly forecastable way (e.g. age). Then $s_{j,t}$ is simply

$$s_{j,t} = \frac{Labor\ Wealth_{j,t}}{\sum_{j=1}^2 Labor\ Wealth_{j,t}}.$$

The retirement wealth, which we define as the sum of discounted retirement benefits, for a household at age t is

$$Retirement\ Wealth_t = \frac{1}{(1+r)^{R-t}} \sum_{k=1}^{T-R} \frac{b}{(1+r)^k}.$$

We then define the human wealth of a household at age t as

$$Human\ Wealth_t = Retirement\ Wealth_t + \sum_{j=1}^2 Labor\ Wealth_{j,t}.$$

So q_t is

$$q_t = \frac{Retirement\ Wealth_t}{Human\ Wealth_t},$$

and π_t is

$$\pi_t = \frac{Assets_t}{Assets_t + Human\ Wealth_t}.$$

E.2.3 Frisch Elasticities and Outside Insurance

For the Frisch elasticities and the “outside insurance” coefficient, we follow BPS and use the 31 moment conditions in Figure 8 of the original BPS paper (and Table 1 of the BPS online Appendix) to conduct a GMM estimation with an identity weighting matrix. The moment conditions include a set of second-order moments of Δc_t , $\Delta y_{j,t}$, $\Delta w_{j,t}$, and the lag of them. The formulas for these moment conditions are derived based on the BPS formulas for Δc_t , $\Delta y_{j,t}$ and $\Delta w_{j,t}$. For example,

$$\begin{aligned} E(\Delta c_t^2) &= E[(\kappa_{c,u1} \Delta u_{1,t} + \kappa_{c,u2} \Delta u_{2,t} + \kappa_{c,v1} v_{1,t} + \kappa_{c,v2} v_{2,t})^2] \\ &= E[\kappa_{c,u1}^2 (2\sigma_{u1}^2) + \kappa_{c,u2}^2 (2\sigma_{u2}^2) + 2(\kappa_{c,u1} \kappa_{c,u2}) (2\sigma_{u1u2}) \\ &\quad + \kappa_{c,v1}^2 (\sigma_{v1}^2) + \kappa_{c,v2}^2 E(\sigma_{v2}^2) + 2(\kappa_{c,v1} \kappa_{c,v2}) (\sigma_{v1v2})] \\ &= p \lim \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \{ \kappa_{c,u1}^2 (i, t) (2\sigma_{u1}^2) + \kappa_{c,u2}^2 (i, t) (2\sigma_{u2}^2) + 2[\kappa_{c,u1}(i, t) \kappa_{c,u2}(i, t)] (2\sigma_{u1u2}) \\ &\quad + \kappa_{c,v1}^2 (i, t) (\sigma_{v1}^2) + \kappa_{c,v2}^2 (i, t) (\sigma_{v2}^2) + 2[\kappa_{c,v1}(i, t) \kappa_{c,v2}(i, t)] (\sigma_{v1v2}) \}. \end{aligned}$$

The results for other moment conditions can be derived in a similar way. We also impose the symmetry assumptions in the form of $\eta_{h_j,p} = -\eta_{c,w_j} \frac{pc}{w_j h_j}$, $j = 1, 2$ as the original BPS paper.

E.2.4 Transmission Coefficients

Collecting the estimation results from previous steps, the transmission coefficients for each household at each age are calculated using the formulas derived in Appendix E.1. The reported transmission coefficients are the sample averages of them.

E.3 Additional Results on the Performance of the BPS Method

In this subsection we provide further results documenting how well the BPS method performs on model-simulated data. In particular, Figure 10 and Figure 11 report the true and estimated transmission coefficients to household consumption of permanent wage shocks over the life cycle for the additively separable and non-separable preferences respectively.

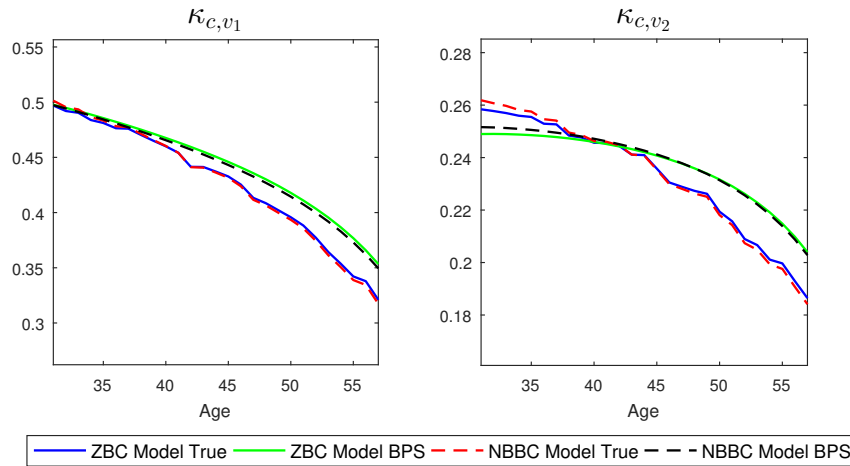


Figure 10: Age Profiles of κ_{c,v_j} : Model True vs. Model BPS




Figure 11: Age Profiles of κ_{c,v_j} : Model True vs. Model BPS (Non-separable Preferences)

F Computation Method

The household optimization problem is solved backwards using the endogenous grid method proposed by Carroll (2006). When the extensive margin of female labor supply is allowed, for each iteration and each household state, the optimization problem is solved twice under two alternative scenarios: the current period female labor supply is strictly positive or zero. The final optimal policy is obtained by comparing the discounted utility achieved in these two scenarios.

The grid for asset has 100 grid points, and the distance between two adjacent grid points increases with the asset level such that the grid points are denser around the low asset levels where borrowing constraints are more likely to bind. The range of the asset grid is age-dependent and eventually endogenously determined by the model to have a better coverage of the more relevant state space.

The joint process of the two earners' permanent wage components is approximated by a discrete Markov process with age-dependent sets of states and transition matrices, and each state corresponds to one possible realization of the two permanent components. The number of states is fixed, but the values of them vary across ages to match the unconditional dispersion of the joint distribution over the life cycle. The grid points and transition matrices are constructed in the same spirit as

Tauchen (1986), and try to mimic the joint unit-root process. The grid for the two permanent components has 11 points in each dimension, so there are in total 121 grid points at each age. The discretization of the transitory components is similar and simpler. Since the transitory components are *iid* across ages, the grid no longer needs to be age-dependent, and no transition matrix is required. The grid for the transitory components has 5 points in each dimension, so there are in total 25 grid points.

In the following, we briefly describe how we solve the household optimization problem when female labor supply is strictly positive. When female labor supply is zero, one only needs to impose this restriction properly in the first-order conditions, and follows the same steps.

F.1 Additively Separable Preferences

For a working household, the first-order conditions are

$$\begin{aligned} C^{-\sigma} &= \lambda, \\ \psi_1 H_1^{\eta_1 - 1} &= \lambda e^{h_{1,t} + F_1 + u_1}, \\ \psi_2 H_2^{\eta_2 - 1} &= \lambda e^{h_{2,t} + F_2 + u_2}, \end{aligned}$$

and the Euler equation is

$$C^{-\sigma} \geq \frac{1+r}{1+\delta} \sum_{(F'_1, F'_2)} \pi(F'_1, F'_2 | F_1, F_2) \sum_{(u'_1, u'_2)} \pi(u'_1, u'_2) C'^{-\sigma}$$

with equality when $A' > \underline{A}_{t+1}$.

For a retired household, the Euler equation is

$$C^{-\sigma} \geq \frac{1+r}{1+\delta} C'^{-\sigma} \text{ with equality when } A' > \underline{A}_{t+1}.$$

When the borrowing constraint is not binding, for any $A' > \underline{A}_{t+1}$, given the

policy function for the next period $c(A', F'_1, F'_2, u'_1, u'_2, t + 1)$, we can solve for the current consumption (C) directly using the Euler equation. Then we can pin down the corresponding male and female labor supply (H_1 and H_2) using the first-order conditions if it is a working household, and finally the current asset level (A) using the budget constraint. For the current asset levels such that the borrowing constraint is binding, i.e., $A' = \underline{A}_{t+1}$, we can first denote the male and female labor supply in terms of the current consumption as implied by the first-order conditions, and then solve for the current consumption using the budget constraint.

F.2 Non-separable Preferences

For a working household, the first-order conditions are

$$\begin{aligned}\Delta^{\frac{1-\sigma}{\gamma}-1} \alpha C^{\gamma-1} &= \lambda P, \\ \Delta^{\frac{1-\sigma}{\gamma}-1} (1-\alpha) \Gamma^{-\frac{\gamma}{\theta}-1} \xi H_1^{\theta-1} &= \lambda W_1, \\ \Delta^{\frac{1-\sigma}{\gamma}-1} (1-\alpha) \Gamma^{-\frac{\gamma}{\theta}-1} (1-\xi) H_2^{\theta-1} &= \lambda W_2,\end{aligned}$$

where $\Delta \equiv \alpha C^\gamma + (1-\alpha)[\xi H_1^\theta + (1-\xi)H_2^\theta]^{-\frac{\gamma}{\theta}}$ and $\Gamma \equiv \xi H_1^\theta + (1-\xi)H_2^\theta$. The Euler equation for a working household is then

$$\Delta^{\frac{1-\sigma}{\gamma}-1} C^{\gamma-1} \geq \frac{1+r}{1+\delta} \sum_{(F'_1, F'_2)} \pi(F'_1, F'_2 | F_1, F_2) \sum_{(u'_1, u'_2)} \pi(u'_1, u'_2) \Delta'^{\frac{1-\sigma}{\gamma}-1} C'^{\gamma-1}$$

with equality when $A' > \underline{A}_{t+1}$.

The Euler equation for a retired household is

$$C^{-\sigma_R} \geq \frac{1+r}{1+\delta} C'^{-\sigma_R} \text{ with equality when } A' > \underline{A}_{t+1}.$$

For an age R household, the Euler equation is

$$\alpha \Delta^{\frac{1-\sigma}{\gamma}-1} C^{\gamma-1} \geq \frac{1+r}{1+\delta} \Psi C'^{-\sigma_R} \text{ with equality when } A' > \underline{A}_{R+1}.$$

From the first-order conditions and the utility function, we know H_1 is always

positive for strictly positive wages. So if female labor supply is also strictly positive, we have

$$\begin{aligned}\frac{\xi H_1^{\theta-1}}{(1-\xi)H_2^{\theta-1}} &= \frac{W_1}{W_2}, \\ \frac{\alpha C^{\gamma-1}}{(1-\alpha)\Gamma^{-\frac{\gamma}{\theta}-1}(1-\xi)H_2^{\theta-1}} &= \frac{P}{W_2}.\end{aligned}$$

Then after some algebra, we can denote H_1 and H_2 in terms of C :

$$\begin{aligned}H_1 &= \left\{ C^{1-\gamma} P \frac{(1-\alpha)}{\alpha} \left\{ \xi [(1-\xi)W_1]^{\frac{\theta}{\theta-1}} + (1-\xi)(\xi W_2)^{\frac{\theta}{\theta-1}} \right\}^{-\frac{\gamma}{\theta}-1} \frac{\xi}{W_1} [(1-\xi)W_1]^{\frac{\gamma+\theta}{\theta-1}} \right\}^{\frac{1}{\gamma+1}} \\ &\equiv \phi_1 C^{\frac{1-\gamma}{\gamma+1}}, \\ H_2 &= \left\{ C^{1-\gamma} P \frac{(1-\alpha)}{\alpha} \left\{ \xi [(1-\xi)W_1]^{\frac{\theta}{\theta-1}} + (1-\xi)(\xi W_2)^{\frac{\theta}{\theta-1}} \right\}^{-\frac{\gamma}{\theta}-1} \frac{(1-\xi)}{W_2} (\xi W_2)^{\frac{\gamma+\theta}{\theta-1}} \right\}^{\frac{1}{\gamma+1}} \\ &\equiv \phi_2 C^{\frac{1-\gamma}{\gamma+1}}.\end{aligned}$$

Substitute H_1 and H_2 in the marginal utility of consumption with the formulas above, we can write the marginal utility in terms of only C :

$$\begin{aligned}MU(C, H_1, H_2) &= \alpha \Delta^{\frac{1-\sigma}{\gamma}-1} C^{\gamma-1} \\ &= \alpha \left\{ \alpha + (1-\alpha) \left[\xi \phi_1^\theta + (1-\xi) \phi_2^\theta \right]^{-\frac{\gamma}{\theta}} C^{\frac{-2\gamma}{\gamma+1}} \right\}^{\frac{1-\sigma-\gamma}{\gamma}} C^{-\sigma}.\end{aligned}$$

After retirement,

$$MU(C) = \Psi C^{-\sigma}.$$

Now we can solve the household optimization problem in a way similar to the separable preference case. When the borrowing constraint is not binding, for any $A' > \underline{A}_{t+1}$, given the next period policy function $c(A', F'_1, F'_2, u'_1, u'_2, t+1)$, we can calculate the marginal utility of consumption in the next period and then solve for the current consumption using the Euler equation. If the borrowing constraint is binding, we can again pin down the current consumption using the budget constraint.