American Economic Journal: Macroeconomics 2021, 13(1): 79–113
https://doi.org/10.1257/mac.20180125

Consumption Insurance against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation†

By Chunzan Wu and Dirk Krueger*

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 US data. In the model, 35 percent of male and 18 percent of female permanent wage shocks pass through to consumption, compared to the empirical estimates of 32 percent and 19 percent. Most of the consumption insurance against permanent male wage shocks is provided through the presence and labor supply response of the female earner. Abstracting from this private intrahousehold income insurance mechanism strongly biases upward the welfare losses from idiosyncratic wage risk as well as the desired extent of public insurance through progressive income taxation. Relative to the standard one-earner life cycle model, the optimal degree of tax progressivity is significantly lower and the welfare gains from implementing the optimal system are cut roughly in half. (*JEL* D15, H21, H24, J16, J22, J31)

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 breadwinner with exogenous labor supply 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 influential papers, Blundell, Pistaferri, and Preston (2008) and Kaplan and Violante (2010) measure the magnitude of the household consumption response to earnings shocks with given persistence by consumption insurance coefficients, defined as the fraction of the variance of the shock that does *not* translate into a corresponding change

*Wu: Peking University, 5 Yiheyuan Road, Haidian District, Beijing 100871, China (email: chunzan@sas.upenn.edu); Krueger: University of Pennsylvania, Department of Economics, 133 South 36th Street, Philadelphia, PA 19104, and CEPR, CFS, NBER, and Netspar (email: dkrueger@econ.upenn.edu). Virgiliu Midrigan was coeditor for this article. We thank Greg Kaplan and Luigi Pistaferri for useful conversations at an early stage of this project, as well as Itay Saporta-Eksten for providing important details about the implementation of the BPS method. Krueger gratefully acknowledges financial support from the NSF under grant SES-0820494. The collection of data used in this study was partly supported by the National Institutes of Health under grant number R01 HD069609 and R01 AG040213, and the National Science Foundation under award numbers SES 1157698 and 1623684.
†Go to https://doi.org/10.1257/mac.20180125 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

in log consumption.[1] That is, if the consumption insurance coefficient for a given earnings shock is one, household consumption growth is completely insulated from the earnings shock, and if it is zero, the earnings shock translates one for one into consumption growth. Blundell, Pistaferri, and Preston (2008) empirically estimate these consumption responses to transitory and permanent earnings shocks on US 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 percent.[2] Kaplan and Violante (2010) evaluate whether a calibrated single-earner incomplete markets life cycle model with transitory and permanent earnings shocks is consistent with the empirical estimates of 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-implied consumption insurance coefficient ranges between 7 percent and 22 percent, depending on the tightness of the borrowing constraints.

In Blundell, Pistaferri, and Preston (2008) and Kaplan and Violante (2010), household earnings are treated as exogenous, and the key mechanism through which consumption insurance is achieved is asset accumulation. This literature is therefore silent about the underlying shocks behind earnings fluctuations, as well as the alternative mechanisms through which households respond to these underlying shocks. The current paper instead models the fundamental sources of consumption risk as idiosyncratic shocks (of various persistence) to wages of the male and female earners in an otherwise standard incomplete markets economy with two-member households. In this paper we seek to make three contributions. Our first 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 household members and participation of the female earner, as well as precautionary savings and progressive income taxation and social security) by which consumption insurance occurs in the model.

Second, we assess 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 important recent paper by Blundell, Pistaferri, and Saporta-Eksten (2016), henceforth BPS.[3] 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)

[1] Formally, denote by c_{it} the log of consumption of household i at time or age t, and define the consumption insurance coefficient for earnings shock x_{it}^{n} of type n as

$$\phi_t^n = 1 - \frac{\operatorname{cov}_i\left(\Delta c_{it}, x_{it}^n\right)}{\operatorname{var}_i\left(x_{it}^n\right)},$$

where cov_i and var_i are the cross-sectional (co-)variances across households i at time (age) t.

[2] Also see Santaeulàlia-Llopis and Zheng (2018) for a recent application of the same method to Chinese data.

[3] In Blundell, Pistaferri, and Saporta-Eksten (2018), the authors extend their analysis to a model with children; we discuss this paper in greater detail below.

households.[4] Our findings suggest that a standard Bewley model with two-earner households and endogenous labor supply can explain virtually all of the insurance of household consumption and labor income to wage shocks estimated by BPS. This is 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 35 percent of male permanent wage shocks and 18 percent of female permanent wage shocks pass through to household consumption. The corresponding empirical estimates of BPS are 32 percent and 19 percent, respectively. 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 data counterparts are slightly negative, but economically and statistically close to 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, this is caused by an improved asset position as households age.[5] A decomposition of consumption insurance against *male* permanent wage shocks in the model shows that the presence and labor supply response 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.

Motivated by this last result concerning the importance for consumption insurance of the secondary earner, as our third contribution we demonstrate that the welfare losses from idiosyncratic wage risk, and the desired social insurance via progressive income taxes, are significantly overstated when this private household insurance mechanism is not accounted for. Relative to the standard one-earner life cycle model, the overall welfare cost from idiosyncratic wage risk is reduced by 35 percent (15.0 percent instead of 23.2 percent, measured as consumption equivalent variation). Finally, the optimal[6] degree of tax progressivity is significantly lower and the welfare gains from implementing the optimal system are cut in half in the two-earner model with endogenous family labor supply adjustments, suggesting that modeling this margin of adjustment explicitly is of first-order importance for the evaluation of social insurance polices.

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

[5] In the paper, we also evaluate the empirical approach of BPS using model-simulated data. The estimation equations BPS employ are derived from a theoretical model with endogenous labor supply and incomplete asset markets and require interior solutions of the household maximization problem, which is not 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 strongly by the violation of these assumptions as long as one restricts attention to households aged 30 to 57, as they (and we) do. Most households in this age range are no longer impacted by even a tight borrowing constraint (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 nonparticipation rate is only moderate, both in our simulated data as well as in the original BPS dataset.

[6] Optimality is defined as maximizing expected lifetime utility of a given cohort of households, with all policy reforms required to be revenue-neutral.

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); Holter, Krueger, and Stepanchuk (2019); and Park and Shin (2018). However, their applications mainly focus on inequality and fiscal policy rather than the private consumption insurance question we address here. As we do, Karahan and Ozkan (2013) measure the welfare cost of idiosyncratic earnings risk but do not endogenize labor supply of the household. Heathcote, Storesletten, and Violante (2014) 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, but they abstract from the insurance provided by a second earner in the household. Finally, perhaps closest to this paper is the study by Blundell, Pistaferri, and Saporta-Eksten (2018). In their structural model, households also choose family labor supply (along the intensive and extensive margin) and time spent with children.[7] Their main focus is on how consumption insurance is impacted by the presence of children in the household, and how public income transfers to families with children change this impact. They do not consider the analysis of (optimal) progressive income taxation in the presence of private consumption insurance through spousal labor supply, one main focus of our work.

The paper is also related to the model-based literature on within-household risk sharing and the role female labor supply plays in this context.[8] 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.

The rest of this paper proceeds as follows: Section I sets up the model, and Section II discusses its calibration. Section III reports our main positive results, with main focus on the consumption insurance against wage risk in the model and the relevant mechanisms to achieve that insurance. It also provides an assessment of the potential bias of the BPS estimates. In Section IV, we turn to the normative evaluation of wage risk and the optimal degree of progressive taxation against that risk. Section V provides a sensitivity analysis with respect to the tightness of borrowing constraints and the separability of household utility between consumption and labor

[7] Related to Blundell, Pistaferri, and Saporta-Eksten (2018), both Keane and Wasi (2016) and Eckstein, Keane, and Lifshitz (2019) also develop and structurally estimate life cycle models with spousal labor supply decisions along the extensive margin and human capital accumulation.

[8] An empirical literature also investigates the sources of household consumption insurance against idiosyncratic earnings risk. See, e.g., Dynarski and Gruber (1997) and Cullen and Gruber (2000).

supply of both spouses, and Section VI concludes. The online Appendix contains supplementary results as well as the implementation of the BPS methodology.

I. The Model

In this section, we first describe the physical environment of our model and then state the household maximization in recursive formulation.

A. *Environment*

We study a partial equilibrium life cycle model with idiosyncratic wage risk and endogenous household labor supply. We follow a cohort of a continuum of measure one households over their life cycle. These 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 of equal age: a male and a female. Generically, we denote by $X_{j,t}$ the variable X of earner $j \in \{1,2\}$ at age t, with $j = 1$ ($j = 2$) indicating the male (female) member 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. 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 introducing a fixed per period utility cost f whenever female hours worked is strictly positive. Households discount the future utility at the constant rate δ, so that $1/(1+\delta)$ is the household time discount factor.

The two members of each household are assumed to make joint decisions on consumption and labor supply. Members of a household can work at wages $W_{j,t}$ determined by their labor productivity. The log wages of both household members are stochastic and represent the sum of a deterministic life cycle component $g_{j,t}$, a transitory component $u_{j,t}$, and a permanent component $F_{j,t}$:

$$
\begin{aligned}
\ln W_{j,t} &= g_{j,t} + F_{j,t} + u_{j,t}, \\
F_{j,t} &= F_{j,t-1} + v_{j,t}, \\
\begin{bmatrix} v_{1,t} \\ v_{2,t} \end{bmatrix} &\sim \text{i.i.d. } N\left(0, \begin{bmatrix} \sigma_{v_1}^2 & \sigma_{v_1,v_2} \\ \sigma_{v_1,v_2} & \sigma_{v_2}^2 \end{bmatrix}\right), \\
\begin{bmatrix} u_{1,t} \\ u_{2,t} \end{bmatrix} &\sim \text{i.i.d. } N\left(0, \begin{bmatrix} \sigma_{u_1}^2 & \sigma_{u_1,u_2} \\ \sigma_{u_1,u_2} & \sigma_{u_2}^2 \end{bmatrix}\right).
\end{aligned}
$$

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 falls 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 in which she is alive.

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 $\underline{A}_t$.

B. *Household Optimization Problem*

A working household's problem can be written in recursive form as

$$\begin{aligned} V(A, F_1, F_2, u_1, u_2, t) \; = \; & \max_{C, A', H_1, H_2} u(C, H_1, H_2) - \mathbf{1}(H_2 > 0) f \\ & + \frac{1}{1+\delta} \sum_{(F_1', F_2')} \pi(F_1', F_2' \,|\, F_1, F_2) \sum_{(u_1', u_2')} \pi(u_1', u_2') V(A', F_1', F_2', u_1', u_2', t+1), \end{aligned}$$

subject to

$$C + A' = Y - \tilde{T}(Y) - \tau_{ss} Y + (1+r)A,$$

$$Y = W_{1,t} H_1 + W_{2,t} H_2,$$

$$C, H_1, H_2 \geq 0, \quad A' \geq \underline{A}_{t+1},$$

where $\pi(\cdot \,|\, \cdot)$ governs the transition probabilities of the wage shocks,[9] and $\mathbf{1}(H_2 > 0)$ equals 1 if the female hours variable H_2 is positive. Female hours of $H_2 = 0$ corresponds to nonparticipation. The term $\tilde{T}(Y)$ is the income tax function that determines the tax liability of a household with before-tax income Y, and τ_{ss} is a flat payroll tax representing the Social Security and Medicare taxes. The dynamic programming problem of a retired household is given by

$$V^R(A, t) = \max_{C, A'} u^R(C) + \frac{1}{1+\delta} V^R(A', t+1),$$

subject to

$$C + A' = b + (1+r)A,$$

$$C \geq 0, \quad A' \geq \underline{A}_{t+1}.$$

[9] Since 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.

When in working age, the household has an additively separable utility function of 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 key advantage of using this preference structure is that the intertemporal and Frisch labor supply elasticities are exclusively determined by exogenous parameters which are therefore directly interpretable. In Section VB, we explore the robustness of our results to using a nonseparable utility specification that is more flexible in terms of substitution patterns between consumption and hours worked of both spouses. The period utility function for a retired household is given by

$$u^R(C) = u(C,0,0) = \frac{C^{1-\sigma}}{1-\sigma}.$$

II. Calibration

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

A. *Data*

In BPS, the data used are from the 1999 to 2009 waves of the Panel Study of Income Dynamics (Institute for Social Research 1999–2009). The PSID collects data from two groups of households: one group representative of the US population, the other from low-income households in the Survey of Economic Opportunity (SEO). The 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 one goal of our study is to compare the degree of consumption insurance implied by our model to 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 of fitting the BPS estimates. Consequently, if we still find significant differences between our model and the BPS empirical results, they are likely caused by model misspecification rather than inappropriate parameter values.

B. *Demographic and Initial Conditions*

Households are born at age 21, retire at age 65, and die at age 80. Therefore, age 1 in the model corresponds to age 21 in the data, and consequently $R = 45$

and $T = 60$. Households start their life with zero assets and a permanent and transitory components of log wages equal to zero.[10]

C. *Wage Process*

The deterministic life cycle profile of wages is taken directly from the paper by Rupert and Zanella (2015). They estimate this wage profile on PSID data from 1967 to 2008. The original wage profile starts from age 23 and has only biennial values after age 52. Therefore, we interpolate their profiles to annual observations 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 over the life cycle but different levels.[11] The life cycle average of the male wage trend is normalized to be 1. In the BPS data, average annual earnings of working females is 0.491 times that of males; we calibrate the level of female wages to match this earnings ratio.[12]

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 i.i.d. across time but potentially positively correlated between the two earners of a household:

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

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

[10] The median age at first marriage in the United States between 2000 and 2010 is 27.5 for males and 26 for females, according to US census data. The starting age of 21 of married households in the model is therefore younger than in the data. There are mainly two reasons why we made this choice. First, the initial age at which a couple starts cohabitating and thus sharing wage risk is likely much younger than that of eventual marriage. Second, at the time of official marriage, couples 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 aged 30–57 are used (the same age group employed by BPS), and therefore the exact choice of initial conditions for the model is not critical for our results.

[11] Since 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.

[12] One unit of income in the model represents $61,597 in 2000 dollars, which is the average male labor income for age 30–57 households in the BPS dataset. This number is slightly different from the one reported in Table 1 of BPS because BPS report the average of nominal income without adjusting for inflation.

D. *Borrowing Limit*

Since the tightness of the borrowing limit is an important determinant of consumption smoothing opportunities, especially early in life, we calibrate it such that the model matches the debt-to-income ratio of young households in the data. In particular, we set the borrowing limit as $\underline{A}_1 = \underline{A}$, $\underline{A}_{t+1} = (1+r)\underline{A}_t$ if $1 \leq t \leq T-1$, and $\underline{A}_{T+1} = 0$ such that households can borrow up to a limit $\underline{A}$ at age 21 and can roll over that debt (at the fixed interest rate r) until the end of life T. The value for $\underline{A}$ is calibrated such that the median debt-to-income ratio of households aged 21–30, conditional on having any debt, in the model matches its counterpart in the extended BPS dataset, which is 16.3 percent.[13]

E. *Discount Factor and Interest Rate*

We adopt the BPS choice of the real risk-free interest rate of $r = 2$ percent per year. The discount rate δ, a key determinant of household precautionary (and life cycle) saving, is calibrated such that the average net worth for households aged 30 to 57 in the model equals 4.188 times average male labor income, as measured in the BPS dataset. This delivers a time discount rate of 0.5 percent per annum.

F. *Income Tax Function*

We permit the labor income tax function $T(Y)$ to be progressive and, following Bénabou (2002), use a two-parameter tax function (also employed by BPS) of the form

$$\tilde{T}(Y) = Y - (1-\chi)Y^{1-\mu},$$

where μ and χ are two parameters governing the progressivity and the level of the income tax, respectively. It implies that after-tax labor income $Y - \tilde{T}(Y) = (1-\chi)Y^{1-\mu}$ is a concave function of pre-tax labor income. We estimate the income tax function parameters by running the following OLS regression on the BPS dataset:

$$\ln\left(Y - \tilde{T}(Y)\right) = \ln(1-\chi) + (1-\mu)\ln(Y).$$

Tax liabilities $\tilde{T}(Y)$ are defined as federal income taxes minus eligible amounts of EITC and food stamp benefits, all computed by BPS. The estimated tax parameters are $\mu = 0.1327$ and $\chi = 0.1575$.

[13] The extended BPS dataset is generated by the same code provided by BPS, except that it expands the age range to include age 21–30 households. In the data, we only consider noncollateralized debt and abstract from mortgage debt, car loans, and other collateralized debt.

G. *Payroll Tax and Retirement Benefit*

The flat payroll tax in the model τ_{ss} is set to 7.65 percent, based on the actual Social Security and Medicare tax rates on pre-tax income of employees. In the United States, social security benefits are piecewise linear functions of average monthly past earnings over the working life. Additional rules govern benefits for spouses. A full representation of the US social security system is costly in terms of computation, since it adds a continuous state variable to the recursive formulation of the problem. Hence we model the progressivity of the US social security benefit formula starkly by assuming that benefits b per household are independent of past contributions. We calibrate b to the average social security benefits for married retired households aged 62 and older between 1999 and 2009 in the Current Population Survey (CPS), given by \$18,484 in 2000 dollars. Since the benefits from Medicare are difficult to measure directly, we assume that they are proportional to the social security benefits, based on the ratio of Medicare tax rate to Social Security tax rate. Therefore, the retirement benefit b in the model is calibrated to $\$18{,}484 \times 7.65\%/6.2\% = \$22{,}807$ in 2000 dollars. This implies a retirement benefit relative to average earnings of age 30–57 households of 27 percent, somewhat lower than the typical replacement rate in life cycle models since average earnings of married, working households aged 30–57 are significantly higher, and replacement rates lower, than in the overall population.

H. *Fixed Utility Cost of Female Participation*

In the data, both male and female workers have significant labor market nonparticipation rates. However, since the BPS results are based on a sample of households with working male members, we do not include an extensive margin of male labor supply decisions in the model. To generate empirically plausible female labor supply decisions along the extensive margin in the model, the fixed utility cost of female labor market participation f is chosen such that the average female nonparticipation rate of households aged 30 to 57 is 20 percent, as in the BPS dataset.

I. *Preference Parameters*

With our preference specification, 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 value of parameter ψ_1 scaling the disutility of male labor is calibrated such that the average male labor income of age 30–57 households in the model equals 1 (normalization). The value of ψ_2 for female labor is calibrated to match the ratio of average female labor supply to average male labor supply conditional on working, which is 0.733 in the BPS dataset.

[14] Following 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—Calibrated Model Parameters

Parameter	Governing	Value
Panel A. Preferences		
δ	Discount rate of utility	4.69×10^{-3}
ψ_1	Disutility of male labor supply	2.538
ψ_2	Disutility of female labor supply	1.953
σ	Inverse of consumption Frisch elasticity	$1/0.578$
η_1	Male labor supply Frish elasticity	0.528
η_2	Female labor supply Frisch elasticity	0.850
f	Fixed utility cost of female participation	0.0306
Panel B. Wage process		
$e^{g_{2,t}-g_{1,t}}$	Female-male wage trend ratio	0.485
$\sigma^2_{v_1}$	Variance of male permanent shocks	0.0303
$\sigma^2_{v_2}$	Variance of female permanent shocks	0.0382
σ_{v_1,v_2}	Covariance of permanent shocks	0.0027
$\sigma^2_{u_1}$	Variance of male transitory shocks	0.0275
$\sigma^2_{u_2}$	Variance of female transitory shocks	0.0125
σ_{u_1,u_2}	Covariance of transitory shocks	0.0058
Panel C. Redistribution system		
μ	Income tax progressivity	0.1327
χ	Income tax level	0.1575
τ_{ss}	Payroll tax	0.0765
b	Retirement benefits	0.3703
Panel D. Financial market		
r	Risk-free interest rate	0.02
$\underline{A}$	Borrowing constraints	-0.126

Note: This table reports the values of parameters in the model with additively separable preferences.

Table 1 summarizes the calibration of the model, and Table 2 reports the empirical calibration targets as well as the (near-perfect) fit of the model along these dimensions.

III. Quantitative Results

We now report, in Section IIIB, how well households are insured against individual wage shocks of both earners in the model, how this extent of insurance compares to the empirically estimated BPS transmission coefficients, and then evaluate the relative importance of various insurance mechanisms (labor supply adjustments, savings, progressive income tax, etc.). Prior to doing so, we briefly document, in the next section, the model-implied mean and variance profiles over the life cycle of consumption, asset, and labor supply, and compare them to the data.[15]

[15] To generate these profiles, 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 ages 21 to 80, although we only use observations from age 30 to 57 when comparing model implications to the BPS estimates, consistent with the dataset they use. Details about the numerical method are provided in online Appendix G.

Table 2—Empirical Targets Matched

Empirical targets	Data	Model
Average male labor income	1	1.000
Average female labor income | work	0.491	0.491
Female-male ratio of average labor supply | work	0.733	0.733
Average female nonparticipation rate	0.20	0.200
Average net worth	4.188	4.189
Median debt-to-income ratio | debt (age 21–30)	0.163	0.163

Notes: This table reports the empirical moments matched by the model with additively separable preferences. Moments are for age 30–57 households unless specified otherwise.

A. *Life Cycle Profiles in the Benchmark Economy*

Figure 1 displays average consumption, asset, and labor supply profiles from age 30 to 57, both for the model (solid blue lines) as well as the PSID data used by BPS (dotted lines with shaded 95 percent confidence intervals). As is common in life cycle models, assets rise over the life cycle as households accumulate wealth to fund retirement consumption (and because $r > \delta$), but also for precautionary reasons to hedge against stochastic wage fluctuations. Assets in the model peak at retirement (age 65 in the model) and are then drawn down to fund retirement consumption (we do not display the latter part of the life cycle in the figure because there are no corresponding BPS data to compare to). Note in the model and in the data that by age 30, households on average have accumulated significant assets (see panel B of Figure 1).

Figure 2 plots the share of households that are borrowing constrained by age, and demonstrates that by age 30 essentially nobody in the economy is directly at the constraint, rendering the constraint fairly unimportant for the consumption and labor supply responses to wage shocks during the prime working years (30–57) of households. This is an important observation to keep in mind for the assessment of the potential bias of the BPS estimates (which rely on the assumption of interior allocations) in Section IIIB.

Consumption in the model rises over the life cycle since wage and earnings risk, and the associated precautionary saving, as well as a fairly high degree of patience ($r - \delta \approx 1.5$ percent) lead to low consumption early in life and subsequent positive consumption growth. Overall, the model captures well the growth of consumption and assets over the life cycle in the data.[16]

Panels C and D of Figure 1 show average hours worked, separately for males and females, over the life cycle. For ease of comparison, we place average hours of males and females on the same scale. As will be clear from Figure 3, lower average female hours originate both from lower hours conditional on working, but also from a very significant nonparticipation rate (on average 20 percent of the

[16] Since the consumption data do not include all types of consumption expenditures, in Figure 1, the life cycle consumption profile from the data is scaled up by a constant factor such that average consumption in the data is identical to that implied by the model.

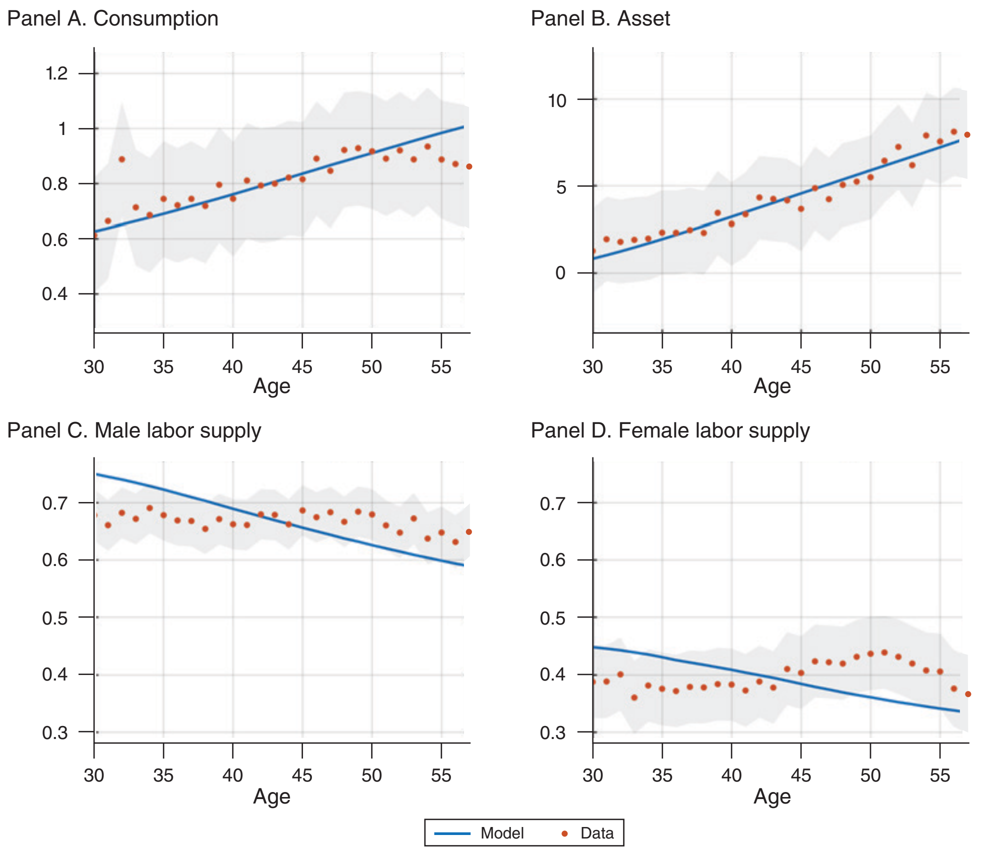

Figure 1. Life Cycles of Cross-Sectional Means

Notes: This figure shows the life cycles of cross-sectional means in the benchmark model with additively separable preferences (solid lines) and in the data (dotted lines) together with the 95 percent confidence interval (grey bands). The data are from the BPS dataset including age 30–57 households. The consumption life cycle from the data is scaled up to match the life cycle average of consumption in the model.

population across all ages), both in the model as well as in the data (see panel B of Figure 3).

Figure 3 displays how consumption and hours dispersion, as measured by the variance of logs, evolve over the life cycle and also shows female nonparticipation rates.[17] In the model, consumption inequality is increasing over the life cycle as permanent wage inequality rises strongly with age, on account of the very substantial exogenous permanent annual wage shocks (with variances of 0.030 and 0.038, respectively) hitting both males and females. The model captures well not

[17] Since in the model households are ex ante identical, the variances of all variables are zero at the start of the life cycle at age 21. And since the model abstracts from other sources of household heterogeneity that might contribute to the variance of consumption and hours worked in the data, it is impossible for the model to match the *level* of the variances in the data. In Figure 3, the life cycle variance profiles from the model are shifted by (variable-specific) constants such that, on average, the model variances match their counterparts in the data. The key question therefore is whether the model implies empirically plausible *changes* in the variances over the life cycle.

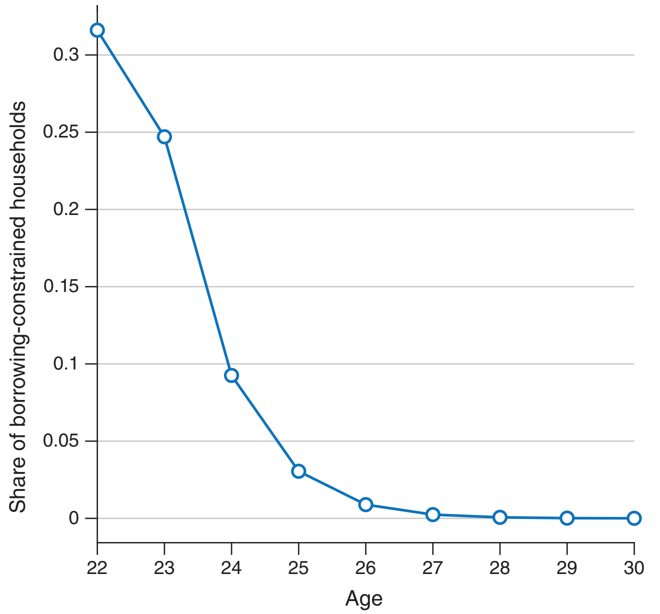

Figure 2. Share of Borrowing-Constrained Households

Note: This figure plots the share of young households on the borrowing constraints in the benchmark model with additively separable preferences.

only the average consumption profile in the data, but also the evolution of consumption inequality over the life cycle. The model also matches the fairly flat life cycle variance profiles of the (model-endogenous) hours choices quite well (see panels C and D of Figure 3).

Overall, even though our model does not give a perfect picture of averages and inequality in the key endogenous economic choices over the life cycle, it especially captures very well both mean consumption and consumption inequality by age, crucial for a paper that focuses on the degree of, and mechanisms for, consumption insurance against wage shocks. In the next section, we first describe how we measure this degree of consumption insurance, before documenting how well consumption is insulated from wage shocks in our model and what roles the different mechanisms for providing that insurance play.

B. *The Transmission of Wage Shocks*: *Model versus Data*

Measuring the Transmission of Wage Shocks.—The main applied purpose of our paper is to quantify how household labor supply, income, and consumption respond to wage shocks of its two members in our model, and how public insurance through the tax system shapes these responses. To measure these responses, Blundell, Pistaferri, and Preston (2008) and BPS introduce the concept of transmission coefficients and estimate them from the PSID data. Let lowercase letters denote the logarithms of variables, so that Δc_t and $\Delta y_{j,t}$ are the growth rates of household consumption and

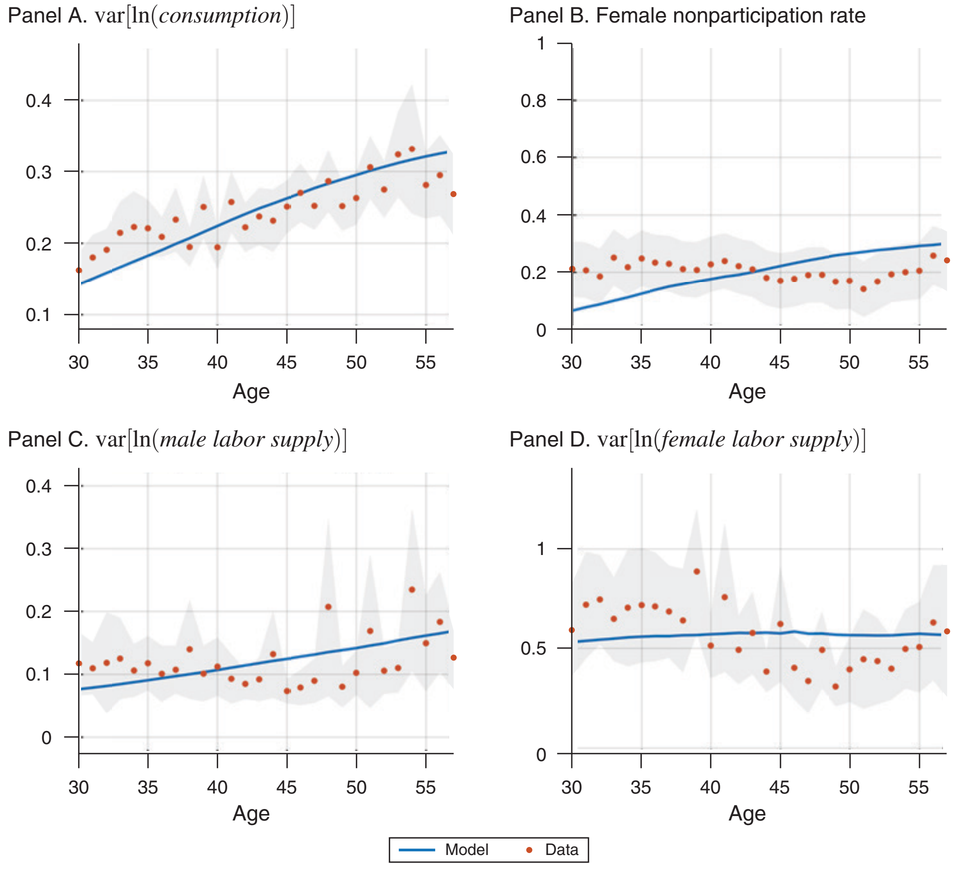

FIGURE 3. LIFE CYCLES OF CROSS-SECTIONAL VARIANCES AND FEMALE NONPARTICIPATION

Notes: This figure shows the life cycles of cross-sectional variances and female nonparticipation rates in the benchmark model with additively separable preferences (solid lines) and in the data (dotted lines) together with the 95 percent confidence interval (grey bands). The data are from the BPS dataset including age 30–57 households. The life cycles of variances in the model are shifted to match the life cycle averages of variances in the data.

labor income of household member j.[18] The responses of these household variables to permanent and transitory wage shocks are captured by the transmission coefficients κ in the following equation:

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

The transmission coefficients therefore measure how consumption and labor income of household member $j \in \{1,2\}$ respond to transitory and permanent wage shocks. For example, κ_{c,v_j} is the household consumption response to earner j's permanent

[18] Empirically, 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.

wage shock. A value of $\kappa_{c,v_j} = 0.4$ means that 40 percent of the shock passes through to household consumption, and hence 60 percent of it is insured. In the next section, we compute the transmission coefficients of shocks κ from equation (1) in the model and contrast them to the empirical estimates by BPS, thereby assessing how much insurance against wage risk households obtain in the model and whether the model captures well the empirically observed transmission of these shocks.

Transmission of Shocks in the Model and in the Data.—In model simulations, realizations of transitory and permanent wage shocks are known separately, in addition to household consumption and labor income. Thus, the transmission coefficients can be obtained directly from OLS regressions of equation (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,"[19] whereas results estimated from the PSID data by BPS are labeled as "Data BPS."

Table 3 reports the "Model True" transmission coefficients, together with the "Data BPS" results. 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.

Comparing the consumption transmission coefficients in the model economy with the BPS estimates, we observe that the model implies almost perfect consumption insurance against transitory wage shocks. Only 1 percent of male and 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 the 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 35 percent of male and 18 percent of female *permanent* wage shocks pass through to household consumption. The corresponding estimates from BPS are 32 percent for shocks to the male and 19 percent for shocks to the female wage. Thus the model-implied consumption insurance against permanent wage shocks is quantitatively very close to the empirical estimates. Taking the BPS point estimates, the model can explain about 96 percent and 101 percent of the consumption insurance against permanent wage shocks.[20]

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). 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.[21]

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

[20] Recall 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.

[21] The 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 coefficient to labor income.

Table 3—Transmission Coefficients in the Data and the Model

	Data BPS (1)	Model True (2)
κ_{c,u_1}	−0.14 (0.07)	0.01
κ_{c,u_2}	−0.04 (0.07)	0.01
κ_{c,v_1}	0.32 (0.05)	0.35
κ_{c,v_2}	0.19 (0.03)	0.18
κ_{y_1,u_1}	1.58 (0.16)	1.44
κ_{y_1,u_2}	0.11 (0.06)	−0.05
κ_{y_1,v_1}	0.92 (0.08)	1.16
κ_{y_1,v_2}	−0.22 (0.04)	−0.19
κ_{y_2,u_1}	0.17 (0.11)	−0.12
κ_{y_2,u_2}	1.88 (0.23)	1.76
κ_{y_2,v_1}	−0.75 (0.14)	−0.51
κ_{y_2,v_2}	1.42 (0.08)	1.46

Notes: The numbers inside parentheses are standard errors from BPS. Only age 30–57 households are included.

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 also explains the slightly negative transmission coefficient $\kappa_{y_j,u_{-j}}$ to labor income of one spouse from a wage shock of the other spouse.[22]

The labor income transmission coefficients of one's own permanent wage shock, 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 that permanent shocks induce. In contrast, when one spouse receives a permanent wage shock, labor supply of the other spouse responds more strongly compared to a transitory shock, again on account of the larger wealth effect on labor supply. The value of $\kappa_{y_1,v_2} = -0.19$ indicates that male labor supply increases by 0.19 percent in response to a permanent 1 percent decline in the female wage, and $\kappa_{y_2,v_1} = -0.51$ implies a strong positive response of female hours of 0.51 percent to a 1 percent permanent reduction in the male wage. These results suggest that the labor supply adjustment of spouses, and especially that of females, is a crucial adjustment mechanism for a household dealing with reductions of male wages, especially permanent ones. Crucially, comparing the transmission coefficients in the model economy (column 2 of Table 3) with their empirical counterparts, "Data BPS" (column 1), the model overall reproduces the main patterns in the data well, with qualitative or significant quantitative deviations mainly observed only in the magnitude of the cross-income response of one household member to temporary wage shocks of the other member (which are negative but small in the model, and positive but statistically insignificant in the data).

[22] Since male wages are on average larger than female wages, so is the wealth effect and induced transmission coefficient on female earnings.

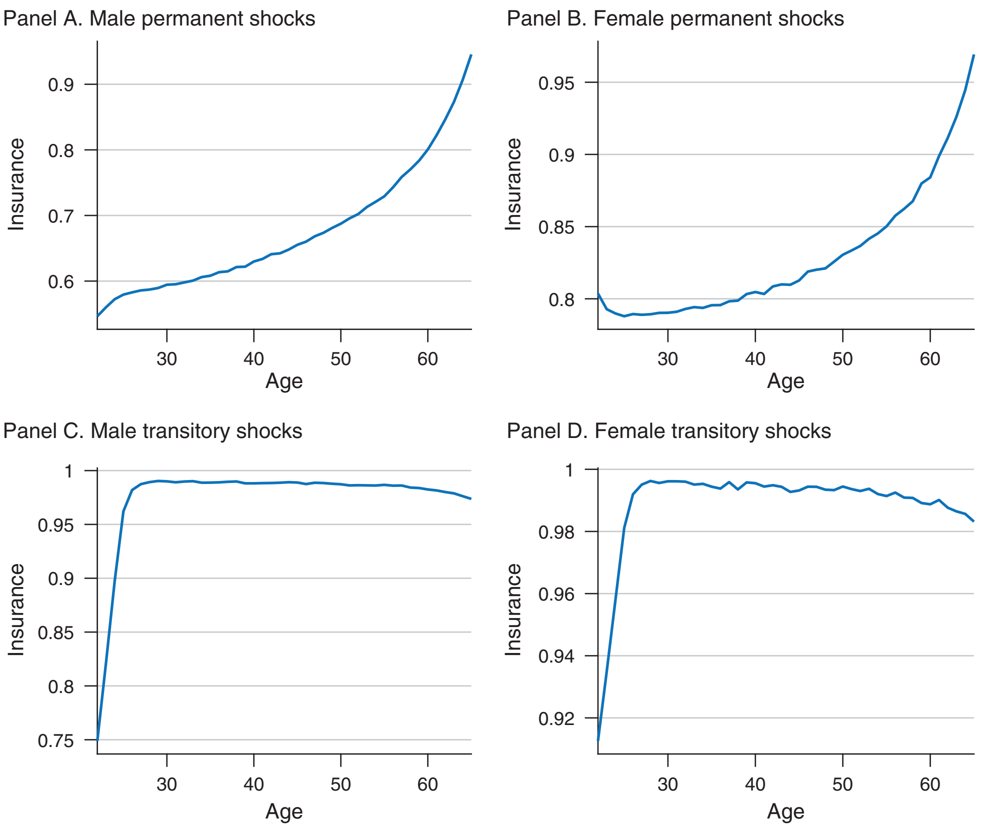

Figure 4. Age Profiles of Consumption Insurance

Note: This figure plots consumption insurance over the life cycle against male (left) and female (right) permanent (top) and transitory (bottom) wage shocks in the benchmark model with additively separable preferences.

Age Profiles of Consumption Insurance.—The transmission coefficients in Table 3 are sample averages for all households aged 30 to 57. Figure 4 documents that there is very substantial age heterogeneity in the response of consumption to wage shocks by household age. Perhaps not surprisingly, consumption is very well insured against transitory wage shocks of both earners, with insurance close to 100 percent (transmission κ_{c,u_j} close to zero) after age 25. At very young ages (21–25), borrowing constraints are binding for a subset of households, and thus household consumption responds more strongly even to transitory shocks, especially those to male wages (see the lower panel of Figure 4). The transmission κ_{c,v_j} of permanent wage shocks to consumption displays much more significant variation over the life cycle, with the amount of consumption insurance against permanent shocks strongly rising over the life cycle, as the top panel of Figure 4 indicates. Better consumption insurance is the result of increased asset accumulation and declining human wealth with household age so that permanent wage shocks become less important for consumption 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 is consistent with the empirical BPS estimates.

As Figure 10 in online Appendix C shows, 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.[23] Most notably, whereas in young ages female labor supply, and thus female earnings, respond quite strongly to an adverse permanent male wage shock ($\kappa_{y_2,v_1} \approx -75$ percent), with age this adjustment mechanism becomes less potent (κ_{y_2,v_1} declines in absolute value) as older households primarily rely on assets to smooth consumption when confronted with a negative, permanent decline in male wages.

Does the BPS Method Recover Well the Transmission of Wage Shocks in the Data?.—Thus far we have treated the BPS estimates of the transmission coefficients κ as accurate representation of the true data. The purpose of this section is to evaluate whether this interpretation is justified by assessing whether their methodology produces unbiased estimates of the true amount of insurance. We of course do not observe the true empirical data-generating process, and therefore we conduct this analysis using simulated data from our model, for which we do, in fact, perfectly know the true transmission coefficients.

To briefly recap the BPS method,[24] they 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 I, and assumes interior solutions (thus abstracting from binding borrowing constraints and extensive margin labor supply decisions), then the transmission coefficients κ are given in closed form as functions $\kappa(\eta, \pi_t, s_{j,t}, \beta)$ of Frisch elasticities η, wealth shares $(\pi_t, s_{j,t})$, and an "outside insurance" coefficient β.[25] The Frisch elasticities themselves are functions exclusively of the deep preference parameters in the household utility function if the latter is separable, but also depend on endogenous choices if it is not. The wealth share π_t measures the share of financial wealth in total (human and financial) wealth, and the $s_{j,t}$ captures the share of household human wealth (present discounted value of future earnings) accruing to each earner $j \in \{1,2\}$.

The BPS method for estimating the transmission coefficients then encompasses four steps: (i) Estimate the variance-covariance matrices of the permanent and transitory wage shocks directly from wage data (with results that were documented in Section II). (ii) Measure the wealth shares π_t and $s_{j,t}$ directly from the asset and labor income data. (iii) Conditional on the results obtained from the first two steps, and using the empirical second-order moments of Δc_t, $\Delta y_{j,t}$, and $\Delta w_{j,t}$, employ

[23] Note that permanent wage shocks are only permanent until retirement, after which social security benefits set in. Consequently, permanent shocks later in life induce less of a wealth effect.

[24] Online Appendix D describes their method as implemented in this paper in detail.

[25] BPS introduce this parameter to capture sources of household insurance that are not explicitly present in their (and our) model, such as insurance provided by networks of relatives and friends. Their baseline results are estimated under the restriction of $\beta = 0$.

a generalized method of moments (GMM) strategy to jointly estimate the Frisch elasticities η and the "outside insurance" coefficient β (unless the latter is restricted to zero).[26] (iv) Calculate the estimates of the transmission coefficients based on the closed-form formulas $\kappa(\eta, \pi_t, s_{j,t}, \beta)$ stated explicitly in online Appendix D. Since we control the data-generating process, we know the true Frisch elasticities and implied transmission coefficients in the model, which we label as "Model True" in the tables below. The estimates based on model-simulated data, to which the BPS methodology is then applied, are denoted as "Model BPS."[27]

To clarify the potential sources of the bias in the BPS estimates, recall that the BPS method imposes several assumptions to obtain a transparent and empirically operational methodology, and violations of these assumptions may result in biased estimates of the transmission coefficients. Their method is based on the log-linearization of the household optimality conditions, and thus it requires interiority of saving and labor supply decisions, and relies on a log-linear approximation of the true nonlinear policy functions being accurate. These assumptions are systematically violated if borrowing constraints are frequently binding or if nonparticipation of one household member is ubiquitous, and the question is whether these violations are severe enough to spill over into significantly biased estimates of the wage shock transmission coefficients.

Table 4 reports the results for the transmission coefficients. The BPS method captures the degree of income and consumption insurance against transitory wage shocks almost perfectly. For permanent shocks, it captures the transmission of female wage shocks to earnings and consumption well. It does, however, somewhat underestimate the consumption insurance against male wage shocks (overestimates the consumption transmission coefficient) and understate the transmission of male wage shocks into male and female earnings, with the latter bias likely due to the selection problem caused by the extensive margin of female labor supply.[28] However, the main message of the table is that the *magnitude* of the bias in all cases

[26] For the baseline BPS estimates, 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 and turns the Frisch elasticities into deep preference *parameters* in the utility function. In column 1 of Table 11 in online Appendix C, we report the BPS estimates of these Frisch elasticities under the separability assumption, which we use as parameter values in our model, and thus labeled as "Model True"; they of course coincide with the values from the calibration Table 1. We also report the estimates of the Frisch elasticities if one applies the BPS methodology to model-simulated data. It confirms that, at least for model-generated data, the BPS method recovers the true elasticities well, especially if the outside insurance coefficient is permitted to be positive. The intertemporal elasticity of substitution in consumption and in male labor supply are very close to their true values. There is a slight downward bias in the female labor supply elasticity (0.81 versus the true 0.85), likely because the maintained assumption of interior female hours is violated in our model for a nontrivial share of observations.

[27] As an important aside, the description of the BPS method also allows us to clarify why the "Model True" transmission estimates might well deviate from the empirical estimates of BPS even though we employ their estimates of the Frisch elasticities (and thus preference parameters) and the stochastic wage process. Their estimates of the insurance coefficients depend on the joint distribution of the wealth shares π_t and $s_{j,t}$ which they derive directly from the data, whereas they are the outcome of household saving and labor supply decisions in our model. Nothing guarantees that our model fits the data along these dimensions, and only if it does will the transmission coefficients estimated by BPS and from our model line up closely. Our results above that the model fits the empirical BPS estimates well are therefore informative about whether our model is a good approximation of the true data-generating process, at least for the aspects related to the consumption and labor income dynamics of prime-age households.

[28] The BPS "Baseline" method does not take into account consumption insurance through the social security system, which taxes labor income proportionally and pays benefits linked only imperfectly to past earnings. Ignoring social security underestimates consumption insurance. Column "SS" in Table 10 in online Appendix C shows that including social security benefits when calculating the smoothing parameters $(\pi_t, s_{j,t})$ cuts the bias in consumption

Table 4—Estimation of Transmission Coefficients

	Model True	Model BPS
κ_{c,u_1}	0.01	−0.01
κ_{c,u_2}	0.01	0.02
κ_{c,v_1}	0.35	0.42
κ_{c,v_2}	0.18	0.22
κ_{y_1,u_1}	1.44	1.47
κ_{y_1,u_2}	−0.05	−0.03
κ_{y_1,v_1}	1.16	1.10
κ_{y_1,v_2}	−0.19	−0.20
κ_{y_2,u_1}	−0.12	−0.10
κ_{y_2,u_2}	1.76	1.76
κ_{y_2,v_1}	−0.51	−0.61
κ_{y_2,v_2}	1.46	1.53

Note: Based on households aged 30–57.

is fairly small, and thus based on simulated model data we would conclude that the BPS estimates are likely good approximations to true insurance in the data, at least if the data are generated by a process close to our model.[29]

Inspecting the Mechanisms: How Do Households Insure against Wage Shocks?.—After having documented that households obtain substantial consumption insurance even against permanent wage shocks in the model, and showing that the extent of insurance accords well with the data (and demonstrating in the previous section that this is likely not an artifact of the BPS method), in this section we seek to better understand what mechanisms, quantitatively, are important for this finding. In our economy, households can smooth wage shocks through four basic mechanisms, two of which that are exogenous to the household and two that 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. Similarly, the progressive income tax code partially insulates after-tax income from pre-tax wage and thus earnings shocks. The two key margins along which households can adjust behavior in response to wage shocks are the labor supply of both members, both along the intensive and along the extensive margin, as well as the accumulation

insurance by more than half. Permitting outside insurance through the parameter β (column "Outside") has the same effect.

[29] In the model, we can base the estimation of transmission coefficients on an arbitrarily large sample, and therefore the model-based estimates we report have standard errors virtually equal to zero. Table 12 in online Appendix C demonstrates that even if we estimate the transmission coefficients on a model sample size comparable to the BPS data, the resulting standard errors are very small (and an order of magnitude smaller than those in the data), suggesting that (i) the actual data are much noisier than those generated from the model, and (ii) the results comparing model estimates to the data and assessing the performance of the BPS method, are equally valid when using the smaller sample of simulated data.

Table 5—Consumption Insurance Decomposition (Male Shocks)

Economy	Insurance provided by: Male earner (%) (1)	Female earner (%): Composition (2)	Extensive (3)	Intensive (4)	Income tax (%) (5)	Savings + social security (%) (6)	Total insurance (%) (7)
Panel A. Permanent shock							
One-earner, exogenous income	—	—	—	—	13.3	33.7	47.0
+ male intensive margin	−1.8	—	—	—	13.5	41.2	52.9
+ female exogenous income	−10.9	31.4	—	—	10.5	29.8	60.9
+ female extensive margin	−11.5	27.4	5.1	—	10.5	30.1	61.6
+ female intensive margin	−17.2	34.5	0.5	15.5	8.9	24.6	66.7
Panel B. Transitory shock							
One-earner, exogenous income	—	—	—	—	13.3	84.6	97.9
+ male intensive margin	−40.0	—	—	—	18.6	119.4	98.0
+ female exogenous income	−42.8	38.6	—	—	13.8	88.4	98.0
+ female extensive margin	−42.8	35.0	1.8	—	14.1	90.0	98.1
+ female intensive margin	−43.2	42.1	0.3	4.2	12.8	81.4	97.7

Notes: This table reports the decomposition results of consumption insurance against male permanent and transitory wage shocks in a sequence of economies with different sets of insurance channels available. Households aged 21–65 are included. Total insurance $= \sum_m Insurance(m)$. Details about the decomposition method are in online Appendix A.

(and decumulation) of assets. Figure 8 in online Appendix A shows how wage shocks of the male earner are mitigated through the various mechanisms before they end up in household consumption. Male wage shocks map male wages into male earnings, and the addition of female earnings (including female labor supply reactions) turns this into household pre-tax earnings. The progressive tax system maps pre-tax into after-tax earnings, and precautionary savings as well as the social security system shape the mapping between household after-tax earnings and consumption. Table 5 breaks down the extent of insurance achieved in each step of this mapping from male wages to consumption.

Table 5 displays the transmission coefficient to consumption of permanent (panel A) and transitory (panel B) male wage shocks in a sequence of economies that differ in their availability of insurance mechanisms. 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).[30] 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 earner, the female earner, household income insurance through the progressive income tax system, as well as through asset accumulation and the progressive social security system (combined).[31] For example, the

[30] We 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. All households aged 21–65 are included in the calculations of consumption insurance. In order to maximize comparability across economies we retain the same calibration across all models.

[31] Since the mapping between after-tax income and consumption is influenced both by private precautionary saving and the social security system, we cannot measure their insurance contribution in the model separately, short

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, of a 1 percent permanent male wage shock 53 percent are transmitted to consumption, and 47 percent are insured. Since with exogenous labor supply and only one earner a male wage shock translates into an equally large household income shock, the only sources of insurance are the progressive income tax system (which insures 13 percent of the pre-tax income decline) and consumption insurance through precautionary asset accumulation and redistributive public social security, which provide a further 34 percent consumption insurance.[32] Fixing a row, the different columns decompose consumption insurance for a given economy, and moving across rows, the last column displays how much extra consumption insurance is achieved by activating an additional adjustment mechanism.

Comparing total consumption insurance (column 7) across rows, we see that whereas in the one-earner model with exogenous labor supply, more than 50 percent of a permanent wage shock transmits to consumption, the presence of a second earner and endogenous labor supply adjustments of both spouses drive up that insurance to about two-thirds of the shock (66.7 percent). Thus consumption insurance increases by about 20 percentage points due to these mechanisms. The column also demonstrates that it is the presence and labor supply adjustment of the second earner, rather than the labor supply response of the primary earner, that is responsible for the better insurance. The latter improves insurance by 5.9 percentage points, whereas the female earner contributes 13.8 percentage points of extra insurance, due to the fact that (i) holding labor supply constant, she provides an independent source of income (the composition effect, supplying 8 percentage points of extra insurance), and (ii) she increases labor supply along the intensive and extensive margin (generating 5.8 percentage points of extra insurance). The intensive margin is relatively more important than the extensive margin, since in the benchmark model 80 percent of female individuals already participate (by calibration of the fixed cost of participation), therefore limiting the quantitative scope for adjustment along this margin.

The point that labor earnings responses of the second earner are crucial for consumption insurance against permanent wage shocks of the male earner is reinforced by decomposing the sources of insurance in the benchmark model, in the last row of panel A of Table 5. With all insurance mechanisms present, male labor supply actually falls in response to a negative permanent male wage shock, and thus a 1 percent decline in male wages leads to a 1.17 percent decline in male earnings. The fact that household consumption only falls by 0.33 percent again stems mainly from the presence of female earnings and labor supply adjustments of the female worker (overall insuring approximately 50 percent of the male earnings decline), and to a significant

of solving a counterfactual model with one of the two elements absent.

[32] This extent of consumption insurance through asset accumulation and social security is slightly higher than that documented in Kaplan and Violante (2010) (34 percent versus their 23 percent). Kaplan and Violante (2010) measure consumption insurance against shocks to after-tax income, whereas we quantify insurance against shocks to pre-tax income. Since we do not recalibrate, the one-earner, exogenous income economy has a higher asset-to-income ratio than the one in Kaplan and Violante (2010) (and a higher one than our benchmark economy as well as the data), implying better consumption insurance against permanent income shocks.

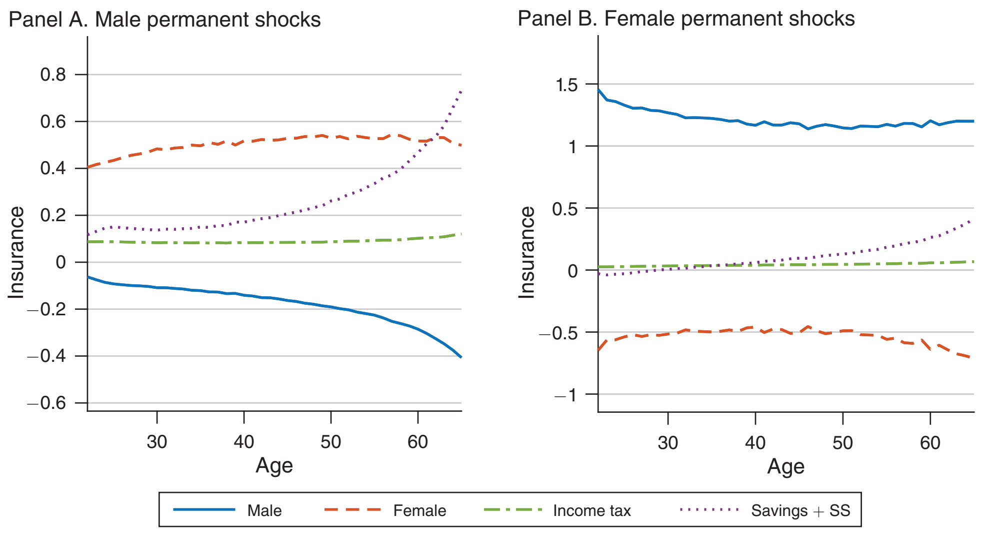

Figure 5. Age Profiles of Consumption Insurance Decomposition

Notes: This figure plots consumption insurance by source over the life cycle against male (panel A) and female (panel B) permanent wage shocks in the benchmark model with additively separable preferences. The sources are the male earner (blue solid line), the female earner (red dash line), progressive income tax (green dash-dot line), and savings plus social security (purple dotted line).

but secondary part from savings adjustment (24.6 percent) and public income insurance through the progressive income tax system (8.9 percent). This last observation also suggests that the presence of a secondary earner and the active adjustment of labor supply in response to male shocks will alter very significantly the welfare cost of these shocks and the demand for public insurance against them. We will return to this point in the next section of the paper. Of course, for transitory wage shocks, the lower panel of Table 5 shows that savings responses are the primary vehicle for providing consumption insurance (with female earners being an important secondary contributor), and that these shocks are almost perfectly insured, as standard permanent income logic suggests.

To briefly investigate whether the contributions of the insurance mechanisms we have highlighted in the previous table vary over the life cycle, in Figure 5 we plot the consumption insurance contributions for a permanent wage shock against household age. Panel A pertains to a permanent shock to male wages, whereas panel B 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 5, the labor supply 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.

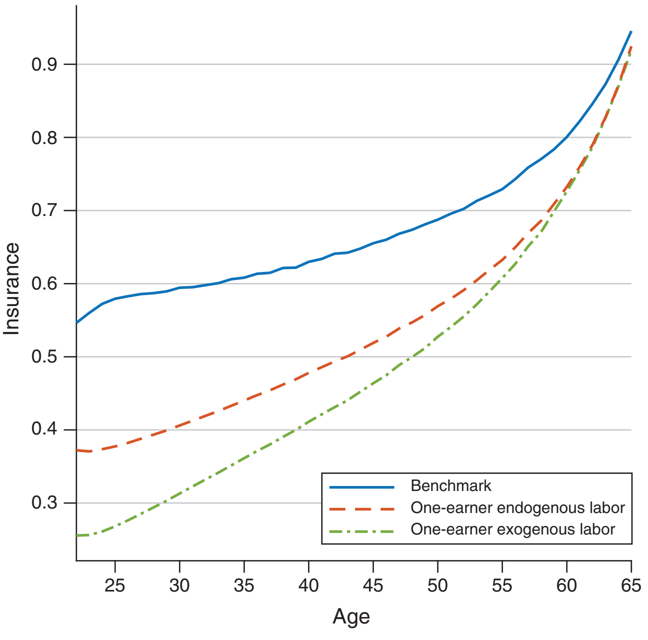

Figure 6. Age Profiles of Consumption Insurance in Three Models

Note: This figure plots consumption insurance over the life cycle against male permanent wage shocks in three models with additively separable preferences.

It is important to keep in mind that the degree of consumption insurance documented empirically by BPS pertains to a very specific (but large) subset of the overall population, the group of married two-earner households in prime working ages, and only to shocks to wages (as opposed to unemployment shocks, health shocks, shocks to family composition, etc.). In light of the importance of labor income of the secondary earner documented thus far, it is plausible to conjecture that households with other characteristics, especially single-earner households, could be subject to significantly less consumption insurance to the same type of wage shocks. Figure 6 verifies this conjecture from the perspective of the model. It shows the total extent of consumption insurance (against permanent male wage shocks) for single-earner households over the life cycle and contrasts it to that of the benchmark model. We observe that although both types of households display very similar consumption insurance in older ages, the consumption response to male wage shocks is significantly larger for single-earner households than for two-earner households in the model, especially if these single earners work in jobs where adjusting hours is difficult (the exogenous earnings model). Whereas even young (30-year-old) households in the benchmark model can insure 60 percent of a male permanent wage shock, a single-earner 30-year-old with flexible hours attains only 40 percent of consumption insurance, and only 30 percent if working in professions with fixed hours. This point reinforces that there might be very significant heterogeneity in the population with respect to the consumption response to idiosyncratic wage shocks, and that the presence and response of a second active earner in the household is a crucial dimension of heterogeneity to consider.

IV. The Welfare Cost of Idiosyncratic Risk and Its Optimal Insurance Revisited

In this section, we present our main economic application of the model by revisiting the welfare cost of wage risk, as well as the optimal social insurance against this risk. In our model, the household can self-insure against this risk not only by engaging in precautionary saving and changing labor supply of the primary earner (as is common in the literature), but also by adjusting labor supply along the extensive and intensive margin of the second earner of the family. We expect that this new margin reduces both the welfare cost of wage risk and the desirable degree of public insurance in the form of labor income tax progressivity. We now measure the extent to which this is true.

A. *Welfare Cost of Idiosyncratic Wage Risk*

We first quantify how much households are willing to pay to be completely insulated from idiosyncratic wage risk in our benchmark model and in two comparison economies in which labor supply either cannot respond to the risk at all (column 3 of Table 6) or the labor supply response is limited to the primary earner (column 2 of Table 6). In both these comparison economies, female labor income is absent (and so is the composition effect, which, as we demonstrated above, is an important source of consumption insurance) and so is the opportunity of the household to adjust female labor supply in response to adverse male wage realizations.

Since all households are ex ante identical, we measure the welfare cost of wage risk as the percentage reduction in consumption (at each age, in each contingency) such that expected lifetime utility is identical in the absence and in the presence of wage risk.[33] When changing the amount of idiosyncratic wage risk, we keep the life cycle profile of average wages constant across all economies. Since we study partial equilibrium models, there is no interaction across different age cohorts on labor or capital markets and no impact of changing wage risk on aggregate factor prices.

Table 6 presents the results. It shows that, as is well known from the literature, even in the presence of self-insurance through saving, the welfare losses from idiosyncratic income risk are large, in the order of 25 percent of lifetime consumption (see the first row of column 3 in the table).[34] These losses stem primarily from

[33] Specifically, let $(\mathbf{c}^0, \mathbf{h}_1^0, \mathbf{h}_2^0)$ and $(\mathbf{c}^1, \mathbf{h}_1^1, \mathbf{h}_2^1)$ denote the allocation of consumption and labor supply before and after a change in wage risk, and $W(\mathbf{c}, \mathbf{h}_1, \mathbf{h}_2)$ be the welfare function that gives the lifetime utility under a particular consumption and labor supply allocation. The welfare cost of this change, in consumption-equivalent variation, *CEV*, is defined by

$$W\Big(\big(1 + CEV\big)\,\mathbf{c}^0, \mathbf{h}_1^0, \mathbf{h}_2^0\Big) \;=\; W\Big(\mathbf{c}^1, \mathbf{h}_1^1, \mathbf{h}_2^1\Big).$$

[34] The absolute magnitude of these losses is at the high end of the numbers reported in the literature. For example, including initial risk at age 24, Karahan and Ozkan (2013) report a somewhat smaller welfare cost of 16.8 percent. The main reason is that they estimate a smaller variance of permanent wage shocks than BPS (0.0113 versus 0.0303 in BPS). Another reason is that Karahan and Ozkan (2013) consider a shorter working life (24–60 versus 21–65 in our model), and hence idiosyncratic wage risk is less important.

Table 6—Welfare Cost of Idiosyncratic Wage Risks

	Benchmark model (%) (1)	One-earner household (%) Endogenous labor (2)	Exogenous labor (3)
Panel A. Male wage risks			
Total welfare change	−15.02	−23.16	−24.54
Consumption	−12.40	−20.60	−24.54
Level	1.22	−0.76	0.58
Distribution	−13.45	−19.99	−24.97
Male labor	1.19	−3.23	—
Level	2.95	−1.30	—
Distribution	−1.71	−1.96	—
Female labor	−4.13	—	—
Level	−4.23	—	—
Distribution	0.10	—	—
Panel B. Female wage risks			
Total welfare change	0.55	—	—
Panel C. All wage risks			
Total welfare change	−14.63	—	—

Notes: This table reports the welfare changes due to the introduction of idiosyncratic wage risks (permanent and transitory) to different model economies. For the benchmark model, female wage risks are introduced first before male wage risks. Welfare changes are reported in consumption-equivalent variations (CEV), and details about the decomposition of welfare changes are in online Appendix B.

the substantial permanent wage shocks that are difficult to insure against through precautionary saving. As the decomposition in the next rows shows,[35] essentially all the welfare losses are due to the fact that the higher income risk feeds into larger consumption risk. These conclusions are largely unchanged if labor supply of the primary earner is permitted to adjust in response to wage risk (compare columns 2 and 3 of Table 6), although household (i.e., male) average labor supply does increase in response to larger male wage risk, providing some consumption insurance against that risk. Overall, although the welfare loss from consumption risk falls by about 5 percentage points (compare the fourth row in columns 2 and 3), this is achieved by on average larger, more dispersed, and therefore welfare-reducing labor supply, as rows 5–7 in column 2 display. Therefore, the overall welfare cost of wage risk is only modestly mitigated by the labor supply margin of the primary earner.

[35] Details about the decomposition of welfare changes are in online Appendix B. As an example, the welfare change due to consumption change, CEV_C, is defined by

$$W\Big(\big(1 + CEV_C\big)\,\mathbf{c}^0, \mathbf{h}_1^0, \mathbf{h}_2^0\Big) \;=\; W\Big(\mathbf{c}^1, \mathbf{h}_1^0, \mathbf{h}_2^0\Big).$$

Further, CEV_C can be decomposed into level and distribution effects, CEV_{CL} and CEV_{CD}, defined by

$$W\Big(\big(1 + CEV_{CL}\big)\,\mathbf{c}^0, \mathbf{h}_1^0, \mathbf{h}_2^0\Big) \;=\; W\Big(\frac{\bar{\mathbf{c}}^1}{\bar{\mathbf{c}}^0}\mathbf{c}^0, \mathbf{h}_1^0, \mathbf{h}_2^0\Big),$$

$$W\Big(\big(1 + CEV_{CD}\big)\big(1 + CEV_{CL}\big)\,\mathbf{c}^0, \mathbf{h}_1^0, \mathbf{h}_2^0\Big) \;=\; W\Big(\mathbf{c}^1, \mathbf{h}_1^0, \mathbf{h}_2^0\Big),$$

where $\bar{\mathbf{c}}^0$ and $\bar{\mathbf{c}}^1$ are the average consumption before and after the change.

The comparison between our benchmark model and endogenous labor supply response opportunities of the secondary earner (column 1) demonstrates that this is a very effective mechanism for dealing with wage risk of the primary earner, with the overall welfare cost from that risk being reduced by 39 percent (15.0 percent instead of 24.5 percent; see row 1 of columns 1 and 3). Most of this reduction stems from better consumption insurance afforded by the labor supply response of the secondary earner (comparing rows 2 and 4 across the three columns of Table 6). Importantly, as the remaining rows of column 1 show, now male labor supply can fall when male wage realizations are low as the female member of the household can start to work, or work longer hours, to compensate the male income loss.[36] The implied welfare losses from extended female hours are partially offset by the welfare gains of shorter male hours. Overall, however, uninsured consumption risk remains the largest cost of idiosyncratic wage risk of the primary earner but with the secondary labor supply margin acting as a quantitatively very important mitigating factor.

This improved private insurance against wage risk can also plausibly be expected to reduce the demand for public income insurance. In the next subsection we demonstrate this point by revisiting the optimal degree of labor income tax progressivity in the presence of joint household labor supply decisions.

B. *Optimal Public Insurance through Progressive Income Taxation*

To determine the optimal degree of tax progressivity we maximize expected lifetime utility of a newborn household with respect to the tax progressivity parameter μ, and adjusting the tax level parameter χ such that the present discounted value (at the fixed interest rate r) of taxes paid by the cohort over its life cycle remains constant, and thus all potential policy reforms are revenue neutral. We conduct this thought experiment both for the benchmark economy with endogenous female labor supply and for the one-earner economy with endogenous labor supply from the previous section.[37] Our objective is to quantify how the optimal degree of public insurance changes in response to the better private household insurance afforded by family labor supply.

We summarize our main results in Figure 7 and Table 7. To interpret these results, recall that the benchmark tax system is given by progressivity parameter $\mu = 0.13$ and level parameter $\chi = 0.16$. The table displays the optimal tax system (for each economy), as well as changes in aggregate variables as well as in welfare, relative to the initial status quo tax system (including a decomposition of the welfare gains). The figure plots, against the degree of tax progressivity, the change in welfare (measured as percent consumption equivalent variation) and consumption insurance against male permanent wage shocks, relative to the benchmark system ($\mu = 0.13$).

[36] As panel B of the table shows, introducing female wage risk is actually welfare improving, since it increases the option value of female labor supply: at high female wage realizations the female worker participates, whereas at low wage realizations it is not worth incurring the fixed cost of participation.

[37] If labor supply and thus earnings are exogenous as, for example, in Blundell, Pistaferri and Preston (2008), then the optimal tax problem is trivial, since the government can provide full consumption insurance by taxing earnings at a confiscatory rate and redistributing the receipts in a lump-sum fashion.

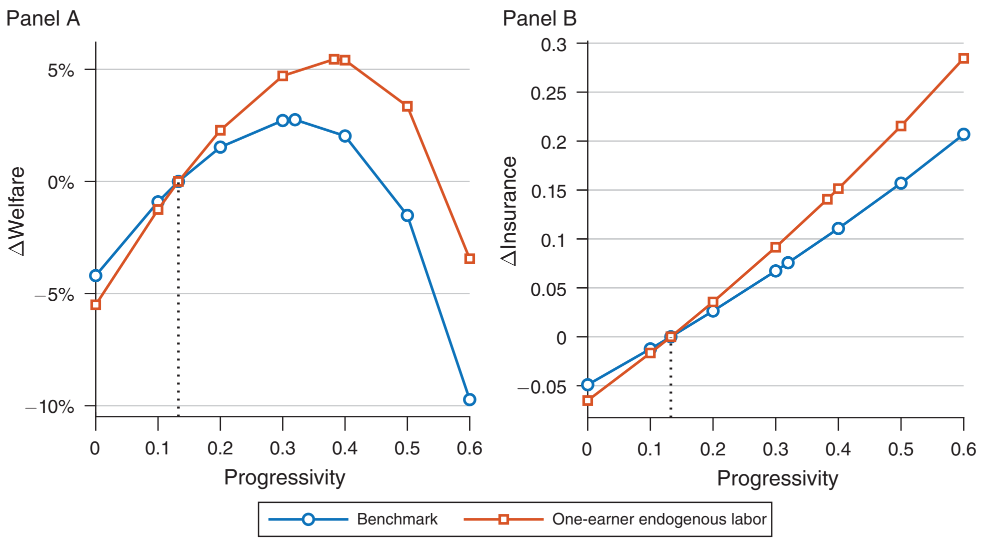

Figure 7. Welfare and Insurance Effects of Tax Progressivity

Notes: This figure shows how household welfare (panel A) and consumption insurance (panel B) change with income tax progressivity in two economies with additively separable preferences. Progressivity is measured by parameter μ in the tax function. Welfare changes are reported in consumption-equivalent variations. Consumption insurance is measured as one minus the transmission coefficient to consumption of male permanent wage shocks. The plots display the changes, relative to the benchmark tax system (i.e., $\mu = 0.1327$).

Panel A of Figure 7 summarizes the two key results of this section. First, relative to the benchmark tax system a very significant increase in tax progressivity strongly improves welfare. More importantly, however, both the magnitude of the welfare gains as well as the optimal degree of tax progressivity fall very significantly in the presence of better private insurance against male wage risk due to family labor supply. Concretely, the optimal tax progressivity parameter falls from 0.38 to 0.32, and the welfare gains from implementing the optimal (within the class of tax functions considered here) are cut in half, from 5.45 percent of lifetime consumption to 2.75 percent of lifetime consumption. The sources of the welfare gains of a more progressive tax system are, as Table 7 clarifies, better consumption insurance and a reduction of utility-reducing male (and if endogenous, female) labor supply, which have to be traded off against lower average consumption associated with larger tax progressivity.

In the previous section, we showed that private consumption insurance improves in the presence of female labor supply. Panel B of Figure 7 provides the corollary: while expanding public insurance through tax progressivity improves overall consumption insurance, it does less so in the benchmark economy with flexible labor supply of the secondary earner. Consequently, abstracting from the joint decision of family labor supply has the potential of very significantly biasing the optimal degree of tax progressivity and the welfare benefits associated with it.

Table 7—Optimal Income Tax

	Benchmark model	One-earner household endogenous labor
Panel A. Optimal income tax		
Progressivity (μ)	0.3197	0.3825
Level (χ)	0.1417	0.1807
Panel B. Changes in aggregate variables		
Consumption	−11.13%	−14.13%
Asset	−17.88%	−26.48%
Male labor supply	−7.31%	−10.76%
Female labor supply	−12.78%	—
Male labor income	−6.76%	−10.03%
Female labor income	−10.58%	—
Female nonparticipation[a]	2.87%	—
Panel C. Welfare gains		
Total welfare gain	2.75%	5.45%
Consumption	−3.73%	−2.75%
Level	−11.13%	−14.13%
Distribution	8.33%	13.25%
Male labor	4.03%	8.43%
Level	3.77%	7.73%
Distribution	0.24%	0.65%
Female labor	2.60%	—
Level	2.75%	—
Distribution	−0.14%	—

Notes: This table reports the optimal income tax policy and the effects of moving from the actual income tax to the optimal one. Welfare changes are reported in consumption-equivalent variations (CEV), and details about the decomposition of welfare changes are in online Appendix B.

[a]The number reported is the actual change in female nonparticipation rate.

V. Robustness

In this section we briefly revisit two important assumptions made in the benchmark model that might affect the results in a quantitatively important way. First, we vary the tightness of borrowing constraints, and second, we relax the assumption of preferences that are additively separable between consumption and labor.

A. *Importance of the Tightness of Borrowing Constraints*

In the benchmark model, we calibrated the tightness of borrowing constraints in such a way that the model matches the debt-to-income ratio of young (aged 21–30) households in the data. Table 8 displays results for the degree of consumption insurance against transitory and permanent male wage shocks (and its decomposition) in the benchmark economy, and in two economies in which borrowing is either ruled out altogether ("Zero BC") and in which the constraints are set so large as to never be binding ("Nonbinding BC"). As in Table 5, households aged 21–65 are included in the calculation of consumption insurance in Table 8.

As the table demonstrates, the degree of consumption insurance against permanent wage shocks, as well as its sources, are virtually unaffected by the tightness of the constraints. The impact on insurance against transitory shocks is more

Table 8—Borrowing Constraints and Consumption Insurance

Economy	Insurance provided by: Male earner (%)	Female earner (%): Composition	Intensive	Extensive	Income tax (%)	Savings + social security (%)	Total insurance (%)
Panel A. Permanent shock							
Zero BC	−17.1	34.5	15.5	0.4	8.9	24.5	66.7
Benchmark	−17.2	34.5	15.5	0.5	8.9	24.6	66.7
Nonbinding BC	−17.3	34.5	15.5	0.4	8.9	25.0	66.9
Panel B. Transitory shock							
Zero BC	−42.9	42.0	4.4	0.2	12.8	80.8	97.4
Benchmark	−43.2	42.1	4.2	0.3	12.8	81.4	97.7
Nonbinding BC	−44.0	42.3	3.8	0.4	12.9	83.3	98.7

Notes: This table reports the decomposition results of consumption insurance against male permanent and transitory wage shocks in three-model economies: a model with zero borrowing constraints (BC), the benchmark model, and a model with nonbinding borrowing constraints (20 times the borrowing constraints in the benchmark model). Households aged 21–65 are included. Total insurance $= \sum_m Insurance(m)$. Details about the decomposition method are in online Appendix A.

noticeable, and as expected, consumption insurance declines as the borrowing constraints tighten. The magnitude of this change is moderate, however, with 98.7 percent of the shock insured if borrowing constraints are not binding, and 97.4 percent if households cannot borrow at all. Interestingly, there is some substitution in the sources of insurance: as borrowing constraints tighten, the importance of savings for consumption insurance declines, and female labor supply adjustments play a larger role in dealing with the shock. We conclude that our benchmark results are qualitatively and, to a large degree, quantitatively robust to the specification of borrowing constraints.

B. *Nonseparable Preferences*

The key advantage of using the preference structure thus far is that the intertemporal and Frisch labor supply elasticities are exclusively determined by exogenous parameters that are therefore directly interpretable. However, it is restrictive in that it does not permit hours worked to affect the marginal utility of consumption (and vice versa). When BPS relax the assumption of separability in their estimation, they find important Frisch complementarity between consumption and leisure. In addition, introducing nonseparable preferences into the model could, in principle, help it better capture aspects of the empirically estimated transmission coefficients the benchmark model had difficulty with, for example, the negative consumption response to positive transitory wage shocks ($\kappa_{c,u_j} < 0$). In this section, we therefore modify the utility function to a nonseparable (between consumption and labor of the two spouses) form:

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

Table 9—Transmission Coefficients with Nonseparable Preferences

	Data BPS (1)	Model True: Separable (2)	Model True: Nonseparable (3)
κ_{c,u_1}	−0.14 (0.07)	0.01	−0.15
κ_{c,u_2}	−0.04 (0.07)	0.01	−0.07
κ_{c,v_1}	0.32 (0.05)	0.35	0.24
κ_{c,v_2}	0.19 (0.03)	0.18	0.13
κ_{y_1,u_1}	1.58 (0.16)	1.44	1.70
κ_{y_1,u_2}	0.11 (0.06)	−0.05	0.09
κ_{y_1,v_1}	0.92 (0.08)	1.16	0.98
κ_{y_1,v_2}	−0.22 (0.04)	−0.19	−0.27
κ_{y_2,u_1}	0.17 (0.11)	−0.12	0.18
κ_{y_2,u_2}	1.88 (0.23)	1.76	1.62
κ_{y_2,v_1}	−0.75 (0.14)	−0.51	−0.47
κ_{y_2,v_2}	1.42 (0.08)	1.46	1.16

Notes: The numbers inside parentheses are standard errors from BPS. Only households aged 30–57 are included.

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 between consumption and labor supply of both spouses. However, now the simple mapping between the preference parameters and the Frisch elasticities BPS estimate is lost, in that the Frisch elasticities are no longer deep parameters, but rather depend on the endogenous choices by households as well as the parameters $(\alpha, \xi, \gamma, \theta, \sigma)$.

The purpose of this section is to document how the degree of consumption insurance in the model is impacted by the nonseparable utility specification, and to investigate the extent to which it helps the model match the empirically estimated transmission coefficients. Table 9 (transmission coefficients, equivalent of Table 3) summarize the most relevant results. The complete set of findings, including the calibration of the model, the model-implied life cycle profiles, estimates of model-implied Frisch elasticities, decomposition of insurance into various mechanisms and an evaluation of the biases of the BPS method with the nonseparable utility are available in online Appendix E.

Relative to the benchmark (column 2 of Table 9), in the model with nonseparable preferences (column 3) consumption and leisure are Frisch complements, and thus higher labor supply (lower leisure) reduces the marginal utility of consumption, and higher consumption increases the marginal disutility of labor supply. This change mainly impacts the labor supply and consumption response to transitory wage shocks. These shocks have only a small wealth effect, and thus their impact is largely determined by the substitution effect. But now a higher wage, inducing larger labor supply (lower leisure), drives down consumption due to the complementarity between consumption and leisure. As a result, household consumption responds mildly negatively to transitory wage shocks by both spouses ($\kappa_{c,u_1} < 0, \kappa_{c,u_2} < 0$), as in the BPS estimates of the data. The same basic mechanism applies to permanent wage shocks (which have a much stronger wealth effect, though), which explains why now permanent wage shocks transmit to consumption even less strongly than in the separable case. Furthermore, since labor (leisure) of both spouses now are complements, positive transitory wage shocks of one spouse now induce a positive hours and thus earnings response of the other spouse ($\kappa_{y_1,u_2} > 0, \kappa_{y_2,u_1} > 0$), again something found in the BPS estimates and hard to rationalize in the benchmark model.

Broadly speaking, the model with nonseparable preferences matches the BPS estimates of the transmission coefficients of *transitory* shocks better, both quantitatively, but also qualitatively (in terms of their signs). Note, however, that most of these estimates are at most marginally statistically significant. On the other hand, it significantly overstates the degree of consumption insurance against permanent wage shocks, especially those of the primary earner (24 percent transmission), relative to the empirical estimates (32 percent transmission), and relative to the benchmark model (35 percent transmission). Since permanent wage shocks are the main sources of welfare losses from incomplete private insurance, and the main argument for the provision of public insurance, and the separable model fares better relative to the BPS estimates in this regard, we decided to conduct our analyses in Section IV with the benchmark model rather than the model with nonseparable preferences.[38] Overall, however, our main conclusion from the previous section remains intact: the Bewley model with endogenous dual earner labor supply implies roughly as much insurance (if not more) against wage shocks, especially permanent wage shocks, as the empirical BPS estimates appear to exhibit.[39]

VI. Conclusion

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 quantitatively replicates well the extent of consumption insurance against permanent wage shocks estimated from US micro household consumption data by Blundell, Pistaferri, and Saporta-Eksten (2016). These results 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 have demonstrated this for the

[38] Tables 16 and 17 in online Appendix E document the biases in the estimates of the Frisch elasticities and transmission coefficients, again based on simulated data. Overall, as in the separable case, these biases are small, indicating that even with nonseparable preferences the BPS method recovers the true Frisch elasticities and transmission coefficients well.

[39] Table 19 in online Appendix E (counterpart of Table 5) again decomposes consumption insurance into the various mechanisms. The main difference to the separable case is that now reducing male labor supply and increasing female labor supply at the same time is very costly (due to the complementarity of leisure), and thus male labor supply and, consequently, earnings do not fall much as in response to a negative male wage shock. The extent to which female labor income and savings are used to insure the lower male earnings is fairly similar across the two versions of the model, however.

case of progressive income taxation. Explicitly modeling spousal adjustments of earnings strongly reduces the welfare losses from wage risk of the primary earner as well as the desired extent of tax progressivity.

Given the importance of this mechanism, a next plausible step in this research agenda would be to investigate, in the context of this class of models, the optimal design of progressive taxation among both earners of the family, including the question of whether to tax both partners jointly or separately, and whether to subject the primary and the secondary earner to systems with different degrees of progressivity.

REFERENCES

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (3): 659–84.

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): 755–64.

Bénabou, Roland. 2002. "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica* 70 (2): 481–517.

Bewley, Truman F. 1986. "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers." In *Contributions to Mathematical Economics in Honor of Gerard Debreu*, edited by Werner Hildenbrand and Andreu Mas-Colell, 79–102. Amsterdam: North-Holland.

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.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2018. "Children, Time Allocation, and Consumption Insurance." *Journal of Political Economy* 126 (S1): S73–S115.

Carroll, Christopher D. 2006. "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems." *Economics Letters* 91 (3): 312–20.

Cullen, Julie Berry, and Jonathan Gruber. 2000. "Does Unemployment Insurance Crowd Out Spousal Labor Supply." *Journal of Labor Economics* 18 (3): 546–72.

Dynarski, Susan, and Jonathan Gruber. 1997. "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity* 27 (1): 229–303.

Eckstein, Zvi, Michael Keane, and Osnat Lifshitz. 2019. "Career and Family Decisions, Cohorts Born 1935–1975." *Econometrica* 87 (1): 217–53.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2010. "The Macroeconomic Implications of Rising Wage Inequality in the United States." *Journal of Political Economy* 118 (4): 681–722.

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.

Holter, Hans A., Dirk Krueger, and Serhiy Stepanchuk. 2019. "How Does Tax Progressivity and Household Heterogeneity Affect Laffer Curves?" *Quantitative Economics* 10 (4): 1317–56.

Huggett, Mark. 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17 (5–6): 953–69.

Imrohoroğlu, Ayşe. 1989. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *Journal of Political Economy* 97 (6): 1364–83.

Integrated Public Use Microdata Series (IPUMS). 1999–2009. "Current Population Survey." Minneapolis, MN: IPUMS. https://cps.ipums.org/cps/ (accessed April 2018).

Kaplan, Greg, and Giovanni L. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics* 2 (4): 53–87.

Karahan, Fatih, and Serdar Ozkan. 2013. "On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory, and Implications." *Review of Economic Dynamics* 16 (3): 452–76.

Keane, Michael P., and Nada Wasi. 2016. "Labour Supply: The Roles of Human Capital and the Extensive Margin." *Economic Journal* 126 (592): 578–617.

Ortigueira, Salvador, and Nawid Siassi. 2013. "How Important Is Intra-Household Risk Sharing for Savings and Labor Supply?" *Journal of Monetary Economics* 60 (6): 650–66.

Institute for Social Research. 1999–2009. "Panel Study of Income Dynamics." Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI. https://psidonline.isr.umich.edu/ (accessed August 2015).

Park, Seonyoung, and Donggyun Shin. 2018. "Welfare Consequences of Rising Wage Risk in the United States: Self-Selection into Risky Jobs and Family Labor Supply Adjustments." Victoria University of Wellington, School of Economics and Finance Working Paper 7967.

Rupert, Peter, and Giulio Zanella. 2015. "Revisiting Wage, Earnings, and Hours Profiles." *Journal of Monetary Economics* 72: 114–30.

Santaeulàlia-Llopis, Raül, and Yu Zheng. 2018. "The Price of Growth: Consumption Insurance in China 1989–2009." *American Economic Journal: Macroeconomics* 10 (4): 1–35.

Wu, Chunzan, and Dirk Krueger. 2021. "Replication data for: Consumption Insurance against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E116881V1.