

Online appendix

Singles, Couples, and Their Labor Supply: Long-Run Trends and Short-Run Fluctuations

Jonna Olsson*

March 23, 2024

A. Appendix material	2
A.1. Changes on the intensive vs. extensive margin over time	2
A.2. Volatility calculations	5
A.3. More about sectors	7
A.4. International perspective on employment recoveries and female labor force participation	8
A.5. Equilibrium definition	10
A.6. Unemployment by subgroup	12
A.7. Couple insurance	13
A.8. Asset holdings couple and single households	16
A.9. Robustness: Equal discount factors for singles and couples	17
A.10. Robustness: A model with children	18
A.11. Comparing steady-state models	24
A.12. Asset holdings for single vs married women, model vs data	27
A.13. Aggregate shocks during a transition path in an RA model	30
A.14. Robustness: Extended model	33
A.15. Robustness: A model with progressive taxation	39

*jonna.olsson@nhh.no, NHH Norwegian School of Economics, Helleveien 30, 5045 Bergen, Norway.




Figure A.1: Intensive margin labor choice over time, by subgroup. Population aged 25-64.
Source: CPS.

A. Appendix material

A.1. Changes on the intensive vs. extensive margin over time

Figure A.1 shows the intensive margin by subgroup and its evolution over time. Panel (a) shows average weekly hours per worker (using AHRSWORKT, the total number of hours worked over all jobs in the last week). Panel (b) shows the annual hours worked per worker, where “worker” is defined as having a positive annual hours worked. Annual hours worked is defined as hours worked last week times the number of weeks worked last year.

There is a more pronounced upward trend in the intensive margin for married women when looking at the annual hours per worker. This implies that at least some married women increased the number of weeks worked during a year.

Figure A.2 shows annual hours per capita for the four subgroups. To be more precise in the contribution of intensive vs. extensive margin to the long-run changes in annual hours per capita I do the following decomposition:

$$\log \left(\frac{H_t^i}{P_t^i} \right) = \underbrace{\log \left(\frac{E_t^i}{P_t^i} \right)}_{\text{extensive margin}} + \underbrace{\log \left(\frac{H_t^i}{E_t^i} \right)}_{\text{intensive margin}}$$

with H_t^i denoting total hours worked in year t by subgroup i , P_t^i total population of subgroup i , and E_t^i the employment level (defined as having positive hours worked for the year).

The intensive vs. extensive margin contribution to the growth in hours per capita

since 1964,

$$\log \left(\frac{H_t^i}{P_t^i} / \frac{H_{64}^i}{P_{64}^i} \right) = \underbrace{\log \left(\frac{E_t^i}{P_t^i} / \frac{E_{64}^i}{P_{64}^i} \right)}_{\text{extensive margin}} + \underbrace{\log \left(\frac{H_t^i}{E_t^i} / \frac{H_{64}^i}{E_{64}^i} \right)}_{\text{intensive margin}}$$

is shown in Figure A.3. Year 1964 is chosen as the baseline year since it is the first year with reliable data on weeks worked per year. Panel (d) shows the cumulative growth of hours per capita for married women and its decomposition in extensive margin and intensive margin, and as the figure shows, the extensive margin changes have been much more important than the intensive margin changes (panels (a), (b), and (c) show the other demographic subgroups for completeness).

A few minor comments about the choice of variables:

- (a) There is an alternative choice for weekly hours, “usual hours worked per week”. I prefer the actual hours worked per week, since that should better capture the work being done (since I am aggregating into reasonably large groups and want to capture potential differences over time in for instance sick leave and vacations), and because usual hours worked is only available from 1976 and onward.
- (b) For the number of weeks worked, I use the variable WKSWORK2, which is number of weeks worked last year, intervalled. Respondents were prompted to include weeks in which they worked for even a few hours and to include paid vacation and sick leave as work. 60 percent of the answers with positive number of weeks are in the highest bin, 50-52 weeks, which i assign as full year. For the remaining observations, I assign the mid point of the interval. The alternative variable, weeks worked last year (not intervalled), is only available from 1976 and onward.




Figure A.2: Annual hours per capita by subgroup. Population aged 25-64. Source: CPS.




Figure A.3: Decomposition of cumulative growth in annual hours per capita since 1964. Population aged 25-64. Source: CPS.

A.2. Volatility calculations

I define the volatility related to aggregate fluctuations in the following way. I compute the cyclical component of the labor series for each demographic subgroup as the residual after applying an HP filter (using a smoothing parameter of 6.25). Then I regress this cyclical component on the cyclical component of GDP (which is also the residual after applying an HP filter with a smoothing parameter of 6.25).¹ Hence, the equation I estimate is:

$$\hat{x}_{c,t,i} = \beta \hat{Y}_{c,t} + \varepsilon_t \quad (\text{A.1})$$

with $\hat{x}_{c,t,i}$ denoting the cyclical component from HP filtering of yearly series of $\log(x_i)$ (with $i \in \{\text{married men, married women, single men, single women}\}$) and $\hat{Y}_{c,t}$ the cyclical component from HP filtering of yearly series of $\log$ GDP. x_i is weekly hours per capita for each demographic subgroup. The coefficient of interest, β , captures how responsive aggregate hours for this particular subgroup are to fluctuations in GDP. After estimating this equation, I create a series of predicted hours by subgroup, and then finally measure the percentage standard deviation of these series. Hence, one can think of this measure of volatility as capturing the component of hours volatility that is related to aggregate economic fluctuations.²

A.2.1. Volatility by subgroup, only prime-aged individuals

Figure 2(a) in the main text shows the volatility for the population aged 25-64. In contrast, Figure A.4 shows the same graph, but restricting the sample to only prime-aged individuals aged 35 to 54. The same pattern emerges as for the full sample: hours worked by men are more volatile than hours worked by women (comparing single men to single women, and married men to married women), and hours worked by singles are more volatile than hours worked by married individuals (comparing single men to married men, and single women to married women). Hence, the fact that singles are more likely to be young is not what is driving the difference in volatility between singles and married individuals.

¹This method follows Jaimovich and Siu (2009). The results are robust to including additional RHS variables such as lagged GDP. Real GDP per capita is constructed by U.S. Bureau of Economic Analysis (2024a) and U.S. Bureau of Economic Analysis (2024b).

²As a robustness check, I rerun the analysis using different filtering techniques (Baxter-King, Christiano-Fitzgerald, and Butterworth filters) and the conclusions remain unchanged.




Figure A.4: Volatility in hours worked by demographic subgroup for the time period 1962-2017. Population aged 35-54. See text in section 2.2 for the definition of total volatility and volatility related to aggregate fluctuations. Source: CPS, FRED.

A.3. More about sectors

Figure A.5 shows the fraction of individuals in the labor force working in the manufacturing sector, while Figure A.6 shows the responsiveness of employment to fluctuations in GDP by sector and subgroup.




Figure A.5: The fraction working in manufacturing (age 25-64). Source: CPS. The graph includes everyone with a defined sector, i.e., also unemployed with sector definition.




Figure A.6: Responsiveness by sector and subgroup (age 25-64). Result from regressing the cyclical component of hours by type in sector on the GDP cyclical component. Source: CPS.




Figure A.7: Ratio of female to male labor force participation, normalized to the 2016 level (age 25-54). Source: OECD (2018).

A.4. International perspective on employment recoveries and female labor force participation

The phenomenon of slow employment recoveries in more recent years is not as prevalent outside the U.S. Graetz and Michaels (2017) look at evidence from a broader set of countries, and compile data from 17 countries for the years 1970 to 2011. Their conclusion is that recent recoveries have generally not involved any significantly slower recovery of employment. In this respect, the slower employment recoveries in the U.S. seem to be the exception.

Figure A.7 shows the ratio of female to male labor force participation for the G7 countries, normalized to the 2016 level. As can be seen, the U.S. was the first country in this group to reach its current level, while the other countries have been on an upward trajectory up until now.

Here I focus on the trend and therefore normalize male/female labor force participation, but both absolute and relative labor force participation for men and women respectively differ substantially across countries. One reason is differences in taxation, see Bick and Fuchs-Schundeln (2017) and Bick et al. (2019) for estimates of how much differences in non-linear income and consumption taxes contribute to the level differences in married women's work across countries.

Figure A.8 shows the labor force participation by gender for the population aged 25-54 for all OECD countries for which data is available for at least some time period. As can be seen, there is a common trend of falling labor force participation among men, and increasing labor force participation among women. Figure A.9 shows the ratio of female to male labor force participation, corresponding to Figure A.7 but for all OECD countries.




Figure A.8: Labor force participation by gender, population aged 25-54. The thick black line indicates the US. Source: OECD.




Figure A.9: Ratio of the female to male labor force participation rate for all countries for which data is available from OECD. Figures are normalized to the 2016 ratio. Population aged 25-54. The thick black line indicates the US. Source: OECD.

A.5. Equilibrium definition

Denote a household's state vector by (a, b) where assets $a \in \mathcal{A}$ and type $b \in \mathcal{B}$, with $\mathcal{A}$ and $\mathcal{B}$ given by:

$$\mathcal{A} = [\underline{a}, \bar{a}]$$

$$\mathcal{B} = \left\{ \left(c, (\omega_m, \omega_f), (u_m, u_f) \right), \left(m, \omega_m, u_m \right), \left(f, \omega_f, u_f \right) \right\}$$

using c to denote couple households, m single males, and f single females. Further, define $\mathcal{C}$, $\mathcal{M}$, and $\mathcal{F}$ as:

$$\mathcal{C} = \{ \mathcal{A} \times (x, \cdot) \mid (x, \cdot) \in \mathcal{B}, x = c \}$$

$$\mathcal{M} = \{ \mathcal{A} \times (x, \cdot) \mid (x, \cdot) \in \mathcal{B}, x = m \}$$

$$\mathcal{F} = \{ \mathcal{A} \times (x, \cdot) \mid (x, \cdot) \in \mathcal{B}, x = f \}$$

and

$$\mu_c = \int_{\mathcal{C}} d\Gamma, \quad \mu_m = \int_{\mathcal{M}} d\Gamma, \quad \mu_f = \int_{\mathcal{F}} d\Gamma.$$

Let $\hat{\omega}_i$ define the realized market productivity: $\hat{\omega}_i = (1 - \tau_i)\omega_i \quad \forall i \in \{m, f\}$.

A recursive competitive equilibrium is given by a set of prices $\{r, w\}$, decision rules $C(a, b)$, $E_m(a, b)$, $E_f(a, b)$ and $A(a, b)$, a tax level T , and a stationary distribution Γ such that:

1. The decision rules solve the households' problem for all (a, b) .
2. Firms optimize, i.e., factor prices are given by:

$$r = F_1(K, L) - \delta \quad \text{and} \quad w = F_2(K, L)$$

3. The government budget balances:

$$\left(\frac{\pi_w}{\pi_r + \pi_w} \right) T (\mu_c 2 + \mu_m + \mu_f) =$$

$$w\varphi \left(\int_{\mathcal{C}_{um}} \hat{\omega}_m d\Gamma + \int_{\mathcal{C}_{uf}} \hat{\omega}_f d\Gamma + \int_{\mathcal{M}_u} \hat{\omega}_m d\Gamma + \int_{\mathcal{F}_u} \hat{\omega}_f d\Gamma \right)$$

$$+ \left(\frac{\pi_r}{\pi_r + \pi_w} \right) \kappa \left(\mu_c (2 - \tau_f^r - \tau_m^r) + \mu_m (1 - \tau_m^r) + \mu_f (1 - \tau_f^r) \right)$$

where $\mathcal{C}_{um} = \{ \mathcal{A} \times (x, \cdot, (u, \cdot)) \mid (x, \cdot, (u, \cdot)) \in \mathcal{B}, x = c \text{ and } u = 1 \}$ etc.

4. Capital and labor markets clear:

$$K' = \int_{\mathcal{A} \times \mathcal{B}} A(a, b) \, d\Gamma$$

$$L = \int_{\mathcal{C}} (\hat{\omega}_m E_m(a, b) + \hat{\omega}_f E_f(a, b)) \, d\Gamma + \int_{\mathcal{M}} \hat{\omega}_m E_m(a, b) \, d\Gamma + \int_{\mathcal{F}} \hat{\omega}_f E_f(a, b) \, d\Gamma$$

5. For all relevant Borel sets B

$$\Gamma(B, b) = \sum_{\bar{b}} \pi(b|\bar{b}) \int_{a: A(a, \bar{b}) \in B} \Gamma(da, \bar{b})$$

Note that when I solve for the deterministic transition path, I do not use the stationary equilibrium but a sequential equilibrium definition in which the equilibrium objects depend on time. To facilitate the understanding, I nevertheless spell out the stationary version here.

A.6. Unemployment by subgroup

Figure A.10(a) shows unemployment figures by subgroup for individuals aged 25-64. Figure A.10(b), which shows the unemployment rates restricting the sample to individuals aged 35-64, confirms that the pattern is extremely similar even though the youngest in the workforce are excluded. Hence, the high unemployment rate among single men is not purely driven by them being younger.

Table A.1 shows the average unemployment rate over the period by subgroup.




Figure A.10: Unemployment by subgroup. Source: CPS.

	Married men	Married women	Single men	Single women
Average unemployment	3.9%	4.2%	8.9%	6.2%

Table A.1: Average unemployment rate 1962-2018, individuals aged 25-64. Source: CPS.

A.7. Couple insurance

Of the three insurance mechanisms available to couples – the productivity risk, the unemployment risk, and the labor supply mechanism – the lower joint productivity risk is most important. If we assume that the husband's and the wife's productivity processes were uncorrelated within the couple, the couple's need for precautionary savings would decrease substantially. The reason is simply that the variance of their joint productivity would decrease by half compared to the perfectly correlated case.

For a simple intuition, assume that the individual perceives a productivity process with the following first two moments: $E[\omega_i] = \mu$ and $\text{Var}(\omega_i) = \sigma^2$. Then a couple with perfectly correlated productivity processes would have the following two first moments: $E[\omega_i + \omega_j] = 2\mu$ and $\text{Var}(\omega_i + \omega_j) = 4\sigma^2$. However, if the productivity processes are uncorrelated, the first moment is unchanged, while the second moment is $\text{Var}(\omega_i + \omega_j) = 2\sigma^2$.

The opportunity to supply labor in an uncorrelated way is approximately one third as important in terms of how much less the couple wants to save, while the uncorrelated risk of unemployment affects the precautionary savings very little (given the calibration in this model, assuming a reasonably high replacement rate in case of unemployment).³

To more formally evaluate the importance of the different insurance mechanisms for couple households I compare the *average assets held by couple households* to average assets held by single households in different model settings. The model used for this is the simplest possible, in which a couple household consists of two individuals, both with the same productivity process as a single (for simplicity, I have used the male productivity process). Thus, if the couple household had perfectly correlated productivity shocks, unemployment shocks, and labor choices, they would behave as a twice as large single household (I have also set the consumption scale parameter for couple households $\zeta_1 = 2$).

This is shown in the right-most bar in Figure A.11. If couple households have perfectly correlated productivity shocks, unemployment shocks, and are forced to both work or both enjoy leisure in a given period, their average savings are twice as large as the savings of the average single household in the economy.

The first bar in the graph shows the average asset holdings by couple households if they have access to all three insurance mechanisms. The experiment is done as a partial-equilibrium exercise, thus I hold the interest rate constant across the

³One way to interpret the effect of the opportunity to supply labor in an uncorrelated way within the couple in its cleanest form is to think about it as if the couple is one step closer to an intensive margin labor supply decision. See Pijoan-Mas (2006) for a detailed discussion how work effort on the intensive margin is used as a consumption smoothing mechanism in a standard heterogeneous-agent model.




Figure A.11: Average asset holdings of couple households for different levels of intra-household insurance. See text for description.

different scenarios. If couples face uncorrelated productivity and unemployment shocks, and can supply labor in an uncorrelated fashion, their savings level is substantially lower. Since they can self-insure much more efficiently, they need to hold less wealth for precautionary reasons.

Then bar two, three, and four show how the average asset holdings of couple households change if we add one insurance mechanism at a time. As can be seen, the fact that couple households have uncorrelated unemployment has a negligible effect on their desired savings level (average wealth increases by 1%). This is due to the fact that unemployment is a transitory shock, and that the replacement rate in this model is relatively high (50%).

If we remove the possibility to supply labor in a non-correlated fashion from the couple household (while keeping unemployment and productivity shocks uncorrelated), average savings double, as the second bar shows.

If we let the two individuals receive perfectly correlated productivity shocks (while holding unemployment shocks uncorrelated and let them supply labor in an uncorrelated fashion), the average asset holdings increase substantially, as bar four shows.

Thus, of the three insurance mechanisms mentioned above – the productivity risk, the unemployment risk, and the labor supply mechanism – the lower joint productivity risk is most important, given any reasonable calibration of the model.

A.7.1. Calibration of a stylized economy

In the stylized economy shown in Table 2, I give all individuals the productivity process estimated by Krueger et al. (2016) (thus not using the male and female specific estimates). The wage process is consequently characterized by the parameters (ρ, σ^2) which are set to $(0.9695, 0.0384)$. The production process is assumed to be uncorrelated between the male and the female within couples. Both couples and singles are given a discount factor (β) of 0.96 and all individuals are given an equal disutility of labor (ψ) of 1.0.

To focus on the importance of the couple insurance mechanisms, I remove the retirement phase from the model. The consumption scale factor (ζ_1) is set to 2 for couple households (and 1 for single households).

The economy with both singles and couples consist of a unit mass of each type. The production technology parameters are kept identical to the baseline model.

A.8. Asset holdings couple and single households




Figure A.12: Asset holdings by household type. Source: PSID 2008, referring to the previous year. All values measured in 2008 USD. House value is the value of first residence minus mortgages.





Figure A.13: Asset holdings over the life cycle. Source: PSID 2008, referring to the previous year. All values measured in 2008 USD. Robust standard errors.

A.9. Robustness: Equal discount factors for singles and couples

The upper half of Table A.2 shows the steady-state results from the baseline model (numbers correspond to Table 4 and Table 5 in the main paper), while the lower half shows the results from a model where I have assigned the same discount factor to couple and single households. As can be seen, the main difference is the level of savings: in the model with equal discount factors, the couple households save substantially less assets than in the baseline model and than in the data. The steady-state level of employment is similar in the two models, with the couple households working slightly more and the single households working less in the “equal-discount-factor model” than in the baseline model. However, the employment levels could have been adjusted with a re-calibration of ζ_1 , the consumption utility scale parameter for couple households and ψ , the general disutility of work (I keep all preference parameters except the discount rate constant in this comparison to facilitate the exposition). Importantly, the results for the short run responses in employment to a TFP shock remain very similar: the values for married men and married women are close to the baseline results, and substantially smaller than for single households.

	Couples	Single men	Single women
<i>Panel A: Baseline model ($\beta_s = 0.949$; $\beta_c = 0.96$)</i>			
Employment (%)	86.7, 67.7	75.5	74.6
Consumption	1.9	1.0	0.7
Assets	2.1	1.0	0.6
<i>Immediate impact on employment from a TFP shock of one std</i>			
Percent (%)	0.4, 1.4	4.2	3.4
<i>Panel B: Model with $\beta_s = \beta_c = 0.949$</i>			
Employment (%)	87.3, 68.1	70.8	69.3
Consumption	1.9	1.0	0.7
Assets	1.1	1.0	0.6
<i>Immediate impact on employment from a TFP shock of one std</i>			
Percent (%)	0.4, 1.3	4.4	4.3

Table A.2: Model results: comparing a model with different discount factors (Panel A) with a model with same discount factors (Panel B) for singles and couples.



(a) Employment-to-population ratio, women with and without kids. (b) Fraction of women having kids at home, single vs. women in couple.

Figure A.14: Employment-to-population ratio for women with and without children and fraction of women having children at home. Source: CPS, age group 25-64. Having children is defined as having one or more own children residing together with respondent, and includes biological children as well as step-children and adopted children.

A.10. Robustness: A model with children

Fertility is an important factor to consider in order to fully understand the increase in female labor force participation and the interactions with savings and consumption. Figure A.14(a) shows the employment-to-population ratio among women with and without children, and as the graph shows, the increase in employment has been faster for women with children. As the graph also shows, the working gap between women with and without children has vanished: a woman with children is today equally likely to be working as a woman without. Figure A.14(b) shows the fraction of women having children, and as the figure shows, it is more likely for couples to have children.

Both for married and single women the employment has increased faster for mothers. For married women the employment gap between women with and without children closed around 1980 (Figure A.15(a)), while for single women it did not close until around 2000 (Figure A.15(b)).

In the baseline model, children are only captured in a reduced form: the different burden of taking care of the children is likely one of the reasons for the higher disutility of labor for women than for men, and the higher probability of married households to have children is likely part of what gives the couple households a higher consumption scale weight and a higher discount factor. However, these differences are just averages and the baseline model does not distinguish between or give different predictions for married women with and without children or single women with or without children.

To investigate the importance of having children for savings behavior in the context of this model, I create a version where the event of having a child is an




Figure A.15: Employment-to-population ratio for women aged 25-64 with and without kids. Source: CPS. Having kids is defined as having one or more own children residing together with respondent, and includes both step-children and adopted children as well as biological children.

exogenous shock. To have children in practice in the model means a demand shock (the household increases the weight on consumption with more mouths to feed in the family) and a simultaneous disutility-of-work shock (the disutility of work increases for the woman in the presence of children).

A.10.1. Including children in the model

The child shock is more likely to occur for married couples than for singles, and has an impact on couple households and on single women (the fraction of single men having children in their home is actually 13%, but I approximate it with zero, thus nothing changes in terms of preferences for single men that draw a child shock).

The exogenous shock is modeled such that:

- when a household is reborn after retirement, it is reborn without children,
- a household without kids might get hit by a child shock,
- if the household gets hit by a child shock, the expected time of staying in having-children state is 25 years,
- the expected time in working-age is, just as before, 45 years,
- the expected time in retirement is, just as before, 15 years,
- over a lifetime, 15% of women will remain without children,
- it is 1.54 times more likely for a married woman to have children than for a single woman.⁴

These assumptions jointly pin down the transition matrices shown in Table A.3. Finally, the productivity shocks are assumed to be independent of the having-children shocks.

⁴Pooled average over the years 1968 and 2017. Source CPS.

	Single women			Couples		
	No children	Children	Retired	No children	Children	Retired
No children	0.960	0.024	0.016	0.941	0.054	0.005
Children	0	0.96	0.04	0	0.96	0.04
Retired	0.067	0	0.933	0.067	0	0.933
<i>Ergodic distribution</i>	0.46	0.27	0.27	0.31	0.42	0.27

Table A.3: Transition matrices for the model with exogenous shocks to family composition.

The differences between having children or not are the following in the model: (a) women with children face a higher disutility of work, and (b) households with children have a higher utility scale parameter on consumption in their utility function (ζ_1). The calibration of these additional parameters (and re-calibration of the relevant existing parameters) are pinned down according to the following logic:

- Having children increases ζ_1 (the consumption utility scale parameter) for a single women from 1 to 1.5 (following the OECD equivalence scale where an additional child counts as 0.5).
- The fact that women with children work as much as women without pins down the relative extra utility cost of working for women with children (approximately 32% higher utility cost from work).
- The fact that single women in total work approximately as much as single men pins down the general higher disutility of work for women vs. men (assuming the observed wage gap), in this new calibration it is 16% more costly in utility terms to work for females than for males (previously 8%).
- The fact that married women with children work approximately the same as married women without kids pins down the difference in ζ_1 (the consumption scale parameter) between couple households with and without kids: couples with children have a 50% higher weight on consumption.

There are a couple of things to note here. First, it is not evident that the increase in disutility of work from having children is equal for single women and married women. However, without further evidence, assuming equal disutility for those two groups is a reasonable starting point. Second, in the same vein, one could argue that the disutility of labor for married men with children should increase, and the consumption utility scale parameter should increase for the single man with kids (there should be approximately as many single men who are fathers and perhaps paying child support, even though the children are primarily living with their mothers). Again, due to lack of better evidence, I take a simplified view, where

	Couples	Single men	Single women
<i>Panel A: Baseline model ($\beta_s = 0.949; \beta_c = 0.96$)</i>			
Employment (%)	86.7, 67.7	75.5	74.6
Consumption	1.9	1.0	0.7
Assets	2.1	1.0	0.6
<i>Panel B: Model with exogenous children shocks ($\beta_s = 0.949; \beta_c = 0.96$)</i>			
Employment (%)	89.5, 63.3	77.2	72.8
Consumption	1.9	1.0	0.7
Assets	2.2	1.0	0.6

Table A.4: Model results, comparing the baseline model (Panel A) with a model with exogenous children shocks (Panel B). Note: the model with children shocks is solved with only 5 productivity states and is not calibrated to match the data moments exactly, the purpose is to show the magnitude of the differences.

these two parameters do not depend on children, as a starting point. Importantly, I keep the discount factors for singles and couples at the same levels as in the benchmark calibration.

Table A.4 illustrates the steady-state results from the “exogenous children shock” model, assuming the gender wage gap of today (20%). As can be seen, there are no substantial changes to savings behavior. The changes in consumption needs and the disutility of labor that children imply has an effect on the intertemporal choices of the households, but given that we want the model to replicate the observed actual employment levels, this necessarily means that other unobserved preference parameters have to adjust. Thus, given the observed employment and wage levels, total income for the different groups is the same in the model with exogenous child shocks, and the difference in risk exposure due to this shock is not large enough to induce any quantitatively important differences in asset accumulation. Thus, what matters for the savings behavior is primarily the discount rate.

A.10.2. Additional childrelated mechanisms

The model described in this section includes a child shock, but it is still very reduced form, and primarily serves the purpose of showing the importance of an exogenous child shock for the savings decision. There are additional mechanisms related to children that could be important to understand the evolution of female employment over time.

One such mechanisms is childcare costs. The inclusion of childcare costs could potentially increase the effect of the shrinking gender wage gap, and to some extent explain the faster increase in employment among married women with children than without children. The current model shows that the shrinking gender wage

gap can explain 2/3 of the increase in labor supply among married women. The remaining third is driven by various mechanisms not in the model, such as decrease in discrimination, changing social norms, and more evenly distributed household chores, but also the inclusion of childcare costs. A small static model of a couple household shows the point. The couple faces the following problem:

$$\max_{c, \ell_m, \ell_f} u(c, \ell_m, \ell_f) \quad (\text{A.2})$$

$$\text{s.t. } c = \ell_m w_m + \ell_f w_f - \ell_f \mathbb{1}_{kids > 0} \kappa \quad (\text{A.3})$$

$$\ell_i \in \{0, 1\} \quad \forall i \in \{m, f\} \quad (\text{A.4})$$

where κ represents the childcare cost the couples with children incur if the female is working. It is thus assumed that if the woman is not working, the childcare cost is effectively zero. The budget constraint for a couple without children can thus be rewritten as

$$c = \ell_m w_m + \ell_f w_f \quad (\text{A.5})$$

while the budget constraint for a couple with children is given by:

$$c = \ell_m w_m + \ell_f (w_f - \kappa) \quad (\text{A.6})$$

For the couple with children, the cost of childcare reduces the effective wage of the woman.

We can then compare the *change in relative wage* for married women with and without children between 1960 and 2000. The change in relative wage for married women without children is given by:

$$\frac{0.8 - 0.55}{0.55} = 45\% \quad (\text{A.7})$$

Thus, in the 1960s, the female received 55% of the wage of a male, while in 2000 this had increased to 80%. This is an increase of 45% (not percentage points).

For the couple with children the analogous calculation is given by (normalizing κ so that it is expressed as a fraction of male wages):

$$\frac{(0.8 - \kappa) - (0.55 - \kappa)}{0.55 - \kappa} = \frac{0.8 - 0.55}{0.55 - \kappa} > 45\% \quad (\text{A.8})$$

Thus, the increase in relative wage between 1960 and 2000 is higher for women with than without children if we include a cost of childcare if the woman is working. This could, within this framework, be part of the explanation why the labor force participation has risen faster for married women with children.

A related more general question is if we with the help of this (type of) model can

understand employment patterns over time for women with vs. without children better. To capture the historical employment patterns for these two groups separately is outside the scope of the current model/paper, but having children is an important part of the labor supply decision.

An extended model focused on evaluating the labor supply of women with vs. without children would give rise to a number of interesting issues. For example, the selection effect would be stronger among women with children than women without: faced with an additional childcare cost if the woman decides to work, the effective wage would be lower for mothers, and for some potential low-earners, it would not pay off to work if the net income, after childcare costs, is too low. Furthermore, one could argue that the extra disutility of labor for women with children has decreased over time, and that has been one of the factors behind the rise in female labor force participation. As shown in the Appendix section A.11, the part of the increase in labor force participation that is not explained by the shrinking gender wage gap could be explained by an on average 20% higher disutility of work for women in the 1960s than today. Assuming that the disutility of work has fallen in particular for women with children, and further allowing for differential speed of decline for married and single women, it would help explaining also the faster increase in employment among single women with children.

The aim of the current paper is to understand both long- and short-run labor supply patterns. Thus, an extended model focusing on the children vs. no-children margin should also speak to short-run fluctuations for those groups both empirically and in the model. These questions would be very interesting, and are left for future research.

A.11. Comparing steady-state models

In this paper, I focus on a shrinking productivity gap between men and women as the driver for the increase in labor supply among married women. However, an alternative explanation often mentioned is that the disutility of work has decreased for women, for instance as a result of less discrimination and harassment in the workplace, or more permitting social norms in society.

To shed some light on the effect of these two parameters – the gender productivity gap (τ_f) and the female disutility of labor (ψ_f) – I compare steady-state results from models with different calibrations of those two parameters. Table A.5 describes four models. Model 1 refers to the baseline model, calibrated so that labor force participation by subgroup and the observed female wage gap are what we observe in the data today, as described in section 4. Model 2 refers to a model where ψ_f is the same as in the baseline model, but the female wage gap is what we observed at the beginning of the 1960s.

Model 3 is an alternative specification in which the productivity gap, τ_f , is the same as in the baseline model, but ψ_f , the disutility of labor for females, is adjusted upwards by 20%. Model 4, finally, is a model in which I change both τ_f and ψ_f as compared to the baseline model.

Figure A.16 shows the resulting employment rates. First of all, in model 2, a higher productivity wedge, τ_f , makes the married women work substantially less.⁵ What might at a first glance be more surprising is the small (and even positive) effect on single women's employment rate. However, for single women, a change in the productivity wedge is similar to a change in the overall wage level, and such a change has a small impact on labor supply with balanced growth preferences. The income effect and the substitution effect cancel, and therefore the net effect on single women is negligible. Hence, the model delivers a theory of why married and single women display different responses to changes in wages in the long run.

A second observation is that the increase in married men's employment is not nearly as large as the decrease in married women's employment. There is very little "crowding-out" within the couple, in line with the data.

A third observation is the negligible effect on single men, which is not surprising. The only way in which this higher productivity wedge affects single men is via the general equilibrium channel with changes in the interest rate, the wage and the lump-sum tax.

In model 3, an increase in women's disutility of labor by 20% leads to decreasing employment figures for both married women and single women, while again single

⁵ A related observation is that the selection effect is stronger in this model: for married women, the average wage among working individuals is 16% higher than the offered wage.

	Model 1 baseline	Model 2 $\tau_f \uparrow$	Model 3 $\psi_f \uparrow$	Model 4 $\tau_f \uparrow, \psi_f \uparrow$
τ_f	20%	45%	20%	45%
ψ_f	1.513	1.513	1.816	1.816
Observed wage gap (mean)	22%	44%	19%	42%
Observed wage gap (median)	20%	45%	20%	45%

Table A.5: Comparison of four models with different calibrations of female disutility of labor (ψ_f) and productivity wedge (τ_f) and the resulting observed wage gap.




Figure A.16: Steady-state comparison. Employment rates by subgroup from four steady-state models with different calibrations as described in Table A.5.

men are, of course, unaffected. The upward adjustment of married men's labor supply to compensate for their wives' lower labor supply is negligible.

In model 4, finally, both the productivity gap and the disutility of labor are increased as compared to the benchmark. In this model, single women work slightly less due to the increased disutility, while married women work substantially less, due to both the increased disutility and the increased productivity gap. Once more, married men work slightly more to compensate for their wives' lower earnings, and single men are hardly affected. Model 4 is an example of a parametrization of the model for which the employment rates by subgroup in 1962 are hit very closely.

The comparison between these steady-state models shows that the increase in labor supply among married women is unlikely to be mainly driven by a general decrease in the disutility of labor for women. If this had been the case, single women would also have been strongly affected. A change in the gender wage gap (in line with what has been observed in the data), on the other hand, affects the labor supply of married women, but not that of single women. The effect on married

women is approximately two thirds of the increase we have seen historically. The remaining third could be modelled as an additional change in the disutility of work, which would affect both single and married women.

A.12. Asset holdings for single vs married women, model vs data

The model has predictions about the correlation between asset holdings, productivity levels, and work decisions. Figure A.17 shows model results for asset holdings against the productivity level for single and married women, respectively, *conditional on working*. The reason for selecting only working women is that this is what can be compared with the data: for non-workers, productivity is not observed. Looking at single women, it is clear that women of all productivity levels work, but there is a strong connection between wealth, productivity, and the work decision: single women with high wealth and low productivity do not work.

Figure A.17(b) shows the corresponding graph for married women from the model (unconditional on the spouse's productivity). As can be seen, married women of all productivity levels are in the work force. However, for a particular (low) productivity level, a married woman can still be in the workforce, despite having a higher household asset level (compare, e.g., a single woman with $\log(\text{productivity})$ of -0.5 to a married woman with this same productivity level).

Hence, there are three predictions from the model. First, there should not be any single women with high assets and low productivity working. Second, there should be relatively more single women close to the borrowing constraint. Third, there should be a higher correlation between the single woman's productivity and her asset level (conditional on not being on the borrowing constraint).

Figure A.18 shows the corresponding figures from the data. I use PSID data from the years between 1998 and 2008 (PSID, 2023), and pool all observations.⁶ As can be seen, there are hardly any single women with high assets and low productivity working, compared to the same asset level for married women. I interpret "being close to the borrowing limit" as having negative assets. 23% of the single working women have negative assets. The corresponding figure for married women is 10%.

Formally, to verify if the correlation between (positive) asset holdings and observed productivity is higher for single women than for married women, I run the following regression:

$$\log(a_{ij}) = \alpha_j + \beta_j \log(\omega_i) + \gamma \mathbf{X}_i + \varepsilon_{ij} \quad (\text{A.9})$$

where i indicates the individual observation and j is the civil status: single or couple ($j \in \{s, c\}$). a is the asset holdings, ω is the observed hourly wage, and $\mathbf{X}_i$ is a set of control variables: age, age squared, and education level.⁷ β_j informs us about the correlation between the observed wage and the asset level for singles and married,

⁶ Assets include cash, bonds, stocks, real estate, cars, and is net of mortgages and other debt and are CPI adjusted using OECD (2024).

⁷ I divide the education level into three distinct groups: below 12 years, exactly 12 years, or above 12 years of schooling.

	Assets	
	(1)	(2)
$\log(\omega) \times \text{singles}$	1.296 (0.052)	0.904 (0.051)
$\log(\omega) \times \text{couple}$	0.983 (0.023)	0.692 (0.022)
Controls	No	Yes
N	16,400	16,280
$\beta_s = \beta_c$ (p-value)	0.00	0.00

Table A.6: Regression results estimating eq. (A.9). Source: PSID 1998-2008, pooling all years. The regression only includes women.

respectively. Table A.6 shows the results. As can be seen, when controlling for the age profile and education, the correlation between asset holdings and productivity becomes lower. However, the correlation is higher for single women, and the difference between the coefficients for singles and married is statistically significant for both specifications.




Figure A.17: Model results for asset holdings vs productivity level, conditional on working.



Figure A.18: Asset holdings vs productivity level for single vs married women in the data. Source: PSID 1998-2008, pooling all years. Assets include cash, bonds, stocks, real estate, cars, and are net of mortgages and other debt. I leave out asset holdings larger than USD 7 million (0.2% of the observations) and asset holdings smaller than USD -300.000 (0.1% of the observations).

A.13. Aggregate shocks during a transition path in an RA model

A natural question to ask is if the linearization method proposed by Boppart, Krusell, and Mitman (2018) (the BKM method) reproduces the results from a model with true uncertainty, if the shocks take place during a transition (as is the case in this paper). To investigate this, I make a simple example in a representative agent setting with savings and endogenous labor. I assume a utility function of the McCurdy type, with log preferences for consumption. I do the following three steps:

Step 1: Set up a model with “true” uncertainty

- Set up an RA model with z (TFP) as a state variable, solve nonlinearly
- Solve for a transition path (letting θ , the Frisch elasticity, go from 2 to 1 during 40 years, and thereafter remain)
- Simulate the model with a shock sequence for z

Step 2: Create the model analogue with the BKM method

- Set up an RA model w/o uncertainty (otherwise the same as above)
- Solve for the same transition path as above (letting θ , the Frisch elasticity, go from 2 to 1 during 40 years, and thereafter remain)
- For each year during the transition: solve for IRFs from a shock
- Simulate the model with the same shock sequence as above (using the appropriate IRFs from the right year during the simulation)

Step 3: Compare the results from the two models

It turns out that the BKM method works very well and close to perfectly replicates the results from the model with “true” uncertainty. In this particular example, the impulse response functions (IRFs) vary along the transition. The effect on impact on labor supply is approximately 40% stronger if the shock hits exactly when the transition starts, than if it happens once the economy has reached its new steady state, which in this example underlines the importance of using the IRFs from the right year.




Figure A.19: Compare results




Figure A.20: Comparing IRFs over the transition




Figure A.21: Difference if one uses final steady-state IRFs for the whole transition. Note the scale on the x-axis (otherwise one cannot see any difference). If the final steady-state IRFs are used for the transition path, the impact of the shock early during the transition path is underestimated.




Figure A.22: Difference if one uses beginning-of-transition IRFs for the whole transition. If the beginning-of-transition IRFs are used for the whole time period, the impact of shocks is overestimated later.

A.14. Robustness: Extended model

In this section I present a model with additional features designed to better capture the quantitative cyclical properties across subgroups.

Below I refer to the model specification and calibration given in the main body of the paper as the “Baseline model” and the model with various adjustments to quantitatively capture cyclical properties in the data better as outlined below as the “Adjusted model”.

A.14.1. Quantitative calibration targets for cyclical properties

The relative volatility in hours worked in the data across demographic subgroups is given in the first row of Table A.7. These numbers correspond to Figure 2(a), volatility related to aggregate fluctuations, normalized to the value for single men.

The row labeled “Baseline model” shows the steady state employment responses to a TFP shock by group for the model specification in the main body of the paper. As in the data, the model predicts higher volatility for singles than for married individuals, however, the quantitative difference is larger than what is observed in the data. Moreover, the model overestimates the volatility of women vs. men in the aggregate: the baseline model captures the micro-level evidence of women having higher labor supply elasticity rather than the macro evidence of men being more volatile in the aggregate.

A.14.2. Adjustments to the model

I first do three adjustments to the baseline model that a priori seem reasonable and (in a reduced form) capture more general phenomena. This brings the model responses in employment to a TFP shock substantially closer to the data. Thereafter, a residual adjustment is done with a group-specific unemployment shock for married men only, a much smaller adjustment than would have been necessary otherwise.

These adjustments are not time-varying, but rather capture some permanent differences across subgroups. The closing of the gender wage gap is still the *only* component of the model that is changing over time, to be able to analyze that effect in isolation.

Adjustment 1: Steady-state unemployment level The steady-state unemployment rate for each subgroup is now the average unemployment rate for that group during the period 1962 to 2018. The figures are given in Table A.1.

Since baseline unemployment is higher for singles than for married individuals, this dampens the response of singles to aggregate TFP shocks slightly.

	Married men	Married women	Single men	Single women
<i>Data:</i>				
Volatility related to agg. fluct. (normalized to single men)	0.6	0.4	1.0	0.5
<i>Baseline model:</i>				
Employment response to TFP shock (normalized to single men)	0.4% 0.1	1.4% 0.3	4.2% 1.0	3.4% 0.8
<i>Adjusted model:</i>				
Adjustment 1, 2, 3 (normalized to single men)	0.7% 0.4	0.9% 0.5	1.8% 1.0	1.0% 0.6
Adjustment 1, 2, 3, 4 (normalized to single men)	1.2% 0.6	0.7% 0.4	1.8% 1.0	1.0% 0.5

Table A.7: Volatility over the business cycle in the data and in the model. The data refers to volatility related to aggregate fluctuations (see section 2.2 in the paper). Model responses refer to immediate impact on employment from a TFP shock.

Adjustment 2: Stochastic disutility shocks In each period, households might receive an idiosyncratic temporary disutility-of-work shock. With probability one third, disutility is 18 percent lower than usual and with probability one third disutility is 18 percent higher than usual. This is to capture anything temporary that might happen in a time period making work more or less costly in disutility terms. The shock is assumed to be iid.

The implication of this disutility shock is that the bunching around the working frontier becomes less extreme. Since the bunching phenomenon is more pronounced for the single households, this lowers the response for single households' labour supply to a TFP shock relative to married individuals.

Adjustment 3: Gender differences in sector exposure In section 2.2 I show that the sector composition is very different for men and for women: men work in sectors that are much more volatile over the business cycle.⁸ To capture the sector differences between men and women (in a very reduced form), I let the TFP shock hit men (single as well as married) harder than it hits women.

To evaluate the volatility by demographic subgroup related to sectoral composition I calculate the following: for each sector I measure the sector-specific responsiveness to aggregate fluctuations (as indicated by the x-axes in Figure 4 in the paper). Then I create a weighted average for each demographic subgroup, using

⁸Up until the Covid recession, which was very different. See Alon et al. (2020) for a description of how gender patterns were completely reversed in this recession compared to previous experiences.

the fraction of the group working in each sector as the weights (as indicated by the size of the bubbles in Figure 4 in the paper). Normalizing the resulting values to single men gives:

- Single men: 1.0
- Married men: 1.0
- Single women: 0.6
- Married women: 0.6

Thus, given this measure of sector-based cyclical exposure, there is no difference between singles and couples, and I therefore use a gender-specific adjustment, but do not distinguish between couples and singles.

The wage for women before the adjustment is given by (omitting the time subscripts for readability):

$$(1 - \tau_f)\omega_f w = (1 - \tau_f)\omega_f z(1 - \alpha)K^\alpha L^{-\alpha} \quad (\text{A.10})$$

with τ_f denoting the female productivity wedge, ω_f the woman's current underlying productivity, z current TFP state, K total capital stock, and L total efficiency units of labor. To make women less affected by a TFP shock than men, I instead give the women the following wage (and marginal productivity):

$$(1 - \tau_f)z_f(1 - \alpha)K^\alpha L^{-\alpha} \quad (\text{A.11})$$

$$\log z_f = \vartheta_z \log z \quad (\text{A.12})$$

where $\vartheta_z \in [0, 1]$ represents the dampened effect a TFP shock has on women compared to men.

In practice in the model, I adjust current τ_f , the gender wage gap, by a factor ϑ_τ so that it is equivalent to the effect of a TFP shock on women being dampened by a factor of ϑ_z (still omitting the time subscripts for readability):


$$(1 - \tau_f^f \vartheta_\tau)z\alpha K^\alpha L^{-\alpha} = (1 - \tau_f)z_f\alpha K^\alpha L^{-\alpha} \quad (\text{A.13})$$

Thus, the dampening factor is given by:

$$\vartheta_\tau = \frac{1}{\tau_f} - \frac{z_f}{z} \cdot \frac{1 - \tau_f}{\tau_f} \quad (\text{A.14})$$

I set $\vartheta_z = 0.7$. Thus, if the TFP shock hitting males is 1 percent, the females only experience a TFP shock of 0.7 percent, and that is constructed by choosing ϑ_τ accordingly.

This adjustment of course also affects the total efficiency units of labor, L , in the economy.



(a) Productivity gap ($1 - \tau_f$), and the median wage gap in the data.

(b) Resulting transition path for employment by subgroup.

Figure A.23: Transition path in Adjusted model (adjustments 1, 2, 3, and 4).

Adjustment 4: Specific unemployment shock for married men The three adjustments described above take us reasonably far matching volatility patterns in steady state, as shown in Table A.7. However, married men are still slightly less volatile than their wives in the model. There could be a number of reasons for why the volatility of married men—relative to other subgroups—is proportionally less driven by supply choices captured by the model.

To mechanically induce the right volatility patterns in the model in the steady-state, I finally introduce a specific unemployment shock for married men, assuming that the flow into unemployment changes by a small fraction inversely proportional to the TFP shock. For every percent z , the TFP level, is above its normal state, the flow into unemployment is 0.18 percentage points lower than in normal times.

With this final adjustment, the relative size of the immediate response to a TFP shock across demographic subgroups is close to what we observe in the data. Importantly, men are more volatile than women (both married and singles) and singles are more volatile than individuals in couples (both men and women) and the magnitude of the difference is in line with the data. See Table A.7 for an overview.

Recalibration These changes also requires adjustments of the preference parameters to hit the relevant employment targets in steady state. ζ_1 , the consumption weight for couples, and ψ , the baseline disutility of labor, are decreased by 5% each.

A.14.3. Long-run behavior of the adjusted model

As before, the shrinking gender wage gap is the only thing that varies over time. The resulting employment paths for the different subgroups are shown in Figure A.23(b).

The main take-aways for the long-run behavior remain unchanged compared to the baseline model: married women's employment rate increases the most, married

men's employment rate falls slightly, but the husbands are far from crowded out one-for-one by their wives entering the labor market. Single women are hardly affected by the shrinking gender wage gap since for this group, income and substitution effects cancel. Lastly, single men are only affected by general equilibrium forces. Therefore, their labor supply is, not surprisingly, hardly affected.

A.14.4. Cyclical behavior of the adjusted model

Then it is time to turn to the cyclical properties of the model over time. Table A.8 shows the immediate impact on employment from a TFP shock of one standard deviation (on men, for women the shock is smaller, as discussed above). As the table shows, the main picture from the baseline model remains. Married men as a group have increased their response over time. In 1970, the effect of an aggregate shock to their employment level was two thirds of what is estimated for steady-state "today". The response of married women on the other hand has decreased over time. In 1970, their response was 15% stronger than today. This decreasing response is qualitatively the same as in the baseline model, but the decrease is somewhat smaller quantitatively.

In Table A.8, the employment responses are reported compared to the underlying deterministic trend. Now I turn to actual employment *levels*. If the economy is hit by a negative TFP shock, how does employment react, and how does the reaction depend on where the economy is in the transition path? To see if the new adjusted model still captures the historical facts of employment effects after a recession, I run the same experiment as in the main paper: The economy is shocked with a negative TFP shock of two standard deviations during two points in time. Those two points are chosen to be 1970, to capture the economy's response in the middle of the transition, and 2010, to capture the response when the economy has almost finalized the transition. Figure A.24 shows the resulting employment figures, normalized to one in the year before the shock hits. Thus, Figure A.24 corresponds to Figure 14 in the main paper.

As Figure A.24 shows, the insights from the main paper remain. As can be seen from Figure A.24(a), when the economy is hit by a negative TFP shock in 1970, total employment falls. Single men are most affected immediately on impact. Four years later, the economy has returned to its initial total employment level (even slightly above). The subgroup driving this recovery is married women, and to some extent single women. After four years, neither single men nor married men have returned to their pre-recession employment level. Figure A.24(b) shows the same type of TFP shock, but in the year 2010, i.e., when the employment figure for married women has stabilized at its current level. Now, four years after the shock, the employment

	Married men	Married women	Single men	Single women	Total employment
End steady-state	1.2%	0.7%	1.8%	1.0%	1.1%
<i>Response during the transition:</i>					
1970s	0.8%	0.8%	2.0%	0.6%	0.9%
1990s	1.0%	0.8%	2.0%	0.8%	1.0%
2010s	1.2%	0.7%	1.9%	1.0%	1.1%

Table A.8: Immediate impact on employment from a TFP shock of one standard deviation. Impact measured as the deviation from the deterministic underlying trend.




Figure A.24: The economy's response to a negative TFP shock (two standard deviations) in terms of employment. The employment figures are normalized to one in the year before the shock hits.

rates have still not returned to their pre-shock level. The main difference is that married women are no longer driving up the employment figures. Again, as in the data, single men are the group most severely hit by the negative shock in terms of employment.

A.15. Robustness: A model with progressive taxation

In this section I evaluate how the results would change with a progressive tax schedule. Otherwise the model is identical to the baseline model (except for recalibration of certain parameters discussed below). Thus, the results reported in this section is for a utility function of the King, Plosser, and Rebelo (1988) type which features additive separability between consumption and labor.

Following Heathcote et al. (2014), I define the net tax and transfers to be paid as a function of total earnings in the family:

$$T(\iota) = \iota - \lambda \iota^{1-\varrho}$$

where total earnings $\iota = w_f e_f + w_m e_m$ for couple households and $\iota = w_i e_i$, $i \in \{m, f\}$ for single households. Loosely speaking, λ defines the level of taxation, while ϱ defines the progressivity of the tax system. I follow Brinca et al. (2016) and use $\varrho = 0.137$ while I fix λ such that on average over time, the taxes paid in approximately equals the taxes paid out.⁹

In practice, this means that for single men and women, the difference compared to the model with a lump-sum tax is that the tax is zero if individuals are not working, while when they are working, the variation in net income is compressed (less risky).

For married couple, the total taxes paid is more complicated and a function of both productivities and work choices by the husband and wife respectively.

Table A.9 compares the outcome from the baseline model with lumpsum taxes with a model with progressive household taxes. To make a fair comparison of the two models, the preference parameters are recalibrated. To (approximately) hit the employment levels, I do the following changes in the progressive-tax model: a) the consumption scaling for couple households is increased by 13%; b) the baseline level of disutility of work is decreased by 13%; c) the disutility for females is now 3.8% higher than for males (was previously 7.6% higher).

As the table shows, the main insights remain: the short-run responses to a temporary shock are similar (even though the single women respond slightly more in the progressive-tax version of the model). This shows that the short-run responses on the extensive margin to temporary wage changes are more sensitive to how many households are close to their reservation wage than small changes in the size of the wage change.

⁹To have government budget balance in each period would require a time-shifting λ , and that in each period the model must be solved for the correct value, since the work hours and thus the taxes paid into the system is an endogenous outcome of the tax level. I don't think this is necessary to get the intuition for how the model does (not) change with this type of taxes and since I don't do any welfare calculations it is not crucial, and therefore I use this approximation.

	Couples	Single men	Single women
<i>Panel A: Baseline model with lumpsum tax</i>			
Employment (%)	86.7, 67.7	75.5	74.6
Consumption	1.9	1.0	0.7
Assets	2.1	1.0	0.6
<i>Immediate impact on employment from a TFP shock of one std</i>			
Percent (%)	0.4, 1.4	4.2	3.4
<i>Panel B: Model with $T(\iota) = \iota - \lambda \iota^{1-\epsilon}$</i>			
Employment (%)	85.5, 68.4	74.2	70.6
Consumption	1.9	1.0	0.8
Assets	2.4	1.0	0.7
<i>Immediate impact on employment from a TFP shock of one std</i>			
Percent (%)	0.5, 1.4	4.7	5.1

Table A.9: Model results: comparing a model with lumpsum taxation (Panel A) with a model with progressive taxation (Panel B).




Figure A.25: Resulting transition path for employment by subgroup.

Turning to the long-run behavior, Figure A.25 shows the resulting transition paths for employment by subgroup for the baseline model and the model with progressive taxation. As can be seen, the results are very similar.

To understand why the changes are so small in the long run, there are two important points.

First, what is relevant for the long-run response is the increase in the ratio of “male-earnings-only” compared to “earnings-if-both-work” between the beginning and the end of the transition. This does not change very much with a progressive tax schedule. Think about a simple example with a couple, in which the male earns 1. The woman in the beginning of the transition earns 0.55, and in the end of the transition earns 0.8. Thus, this couple represents the average.

In the beginning of the transition, the ratio of “earnings-if-both-work”-to-“male-

earnings-only" is $1.55/1 = 1.55$. In the end of the transition, the ratio is $1.8/1 = 1.8$. Thus, the change in this ratio was $1.8/1.55 = 1.16$, i.e. 16%.

Now introduce a progressive tax, assume $\lambda = 1$ and $\varrho = 13.7\%$. Then the income if only the male works is still 1 in all scenarios. However, the income if both work is now 1.46 in the beginning of the transition and 1.66 in the end of the transition. Thus, the change in ratio is given by $1.66/1.46 = 1.14$, which is only slightly lower than the ratio in the example without progressive taxes.

Second, there is an additional more subtle point: with the lump-sum tax formulation in the baseline formulation of the model, each household has a (negative) non-labor income. Remember that for married women an increase in female wages has an effect (as opposed to for single women) because the income effect of a permanent wage change is muted by the income of her husband. Thus, the negative non-labor income in form of a lump-sum tax muted the married women's employment response. If the negative lump-sum transfer is removed, this makes the married women's non-labour income slightly larger, and thus the income effect is even more muted, and the employment response is slightly stronger. This effect is relatively small, but close to perfectly counteracts the small increase in the earnings ratios discussed in the previous point.