



Full length article

The aggregate implications of gender and marriage



Margherita Borella^a, Mariacristina De Nardi^b, Fang Yang^{c,*}

^aUniversity of Torino and CeRP-Collegio Carlo Alberto, Italy

^bUCL, Federal Reserve Bank of Chicago, CEPR, IFS, and NBER, United Kingdom

^cLouisiana State University, United States

ARTICLE INFO

Article history:

Available online 2 February 2017

Keywords:

Gender
Marriage
Wage

ABSTRACT

Wages, labor market participation, hours worked, and savings differ by gender and marital status. In addition, women and married people make up a large fraction of the population and of labor market participants, total hours worked, and total earnings. For the most part, macroeconomists have been ignoring women and marriage in setting up structural models and in calibrating them using data on males only. In this paper, we ask whether ignoring gender and marriage in both models and data implies that the resulting calibration matches well the key economic aggregates. We find that it does not and we ask whether there are other calibration strategies or relatively simple models of marriage that can improve the fit of the model to aggregate data.

© 2017 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Introduction

Wages, labor market participation, hours worked, and savings differ by gender and marital status. In addition, women and married people¹ make up a large fraction of the population and of labor market participants, total hours worked, and total earnings.

For the most part, macroeconomists have been ignoring women and marriage in setting up life-cycle structural models and in calibrating them using data on males only. In this paper, we ask whether ignoring gender and marriage in both models and data implies that the resulting calibration matches well the key economic aggregates of labor participation, hours worked, labor income, and net worth. We find that it does not and we ask whether there are other calibration strategies or relatively simple models of marriage that can improve the model fit along these important dimensions.

To investigate the aggregate importance of gender and marriage and to determine what might be the simplest model that best captures the most important aggregates, we construct and calibrate four different economies. Economy 1 is a “No marriage, only men”, economy that adopts a standard one-gender, no marriage, life-cycle framework, and only uses data on men for calibration purposes, as usually done in quantitative macro models. We find that this model economy (and calibration) misses the observed economy’s aggregate outcomes, including labor supply, earnings, and

hours by a large amount over all of the working period. More specifically, this economy drastically overestimates participation by about 10–20 percentage points, overestimates average hours by over one-third of actual aggregate hours, and also overestimates average earnings by over one-third over the entire working period. It also underestimates retirement savings.

Economy 2 is a “No marriage, men and women together” economy that uses the same model as Economy 1, but is calibrated using data on both men and women together, as individual-level data, thus ignoring any gender differences and whether individuals are in couples or not. If this kind of “aggregation” were to match the aggregates well, this would be a possible way to take gender and marriage into account without writing a more complicated model. The biggest success of this calibration compared with Economy 1 is the match in labor earnings over the life cycle. Unfortunately, however, this calibration still misses the observed patterns of participation and hours over most of the life cycle, with the exception of the period between ages 45 and 55. For instance, at age 30, aggregate participation is about 80%, while the model predicts close to 98%, and aggregate average hours are about 1,600 a year, while the model predicts close to 2,000. In this case, too, the model underpredicts retirement savings and generates faster asset decumulation after retirement than in the data. These discrepancies hold even when we allow for an age-varying fixed cost of participating in the labor market, calibrated to better match participation. This calibration thus constitutes an improvement over Economy 1, but is still lacking in many respects.

Economy 3 is a “No marriage, household-level calibration for couples only” economy that uses the same model as Economies 1 and 2, but in the calibration we aggregate the data at the house-

* Corresponding author.

E-mail address: fyang@lsu.edu (F. Yang).

¹ We use the terms marriage and couples interchangeably. In the data work, we consider people married or in couples if they are cohabiting and/or married.

hold level and we only keep couples. This is in contrast with Economies 1 and 2, in which we were only looking at individual-level data, ignoring that some people are in couples and some others are not. We find that this calibration has a similar fit to labor earnings, participation (performing better before age 45 and worse after), and hours worked (fitting well only during the ages 40–55) as Economy 2. Thus, if one wants to avoid modeling marriage, we conclude that it is best to adopt the calibration in Economy 2 or 3 and take into account both gender and possibly household structure in the calibration. This being said, Economies 2 and 3, even with an improved calibration, still miss key aspects of participation and labor supply over the life cycle.

Economy 4 is a “Marriage and singles, men and women” economy in which we explicitly model married and single men and women over their whole life cycle. The model is calibrated to the observed data for the four relevant groups: married men, married women, single men, and single women. This model does much better than all the other economies that we consider in matching the aggregate data. In terms of discrepancies between the model aggregates and actual aggregates, Economy 4 tends to underpredict labor participation between ages 40 and 50 by less than 8 percentage points and tends to overpredict labor income by less than 3 percentage points over all of the life cycle. Compared with the economies without two-agent couples and gender differences, these misses are small. Thus, this comparison indicates that modeling gender and marriage and the related economic incentives is important in explaining key economic aggregates over the life cycle.

Based on these findings, we thus conclude that even macroeconomists not interested in heterogeneity in marriage and gender per se should start taking marriage and gender differences into account in the context of quantitative structural models. In particular, modeling gender and marriage explicitly would yield the best results in terms of matching the aggregates, but if this is not possible given the question and complications at hand, calibrating (or estimating) the model including both men and women in the data or keeping track of households to determine household-level wages and hours (or earnings), participation, and assets will help the model better match the aggregates.

Our paper provides two main contributions. First, it documents important aspects of the data, both in the aggregate and over the life cycle, for single and married men and women. Because marriage and female labor supply patterns have changed a lot over the last seventy years, we apply our model to data from the Panel Study of Income Dynamics (PSID) and the Health and Retirement Survey (HRS) for the 1941–45 cohort. Women and married people make up for a large fraction of workers, hours, and earnings in the aggregate economy for our cohort of interest. The fraction of workers who are women increases from 37% at age 25 to 44% at age 65. The fraction of hours worked by women as a fraction of total hours rises from 28% at age 25 to 40% at age 65, while the fraction of earnings by women rises from 24% at age 25 to 27% at age 65. Married people earn over 84% of total earnings and contribute over 85% of the hours worked during the whole life cycle of this cohort. Over the life cycle, we find that, for this cohort, married men display on average the highest labor participation rate, over 98%, until they turn 40, and only slowly decrease participation until age 50, while the participation of single men starts dropping fast after age 40. The participation of single women is a shifted-down version of that of married men's by about 10 percentage points. Married women have an even lower participation, which is hump-shaped over the life cycle and peaks at 50% around age 45. In addition, women not only are less likely to participate in the labor market than men, but also display lower average hours conditional on participation.

Second, besides considering a standard life-cycle model and calibrating it either to males or all people regardless of their gender,

or households, our paper also constructs a structural and dynamic life-cycle model that explicitly models single and married men and women and calibrates it to both PSID and HRS data, and compares the aggregate implications of our four models. We calibrate the parameters of each model economy to match the observed data as well as possible for the relevant group in consideration (as done in the previous literature), and we then investigate how well each calibrated model can match the aggregate economy, which includes single and married men and women.

Related literature

The vast majority of papers studying quantitative life-cycle macroeconomic questions use data on males only. Notable exceptions to this include Tertilt (2005), which studies the effects of marriage institutions (and polygyny more specifically) on aggregate output; and Doepke and Tertilt (2016), which advocates modeling the family and gender to understand various historical changes, both in the short and the long run. Our paper focuses on gender and marriage for the purpose of understanding the economy's aggregates at a point in time, using U.S. data for one cohort.

Our work builds on two main branches of the literature. One such branch studies the determinants of life-cycle female labor supply, typically assuming that male labor supply is fixed and sometimes abstracting from savings. Attanasio et al. (2008) and Eckstein and Lifshitz (2011) point to the importance of changing wages and child care costs to explain increases in female labor supply over time. Eckstein et al. (2016) examine the changes over time in the selection of married women working and find that it accounts for 75% of the observed increase in the marriage-wage premium (the differential in salary for married versus single women) increase over time. Guner et al. (2012) find that gender-based taxes, implying that women face lower and proportional income tax rates, increase output and female labor participation and improve welfare. Kaygusuz (2015) studies the effects of the Economic Recovery Tax Act of 1981 on married female labor force participation. Nishiyama (2015) finds that removing spousal and Social Security survivor benefits would increase female labor participation, female hours worked, and aggregate output. Low et al. (2016) study how marriage, divorce, and female labor supply are affected by welfare programs in the U.S. Blundell et al. (2016) study how the U.K. tax and welfare system affects the career of women. Compared with this set of papers, we allow for savings and both intensive and extensive labor supply decisions for both men and women, and we take our model to data by using the PSID and the HRS. In addition, we study the implications of gender and marriage on the economy's main aggregates by using a set of models that take different stances on marriage and gender.

Another branch of the literature models the joint retirement behavior of couples (Blau, 1998; Blau and Gilleskie, 2006; van der Klaauw and Wolpin, 2008; Casanova, 2012). Although we do allow for endogenous labor supply and participation and people in our models start reducing work efforts as they age, we take the maximum retirement age to be exogenous and leave the question of benefit claiming for married and single people for future work.

The data

We use both PSID and HRS data for the cohort of men and women born between 1941 and 1945.² We pick one cohort to

² See Appendix A for a discussion of these data sets and details about our computations.

Table 1
Fraction of married and single men and women by age and during the working period.

Group	Age 25	Age 35	Age 45	Age 55	Age 65	Ages 25–65
Married women	0.43	0.42	0.40	0.39	0.37	0.40
Married men	0.43	0.46	0.44	0.43	0.44	0.45
Single women	0.07	0.07	0.10	0.12	0.13	0.09
Single men	0.07	0.05	0.06	0.06	0.06	0.06

Table 2
Top panel: Fraction of women, fraction of hours worked by women, and fraction of earnings by women in a given age group or the whole working period. Bottom panel: Fraction of married people, fraction of hours worked by married people, and fraction of earnings earned by married people in a given age group.

Group	Age 25	Age 35	Age 45	Age 55	Age 65	Ages 25–65
Women among workers	0.37	0.40	0.46	0.46	0.44	0.42
Hours worked by women	0.28	0.31	0.39	0.40	0.40	0.35
Earnings by women	0.24	0.22	0.30	0.27	0.27	0.26
Married among workers	0.86	0.85	0.84	0.82	0.78	0.84
Hours worked by married	0.86	0.86	0.84	0.83	0.80	0.85
Earnings by married	0.88	0.87	0.86	0.87	0.85	0.87

abstract from changes in behavior across cohorts. We pick this particular cohort so that their entire adult life is first covered by the PSID, which starts in 1968 (from which we use the rich information about wages for men and women), and then by the HRS, which starts covering people at age 50 in 1994 (and has rich information on mortality by many observable characteristics). Thus, this is a cohort for which we have excellent data over most of their working life and retirement.

Gender and marital status in the aggregate

Table 1 reports the fraction of married and single men and women by age group and over the working period. For instance, in this cohort, at age 25, 43% of the people are married men and 43% married women, while 7% are single men and 7% are single women. At age 65, 37% of people are married women, 44% are married men, 13% are single women, and 6% are single men. As we can see from this table, in each age group, the vast majority of people are married and the fraction of married women in the population is large.

We now turn to key economic outcomes pertaining to our four groups. Table 2 shows that women and married people not only are numerous, but that many of them also work and they make up for a sizeable fraction of the aggregate economy. The first line reports the fraction of workers who are women by age. For the cohort that we consider, this number increases from 37% at age 25 to 44% at age 65. The next two lines report the fraction of hours worked by women as a fraction of total hours and the fraction of total earnings by women. Both are sizeable. The fraction of hours worked by women rises from 28% at age 25 to 40% at age 65, while the fraction of earnings by women raises from 24% at age 25 to 27% at age 65. The bottom panel of the table shows that married people make up for the majority of the population at each age group, going from 86% at age 25 to 78% at age 65. The fraction of hours worked by married people follows the same pattern from 86% at age 25 to 80% at age 65, while the fraction of earnings by married people stays around 87% and declines only very slightly at age 65.

Thus, the aggregates (earnings, hours, and number of people working) are comprised of large fraction of women and married people, who are not typically included in the calibration, let alone modeled explicitly, in standard macro models. Besides noting that excluding these households in the calibration will very likely miss key aggregates, we should also note that the savings and labor supply of people in couples might be very different from those of the single decision maker that is typically modeled. In fact,

single decision makers are a minority in the data and also earn a relatively small fraction of total earnings at all ages. The next section turns to comparing and describing aggregate and disaggregated patterns of important variables over the life cycle to shed more light on behavior and aggregation.

Life-cycle patterns

We now turn to the life-cycle patterns of participation, hours, labor income, and wealth drawn from the data for the cohort that we focus on. We start by displaying figures for the disaggregated profiles by gender and marital status, which we use in Economy 4 and then turn to showing the profiles for the aggregates and all men only, both married and single, (with the latter variable and set of profiles being the relevant variable used in most macro models).³

Single and married men and women over the life cycle

To start understanding the role of gender and marital status in affecting life cycle labor market outcomes and assets accumulation, we study life-cycle profiles by gender and marital status. Fig. 1, panel (a), displays average labor participation for single and married men and women.

Married men display on average the highest participation rate, over 98%, and only slowly decrease their participation from age 40 to 50. The participation of single men, instead, drops fast starting at age 40. The participation of single women displays a very similar shape to that of married men, but its level is shifted down by about 10 percentage points. Married women have the lowest participation. It starts around 50% at age 25, then increases to 78% at age 40, holds there until age 50, and then gradually declines at a similar rate to that of the other three groups.

Average hours are shown in panel (b) of Fig. 1 and follow broadly similar patterns. Married men on average work more hours than everyone else. Women not only have a participation

³ To construct these graphs, we take averages using the PSID data for individuals born between 1936 and 1950; that is, we include two five-years cohorts adjacent to the cohort of interest (the one born in 1941–45) in order to have enough observations in each age cell and at the same time minimize the impact of cohort effects. To characterize the life-cycle pattern of assets, we adopt a somewhat different procedure because in the PSID assets are observed only in a few waves. For that, we select individuals born between 1931 and 1955 to have enough observations and estimate a fourth-order polynomial in age with ordinary least squares following individuals up to age 75 (age 70 in the case of single men and single women). It is this estimated polynomial in age that is plotted in our graphs.

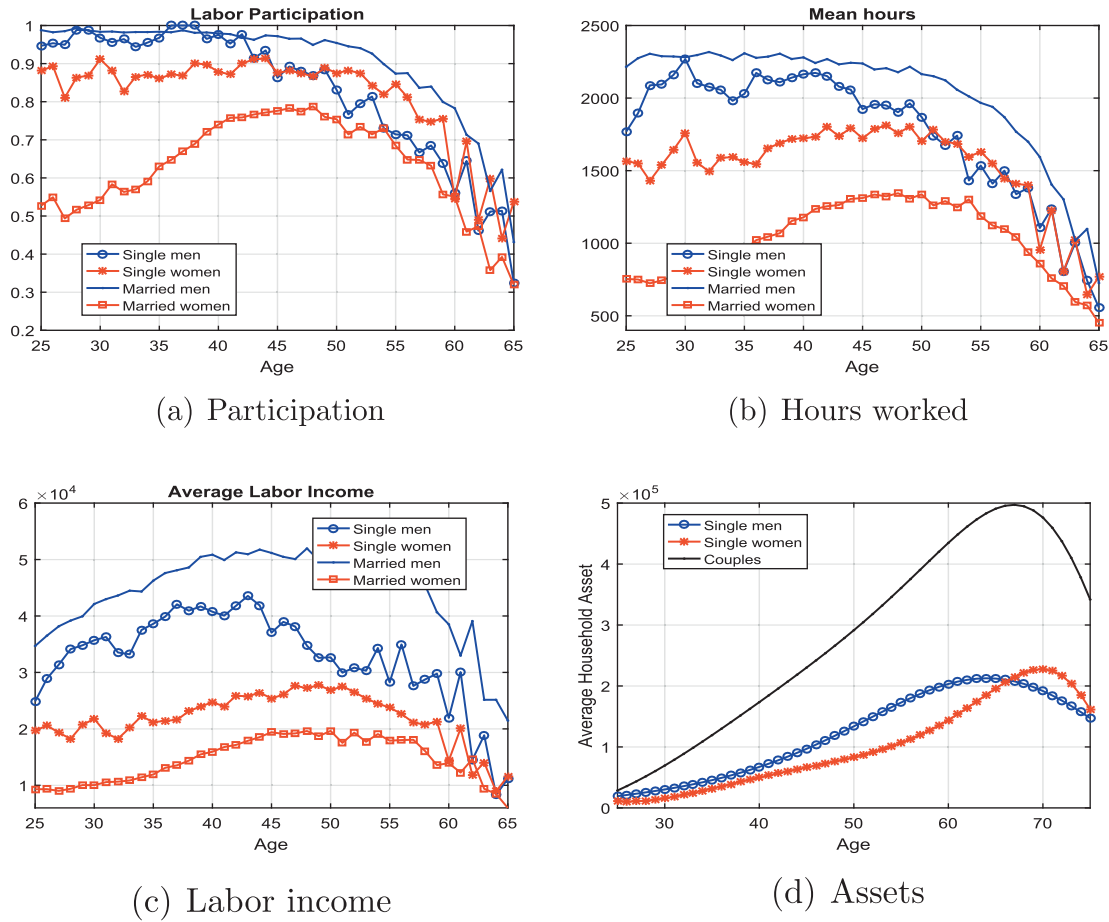


Fig. 1. Life-cycle profiles by gender and marital status.

rate lower than men on average, but also display lower average hours, even conditional on participation.

Turning to labor income and assets for single and married men and women, panel (c) shows that married men have much higher labor income than any other group, because of their higher average number of hours and hourly wage (which we document in the calibration section). Women, both single and married, have a much lower average income at all ages, due to both their low average number of hours and low hourly wage. While for single men, average labor income drops after age 45, for women, both single and married, it increases until age 50, due to their increasing participation profile and hours. In panel (d), we report average assets for married and single men and women. Average assets increase until age 65 for all groups, with women accumulating the lowest amount and showing no sign of a slowdown in accumulation before age 70.

The aggregates and men over the life cycle

The graphs in Fig. 2 display two lines, the solid line corresponds to the aggregates in the economy, including men and women, while the circled line corresponds to the aggregates for males only.

Graph (a) shows the lifecycle pattern of participation. The participation of men stays around 98% until age 40 and then starts declining. The rate of decrease is faster starting at age 50, with male participation dropping to 40% by age 65. Comparing it with the lines for women, reveals that the gap in participation rate between the men only and everyone of working age is over twenty percentage points at age 25 and only gradually shrinks over the life cycle, because married women participate more later on and then both men and women start to reduce their participation as they

age. This graph shows that only considering male participation in the labor market drastically overestimates overall aggregate participation.

Average hours (panel (b) in the same figure) follow a similar pattern, starting at around 2200 hours at age 25 for men, increasing to 2300 hours by age 27, and then declining slowly after age 40, and much faster after age 50, reflecting decreasing participation. In contrast, average hours among people of working age, including men and women, start out at a much lower level, around 1500 a year at age 25 (thus 800 hours lower than for men) and only gradually increase, to peak at about 1700 hours around age 45. Thus, the gap in average hours worked also shrinks over the life cycle but remains large over most of it and only looking at average hours worked by men overestimates average hours worked.

Panel (c) highlights that average labor income for men increases gradually to \$50,000 by age 40 due to the increasing average hourly wage shown in Fig. 3, stays roughly constant until age 55, and then drops sharply after that age due to drops in both average hourly wage and hours worked. The corresponding labor income for the aggregate population including men and women is largely a shifted-down version of the one for men by about \$12,000 a year until age 55. At age 25, this corresponds to overestimating average labor earnings in the population by about one-third.

As panel (d) of the figure shows, average assets for men tend to increase until age 65, reaching \$250,000, although after age 50 the growth rate gradually slows down. Those for the overall population turn out to be very similar. This is unsurprising, given that most people are married and we split asset ownership between partners equally to plot data for men and women separately.

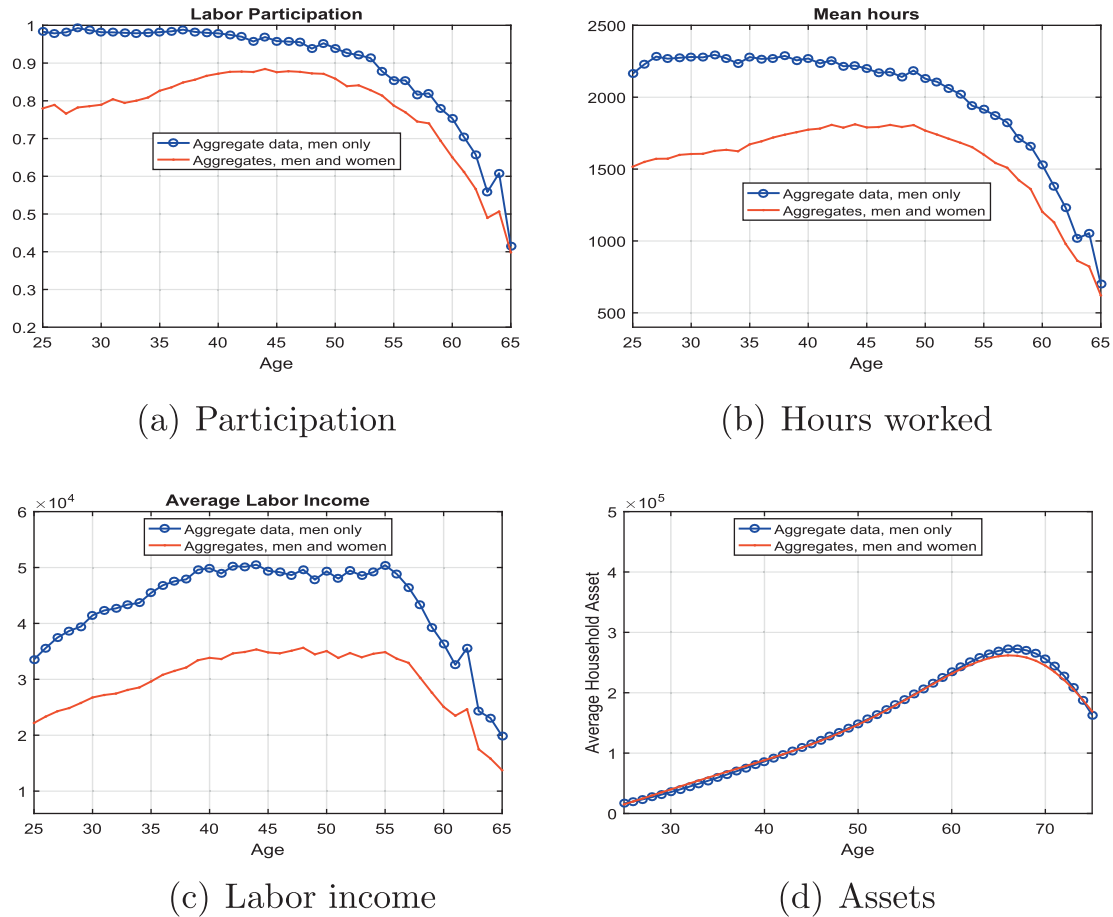


Fig. 2. Life-cycle profiles for men.




Fig. 3. Potential wages over the life cycle for men and women born in the 1941–45 cohort (PSID data), for men and women together regardless of gender, and at the household level (average between partners) for people in couples.

The full model with marriage and singles

In this section, we describe the full model with couples and singles that we use in Economy 4. Economies 1, 2, and 3 only include single people and are thus special cases of this model.

Our model period is one year. We explicitly model the working and retirement stages of the life cycle. Let t be age $\in \{t_0, t_1, \dots, t_r, \dots, t_d\}$, with t_r being retirement time and t_d being the maximum possible lifespan. People start their economic life

at the age of 25 and live up to the maximum age of 100. They retire at age 66 and from that time on, they face mortality risk.

During their **working stage**, people are alive for sure, face shocks to their wages, and are either single or married. Each household, whether single or married, chooses how much to save for next period and how much to work, with married people choosing the labor supply of both partners. As in French (2005), we introduce a fixed time cost of working for each person, which implies that, consistently with the observed data, most people will not choose to work just a few hours. This fixed time cost of working includes commuting time and time spent getting ready for work.

After **retirement**, the only control variable is savings and each person faces an exogenous probability of death, which depends on their gender and marital status. Thus, because of mortality risk after retirement, married people may lose their spouse during that time. We use the superscript i to denote gender, with $i = 1$ being a man and $i = 2$ being a woman. We use the superscript j to denote marital status, with $j = 1$ for singles and $j = 2$ for couples. We allow the fixed cost of working to differ by gender and marital status and, for married women, also by age, to incorporate the costs of raising children in a parsimonious way. We allow the earnings processes to depend on gender.

The government

The government taxes labor income using a proportional tax to finance old-age Social Security. Social Security benefits are taken from the 1997 average payments for the groups that we study and thus depend on gender and marital status. We balance the Social Security budget for the cohort that we consider.

Single men and women

Consider single people of working age. They have preferences given by

$$v(c_t, l_t) = \frac{(c_t^\omega l_t^{1-\omega})^{1-\gamma} - 1}{1-\gamma} \quad (1)$$

where c_t is consumption and l_t is leisure, which is given by

$$l_t = 1 - n_t - \phi_t^{i,1} I_{n_t}, \quad (2)$$

that is, total time endowment less n_t , hours worked in the labor market, less the fixed time cost of working. I_{n_t} is an indicator function that equals 1 when hours worked are positive and zero otherwise. The term $\phi_t^{i,1}$ represents the fixed time cost of working for singles of gender i at age t .

Let e_t^i be a deterministic age-efficiency profile, which is a function of the individual's age and gender. Let ϵ_t^i be a persistent earnings shock that follows a Markov process. The product of e_t^i and ϵ_t^i determines an agent's units of effective labor per hour worked during the period.

$$\ln \epsilon_{t+1}^i = \rho^i \ln \epsilon_t^i + v_t^i, \quad v_t^i \sim N(0, \sigma_v^2). \quad (3)$$

Let a_t be assets that earn interest rate r . The state variables for a single individual are age t , gender i , asset a_t^i , and the persistent earnings shock ϵ_t^i . From the first period of working age until retirement, the recursive problem of the single person of gender i can thus be written as

$$W_t^{s,i}(a_t^i, \epsilon_t^i) = \max_{c_t, a_{t+1}, n_t} \left[v(c_t, 1 - n_t - \phi_t^{i,1} I_{n_t}) + \beta E_t W_{t+1}^{s,i}(a_{t+1}^i, \epsilon_{t+1}^i) \right] \quad (4)$$

$$Y_t = e_t^i \epsilon_t^i n_t \quad (5)$$

$$c_t + a_{t+1}^i = (1+r)a_t^i + (1-\tau_{SS})Y_t \quad (6)$$

$$a_t \geq 0, \quad n_t \geq 0, \quad \forall t \quad (7)$$

The expectation operator is taken with respect to the distribution of ϵ_{t+1}^i conditional on ϵ_t^i .

After retirement, single people face a positive probability of dying every period. The retired individual's recursive problem can be written as

$$R_t^{s,i}(a_t) = \max_{c_t, a_{t+1}} \left[v(c_t, 1) + \beta s_t^{s,i} R_{t+1}^{s,i}(a_{t+1}) \right] \quad (8)$$

$$c_t + a_{t+1} = (1+r)a_t + (1-\tau_{SS})Y_r^{ij} \quad (9)$$

$$a_t \geq 0, \quad \forall t \quad (10)$$

The term s_t^i is the survival probability, which is a function of age and gender.

Married couples

Couples maximize their joint utility function and their utility from total consumption and from the leisure of each household member is given by⁴

$$w(c_t, l_t^1, l_t^2) = \frac{\left(\left(\frac{c_t}{2}\right)^\omega (l_t^1)^{1-\omega}\right)^{1-\gamma} - 1}{1-\gamma} + \frac{\left(\left(\frac{c_t}{2}\right)^\omega (l_t^2)^{1-\omega}\right)^{1-\gamma} - 1}{1-\gamma}, \quad (11)$$

During the working period, each of the spouses is affected by a wage shock, which is realized and known at the beginning of each working period. As for singles, the superscript $i = 1$ refers to men, while the superscript $i = 2$ refers to women. Spouses differ in their earnings processes and initial wage shocks. The state variables for married couples are $a_t, \epsilon_t^1, \epsilon_t^2$.

The recursive problem for the married couple of working age can be written as

$$W_t^c(a_t, \epsilon_t^1, \epsilon_t^2) = \max_{c_t, a_{t+1}, n_t^1, n_t^2} \left[w(c_t, 1 - n_t^1 - \phi_t^{1,2} I_{n_t^1}, 1 - n_t^2 - \phi_t^{2,2} I_{n_t^2}) + \beta E_t W_{t+1}^c(a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2) \right] \quad (12)$$

$$Y_t^i = e_t^i \epsilon_t^i n_t^i \quad i = 1, 2 \quad (13)$$

$$c_t + a_{t+1} = (1+r)a_t + (1-\tau_{SS})(Y_t^1 + Y_t^2) \quad (14)$$

$$a_t \geq 0, \quad n_t^1, n_t^2 \geq 0, \quad \forall t \quad (15)$$

The expected value of the couple's value function is taken with respect to the conditional probabilities of the two ϵ_{t+1}^i , given the current values of ϵ_t^i for each of the spouses (we assume independent draws).

During retirement, that is from age t_r onwards, each of the spouses is hit with a realization of the probability s_t^i . We assume that the death of each spouse is independent from that of the other. The married couple's recursive problem during retirement can be written as

$$R_t^c(a_t) = \max_{c_t, a_{t+1}} \left[w(c_t, 1, 1) + \beta s_t^{c,1} s_t^{c,2} R_{t+1}^c(a_{t+1}) + \beta s_t^{c,1} (1 - s_t^{c,2}) R_{t+1}^{s,1}(a_{t+1}) + \beta s_t^{c,2} (1 - s_t^{c,1}) R_{t+1}^{s,2}(a_{t+1}) \right] \quad (16)$$

$$c_t + a_{t+1} = (1+r)a_t + (1-\tau_{SS})(Y_r^{1,c} + Y_r^{2,c}) \quad (17)$$

$$a_t \geq 0 \quad \forall t \quad (18)$$

Calibration

In all of our four model economies, we assume that the U.S. is an open economy and we set the interest rate r to 4% and the risk aversion parameter γ to 2. The Social Security tax, τ_{SS} , on workers is set to be 3.8%, which was the worker's tax rate for the Old-Age, Survivors, and Disability Insurance program in 1968 (See Table 3). We choose this year because we focus on birth cohorts of 1941–45, who were age 23–27 in 1968, and thus had recently started working in that year.

Other inputs needed by our model economies include potential wages, the stochastic processes for wages, and mortality probabilities. We now turn to describing their estimation.

Turning to potential wages, Fig. 3 displays four lines: mean potential wage by gender (solid line for men and starred line for women), mean potential wage obtained ignoring gender and marital status differences and thus pooling all men and women together (dotted line), and average potential wage at the household level for couples, expressed in per-capita terms (dashed line).

Table 3
Calibration of the interest rate, risk aversion, and Social Security tax rate.

Parameters		Value
r	Interest rate	4%
γ	Risk aversion coefficient	2
τ_{SS}	Social Security tax rate on employees	3.8%

⁴ An alternative specification is to use the collective model and solve the Pareto-efficiency intro-household allocation along the line of, for example, Chiappori (1988, 1992, 1998).

We compute these from the PSID, should everyone in our sample choose to work. These lines display several interesting facts. First, for this cohort, the wages of women are significantly lower than those for men for their whole working lives. For instance, at age 25, men start their working life with an average potential wage of over \$14 an hour (which we express in 1998 dollars), while women's potential hourly wage is just above \$10. Thus, the gap between men and women at age 25 is \$4 an hour, that is, 40% of women's wages. Second, men's wages grow faster than women's until age 35, while women's wages grow faster after age 35–40. The result is that even at age 55–60, when women's potential wages are peaking, the potential wage gap across genders is over \$7, which is more than one-half of what an average woman would make per hour, should she choose to work. The profile generated ignoring gender differences is, unsurprisingly, almost in the middle of those for men and women and its shape is similar to both of the other profiles. The profile for the average potential wage of someone in a couple is a bit higher than the profile that ignores marital status, partly because single people have lower wages.

Turning to the stochastic processes for wages, Table 4 reports estimates for the stochastic component of earnings, which we assume to be an AR(1) (as specified in Eq. (3)) and we estimate from PSID data. It shows that women experience slightly smaller persistence in wages than men, 0.96 compared with 0.97, but similar variances. The third column reports data for the case in which we pool all data for men and women together and ignore gender differences (used in Economy 2), while the last column treats the household as a unit and computes average potential wages between members of the couple. The table shows that ignoring gender differences and putting men and women together when estimating the stochastic component for wages generates a persistence very similar to that of men, but a variance of the productivity shock that is twice as large as that for men and an initial variance that is 1.3 times larger than that of men. In contrast, for the average member of a couple, while the persistence is also very similar to that of men's wages, both the variance of the productivity shock and especially the initial variance are much lower than the corresponding statistics in the other samples.

Turning to the calibration of our mortality probabilities, Table 5 compares their implied life expectancies at age 70, 80, and 90 from the 2013 U.S. life table (source: Social Security Administration (SSA)) and our calculations from the HRS. The numbers refer to different populations, as the U.S. life table is a cross-sectional life table for the year 2013, while our calculations use the period 1998–2013, control for cohort effects, and refer to the 1941–45 cohort.

Despite these differences, our HRS sample generates very similar life expectancies for women and underestimates the life expectancy for men by only nine months at age 70. At more advanced ages, our estimates become even more precise. According to our HRS sample, there is a large gap in life expectancy at age 70 between singles and married individuals: married men expect to live 2.5 years longer than single men and married women expect to live 2.1 years longer than single women. These gaps shrink at older ages. The appendix reports details on how we use the PSID and the HRS data to compute these processes.

Table 4
Estimated processes for the wage shocks for men and women (PSID data).

Parameter	Men	Women	Men and Women together	Married households
Persistence	0.973	0.963	0.973	0.972
Variance prod. shock	0.016	0.014	0.021	0.013
Initial variance	0.128	0.122	0.163	0.078

Table 5

Life expectancy at ages 70, 80, and 90 in years. First column: U.S. life tables from Social Security Administration. Other columns: Our computations based on HRS data.

Gender	U.S. life tables	All	Single	Married
<i>At age 70</i>				
Women	16.4	16.4	15.4	17.5
Men	14.2	13.5	11.5	14.0
<i>At age 80</i>				
Women	9.6	9.7	9.5	10.3
Men	8.2	8.0	7.3	8.2
<i>At age 90</i>				
Women	4.8	4.6	4.6	4.8
Men	4.0	3.8	3.7	3.9

Results

We now turn to discussing the results for our four model economies and their aggregate implications.

Economy 1, a no-marriage economy calibrated to men

For the calibration of Economy 1, a no-marriage economy calibrated to men, we use the survival probabilities and the wage process for men that we just described in the calibration section. The top panel of Table 6 reports the parameters that we use to match our targets, that we list in the bottom part of the same table. The Social Security benefit is chosen to match the government budget constraint for this cohort. All of the parameters are consistent with those used in the literature. In particular, the value of labor participation cost is very close to that in French (2005).

The model's fit to the data for men

Fig. 4 displays labor supply participation, average hours, labor income, and savings over the life cycle in the actual data and the singles economy. The model matches well the main features of the data for men but tends to overpredict participation and hours after age 55, where it might be missing the role of health shocks and endogenous benefit and pension claiming. Overall, however, the model matches the most important features of the data for men over the life cycle well.


The model's implications compared with the aggregate data

To evaluate the extent to which a model calibrated to only men misrepresents key economic aggregates, we now turn to comparing the outcomes of this calibrated model to those in the aggregates in terms of participation, hours, earnings, and assets in the data.


Table 6

Economy 1, no-marriage economy calibrated to men. Top panel: Parameters. Bottom panel: Target moments. Data: our computations from PSID and HRS. The SS budget deficit is expressed as the ratio to SS budget for this cohort.


Parameters	Value	
β	Discount factor	0.957
ω	Consumption weight	0.510
$\phi_t^{t=1,j}$	Labor participation cost	0.283
$y_r^{t=1,s}$	Social Security benefit	\$8023
Moments	Data	Model
SS budget deficit	0.000	0.002
Average assets, single men at 50	148710	149017
Average hours, single men at 50	2129	2120
Participation, single men at 50	0.939	0.964




(a) Participation



(b) Hours worked

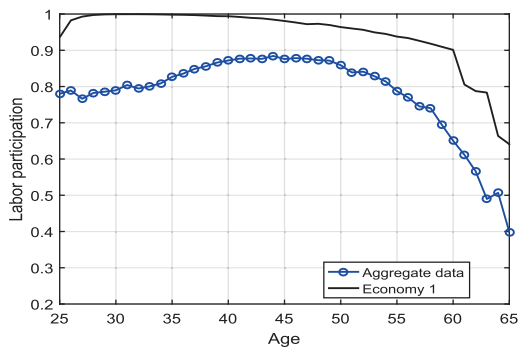


(c) Labor income




(d) Assets


Fig. 4. Economy 1, a no-marriage economy calibrated to men. Model fit to data for men.



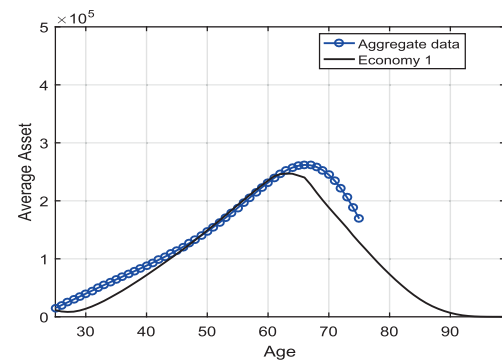
(a) Participation



(b) Hours worked



(c) Labor income



(d) Assets

Fig. 5. Economy 1, a no-marriage economy calibrated to men. Model fit to aggregate data.

Fig. 5 compares the aggregate life-cycle profile by age for the observed data and Economy 1. It shows that only modeling men misses important aspects of aggregate behavior over the life cycle. First, the economy with only men drastically overestimates participation by about 10 percentage points over the life cycle of this cohort. Second, it overestimates average hours over the life cycle by about 500 hours, which is almost one-third of actual aggregate hours in this cohort. Third, it overestimates average earnings by age. For instance, at age 45, average earnings are close to \$35,000 a year, while the singles economy predicts over \$45,000, thus over-predicting earnings by almost one-third. Finally, the model under-predicts average assets at retirement and predicts much faster asset decumulation after retirement than in the data, partly because the life expectancy of men is lower than that of women.

Table 7
Economy 2, no-marriage economy, men and women together. Top panel: Parameters. Bottom panel: Targeted moments. Data: our computations from PSID and HRS. The SS budget deficit is expressed as the ratio to SS budget for this cohort.

Parameters	Value	
β	Discount factor	0.958
ω	Consumption weight	0.471
$\phi_t^{i=1,j}$	Labor participation cost	0.302
$y_r^{i=1,s}$	Social Security benefit	\$5006
Moments	Data	Model
SS budget deficit	0.000	-0.001
Average assets, individuals at 50	147134	147530
Average hours, individuals at 50	1768	1758
Participation, individuals at 50	0.859	0.872

Economy 2, a no-marriage economy calibrated to men and women together

The goal of Economy 2 is to see whether we can use data for men and women together to match aggregate outcomes without having to write, solve, and calibrate more complicated models with gender and marriage. Thus, we put all men and women in the same sample, both married and single, and we estimate the survival probabilities and the wage process for this mix of people, treating them as individuals, and ignoring any gender and marital status differences.

The top panel of Table 7 reports the parameters that we use to match our targets, which we list in the bottom part of the table. The calibrated discount factor is very similar to that in the singles economy, because average assets at age 50 are very similar in these two economies. The weight on consumption is lower in Economy 2 than in Economy 1 because average working hours at age 50 are lower in Economy 2, as women tend to work less than men on average. The Social Security benefit chosen to match the government budget constraint for this cohort is lower in Economy 2 because total labor income that pays Social Security tax is also lower due to the inclusion of women in the sample. The participation cost is higher than in Economy 1 to match the lower labor participation rate.

The model's fit to the data for men and women pooled together, and thus the aggregates

Fig. 6 displays labor supply participation, average hours, labor income, and savings over the life cycle in the actual data and in Economy 2. The model has a harder time fitting the profiles of par-




Fig. 6. Economy 2, a no-marriage economy calibrated to men and women together. Model fit to targeted data (which are also the aggregate data).

ticipation and hours than Economy 1 does, because it does not explicitly take into account the differences between men and women. In particular, the model overpredicts participation and hours for both younger and older individuals and only does well on both dimensions during the prime age period (age 40 to 55) and largely overpredicts labor earnings.

Unlike in Economy 1, where we target data for males, for this economy, the targets that we are trying to match also coincide with the true aggregates in the economy and thus also speak to the success and failure of the model in that regard. Thus, calibrating the model to wages and life expectancy of men and women together, ignoring gender and marital status differences, yields a better match than just using data for men, as done in Economy 1. The biggest success of this calibration compared with Economy 1 is the match in labor earnings over the life cycle, which is very

close. Unfortunately, however, it still implies patterns of participation and hours over the life cycle that miss over most of the life cycle, with the exception of the period between ages 45 and 55. In addition, and as in Economy 1, Economy 2 underpredicts assets at retirement and generates faster asset decumulation after retirement than in the data.

In Economy 4, we allow the fixed cost of working for married women to be changing by age to take into account child rearing costs in a simple way. Thus, one might think that the two models are not on an even ground because Economy 4 has more parameters that allow for a better fit of the data. Thus, in Appendix D we report results for Economy 2 in the case of age-varying fixed costs of working and show that it fits the data very similarly to the model without an age-varying fixed cost of working (we call it Economy 2b).

Table 8

Economy 3, a no-marriage, household-level calibration for couples only. Top panel: Parameters. Bottom panel: Target moments. Data: our computations from PSID and HRS. The SS budget deficit is expressed as the ratio to SS budget for this cohort.


Parameters		Value
β	Discount factor	0.964
ω	Consumption weight	0.412
$\phi_t^{i=1,j}$	Labor participation cost	0.218
$\gamma_t^{i=1,s}$	Social Security benefit	\$5070
Moments		
	Data	Model
SS budget deficit	0.000	0.001
Average assets, households at 50	150030	149042
Average hours, households at 50	1780	1819
Participation, households at 50	0.982	0.910

Economy 3, a no-marriage, household-level calibration for couples only


Economy 3 is a “no-marriage, household-level calibration for couples only” economy that uses the same model as Economies 1 and 2, but in the calibration we aggregate the data at the household level and we only keep couples. More specifically, for members of couples we compute life expectancy regardless of gender and we compute the average potential wage of the two members of the household. This is in contrast with Economies 1 and 2, in which we were only looking at individual-level data, ignoring that some people are in couples and some others are not. In Appendix D we show the results for a household-level calibration that also includes singles (we call it Economy 3b); the results are very similar to those in Economy 3 that we report here.




(a) Participation



(b) Hours worked



(c) Labor income



(d) Assets

Fig. 7. Economy 3, a no-marriage economy calibrated to married people, at the household level. Model fit to household-level data for married people.

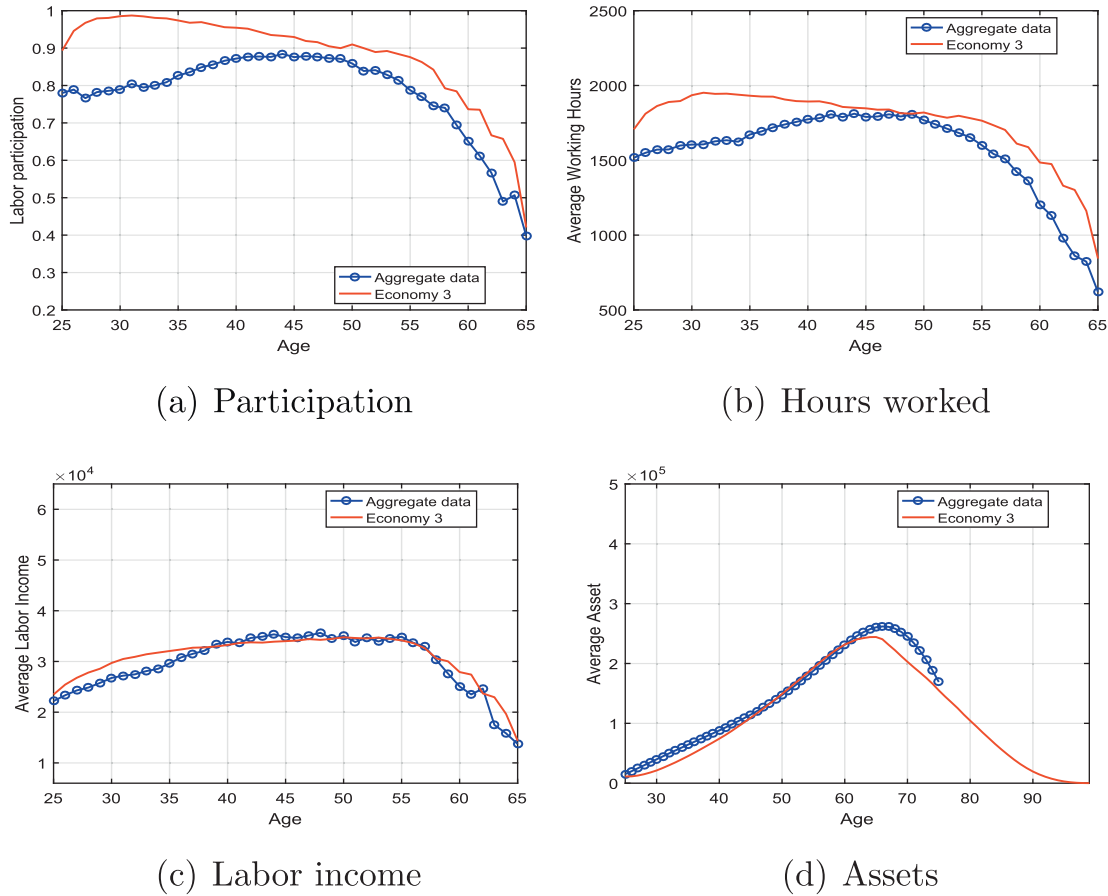


Fig. 8. Economy 3, a no-marriage economy calibrated to married people, at the household level. Model fit to aggregate data.

The top panel of Table 8 reports the parameters that we use to match our targets, which we list in the bottom part of the table. The calibrated discount factor is higher than in Economy 1 because average assets at age 50 for households are similar, but the wages are lower. The weight on consumption is lower than in Economy 1 because average working hours at the individual level in a household at age 50 are lower in Economy 3, because women tend to work less than men on average. The Social Security benefit chosen to match the government budget constraint for this cohort is lower than in Economy 1, because average labor income that pays Social Security tax is also lower at the household level than for males. The participation cost is lower than in Economy 1 to match a much higher labor participation rate at the household level (we count a household participating in the data if at least one of the people in the household participates in the labor market).

The model's fit to the data for married people, at the household level


Fig. 7 displays labor supply participation, average hours, labor income, and savings over the life cycle in the actual data for married people, at the household level, and in Economy 3. It is important that this is a household-level calibration. Thus, in the data, we count a married household as participating in the labor market if at least one of the two people in the household is working. As a result, the empirical profile for participation is very high and only starts declining at age 50. The model has a hard time matching the participation profile for households in couples over the life cycle and underestimates participation at all ages, except between age 30 and 35. In addition, the gap between the model-predicted participation and actual participation in the data keeps increasing with age starting at age 35. For the data, we compute hours worked as

Table 9


Economy 4, a marriage and singles economy. Top panel: Parameters. Bottom panel: Target moments for the marriage economy. The data moments come from our computations on PSID and HRS data. The SS budget deficit is expressed as the ratio to the SS budget for this cohort.

Parameters		Value
β	Discount factor	0.961
ω	Consumption weight	0.496
$\phi_t^{i=1,j}$	Men participation cost	0.316
$\phi_t^{i=2,j=1}$	Single women part. cost	0.376
$\phi_t^{i=2,j=2}$	Married women part. cost	See text.
$\gamma_r^{i=1,s}$	Single men SS benefit	\$6,731
Moments		
SS budget deficit	0.000	0.000
Avg. assets, single men at 50	133821	167966
Avg. assets, single women at 50	83156	82440
Avg. assets, couples at 50	291433	225041
Avg. hours, single men at 50	1869	1867
Avg. hours, single women at 50	1703	1689
Avg. hours, married men at 50	2165	2011
Avg. hours, married women at 50	1337	1542
Part., single men at 50	0.831	0.910
Part., single women at 50	0.875	0.891
Part., married women at 35	0.630	0.637
Part., married women at 45	0.776	0.694
Part., married women at 55	0.683	0.649


the average, per-person, in the couple. Comparing the model-generated data with the actual data for married households, the model cannot match hours over the life cycle, except for ages 45–50. One bright spot is that this model matches average labor income over the life cycle much better than Economy 1 and as well




(a) Participation



(a) Labor Income



(b) Hours



(b) Assets

Fig. 9. Economy 4, a marriage and singles economy. Data (Left panel) and model (Right panel). Model fit.

as Economy 2. As in the previously studied model economies, this one also understates retirement savings.

The models' implications compared with the aggregate data

Fig. 8 compares the aggregates generated over the life cycle by our economy with married people only to those in the aggregate economy, including everyone at every age. This model economy does much better than Economy 1 and as well as Economy 2 in terms of matching the aggregate data, and actually matches participation, hours worked, labor income, and assets reasonably well over the whole life cycle. In terms of discrepancies, Economy 3 tends to overpredict labor participation and hours outside of the 40–50 age group.

The marriage and singles, men and women economy

For the marriage economy, we use the model that explicitly models households with two decision makers and in the calibration stage we use the survival probabilities for single men, single women, and married men and women, and the wage processes for men and women that we have estimated and described in Section “Calibration”. Appendix D reports the results for an Economy

Fig. 10. Economy 4, a marriage and singles economy. Data (Left panel) and model (Right panel). Model fit.

with only married couples (Economy 4b), which are very similar to those that we report here.

The top panel of Table 9 reports the parameters that we choose for the marriage economy to match the corresponding target moments listed in the bottom panel of the Table.

The Social Security benefit for single men is chosen to match the government budget constraint for this cohort. We pin down the Social Security benefits for the other groups by fixing the ratios of benefits by group relative to single men to match the data reported in 1997. According to SSA, in 1997, median income for elderly unmarried women was 98.8% of that for unmarried men, and median income for elderly married couples was 183.0% that for unmarried men.⁵

⁵ We calculated the ratios using the 1997 SSA report, which reads “In 1997, median income for elderly unmarried women (widowed, divorced, separated, and never married) was \$11,161, compared with \$14,769 for elderly unmarried men and \$29,278 for elderly married couples. . . . Elderly unmarried women – including widows – get 51 percent of their total income from Social Security. Unmarried elderly men get 39 percent, while elderly married couples get 36 percent of their income from Social Security.” Thus, we set SS benefits to \$5,692 for single women, to \$5,760 for single men, and \$10,540 for couples (<https://www.ssa.gov/history/reports/women.html>).

We use participation costs to match the labor participation rate and hours worked at age 50 for a given group. For single men, married men, and single women, we assume that these costs are constant over the life cycle. For married women, to capture the role of child rearing in a simple way, we assume that married women face a time-varying participation cost over their life cycle. The corresponding coefficients for $\phi_t^{i=2,j=2}$ are 0.0009, -0.0233 , and -1.0297 . The additional targets for married women are average hours and participation at ages 35, 45, and 55.

Fig. 14 in Appendix C reports the age-varying costs of working for each group in this economy. As it turns out, the participation cost for married women is relatively high during the peak child-bearing years, but then decreases to reconcile their labor income patterns, due to their lower wage and higher family assets. The participation cost for single women turns out to be higher than that of both married women and men to match their low participation rate. In reality, they might be stuck in low-paying jobs expecting to get married. Thus, the participation cost is a stand in for more than just commuting costs.

The model's fit to the data for married and single men and women

The top panel of Fig. 9 displays the life cycle profiles of labor participation rate by gender in the data (left panel) and the model (right panel). Here, too, the model captures the main aspects of the data. Single women are more likely to participate in the labor market than married women, despite the fact that single women face a larger participation cost. Married women have a lower labor participation rate than single women because of division of labor

among couples in the presence of a fixed cost of working and lower wages for women than men. Single men participate less than married men in the labor market when older.

The bottom panel of Fig. 9 displays average hours worked. As in the data, the model delivers a large difference in hours worked by gender among singles: Single women work less on average than single men because single women face a larger participation cost. As in the data, married men work much more than married women because married couples care about total household income and maximize household utility; and because women on average earn lower wages than men, they enjoy more leisure.

The bottom panel of Fig. 10 displays life-cycle profiles of average household assets by gender and marital status. Before age 70, assets among single women are lower than assets among single men because single women have lower earnings capacity. However, women hold more assets than men late in life because women live longer and thus need to shift more resources to the future to smooth consumption. Average assets among married couples are much higher than for singles but still fall short of what we observe in the data. This is because the model struggles to match savings for our three age groups; in particular, it overestimates savings for single men while underestimating savings for couples. This could possibly reflect some preference heterogeneity across groups that we are abstracting from.

The models' implications compared with the aggregates in the data

Fig. 11 compares the aggregates generated over the life cycle by our economy with married people and single people to those in the

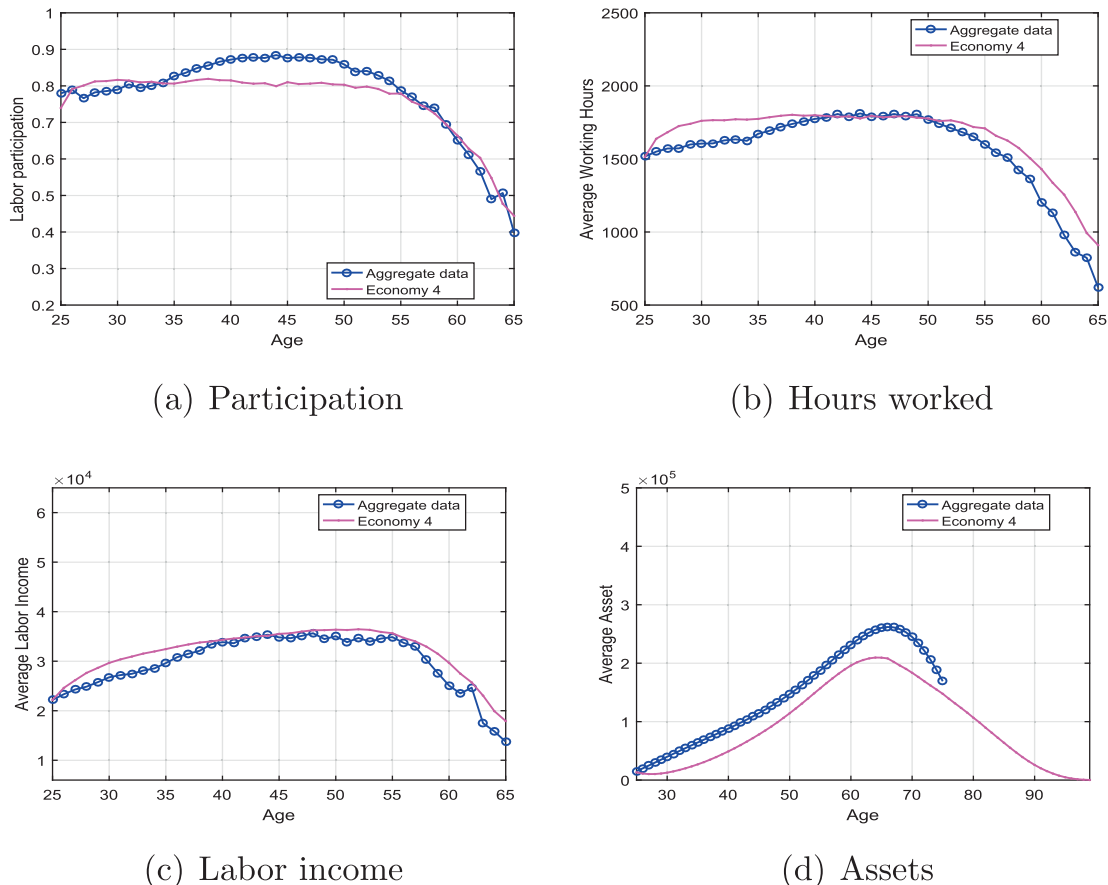


Fig. 11. Economy 4, a marriage and singles economy. Data and model. Model and aggregate data.

aggregate economy, including everyone at every age. This model economy does much better than Economies 1, 2, and 3 in terms of matching the aggregate data; and actually matches participation, hours worked, labor income, and assets well over the whole life cycle. In terms of discrepancies, Economy 4 also tends to underpredict labor participation between ages 40 and 50 (by less than 10 percentage points) and tends to overpredict hours and labor income a bit over all of the life cycle. However, compared with the economies without two-agent couples, these misses between model and data are minor.

Conclusions

We now turn to summarizing our results for the aggregates in the data and for the four economies that we have constructed. We do so by aggregating the age groups of our cohort, which is equivalent to considering the economy in a steady state with no population growth.

Table 10 reports the average participation, hours, and earnings in the data and in Economy 1–4, over the age range 25–65. In the top panel, the first row reports the fraction of people aged 25–65 participating in the labor market. A comparison of our data and model economies shows that average participation over the life cycle is 78% for our cohort and 94%, 86%, and 76% respectively in Economies 1–4. The second row in the table shows that average yearly working hours are 1,562 in our data, and 2,050, 1,725, 1,756, and 1,656 respectively in Economies 1–4. Lastly, the third line shows that average earnings over the working period are \$29,628 in our data, and 42,566, 30,392, 30,702, and 31,652 respectively in Economies 1–4.

The bottom panel of the table summarizes the success and misses of the model economies. Economy 1 overpredicts average labor participation over the life cycle by 16 percentage points, overpredicts hours by 31%, and overpredicts earnings by 44%. Economies 2 and 3 respectively overpredict participation by 12 and 10 percentage points and underpredict earnings by only 2.6 and 3.6%, thus coming somewhat closer to the data than Economy 1. Finally, Economy 4 does better than all of the economies that we consider by underpredicting participation by only 2% and overpredicting hours by 6% and earnings by 7%.

We thus conclude that even macroeconomists not interested in heterogeneity in marriage and gender per se should start taking marriage and gender differences more seriously in the context of quantitative structural models. In particular, modeling gender and marriage explicitly would yield the best results in terms of matching the aggregates, but if this is not possible given the question and complications at hand, calibrating (or estimating) the model including both men and women in the data or keeping track of households to determine household-level wages and hours (or earnings), participation, and assets will help the model better match the aggregates.

Table 10

Average participation, hours, and earnings over the ages of 25 to 65 in the data and in the model economies. Data and misses.

Variable	Data	Economy 1	Economy 2	Economy 3	Economy 4
Participation	0.78	0.94	0.86	0.88	0.76
Hours	1,562	2,050	1,725	1,756	1,656
Earnings	29,628	42,566	30,392	30,702	31,653
Differences Participation	–	+0.16	+0.12	+0.10	–0.02
Hours	–	+31.2%	+10.4%	+12.4%	+6.0%
Earnings	–	+43.7%	+2.6%	+3.6	+6.8%

Acknowledgement

De Nardi gratefully acknowledges support from the ERC, Europe, Grant 614328 “Savings and Risks”. Yang gratefully acknowledges the hospitality of the Federal Reserve Bank of Chicago. We thank Marco Bassetto, Juan Carlos Conesa, Helen Koshy, Rory McGee, Alfonso Sousa-Poza, seminar participants at various institutions, and an anonymous referee for useful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, the CEPR, any agency of the federal government, the Federal Reserve Bank of Chicago, or the IFS.

Appendix A. The data

In this appendix we describe the data sets and techniques used to compute the inputs for the models and the data outputs that the model is calibrated to match.

The PSID

The Panel Study of Income Dynamics (PSID) is a longitudinal study of a representative sample of the U.S. population. In the first year of the survey, 1968, about 5000 families were first interviewed, with information gathered on these families and all of their descendants from that time onwards. Individuals are followed over time to maintain a representative sample of families. To accomplish this, the PSID sample persons include all persons living in the PSID families in 1968 plus anyone subsequently born to or adopted by a sample person. All sample members are followed even when leaving to establish separate family units. PSID families also include many non-sample persons, typically individuals who married sample persons after the beginning of the study in 1968. Information on non-sample persons such as spouses is collected while they are living in the same household as an individual in the original sample. However, once they stop living with a sample person, they are not followed further.

The original 1968 PSID sample was drawn from an over-sample of 1,872 low-income families from the Survey of Economic Opportunity (the SEO sample) and from a nationally representative sample of 2,930 families designed by the Survey Research Center at the University of Michigan (the SRC sample).

Sample selection

We are interested in the cohort born in 1941–45. To gather relevant information for this cohort, we select individuals born before or after the period 1941–45, control for cohort effects, and use results relative to the cohort of interest. More specifically, we select all individuals in the SRC sample who are interviewed at least twice in the sample years 1968–2013, select only heads and their wives, if present, and keep individuals born between 1931 and 1955. Thus, as we consider five-year-of-birth cohorts, we are

Table 11
Sample Selection in the PSID.

Selection	Individuals	Observations
Initial sample (observed at least twice)	30,587	893,420
Heads and wives (if present)	18,304	247,203
Born between 1931 and 1955	5,137	105,381
Age between 20 and 70	5,129	103,420

including two younger (1931–35 and 1936–40) and two older cohorts (1946–50 and 1951–55) than our cohort of interest. The resulting sample, used to estimate potential wage, includes 5,129 individuals aged 20 to 70, for a total of 103,420 observations (Table 11). When estimating the life-cycle wage rate profiles, however, to minimize the impact of cohort effects, we restrict the sample to 2,795 individuals born between 1936 and 1951, for a total of 59,885 observations, as we detail below.

The wage rate is defined as annual earnings divided by annual hours worked. Gross annual earnings are defined as previous year's income from labor, while annual hours are previous year's annual hours spent working for pay.

Assets are only observed in 1984, 1989, 1994, and 1999 and then every two years until 2013. We use total assets defined as the sum of all assets types available in the PSID, net of debt and plus the value of home equity. All monetary values are expressed in 1998 prices.

Wage profiles

We need an estimate of the potential wage profiles and wage processes by age and gender for all individuals, including those non-participating in the labor market. We first use an imputation procedure to estimate potential wage for non-participants. We then estimate the average wage profile by age and the persistence and variance of its unobserved component. In this second set of estimates, we either pool men and women together or consider them separately, depending on the economy we are calibrating.

Imputation. In the data, wages may be missing for two reasons: (1) An individual may have been active in the labor market but earnings or hours (or both) are missing; (2) An individual has not been active in the labor market.

Because estimated variances are very sensitive to outliers, we set to missing observations with an hourly wage rate below half the minimum wage and above \$250 (in 1998 values). We then impute all the missing values with the same procedure.

In our sample, out of a total of 103,420 observations, we observe the wage rate for 79,956 observations and we impute it for the remaining 23,464. Of those missing, 23% refer to individuals who were active in the labor market but have earnings and/or hours information missing. The remaining 77% of the missing refer to individuals who were not active in the labor market in a particular year.

Fig. 12 reports the fraction of people in a given age and demographic group for whom the actual wage is observed. This figure displays a pattern very similar to the participation profile. Thus, married women, being the group with the lowest participation rate, is also the group for which the wage imputation procedure has to be applied the most.

We impute wage values using coefficients from OLS regressions run separately for men and women and for single and married individuals (four separate groups). To avoid endpoint problems with the polynomials in age, we include individuals aged 20–70 in the sample. The dependent variable in our regressions is the logarithm of the hourly wage rate. As explanatory variables we include: a fourth-order polynomial in age, cohort dummies, time dummies (adding up to zero and orthogonal to a time trend), fam-




Fig. 12. Fraction of individuals for whom actual wage is observed.

ily size (dummies for each value), education (dummies for three levels, less than high school, high school, and college), number of children (dummies for each value), age of youngest child, and an indicator of partner working if married. As an indicator of health, we use a variable recording whether bad health limits the capacity of working, as this is the only health indicator available for all years (self-reported health starts in 1984 and is not asked before). However, this health indicator is not collected for wives, so we do not include it in the regression for married women. All regressions also include interaction terms between the explanatory variables. Using the estimated coefficients, the predicted value of the (logarithm of the) wage is taken as a measure of the potential wage for observations with a missing wage. When the wage is observed, we use the actual wage.

Alternatively, we could estimate fixed-effect regressions separately for men and women and include marital status as an explanatory variable because it is not constant over time. From these estimates, we would obtain an estimate of the fixed effect for each individual and use it along with the estimated coefficients to impute missing values. With this procedure, however, we would obtain an estimate of the fixed effect only for individuals with at least two non-missing wages, while with OLS we can estimate the potential wage for all individuals in our sample.

Estimation. We use the potential wage to estimate the deterministic age-efficiency profile e_t^i as a function of age. The sample is restricted to individuals born between 1936 and 1951 to account for changes in the age-efficiency profile. After this selection, there are 2,795 individuals born between 1936 and 1951, for a total of 59,885 observations. To estimate the wage profile, we run a fixed-effect regression of the logarithm of the wage rate on a fourth-order age polynomial, separately for men and women or pooling them together, depending on the economy we are calibrating. We then regress the residuals on time and cohort dummies to compute the average effect for the cohort born in 1941–45 and use that estimate to fix the constant of the wage profile. The estimated potential wage profile for all cases is shown in the main text in Fig. 3.

To show the effect of our imputation procedure we plot in Fig. 13, for men and women separately, both the estimated potential wage profile, which includes our imputation for missing values, and the estimated observed wage profile, which includes only observations for which the wage rate is actually observed. As shown in the picture, indeed the potential wage is lower than the observed one for men, while for women the difference is evident only after age 55. This is consistent with people with lower productivity choosing not to participate in the labor market and thus potential wage being lower than the actual observed wage.

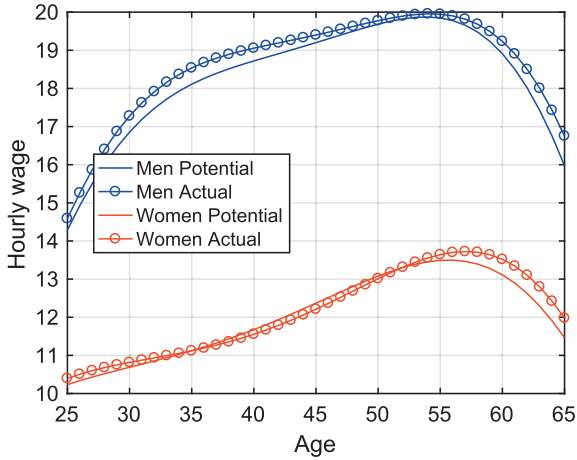


Fig. 13. Actual and Potential wages over the life cycle for men and women born in the 1941–45 cohort (PSID data).

To estimate the variance structure of the log wage rate, we use residuals of the log wage rate computed with the same procedure just described; that is, we first run a fixed-effects regression and then we regress the residuals on time and cohort dummies.⁶ In this case, to have enough observations to estimate the variances, we include all cohorts from 1931 to 1955. When computing the variances, we limit the age range between 25 and 55 and we use waves only up to 1997, the last wave in which yearly data are available. As we rely on residuals also taken from imputed wages, we drop the highest 0.5% residuals both for men and women, in order to avoid large outliers to inflate the estimated variances (however, the effect of this drop is negligible on our estimates).

The shock in log wage is modeled as the sum of a persistent component plus white noise, which captures measurement error:

$$\ln u_{t+1}^i = \ln \epsilon_{t+1}^i + \zeta_{t+1}^i$$

$$\ln \epsilon_{t+1}^i = \rho^i \ln \epsilon_t^i + v_{t+1}^i,$$

where i indicates men and women, and ζ_{t+1}^i and v_{t+1}^i are white-noise processes with zero mean and variances equal to σ_{ζ}^2 and σ_{iv}^2 , respectively. This last variance together with the persistent parameter ρ^i characterize the AR process in the model. Estimation is carried out using minimum distance techniques, standard in the literature of earnings dynamics. Results for the various cases (men, women, men and women together, and married households) are shown in the main text in Table 4.

Initial assets and wage shocks

We parameterize the joint distribution of initial assets and wage shocks at age 25 as a joint log normal in the logarithm of assets and the wage shock. We set a shift parameter δ_a for assets to have only positive values and to be able to take logs. The wage shock is the residual from the wage regressions described above, re-scaled in order to remove the white-noise component.

$$\begin{pmatrix} \ln(a + \delta_a) \\ \ln \epsilon \end{pmatrix} \sim N \begin{pmatrix} \mu_a + \delta_a \\ \mu_\epsilon \end{pmatrix}, \Sigma_s,$$

⁶ In the second step, we add time dummies (as in Deaton and Angus, 1994), constrained to add up to zero and orthogonal to a time trend. Adding this sort of time dummy is not strictly necessary, as those business cycle shocks would be captured by the white-noise component in the estimation of the variance of the unobserved component of wage.

where Σ_s is a 2x2 covariance matrix. This is done separately for all men, or all individuals, to calibrate the economies where we do not condition on marital status, or for single men and single women (in Economy 4, where we condition on marital status).

For couples, in Economy 4, we compute:

$$\begin{pmatrix} \ln(a + \delta_a) \\ \ln(\epsilon_h) \\ \ln(\epsilon_w) \end{pmatrix} \sim N \begin{pmatrix} \mu_a + \delta_a \\ \mu_{eh} \\ \mu_{ew} \end{pmatrix}, \Sigma_c,$$

where Σ_c is a 3x3 covariance matrix computed on the data for married or cohabitating couples.

The HRS

We use the Health and Retirement Study (HRS) to compute inputs for the retirement period, because this data set contains a large number of observations and high-quality data for this stage of the life cycle. In fact, the HRS is a longitudinal data set collecting information on people aged 50 or older, including a wide range of demographic, economic, and social characteristics, as well as physical and mental health, and cognitive functioning.

The HRS started collecting information in 1992 on individuals born between 1931 and 1941, the so-called initial HRS cohort, which was then re-interviewed every two years. Other cohorts were introduced over the years, the AHEAD (Assets and Health Dynamics Among the Oldest Old) cohort, born before 1924, was first interviewed in 1993, while the CODA (Children of the Depression) cohort, born 1924 to 1930, and the WB (War-Baby) cohort, which includes individuals born 1942 to 1947, were introduced in 1998 and subsequently interviewed every two years. Younger cohorts, the EBB (Early Baby Boomer), born 1948 to 1953, and the MBB (Mid Baby Boomer), born 1954 to 1959, were first interviewed in 2004 and 2010 respectively.

To estimate input probabilities and outputs to be matched, we need information for individuals born during the 1941–45 period. To this end, we include all individuals of all cohorts starting from wave 3 (year 1996). We estimate the relevant probabilities using all individuals born between 1900 and 1945, controlling for cohort effects and using coefficients relative to the cohort born in 1941–45.

Our data set is based on the RAND HRS files and the EXIT files to include information on the wave right after death. Our sample selection is as follows. Of the 37,317 individuals initially present, we drop individuals for whom marital status is not observed (2,275 individuals) because marital status is crucial information in our analysis. This sample consists of 35,042 individuals and 176,698 observations. We then select individuals in the age range 66–100 born in 1900 to 1945, obtaining a sample of 15,072 individuals and 67,744 observations.

Survival probabilities

We model the probability of being alive at time t as a logit function:

$$\pi_{st} = \text{Prob}(S_t = 1/X_t^s) = \frac{\exp(X_t^s)}{1 + \exp(X_t^s)}.$$

We include as explanatory variables a fourth-order polynomial in age, gender, and interactions between these variables. When we calibrate economies that ignore gender differences, we do not control for gender. For Economy 4, when we consider marriage, we also include marital status and interactions of the previous variables with marital status. We also include cohort dummies and use coefficients relative to the 1941–45 cohort in the model. As

the HRS is collected every two years, we transform the biennial probability of surviving into an annual probability by taking the square root of the biennial probability. Having estimated survival probabilities in the age range 66–100, we apply standard formulas to compute life expectancy at various ages.

Appendix B. Sketch of the solution algorithm

The state space for average lifetime earnings and asset holdings is discretized into unevenly spaced grids. We use a regular grid search to solve the model.

In our full model, we need to compute two sets of value functions by age. The value function of being single and the value function of the two people in the married couple.

The value function of the couples depends on their own future continuation value and the one of the singles, in case of death of a spouse. Thus, this is how the solution is computed during retirement:

1. Compute the value function of the retired single person for all time periods after retirement, doing the usual backward iteration starting from the last period.
2. Compute the value function of the retired couple for all time periods after retirement, which uses the previous one in case of death of one of the spouses, doing the usual backward iteration starting from the last period.

For each period, working backwards over the life cycle, we apply the following solution strategy:

1. For any given time period, compute the value function of being single.
2. Given the value function of being single, compute the value function of the couple.

After we solve the value functions and policy functions, we simulate as follows. We take initial distributions of people by marital status and wages from the PSID data. We then use the model to simulate our cohort until the end of its life cycle.

Appendix C. Additional calibration material

Fig. 14 reports the life-cycle labor participation costs that we use for each group in Economy 4.

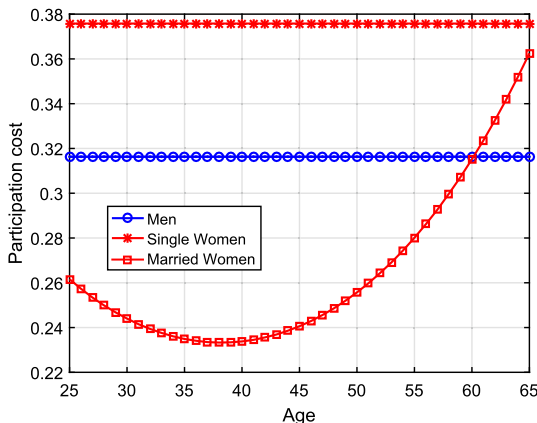


Fig. 14. Estimated lifecycle labor participation cost in time in the marriage and singles, men and women economy.

Appendix D. Three more economies

Economy 2b, a no-marriage economy calibrated to men and women together with an age-varying fixed cost of working

In Economy 4, we allow the fixed cost of working for married women to be changing by age to take into account child-rearing costs in a simple way. Thus, one might think that the two models are not on an even ground because Economy 4 has more parameters that allow for a better fit of the data. Thus, in Fig. 15, we also allow for an age-varying fixed cost of working in Economy 2. We use the following function, which is parameterized by three parameters, chosen to match participation at ages 35, 45, and 55, as done in Economy 4, $\phi_t^{i=2,j=2} = \frac{\exp(at^2-bt-c)}{1+\exp(at^2-bt-c)}$. As we can see here, even when we allow for more parameters for Economy 2 to match the data, the fit of labor participation only marginally improves between the ages of 45 and 55, and the economy still has similar misses as Economy 1 without the age-varying participation cost. Table 12 reports the calibration table for Economy 2 in the case of age-varying fixed costs of working.

Economy 3b, a no-marriage, household-level calibration for couples and singles

Economy 3b is a “no-marriage, household-level calibration” economy that uses the same model as Economies 1, 2, and 3, but in the calibration we aggregate the data at the household level, for both singles and couples. Compared with Economy 3, it also contains single households.

The top panel of Table 13 reports the parameters that we use to match our targets, which we list in the bottom part of the table.

The model's fit to the data for married people, at the household level

Fig. 16 displays labor supply participation, average hours, labor income, and savings over the life cycle in the actual data for married people, at the household level, and in Economy 3b. Consistently with the observation that there are very few people, the results for this model economy are very similar to those in Economy 3.

The models' implications compared with the aggregates in the data

Fig. 17 compares the aggregates generated over the life cycle by our economy with married people only to those in the aggregate economy, including everyone at every age. Here, too, the results for this model economy are very similar to those in Economy 3.

Economy 4b, Marriage only, men and women

Economy 4b is a slightly simpler case of Economy 4 that ignores the working singles. It is a “marriage only, men and women” economy in which everyone is married, and married men and women are calibrated using the corresponding observed data for their own group. Hence, there are no single people in this economy during the working period (retired people might become single because of spousal death), but gender and marriage are modeled explicitly. Since the model largely abstracts from singles, it misses them from the data, but it is bit simpler than the next model that we consider. As we will see, since there are few single people, the results and conclusions are very similar to those of Economy 4.

For this economy, we use the survival probabilities for married men and women, and the wage processes for men and women that we have estimated and described in Section “Calibration”. The difference from Economy 4, is that there are no single people of working age in in this economy. Since people in a couple can die after age 65, we have some people turning single after that time, due

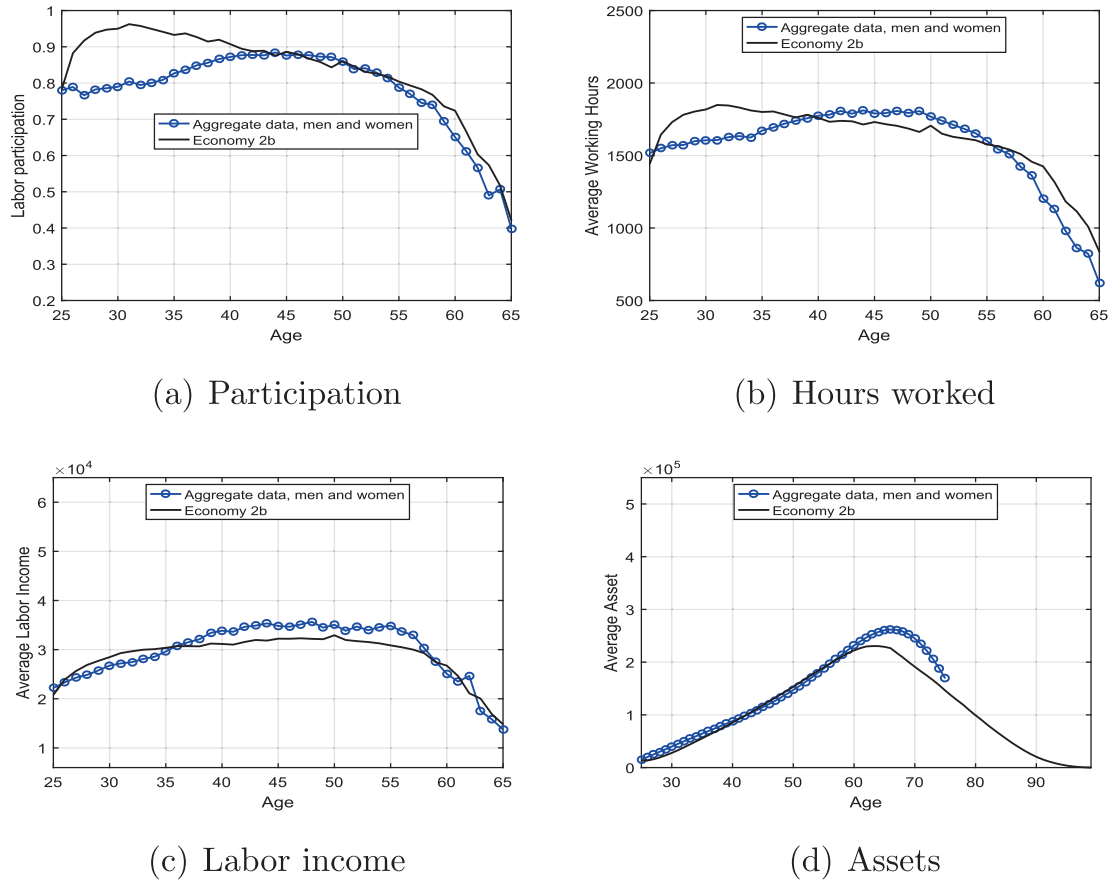


Fig. 15. Economy 2b, a no-marriage economy calibrated to men and women together, with an age-varying cost of working.

Table 12

Economy 2b, no-marriage economy, men and women together, age-varying fixed cost of working. Top panel: Parameters. Bottom panel: Target moments. Data: our computations from PSID and HRS. The SS budget deficit is expressed as the ratio to SS budget for this cohort.

Parameters		Value
β	Discount factor	0.961
ω	Consumption weight	0.484
$Y_r^{i=1,s}$	Social Security benefit	\$4734
$\phi_t^{i,j}$	Labor participation cost (2nd)	0.000
$\phi_t^{i,j}$	Labor participation cost (1st)	-0.008
$\phi_t^{i,j}$	Labor participation cost (constant)	-0.493
SS budget deficit		0.000
Average assets, men and women at 50	147134	153201
Average hours, men and women at 50	1768	1706
Participation, men and women at 35	0.827	0.933
Participation, men and women at 45	0.876	0.886
Participation, men and women at 55	0.787	0.804

to spousal death. We ignore all single people for calibration purposes because this economy focuses on understanding married couples.


The top panel of Table 14 reports the parameters that we choose for Economy 3 to match the corresponding target moments that are listed in the bottom panel of the table. The Social Security benefit for couples is chosen to match the government budget constraint for this cohort, which is very close to the number reported by SSA. According to SSA, in 1997, median income for elderly married couples was \$10,540.⁷

⁷ See footnote 5 for detailed explanation.


Table 13

Economy 3b, a no-marriage, household-level calibration for couples and singles. Top panel: Parameters. Bottom panel: Target moments. Data: our computations from PSID and HRS. The SS budget deficit is expressed as the ratio to SS budget for this cohort.


Parameters		Value
β	Discount factor	0.960
ω	Consumption weight	0.431
$\phi_t^{i=1,j}$	Labor participation cost	0.247
$Y_r^{i=1,s}$	Social Security benefit	\$5019
SS budget deficit		0.000
Average assets, households at 50	134752	134071
Average hours, households at 50	1774	1814
Participation, households at 50	0.947	0.891



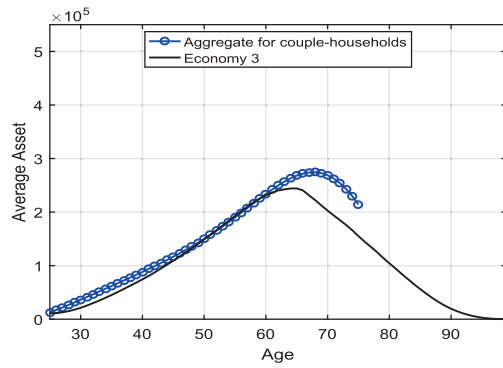
(a) Participation



(b) Hours worked




(c) Labor income




(d) Assets


Fig. 16. Economy 3b, a no-marriage economy calibrated to married and single people at the household level. Model fit to household-level data.



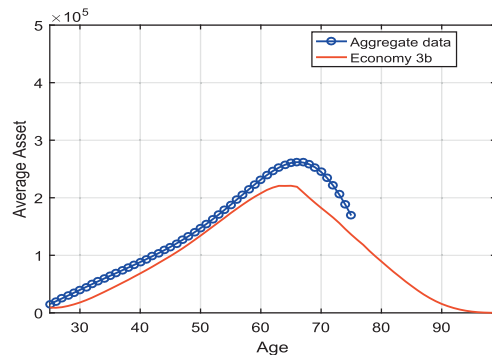
(a) Participation



(b) Hours worked



(c) Labor income



(d) Assets

Fig. 17. Economy 3b, a no-marriage economy calibrated to married and single people at the household level. Model fit to aggregate data.

Table 14

Top panel: Parameters used in the marriage only, men and women economy. Bottom panel: Target moments for the marriage only, men and women economy. The data moments come from our computations on PSID and HRS data. The SS budget deficit is expressed as the ratio to the SS budget for this cohort.

Parameters	Value	
β	Discount factor	0.965
ω	Consumption weight	0.492
$\phi_t^{i=1,j}$	Married men participation cost	0.286
$\phi_t^{i=2,j=2}$	Married women part. cost	See text.
$\gamma_r^{i=1,s}$	Couples SS benefit	\$11,534
Moments	Data	Model
SS budget deficit	0.000	-0.018
Avg. assets, couples at 50	291433	288987
Avg. hours, married men at 50	2165	2138
Avg. hours, married women at 50	1337	1454
Part., married men at 50	0.954	0.914
Part., married women at 35	0.630	0.635
Part., married women at 45	0.776	0.683
Part., married women at 55	0.683	0.608

We use participation costs to match the labor participation rate and hours worked at age 50 for married men and women. For married men, we assume that these costs are constant over the life cycle. For married women, to capture the role of child rearing in a simple way, we assume that married women face a time-varying participation cost over their life cycle, which is a quadratic function of age $\phi_t^{i=2,j=2} = \frac{\exp(0.0009t^2 - 0.0236t - 1.0708)}{1 + \exp(0.0009t^2 - 0.0236t - 1.0708)}$. The additional targets for married women are average hours and participation at ages 35, 45, and 55. Fig. 14 in Appendix C displays the life-cycle labor participation costs that we use for each group. As it turns out, the participation cost for married women is relatively high during the peak childbearing years, but then decreases to reconcile their labor income patterns, due to their lower wage and higher family assets.

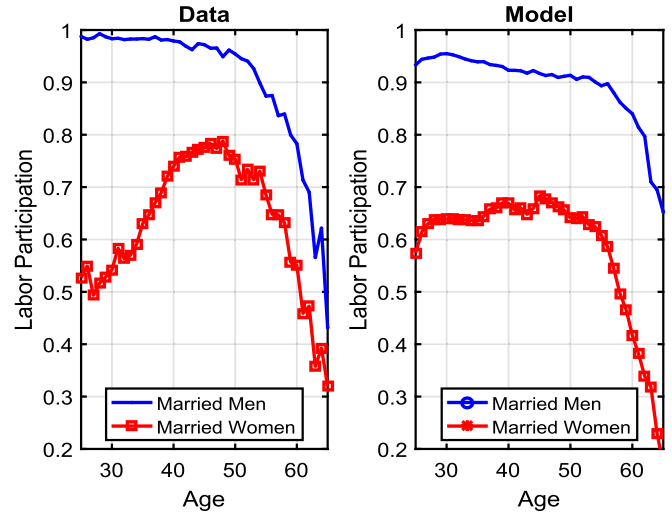
The model's fit to the data for married men and women, excluding singles

The top panel of Fig. 18 displays the life-cycle profiles of labor participation rate by gender in the data (left panel) and the model (right panel). As in the data, the model delivers a large difference in hours worked by gender. As in the data, married men participate much more than married women because women on average earn lower wages than men.


The second panel of Fig. 18 displays average hours worked. As in the data, the model delivers a large difference in hours worked by gender. As in the data, married men work more than married women because married couples care about total household income, maximize household utility, and women on average earn lower wages than men. As a result, in our model they enjoy more leisure. The third panel displays labor income. The model matches the main data patterns by gender well but overestimates labor income of married women.

The models' implications compared with the aggregates in the data

Fig. 19 compares the aggregates generated over the life cycle by our economy with married people only to those in the aggregate economy, including everyone at every age. This model economy does much better than Economies 1, 2, and 3 in terms of matching the aggregate data, and actually matches participation,



(a) Participation



(b) Hours

Fig. 18. Labor participation and hours in the data (Left panel) and in the marriage only, men and women economy (Right panel).

hours worked, labor income, and assets well over the whole life cycle. In terms of discrepancies, Economy 4b tends to underpredict labor participation between ages 40 and 50 (by less than 10 percentage points) and tends to overpredict labor income a bit over all of the life cycle. However, compared with the economies without two-agent couples, these misses between model and data are minor.

Thus, as with Economy 4, Economy 4b (marriage, married men and women only) does much better than Economies 1 and 2 in terms of matching the aggregate data, and actually matches participation, hours worked, labor income, and assets well over the whole life cycle. The fit of Economy 4b is close to that of Economy 4, as the fraction of single people by age is small.

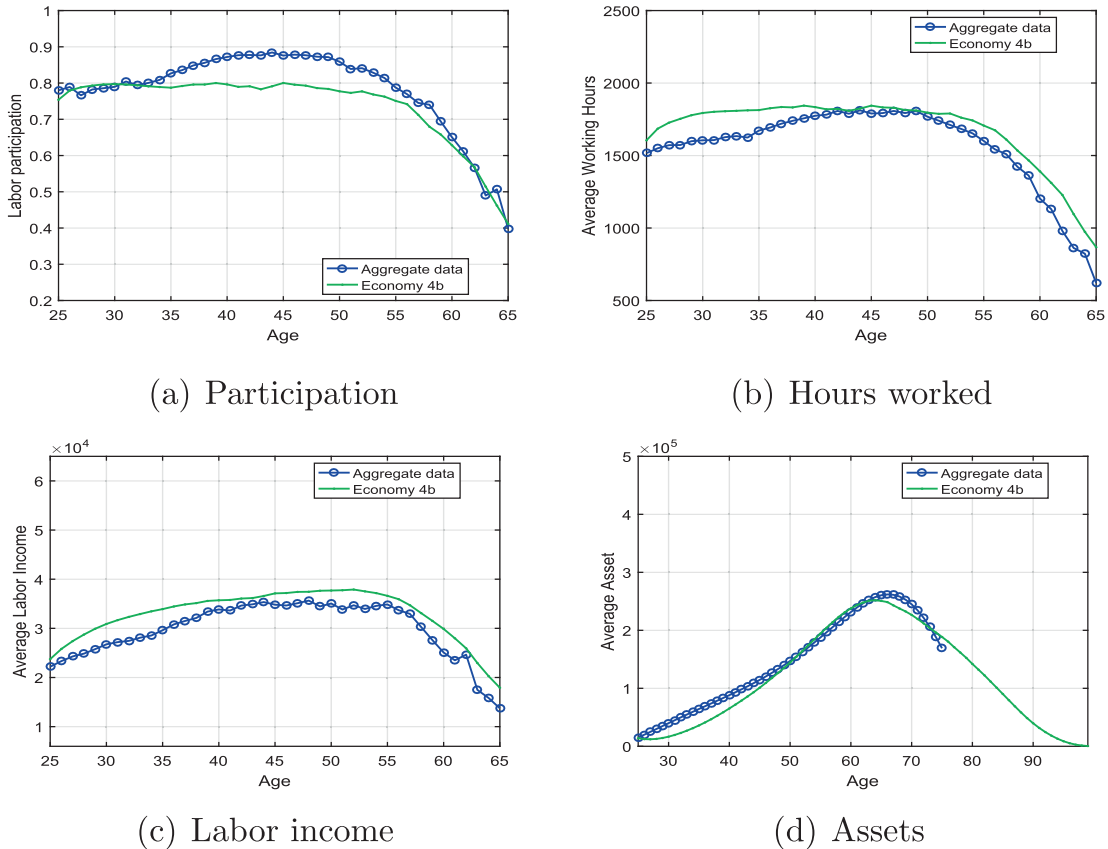


Fig. 19. Life-cycle profiles for the aggregate economy compared with those from Economy 4b, a marriage only, men and women economy.

References

- Attanasio, Orazio, Low, Hamish, Sánchez-Marcos, Virginia, 2008. Explaining changes in female labor supply in a life-cycle model. *Am. Econ. Rev.* 98 (4), 1517–1552.
- Blau, David M., 1998. Labor force dynamics of older married couples. *J. Labor Econ.* 16 (3), 595–629.
- Blau, David M., Gilleskie, Donna, 2006. Health insurance and retirement of married couples. *J. Appl. Econometrics* 21 (7), 935–953.
- Blundell, Richard, Dias, Monica, Costa, Meghir, Costas, Shaw, Jonathan, 2016. Female labour supply, human capital and welfare reform. *IFS Econometrica*, *Econometric Society*, vol. 84, pp. 1705–1753.
- Browning, Martin, Chiappori, Pierre-André, 1998. Efficient intrahousehold allocation: a characterisation and tests. *Econometrica* 66 (6), 1241–1278.
- Casanova, Maria, 2012. Happy together: a structural model of couples' joint retirement choices. Mimeo.
- Chiappori, Pierre-André, 1988. Rational household labor supply. *Econometrica* 56, 63–89.
- Chiappori, Pierre-André, 1992. Collective labor supply and welfare. *J. Political Econ.* 100 (3), 437–467.
- Deaton, Angus, Paxson, Christina, 1994. Saving, growth, and aging in Taiwan. NBER chapters. In: *Studies in the Economics of Aging*. National Bureau of Economic Research Inc, pp. 331–362.
- Eckstein, Zvi, Lifshitz, Osnat, 2011. Dynamic female labor supply. *Econometrica* 79 (6), 1675–1726. November.
- Eckstein, Zvi, Keane, Michael, Lifshitz, Osnat, 2016. Sources of change in the life-cycle decisions of American men and women. 1962–2014. Mimeo.
- French, Eric, 2005. The effects of health, wealth, and wages on labor supply and retirement behavior. *Rev. Econ. Stud.* 72 (2), 395–427.
- Guner, Nezih, Kaygusuz, Remzi, Ventura, Gustavo, 2012. Taxation and household labour supply. *Rev. Econ. Stud.* 79, 1113–1149.
- Remzi, Kaygusuz, 2015. Social security and two-earner households. *J. Econ. Dyn. Control* 59, 163–178.
- Low, Hamish, Meghir, Costas, Pistaferri, Luigi, Voena, Alessandra, 2016. Marriage, Social Insurance and Labour Supply. Mimeo.
- Nishiyama, Shinichi, 2015. The Joint Labor Supply Decision of Married Couples and the Social Security Pension System. Lancaster University, Mimeo.
- Tertilt, Michele, 2005. Fertility, and Savings. *J. Political Econ.* 113 (6), 1341–1371.
- van der Klaauw, Wilbert, Wolpin, Kenneth I., 2008. Social security and the retirement and savings behavior of low-income households. *J. Econometrics* 145, 21–42.