



RUHR

ECONOMIC PAPERS

Household Taxation, Consumption Insurance, and Welfare

Christian Bredemeier, Jan Gravert, and Falko Juessen



BERGISCHE
UNIVERSITÄT
WUPPERTAL

#1228

Imprint

Ruhr Economic Papers

Ruhr Economic Paper #1228 “Household Taxation, Consumption Insurance, and Welfare”

Responsible Editor: Hendrik Jürges

Jointly published by

RWI – Leibniz-Institut für Wirtschaftsforschung e.V.
Hohenzollernstr. 1-3, 45128 Essen, Germany

Ruhr-Universität Bochum (RUB), Department of Economics
Universitätsstr. 150, 44801 Bochum, Germany

Technische Universität Dortmund, Department of Economic and Social Sciences
Vogelpothsweg 87, 44227 Dortmund, Germany

Universität Duisburg-Essen, Department of Economics
Universitätsstr. 12, 45117 Essen, Germany

Bergische Universität Wuppertal, Schumpeter School of Business and Economics
Gaußstraße 20, 42119 Wuppertal, Germany

Series Coordination:

RWI Press Office, presse@rwi-essen.de, phone +49 (0) 201 8149-213

RWI – Leibniz Institute for Economic Research, Hohenzollernstr. 1-3, 45128 Essen, Germany
www.rwi-essen.de

RWI is funded by the Federal Government and the federal state of North Rhine-Westphalia.
The Institute has the legal form of a registered association; Vereinsregister, Amtsgericht Essen VR 1784

The working papers published in the series constitute work in progress circulated to stimulate discussion and critical comments. Views expressed represent exclusively the authors' own opinions and do not necessarily reflect those of the editors and institutions.

Household Taxation, Consumption Insurance, and Welfare

Christian Bredemeier*, *University of Wuppertal and IZA*

Jan Gravert, *University of Wuppertal*

Falko Juessen, *University of Wuppertal and IZA*

Abstract

We reexamine joint taxation of couples in the presence of wage risk and incomplete insurance markets. By linking spouses' marginal tax rates, joint taxation changes how households use labor supply to self-insure consumption against wage shocks. This creates an insurance advantage of joint taxation worth 1.61 percent of consumption in a calibrated life-cycle model. It largely offsets the disadvantages of joint taxation, though not completely in the current economic and policy environment. Joint taxation can become welfare-preferred at higher tax progressivity and, when progressivity is optimized separately under each system, joint and separate taxation deliver essentially identical welfare.

JEL classification: D15, D52, H21, H24, J22

Keywords: Joint versus separate taxation; Progressivity; Earnings risk; Consumption insurance; Welfare

*Corresponding author: Christian Bredemeier, Schumpeter School of Business and Economics, University of Wuppertal. Email: bredemeier@wiwi.uni-wuppertal.de.

1 Introduction

A central question in the design of income taxation is whether married couples should be taxed as one economic unit or as two separate individuals. A standard static Ramsey logic gives a clear answer. If spouses differ in labor-supply elasticities, efficiency calls for lower marginal tax rates on the more responsive earner. Since secondary earners are typically more responsive, this inverse-elasticity logic favors separate taxation and is one reason why joint taxation is often viewed as inefficient (Boskin and Sheshinski, 1983; Piggott and Whalley, 1996, Apps and Rees, 1999, Gottfried and Richter, 1999, Piggott and Whalley, 1999). This argument, however, abstracts from risk and insurance. Earnings are risky, markets do not provide complete insurance, and spouses can insure each other (e.g., Ortigueira and Siassi, 2013, Blundell et al., 2016, Wu and Krueger, 2021).

We reexamine the choice between joint and separate taxation through this insurance lens. Spouses share resources and adjust labor supply in response to earnings shocks, and progressive taxation interacts with these private insurance arrangements. Under separate taxation, one spouse's earnings directly affect only that spouse's marginal tax rate, whereas under joint taxation they also affect the partner's marginal rate. The tax unit therefore changes the incentives governing how spouses adjust labor supply after a shock.

Consider a negative wage shock to one spouse. Under separate taxation, the resulting decline in earnings lowers that spouse's marginal tax rate but leaves the partner's marginal tax rate unaffected. Under joint taxation, the fall in joint taxable income lowers the marginal tax rates faced by both spouses. This linkage tends to shift labor-supply adjustment toward the partner whose wage has not fallen, allowing the affected spouse's labor supply to follow productivity more closely. Joint and separate taxation can therefore induce different patterns and costs of household self-insurance, suggesting an insurance value of joint taxation.

We quantify this insurance value of joint taxation, identify the adjustment margins that generate it, and assess its contribution to the welfare comparison. Better insurance need not imply higher overall welfare, as the two forms of household taxation also differ in their effects on tax distortions.

To these ends, we study an incomplete-markets model with wage risk and endogenous labor supply in married-couple households. A simplified version of the model provides analytical results on household self-insurance. While the effect of the tax unit on consumption smoothing itself is ambiguous, joint taxation tends to shift labor-supply adjustment toward the spouse not directly affected by a wage shock, allowing the affected spouse's labor supply to follow productivity more closely.

To assess the magnitude of these effects and the resulting welfare comparison, we use the full quantitative model, which incorporates permanent and transitory wage risk and endogenous labor supply at the intensive and extensive margins over the life cycle. We calibrate the model to the U.S. economy using married-couple data from the Panel Study of Income Dynamics. The model follows the structure of Wu and Krueger (2021), and its implications for consumption insurance are checked against the empirical evidence in Blundell et al. (2016).

The quantitative model shows that joint taxation's insurance advantage is economically meaningful. In the baseline economic and policy environment, its insurance value amounts to about 1.6 percent of consumption, while the baseline welfare cost of wage risk is less than 18 percent. This sizeable insurance value arises even though consumption responses to wage shocks are very similar under joint and separate taxation. Labor-supply adjustment, by contrast, differs substantially. Following a permanent wage shock to the husband, he adjusts his hours by about a quarter more under joint than under separate taxation, and his wife changes her hours by about a quarter more in the opposite direction. Her participation also responds more strongly under joint taxation. Restricting households' labor-supply adjustment options confirms that these adjustments account for most of the insurance value, primarily through the intensive margin.

A countervailing fiscal effect prevents this insurance advantage from translating into higher welfare in the baseline economic and policy environment. Progressive taxation raises more revenue from wage dispersion when spouses are taxed separately. Labor-supply responses reinforce this effect. Under a common revenue requirement, the resulting tax adjustments favor separate taxation. The insurance value offsets the resulting advantage of

separate taxation to a large degree, but not completely in the baseline environment, leaving a welfare loss from joint taxation equivalent to 0.74 percent of baseline consumption.

The balance between these forces depends on the economic and policy environment. Over a broad range surrounding the calibrated economy, the insurance value of joint taxation increases with the magnitude of wage risk and as spouses' earnings capacities become more similar. These insurance gains are substantial, but do not suffice to make joint taxation welfare-preferred in the range of economic environments we consider. The picture changes when we consider broader tax reforms. Greater progressivity strengthens the insurance value of joint taxation and, at sufficiently high common progressivity, makes joint taxation welfare-preferred to separate taxation. When progressivity is optimized separately under joint and separate taxation, the two forms of household taxation deliver essentially identical welfare: the insurance advantage of joint taxation virtually offsets its other disadvantages.

Related literature. Our first contribution is to the literature on the optimal tax unit for married couples. The classic result, based on the inverse-elasticity rule, is that spouses should generally be taxed separately when their labor-supply elasticities differ from one another (Boskin and Sheshinski, 1983). This insight has been developed in models with household production (Piggott and Whalley, 1996, Apps and Rees, 1999, Gottfried and Richter, 1999, Piggott and Whalley, 1999, Alesina et al., 2011), endogenous human capital (Alesina et al., 2011, Guner, Kaygusuz, and Ventura, 2012a, 2012b), household bargaining (Alesina et al., 2011), non-cooperative decision making within the household (Meier and Rainer, 2015), and intertemporal household decisions (Guner, Kaygusuz, and Ventura, 2012a, 2012b). We contribute by showing that the tax unit also has an insurance role. This role is absent from static complete-markets arguments but matters when wage risk is incompletely insured and labor supply is endogenous. Separating insurance, deterministic efficiency, and fiscal effects clarifies how that role enters the welfare comparison.

Within this literature, our paper is closest to Corneo (2013), who shows that joint taxation can provide insurance against unemployment risk by linking one spouse's tax rate to the other spouse's labor-market state. We extend this insight to continuous wage

risk and endogenous labor supply in a quantitative life-cycle incomplete-markets model. This setting introduces a distinct self-insurance margin: joint and separate taxation differ not only in their direct tax-insurance effects, but also in how they allocate labor-supply adjustment across spouses.

Our second contribution is to the literature on consumption insurance and household labor supply. A large empirical literature documents spousal labor-supply responses to changes in a partner’s labor-market outcomes, including the added-worker effect following unemployment and cross-wage responses to earnings shocks (e.g., Stephens, 2002, Blundell et al., 2016, Blundell et al., 2018). These responses provide an important margin of household insurance against labor-market risk (Ortigueira and Siassi, 2013; Attanasio et al., 2005, Damiano and Pistaferri, 2020, Asdrubali et al., 2020, Theloudis, 2021, Wu and Krueger, 2021). Wu and Krueger (2021) is particularly close. They show that a calibrated incomplete-markets model with permanent and transitory wage shocks, endogenous labor supply, and two-earner households can replicate the empirical consumption-insurance estimates of Blundell et al. (2016). They also show that progressivity-enhancing tax reforms yield welfare gains, but these gains are reduced because households already self-insure through spousal labor supply. We build on this framework and ask a different question: how does the tax unit shape the household’s self-insurance technology? Our results show that joint taxation changes how households use labor supply for self-insurance by reallocating adjustment across spouses. By restricting adjustments in hours and participation separately, we further establish that the ability to adjust hours is particularly important for the resulting insurance value.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 derives analytical results. Section 4 describes the parameterization. Section 5 presents the quantitative analysis. Section 6 concludes.

2 Model

We use a partial-equilibrium incomplete-markets model with endogenous labor supply and progressive income taxes, building on the consumption-insurance framework of Wu and Krueger (2021). Households insure against wage risk through saving and both spouses' labor supply. The tax unit affects this private insurance technology by changing how a shock to one spouse's wage alters the returns to work of both spouses.

2.1 Environment

Demographics and preferences. The economy is populated by a continuum of ex-ante identical married couples. Each couple consists of one man and one woman. An agent is denoted by the index gj , where $g \in \{m, f\}$ denotes gender and j identifies the household. Spouses are of the same age, denoted by h . They enter the economy at age $\underline{h}$, work until age $h_{\text{ret}} - 1$, retire at h_{ret} , and die at age $\bar{h}$.

Agent gj has preferences over consumption that is public within the household, c_j , and individual labor supply, n_{gj} :

$$u_g(c_j, n_{gj}) = \frac{1}{2} \frac{c_j^{1-\sigma}}{1-\sigma} - \psi_g \frac{n_{gj}^{1+1/\eta_g}}{1+1/\eta_g} - \mathbf{1}(n_{gj} > 0) \Phi_g. \quad (1)$$

Here, σ is risk aversion, η_g is the gender-specific Frisch elasticity, and ψ_g governs the disutility of work. The indicator $\mathbf{1}(n_{gj} > 0)$ equals one when the agent works, so Φ_g is a fixed utility cost of participation.

We assume unitary decision making in the household. Period utility is

$$v(c_j, n_{mj}, n_{fj}) = u_m(c_j, n_{mj}) + u_f(c_j, n_{fj}), \quad (2)$$

and future utility is discounted by β .

Asset markets. At the beginning of each period, households hold wealth a_j in a non-state-contingent asset paying the risk-free interest rate r . They face the borrowing constraint

$$a'_j \geq \underline{a}(h_j + 1), \quad (3)$$

where a prime denotes next-period values and $\underline{a}(h)$ is the lower bound on wealth held at the beginning of age h . The limit evolves according to $\underline{a}(h + 1) = (1 + r)\underline{a}(h)$. All households enter with zero wealth.

Wages. Wage heterogeneity reflects deterministic differences by age and gender as well as permanent and transitory shocks. Let $w_g(h)$ denote the expected wage offer for gender g at age h , and w_{gj} the wage of individual spouse g in household j . Thus, the absence of a household index in $w_g(h)$ indicates an expectation across households. With permanent component z_{gj} , transitory component ϵ_{gj} , and deterministic log-wage component $\ell_g(h_j; w_g, \Sigma)$, wages are given by

$$\begin{aligned} \log w_{gj} &= z_{gj} + \epsilon_{gj} + \ell_g(h_j; w_g, \Sigma), \\ z'_{gj} &= z_{gj} + \varepsilon'_{gj}, \end{aligned} \quad (4)$$

where Σ is the variance-covariance matrix of wage shocks. The permanent innovations ε_{gj} and transitory shocks ϵ_{gj} have mean zero, are jointly normally distributed within each period, and are independent across households and over time. Their within-household covariance matrices are

$$\Sigma_\varepsilon = \begin{pmatrix} \sigma_{\varepsilon_m}^2 & \sigma_{\varepsilon_m, \varepsilon_f} \\ \sigma_{\varepsilon_m, \varepsilon_f} & \sigma_{\varepsilon_f}^2 \end{pmatrix}, \quad \Sigma_\epsilon = \begin{pmatrix} \sigma_{\epsilon_m}^2 & \sigma_{\epsilon_m, \epsilon_f} \\ \sigma_{\epsilon_m, \epsilon_f} & \sigma_{\epsilon_f}^2 \end{pmatrix}.$$

Permanent and transitory shocks are mutually independent. Their variance-covariance matrices enter Σ as

$$\Sigma = \text{diag}(\Sigma_\varepsilon, \Sigma_\epsilon),$$

where the male shock is ordered first in each block.

We treat expected wage levels $w_g(h)$ and the shock process as described by Σ as primitives of the wage process. For any risk environment, the deterministic log-wage component $\ell_g(h; w_g, \Sigma)$ is chosen so that expected wage offers equal $w_g(h)$ at every working age. Changes in risk therefore affect wage dispersion without mechanically changing expected wages over the life cycle.

We abstract from ex-ante wage heterogeneity due to ability differences and from endogenous human-capital accumulation. Ex-ante heterogeneity can raise the demand for progressivity for redistributive purposes, while distorted human-capital accumulation provides a force toward lower progressivity; see, e.g., Heathcote et al. (2010), Holter et al. (2019), and Holter et al. (2023). Our main analysis asks how the tax unit shapes insurance at a given degree of progressivity. These omitted margins are more directly relevant for the final comparison of tax systems at their respective welfare-maximizing progressivity levels, which should be interpreted conditional on the model's sources of heterogeneity and adjustment.

Taxes and pensions. The government provides a total pension benefit b to each retired couple and taxes labor earnings.¹ Social security is financed through a flat payroll tax at rate τ_{ss} .

We use a smooth tax function governed by a scale parameter χ , progressivity μ , and a tax-unit indicator $\theta \in \{0, 1\}$, with $\theta = 1$ for joint and $\theta = 0$ for separate taxation. This follows Gayle and Shephard (2019) and Holter et al. (2023), who extend the two-parameter Benabou (2002) tax function to allow the tax base to depend on both spouses' incomes. With individual gross labor income $y_{gj} = w_{gj}n_{gj}$, household income-tax liabilities are

$$T(y_{mj}, y_{fj}) = y_{mj} + y_{fj} - (1 - \chi) \left(y_{mj}^{1-\mu(1-\theta)} + y_{fj}^{1-\mu(1-\theta)} \right)^{\frac{1-\mu}{1-\mu(1-\theta)}}. \quad (5)$$

For $\theta = 1$, the tax unit is the household:

$$T(y_{mj}, y_{fj} \mid \theta = 1) = y_{mj} + y_{fj} - (1 - \chi)(y_{mj} + y_{fj})^{1-\mu}.$$

¹Tax liabilities depend solely on labor income, as is common in the quantitative consumption-insurance literature; see, for example, Blundell et al. (2016) and Wu and Krueger (2021).

Then, taxes depend on total household earnings, not on their distribution between spouses. For $\theta = 0$, the tax unit is the individual:

$$T(y_{mj}, y_{fj} \mid \theta = 0) = y_{mj} - (1 - \chi)y_{mj}^{1-\mu} + y_{fj} - (1 - \chi)y_{fj}^{1-\mu}.$$

Household taxes are then the sum of individual liabilities, and one spouse's marginal tax rate does not depend directly on the other's earnings.

Higher μ implies greater progressivity. At $\mu = 0$, taxation is linear and net income is proportional to gross income. As μ approaches one, net income becomes insensitive to gross income. The scale parameter χ equals the average income-tax rate at normalized income one—at the household level under joint taxation and at the individual level under separate taxation.

Joint taxation can equivalently be written in income-splitting form:

$$T(y_{mj}, y_{fj} \mid \theta = 1) = 2 \left[\frac{y_{mj} + y_{fj}}{2} - 2^{-\mu}(1 - \chi) \left(\frac{y_{mj} + y_{fj}}{2} \right)^{1-\mu} \right].$$

Thus, joint filing splits household income equally between spouses, applies an individual schedule to each half, and adds their tax liabilities. If separate filing under the same underlying individual schedule is also available, its liability is $y_{mj} + y_{fj} - 2^{-\mu}(1 - \chi)(y_{mj}^{1-\mu} + y_{fj}^{1-\mu})$. Couples weakly prefer joint filing under a progressive schedule.

This filing choice is distinct from our policy comparison. Removing the joint-filing option while holding the underlying schedule fixed would mechanically raise taxes by removing the splitting advantage. Instead, when comparing joint and separate taxation, we adjust the separate-tax scale to match revenue in the corresponding joint-tax economy. The comparison therefore captures the behavioral, efficiency, and insurance consequences of the tax unit at equal public revenue, rather than the revenue increase from simply withdrawing a filing option.

Budget constraints. Let $y_j = y_{mj} + y_{fj}$ denote household gross labor income and $y_j^{\text{net}} = y_j - T(y_{mj}, y_{fj})$ income net of income taxes but before payroll contributions. The working-age budget constraint is

$$c_j + a'_j = (1 + r)a_j + (1 - \tau_{ss})(y_{mj} + y_{fj}) - T(y_{mj}, y_{fj}), \quad (6)$$

and the retirement budget constraint is

$$c_j + a'_j = b + (1 + r)a_j. \quad (7)$$

Retired households do not work.

The smooth tax function abstracts from the bracket structure of the actual tax code. This approximation can affect the strength of the insurance channel in either direction. Under piecewise linear taxation, a shock to one spouse may leave the partner's marginal tax rate unchanged if the household stays in the same bracket, but change it sharply if the household crosses a threshold. The parametric function captures the curvature of a smooth approximation to the tax schedule.

2.2 Household decisions

After observing wage shocks, the household chooses consumption, labor supply, and next-period wealth to solve

$$\begin{aligned} V(a_j, z_{mj}, z_{fj}, \epsilon_{mj}, \epsilon_{fj}, h_j) = & \max_{c_j, n_{mj}, n_{fj}, a'_j} \{v(c_j, n_{mj}, n_{fj}) \\ & + \beta \mathbb{E} [V(a'_j, z'_{mj}, z'_{fj}, \epsilon'_{mj}, \epsilon'_{fj}, h_j + 1)]\}, \end{aligned} \quad (8)$$

subject to nonnegative hours, the borrowing constraint (3), and the relevant budget constraint. Labor supply is zero in retirement. The expectation is over next period's wage shocks conditional on current permanent components. Within each working period, households make their labor-supply and saving decisions after wage shocks have been realized.

Conditional on participation and an interior hours choice, spouse g 's labor-supply condition is

$$\psi_g n_{gj}^{1/\eta_g} = (1 - \tau_{ss} - \text{mtr}_{gj}) w_{gj} c_j^{-\sigma}, \quad (9)$$

where $\text{mtr}_{gj} = \partial T / \partial y_{gj}$ is the marginal income-tax rate. The net-of-tax rate is

$$1 - \text{mtr}_{gj} = (1 - \chi)(1 - \mu) y_{gj}^{-\mu(1-\theta)} \left(y_{mj}^{1-\mu(1-\theta)} + y_{fj}^{1-\mu(1-\theta)} \right)^{-\frac{\mu\theta}{1-\mu(1-\theta)}}. \quad (10)$$

Because Φ_g is a fixed cost, participation is determined by comparing nonparticipation with the best positive-hours choice.

The intertemporal consumption-saving condition is

$$c_j^{-\sigma} = \beta(1 + r) E[(c'_j)^{-\sigma}] + \xi_j, \quad (11)$$

where $\xi_j \geq 0$ is the multiplier on the borrowing constraint, with $a'_j - \underline{a}(h_j + 1) \geq 0$ and $\xi_j[a'_j - \underline{a}(h_j + 1)] = 0$. The same condition applies in retirement.

2.3 Consumption responses to shocks

Consumption insurance is reflected in how strongly household consumption responds to wage shocks. For a permanent innovation or transitory shock $x_{gj} \in \{\varepsilon_{gj}, \epsilon_{gj}\}$ and a positive outcome q_j , define the local transmission coefficient

$$\kappa_{q,x_g} = \frac{\partial \log q_j}{\partial x_{gj}}.$$

Lower consumption transmission, κ_{c,x_g} , means greater consumption insurance, but does not reveal how that insurance is achieved. To distinguish the underlying adjustments, consider first an interior allocation with both spouses employed and participation held fixed. Let $-g$ denote the other spouse. Then

$$\kappa_{c,x_g} = (1 - \mu) \left[s_{gj}^\theta (1 + \kappa_{n_g,x_g}) + s_{-gj}^\theta \kappa_{n_{-g},x_g} \right] + \kappa_{c,x_g}^{resi}, \quad (12)$$

where $\kappa_{c,x_g}^{resi} = \kappa_{c,x_g} - \kappa_{y^{net},x_g}$ captures the difference between consumption and net-income responses due to saving and payroll contributions. Under joint taxation, the weights are gross-income shares, $s_{gj}^1 = y_{gj}/y_j$. Under separate taxation, they are net-income shares, $s_{gj}^0 = y_{gj}^{net}/y_j^{net}$, where $y_{gj}^{net} = (1 - \chi)y_{gj}^{1-\mu}$.

Participation changes generate additional discrete changes in household earnings and resources. We therefore examine participation responses separately alongside hours responses conditional on employment in the quantitative analysis. Equation (12) organizes the continuous margins; it does not assign consumption stabilization to mutually exclusive welfare channels.

The same consumption response can arise from very different combinations of the affected spouse's labor supply, the partner's labor supply, taxes, and intertemporal smoothing. These adjustments need not have the same welfare cost. In particular, cushioning an adverse wage shock by increasing the affected spouse's work shifts labor toward a state in which that spouse is relatively unproductive. Adjustment through the partner avoids part of this cost. Comparing tax units therefore requires more than comparing consumption pass-through: it requires understanding how the household distributes its adjustment across spouses and margins.

3 Analytical results

The tax unit affects household insurance through both labor-supply incentives and spouses' relative earnings. A simplified version of the model separates these two channels. We abstract from payroll taxes, set $\tau_{ss} = 0$, impose a no-borrowing constraint, and consider a working-age household with $a_j = a'_j = 0$. Preferences have log consumption utility and quadratic labor disutility, $\sigma \rightarrow 1$ and $\eta_m = \eta_f = 1$. We set within-household shock covariances to zero. Both spouses work, and we consider local wage changes that leave participation and the binding borrowing constraint unchanged.

The household's problem is then locally static: consumption equals net income and $\kappa_{c,x_g}^{resi} = 0$. Permanent and transitory innovations enter the current wage in the same way,

so the following impact responses apply to either shock. Consumption satisfies

$$\frac{\partial \log c_j}{\partial \log w_{gj}} = (1 - \mu) \left[s_{gj}^\theta \left(1 + \frac{\partial \log n_{gj}}{\partial \log w_{gj}} \right) + s_{-gj}^\theta \frac{\partial \log n_{-gj}}{\partial \log w_{gj}} \right]. \quad (13)$$

The labor-supply conditions imply

$$\frac{\partial \log n_{gj}}{\partial \log w_{gj}} = 1 + \frac{\partial \log(1 - \text{mtr}_{gj})}{\partial \log w_{gj}} - \frac{\partial \log c_j}{\partial \log w_{gj}}, \quad (14)$$

$$\frac{\partial \log n_{-gj}}{\partial \log w_{gj}} = \frac{\partial \log(1 - \text{mtr}_{-gj})}{\partial \log w_{gj}} - \frac{\partial \log c_j}{\partial \log w_{gj}}. \quad (15)$$

Under joint taxation, both spouses face the same marginal tax rate, so their net-of-tax rates respond identically to a change in either spouse's wage. The corresponding elasticity is given by

$$\frac{\partial \log(1 - \text{mtr}_{gj})}{\partial \log w_{gj}} = \frac{\partial \log(1 - \text{mtr}_{-gj})}{\partial \log w_{gj}} = -\mu \left[s_{gj}^1 \left(1 + \frac{\partial \log n_{gj}}{\partial \log w_{gj}} \right) + s_{-gj}^1 \frac{\partial \log n_{-gj}}{\partial \log w_{gj}} \right]. \quad (16)$$

Under separate taxation, each spouse's marginal tax rate responds only to own earnings:

$$\begin{aligned} \frac{\partial \log(1 - \text{mtr}_{gj})}{\partial \log w_{gj}} &= -\mu \left(1 + \frac{\partial \log n_{gj}}{\partial \log w_{gj}} \right), \\ \frac{\partial \log(1 - \text{mtr}_{-gj})}{\partial \log w_{gj}} &= -\mu \frac{\partial \log n_{-gj}}{\partial \log w_{gj}}. \end{aligned} \quad (17)$$

Solving gives, under joint taxation,

$$\frac{\partial \log n_{gj}}{\partial \log w_{gj}} = s_{-gj}^1, \quad \frac{\partial \log n_{-gj}}{\partial \log w_{gj}} = -s_{gj}^1, \quad \frac{\partial \log c_j}{\partial \log w_{gj}} = (1 - \mu)s_{gj}^1, \quad (18)$$

and, under separate taxation,

$$\frac{\partial \log n_{gj}}{\partial \log w_{gj}} = \frac{1 - \mu}{1 + \mu} s_{-gj}^0, \quad \frac{\partial \log n_{-gj}}{\partial \log w_{gj}} = -\frac{1 - \mu}{1 + \mu} s_{gj}^0, \quad \frac{\partial \log c_j}{\partial \log w_{gj}} = (1 - \mu)s_{gj}^0. \quad (19)$$

Adjustment at given income weights. At given effective income weights, consumption responds identically under joint and separate taxation, but both spouses' labor-supply responses are larger in absolute value under joint taxation. Following a negative wage

shock, the affected spouse reduces hours more strongly, reinforcing the earnings loss. The partner increases hours more strongly, cushioning the decline in household resources. Separate taxation dampens both adjustments through the factor $(1 - \mu)/(1 + \mu)$. The same consumption insurance can therefore be achieved through different allocations of work across spouses.

Endogenous income weights. This comparison is partial because the weights are endogenous. When taxes are progressive, separate taxation narrows the earnings gap between spouses, so the primary earner has $s_g^0 < s_g^1$, while the secondary earner has $s_g^0 > s_g^1$. Consequently, separate taxation reduces consumption exposure to primary-earner shocks and increases exposure to secondary-earner shocks in this simplified environment.

Changes in income weights affect own- and partner-hours responses differently. For a primary-earner shock, the larger partner weight under separate taxation offsets the attenuation of the primary earner's own-hours response; sufficiently large share differences can reverse its ranking across regimes. In contrast, the smaller primary-earner weight reinforces the attenuation of the partner's response. For a secondary-earner shock, share compression reinforces the attenuation of the affected spouse's own-hours response but offsets the attenuation of the partner's response. Thus, the relative strength of household adjustment cannot in general be inferred from the tax multiplier alone.

When effective income weights are similar across tax regimes, differences in labor-supply responses arise primarily from the tax multiplier rather than from differences in earnings shares. Consumption pass-through can therefore be similar even when the allocation of work differs substantially. Since these adjustments have different utility costs, consumption insurance alone does not determine the welfare consequences of the tax unit.

The relative importance of these mechanisms depends on the distribution of wage risk between spouses, their endogenous earnings shares, and the scope for saving and participation to absorb shocks. The calibrated model incorporates these margins to assess the insurance gains from joint taxation and their contribution to overall welfare.

4 Parameterization

We calibrate the model to the United States using the Panel Study of Income Dynamics (PSID). The parameterization combines externally set preference and institutional parameters, directly estimated wage and tax inputs, and five internally calibrated parameters governing labor supply, saving, and borrowing.

Data. We use the 1999–2019 PSID survey waves. Retrospective income, hours, and policy variables generally refer to the preceding calendar year. Following Blundell et al. (2016), we exclude the Survey of Economic Opportunity oversample and households with extreme observed wealth, labor income, and total transfers. The main sample consists of different-sex married couples with both spouses aged 30–57 in the retrospective reference year, positive male labor income and hours, and internally consistent income and hours records for both spouses.²

To measure pensions, life-cycle wages, and borrowing, we extend the sample to the relevant ages. Pension benefits are estimated from couples in which both spouses are at least 65; expected wage profiles from couples with both spouses aged 21–65; and the borrowing-limit target from couples with husbands aged 21–30, with no additional age restriction on wives. Empirical quantities use PSID family weights, except for the stochastic wage-process estimation, which follows the wage-risk literature and does not use survey weights.

Externally set parameters. We use the preference elasticities estimated by Blundell et al. (2016): risk aversion $\sigma = 1.73$, and Frisch elasticities $\eta_m = 0.528$ and $\eta_f = 0.85$. The real interest rate is $r = 2\%$ and the payroll tax rate is $\tau_{ss} = 7.65\%$, as in Wu and Krueger (2021). Labor-market entry occurs at age 21, retirement begins at 66, and terminal age is 80.

²Following Blundell et al. (2016), we split household histories into stable couple spells at changes in family composition. Wage changes are constructed only within spells; the change year can be retained as the first observation of a new spell.

Wage profiles and wage risk. We construct two distinct wage inputs. First, we estimate expected life-cycle wage profiles for married men and women by fitting second-order fractional polynomials to average real hourly wages by age. Both profiles are divided by the arithmetic mean of the predicted male profile over ages 21–65. Income and wealth are measured in units of average real gross labor income of husbands; one income unit corresponds to 97,572 real 2019 dollars. These normalizations also determine the scale of model hours. Appendix A reports the normalized profiles.

Second, we estimate permanent and transitory wage risk from residual wage growth. The residual regressions include year effects and controls for birth cohort, education, race, state, large-city residence, changes in large-city residence, and interactions of education, race, and the city variables with year. We estimate the joint residual process for husbands and wives by GMM, using wage-growth variances, own lag autocovariances, and cross-spouse covariances. Measurement error is treated using the validation-study approach from the earlier literature; see Blundell et al. (2016). The estimation details are provided in Appendix B.

Before the first wage shocks are realized, permanent components are common across households and normalized to zero. Wage dispersion at age 21 therefore comes from the first realizations of the estimated permanent and transitory shocks, not from an additional distribution of persistent heterogeneity. The deterministic log-wage component adjusts for the distribution of these shocks so that the stochastic process preserves the expected wage profiles. The same convention applies when risk is varied or removed. In particular, the no-risk wage at a given age equals the expected wage offer in the corresponding risk economy.

Pensions and taxes. For pension benefits, we use married couples in which both spouses are at least 65. We average valid Head-plus-Spouse Social Security income in the 2005–2019 waves, convert it to real 2019 dollars, and scale it by the ratio of the employee

FICA rate including Medicare to the OASDI rate, 0.0765/0.062. Expressed in the model’s income unit, this gives $b = 0.3501$.³

We estimate the tax-function parameters χ and μ under joint taxation, $\theta = 1$, consistent with the filing status of the vast majority of married couples in the United States.⁴ Let $\bar{y}_m$ denote mean gross real labor income of husbands in the main sample and e_j the regression residual. We estimate

$$\log(y_j^{\text{net}}/\bar{y}_m) = \log(1 - \chi) + (1 - \mu) \log(y_j/\bar{y}_m) + e_j,$$

excluding observations with nonpositive gross or net income.⁵ The estimates are $\chi = 0.0971$ and $\mu = 0.1412$. The progressivity estimate lies within the range of previous estimates, including Guner et al. (2014), Blundell et al. (2016), and Holter et al. (2019).

Internally calibrated parameters. We jointly calibrate the discount factor β , labor-disutility weights ψ_m and ψ_f , the fixed participation cost Φ_f , and the initial borrowing limit $\underline{a}$ (21). The calibration strategy follows Wu and Krueger (2021).

Average household net worth relative to average male labor income is 4.5991 in the PSID, yielding $\beta = 0.9838$. Men’s labor-disutility weight makes average male labor income equal to one in model units, giving $\psi_m = 0.9338$. Women’s labor-disutility weight matches the ratio of mean hours among working wives to mean hours among husbands, 0.7577, giving $\psi_f = 1.1041$. The fixed participation cost matches the female nonparticipation rate of 17.59%, giving $\Phi_f = 0.0332$. We set the male fixed cost to zero because the calibration sample requires the husband to work.

To determine the initial borrowing limit, we use young couples with husbands aged 21–30. We match the median ratio of non-mortgage, non-vehicle debt to combined spousal

³Using average male labor income over 2005–2019 instead gives a same-period ratio of 0.3415. We use the full 1999–2019 main-sample normalization because benefits are expressed in the common model income unit.

⁴About 95% of married couples filed jointly in 2018 according to the Internal Revenue Service.

⁵Gross income is joint labor earnings reported in the PSID; net income is constructed. Using year-specific policy rules and PSID earnings and household characteristics, we calculate federal income-tax liability net of the nonrefundable Child Tax Credit, and subtract refundable ACTC and EITC amounts and regular SNAP benefits. Net income is gross income minus these net tax-transfer payments. These are statutory liabilities and eligibility-based credits and benefits, not observed payments or take-up. Payroll taxes are excluded because the model applies τ_{ss} separately. A calculated payroll-tax amount enters only the statutory ACTC alternative for families with at least three children.

labor income, conditional on positive debt.⁶ The target is 0.1923, giving $a(21) = -0.1609$. Table 1 summarizes the parameters, and Table 2 compares the empirical and model moments.

Consumption-insurance fit. Of particular importance is whether the model reproduces the extent and composition of household consumption insurance. Table 3 reports insurance against permanent shocks to the husband’s wage, following Wu and Krueger (2021).⁷ This shock is the standard benchmark in Blundell et al. (2016) and Wu and Krueger (2021) and connects directly to the household-adjustment mechanism.

The model insures 70.15% of a permanent shock to the husband’s wage, close to the empirical estimate of 68% in Blundell et al. (2016). This fit is obtained without targeting consumption insurance in the calibration.

The partner’s earnings are central to consumption insurance. Their contribution reflects both the additional source of household resources provided by the wife’s earnings and her labor-supply responses. Own labor supply contributes negatively because the husband’s hours move procyclically with his wage. Direct tax insurance is modest, and saving together with social security accounts for the residual. The prominence of partner earnings makes the incentives governing spouses’ joint labor-supply decisions particularly relevant for the comparison of tax units.

⁶Non-mortgage, non-vehicle debt (“other debt” in the PSID) includes credit-card debt, student loans, medical and legal bills, and loans from relatives.

⁷We follow the simulated-data decomposition in their Appendix A and combine the three additive contributions associated with the wife’s earnings—income composition and the intensive and extensive labor-supply margins—into one partner component.

Table 1: Parameter values, baseline model

Parameter	Symbol	Value
<i>Externally set</i>		
Risk aversion	σ	1.730
Frisch elasticity, men	η_m	0.528
Frisch elasticity, women	η_f	0.850
Interest rate	r	0.020
Payroll tax rate	τ_{ss}	0.077
<i>Estimated from the PSID</i>		
Expected wage profiles, men and women	$w_m(h), w_f(h)$	Fig. 5
Variance, permanent shock, men	$\sigma_{\varepsilon_m}^2$	0.0319
Variance, permanent shock, women	$\sigma_{\varepsilon_f}^2$	0.0329
Covariance, permanent shocks	$\sigma_{\varepsilon_m, \varepsilon_f}$	0.0036
Variance, transitory shock, men	$\sigma_{\epsilon_m}^2$	0.0153
Variance, transitory shock, women	$\sigma_{\epsilon_f}^2$	0.0163
Covariance, transitory shocks	$\sigma_{\epsilon_m, \epsilon_f}$	0.0040
Pension benefit	b	0.3501
Progressivity	μ	0.1412
Tax scale	χ	0.0971
<i>Internally calibrated</i>		
Discount factor	β	0.9838
Labor disutility, men	ψ_m	0.9338
Labor disutility, women	ψ_f	1.1041
Participation cost, women	Φ_f	0.0332
Initial borrowing limit, age 21	$\underline{a}(21)$	-0.1609

Notes: Risk aversion and Frisch elasticities are from Blundell et al. (2016); the interest rate and payroll tax rate follow Wu and Krueger (2021). The baseline uses joint taxation, $\theta = 1$. The final five parameters jointly match the moments in Table 2.

Table 2: Empirical targets and model moments

Moment	Target	Model
Average male labor income	1	0.9997
Mean hours of working wives / mean hours of husbands	0.7577	0.7570
Female nonparticipation rate	0.1759	0.1759
Average household net worth / average male labor income	4.5991	4.5986
Median ratio of debt to combined spousal labor income, conditional on positive debt (husband age 21–30)	0.1923	0.1923

Notes: The first four moments use ages 30–57. Empirical debt excludes mortgages and vehicle debt.

Table 3: Consumption insurance against male permanent wage shocks

	BPS	Baseline model
Total insurance (percent)	68	70.15
Contribution from		
Own labor supply		−20.75
Partner earnings		63.12
Income tax		8.14
Savings and social security		19.65

Notes: BPS: Blundell et al. (2016). The insurance coefficient ϕ_{ε_m} is one minus the consumption elasticity with respect to the male permanent shock. The four contributions sum to this coefficient, up to rounding. The partner contribution combines income composition and the intensive and extensive margins of the wife’s labor supply; it is not a pure labor-response effect. All entries are expressed in percent.

5 Quantitative analysis

The quantitative analysis links households’ responses to wage shocks to the welfare consequences of joint and separate taxation. We first study the effects of replacing joint with separate taxation at the estimated degree of progressivity. Restrictions on labor choices then distinguish the role of household adjustments in shaping insurance and welfare. We vary wage risk, intra-household earnings inequality, and progressivity to assess how these forces change across economic and policy environments.

In the main policy comparisons, joint and separate taxation raise the same discounted income-tax revenue. The scale parameter under separate taxation is adjusted to match the revenue of the corresponding joint-taxation economy; the payroll-tax rate and pension benefits are unchanged. The comparison concerns two tax systems that raise the same revenue, rather than an individual couple’s choice of filing status under a given tax code.

5.1 Consumption insurance and adjustment margins

Table 4 compares households’ responses to a one-percent permanent increase in the husband’s wage under joint and separate taxation. Corresponding results for permanent shocks to the wife’s wage are reported in Appendix C. We focus on permanent wage shocks because the literature documents that households remain incompletely insured against

them and rely importantly on labor-supply adjustments to cushion their effects (Blundell, Pistaferri, and Saporta-Eksten 2016; Wu and Krueger 2021). By contrast, transitory wage shocks are almost fully insured, predominantly through saving and borrowing. Our model reproduces these patterns, making permanent shocks particularly informative about the labor-supply mechanisms through which the choice of tax unit shapes household self-insurance. The two economies raise the same income-tax revenue and are otherwise identical. Responses are measured as the change in a variable associated with a one-percent wage increase, controlling for the other wage innovations.

For the husband’s wage shock, about 30 percent of the wage increase passes through to household consumption under both joint and separate taxation. If anything, separate taxation provides slightly more consumption insurance against this shock: pass-through is 29.2 percent, compared with 29.9 percent under joint taxation, consistent with the analytical comparison for shocks to the primary earner. The similar consumption responses, however, conceal sizable differences in how households adjust. For every one-percent increase in the husband’s wage, his hours increase by 0.21 percent under joint taxation, compared with 0.17 percent under separate taxation. Conditional on participation, the wife’s hours fall by 0.46 percent under joint taxation, compared with 0.37 percent under separate taxation; her participation probability falls by 0.16 rather than 0.07 percentage points. Gross-income responses are also similar under joint and separate taxation, as the spouses’ labor-supply adjustments largely offset each other. Taxes and intertemporal smoothing further dampen the transmission to consumption.

The corresponding results for permanent shocks to the wife’s wage, reported in Appendix C, reveal a similar pattern. Consumption pass-through is again nearly identical under joint and separate taxation (20.3 and 20.7 percent, respectively), with joint taxation now providing slightly more consumption insurance. Labor-supply responses, however, are again substantially stronger under joint taxation: the wife’s hours and participation increase more strongly, while the husband’s hours decline more strongly. Thus, for shocks to either spouse, the affected spouse’s labor supply responds more strongly to changes in

Table 4: Responses to a one-percent permanent increase in the male wage

	Joint taxation	Separate taxation
Household consumption	0.299	0.292
Male hours	0.208	0.165
Female hours, conditional on participation	−0.460	−0.367
Female participation	−0.160	−0.073
Household gross labor income	0.576	0.583
Household net labor income	0.495	0.489

Notes: Responses are in percent, except female participation, which is in percentage points. Progressivity equals its estimated value; income-tax revenue is matched across joint and separate taxation. We regress changes in log outcomes, or the change in the participation indicator, jointly on the changes in the permanent and transitory log-wage components of both spouses after removing age profiles. Female log-hours responses use women with positive hours in both adjacent periods. Net labor income is gross labor income minus net income-tax liabilities, including transfers represented by the tax function, before social security contributions. Coefficients are normalized to a one-percent male permanent wage innovation.

productivity under joint taxation, accompanied by larger offsetting adjustments by the partner.

These responses reflect the coupling of spouses’ marginal tax rates under joint taxation. Under joint taxation, the effect of a wage shock on marginal tax rates is spread across both spouses. Compared with separate taxation, the affected spouse’s marginal tax rate responds less strongly, while the partner’s marginal tax rate also changes. Consequently, the affected spouse’s labor supply follows changes in productivity more closely, while the partner makes larger offsetting adjustments. Following a negative shock, the same connection strengthens the partner’s incentive to expand work. Joint taxation thus shifts self-insurance toward the partner and away from work by the affected spouse in low-productivity states. Consumption responses can therefore be similar even when the allocation of work, and the utility cost of achieving that consumption path, differs substantially.

5.2 The insurance value of joint taxation

To measure the welfare consequences of these adjustments, we consider how much the welfare loss from wage risk differs between joint and separate taxation. Formally, let $\mathcal{V}_\theta^\Sigma(\chi, e)$ denote welfare, measured by expected lifetime utility, under tax system θ , wage

risk Σ , tax-scale parameter χ , and economic and policy environment e .⁸ This notation separates the central dimensions of comparison, θ and Σ , from the tax-scale parameter χ , whose adjustment depends on how revenue is treated in a given experiment, and from the remaining economic and policy environment e . Welfare is evaluated before the first wage shocks are realized. All households are identical at that point and hold zero assets; the criterion averages over both spouses' permanent and transitory shocks from the first working period onward.

The *welfare loss from wage risk* is

$$\mathcal{L}_\theta^\Sigma(\chi, e) = \mathcal{V}_\theta^{\mathbf{0}}(\chi, e) - \mathcal{V}_\theta^\Sigma(\chi, e), \quad \theta \in \{0, 1\}, \quad (20)$$

where $\mathbf{0}$ denotes a zero matrix conformable with Σ . This object gives the value of eliminating wage risk while preserving expected wages and the tax schedule. It measures the average loss in expected utility that households incur due to wage risk, relative to a hypothetical life without such risk in an otherwise identical economic and policy environment.

The *insurance value of joint taxation* is the reduction in that welfare loss relative to separate taxation:

$$\mathcal{I}^\Sigma(e) = \mathcal{L}_0^\Sigma(\chi_0^\Sigma, e) - \mathcal{L}_1^\Sigma(\chi_1^\Sigma, e). \quad (21)$$

The comparison holds the environment e fixed and applies the same revenue target to both tax systems, with χ_0^Σ and χ_1^Σ denoting the corresponding tax-scale parameters.

We report welfare differences both in lifetime-utility units and as consumption equivalents relative to a common reference: the consumption plan in the calibrated joint-taxation economy with baseline wage risk. Appendix D provides the formal definition.

Table 5 reports the welfare losses from wage risk and the insurance value of joint taxation. Wage risk generates substantial welfare losses under both tax systems, but these losses are smaller under joint taxation. Eliminating wage risk is worth 17.9 percent of baseline consumption under joint taxation and 20.1 percent under separate taxation. The resulting insurance value of joint taxation amounts to 1.61 percent of baseline consumption.

⁸See Appendix D for a formal welfare definition.

Joint taxation thus eliminates a sizable share of the welfare cost of wage risk—about 9 percent when measured in utility units.

Importantly, this insurance advantage does not arise from a smaller consumption response to the male permanent wage shock: in fact, Table 4 shows a slightly larger response under joint taxation. Instead, the pronounced differences in spouses’ labor-supply responses point to labor-supply adjustment as the central mechanism behind the insurance advantage. The experiments below confirm this interpretation, showing that households’ ability to adjust their labor supply accounts for most of the insurance advantage. The quantitative magnitude of this effect is noteworthy. A change in the tax treatment of couples that modifies households’ adjustment incentives reduces the welfare cost of wage risk by almost one tenth, despite barely affecting the pass-through of wage shocks to consumption. Thus, a substantial part of the welfare consequences of wage risk depends on how households self-insure through behavioral adjustments, rather than only on the degree to which consumption is ultimately shielded from wage shocks.

5.3 Implications for welfare

The welfare consequences of replacing separate with joint taxation are not limited to its insurance value. We define the *welfare value of joint taxation* as

$$\mathcal{W}^\Sigma(e) = \mathcal{V}_1^\Sigma(\chi_1^\Sigma, e) - \mathcal{V}_0^\Sigma(\chi_0^\Sigma, e). \quad (22)$$

Table 5: Welfare loss from wage risk and the insurance value of joint taxation

	Utility units	CEV (percent)
Welfare loss, separate taxation $\mathcal{L}_0$	7.960	20.07
Welfare loss, joint taxation $\mathcal{L}_1$	7.220	17.92
Insurance value $\mathcal{I} = \mathcal{L}_0 - \mathcal{L}_1$	0.740	1.61

Notes: Estimated progressivity and wage risk. Risk economies raise the same income-tax revenue; each no-risk counterpart retains the tax schedule of its risk economy. The subtraction indicated by the horizontal rule applies to utility units only. Each CEV is a separate transformation relative to the common calibrated joint-taxation consumption plan and is therefore not obtained as the difference between the CEVs in the two rows above it.

The difference between the welfare and insurance values of joint taxation is the welfare difference between the two no-risk economies. Accordingly, we decompose

$$\mathcal{W}^\Sigma(e) = \mathcal{I}^\Sigma(e) + \mathcal{N}^\Sigma(e), \quad (23)$$

where

$$\mathcal{N}^\Sigma(e) = \mathcal{V}_1^0(\chi_1^\Sigma, e) - \mathcal{V}_0^0(\chi_0^\Sigma, e) \quad (24)$$

collects the non-insurance component of the welfare comparison. This component compares two no-risk economies that differ both in whether spouses are taxed jointly or separately and in the tax scale, χ . The tax scales are inherited from the corresponding risky economies, where they are constructed such that joint taxation with χ_1^Σ and separate taxation with χ_0^Σ generate the same total tax revenue. Once wage risk is removed, however, these same tax schedules generally imply different tax burdens and generate different revenues.

To separate these two differences, let χ_θ^0 denote the tax scale that, in the absence of wage risk, generates the same revenue as the baseline economy with joint taxation and baseline risk. The equal-revenue no-risk welfare difference is

$$\mathcal{E}(e) = \mathcal{V}_1^0(\chi_1^0, e) - \mathcal{V}_0^0(\chi_0^0, e). \quad (25)$$

We can then decompose the non-insurance component as

$$\mathcal{N}^\Sigma(e) = \mathcal{E}(e) + \mathcal{F}^\Sigma(e), \quad (26)$$

which defines $\mathcal{F}^\Sigma$ as the fiscal effect of applying tax scales inherited from the risky economies.

The equal-revenue component $\mathcal{E}$ compares joint and separate taxation without wage risk and at equal revenue. It contains several effects, including distortions to hours and participation decisions and endogenous saving. The fiscal component $\mathcal{F}$ connects this comparison to the no-risk comparison implicitly embedded in $\mathcal{I}$, in which tax schedules are inherited from the risky economies and revenues therefore need not be equal. It

captures the welfare consequences of the resulting differences in resources transferred from households to the government, including households' endogenous responses to these resource differences.

Table 6 shows the welfare value of joint taxation and its decomposition under the baseline parameterization. Joint taxation lowers welfare by the equivalent of 0.74 percent of baseline consumption. Thus, the insurance value of joint taxation of 1.61 percent documented above is more than offset by the non-insurance component of the welfare comparison. The decomposition of this component shows that it is overwhelmingly fiscal: the fiscal component amounts to -2.21 percent, whereas the equal-revenue component is only -0.08 percent. The fiscal component therefore accounts for nearly all of the negative non-insurance component.

The large fiscal component arises because tax revenue responds more strongly to wage risk under separate than under joint taxation. Removing wage risk reduces income-tax revenue by about 20 percent under joint taxation and 34 percent under separate taxation, see Table 10 in Appendix E. With progressive taxation, greater dispersion in taxable income raises average tax payments. Separate taxation applies progressivity to each spouse's income individually, whereas joint taxation partly pools earnings across spouses before progressivity is applied. Consequently, the wage inequality generated by risk raises

Table 6: The welfare value of joint taxation and its decomposition

	Utility units	CEV (percent)
Welfare value $\mathcal{W}$	-0.347	-0.74
Insurance value $\mathcal{I}$	0.740	1.61
Non-insurance component $\mathcal{N}$	-1.087	-2.29
Fiscal component $\mathcal{F}$	-1.049	-2.21
Residual $\mathcal{E}$	-0.038	-0.08

Notes: Positive values favor joint taxation. All comparisons use the estimated progressivity and expected wage profiles. The risk economies and the independently revenue-matched no-risk benchmarks raise baseline income-tax revenue. The decomposition is additive in utility units, up to rounding, but not in the separately transformed CEV column.

revenue more under separate taxation. Endogenous labor-supply responses reinforce this revenue difference.⁹

The stronger revenue response under separate taxation implies a fiscal advantage in an economy with wage risk. Under a common government revenue requirement, the resulting tax adjustments favor separate taxation in our model. The negative fiscal component captures the importance of this advantage for the welfare comparison.

This fiscal advantage accounts for the welfare advantage of separate taxation in the baseline model. The insurance value of joint taxation largely offsets this advantage, but does not eliminate it in the baseline economic and policy environment. Other differences between joint and separate taxation, including their differential distortions of labor supply within the household, have a quantitatively small net welfare effect.

5.4 The importance of labor-supply adjustments

The sizable differences in labor-supply responses documented above point to labor-supply adjustments as a potentially important determinant of the insurance effects of joint taxation. To quantify their importance, we shut off these adjustment margins selectively while leaving the economic and policy environment unchanged. When a labor-supply margin is shut off, it is fixed at its age-specific level in the corresponding no-risk economy, so that wage shocks cannot induce adjustment along that margin.¹⁰

We consider four cases: both hours and participation can adjust as in the baseline model; both are fixed; only hours can adjust; and only female participation can adjust.¹¹

Table 7 reports the insurance and welfare values of joint taxation for these four cases. Allowing both labor-supply margins to adjust (first line) reproduces the baseline model.

⁹Appendix E decomposes the revenue asymmetry into the direct effect of wage dispersion and the additional effect of endogenous labor-supply adjustment.

¹⁰For hours, this fixes the age-specific work schedule. In the deterministic reference allocations, female participation equals one at every working age. Fixing participation therefore requires the wife to work throughout the working life, rather than imposing a fixed aggregate participation rate.

¹¹When participation is adjustable but hours are fixed, the household chooses between the wife not working and her supplying the age-specific female hours from the corresponding no-risk economy. Male hours remain fixed at their no-risk levels. Consumption and saving remain endogenous throughout.

Table 7: Labor-supply adjustment: insurance and welfare values of joint taxation

Hours	Participation	$\mathcal{I}$ (% CEV)	$\mathcal{W}$ (% CEV)
Adjustable	Adjustable	1.61	−0.74
Fixed	Fixed	0.27	−2.04
Adjustable	Fixed	1.49	−0.86
Fixed	Adjustable	0.65	−1.67

Notes: CEVs are relative to common baseline consumption; positive values favor joint taxation. Hours refer to both spouses; participation refers to the wife. Fixed choices follow their age-specific levels in the corresponding no-risk economy. Tax schedules are inherited from the unrestricted economies and held fixed across adjustment scenarios; income-tax revenue is therefore not equalized across labor-supply restrictions.

The results show that the insurance value of joint taxation depends strongly on households' ability to adjust labor supply in response to wage shocks. When both labor-supply margins are fixed (second line), $\mathcal{I}$ falls from 1.61 to 0.27 percent of baseline consumption, a decline of 1.35 percentage points. Thus, more than four fifths of the baseline insurance value disappears when households cannot adjust labor supply to wage shocks. Much of the insurance value of joint taxation therefore arises from how joint rather than separate taxation affects households' labor-supply responses to wage shocks. Correspondingly, the welfare value falls by 1.29 percentage points, from −0.74 to −2.04 percent, so that the 1.35-percentage-point decline in the insurance value translates almost one-for-one into a declining welfare value in CEV terms.

Even with both labor-supply margins fixed, joint taxation retains a positive insurance value of 0.27 percent of baseline consumption. Labor-supply adjustment therefore amplifies, but is not necessary for, its insurance advantage. With fixed labor supply, joint taxation still links spouses' marginal tax rates: a negative wage shock to the secondary earner, for example, reduces the tax burden on the primary earner and thereby cushions its effect on household net income. Thus, beyond the larger insurance channel operating through labor supply, joint taxation provides insurance through how tax payments respond to spouses' wage realizations along the tax schedule.

The lower two lines of Table 7 show that most of the effect of labor-supply adjustability operates through hours rather than participation. Allowing hours to adjust raises $\mathcal{I}$ by

between 0.96 and 1.22 percentage points, depending on whether participation is adjustable. By contrast, the corresponding contribution of participation ranges only from 0.13 to 0.38 percentage points. There is some negative interaction between the margins, with the effect of having one option being larger when the other remains shut off. The quantitative importance of hours compared to participation also links these results closely to the analytical mechanism in Section 3, which considers the intensive-margin responses of two working spouses.

Labor-supply adjustment also contributes to the fiscal difference between joint and separate taxation. Appendix E uses a fixed-work counterfactual similar to the experiments above to separate this contribution from the direct effect of wage dispersion. Rather than restricting households' adjustment options, it evaluates tax revenue at risky wage realizations while holding work choices at their no-risk levels and compares it with revenue under households' optimal risky-economy work choices. A little more than half of the differential revenue response under joint and separate taxation arises directly from wage dispersion, while a little less than half results from endogenous labor-supply adjustment. Thus, the fiscal difference is present even without adjustment of labor supply, but such adjustment substantially amplifies it.

5.5 The role of the risk magnitude

Greater wage risk increases the potential value of insurance, but also affects the fiscal comparison between joint and separate taxation. To examine how these effects depend on the magnitude of risk, we scale the standard deviations of wage innovations for both spouses by a common factor q . We conduct three distinct experiments, varying all wage risk, permanent risk alone, or transitory risk alone.

As in the preceding analysis, joint and separate taxation are compared at equal income-tax revenue. At each value of q , joint taxation retains the estimated tax schedule, while the tax scale under separate taxation is adjusted to match the resulting joint-taxation revenue. The corresponding no-risk economies inherit the tax schedules of their risky counterparts. This approach continues our focus on isolating the mechanisms. To also


assess changes in wage risk as an economic fundamental, Appendix F considers a fixed government revenue constraint across risk magnitudes, such that tax scales need to be adjusted to changes in fundamentals. The results for $\mathcal{I}$ and $\mathcal{W}$ remain largely unchanged.

Figure 1 shows how the welfare components respond to changes in wage risk. The insurance value of joint taxation increases substantially with the magnitude of risk. When all wage innovation standard deviations are halved ($q = 0.5$), $\mathcal{I}$ falls from 1.61 to about 0.6 percent of baseline consumption; increasing them by 50 percent ($q = 1.5$) raises it to about 2.3 percent. This sensitivity is driven primarily by permanent risk: varying permanent risk alone produces a profile close to that obtained when all risk is scaled, whereas varying transitory risk has a much smaller effect. Persistent wage shocks make private self-insurance costly over a longer horizon, increasing the insurance value of household labor adjustment.

At high levels of permanent or combined risk, the insurance value peaks around $q = 2$ and then declines. At this magnitude, innovation standard deviations are doubled and their variances quadrupled, generating much greater accumulated wage dispersion. Extreme dispersion increasingly places households in states with either negligible taxable earnings or very low marginal utility of consumption, which may explain the declining additional insurance value of joint taxation. Around the calibrated economy ($q = 1$), the relationship is already concave but retains a substantial positive slope. Thus, greater exposure to persistent wage risk strengthens the insurance role of joint taxation in the empirically relevant neighborhood of the estimated process.

The overall welfare comparison between joint and separate taxation moves in the opposite direction. As (persistent) wage risk increases, $\mathcal{W}$ declines, even though the insurance value generally rises. Greater wage risk also implies more pronounced wage dispersion, which in turn amplifies the revenue advantage of separate taxation discussed in Section 5.3. The rising insurance value of joint taxation counteracts this fiscal consequence of greater wage risk to a large degree, but not completely. At $q = 2$ in the all-risk experiment, $\mathcal{I}$ reaches 2.54 percent of baseline consumption, while $\mathcal{W}$ falls to -2.07 percent. Both effects are driven primarily by permanent risk. Thus, greater wage risk

Figure 1: Wage risk: insurance and welfare values of joint taxation



Notes: Values are consumption-equivalent variations (CEVs), in percent of common baseline consumption; positive values favor joint taxation. q scales both spouses' innovation standard deviations for the indicated risk component; $q = 1$ corresponds to the estimated process. In panels (b) and (c), the other risk component remains at its estimated level. The joint-taxation tax scale remains at its calibrated value, while the tax scale under separate taxation is adjusted at each q so that joint and separate taxation raise equal income-tax revenue by construction. All other model parameters remain fixed.

generally strengthens the insurance case for joint taxation, even as separate taxation becomes increasingly favored in the overall welfare comparison.

5.6 The role of intra-household earnings inequality

Spouses' relative earnings capacities shape their ability to insure each other through labor-supply adjustments. To examine how the insurance mechanism depends on intra-household earnings inequality, we vary how similar spouses are in their earnings capacities. Specifically, we scale the wife's expected wage at every age by a common factor q_{wf} . Hours and participation adjust endogenously, determining how changes in relative wage opportunities translate into the wife's share of household earnings.

As in the preceding experiments, the joint-taxation tax schedule remains fixed, while the separate-taxation tax scale is adjusted at each q_{wf} to match joint-taxation revenue. The corresponding no-risk economies inherit the tax schedules of their risky counterparts. Appendix F additionally examines the results under a fixed government revenue requirement across values of q_{wf} .

The variation in expected wages translates into substantial changes in actual intra-household earnings inequality. As wives' expected wages increase from 40 percent below to 50 percent above their estimated levels, wives' endogenous share of total household earnings rises monotonically from about 21–22 percent to about 47 percent under joint and separate taxation. At baseline, the shares are 34.4 percent under joint taxation and 35.1 percent under separate taxation. Figure 2 shows that this convergence in spouses' earnings capacities strengthens the insurance value of joint taxation: $\mathcal{I}$ rises from about 0.56 to 1.90 percent of baseline consumption, with the steepest gains occurring when the wife's earnings capacity is initially relatively low. As spouses become more similar in their potential earnings, labor-supply adjustments provide more effective insurance, making joint taxation's encouragement of such adjustments increasingly valuable.


This contrasts with the familiar tax-saving advantage of income splitting under a common underlying individual tax schedule, which is largest for couples with unequal earnings and vanishes at equal earnings. A narrowing gender earnings gap can thus strengthen the insurance role of joint taxation while reducing its conventional tax-saving advantage. Equal earnings capacities do not necessarily maximize the insurance value, however, as spouses still differ in preferences, participation, and life-cycle wage profiles.

Unlike in the permanent-risk experiment, $\mathcal{I}$ and $\mathcal{W}$ move in the same direction as intra-household earnings inequality declines. The overall welfare value of joint taxation rises from -1.07 to -0.50 percent of baseline consumption, largely reflecting the substantial increase in its insurance value. Nevertheless, joint taxation remains welfare-inferior to separate taxation throughout the experiment. Thus, while relative earnings capacities materially affect the insurance role of joint taxation, the resulting gains do not reverse the overall welfare comparison.

5.7 Progressivity and the tax unit under optimal taxation

The preceding analysis has examined the welfare consequences of joint vs. separate taxation under different economic conditions. We now turn to broader tax reforms that change not only whether spouses are taxed jointly or separately, but also the progressivity of

Figure 2: Intra-household earnings inequality: insurance and welfare values of joint taxation



Notes: Values are consumption-equivalent variations (CEVs), in percent of common baseline consumption; positive values favor joint taxation. q_{w_f} scales the wife's expected wage profile; $q_{w_f} = 1$ corresponds to the baseline. The joint-taxation tax scale remains at its calibrated value, while the tax scale under separate taxation is adjusted at each q_{w_f} so that both tax units raise equal income-tax revenue by construction. All other model parameters remain fixed.

the tax schedule. We first examine how the insurance and overall welfare values of joint taxation depend on progressivity. We then compare the best joint-taxation system to the best separate-taxation system when progressivity can be chosen separately.

For the first exercise, we vary progressivity μ under both joint and separate taxation. At each value of μ , we adjust the tax scales under joint and separate taxation so that both risky economies raise the income-tax revenue of the calibrated baseline. The corresponding no-risk economies inherit the tax schedules of their risky counterparts. All other model parameters remain fixed.

Figure 3 shows the insurance and welfare values of joint taxation as a function of tax progressivity. The insurance value rises strongly with progressivity. Under a linear tax ($\mu = 0$), the tax unit is irrelevant and both $\mathcal{I}$ and $\mathcal{W}$ are zero. At estimated progressivity (vertical line), the insurance value is 1.61 percent of baseline consumption; at $\mu = 0.6$, it reaches 4.87 percent. This pattern is consistent with the labor-adjustment mechanism: greater progressivity strengthens the connection between one spouse's earnings and the partner's marginal tax rate under joint taxation. Figure 8 in Appendix G shows that

the associated differences in labor responses also widen with progressivity. Notably, progressivity strengthens the partner’s participation response under joint taxation while weakening it under separate taxation, where pronounced progressivity already insures the affected spouse’s net income strongly against the wage shock.

The overall welfare comparison does not move one-for-one with the insurance value. Separate taxation remains preferable near the baseline, and its advantage initially widens. At higher progressivity, the growing insurance value comes to outweigh the other effects. The welfare value turns positive around $\mu = 0.35$. Joint taxation can therefore be welfare-preferred at sufficiently high tax progressivity.


Welfare-maximizing progressivity. The preceding analysis compares joint and separate taxation at a common degree of progressivity. A broader reform perspective allows progressivity to change as well, with welfare-maximizing tax schedules chosen separately under joint and separate taxation. How much does the choice between joint and separate taxation matter in this case?

Figure 4 shows welfare under both tax units as progressivity varies, with tax scales adjusted to maintain the common baseline revenue requirement. Both welfare profiles exhibit interior maxima, with joint taxation reaching its maximum at somewhat higher progressivity than separate taxation. Remarkably, the two maxima are very close in welfare terms.

Comparing the welfare profiles with Figure 3 helps explain this near-equality. The maxima occur in a region where the overall welfare value of joint taxation is small, despite its substantial insurance value. The insurance mechanism largely offsets the combined welfare disadvantage arising from the other effects of joint taxation.


Table 8 summarizes the numerical results and contrasts the welfare comparison at calibrated progressivity with that between the separately optimized tax systems. Welfare-maximizing progressivity is about 36.4 percent under joint taxation and 32.1 percent under separate taxation. The resulting welfare difference is only about -0.04 percent of baseline consumption.

Figure 3: Tax progressivity: insurance and welfare values of joint taxation



Notes: Values are consumption-equivalent variations (CEVs), in percent of common baseline consumption; positive values favor joint taxation. μ denotes tax progressivity; the vertical line marks its calibrated value. At each μ , the tax scales under both joint and separate taxation are adjusted to maintain the common baseline level of income-tax revenue. All other model parameters remain fixed.

Figure 4: Welfare under joint and separate taxation across progressivity



Notes: Welfare under joint and separate taxation, expressed as differences to the baseline economy in consumption-equivalent variations (CEVs), in percent of common baseline consumption. Both profiles refer to economies with wage risk. The vertical line marks calibrated progressivity; markers indicate the respective welfare maxima. At each μ , the tax scales under both joint and separate taxation are adjusted to maintain the common baseline level of income-tax revenue. All other model parameters remain fixed.

The distinction between the two policy perspectives is consequential. At calibrated progressivity, joint taxation provides substantial insurance, but separate taxation remains welfare-preferred. When progressivity can also be reformed, however, the choice of tax unit has remarkably little effect on attainable welfare. The insurance value of joint taxation is central to this result, almost entirely offsetting its other welfare disadvantages.

Table 8: The welfare effect of the tax unit at estimated and optimal progressivity

	Progressivity μ		Welfare difference
	Joint	Separate	Joint – Separate
Estimated progressivity	0.141	0.141	–0.742
Optimal progressivity	0.364	0.321	–0.036

Notes: The last column reports the CEV of the welfare difference between joint and separate taxation, in percent of common baseline consumption; negative values favor separate taxation. In the first row, both tax systems use the estimated degree of progressivity. In the second row, progressivity is chosen separately to maximize welfare under each tax system. Every economy raises baseline income-tax revenue. Formally, the optimal progressivity levels are $\mu_{\theta}^* = \arg \max_{\mu} \mathcal{V}_{\theta}^{\Sigma}(\chi_{\theta}^{\Sigma}(e(\mu)), e(\mu))$, where $e(\mu)$ denotes the baseline economic and policy environment with tax progressivity set to μ , holding all other components of e fixed at their baseline values, and $\chi_{\theta}^{\Sigma}(e(\mu))$ is the tax-scale parameter that maintains the common baseline level of income-tax revenue.

6 Conclusion

This paper studies the insurance value of joint taxation in the presence of incomplete financial markets and endogenous labor-supply responses to wage shocks. The standard case for separate taxation emphasizes efficiency in a deterministic context: when spouses differ in labor-supply elasticities, the more responsive earner should face a lower marginal tax rate. We show that this argument leaves out an important role of the tax unit: joint taxation affects the welfare costs of risk by shaping how spouses adjust their labor supply to wage shocks.

Whether spouses are taxed jointly or separately changes how a wage shock to one spouse affects marginal tax rates within the household. Under joint taxation, the shock changes the marginal tax rates faced by both spouses and thereby encourages adjustment through the partner’s labor supply. Because the partner is not directly exposed to the

same wage shock, shifting adjustment toward the partner can provide a less costly margin of self-insurance.

We quantify this mechanism in a life-cycle model of married couples with incomplete markets, wage risk, as well as endogenous labor supply at the intensive and extensive margins, calibrated to the U.S. economy. We find that the insurance value of joint taxation equals 1.61 percent of baseline consumption. It generally increases with greater wage risk, more similar earnings capacities between spouses, and greater tax progressivity. Restricting options to adjust labor supply in response to shocks confirms that this adjustment margin is central to the insurance value of joint taxation.

The insurance value, however, does not directly determine the overall welfare comparison between joint and separate taxation. Wage dispersion affects tax revenues differently under joint and separate taxation. Under a common government revenue requirement, the resulting tax-scale adjustments affect distortions, with endogenous labor-supply responses reinforcing this effect. At current tax progressivity, this fiscal advantage of separate taxation more than offsets the insurance value of joint taxation, so that separate taxation yields higher overall welfare. Labor-market risk therefore strengthens the insurance case for joint taxation while, in the current economic and policy environment, making its overall welfare comparison less favorable.

The perspective changes when we move from changing the tax unit at a given degree of progressivity to broader income-tax reform. Greater progressivity strengthens the insurance value of joint taxation and can eventually make joint taxation welfare-preferred. More importantly, when progressivity is optimized separately under joint and separate taxation, the two tax units deliver essentially identical welfare. The insurance advantage of joint taxation then virtually offsets its disadvantages.

References

- Alesina, A., A. Ichino, and L. Karabarbounis (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy* 3(2), 1–40.
- Apps, P. F. and R. Rees (1999). Individual versus joint taxation in models with household production. *Journal of Political Economy* 107(2), 393–403.
- Asdrubali, P., S. Tedeschi, and L. Ventura (2020). Household risk-sharing channels. *Quantitative Economics* 11(3), 1109–1142.
- Attanasio, O., H. Low, and V. Sánchez-Marcos (2005). Female labor supply as insurance against idiosyncratic risk. *Journal of the European Economic Association* 3(2-3), 755–764.
- Benabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica* 70(2), 481–517.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy* 126(S1), S73–S115.
- Boskin, M. J. and E. Sheshinski (1983). Optimal tax treatment of the family: Married couples. *Journal of Public Economics* 20(3), 281–297.
- Corneo, G. (2013). A note on the taxation of couples under income uncertainty. *Finanz-Archiv/Public Finance Analysis* 60(1), 129–134.
- Daminato, C. and L. Pistaferri (2020). Family labor supply and asset returns. *European Economic Review* 124, 103389.
- Gayle, G.-L. and A. Shephard (2019). Optimal taxation, marriage, home production, and family labor supply. *Econometrica* 87(1), 291–326.
- Gottfried, P. and W. F. Richter (1999). The tax unit and household production: Comment. *Journal of Political Economy* 107(2), 404–409.


- Guner, N., R. Kaygusuz, and G. Ventura (2012a). Taxation and household labour supply. *Review of Economic Studies* 79(3), 1113–1149.
- Guner, N., R. Kaygusuz, and G. Ventura (2012b). Taxing women: A macroeconomic analysis. *Journal of Monetary Economics* 59(1), 111–128.
- Guner, N., R. Kaygusuz, and G. Ventura (2014). Income taxation of US households: Facts and parametric estimates. *Review of Economic Dynamics* 17(4), 559–581.
- Heathcote, J., K. Storesletten, and G. L. Violante (2010). The macroeconomic implications of rising wage inequality in the United States. *Journal of Political Economy* 118(4), 681–722.
- Holter, H. A., D. Krueger, and S. Stepanchuk (2019). How do tax progressivity and household heterogeneity affect Laffer curves? *Quantitative Economics* 10(4), 1317–1356.
- Holter, H. A., D. Krueger, and S. Stepanchuk (2023). Till the IRS do us part: (Optimal) taxation of households. Working paper.
- Meier, V. and H. Rainer (2015). Pigou meets Ramsey: Gender-based taxation with non-cooperative couples. *European Economic Review* 77, 28–46.
- Ortigueira, S. and N. Siassi (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics* 60(6), 650–666.
- Piggott, J. and J. Whalley (1996). The tax unit and household production. *Journal of Political Economy* 104(2), 398–418.
- Piggott, J. and J. Whalley (1999). Reply to Apps/Rees and Gottfried/Richter. *Journal of Political Economy* 107(2), 410–418.
- Rupert, P. and G. Zanella (2015). Revisiting wage, earnings, and hours profiles. *Journal of Monetary Economics* 72, 114–130.
- Stephens, M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics* 20(3), 504–537.

- Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters* 20(2), 177–181.
- Theloudis, A. (2021). Consumption inequality across heterogeneous families. *European Economic Review* 136, 103765.
- Wu, C. and D. Krueger (2021). Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. *American Economic Journal: Macroeconomics* 13(1), 79–113.

Appendix A Expected life-cycle wage profiles

We estimate expected life-cycle wage profiles separately for married men and women. We apply the common married-couple sample restrictions, extend the age range to couples with both spouses aged 21–65, and use positive observed hourly wages for each gender. Following Rupert and Zanella (2015), we fit second-order fractional polynomials to PSID-family-weighted mean real hourly wages by age. Both profiles are normalized by the arithmetic mean of the predicted male profile over ages 21–65. Figure 5 shows the resulting profiles.

Figure 5: Normalized expected life-cycle wage profiles



Notes: Expected hourly wages of married men and women, normalized by the mean predicted male wage over ages 21–65.

Appendix B Estimation of the stochastic wage process

Our estimation follows the permanent–transitory approach widely used in the consumption–insurance literature, including Blundell et al. (2016). We use this distinction to estimate the two innovation covariance matrices from the second moments of husbands’ and wives’ residual wage growth.

We use the married-couple sample described in Section 4, with both spouses aged 30–57 in the income year. The 1999–2019 PSID interviews provide earnings and hours

for 1998–2018. Real hourly wages are annual labor income divided by annual hours and the price deflator. We construct two-year log-wage changes within stable couple spells, requiring usable observations at both endpoints and excluding observations flagged by sharp wage or earnings reversals. Female wage changes are observed only when the wife works at both endpoints; a participation-selection correction leaves the estimates virtually unchanged, so we report the uncorrected results.

Separately for men and women, we regress these changes on year and birth-year indicators, education, race, state, large-city residence, and changes in large-city residence. Education, race, and the city variables are also interacted with year. Let d_{gjt} denote the resulting residual for spouse g in couple j , where t indexes calendar income years. Thus d_{gjt} measures residual growth from $t - 2$ to t , not annual growth.

In addition to the permanent and transitory components of the model, observed log wages contain measurement error ν_{gjt} with variance $\sigma_{\nu_g}^2$. We treat this error as serially uncorrelated, uncorrelated between spouses, and orthogonal to the underlying wage process. Its variance is set to an externally specified fraction of each gender’s observed log-wage variance, using the validation-study approach discussed in the main text. It is not separately estimated from the growth moments.

Under the random-walk permanent component and serially independent transitory shocks, the seven moment restrictions are, suppressing the couple index,

$$\mathbb{E}[d_{gt}^2] = 2\sigma_{\varepsilon_g}^2 + 2\sigma_{\epsilon_g}^2 + 2\sigma_{\nu_g}^2, \quad g \in \{m, f\}, \quad (27)$$

$$\mathbb{E}[d_{gt}d_{g,t-2}] = -\sigma_{\varepsilon_g}^2 - \sigma_{\nu_g}^2, \quad g \in \{m, f\}, \quad (28)$$

$$\mathbb{E}[d_{mt}d_{ft}] = 2\sigma_{\varepsilon_m, \varepsilon_f} + 2\sigma_{\epsilon_m, \epsilon_f}, \quad (29)$$

$$\mathbb{E}[d_{mt}d_{f,t-2}] = \mathbb{E}[d_{ft}d_{m,t-2}] = -\sigma_{\epsilon_m, \epsilon_f}. \quad (30)$$

The two equalities in (30) jointly describe the cross-spouse lag structure but constitute two separate empirical moments. A two-year change contains two annual permanent innovations but only the difference between its endpoint transitory shocks. After accounting for measurement error, the negative own lag moment identifies transitory variation; the

remaining growth variance identifies permanent risk. The two directional cross-lag moments identify the transitory covariance between spouses, while their contemporaneous moment also contains permanent covariance. The permanent variances and covariance are annual innovation moments; transitory moments describe the wage disturbance at a given date.

We estimate the six parameters jointly by one-step GMM with an identity weighting matrix. Both the first-stage regressions and GMM are unweighted. Each moment uses all observations for which its required residuals are available, rather than imposing a common estimation sample. Standard errors are clustered by stable couple spell. The resulting estimates are included in the parameter overview in Table 1.

Implementation in the quantitative model. The estimated wage processes are discretized using the method of Tauchen (1986) and implemented through discrete Markov transitions. The resulting innovation distributions are applied from the first working period onward. Permanent wage components are initialized at zero immediately before the first draw and subsequently accumulate innovations. First-period wage dispersion therefore arises from the initial realizations of permanent and transitory innovations, without introducing additional ex ante heterogeneity.

Expected wage profiles and wage risk are distinct inputs. For each age and gender, the deterministic log-wage component is set to

$$\ell_g(h; w_g, \Sigma) = \log w_g(h) - \log \mathbb{E}[\exp\{z_g(h) + \epsilon_g(h)\}]. \quad (31)$$

This ensures that $\mathbb{E}[w_{gj}(h)] = w_g(h)$. The expectation is evaluated over the discretized wage process, including the first-period innovations, rather than using the continuous-normal approximation. Removing or scaling wage risk thus changes the deterministic log-wage component while preserving the expected wage profile. This construction allows wage risk and expected earnings capacity to vary independently.

Appendix C Responses to a female permanent wage shock

Table 9 extends the baseline comparison to a permanent increase in the wife’s wage. The wife increases her labor supply through higher participation and longer hours when working, while the husband reduces his hours. All three responses are stronger under joint taxation, even though consumption responds almost identically under joint and separate taxation. Thus, the contrast documented for permanent shocks to the husband’s wage also holds for shocks to the wife’s wage: similar consumption insurance can be achieved through different adjustments in spouses’ work. Consumption responses alone therefore give an incomplete picture of how the tax unit shapes household self-insurance.

Table 9: Responses to a one-percent permanent increase in the female wage

	Joint taxation	Separate taxation
Household consumption	0.203	0.207
Male hours	−0.219	−0.174
Female hours, conditional on participation	0.417	0.338
Female participation	0.167	0.087
Household gross labor income	0.410	0.406
Household net labor income	0.352	0.360

Notes: Responses are in percent, except female participation, which is in percentage points. Progressivity equals its estimated value; income-tax revenue is matched across joint and separate taxation. We regress changes in log outcomes, or the change in the participation indicator, jointly on the changes in the permanent and transitory log-wage components of both spouses after removing age profiles. Female log-hours responses use women with positive hours in both adjacent periods. Net labor income is gross labor income minus net income-tax liabilities, including transfers represented by the tax function, before social security contributions. Coefficients are normalized to a one-percent female permanent wage innovation.

Appendix D Welfare definition and consumption-equivalent variation

Welfare is defined as expected lifetime utility, evaluated before the first wage innovations are realized. Formally, conditional on θ , Σ , χ , and an economic and policy environment e , it equals

$$\mathcal{V}_{\theta}^{\Sigma}(\chi, e) = \mathbb{E}[V(0, \varepsilon_m, \varepsilon_f, \epsilon_m, \epsilon_f, \underline{h}) \mid \theta, \Sigma, \chi, e]. \quad (32)$$

Consumption equivalence. Let $\mathbf{c}^{\text{base}}$ and $\mathbf{n}_g^{\text{base}}$ be the contingent consumption and labor plans of the calibrated joint-taxation risk economy. For a utility difference Δ , the proportional consumption equivalent solves

$$\begin{aligned} \sum_{h=\underline{h}}^{\bar{h}} \beta^{h-\underline{h}} \mathbb{E} \left[v \left((1 + \text{CEV}(\Delta)) \mathbf{c}^{\text{base}}(h), \mathbf{n}_m^{\text{base}}(h), \mathbf{n}_f^{\text{base}}(h) \right) \right] \\ = \sum_{h=\underline{h}}^{\bar{h}} \beta^{h-\underline{h}} \mathbb{E} \left[v \left(\mathbf{c}^{\text{base}}(h), \mathbf{n}_m^{\text{base}}(h), \mathbf{n}_f^{\text{base}}(h) \right) \right] + \Delta. \end{aligned} \quad (33)$$

The reference plans are unchanged across all main-model comparisons, including no-risk environments.

Appendix E Wage risk and tax revenue

This appendix documents the revenue asymmetry discussed in Section 5.3 and examines the contributions of wage dispersion and labor adjustment to this asymmetry.

Table 10 documents the income-tax revenue effects of wage risk under joint and separate taxation. The two risky economies raise the same revenue by construction. The table uses this revenue level as the common normalization (100 for both tax units). Removing wage risk while retaining the respective risky economy's tax schedule reduces revenue by about 20% under joint taxation and by roughly 34% under separate taxation. That wage risk raises revenue more strongly under separate taxation generates the fiscal component $\mathcal{F}$ of the welfare comparison between joint and separate taxation discussed in Section 5.3: under a common revenue requirement, the greater revenue-raising capacity of separate taxation affects household welfare through the tax burden required to raise that revenue.

Table 10: Income-tax revenue with and without wage risk

	Joint taxation	Separate taxation	Difference (JT–ST)
With wage risk	100.00	100.00	0.00
Without wage risk	80.00	65.82	14.19
Change (No risk - risk)	−20.00	−34.18	14.19

Notes: Income-tax revenue is normalized by the common revenue of the two risky economies (100). Each no-risk economy retains the tax schedule of its risky counterpart. Revenue is discounted at the market interest rate over working ages; social security contributions are excluded.

Wage risk affects revenue through both the distribution of wages and households' labor-supply choices. To distinguish these sources of the revenue effect, we first calculate tax revenues at the risky wage realizations while holding labor choices at their corresponding no-risk allocations. In this fixed-work counterfactual, wives participate at every working age, as in the corresponding no-risk economy, while both spouses' hours are held at their age- and gender-specific no-risk levels. Comparing revenues in these fixed-work economies with those in their no-risk counterparts isolates the direct wage effect. The additional change when households use their optimal risky-economy work choices captures labor adjustment, including its interaction with wage realizations. Table 11 reports this decomposition. The table includes a third component because average simulated wages differ slightly from the expected wages held fixed in these experiments. We report this effect separately so that it is not attributed to wage dispersion. Of the 14 percentage-point difference in revenue reductions between joint and separate taxation, a bit more than half arises directly from wages and a little less than half from work adjustment (including rounding). Hence the fiscal asymmetry persists at fixed work and is amplified by endogenous labor supply.

Table 11: Sources of the revenue reduction from removing wage risk

	Joint taxation	Separate taxation	Difference (JT-ST)
Total revenue reduction	20.00	34.18	-14.19
Direct wage dispersion	15.08	22.65	-7.56
Simulated-mean remainder	0.48	0.45	0.03
Work adjustment	4.43	11.08	-6.65

Notes: Baseline wage risk. Entries are percentages of the common revenue of the risky economies. The simulated-mean remainder accounts for small deviations of finite simulated wage paths from their theoretical age- and gender-specific means. Components sum to the total up to rounding.

Appendix F Fiscal adjustment to changes in wage distributions

The variations in wage risk and intra-household earnings inequality in Sections 5.5 and 5.6 help understand the mechanisms underlying the welfare comparison between joint and separate taxation. They can also be interpreted as changes in economic fundamentals: what happens if wages become more risky, or if spouses become more similar in their earnings capacities? From this perspective, potential government revenue requirements

become relevant. Rather than allowing tax revenue to change with economic fundamentals, the government may have to adjust tax schedules to maintain a given revenue target.

We therefore repeat both experiments while holding income-tax revenue fixed at its baseline level. At each parameter value, the tax scales under both joint and separate taxation are adjusted to meet this common target. As in the main analysis, the corresponding no-risk economies inherit the tax schedules of their risky counterparts.


Wage risk. Figure 6 shows how the insurance and overall welfare values respond to changes in wage risk under the fixed revenue requirement. The profiles closely resemble those in Figure 1 in Section 5.5. The insurance value increases substantially with the magnitude of risk, primarily reflecting permanent risk, and eventually becomes hump-shaped at high levels of permanent or combined risk. At $q = 2$ in the all-risk experiment, $\mathcal{I}$ reaches 2.39 percent of baseline consumption, compared with 2.54 percent in the main specification.

The overall welfare comparison also remains qualitatively unchanged. Greater permanent or combined risk reduces $\mathcal{W}$ despite the increasing insurance value. At $q = 2$ in the all-risk experiment, $\mathcal{W}$ is -1.98 percent of baseline consumption, compared with -2.07 percent in Section 5.5. Thus, greater (persistent) wage risk generally strengthens the insurance case for joint taxation but increasingly favors separate taxation overall. These conclusions also hold when the government adjusts tax scales to maintain baseline revenue.

Intra-household earnings inequality. Figure 7 shows that the insurance and overall welfare values also respond similarly to changes in female earnings capacity under the fixed revenue requirement. As q_{wf} increases from 0.6 to 1.5, $\mathcal{I}$ rises from 0.57 to 1.84 percent of baseline consumption, compared with 0.56 to 1.90 percent in Figure 2 in Section 5.6.


The overall welfare comparison likewise improves for joint taxation as spouses become more similar in their earnings capacities, largely reflecting the stronger insurance value. $\mathcal{W}$ rises from -1.11 to -0.49 percent of baseline consumption, remaining negative throughout. Hence, the insurance gains associated with declining intra-household earnings inequality do not reverse the overall welfare comparison, even when government revenue is held fixed across the counterfactual economies.

Figure 6: Wage risk under fixed government revenue



Notes: Values are consumption-equivalent variations (CEVs), in percent of common baseline consumption; positive values favor joint taxation. q scales both spouses' innovation standard deviations for the indicated risk component; $q = 1$ corresponds to the estimated process. Panels (a), (b), and (c) vary all wage risk, permanent risk only, and transitory risk only, respectively. Solid blue lines show the insurance value and dashed orange lines the welfare value. In the latter two experiments, the other risk component remains at its estimated level. All other model parameters remain fixed. At each q , the tax scales under both joint and separate taxation are adjusted to maintain the common baseline level of income-tax revenue.

Figure 7: Female earnings capacity under fixed government revenue




Notes: Values are consumption-equivalent variations (CEVs), in percent of common baseline consumption; positive values favor joint taxation. q_{w_f} scales the wife's expected wage profile; $q_{w_f} = 1$ corresponds to the estimated profile. All other model parameters remain fixed. At each q_{w_f} , the tax scales under both joint and separate taxation are adjusted to maintain the common baseline level of income-tax revenue.

Appendix G Tax progressivity and adjustment margins

Figure 8 shows household responses to a permanent male wage shock under joint and separate taxation at different levels of tax progressivity. The labor-supply responses under joint and separate taxation diverge as progressivity increases. Male and female hours respond more strongly in absolute terms under joint taxation, with the widening differences driven primarily by weaker responses under separate taxation. The participation response displays an even more pronounced difference: under joint taxation, the negative response becomes stronger as progressivity increases, whereas under separate taxation it becomes weaker and eventually virtually disappears.

Figure 8: Household responses to a permanent male wage shock at different levels of tax progressivity



Notes: Responses are to a one-percent permanent male wage innovation and are defined as in Table 4. Blue lines refer to joint taxation and orange lines to separate taxation. μ denotes tax progressivity; the dashed vertical line marks its calibrated value. At each μ , the tax scales under both joint and separate taxation are adjusted to maintain the common baseline level of income-tax revenue. All other model parameters remain fixed.