

Piotr Denderski
Tim Obermeier

25/47

Working paper

Marriage, intra- household inequality, and wage risk

Marriage, Intra-Household Inequality, and Wage Risk*

Piotr Denderski[†]

Tim Obermeier[‡]

September 24, 2025

Abstract

How well are workers insured against idiosyncratic wage risk? We revisit this question using a two-earner life-cycle model with endogenous intra-household allocations and marital status dynamics. In a model calibrated to the US, intra-household consumption inequality explains nearly one-fifth of overall inequality. The welfare cost of wage risk is 50% higher than in the unitary household benchmark. A proportional tax reform induces a sizable welfare loss, with women's welfare falling over four times as much as men's due to stronger labor supply responses. Progressive taxation thus provides disproportionately greater insurance for women once intra-household inequality and marriage dynamics are considered.

Keywords: Wage Risk, Intra-Household Allocations, Marriage Market, Consumption Insurance, Tax Progressivity

JEL codes: D15, E21, J12, H24

*We thank Eckhard Janeba and Michèle Tertilt, as well as Richard Blundell, Monica Costa Dias, Hanno Foerster, Hans-Martin von Gaudecker, Anne Hannusch, Omar Licandro, Tom Potoms, Bastian Schulz and conference and seminar participants at the SED Annual Meeting 2025, the GRAPE Summer Workshop on Macro & Finance 2025, the University of Sussex and the Federal Reserve Bank of Cleveland for their helpful comments. In addition, the bwHPC team provided excellent technical support. This is a substantially revised version of a paper previously circulated as "The Marriage Market, Inequality and the Progressivity of the Income Tax". Email: tim.obermeier@leicester.ac.uk and piotr.denderski@leicester.ac.uk.

[†]University of Leicester and Institute of Economics, Polish Academy of Sciences

[‡]University of Leicester, CEP and IFS

1 Introduction

Workers are exposed to substantial idiosyncratic wage risk that can propagate into consumption when asset markets are incomplete. How well such risk is insured, and through which mechanisms, is central to macroeconomics and to the design of redistributive policies. Beyond savings and government policy, family labor supply is a key insurance mechanism for households (e.g. [Blundell, Pistaferri, and Saporta-Eksten \(2016\)](#)), as the presence of another earner mitigates the impact of individual-level wage shocks. [Wu and Krueger \(2021\)](#) show that a calibrated two-earner life-cycle unitary household model matches well the transmission coefficients of wage shocks into household consumption estimated by [Blundell et al. \(2016\)](#), and that the welfare cost of wage risk is substantially lower compared to a one-earner benchmark.¹

However, while analyzing the impact of wage shocks on household-level consumption, the macroeconomic literature on consumption insurance largely focuses on stable couples and a unitary representation of the household, implicitly assuming away any inequality *within* households. Conceptually, this is a strong assumption: for example, if one spouse gets a raise, they need not share it equally with their significant other. Evidence from the literature on collective household models points to systematic differences in consumption between partners, as well as limited risk sharing where wage shocks can shift intra-household consumption shares (e.g. [Lise and Yamada \(2018\)](#)).² In addition, focusing on stable couples abstracts from the fact that individuals spend an important part of their lives as singles since finding a suitable partner takes time and many marriages end in divorce. As singles are more exposed to wage shocks – due to the absence of a second earner – this is also an important factor for the assessment of the implications of wage risk.

In this paper, to address these issues, we embed a Bewley economy with two-earner labor supply, savings and persistent and transitory wage risk, in the spirit of [Wu and Krueger \(2021\)](#), into an equilibrium marriage market with intra-household bargaining, building on the literature on dynamic collective models (e.g. [Mazzocco, Ruiz, and Yamaguchi \(2013\)](#), [Bayot and Voena \(2014\)](#), [Voena \(2015\)](#), [Low, Meghir, Pistaferri, and Voena \(2024\)](#)). In our model, agents start their lives single and search for partners. The marriage market endogenizes intra-household allocations. There is limited commitment in the marriage and relative consumption and work hours within couples are determined by the Pareto weights. These weights are initially set by Nash bargaining at the time of marriage and can be renegotiated if at least one spouse would otherwise prefer to leave the marriage. Couples can also divorce if no renegotiation is feasible that keeps both spouses in the marriage.

We calibrate the model to US data on men’s and women’s relative earnings and hours, and,

¹Example papers modeling households as single-earners include [Kaplan and Violante \(2010\)](#) and [Heathcote et al. \(2014\)](#).

²See [Chiappori and Meghir \(2015\)](#) and [Chiappori and Mazzocco \(2017\)](#) for surveys of this literature. Various studies have estimated the intra-household sharing rule and find a dependence on relative wages (e.g. [Chiappori et al. \(2002\)](#), [Blundell et al. \(2007\)](#), [Lise and Seitz \(2011\)](#), [Cherchye et al. \(2012\)](#)) or relative incomes (e.g. [Browning et al. \(1994\)](#)). An alternative approach has been to use revealed preference models to back out private consumption from household-level data (e.g. [Cherchye et al. \(2015\)](#), [Cherchye et al. \(2017\)](#))

crucially, on marriage and separation rates over the life cycle and in the cross-section. An important feature of the model is that it allows for consumption inequality both within and between households. In the calibrated model, intra-household inequality accounts for 26.4% of inequality among the married, and 17% in the overall population once singles are included. This is close to the findings of [Lise and Seitz \(2011\)](#), who estimate a model based on British data. Intra-household inequality arises from differences in relative wages, and assets, at the time of marriage - which affect the initial Pareto weight - as well as from renegotiation of the Pareto weight following wage shocks over the course of marriage. Renegotiation is infrequent (23% of couples do so within five years) and small (five-year Pareto weight correlation is approximately 0.9), so weights are highly persistent and largely set by initial relative wages.

Regarding the transmission of wage shocks to consumption, at the household level our model is consistent with the estimates of [Blundell et al. \(2016\)](#): for couples, a 1 pp shock to men's (women's) wages changes household expenditure by 0.34 pp (0.2 pp, respectively). What is more, the model also allows us to examine the transmission at the individual level, thereby capturing asymmetries within couples. For example, a 1 pp shock to men's wages increases men's consumption by 0.36 pp and women's consumption by 0.31 pp, as it can shift the Pareto weights in favor of men. In addition, the model suggests an important role for the initial Pareto weight as a transmission channel of wage risk: persistent wage shocks which singles experience early on in life will also be transmitted into relative consumption in their future marriage and lead to more intra-household inequality.

Our first main result is about revisiting the welfare cost of wage risk from the perspective of our model and demonstrating that it is substantially higher than in a unitary household model. We quantify the welfare cost by shutting down wage shocks and computing the ex-ante consumption-equivalent (CEV) loss relative to this no-risk benchmark for a newborn cohort of men and women. In our model, the CEV loss is 22.5%. For comparison, the loss is 14.2% in a comparison unitary household model. Hence, the unitary household model underestimates the welfare cost of wage risk by a half. We also compute the CEV of approximately 28% in an auxiliary no-marriage model where agents start and remain single. The increase in welfare cost of wage risk in our model relative to the unitary household one is not only due to the presence of singles in our model; the within-couple bargaining and limited commitment matter too. Indeed, in an economy composed of two populations: unitary couples and permanently single agents with shares determined by the marriage rate in our model, based on the calculations for the two auxiliary models, the CEV cost of wage risk would be about 18%.

Our second main result is that we study the impact of income tax progressivity as a policy instrument which provides public insurance against wage risk. In our policy experiment, we eliminate the current US level of progressivity and introduce a proportional tax through a balanced-budget reform. In our model, this policy change leads to a CEV welfare loss of 6.1%, which is stronger than the welfare loss in the unitary comparison model (4%). This is driven by starkly asymmetric welfare effects by gender: women's welfare falls very significantly by 9.5%,

whereas men’s welfare only falls by 2.2%. This asymmetry reverts in the unitary comparison model, where the reform reduces women’s welfare by 3% and men’s welfare by 5%.³ Investigating the mechanisms behind these differences, we find that a particularly important reason for the strong negative welfare impact on women is that the increase in their work hours is significantly more pronounced in our model, which pertains closely to marriage and bargaining.⁴ For example, under proportional taxes, the after-tax income becomes more dispersed, which increases inequality in asset holdings. Because of gender gap in productivity, this benefits men more, and single women increase their labor supply to counter that.

Our paper primarily contributes to the macroeconomic literature on idiosyncratic risk, consumption insurance, and closely intertwined with them cost-benefit analyses of redistributive policies. [Wu and Krueger \(2021\)](#) find that a two-earner life-cycle unitary household model replicates the transmission coefficients from [Blundell et al. \(2016\)](#) well and study the implications for the welfare cost of wage risk and optimal tax progressivity. In addition, [Blundell, Pistaferri, and Saporta-Eksten \(2018\)](#) study consumption insurance in a richer model that accounts for the time allocation to children. Our paper contributes by introducing marriage dynamics and intra-household bargaining and studying their role.⁵ We find that accounting for limited commitment and marital status dynamics matters quantitatively for the assessment of welfare cost of wage risk despite similar wage risk transmission coefficients.⁶ [Heathcote, Storesletten, and Violante \(2017\)](#) and [Conesa and Krueger \(2006\)](#) study (optimal) income tax progressivity, while [Guner, Kaygusuz, and Ventura \(2011\)](#), [Holter, Krueger, and Stepanchuk \(2019\)](#), [Holter, Krüger, and Stepanchuk \(2023\)](#), [Guner, Kaygusuz, and Ventura \(2023\)](#) have analyzed tax and transfer policy in two-earner households. These models capture joint taxation and labor supply responses but abstract from intra-household inequality and the dynamics of marriage and divorce which we introduce. We find that women benefit much more strongly than men benefit from progressive taxation.

Second, our paper is closely related to the collective model literature on couples’ dynamic decision-making. Building on [Mazzocco, Ruiz, and Yamaguchi \(2013\)](#), the limited commitment model, in which intra-household allocations are renegotiated infrequently when participation constraints bind, has become the workhorse model in this literature and has been applied in various contexts, ranging from tax and transfer policy ([Low, Meghir, Pistaferri, and Voena \(2024\)](#), [Bronson, Haanwinckel, and Mazzocco \(2024\)](#)) to divorce laws and prenuptial contracts ([Voena \(2015\)](#), [Bayot](#)

³The magnitude of the welfare decrease captures welfare gains of tax progressivity if one were to take the proportional tax economy as a benchmark. Overall, the welfare drop induced by the proportional tax reform is one-third higher in our model than in the unitary model indicating that the redistributive benefits of progressivity in our model are larger.

⁴We find that the married women are responsible for 42% of the labor supply response to the tax reform along the intensive margin. [Guner et al. \(2012\)](#) study a similar tax reform and find the married women are responsible for 51% of the labor supply response, but half of it is due to the extensive margin.

⁵[Gayle and Shephard \(2019\)](#) study optimal taxation in a collective model with an equilibrium marriage market.

⁶The welfare cost of income or wage risk has been studied by a broader literature (e.g., [Karahan and Ozkan \(2013\)](#), [Heathcote, Storesletten, and Violante \(2008\)](#), [Guvenen, Ozkan, and Madera \(2024\)](#), [De Nardi, Fella, and Paz-Pardo \(2020\)](#)). In addition, several studies have investigated the impact of rising income risk (e.g. [Heathcote, Storesletten, and Violante \(2010\)](#)). [De Nardi, Fella, and Paz-Pardo \(2025\)](#) study the properties of the wage process and the implication for optimal government insurance.

and Voena (2014), Reynoso (2024), Foerster (2025), De Rock et al. (2023))). Lise and Yamada (2018) use Japanese survey data to estimate the extent of commitment empirically and find that their evidence is consistent with limited commitment models.⁷ Compared to this literature, our contribution is to link the limited commitment model to the macroeconomic literature on consumption insurance.

Third, our paper also links to a growing literature that has used equilibrium models of the marriage market (e.g. Shimer and Smith (2000), Knowles (2012), Guvenen and Rendall (2015), Greenwood, Guner, Kocharkov, and Santos (2016), Santos and Weiss (2016), Goussé, Jacquemet, and Robin (2017), Shephard (2019), Ciscato (2019)). Equilibrium search models of marriage with wage dynamics - where the distributions of potential partners are endogenous rather than exogenously imposed - are still relatively rare. Santos and Weiss (2016) study rising income volatility and argue that it explains an important part of the delay in marriage over time. Ciscato (2019) studies changes in the wage structure and their role in explaining marriage and divorce patterns. Most closely related to our model, Shephard (2019) builds an equilibrium marriage market with limited commitment to analyse spousal age gaps. Compared to these papers, our model focuses on consumption insurance and considers an equilibrium marriage market in a Bewley economy with savings.

The rest of this paper is organized as follows. Section 2 describes the model, which is calibrated in Section 3 which is also where the main predictions of the model for marriage dynamics, inequality, and wage risk transmission are presented. Section 4 shows our counterfactual experiments. The last section concludes.

2 Model

In this section, we present the details of the model, starting with the description of preferences and technology. Then, we discuss the role of the government, moving on to specifying the timing of events in the economy. The description of choices and value functions concludes. We embed a Bewley economy with endogenous labor supply for both earners, permanent and transitory wage risk, and savings (as in Wu and Krueger (2021)) in a dynamic collective decision problem with a marriage market.

2.1 Households

Demographics and preferences. In each period, a new cohort of measure one of agents enters the economy at age $t = 1$, and all agents live for T periods. An agent i is either a man or a woman and we distinguish genders with $g(i) \in \{m, f\}$.⁸ The shares of both genders in each cohort are

⁷While Lise and Yamada (2018) also use a dynamic collective household model, our paper takes a different approach by also modeling the marriage market, which endogenizes the Pareto weights within households and adds singlehood to the model.

⁸Unless necessary to avoid confusion, we will scrap the dependence of an individual's gender on their index i for notational convenience.

equal. The preferences of agent i over consumption c_{it} and labor h_{it} in period t , given their gender, are described by the following utility function:

$$u_g(c_{it}, h_{it}) = \frac{c_{it}^{1-\gamma}}{1-\gamma} - \phi_g \frac{h_{it}^{\eta_g}}{\eta_g}.$$

The parameter γ is the coefficient of relative risk aversion and η_g maps to a gender-specific Frisch elasticity. ϕ_g is a gender-specific scaling factor for the disutility of labor.⁹ Agents discount future utilities at the rate β .

An agent can be single or married, and all agents enter as singles. Two singles i_1, i_2 of opposite gender $g(i_1) \neq g(i_2)$, can decide to marry, and we will refer to them as a couple $j = j(i_1, i_2)$. Couples can divorce at every period and are subject to a unilateral divorce regime where it is enough for one of two spouses to wish to divorce for it to happen. Although we postpone the discussion of marriage dynamics to Section 2.4, in what follows we will omit the individual indices and replace them with gender when discussing couples.

Earnings, consumption, and savings. All agents work until $t = t_R > 1$, after which they retire. Hence, the agents work for t_R periods and spend $T - t_R$ in retirement. Agents do not work when retired, $t > t_R \implies h_{it} = 0$. When retired, all agents receive a pension b_g .¹⁰ Before retirement, an agent supplying h_{it} units of labor produces $y_{it} = w_{it}h_{it}$ with the wage w_{it} determined by individual productivity. Provided $t \leq t_R$, worker's wage is pinned down by three factors: a deterministic gender-age-specific trend $\kappa_{t,g}$, a persistent shock $\nu_{i,t}$ and a transitory shock ε_{it} :

$$\log(w_{it}) = \kappa_{t,g} + \nu_{it} + \varepsilon_{it}, \varepsilon_{it} \sim \mathcal{N}(0, \sigma_{\varepsilon,g}^2) \text{ with:} \quad (1)$$

$$\nu_{it} = \nu_{it-1} + \epsilon_{it}, \epsilon_{it} \sim \mathcal{N}(0, \sigma_{\nu,g}^2), \quad (2)$$

The deterministic trend comprises a gender effect and an age effect, $\kappa_{t,g} = \kappa_t + \kappa_g$. The wage shocks are drawn from gender-specific distributions with variances $\sigma_{\nu,g}^2$ and $\sigma_{\varepsilon,g}^2$ for the persistent and transitory shocks, respectively.

Workers can save using a risk-less asset which yields an exogenous return R , or transform their output into consumption. We rule out borrowing, which means that the asset level must be weakly positive in each period. When two agents marry, they pool their savings and upon divorce, they divide them equally.

An important feature of the model is that it allows for economies of scale in consumption, meaning that spouses are able to consume more than they would in singlehood. Define the amount of output transformed into consumption as consumption expenditure C_{it} for single i , and C_{jt} for couple j . Singles' consumption expenditures map one-to-one into consumption, $C_{it} = c_{it}$.

⁹Note that our model focuses on the intensive margin and abstracts from the extensive margin of labor supply. Wu and Krueger (2021) find that the extensive margin contributes very little to insurance against shocks, as labor force participation is high (around 80%) for the period the model is calibrated to. As a result, we abstract from the extensive margin to keep the model more tractable.

¹⁰This captures the consequences of average gender earnings gap on accrued pensions in a reduced form.

For couples, we follow [Voena \(2015\)](#) and assume that for given consumption choices of a couple, c_{ft} and c_{mt} , the level of expenditure required to obtain them is:

$$C_{jt} = \left(c_{ft}^\rho + c_{mt}^\rho \right)^{\frac{1}{\rho}}, \quad (3)$$

where the parameter $\rho \in [1, \infty)$ measures the size of economies of scale. For $\rho = 1$ there are no economies of scale and the larger the ρ , the greater they are.¹¹

Love shocks. Besides consumption and labor supply, the third component of utility is the non-economic utility from the current relationship of the individual (the 'love' shock), which we denote as $\theta_{i,t}$. This is an individual-specific additive utility shifter that married agents receive. Following [Voena \(2015\)](#), we assume that both partners have separate love shock processes, which are independent from each other. The dynamics of these shocks are governed by an AR(1) process:

$$\theta_{i,t} = \rho_\theta \theta_{i,t-1} + \epsilon_{i,t}^\theta,$$

where $\epsilon_{i,t}^\theta$ represents an innovation drawn from a normal distribution with mean zero and variance $\sigma_{\epsilon,\theta}^2$ and ρ_θ is an autoregressive coefficient. The initial values of the love shocks, the ones drawn by prospective partners upon meeting, are drawn from a distribution with mean μ_θ and variance σ_θ^2 .

2.2 The Government

The government collects revenue through a tax on labor income. Singles and couples face different tax regulations in the US. To capture these differences, we assume the taxes imposed on household income y are represented by a marital-status-dependent tax function as in [Heathcote, Storesletten, and Violante \(2017\)](#). This tax function implies that the after-tax income $y - T(y)$ is given by:

$$y - T(y) = \begin{cases} \tau_l^s y^{1-\tau_p^s} & \text{for singles} \\ \tau_l^c (y_m + y_f)^{1-\tau_p^c} & \text{for couples.} \end{cases}$$

The marital-status-dependent parameters τ_l and τ_p capture the level, and the progressivity of the tax system, respectively. The government's budget constraint equates the net present value of the expected life-time revenue from a newly born cohort to an exogenous spending requirement $\bar{G}$:

$$\text{Rev}(\psi_1, \psi_2) = \mathbb{E} \left[\sum_t \beta^t \left(T(y_{i,t}) \right) \right] = \bar{G}$$

¹¹This specification permits pinning down the extent of consumption economies of scale by matching it to equivalence scales. Economies of scale are maximized when consumption of both partners is equal, and zero if one partner consumes the entire budget.

The revenue $\bar{G}$ is used to finance non-valued government expenditure.¹²

2.3 Timing

The timing of the events within a period is as follows. All agents first receive labor productivity shocks which, together with the gender-age-specific trend determine the current-period wage. Then, the marriage market opens which consists of three stages: matching, bargaining, and dissolution. In the matching stage, singles draw a potential partner from the distribution of available singles in that period and an initial realization of the love shock which captures the non-economic utility from the relationship. Based on observing each other's characteristics and the love shock draws, both individuals decide on whether they want to enter a relationship. If they do, they get married and make decisions as a couple. If at least one person prefers staying single, both remain single for this period. Next, the marital bargaining stage takes place. Agents who start the period as married receive the love shocks for that period. Then, they decide about whether to remain married or not following a protocol that we specify in more detail in Section 2.4. Finally, all households make consumption, labor supply, and savings decisions.

2.4 Value Functions and Choices

State variables. The contents of the state vector depend on the labor market and marital statuses. Thus, we set the former as L , either working-age or retired, $L \in \{W, R\}$, and marital status $M \in \{S, C\}$ for singles and couples, respectively. Value functions are also distinguished by agent's gender g , and age t . For working-age individuals, we need to keep track of wage shocks and to economize on notation we set $\omega = (\nu, \varepsilon)$. Whenever we do not specify the labor market status below, the convention is that $\omega = \emptyset$ for the retired. When the labor market status is explicitly stated, we drop ω for the retired. We also need to track assets A , held either individually (by singles), or jointly (by couples).

The individual-level uncertainty is gradually resolved throughout the period and to account for that we set $V_{g,t}^{M,L}(\cdot)$ as individual value function prior to entering the marriage market, but after the individual productivity update, and $U_{g,t}^{M,L}(\cdot)$ stands for individual value after the marriage market has closed. Couples seek to optimize weighted sums of the individual values, and the Pareto weight assigned to wife, λ , is the sufficient statistic for the bargaining positions in the couple. Then, the couple-level value functions are $\mathcal{U}_t^L(\cdot) = \lambda U_{f,t}^{M,L}(\cdot) + (1 - \lambda) U_{m,t}^{M,L}(\cdot)$ and $\mathcal{V}_t^L(\cdot) = \lambda V_{f,t}^{M,L}(\cdot) + (1 - \lambda) V_{m,t}^{M,L}(\cdot)$.

Regarding couples, we follow earlier literature and assume that an agent of age t can only be married to another agent of the same age so that the index t refers to both members of the household.¹³ For married individuals, the state vector contains, apart from the assets of the household and wage shocks for working-age households, the love shock realizations and the Pareto weight

¹²This can also be interpreted as the government producing an economy-wide public good P , which provides each individual with an additive utility $u^P(P)$ and does not affect choices.

¹³A notable exception is [Shephard \(2019\)](#) who analyzes the spousal age gap.

assigned to the wife λ . Hence, we set $V_{g,t}^{C,L}(\lambda, \omega_g, \theta_g, \omega_{g^-}, \theta_{g^-}, A)$ as the value of an individual of gender g and age t who is married before marriage market opens.¹⁴ This individual carries over the Pareto weight λ and love shock realizations from the previous period which might be updated in the current period. Then, $U_{g,t}^{C,L}(\lambda, \omega_g, \theta_g, \omega_{g^-}, \theta_{g^-}, A)$ is the corresponding value of a married individual after the marriage market has closed and the marriage decided to continue in the given period (after updates to love shock value and possible renegotiation of the Pareto weight).

Marriage market. We assume that married individuals do not search for other partners, while all singles enter the marriage market. The marriage market is segmented with respect to age, a single of age t can only meet another single of the same age. Singles randomly meet one potential partner from the pool of available singles of the opposite gender. An individual of gender g and age t meets a prospective partner with characteristics (ω_{g^-}, A_{g^-}) with probability $\Lambda_{g,t}(\omega_{g^-}, A_{g^-})$. Here, the probability $\Lambda_{g,t}(\omega_{g^-}, A_{g^-})$ is the share of (ω_{g^-}, A_{g^-}) agents in the pool of singles of opposite gender g^- at age t as meetings between singles are random.

Upon a meeting, the potential partners each draw an initial realization of the love shock $\theta_{i,t}$. They also observe the love shock draw and the state vector of the potential partner. Then, they bargain over a (normalized) Pareto weight λ that will determine the power of the wife within the household. The Pareto weights are determined by Nash bargaining, with the outside options of exiting the marriage market as single. This can be thought of the value of staying single for an additional period determining the 'bargaining position' of each individual and captures that Pareto weights depend on relative wages, assets and love shocks. Formally, the initial Pareto weight $\tilde{\lambda}$ satisfies:

$$\tilde{\lambda} = \arg \max_{\lambda} \left(U_{f,t}^{C,L}(\lambda, \omega_f, \theta_f, \omega_m, \theta_m, A_f + A_m) - U_{f,t}^{S,L}(\omega_f, A_f) \right) \times \left(U_{m,t}^{C,L}(\lambda, \omega_m, \theta_m, \omega_f, \theta_f, A_f + A_m) - U_{m,t}^{S,L}(\omega_m, A_m) \right).$$

Thus, $\tilde{\lambda}$ maximizes the product of the *gains from marriage* (relative to staying single) of the two spouses. If there is no $\lambda \in (0, 1)$ for which the gains from marriage are positive for both individuals, the match is rejected and the individuals stay single. Otherwise, they get married, pool their individual assets, and leave the population of available singles.

Value functions: singles. An agent who, after the marriage market closes, remains single, solves the following problem when working:

$$U_{g,t}^{S,W}(\omega, A) = \begin{cases} \max_{c,h} u_g(c, h) + \beta \mathbb{E}_{\omega'} V_{g,t+1}^{S,W}(\omega', A') & \text{if } t \leq t_R - 1 \\ \max_{c,h} u_g(c, h) + \beta V_{g,t+1}^{S,R}(A') & \text{if } t = t_R, \end{cases}$$

¹⁴We distinguish partner's characteristics with subscript g^- , e.g., when $g = f$ then $g^- = m$ etc.

subject to the budget constraint $(1 + R)A + wh - T(wh) = c + A'$, the no-borrowing constraint $A' \geq 0$, and the wage dynamics (1) - (2). The retired singles solve:

$$U_{g,t}^{S,R}(A) = \begin{cases} \max_c u_g(c, 0) + \beta V_{g,t+1}^{S,R}(A') & \text{if } t_R < t < T \\ \max_c u_g(c, 0) & \text{if } t = T, \end{cases}$$

subject to the budget constraint $(1 + R)A = c + A'$, and the no-borrowing constraint $A' \geq 0$. The beginning-of-period values of being single account for the opportunity to marry:

$$V_{g,t}^{S,L}(\omega_g, A_g) = U_{g,t}^{S,L}(\omega_g, A_g) + \mathbb{E}_{\theta, \theta_{g-}, \omega_{g-}, A_{g-}} \left[\Lambda_{g,t}(\omega_{g-}, A_{g-}) \max \left\{ U_{g,t}^{C,L}(\tilde{\lambda}, \omega, \theta, \omega_{g-}, \theta_{g-}, A_g + A_{g-}) - U_{g,t}^{S,L}(\omega_g, A_g), 0 \right\} \right]$$

In particular, this expression highlights that the dynamics of the outside options to marriage are determined both by the living standard while single and by the full distribution of potential marriages in the future.

Divorce and Limited Commitment. After the arrival of love shocks, couples make consumption and labor supply decisions cooperatively, subject to limited commitment (see [Chiappori and Mazzocco \(2017\)](#) for a survey). This corresponds to a two-step procedure. In the planning stage, the couple solves the maximization problem of the weighted sum of its members' values, given the current Pareto weight. Second, each partner compares the resulting utility to their outside option of continuing as single right away.¹⁵ We assume that upon divorce household assets are split equally.¹⁶

Starting from the second step, if one partner prefers to divorce rather than accept the consumption and labor decisions planned at current Pareto weight, a renegotiation takes place and the weight can be adjusted. Efficiency requires making the *minimal* adjustment relative to the previous Pareto weight, so that the spouse who wishes to divorce at current λ is just indifferent between divorcing and staying in the marriage. If divorce is prevented, then the household repeats the planning step with updated Pareto weight, and the couple then proceeds with the updated consumption and labor supply decisions. Formally, the participation constraint reads:

$$U_{g,t}^{C,L}(\lambda, \omega_g, \theta_g, \omega_{g-}, \theta_{g-}, A) \geq U_{g,t}^{S,L}(\omega_g, A/2). \quad (4)$$

Value functions: couples. The marriages that end up continuing maximize the weighted sum of its members values. These optimization problems follow the same logic as those of singles, and for the sake of brevity we only present here the case of working-age households who do not face

¹⁵There is two-sided limited commitment and hence spouses can unilaterally file for divorce. Note, newly formed couples by construction are happy with the Pareto weight they agreed upon.

¹⁶This assumption is often made in the literature as an approximation of the US system (see e.g. [Fernandez and Wong \(2017\)](#); [Mazzocco et al. \(2013\)](#)).

retirement next period, relegating the other cases to the Appendix S.A.1. Thus, when $t < t_R - 1$ the household solves:

$$\begin{aligned} \mathcal{U}_t^W(\lambda, \omega_m, \theta_m, \omega_f, \theta_f, A) = & \max_{c_f, c_m, h_f, h_m} \\ & \underbrace{\lambda \left(u_f(c_f, h_f) + \theta_f + \beta \mathbb{E}_{\omega'_f, \omega'_m} V_{f,t+1}^{C,W}(\lambda, \omega'_f, \theta_f, \omega'_m, \theta_m, A') \right)}_{U_{f,t}^{C,W}(\lambda, \omega_f, \theta_f, \omega_m, \theta_m, A)} \\ & + (1 - \lambda) \underbrace{\left(u_m(c_m, h_m) + \theta_m + \beta \mathbb{E}_{\omega'_f, \omega'_m} V_{m,t+1}^{C,W}(\lambda, \omega'_m, \theta_m, \omega'_f, \theta_f, A') \right)}_{U_{m,t}^{C,W}(\lambda, \omega_m, \theta_m, \omega_f, \theta_f, A)}, \quad (5) \end{aligned}$$

subject to the budget constraint:

$$C + A' = w_f h_f + w_m h_m - T(w_f h_f + w_m h_m) + (1 + R)A, \quad (6)$$

the no-borrowing constraint $A' \geq 0$, and the condition capturing economies of scale in household consumption expenditures C , Equation (3). Here, the continuation values – of entering next period married – account for the dynamics of love shocks, the possibility to divorce, and the potential update to the Pareto weight:

$$V_{g,t+1}^{C,W}(\lambda, \omega_g, \theta_g, \omega_{g-}, \theta_{g-}, A) = \mathbb{E}_{\theta'_g, \theta'_{g-}, \lambda'} \max \left\{ U_{g,t+1}^{C,L}(\lambda', \omega_g, \theta'_g, \omega_{g-}, \theta'_{g-}, A), U_{g,t+1}^{S,L}(\omega_g, A/2) \right\}. \quad (7)$$

Pareto weights and within-couple inequality. The first-order conditions for the optimal choice of consumption and hours worked imply the following relationships between the consumption and hours worked of spouses:

$$\frac{c_m}{c_f} = \left(\frac{1 - \lambda}{\lambda} \right)^{1/(\gamma + \rho - 1)} \quad \text{and} \quad \frac{\phi_m h_m^{\eta_m - 1}}{\phi_f h_f^{\eta_f - 1}} = \frac{w_m \lambda}{w_f (1 - \lambda)}. \quad (8)$$

Thus, the Pareto weight is a sufficient statistic for within-couple consumption inequality given preference and economies of scale parameters, and, conditioning further on wages, also for hours worked differentials.¹⁷ All else kept equal, an increase in the bargaining weight assigned to the wife λ , increases higher her consumption, and lowers her working hours, relative to consumption and hours worked by her husband.

Marriage market equilibrium. In equilibrium, singles and couples make optimal choices according to the value functions defined in this section. In addition, marriage and divorce decisions, and Pareto weights are determined according to these value functions. The key equilibrium objects are the distributions of single men and women at each age ($\Lambda_{g,t}(\omega_{g-}, A_{g-})$). These distri-

¹⁷For example, $\lambda = 1 - \lambda = 1/2$ implies $c_m = c_f$ and the more extreme the difference in the bargaining position of spouses $(1 - \lambda)/\lambda$, the greater the gap between their consumption.

butions are endogenous because they depend both on forward-looking decisions (in the case of assets) and on selection into and out of marriage (for both assets and wage shocks). Solving the model therefore requires solving a fixed-point problem where the single distributions are equilibrium outcomes.

3 Quantitative Properties of the Equilibrium

This section contains the details on how we take the model to the data, and what the main quantitative characteristics of its equilibrium are. Understanding the model’s internal mechanisms is an essential step toward interpreting the counterfactual experiments presented in Section 4. There are several key takeaways from this section. First, there is substantial within-household consumption inequality that accounts for one-fourth of consumption dispersion among the married. Second, our model replicates well the results in the earlier literature on transmission of wage shocks to labor supply and household consumption for the married. Importantly, we can go beyond earlier work, decomposing these further into individual consumption, and estimating such transmission coefficients also on singles.¹⁸ Third, we quantify additional channels through which the wage shocks operate, finding asymmetric and gendered effects on marriage formation, bargaining position of spouses, and marriage dissolution.

3.1 Model Calibration and Fit

We begin by outlining the calibration of our model to match the key features of the US economy. Then, we discuss the fit of the model. To facilitate comparison with previous studies and to isolate the impact of our model’s novel elements, we adopt a calibration strategy that closely follows [Wu and Krueger \(2021\)](#) when possible, particularly in the parameterization of preferences and the tax system. By following their calibration strategy, we target selected moments from the Panel Survey of Income Dynamics (PSID) data for the period between 1999 to 2009, which is the sample from [Blundell et al. \(2016\)](#). The main new elements of our calibration are the marriage market parameters (i.e. love shock process), which we use to target moments about marriage and divorce rates from the PSID. We provide an overview of model parameter values in Table 1 and give more details on the calibration below.

We assume that individuals enter the economy at age 21, retire at age 65 and die at 80. Each model period corresponds to one year. This leads to $t_R = 45$ and $T = 60$. Individuals discount the future at rate $\beta = 0.995$, as in [Wu and Krueger \(2021\)](#). The return on assets is 2%, thus, $R = 1.02$. Individuals are born with zero assets and the log of wage upon birth is set to zero for all agents, $\log(w_{i1}) = 0$.

We fix the relative risk aversion coefficient at $\gamma = 1.73$, and the curvature parameters for

¹⁸We elaborate further on the importance of love shocks and analyze the interplay of consumption insurance with marital status dynamics in Supplemental Appendix section [S.B.4](#).

TABLE 1: Model parameters

Param.	Description	Value	Source/Target
Pre-determined			
β	Annual discount factor	0.995	Wu and Krueger (2021)
γ	Consumption, curvature	1.73	Blundell et al. (2016)
η_f	Disutility of work, women, curv.	2.17	Blundell et al. (2016)
η_m	Disutility of work, men, curv.	2.89	Blundell et al. (2016)
R	Annual interest rate	0.02	Wu and Krueger (2021)
ρ	Consumption economies of scale	1.4023	Voena (2015)
$\sigma_{\nu,f}^2$	Perm. wage shock variance, women	0.0382	Blundell et al. (2016)
$\sigma_{\nu,m}^2$	Perm. wage shock variance, men	0.0303	Blundell et al. (2016)
$\sigma_{\varepsilon,f}^2$	Temp. wage shock variance, women	0.0125	Blundell et al. (2016)
$\sigma_{\varepsilon,m}^2$	Temp. wage shock variance, men	0.0275	Blundell et al. (2016)
τ_l^C	Level parameter, tax on couples	0.843	Wu and Krueger (2021)
τ_p^C	Progressivity, tax on couples	0.133	Wu and Krueger (2021)
τ_p^S	Progressivity, tax on singles	0.11	Holter et al. (2019)
b_f	Retirement benefits, women	0.134	Wu and Krueger (2021)
b_m	Retirement benefits, men	0.235	Wu and Krueger (2021)
Estimated			
τ_l^S	Level parameter, tax on singles	0.756	mean adjust. of Holter et al. (2019)
κ_f	Women productivity (rel. to men)	0.572	Gender gap (income)
ϕ_f	Disutility of work, women, scale	6.877	Avg. hours, women
ϕ_m	Disutility of work, men, scale	5.036	Avg. income, married men (norm.)
$\sigma_{\theta}^2, \sigma_{\varepsilon,\theta}^2$	Variance of love shock	0.499	Divorce probability
ρ_{θ}	Autocorrelation of love shock	0.972	Marriage rate over life-cycle, PSID
μ_{θ}	Mean of initial love shock	-1.799	Marriage rate over life-cycle, PSID

Notes: This table shows the calibrated parameters. Some parameters are set exogenously following the literature or to match external targets. Other parameters are calibrated to minimize the distance between a set of model and data moments.

disutility of work at $\eta_f = 2.17$ and $\eta_m = 2.89$ following Blundell et al. (2016).¹⁹ These are the key parameters that determine the labor supply elasticities and the transmission of shocks into the supply and consumption of the labor. To set the parameter for economies of scale, we follow Voena (2015) in setting $\rho = 1.4023$. This value captures the McClements scale, according to which

¹⁹Note that Blundell et al. (2016) (BPS) estimate these parameters based on approximations to a two-earner life-cycle model under the assumption of exogenous Pareto weights. Thus, with endogenous Pareto weights, there could be a bias in the BPS estimates. Intuitively, the reason for this bias is that a positive shock for a married individual would lead to an increase in their Pareto weight, which dampens their labor supply response. In other words, the BPS method could attribute a muted labor supply response to a low labor supply elasticity, while in fact the muted response is driven by a change in Pareto weights. Table 4 shows that in our calibration, the adjustment of the Pareto weight dampens women's labor supply response by 10%, and men's labor supply by 17%, which suggests that the magnitude of the bias is likely small. We therefore use the estimates from BPS to simplify the calibration.

a person living alone spends 61% of what a couple spends to consume the same amount.

To calibrate the wage process, we discipline its life-cycle component, κ_t , using the data from [Rupert and Zanella \(2015\)](#). Then, we set the variances of permanent shocks to values estimated by [Blundell et al. \(2016\)](#), that is, $\sigma_{v,m}^2 = 0.0303$ for men, and $\sigma_{v,f}^2 = 0.0382$ for women.²⁰ The variances of the temporary wage shocks are set to 0.0125 for women and 0.0275 for men.

The gender-specific deterministic component of labor productivity is normalized to unity for men, $\kappa_m = 1$ while κ_f is calibrated internally to match the ratio of (average) female hours to male hours equal to approximately 0.73. This requires $\kappa_f = 0.572$.²¹ We calibrate the scaling factor for the disutility of work for men, ϕ_m , to ensure that the average labor income for married men is equal to unity, as in [Wu and Krueger \(2021\)](#). The counterpart of this parameter for women, ϕ_f is calibrated to match the average labor income of women to be 49% of the average for men. We obtain $\phi_f = 6.877$ and $\phi_m = 5.036$.

We set retirement benefits to 0.134 for women and 0.245 for men. [Wu and Krueger \(2021\)](#) calibrate the pensions to 0.38 for a joint household, and we replicate their average incomes for men and women. Assuming the same replacement rate for both genders, these values carry over the gender gap that we match in incomes. Following their paper, the parameters of the tax function for couples are set to $\tau_1^M = 0.843$ and $\tau_2^M = 0.133$. Regarding the tax function for singles, we use the values provided by [Holter et al. \(2019\)](#) with the caveat that the values they report assume the distribution of taxable income in consideration to have mean normalized to unity. This does not impact the progressivity parameter, which we set to the value they provide, $\tau_2^S = 0.11$.²² The level parameter, however, needs to be calibrated internally because in our model the mean income of singles is not normalized and we arrive at $\tau_1^S = 0.756$.

The remaining three parameters describe the love shocks: the distribution of the initial draws that prospective couples receive, the variance of innovations to the love shock, and its autocorrelation. We discipline these parameters through their impact on marriage and divorce probabilities. We target the overall marriage rate and its life-cycle profile (share of couples in the PSID sample). We further assume that the variance of initial love draws is equal to the variance of the innovations to love, $\sigma_\theta^2 = \sigma_{\epsilon,\theta}^2$. Intuitively, the larger the variance of the distribution of initial love shock draws, the greater the option value of waiting for a more suitable partner. The mean of the initial love shock disciplines the level of the marriage rate and the autocorrelation of love impacts on the likelihood of divorce. We find that the love shocks are quite persistent, $\rho_\theta = 0.972$. As for the mean of initial love shock, we estimate $\mu_\theta = -1.799$. Thus, on average, potential partners are

²⁰Note that we assume that these shocks are uncorrelated between partners. We make this assumption as it allows us to use the same wage process for both singles and coupled individuals. We discretize this process using the method from [Wu and Krueger \(2021\)](#). Allowing for a positive correlation would require specifying a separate wage process for singles, which would likely have a different grid and would imply that getting married directly changes the wage (since the same grid point corresponds to a different wage on the married and single grid). Empirically, the wage process estimated by [Blundell et al. \(2016\)](#) implies a small positive correlation coefficient of 0.08, which suggests that the impact of this assumption is likely small in practice.

²¹For reader's convenience, we present the resulting life-cycle profiles of individual productivity and average wages in Supplemental Appendix [S.B.1](#).

²²It is straightforward to show that if the income is scaled by constant x , the new level parameter is scaled by x^{τ_p} .

undesirable - in this way the model rationalizes low marriage rates among young individuals in the data, as it takes multiple draws to find a good match.

The model fits the targeted moments reasonably well (Table A.1). It slightly under-predicts marriage rates at young ages and over-predicts them at older ages, but overall performs reasonably well (Figure 1a), despite relying on just three parameters to target this profile.²³ Marriage rates rise quickly in early adulthood—by age 30, 70% are in relationships, stabilizing around 80% by age 40.²⁴ In the cross-section, 72% of agents are married. The cross-sectional divorce probability is 0.04 which implies a marriage probability of about 0.1.

3.2 Marriage and Inequality

In this section, we conduct several exercises to better understand the determinants of marriage and consumption inequality in the calibrated model. To show how marriage formation depends on wages, assets and the love shock – and assess the relative importance of these factors – we estimate a logistic regression on singles predicting whether they get married in a given period using model-simulated data. We control for various combinations of wages, assets, and age and report the results of this exercise in Table A.2 in the Appendix. When we control only for wages or only for assets (with or without controlling for age), we find that agents who are wealthier and earn more are more likely to marry. Older agents are less likely to get married. Most importantly, the richest specification - using information on own wages, assets and age - classifies correctly the marital outcomes in about 40% of cases, being wrong about one-third of predicted marriages. Therefore, the characteristics of the randomly drawn partner and love shocks play a significant role in marriage formation.²⁵

Interestingly, we find that the average value of the initial love shock for the newlyweds is negative and equal to -0.094 and this is an average of the initial average love shock for men (0.0) and for women (-0.19). Thus, on average, the just-married women dislike their spouses, but nevertheless marry because of risk sharing and economies of scale in consumption. The average value of the love shock across all couples is much larger, though, and equals 0.54, reflecting selection on love in marriage survival.


The marriages in the model are quite unequal, as can be seen on Figure 1b where we show the cross-section of the wife's Pareto weight λ . This distribution is in a stark contrast to the unitary model that would assign equal weights to both spouses (that is, $\lambda = 1 - \lambda = 1/2$). In our model, the average λ is 0.353 with standard deviation of 0.18. Men being on average more productive than women introduces a right skew in the distribution of the Pareto weight: the share of couples for which λ is below 0.4 equals 0.64 and for which λ exceeds 0.6 is 0.13. Even if one were to take the

²³Over-predicting marriage at older ages naturally follows from assuming identical life spans; in reality, shorter male life expectancy lowers late-age marriage rates.

²⁴For completeness, we show life-cycle profiles of average labor supply, consumption, and assets on Figure A.1.

²⁵Since in the model we observe the randomly met partner's characteristics, we can also run a regression to further disentangle love from search frictions, see Table S.1 in the Supplemental Appendix. When also controlling for partner characteristics, the share of correctly predicted marriages goes up to 48%, which is still substantially below 100%. This illustrates that the love shocks play an important role.

FIGURE 1: Marriage and Inequality



Notes: Panel (a) displays the share of individuals of any given age who are married. Panel (b) shows the distribution of Pareto weights of the woman in across non-retired couples in the benchmark model. Panel (c) shows changes in own Pareto weight over a 5 year period conditional on being with the same partner. Panel (d) shows the decomposition of log-variance of consumption over the life cycle.

unitary household model as a rough approximation of household decision making, and consider the interval $[0.4, 0.6]$ to represent values reasonably close to equal weighting, this would still only describe about a fifth of couples in our model.²⁶

The model also predicts that the re-bargaining is rather rare and changes to λ are moderate.²⁷ The average share of couples who re-bargain across all cohorts in any given year is about 9%, and 23% of couples renegotiate λ within any 5 year window. Conditional on re-bargaining, the changes in the Pareto weight are small, with standard deviation of 0.13, see Figure 1c for illustration. As a result of this, the correlation of individual Pareto weight across a five-year horizon is 0.9.

Thus, since the majority of agents are married, and the Pareto weight exhibits significant per-

²⁶In line with Equation (8) and calibrated values of γ and ρ , this range of λ yields consumption splits between approximately 45/55 and 55/45.

²⁷This is qualitatively consistent with Lise and Yamada (2018), who also find that Pareto weights change infrequently.

TABLE 2: Determinants of initial Pareto weights

	Initial Pareto weight			
	(1)	(2)	(3)	(4)
Female	-0.301*** (0.003)	-0.086*** (0.001)	-0.257*** (0.002)	-0.090*** (0.001)
Wage Ratio		0.957*** (0.003)		0.867*** (0.003)
Asset Ratio			0.943*** (0.010)	0.351*** (0.005)
Age FE	Yes	Yes	Yes	Yes
N	18,886	18,886	18,886	18,886
R^2	0.373	0.906	0.566	0.928
Within- R^2	0.369	0.905	0.564	0.928

Notes: This table shows the determinants of the initial intra-household Pareto weight. The regressions are estimated by OLS using simulated data from the model. The dependent variable is the level of the initial Pareto weight (λ for women and $1 - \lambda$ for men). The intercept has been suppressed. The sample consists of couples at the time of marriage formation. Wage and asset ratios are defined as $x_i/(x_i + x_j)$ for the corresponding outcome. Age fixed effects included in all columns.

sistence, it is important to learn what determines λ upon marriage formation in order to understand inequality in the model. To this end, we use model-simulated data and focus on observations of newly formed marriages, regressing the λ newlyweds choose on various combinations of regressors including gender, relative wages and asset holdings, and age fixed effects. The results of this exercise are in Table 2. The main take-away here is that wages correlates strongly with Pareto weights at the time of marriage formation. Controlling just for gender and relative wages explains about 91% of variation in initial Pareto weight (column 3). By contrast, assets explain a much lower share of variation in λ of the newlyweds.

As the last exercise in this section, we provide information on the dispersion of individual log-consumption in the model. Focusing on the population of married individuals, we can apply a standard within/between-household decomposition of consumption inequality (e.g. [Lise and Seitz \(2011\)](#)):

$$\text{Var}[\log c_{it} \mid \text{married}] = \underbrace{\mathbb{E}[\text{Var}(\log c_{it} \mid i \in j)]}_{\text{within-household}} + \underbrace{\text{Var}[\mathbb{E}(\log c_{it} \mid i \in j)]}_{\text{between-household}}. \quad (9)$$

This formula decomposes the overall variance of (log) consumption into the average within-household variance and the variance of the household-level averages. To get a formula for the population variance - which includes both married individuals and singles - one can further ap-

TABLE 3: Variance decomposition of individual log-consumption

	Value	Fraction (%)
Between Married	0.1290	46.9
Within Married	0.0463	16.9
Between Singles	0.0812	29.6
Between Married and Singles	0.0183	6.65
Population Variance	0.2740	100.0

Notes: The table shows the results for the variance decomposition from Equations (9) and (10) based on model-simulated data. The sample is the non-retired population. Fractions are percentage contributions to the total variance.

ply a within/between decompositions for these two groups:

$$\begin{aligned} \text{Var}[\log c_{it}] = & p_m \text{Var}[\log c_{it} \mid \text{married}] + (1 - p_m) \text{Var}[\log c_{it} \mid \text{singles}] \\ & + \text{Var}[\log c_{it} \mid \text{between married and single}]. \end{aligned} \quad (10)$$

Here, p_m is the share of married individuals in the population. The last component of equation 10 measures how different the means of singles and married individuals are from the population-wide mean:

$$\begin{aligned} \text{Var}[\log c_{it} \mid \text{between married and single}] = & (1 - p_m)(\mathbb{E}[\log c_{it} \mid \text{single}] - \mathbb{E}[\log c_{it}])^2 \\ & + p_m(\mathbb{E}[\log c_{it} \mid \text{married}] - \mathbb{E}[\log c_{it}])^2, \end{aligned}$$

A key part of this component is economies of scale, which increase consumption of married individuals compared to singles. In addition, it also captures selection into marriage, as the type compositions of the groups of married and single individuals may differ. Combining the two within-between decompositions (Equations (9) and (10)), this effectively leads to a decomposition of the population variance that has four terms: (1) within-married inequality, (2) between-married inequality, (3) inequality between singles and (4) inequality between married individuals and singles.

Figure 1d shows the results from this decomposition for consumption inequality over the life cycle. There are five key insights. Firstly, as is standard for life-cycle models, inequality fans out with age, as more individuals are hit by shocks. Secondly, the between-couples component is initially very small, since all agents start single, but increases steadily, and contributes to about half of inequality among older agents.²⁸ Thirdly, the within-couple component, initially also negligible, accounts for about a fifth of dispersion among older couples. Fourthly, the relative contribution of inequality due to singles, while initially the largest, continuously decreases over the life cycle. Fifthly, the component capturing the difference in mean consumption between couples and sin-

²⁸Thus, note that the within and between couple components, as well as the between-singles component, are weighted by the relevant population share in Figure 1d.

gles is rather small, meaning that economies of scale make a limited contribution to inequality in our model.

Averaging across all age groups, the contributions of each component are summarized in Table 3. The within-couples component accounts for 26.5% of consumption dispersion among the married. Using British data and focusing on childless couples, [Lise and Seitz \(2011\)](#) find that intra-household inequality accounts for 25% of consumption inequality among couples. This is strikingly close to what our model implies, suggesting that our model generates a plausible level of intra-household inequality.

3.3 Transmission of Wage Shocks

Next, we investigate the consequences of idiosyncratic shocks to wages.²⁹ We begin by assessing how wage shocks affect consumption and labor, differentiating agents by marital status and gender. Then, we estimate the effects of these shocks on marriage formation, bargaining between spouses, and divorce. We show that our model implies similar transmission of shocks to earlier work where comparable, that is, for married couples, and present novel results for outcomes and populations that earlier work didn't address.

We regress the outcome of interest X_{jt} e.g., consumption or labor supply, on transitory and permanent shocks hitting members of household j :

$$X_{j,t} = f \left(\kappa_0^C + \sum_{g \in \{f,m\}} \kappa_{X,\nu_g}^C \Delta \nu_{gt} + \sum_{g \in \{f,m\}} \kappa_{X,\varepsilon_g}^C \varepsilon_{gt} + \chi_{jt} \right), \quad (11)$$

using model-simulated data.³⁰ We are interested in the signs and magnitudes of estimated coefficients which measure the elasticities of outcomes to shocks. To provide a quantitative assessment of the importance of limited commitment in marriage, we also estimate Equation (11) under the assumption of no adjustment of Pareto weight to shocks.³¹ Then, we also run analogous regression for all singles in the simulated sample:

$$X_{i,t} = f \left(\kappa_0^S + \kappa_{X,\nu_{g(i)}}^S \Delta \nu_{it} + \kappa_{X,\varepsilon_{g(i)}}^S \varepsilon_{it} + \chi_{it} \right). \quad (12)$$

In the main text, we focus on elasticities with respect to permanent labor productivity shocks.³² For external comparison, we benchmark our simulated elasticities against the estimates of [Blundell et al. \(2016\)](#). To make our results comparable to theirs, we condition on couples being together in both periods, as well as on age being between 30 and 57.³³

²⁹We conduct analogous experiments regarding love shocks in Supplemental Appendix [S.B.3.1](#).

³⁰The function $f(\cdot)$ is linear for continuous outcomes, corresponding to OLS, or logistic, for binary outcomes.

³¹That is, for each couple-year observation, we keep the Pareto weight fixed and equal to that implied by the model in the previous period.

³²The results on transmission of temporary productivity shocks are presented in Supplemental Appendix [S.B.3.2](#).

³³In Supplemental Appendix [S.B.3.3](#) we illustrate that our results are quite similar to those of [Wu and Krueger \(2021\)](#) as well.

TABLE 4: Effects of Permanent Wage Shocks on Consumption and Hours

Outcome	Woman Shock			Man Shock		
	Model	No λ adj.	BPS	Model	No λ adj.	BPS
<i>Couples</i>						
Household Consumption	0.20	0.20	0.19	0.34	0.34	0.32
<i>Married Women</i>						
Hours	0.46	0.51	0.42	-0.53	-0.59	-0.75
Consumption	0.23	0.20	-	0.31	0.34	-
<i>Married Men</i>						
Hours	-0.19	-0.21	-0.22	0.14	0.17	-0.08
Consumption	0.18	0.20	-	0.36	0.34	-
<i>Singles</i>						
Hours	0.90	-	-	0.91	-	-
Consumption	0.58	-	-	0.62	-	-

Notes: Reported coefficients from regressions (11) and (12) on permanent wage shocks with the dependent variable being the first-difference of logarithm of the Outcome column, estimated with OLS. The BPS columns reproduce estimates of [Blundell et al. \(2016\)](#). All estimated coefficients are significant at $p = 0.01$ level. We suppress the significance stars and standard errors for readability.

3.3.1 Consumption and labor supply

Taking $X_{j,t}$ in Equation (11) and $X_{i,t}$ in Equation (12) as log-changes in consumption and hours work, we present the results of this exercise in Table 4. The key insight - which we elaborate on below - is that re-bargaining in response to wage shocks (via adjustments in Pareto weights) alters individual labor supply and consumption, but these changes tend to offset each other within the household. As a result, aggregate household consumption responds similarly to the unitary household model of [Blundell et al. \(2016\)](#). In contrast, individual-level outcomes and hours worked are sensitive to the internal reallocation mechanisms induced by limited commitment.

Moving on to the response of hours worked, and, indirectly, labor income, our model is close to the estimates of [Blundell et al. \(2016\)](#), particularly for women's productivity shocks.³⁴ The magnitudes of response to male shocks in our model differ slightly from what they find. In our model, men's productivity shocks lead to a small increase in work hours (0.14 pp.), which reinforces the impact of the shock on labor income, whereas [Blundell et al. \(2016\)](#) find a small decrease in hours (-0.08).³⁵ Next, women's elasticity is -0.53, about two-thirds in absolute value of what they find. The cross-elasticities of hours are important, as they capture the added worker effect. Henceforth, our model implies a weaker added worker effect of female labor supply.

The weaker added worker effect is to a certain extent due to limited commitment. To see this, suppose that a husband is hit by a negative shock. He then reduces his labor supply and his wife

³⁴The own labor income elasticity to own shock equals that of hours plus one. The cross-elasticities of labor income are equal to those of hours'.

³⁵This is similar in [Wu and Krueger \(2021\)](#), who find that men's permanent shocks increase their hours by 0.16 pp.

increase hers. However, some shocks cause the Pareto weight of the man to deteriorate. This implies an increase in his labor supply and a decrease in hers. Therefore, the adjustment of the Pareto weight offsets the labor supply changes implied by unitary models. In the scenario of no adjustment in the Pareto weights, we find that women’s hours increases by 0.59 pp. in response to a 1 pp. negative shock for men, rather than by 0.53 pp. implied by our baseline model. Therefore, limited commitment is responsible for about one-third of the gap between our model and the unitary-model -based results of [Blundell et al. \(2016\)](#). The other side of the same coin is that the man has to cut hours worked by *less* with limited commitment.

It is interesting to observe that the dampening of the added worker effect does not result in a change in overall consumption insurance: the transmission of male and female shocks into consumption is 0.20/0.34 regardless of whether the Pareto weights are allowed to adjust. While limited commitment implies that spousal labor supply provides less insurance, this is offset by an increase in labor supply of the person who is hit by the negative shock. Therefore, our model highlights the role of both men’s and women’s labor supply and demonstrates that the effects of limited commitment on both partners can roughly offset each other in terms of their effect on household consumption insurance.

As for the *individual-level* consumption, male wage shocks benefit men more strongly than women. Conversely, female wage shocks benefit women more than men. Intuitively, this captures imperfect risk sharing due to limited commitment and changes in bargaining position following the shock, as individuals keep part of their income for themselves when experiencing a positive shock. Next, because men work more hours and earn more than women, the elasticities of consumption to men shocks are larger.

To assess the plausibility of these patterns, we compare them to findings from [Voena \(2015\)](#), who estimates a structural model of household behavior under different divorce regimes. She reports model-based transmission coefficients of wage shocks into individual consumption, which reflect the degree of partial insurance within couples. Under a unilateral divorce regime, she finds that the transmission coefficient of men’s shocks into men’s consumption is 0.39, and the corresponding coefficient for women’s consumption with respect to women’s shocks is 0.192. These figures are close to the 0.37 and 0.24 coefficients generated by our model. We conclude that our calibrated model produces a plausible degree of limited commitment, which is in line with previous literature.³⁶

Our model permits us to go beyond the married population, and estimate the transmission coefficients for singles, for whom the insurance due to the presence of another earner is absent by definition. These results are also shown in Table 4. Intuitively, the impacts of permanent productivity shocks on singles are significantly stronger both for hours and consumption.³⁷

³⁶Note that [Voena \(2015\)](#) does not report cross-effects, for example how men’s shocks are transmitted into women’s consumption.

³⁷In Supplemental Appendix S.B.3.4 we exploit the richness of our model further and report the gendered life cycle profiles of consumption insurance for the married and the singles, distinguishing sub-populations according to their marital status dynamics (remaining married, remaining single, marrying, divorcing).

TABLE 5: Permanent Wage Shocks and Marriage-Related Outcomes

	Marriage	Partner wage	Initial Pareto weight	Divorce	Δ Pareto weight
	(1)	(2)	(3)	(4)	(5)
Women					
Single	-0.254*** [-0.033] (0.059)	-0.035 (0.036)	0.689*** (0.032)	–	–
Married	–	–	–	0.696*** [0.027] (0.067)	0.055*** (0.001)
Men					
Single	0.448*** [0.057] (0.066)	-0.011 (0.044)	0.423*** (0.025)	–	–
Married	–	–	–	-0.797*** [-0.030] (0.075)	0.041*** (0.001)

Notes: Reported coefficients on permanent wage shock and, for logit regressions, average marginal effects in square brackets. Coefficients are from regressions (12) (cols 1–3) and (11) (cols 4–5). Columns (1) and (4) estimated by logit; other by OLS. In columns (2) and (3), the dependent variable is logged. In column (5), the dependent variable is log-difference of individual Pareto weight. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

3.3.2 Marriage

Wage shocks can also have effects through the marriage market, as they might affect the probability of getting married or divorced, or influence the Pareto weights within marriage. To explore these channels, we estimate the equation (12) for singles with three different outcomes. Firstly, we take the population of single agents and regress a marriage dummy (equal to one if the agent marries). Secondly, for the sub-population of singles who get married, we regress logarithms of their partner’s wage and of their initial Pareto weight on shocks. Then, we estimate equation (11) with $X_{j,t}$ first as the divorce dummy and secondly, as the log-difference of their Pareto weight. The results of this exercise are presented in Table 5.

Starting with Pareto weights, the results show that wage shocks before marriage are transmitted into the initial Pareto weight at the time of marriage, and to a lesser extent to the Pareto weight of existing couples. For newly formed marriages, our estimates (column 3 of the Table) imply that a 1 percentage point increase in the woman’s wage due to permanent shock increases her Pareto weight by 0.69 percentage point and 0.43 pp. for men. Translated into consumption shares, in a couple that in an absence of shocks would agree, for example, on an equal split of consumption, a one standard deviation positive shock to the woman’s wage would increase her share in total household consumption by 3.4 percentage points. This is equivalent to one day worth of household consumption per month.³⁸ Regarding continuing couples, the effects are tenfold smaller (see Table S.9 in Supplemental Appendix). This is so because not all wage shocks trigger renegotia-

³⁸For a man, the effect is 1.8 pp., see Table S.8 in Supplemental Appendix for more details and results for other values of λ in case of no shock.

tion. Thus, our model implies a quantitatively important transmission channel of wage risk that operates predominantly through the initial Pareto weights.

Besides the impact of the wage shocks on Pareto weights, Table 5 also shows how wage shocks affect marriage and divorce probabilities in the calibrated model. The main result is that these effects are asymmetric by gender: when women are hit with a positive permanent wage shock, they are *less* likely to marry and *more* likely to divorce. The converse is true for men. These effects can be interpreted as wage shocks shifting the value of marriage and the value of singlehood to a different extent. For example, when women experience a positive wage shock, this raises the outside option of being single more strongly compared to the value of marriage (since single women work more than married women), and they become less likely to marry. The quantitative magnitude of the estimated effects is that a one standard deviation permanent wage shock increases the likelihood of getting married for men by about 1 percentage point and decreases it for women by approximately 0.6 pp., while the effect on likelihood of divorce is similar in absolute values for both genders and equal to about half a percentage point. Thus, the estimated effects of on marriage and divorce are quite modest.³⁹ Interestingly, the effect of wage shocks materialize solely through the bargaining position of spouses, and not their selectivity along prospective partner's wage.

4 Counterfactual Exercises

In this section we use our model to conduct two counterfactual experiments. Firstly, we use the model as a measurement device to assess the consequences of the idiosyncratic labor income risk on aggregates and on welfare. Secondly, we investigate the consequences of a balanced-budget introduction of a proportional tax on labor income. We contrast the predictions of our model with the unitary household model. The unitary household model abstracts from limited commitment in marriage and ignores marital status dynamics. Thus, to surgically isolate the insurance benefits of marriage, we also consider a model in which all agents are born and remain single.

In the first exercise, the removal of wage risk substantially decreases consumption inequality and the dispersion of the Pareto weights. We also find that the unitary model severely understates the cost of wage risk, capturing only two-thirds of the cost implied by our model.

Regarding the tax reform, in our model the labor supply of married women is much more elastic than in the unitary model. Because of that, the switch to a flat tax imposes a greater welfare cost on women than it does on men.

4.1 Aggregate and Distributional Consequences of Wage Risk

We consider a version of our model in which there are no wage shocks, the age-gender wage profile preserves the mean wage at each age t for both genders, and other parameters of the model

³⁹We discuss details of patterns of participation constraints binding and the incidence of shocks in Supplemental Appendix S.B.4.

TABLE 6: The Welfare Cost of Wage Risk: Comparison of Models

	Model		
	Benchmark	Unitary Household	No Marriage
Population	−22.5	−14.2	−27.9
Women	−24.4	−17.5	−29.4
Men	−19.9	−10.6	−25.4

Notes: This table shows the welfare cost of wage risk (measured as consumption equivalent variation in percentages) in three alternative models relative to their no-risk versions. Parameters of all models kept at the values calibrated for the benchmark model.

are as in Table 1. Thus, the remaining sources of uncertainty are the love shocks and search frictions in the marriage market. The heterogeneity among prospective partners is now solely induced by their marital histories, leading to heterogeneity in asset holdings.

As a prelude to the assessment of agents’ valuation of the cost of wage risk, we compare consumption, labor supply, and marriage decisions between our benchmark model specification and its no wage risk version. The details on this comparison are in Table A.3 in Appendix A.3.1. On average, hours worked do not respond, and individual consumption increases mildly. In the absence of wage risk, agents save predominantly to consume in retirement so the overall stock of assets goes down. The marriage rate is virtually unaffected, and so is the average Pareto weight λ . However, the presence of wage risk increases consumption inequality roughly threefold.⁴⁰

Next, we ask: starting with the no-risk benchmark, how costly the introduction of wage risk is to agents? To highlight the importance of marital status dynamics and limited commitment, we conduct an analogous exercise of shutting down wage risk in the unitary household model and in a model without marriage.⁴¹ In all three cases, we keep the calibrated values of parameters as in the benchmark model for ease of comparison.⁴²

We measure the cost of wage risk by consumption equivalent variation, a factor that scales consumption in a model without wage risk to keep newborns indifferent between models. Formally, let’s fix one of the three models (benchmark, unitary household, no marriage) and denote the time-zero expected discounted lifetime utility of an agent i in a risk scenario $\mathcal{R}$, and consumption scaling factor χ as $V_i^{\mathcal{S}}(\chi)$. Here we think of $\mathcal{R}$ as model featuring the wage risk or not. Then:

$$V_i^{\mathcal{R}}(\chi) = E \left[\sum \beta^t \left(u(\chi \cdot c_{it}^{\mathcal{R}}, h_{it}^{\mathcal{R}}) + \theta_{it} \cdot \mathbb{1}_{\{\text{married in } t\}}^{\mathcal{R}} \right) \right].$$

⁴⁰Further details, including life-cycle profiles of key equilibrium objects on Figure A.2 are in Appendix A.3.1.

⁴¹It can be shown that the couples’ decision-making in the unitary model is equivalent to a model with a standard household-level utility function in consumption with appropriate rescaling. This is due to the spouses having equal and constant bargaining power, see Appendix A.3.2 for details.

⁴²We have also investigated re-estimating the unitary household model, our main reference point, and found that it essentially did not matter for our results.

TABLE 7: The Welfare Cost of Wage Risk: Decomposition

	Benchmark		Unitary Household		No Marriage	
	Women	Men	Women	Men	Women	Men
(A) Total	−24.4	−19.9	−17.5	−10.6	−29.4	−25.4
(B) Consumption	−17.1	−17.9	−11.7	−11.7	−25.4	−22.6
(B1) Level	0.1	−1.7	4.7	4.7	−2.8	−1.4
(B2) Distribution	−17.2	−16.4	−15.7	−15.7	−23.3	−21.5
(C) Hours	−5.4	−1.3	−6.6	1.3	−5.4	−3.6
(C1) Level	−0.8	1.2	−0.9	6.6	−2.8	−1.5
(C2) Distribution	−4.7	−2.5	−5.7	−5.0	−2.7	−2.1
(D) Love	−3.6	−1.2	-	-	-	-

Notes: This table shows the welfare cost of wage risk (measured as consumption equivalent variation in percentages) in three models. Note that the decomposition adds up multiplicatively: the CEV is defined such $(1 + A) = (1 + B)(1 + C)(1 + D)$, $(1 + B) = (1 + B1)(1 + B2)$ and $(1 + C) = (1 + C1)(1 + C2)$.

The consumption equivalent variation CEV is then defined as:

$$V_i^{\text{no risk}}(1 + CEV) = V_i^{\text{with risk}}(1).$$

Note, this calculation can be done for the whole population of newborns, or separately by gender.

Table 6 summarizes the main results on the cost of wage risk where we report the welfare measure in percentages, $100 \times CEV$. In our model, the agents are willing to give up to around 22 percent of their consumption to avoid wage risk. As for the auxiliary models, it is no surprise that the unitary model and the no-marriage model yield the lower and the upper bound, respectively, for the welfare costs of wage risk, given their extreme stands on the prevalence of marriage. However, the comparison of these two models with ours illustrates that the assumptions underpinning these models lead to quantitatively important biases. The unitary household model misses *a half* of the costs of wage risk while the no-marriage model overestimates them by one-fourth.⁴³ Across all model specifications, the welfare cost of wage risk is consistently higher for women than for men. Interestingly, the largest gender gap appears in the unitary household model, where consumption is equalized among spouses. To better understand the sources of these differences, we decompose the welfare measure into the effects of consumption, hours, and love. The first two components can be further decomposed into their level (change in average value) and distribution (change in dispersion) effects.⁴⁴

⁴³The cost of wage risk in our model is not simply a weighted average of the two extremes based on the marriage rate in our model because risk is not fully pooled within couples. Taking the marriage rate in our model as the weight on the unitary model, a simple weighted average of welfare costs of wage risk in the unitary and no marriage models is roughly 18%.

⁴⁴That the CEV measure can be decomposed like this follows from the homotheticity of preferences with respect to

TABLE 8: Welfare Consequences of a Proportional Tax Reform

	Model		
	Benchmark	Unitary Household	No Marriage
Population	−6.1	−4.0	−6.2
Women	−8.4	−3.0	−8.4
Men	−2.9	−5.0	−2.7

Notes: This table shows the welfare effect (measured as consumption equivalent variation in percentages) of a switch to proportional tax system in three alternative models relative to the progressive taxation case.

The main takeaways from this decomposition, the details of which are presented in Table A.5, are as follows. Firstly, in the no-marriage model, women suffer to a greater extent from wage risk because they face a more volatile income process and they dislike labor supply more.

Compared with the no-marriage model, the welfare cost of wage risk due to consumption is halved in the unitary model. Reduction in consumption cost of wage risk stems from improved risk sharing, and higher consumption at early age. The change in the consumption cost of wage risk is almost fully responsible for the difference in the welfare cost of wage risk between no-marriage and unitary model. However, the gap in cost of wage risk across genders *increases* in the unitary model despite consumption equality between spouses. This is due to relocation of work within the household: women work more in the model with wage risk, while men work less.

Finally, in our benchmark model, consumption cost increases, relative to the unitary household model. The level effect is now negligible, and the distribution of consumption component increases. However, the gender gap in labor supply cost of risk shrinks. These are two sides of the same coin - limited commitment and bargaining. Couples do not share consumption equally, but bargaining limits the extent to which cost of risk can be pushed onto the other spouse. Last, but not least, while utility from love was absent in the no-marriage and unitary household model, it contributes rather little to the increase in cost of wage risk.

4.2 A Proportional Tax Reform

Finally, we study how income tax progressivity affects insurance and welfare in our model. We conduct counterfactual experiment in which we replace the current tax system by a linear tax for both couples and singles ensuring a balanced budget. Thus, given taxable income y , the taxed amount is $T(y) = \tau y$. We then conduct a similar experiment in the unitary household and no marriage models for comparison. In the benchmark model, the balanced-budget switch to proportional taxation requires setting $\tau = 0.207$.⁴⁵

The first step in our analysis is to assess the population- and gender-specific welfare change

consumption. Intuitively, welfare differences arise from the level and distributional effects of consumption and hours, given the concavity of $u(c, h)$, and from the level effects of love shocks, which enter linearly. We provide a formal derivation and discussion of this decomposition in Appendix A.3.3.

⁴⁵In the unitary and no-marriage model, the balanced-budget value of τ stands at 0.204 and 0.163, respectively.

TABLE 9: Proportional Tax Reform: Decomposition of Welfare Effects

	Benchmark		Unitary Household		No Marriage	
	Women	Men	Women	Men	Women	Men
(A) Total	−8.4	−2.9	−3.0	−5.0	−8.4	−2.7
(B) Consumption	−1.4	1.2	0.4	0.4	−3.2	0.7
(B1) Level	6.1	6.6	5.8	5.8	2.9	6.5
(B2) Distribution	−7.0	−5.1	−5.1	−5.1	−5.9	−5.5
(C) Hours	−7.4	−3.7	−3.3	−5.4	−5.4	−3.3
(C1) Level	−6.8	−3.3	−3.2	−4.8	−4.6	−2.7
(C2) Distribution	−0.6	−0.4	−0.1	−0.6	−0.9	−0.6
(D) Love	0.2	−0.3	−	−	−	−

Notes: This table shows the welfare effect of the introduction of a flat tax, i.e. the consumption equivalent variation such that agents are indifferent between the flat tax and the benchmark model. The structure follows Table A.5.

following the reform, employing the methodology outlined in the previous section. We present the results of this exercise in Table 8. We find that overall in all three models agents dislike the switch to the proportional tax system. Conversely, the greater the welfare cost of the reform, the higher the welfare benefits of tax progressivity. In our model, the benefits of tax progressivity are significantly larger than in the unitary household model. In fact, the benchmark case quantitatively resembles more the no-marriage model, not only population-wise, but also across genders.

To understand these results, we again employ a decomposition of the welfare measure, analogously to the first counterfactual exercise, see Table 9 for details. The main finding here is that almost all of the difference between ours and the unitary model comes from the level of labor supply – women work significantly more and men slightly less in our model than in the unitary model – and the distribution of consumption of women becoming more unequal.

Comparing the benchmark to the no-marriage model, we find that the essentially identical evaluation of the tax reform across genders masks lesser level cost for women in the benchmark balanced by higher labor supply cost. For men, the situation is similar - the positive contribution of consumption is greater in the benchmark model, but the labor supply cost is greater too.

Thus, the key to understanding the effects of tax progressivity in our model is in why women value it so much more than men. To this end, we complement the welfare decomposition with the life cycle and cross-sectional effects of the tax reform, see Table A.4 and Figure A.3 in the Appendix. There are four main findings here which shed light on the working of tax progressivity in the model.

Firstly, as a result of the reform, the marriage rate decreases by about one percentage point. While this is not much, it serves as a cautionary note. Because of the differential tax treatment of singles and couples, tax reforms can affect the tax base composition through the marriage mar-

TABLE 10: Composition of Labor Supply Response to Proportional Tax Reform

Group	Benchmark	Unitary Household	No Marriage	Convex Combination
Married Women	42	44	-	29
Married Men	12	56	-	37
Single Women	24	-	62	20
Single Men	22	-	38	12

Notes: This table presents the share (in %) of each group in the total change in hours worked following a proportional tax reform. "Convex Combination" reflects the group's weight in a convex combination of the Unitary and No Marriage models, using the marital status distribution in the Benchmark model.

ket.⁴⁶

Secondly, while overall asset holdings increase, the relative increase is particularly pronounced for single men. This growing asset inequality plays a crucial role in the third key finding: the flat tax significantly raises the dispersion of individual consumption, by about a fourth, as well as the dispersion of Pareto weights, thereby amplifying intra-household inequality. These two outcomes are intertwined. On the one hand, the proportional tax exacerbates the inequality in potential earnings, which directly widens the gap in bargaining positions and thus the distribution of Pareto weights. On the other hand, the resulting asset inequality among singles also affects matching patterns in the marriage market: greater asset dispersion raises the likelihood that individuals with unequal wealth levels pair up, further contributing to within-household inequality through more imbalanced bargaining positions. In this way, asset inequality and intra-household inequality reinforce one another.⁴⁷

The fourth finding links directly to the welfare decomposition which demonstrated that women end up working significantly more than men following the tax reform. We find that this is driven predominantly by the response of labor supply of married women. Following [Guner et al. \(2012\)](#), we zoom in on the labor supply response to the tax reform. Table 10 contains the information on the relative contributions of each marital status sub-populations when taxes are changed from progressive to proportional. The fourth column of this table contains the contributions implied by a model in which the share of the married agents is as in the benchmark, but there is no marriage market. Such a *convex combination* of the unitary and the no-marriage models mimics their assumption of fixed marital statuses. The share of married women in the overall response of labor supply in our model stands at 42% (while their share in overall population is about 35%). In [Guner et al. \(2012\)](#), this value is around 51% but roughly half of it is due to extensive margin. Thus, we conclude that our model, with within-household inequality, implies *an even stronger* response of

⁴⁶For example, we have experimented with imposing the flat tax rate that we computed in the balanced-budget reform solely on couples and found that it reduced the marriage rate by a bit more than 3 percentage points. Essentially, one in twenty marriages would not form in response to such a policy change. Further details are available from the authors upon request.

⁴⁷We decompose the increase in consumption inequality following from the tax reform into the four components outlined in Equations (9) and (10) in Table A.5 in the Appendix and find that while taxes predominantly increase inequality between households, they also increase it *within* couples.

hours worked of married women along the intensive margin. Women work more to improve their outside option upon marrying, but also when married, since on average their bargaining position is weaker than their spouses.

5 Conclusions

In this paper, we propose a model nesting the Bewley framework in an equilibrium marriage market with limited commitment in marriage to study the transmission and costs of idiosyncratic wage risk. We go beyond the unitary representation of stable couples which were the focus of earlier work, accounting for realistic dynamics of marital status and the intra-household division of consumption. We show that the former underestimates the welfare cost of wage risk by about a half, and because of that also underestimates the gains from progressive taxation.

Our paper suggests several important avenues for future research. Extending our framework to include richer dimensions of household time allocation which play an important role in shaping both intra-household inequality and labor supply responses is one such avenue. Another interesting possibility would be to study the interaction between wage risk and wealth accumulation, for example through housing, pensions, or other illiquid assets, which could further amplify or mitigate intra-household inequality.

A Appendix

A.1 Calibration and Model Fit

Below, we provide information on model fit on targeted moments in Table A.1. Average labor income of men is normalized to 1. We take the moments for women’s average labor income and relative hours from Wu and Krueger (2021). Marriage rates and divorce probabilities computed from PSID data. We calibrate the model using the standard method of moments where each moment gets weight 1, except for marriage rates which are weighted with the inverse of the number of years we match. We do not match marriage rates after retirement since marital decisions at these ages are to an important extent driven by forces external to the model (e.g., mortality). We simulate 200,000 agents to compute sample statistics and we use standard settings for the Powell’s method regarding convergence.

TABLE A.1: Model fit


Target	Data	Model	Parameter
Average labour income, men	1.000	0.980	ϕ_m
Average labour income, women	0.491	0.465	ϕ_f
Ratio of female to male hours	0.733	0.709	κ_g
Divorce probability	0.040	0.040	$\sigma_\theta^2, \rho_\theta, \mu_\theta$
Marriage rates by age	See Fig. 1a	See Fig. 1a	$\sigma_\theta^2, \rho_\theta, \mu_\theta$

Below, we plot qualitatively standard average consumption, asset holdings, and labor supply per capita over the life cycle on Figure A.1. Qualitatively these are fairly standard and representative of life-cycle models with idiosyncratic risk and self-insurance. Initial increase in labor supply corresponds to marriage formation and precautionary motive.

A.2 Marriage and Inequality

We display the determinants of marriage decisions in the model in Table A.2 . The regressions are based on model-simulated data where the outcome is a dummy variable for whether an agent gets married. Across the six columns in this Table, we vary the combinations of regressors that correspond to characteristics that can be observed in the data: agents’ age, wages, and assets. We find that overall these observable characteristics contribute to marriage decisions rather moderately. Indeed, specifications in columns (1) and (4) which abstract from assets, result in a model that starkly predicts all singles in the estimation sample to not marry. In the richest specification, the observable characteristics of agents explain about half of marriage formation. Thus, marriage decisions are largely determined by love shocks.

FIGURE A.1: Average consumption, assets, and labor supply over the life-cycle.



A.3 Aggregate and Distributional Consequences of Wage Risk

A.3.1 Consumption, Labor, and Marriage

We present the impacts of wage risk overall in the cross-section by gender and marital status in Table A.3. Agents increase their savings in the face of wage risk and the consumption inequality increases markedly due to wage risk. Otherwise, the effects of wage risk on other cross-sectional moments are fairly small.

Additionally, we demonstrate the life-cycle impacts of wage risk on Figure A.2 to gain further insights in how it affects the equilibrium. Figure A.3b confirms the absence of the precautionary savings motive that was due to wage risk - average assets at each age are lower in the no-risk economy. Nevertheless, agents enter retirement with similar assets. The steep increase of average assets prior to retirement seems to be paid for more by the women than men, reflecting the within-marriage asymmetry, see Figures A.3c and A.3d.

Next, the life cycle-profile of hours indicates reallocation of work over the life cycle. In an economy without wage risk, agents first work less, and then more than in the benchmark, as shown on Figures A.3e and A.3f. With wage risk, a more productive agent has additional precautionary incentives to work more to increase their savings.

Furthermore, the dispersion of the Pareto weights decreases markedly, capturing the marriage market uncertainty and the effects of love shocks, see Figure A.2g. Finally, the life-cycle pattern

TABLE A.2: Marriage Predictors

	Marriage Formation					
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	-1.725*** (0.008)	-2.076*** (0.010)	-1.702*** (0.086)	4.739*** (0.135)	-2.141*** (0.011)	6.183*** (0.146)
Wage	0.024* (0.010)		0.025* (0.011)		-0.539*** (0.012)	-1.459*** (0.020)
Assets		0.881*** (0.009)		2.857*** (0.027)	0.962*** (0.009)	3.583*** (0.032)
Assets (squared)		-0.077*** (0.001)		-0.225*** (0.003)	-0.075*** (0.001)	-0.252*** (0.003)
Age			0.010* (0.005)	-0.247*** (0.007)		-0.303*** (0.008)
Age (squared)			-0.000*** (0.000)	0.001*** (0.000)		0.001*** (0.000)
N	124,843	124,843	124,843	124,843	124,843	124,843
Pseudo R^2	0.000	0.156	0.003	0.307	0.175	0.370
Accuracy	0.000	0.144	0.000	0.263	0.201	0.389
Precision	–	0.420	–	0.565	0.559	0.687

Notes: This table shows the coefficients of a logistic regression explaining the marriage formation dummy in a sample simulated from the model. Accuracy is defined as the share of the sum of true positives and true negatives in the sample. Precision is defined as the share of true positives in all positives.

of consumption inequality changes markedly as presented on Figure A.2h.⁴⁸ Older cohorts are, in fact, *more equal*. This captures the equalising effects of marriage which are much stronger now, due to absence of wage risk. It is still the case that men are more productive than women, which causes the population of singles to remain unequal. However, these differences disappear in marriages: all husbands earn same age-specific wage, and so do wives. Therefore, the between-couples component is almost nil, while inequality between genders is essentially relocated from singles to within-couples, with an additional bit coming from love shocks and marital histories affecting asset holdings. Divorce ameliorates the consumption differences between singles, and so the older generations are more equal.

A.3.2 Comparison of the benchmark and the unitary model

In this appendix, we describe in more detail how our model is related to unitary models once we make the additional assumption of an exogenous Pareto weight $\lambda = 0.5$. In [Wu and Krueger \(2021\)](#), households have the following per-period utility function, which is defined over *household-*

⁴⁸Note, the population-wide variance of log-consumption accounts for variation in consumption levels between cohorts as well, while Figure A.2h presents information on consumption dispersion *within* each cohort.

TABLE A.3: Consequences of Wage Risk in the Cross Section

Variable	Population	Married		Single	
		Women	Men	Women	Men
Consumption	-2.96%	-4.02%	-1.82%	-1.76%	-5.61%
Hours	-0.34%	+1.26%	-2.60%	+1.55%	+1.49%
Assets	+85.9%	+82.9%	+82.9%	+124.5%	+80.8%
Consumption Inequality	+231.3%	+904.5%	+1064.2%	+232.0%	+358.3%
Marriage Rate	+ 0.98%	–	–	–	–

Notes: The table presents changes of cross-sectional averages induced by the introduction of wage risk relative to the counterfactual case of no risk. The “Population” column reflects aggregate effects; other columns report subgroup effects by marital status and gender.

level consumption and the work hours of each partner:

$$u(C, h_f, h_m) = \frac{C^{1-\gamma}}{1-\gamma} - \phi_f \frac{h_f^{\eta_f}}{\eta_f} - \phi_m \frac{h_m^{\eta_m}}{\eta_m}$$

In our model, once we assume that Pareto weights are exogenous, time-invariant and equal to 0.5, the households maximize the per-period utility function:

$$\tilde{u}(c_f, c_m, h_f, h_m) = 0.5 \left(\frac{c_f^{1-\gamma}}{1-\gamma} - \phi_f \frac{h_f^{\eta_f}}{\eta_f} \right) + 0.5 \left(\frac{c_m^{1-\gamma}}{1-\gamma} - \phi_m \frac{h_m^{\eta_m}}{\eta_m} \right)$$

This differs from [Wu and Krueger \(2021\)](#) by using individual, rather than household-level, consumption. However, with fixed Pareto weights, the two specifications are isomorphic, since individual consumption is a fixed fraction of household-level expenditure. More concretely, the specification of economies of scale from [Voena \(2015\)](#) (see equation 3) implies that:


$$K = \left(\frac{\lambda}{1-\lambda} \right)^{\frac{1}{\rho-1+\gamma}}, \quad c_m = \frac{1}{(1+K^\rho)^{\frac{1}{\rho}}} C, \quad c_f = \frac{K}{(1+K^\rho)^{\frac{1}{\rho}}} C.$$

Here, λ is the Pareto weight, ρ is the parameter for economies of scale, and C is household expenditure. Thus, given $\lambda = 0.5$, $\gamma = 1.73$ and $\rho = 1.4$, this simplifies to $c_f = c_m = 0.61C$. With this, the household utility function simplifies to a function of household expenditure:

$$\tilde{u}(C, h_f, h_m) = 0.61^{1-\gamma} \frac{C^{1-\gamma}}{1-\gamma} - 0.5\phi_f \frac{h_f^{\eta_f}}{\eta_f} - 0.5\phi_m \frac{h_m^{\eta_m}}{\eta_m}$$

Thus, fixed Pareto weights imply that our specification is equivalent to the utility specification of [Wu and Krueger \(2021\)](#), up to a rescaling of the preference coefficients. Note that in the unitary version of our model, it is still possible to define individual-level welfare, since the household utility function is derived from the individual-level utility functions and household expenditure maps into individual consumption via the equations above.

FIGURE A.2: Elimination of Wage Risk: Implications for Life-Cycle Outcomes



Notes: This figure shows how the main life-cycle outcomes of men and women differ between the benchmark model and the counterfactual simulation without wage risk.

A.3.3 Welfare Decomposition

To decompose the welfare change into parts due to consumption, the disutility of work and the love shock, the first part is to define CEV_c as the consumption part of the CEV. This corresponds to the CEV if only the consumption utility changes and the remaining components stay identical to the status quo:

$$V_g^0(1 + CEV_c) = UC^1(1) + UD^0 + UL^0$$

The labor supply part of the CEV (CEV_h) is defined via the equation:

$$V_g^0((1 + CEV_h)(1 + \cdot CEV_c)) = UC^1(1) + UD^1 + UL^0$$

Finally, the love part of the CEV is defined via the restriction that multiplying all sub-CEVs results in the overall CEV:

$$(1 + CEV_l)(1 + CEV_h)(1 + \cdot CEV_c) = (1 + CEV)$$

Finally, the consumption and labor supply parts can also further be decomposed into a level and distribution part. Note that this is not relevant for the utility from the love shock, since utility from love is linear. For consumption, the level effect is defined via:

$$V_g^0((1 + \cdot CEV_{cl})) = UC^1\left(\frac{\bar{c}_1}{\bar{c}_0}\right) + UD^0 + UL^0$$

As a result, $CEV_{cl} = \frac{\bar{c}_1}{\bar{c}_0} - 1$. This is useful as it means that CEV_{cl} can directly be interpreted as the level change in consumption between simulations. The distributional change is then defined via $1 + CEV_c = (1 + CEV_{cl})(1 + CEV_{cd})$. Therefore, $CEV_{cd} = \frac{1 + CEV_c}{1 + CEV_{cl}} - 1$. For labor supply, the level effect is similarly defined:

$$V_g^0((1 + \cdot CEV_{hl})(1 + CEV_c)) = UC^1(1) + UD^0\left(\frac{\bar{h}_1}{\bar{h}_0}\right) + UL^0$$

$UD^0\left(\frac{\bar{h}_1}{\bar{h}_0}\right)$ is the status quo disutility from labor supply when hours are scaled up by a fixed factor. Finally, the distributional component can be backed out via $1 + CEV_h = (1 + CEV_{hl})(1 + CEV_{hd})$.

A.4 Proportional Tax Reform

A.4.1 Consumption, Labor, Marriage

Table A.4 contains the summary on changes relative to the progressive taxation benchmark induced by the balanced budget shift to proportional tax reform, while Figure A.3 presents information on its effects over the life cycle.

Regarding cross-sectional effects, agents work and consume more under the proportional tax system. This follows a straightforward logic - progressivity discourages labor supply. With the redistributive effects of progressivity gone, agents work and save more to insure themselves against

TABLE A.4: Consequences of the Tax Reform in the Cross Section

Variable	Population	Married		Single	
		Women	Men	Women	Men
Consumption	+5.7%	+6.3%	+6.2%	+3.0%	+6.5%
Hours	+6.2%	+10.1%	+3.7%	+7.1%	+3.7%
Assets	+14.0%	+13.0%	+13.0%	+15.3%	+20.1%
Consumption Inequality	+25.5%	+30.3%	+25.3%	+32.9%	+23.6%
Marriage Rate	-1.5%	–	–	–	–

Notes: The table presents percentage changes in cross-sectional averages induced by the tax reform relative to the benchmark. The “Population” column reflects aggregate effects; other columns report subgroup effects by marital status and gender. The marriage rate is reported in percentage points.

wage risk. Asset holdings increase much more strongly for singles, especially so for men. We then decompose the change in the consumption inequality following the tax reform, see Table A.5 below. While the bulk of the tax reform contribution to increase in consumption inequality happens through differences between households, a non-negligible part is due to inequality within the married households. This comes from taxes impacting on asset accumulation and on the outside options of spouses.

TABLE A.5: Proportional tax reform and decomposition of consumption inequality


	Post-Tax Reform	Changes	
		Δ to Baseline (%)	Cont. to Total Δ (%)
Between Married	0.1640	27.1	49.3
Within Married	0.0519	12.1	7.9
Between Singles	0.1081	33.0	37.7
Between Married and Single	0.0210	14.8	3.8
Total	0.3450	26.0	100.0

Notes: This table shows inequality within and between couples in the proportional tax simulations. Differences are expressed as percentage changes relative to the benchmark. Contributions are expressed as percentage shares of the overall change.

Over the life cycle, it becomes apparent that the labor supply increases are concentrated in the early years, while agents enjoy greater consumption nearing retirement. Interestingly, the increase in dispersion of Pareto weight initially gradually increases with age. However, this effect disappears as agents near retirement.

While after the reform the distribution of individual consumption is overall more dispersed, the increase in the between-married component is strongest. For example, at age 50, between-marriage dispersion rises from 0.150 to 0.192, an increase of about 28%, whereas within-marriage dispersion increases from 0.0524 to 0.0593, or roughly 13%. This pattern holds across the life cycle: between-marriage dispersion rises by 0.042 points from pre- to post-reform at age 50, compared to just 0.01 for within-marriage.

FIGURE A.3: Proportional Tax Reform: Implications for Life-Cycle Outcomes



Notes: This figure shows how the main life-cycle outcomes of men and women differ between the benchmark model and the counterfactual experiment with a proportional (flat) tax.

References

- BAYOT, D. AND A. VOENA (2014): "Prenuptial contracts, labor supply and household investments," *Unpublished manuscript*.
- BLUNDELL, R., P.-A. CHIAPPORI, T. MAGNAC, AND C. MEGHIR (2007): "Collective labour supply: Heterogeneity and non-participation," *The Review of Economic Studies*, 74, 417–445.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): "Consumption inequality and family labor supply," *American Economic Review*, 106, 387–435.
- (2018): "Children, time allocation, and consumption insurance," *Journal of Political Economy*, 126, S73–S115.
- BRONSON, M. A., D. HAANWINCKEL, AND M. MAZZOCCO (2024): "Taxation and Household Decisions: an Intertemporal Analysis," *Unpublished manuscript*.
- BROWNING, M., F. BOURGUIGNON, P.-A. CHIAPPORI, AND V. LECHENE (1994): "Income and outcomes: A structural model of intrahousehold allocation," *Journal of Political Economy*, 102, 1067–1096.
- CHERCHYE, L., B. DE ROCK, A. LEWBEL, AND F. VERMEULEN (2015): "Sharing rule identification for general collective consumption models," *Econometrica*, 83, 2001–2041.
- CHERCHYE, L., B. DE ROCK, AND F. VERMEULEN (2012): "Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information," *American Economic Review*, 102, 3377–3405.
- CHERCHYE, L., T. DEMUYNCK, B. DE ROCK, AND F. VERMEULEN (2017): "Household consumption when the marriage is stable," *American Economic Review*, 107, 1507–34.
- CHIAPPORI, P.-A., B. FORTIN, AND G. LACROIX (2002): "Marriage market, divorce legislation, and household labor supply," *Journal of Political Economy*, 110, 37–72.
- CHIAPPORI, P.-A. AND M. MAZZOCCO (2017): "Static and intertemporal household decisions," *Journal of Economic Literature*, 55, 985–1045.
- CHIAPPORI, P.-A. AND C. MEGHIR (2015): "Intrahousehold inequality," in *Handbook of Income Distribution*, Elsevier, vol. 2, 1369–1418.
- CISCATO, E. (2019): "The Changing Wage Distribution and the Decline of Marriage," *Unpublished manuscript*.
- CONESA, J. C. AND D. KRUEGER (2006): "On the optimal progressivity of the income tax code," *Journal of Monetary Economics*, 53, 1425–1450.
- DE NARDI, M., G. FELLA, AND G. PAZ-PARDO (2020): "Nonlinear household earnings dynamics, self-insurance, and welfare," *Journal of the European Economic Association*, 18, 890–926.

——— (2025): “Wage risk and government and spousal insurance,” *Review of Economic Studies*, 92, 954–980.

DE ROCK, B., M. KOVALEVA, AND T. POTOMS (2023): “A spouse and a house are all we need? Housing demand, labor supply and divorce over the lifecycle,” Tech. rep., IFS Working Papers.

FERNANDEZ, R. AND J. C. WONG (2017): “Free to leave? A welfare analysis of divorce regimes,” *American Economic Journal: Macroeconomics*, 9, 72–115.

FOERSTER, H. (2025): “Untying the Knot: How Child Support and Alimony Affect Couples’ Decisions and Welfare,” *The Review of Economic Studies*, 92, 3029–3066.

GAYLE, G.-L. AND A. SHEPHARD (2019): “Optimal taxation, marriage, home production, and family labor supply,” *Econometrica*, 87, 291–326.

GOUSSÉ, M., N. JACQUEMET, AND J.-M. ROBIN (2017): “Marriage, labor supply, and home production,” *Econometrica*, 85, 1873–1919.

GREENWOOD, J., N. GUNER, G. KOCHARKOV, AND C. SANTOS (2016): “Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation,” *American Economic Journal: Macroeconomics*, 8, 1–41.

GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2011): “Taxation and household labour supply,” *The Review of Economic Studies*, 79, 1113–1149.

——— (2012): “Taxing women: A macroeconomic analysis,” *Journal of Monetary Economics*, 59, 111–128.

——— (2023): “Rethinking the welfare state,” *Econometrica*, 91, 2261–2294.

GUVENEN, F., S. OZKAN, AND R. MADERA (2024): “Consumption dynamics and welfare under non-gaussian earnings risk,” Tech. rep., National Bureau of Economic Research.

GUVENEN, F. AND M. RENDALL (2015): “Women’s emancipation through education: A macroeconomic analysis,” *Review of Economic Dynamics*, 18, 931–956.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2008): “Insurance and opportunities: A welfare analysis of labor market risk,” *Journal of Monetary Economics*, 55, 501–525.

——— (2010): “The macroeconomic implications of rising wage inequality in the United States,” *Journal of Political Economy*, 118, 681–722.

——— (2014): “Consumption and Labor Supply with Partial Insurance: An Analytical Framework,” *American Economic Review*, 104, 2075–2126.

——— (2017): “Optimal tax progressivity: An analytical framework,” *The Quarterly Journal of Economics*, 132, 1693–1754.

HOLTER, H. A., D. KRUEGER, AND S. STEPANCHUK (2019): “How does tax progressivity and household heterogeneity affect Laffer curves?” *Quantitative Economics*.

- HOLTER, H. A., D. KRÜGER, AND S. STEPANCUK (2023): “Until the IRS Do Us Part:(Optimal) Taxation of Households,” Tech. rep., Mimeo, University of Pennsylvania.
- KAPLAN, G. AND G. L. VIOLANTE (2010): “How much consumption insurance beyond self-insurance?” *American Economic Journal: Macroeconomics*, 2, 53–87.
- KARAHAN, F. AND S. OZKAN (2013): “On the persistence of income shocks over the life cycle: Evidence, theory, and implications,” *Review of Economic Dynamics*, 16, 452–476.
- KNOWLES, J. A. (2012): “Why are married men working so much? An aggregate analysis of intra-household bargaining and labour supply,” *Review of Economic Studies*, 80, 1055–1085.
- LISE, J. AND S. SEITZ (2011): “Consumption inequality and intra-household allocations,” *The Review of Economic Studies*, 78, 328–355.
- LISE, J. AND K. YAMADA (2018): “Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use,” *The Review of Economic Studies*.
- LOW, H., C. MEGHIR, L. PISTAFERRI, AND A. VOENA (2024): “Marriage, Labor Supply and the Social Safety Net,” *Unpublished manuscript*.
- MAZZOCCO, M., C. RUIZ, AND S. YAMAGUCHI (2013): “Labor supply, wealth dynamics, and marriage decisions,” *Unpublished manuscript*.
- REYNOSO, A. (2024): “The impact of divorce laws on the equilibrium in the marriage market,” *Journal of Political Economy*, 132, 4155–4204.
- RUPERT, P. AND G. ZANELLA (2015): “Revisiting wage, earnings, and hours profiles,” *Journal of Monetary Economics*, 72, 114–130.
- SANTOS, C. AND D. WEISS (2016): ““Why not settle down already?” a quantitative analysis of the delay in marriage,” *International Economic Review*, 57, 425–452.
- SHEPHARD, A. (2019): “Marriage market dynamics, gender, and the age gap,” *Unpublished manuscript*.
- SHIMER, R. AND L. SMITH (2000): “Assortative matching and search,” *Econometrica*, 68, 343–369.
- VOENA, A. (2015): “Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?” *The American Economic Review*, 105, 2295–2332.
- WU, C. AND D. KRUEGER (2021): “Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation,” *American Economic Journal: Macroeconomics*, 13, 79–113.

Supplemental Appendix

Marriage, Intra-Household Inequality, and Wage Risk

Piotr Denderski, Tim Obermeier

September 24, 2025

S.A Appendix to Section 2

S.A.1 Remaining value functions for couples

Firstly, observe that Equation (5) has an equivalent representation using only couple-level value functions:

$$\begin{aligned} \mathcal{U}_t^W(\lambda, \omega_m, \theta_m, \omega_f, \theta_f, A) = & \max_{c_f, c_m, h_f, h_m} \lambda (u(c_f, h_f) + \theta_f) + (1 - \lambda) (u(c_m, h_m) + \theta_m) \\ & + \beta \mathbb{E}_{\omega'_f, \omega'_m} \mathcal{V}_{t+1}^W(\lambda', \omega'_m, \theta_m, \omega'_f, \theta_f, A') \end{aligned} \quad (\text{S.1})$$

with $\mathcal{V}_{t+1}^W(\lambda, \omega'_m, \theta_m, \omega'_f, \theta_f, A')$ being the weighted sum of $V_{g,t+1}^{C,W}(\lambda, \omega'_g, \theta_g, \omega'_{g-}, \theta_{g-}, A')$ with weight λ on the value of the wife. Thus, we can similarly define the choice for a couple of age $t = t_R$ as:

$$\begin{aligned} \mathcal{U}_t^W(\lambda, \omega_m, \theta_m, \omega_f, \theta_f, A) = & \max_{c_f, c_m, h_f, h_m} \lambda (u(c_f, h_f) + \theta_f) + (1 - \lambda) (u(c_m, h_m) + \theta_m) \\ & + \beta \mathcal{V}_{t+1}^R(\lambda, \theta_m, \theta_f, A') \end{aligned} \quad (\text{S.2})$$

subject to the budget constraint, Equation (6), the condition capturing economies of scale, Equation (3), and the no-borrowing constraint $A' \geq 0$. The beginning-of-period values in retirement, $\mathcal{V}_{t+1}^R(\lambda, \theta_m, \theta_f, A)$ are defined analogously to Equation (7). Next, when $t_R < t < T$, the couple solves:

$$\begin{aligned} \mathcal{U}_t^R(\lambda, \theta_m, \theta_f, A) = & \max_{c_f, c_m} \lambda (u(c_f, h_f) + \theta_f) + (1 - \lambda) (u(c_m, h_m) + \theta_m) \\ & + \beta \mathcal{V}_{t+1}^R(\lambda, \theta_m, \theta_f, A') \end{aligned} \quad (\text{S.3})$$

subject to no-borrowing constraint $A' \geq 0$, the economies of scale condition, Equation (3), and the budget constraint:

$$c_f + c_m + A' = (1 + r)A + b_m + b_f. \quad (\text{S.4})$$

Finally, at age $t = T$ the couple solves:

$$\mathcal{U}_T^R(\lambda, \theta_m, \theta_f, A) = \max_{c_f, c_m} \lambda (u(c_f, h_f) + \theta_f) + (1 - \lambda) (u(c_m, h_m) + \theta_m). \quad (\text{S.5})$$

subject to $c_f + c_m = (1 + r)A + b_m + b_f$.


S.B Appendix to Section 3

S.B.1 Calibration and Model Fit

We present the age-gender profiles of the deterministic component of productivity $\kappa_{t,g}$ on the left panel and the mean wage profiles (accounting for wage shocks) on the right panel of Figure

S.1.

FIGURE S.1: Productivity



Notes: The figure shows life-cycle trends in wages. The first panel shows the deterministic life-cycle component. For men, this is taken from [Rupert and Zanella \(2015\)](#). For women, it is uniformly scaled downwards according to the calibrated gender gap parameter. The second panel shows average wages taking productivity shocks into account.

S.B.2 Marriage and Inequality

We expand the regression on marriage determinants presented in Table A.2 in the Appendix, and also control for the characteristics of the randomly drawn partner. This exercise aims at dissecting the uncertainty about marriage into the frictions-related component (who agents meet in the marriage market, capturing selection into pool of singles), and the love shock realization. We present the results of this extension in Table S.1. Intuitively, when we only control for wages, partner's wage has a stronger (positive) effect on the likelihood of marriage than own wage. That is the case regardless if we do not (column (1)), or do (column (3)) control for age. The wealthier the agents are, the easier they find it to marry. Once we control for age, a wealthier prospective spouse is also likelier to be accepted (age strongly confounds with assets because of accumulation for retirement). Once we control for wages and assets, the picture becomes a bit more complex - the coefficient on the wage of the potential partner turns negative, which can be due to risk sharing incentives and/or bargaining. The model including the randomly drawn partner's characteristics does a much better job predicting marital status dynamics than the one looking only at individual characteristics presented in Table S.1.

TABLE S.1: Determinants of marriage decisions

	Marriage					
	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	-1.725*** (0.008)	-1.363*** (0.015)	-1.702*** (0.086)	5.227*** (0.137)	-1.360*** (0.015)	6.956*** (0.152)
Wage	0.022* (0.010)		0.023* (0.011)		-0.674*** (0.014)	-1.680*** (0.021)
Wage (potential partner)	0.145*** (0.029)		0.159*** (0.030)		0.083*** (0.025)	-0.867*** (0.032)
Assets		1.227*** (0.012)		2.933*** (0.027)	1.352*** (0.013)	3.785*** (0.033)
Assets (squared)		-0.104*** (0.002)		-0.229*** (0.003)	-0.103*** (0.002)	-0.261*** (0.003)
Assets (potential partner)		-0.861*** (0.017)		0.815*** (0.022)	-0.960*** (0.016)	1.247*** (0.025)
Assets (potential partner, squared)		0.084*** (0.002)		-0.020*** (0.002)	0.095*** (0.002)	-0.031*** (0.002)
age			0.010* (0.005)	-0.262*** (0.007)		-0.325*** (0.008)
age ²			-0.000*** (0.000)	-0.000* (0.000)		-0.001*** (0.000)
<i>N</i>	124,843	124,843	124,843	124,843	124,843	124,843
Pseudo <i>R</i> ²	0.000	0.191	0.003	0.335	0.215	0.415
Accuracy	0.000	0.119	0.000	0.312	0.169	0.480
Precision	–	0.361	–	0.650	0.479	0.776

Notes: This table shows regressions with simulated data from the benchmark model. The dependent variable is a dummy for getting married and the sample consists of singles. Logit estimation.

S.B.3 Effects of Idiosyncratic Shocks

In this section, we provide additional details on the transmission of permanent wage shocks over the life cycle complemented by transmission of transitory wage shocks and love shocks.

S.B.3.1 Love shocks

Next to the wage shocks, fluctuations in love shock value are another source of risk in the model. Thus, we estimate transmission regression analogous to that for the wage shocks but we control for love shocks instead. The results of this exercise are presented in Table S.2.

We find that love shocks have economically and statistically significant effects on individual behavior, with notable asymmetries across gender and shock direction. Positive love shocks - whether experienced directly or by the partner - lead to higher individual consumption and, in the case of men, reduced labor supply. In contrast, negative love shocks elicit more heterogeneous responses. When individuals experience a negative shock themselves, women tend to increase

consumption and reduce hours, while men reduce consumption but slightly increase labor supply. However, when the negative shock concerns the partner, both spouses respond by cutting consumption and increasing work effort.

TABLE S.2: Effect of Love Shocks on Consumption and Hours Changes

	Δ Consumption		Δ Hours	
	Female	Male	Female	Male
Positive love shock	0.002* (0.001)	0.019*** (0.001)	-0.001 (0.001)	-0.011*** (0.001)
Positive love shock (partner)	0.022*** (0.001)	0.004*** (0.001)	-0.011*** (0.001)	-0.002*** (0.001)
Negative love shock	0.014*** (0.001)	-0.011*** (0.001)	-0.013*** (0.001)	0.005*** (0.001)
Negative love shock (partner)	-0.024*** (0.001)	-0.009*** (0.001)	0.014*** (0.001)	0.009*** (0.001)
N	103,556	104,884	103,556	104,884
R^2	0.014	0.008	0.005	0.004

Notes: Model-simulated data on couples. Estimates from OLS regressions with individual-level change in consumption or hours worked as the dependent variable. Intercepts suppressed for compactness. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Then, we investigate the role of love shocks for divorce, regressing the divorce dummy on the sample of couples from model-simulated data. The results of this exercise are displayed in Table S.3.

TABLE S.3: Determinants of Divorce: Role of Love Shocks

	Female	Male
Positive love shock	-0.004** (0.001)	-0.007*** (0.001)
Positive love shock (partner)	-0.006*** (0.001)	-0.004** (0.001)
Negative love shock	0.223*** (0.001)	0.290*** (0.001)
Negative love shock (partner)	0.287*** (0.001)	0.213*** (0.001)
N	107,776	109,124
R^2	0.478	0.468

Notes: OLS estimation of a linear probability model where the dependent variable is a divorce indicator. Sample of married agents in model-simulated data. Columns 1 and 2 show estimates for women and men, respectively. Intercepts and wage change coefficients (own and partner's) are included in the regressions but suppressed for brevity. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The results show that negative love shocks are strong predictors of divorce for both women and men, with particularly large effects when the shock is experienced by the partner. In contrast,

positive love shocks reduce the likelihood of divorce, though the effects are smaller in magnitude. These patterns highlight the asymmetric role of emotional dynamics in relationship stability: deterioration in partner affection increases divorce risk more than improvements reduce it.

S.B.3.2 Temporary Wage Shocks

In the main body of the paper we focus on permanent wage shocks. Here, we report the transmission coefficients estimated in Equations (11) - (12) for temporary productivity shocks in Table S.4. The temporary productivity shocks are essentially insured away by couples and do not impact on C . They also do not affect individual consumption of the married and singles alike. The effect of temporary shocks on married women is quite close to that estimated by [Blundell et al. \(2016\)](#), but other transmission coefficients on hours differ between our model and their results quite significantly.

TABLE S.4: Effects of Temporary Wage Shocks on Consumption and Hours

Outcome	Woman Prod. Shock			Man Prod. Shock		
	Model	No λ adj.	BPS	Model	No λ adj.	BPS
<i>Couples</i>						
Household Consumption	0.01	0.01	-0.04	-0.01	-0.01	-0.14
<i>Married Women</i>						
Hours	0.82	0.84	0.88	-0.08	-0.10	0.17
Consumption	0.02	0.00	-	-0.01	0.00	-
<i>Married Men</i>						
Hours	-0.63	-0.65	0.11	0.19	0.21	0.58
Consumption	0.00	0.02	-	0.00	-0.02	-
<i>Singles</i>						
Hours	0.58	-	-	0.37	-	-
Consumption	0.08	-	-	0.08	-	-

Notes: Reported coefficients from regressions (11) and (12) with the dependent variable being the first-difference of logarithm of the outcome of interest. The BPS columns reproduce estimates of [Blundell et al. \(2016\)](#).

The effects of temporary wage shocks on marriage and related outcomes are presented in Table S.5. The key takeaway here is that even if we find statistically significant effects, their magnitudes tend to be negligible.

S.B.3.3 Comparison to [Wu and Krueger \(2021\)](#)

In this section, we compare our model implications for the transmission of wage shocks with those reported by [Wu and Krueger \(2021\)](#). We compare our and their model's transmission coefficients for permanent and transitory shocks for couples, presenting the results of this exercise in Tables S.6 and S.7. We find that when it comes to transmission of permanent shocks, the predictions of our and their models are very similar. The models differ when it comes to predictions on

TABLE S.5: Temporary Wage Shocks and Marriage-Related Outcomes

	Marriage	Partner wage	Initial Pareto weight	Divorce	Δ Pareto weight
	(1)	(2)	(3)	(4)	(5)
Women					
Single	-0.006 [-0.002]	-0.020	0.144*	–	–
Married	–	–	–	-0.011* [0.003]	-0.010***
Men					
Single	-0.018 [0.001]	-0.003	0.029	–	–
Married	–	–	–	0.008* [0.003]	-0.006***

Notes: Same convention as in Table .

the effects of transitory shocks on men.

TABLE S.6: Effects of Permanent Wage Shocks on Consumption and Hours

Outcome	Woman Shock			Man Shock		
	Model	No λ adj.	WK	Model	No λ adj.	WK
Couples						
Household Consumption	0.21	0.21	0.18	0.34	0.34	0.35
Married Women						
Hours	0.45	0.48	0.46	-0.52	-0.57	-0.51
Consumption	0.24	0.22	–	0.31	0.33	–
Married Men						
Hours	-0.20	-0.23	-0.19	0.13	0.17	0.16
Consumption	0.19	0.22	–	0.37	0.33	–

Notes: Reported coefficients from regressions (11) and (12) on permanent wage shocks with the dependent variable being the first-difference of logarithm of the outcome of interest. The WK columns reproduce estimates of Wu and Krueger (2021).

S.B.3.4 Age profiles of insurance against permanent wage shocks.

In the main body of the paper, we reported transmission coefficients of permanent wage shocks into household consumption, individual consumption, and hours worked. Below, we present age profiles of consumption insurance coefficients, that is, if a transmission coefficient is κ , then an insurance coefficient is $1 - \kappa$.

We start this exercise by looking at couples and present the results for them on Figure S.2. We find that insurance improves with age, as on average the older agents had more time to accumulate savings for precautionary reasons. Furthermore, the married are better insured against women's permanent shocks. This is logical, as the women is, on average, the less productive member of the household and works fewer hours. Hence, her share in overall income is lower than her


TABLE S.7: Effects of Temporary Productivity Shocks on Consumption and Hours

Outcome	Woman Prod. Shock			Man Prod. Shock		
	Model	No λ adj.	WK	Model	No λ adj.	WK
<i>Couples</i>						
Household Consumption	0.01	0.01	0.01	−0.01	−0.01	0.01
<i>Married Women</i>						
Hours	0.82	0.84	0.76	−0.08	−0.10	−0.12
Consumption	0.02	0.00	−	−0.01	0.00	−
<i>Married Men</i>						
Hours	−0.63	−0.65	−0.05	0.19	0.21	0.44
Consumption	0.00	0.02	−	0.00	−0.02	−

Notes: Reported coefficients from regressions (11) and (12) with the dependent variable being the first-difference of logarithm of the outcome of interest. The WK columns reproduce estimates of [Wu and Krueger \(2021\)](#).

husband's. Thus, shocks to women's wage matter less. Next, agents are better insured against cross shocks: the insurance coefficient for women's consumption in case of men's shock is above that for the men's consumption etc. This pattern reflects the limited commitment dynamics and bargaining.

FIGURE S.2: Consumption insurance against permanent wage shocks for couples, by age



Then, we conduct the same exercise for the singles. Women are, as per our results in the main body of the paper, slightly better insured against permanent wage shocks than men. Qualitatively, the life cycle pattern is similar to that for the couples - the older the agents are, the better they can insure against the wage risk. The gap in insurance coefficients between the couples and the singles shrinks with age. This is due to the fact that many older singles have been married and divorced, and through marriage they have improved their asset position.

Next, we study consumption insurance by age conditional on marital status dynamics and present the results on Figure S.4. For each cohort of age k , we distinguish four groups of agents:

FIGURE S.3: Consumption insurance against permanent wage shocks for singles, by age





FIGURE S.4: Transmission of wage shocks by marital transition



Notes: This figure shows the transmission coefficient of wage shocks into consumption depending on different marital transitions. The blue line (stable couples) focuses on individuals who are married in the current and the previous period (i.e. the BPS assumptions). The orange line is individuals who are single in both periods. The green line is individuals who were single in the previous period, but are getting married in the same period as the shock hits. Finally, the violet line is individuals who divorce in the same period as the shock hits and were married in the previous one.

those that are married in this and the next period (*stable couples*), those that are married now, but divorce next period, those that are single, but marry next period, and those that are and remain single in the next period. The shock hits at age $k + 1$. We find that the marital status dynamics introduce significant variation in consumption insurance across the life cycle, albeit these differences decrease with age. The timing of the shock and the marital status change matters: agents who are married before the shock hits, and divorce in the period of the shock, on average fare better than agents who marry when the shock hits.

S.B.3.5 Wage shocks, Pareto weights, and consumption shares

Below, we provide information on how permanent wage shocks affect the consumption split.

TABLE S.8: Impact of one standard deviation permanent wage shock on consumption shares through initial Pareto weight.

λ	Baseline s_f (%)	Δs_f (pp), +1 std. dev. shock	
		Female	Male
0.25	37.4	+2.04	-3.59
0.30	40.2	+2.26	-3.00
0.35	42.8	+2.49	-2.58
0.40	45.3	+2.75	-2.26
0.45	47.6	+3.05	-2.00
0.50	50.0	+3.40	-1.79
0.55	52.4	+3.82	-1.62
0.60	54.7	+4.35	-1.47
0.65	57.2	+5.03	-1.33
0.70	59.8	+5.96	-1.21
0.75	62.6	+7.33	-1.10

Notes: $s_f = c_f / (c_f + c_m)$. Positive Δs_f increases women's share. Results use Equation (8) at calibrated parameters and the estimated coefficients from Table 5, column 3. The value of λ corresponds to what the couple would agree upon in the absence of shocks.

TABLE S.9: Impact of one standard deviation permanent wage shock on consumption shares through changes in Pareto weight of continuing marriages.

λ	Baseline s_f (%)	Δs_f (pp), +1 std. dev. shock	
		Female	Male
0.25	37.4	+0.16	-0.32
0.30	40.2	+0.17	-0.27
0.35	42.8	+0.19	-0.24
0.40	45.3	+0.21	-0.21
0.45	47.6	+0.23	-0.19
0.50	50.0	+0.25	-0.17
0.55	52.4	+0.28	-0.15
0.60	54.7	+0.31	-0.14
0.65	57.2	+0.36	-0.13
0.70	59.8	+0.41	-0.12
0.75	62.6	+0.48	-0.10

Notes: $s_f = c_f / (c_f + c_m)$. Positive Δs_f increases women's share. Results use Equation (8) at calibrated parameters and the estimated coefficients from Table 5, column 5. The value of λ corresponds to the Pareto weight the couple agreed on prior to shock hitting.

S.B.4 Renegotiation and Divorce

Then, we turn our attention to the patterns of binding participation constraints following wage and love shocks. We tabulate the compositions of patterns of shocks conditional on at least one participation constraint binding, and conditional on observing a divorce.

Beginning with Table S.10, we find that conditional on at least one participation constraint binding, female's participation constraint binds minimally more often.

TABLE S.10: Probability of At Least One PC Binding

	Male PC does not bind	Male PC binds
Female PC does not bind	–	0.33
Female PC binds	0.35	0.32

Then, we report on the patterns of wage shocks - either positive, negative, or no shock at all - and love shocks -whether negative shock arrived or not - in Tables S.11 and S.12. While both positive and negative wage shocks can trigger a divorce, since they can increase the wage differential between spouses, positive shocks relax participation constraints.

TABLE S.11: Wage Shock if At Least One PC Binds

		Wage shock (men)		
		Positive	No shock	Negative
Wage shock (women)	Positive	0.0086	0.11	0.036
	No shock	0.054	0.60	0.11
	Negative	0.015	0.058	0.014

TABLE S.12: Love Shock if At Least One PC Binds

		Negative love shock (men)	
		No	Yes
Negative love shock (women)	No	0.66	0.17
	Yes	0.16	0.0071

Although shocks are drawn independently by construction, the distribution of shock realizations in couples where at least one participation constraint binds reveals non-random patterns. Negative love shocks tend to occur asymmetrically within couples with binding participation constraints: it is much more common for only one partner to experience a shock than for both to be affected simultaneously. Similarly, for wage shocks, joint extreme realizations—particularly both spouses receiving negative shocks—are rare among couples who have at least one participation constraint binding. Instead, participation constraints are more likely to bind in states where one partner faces adverse conditions while the other does not. This shows that cross-spousal asymmetries in shocks are a key driver of participation constraints binding in equilibrium.

Then, we further narrow down the sample to divorcing couples. Firstly, we find that the pattern of binding participation constraints is identical to that for the broader population of couples with binding participation constraints, see Table S.13.

Then, we zoom in on the wage shock and love shock patterns for the divorcing couples in

TABLE S.13: Binding PCs conditional on Divorce

	Male PC does not bind	Male PC binds
Female PC does not bind	0.0	0.33
Female PC binds	0.35	0.32

Tables S.14 and S.15. Logically, we find a shift toward more extreme realizations of both wage and love shocks. Among divorcing couples, negative wage shocks for both spouses become slightly more common, and mixed-sign wage realizations—particularly where one partner experiences a negative shock while the other does not—remain prevalent. However, the shift is even more pronounced in the case of love shocks: divorcing couples are far more likely to experience asymmetric negative love shocks, especially configurations where only the man receives a negative love shock (47%) or only the woman does (38%), compared to just 17% and 16%, respectively, in the general PC-binding sample. Joint negative love shocks, while still rare, are also somewhat more frequent among divorcing couples. These patterns suggest that asymmetries in love shocks are a stronger trigger for divorce than wage shocks alone, even though both contribute to participation constraint violations.

TABLE S.14: Wage Shock if Divorce

		Wage shock (men)		
		Positive	No shock	Negative
Wage shock (women)	Positive	0.012	0.082	0.031
	No shock	0.053	0.67	0.083
	Negative	0.0097	0.051	0.013

TABLE S.15: Love Shock if Divorce

		Negative love shock (men)	
		No	Yes
Negative love shock (women)	No	0.13	0.47
	Yes	0.38	0.02

S.C Appendix to Section 4

S.C.1 Consumption Insurance in the Benchmark and Auxiliary Models


Next, we compare the consumption insurance between our benchmark model and the unitary household model for couples and no marriage model for singles. Figure S.1 shows the age profiles of consumption insurance against permanent wage shocks. These results help us better

understand how, *ceteris paribus*, the presence of marriage and divorce affects consumption insurance.

First, focus on panels (a) and (b) which compare the household-level transmission of shocks to men's and women's wages between the benchmark model and the unitary model, focusing on stable couples. While the outcome (household expenditure) is the same, there are several reasons why the transmission of shocks differs between the models: allowing for marriage and divorce, relative to having a unitary model, changes asset accumulation and the expected future values of couples, as individuals anticipate the possibility of divorce. As a result, the response to a shock may differ because of these two factors. Panels (a) and (b) show that that for couples, these effects mainly matter for the response to female shocks, where allowing for marriage and divorce lowers insurance in the beginning of the life-cycle (from around 0.8 to roughly 0.75).

Panels (c) and (d) show similar comparisons for singles, comparing singles in the benchmark model to singles in the no-marriage model. Again, the transmission of shocks differs because of differences in asset accumulation and differences in expected future values, as singles in the benchmark model anticipate the possibility of future marriage. Interestingly, insurance is *lower* in the benchmark model, i.e. despite the possibility of future marriage, shocks to wages are transmitted more strongly into current consumption.

FIGURE S.1: Consumption Insurance Against Permanent Wage Shocks: Comparison Between Models



Notes: This figure compares the transmission coefficients of wage shocks into household consumption between the benchmark model and the unitary model (for couples) and the single-earner model (for singles). Parameters are kept constant across models, except for setting Pareto weights to 0.5 and forcing individuals to always be married in the unitary model, and shutting down marriage in the single-earner model.