The Journal of the Economics of Ageing xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

The Journal of the Economics of Ageing

journal homepage: www.elsevier.com/locate/jeoa

Full length article

The aggregate implications of gender and marriage

Margherita Borella [a], Mariacristina De Nardi [b], Fang Yang [c,*]

[a] University of Torino and CeRP-Collegio Carlo Alberto, Italy
[b] UCL, Federal Reserve Bank of Chicago, CEPR, IFS, and NBER, United Kingdom
[c] Louisiana State University, United States

ARTICLE INFO

Article history:
Available online xxxx

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.

Introduction

Wages, labor market participation, hours worked, and savings differ by gender and marital status. In addition, women and married people[1] 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.

http://dx.doi.org/10.1016/j.jeoa.2017.01.005
2212-828X/

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.

Please cite this article in press as: Borella, M., et al.. The Journal of the Economics of Ageing (2017), http://dx.doi.org/10.1016/j.jeoa.2017.01.005