FOR ONLINE PUBLICATION

A Appendix: Figures and Tables

Figure A.1: Implied log wage intercepts for men and women, by education - 1935 cohort

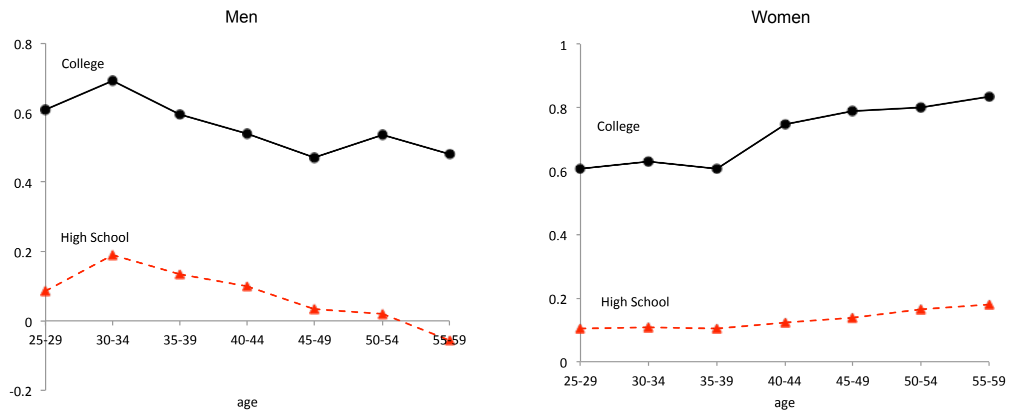

Figure A.2: Proportion of each gender-education group in each marital state during the ages of 25-29, by cohort.

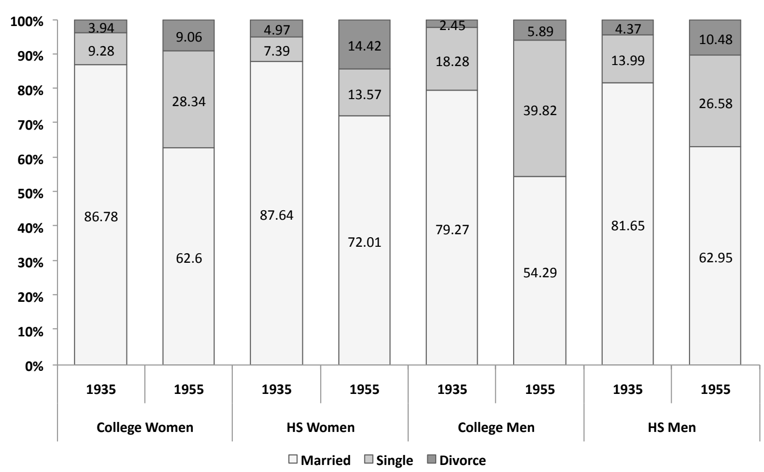

Note: Current Population Survey. Sample consists of all white men and women in each birth cohort. "Married" is defined as "married, with spouse present"; "Divorced" is defined as either "divorced" or "separated". People with at least some college education are those with at least 1 year of college. Those with high school are defined as people with a high school diploma or no more than 12 years of education.

Figure A.3: Linear model: implied log wage intercepts for men and women, by education.

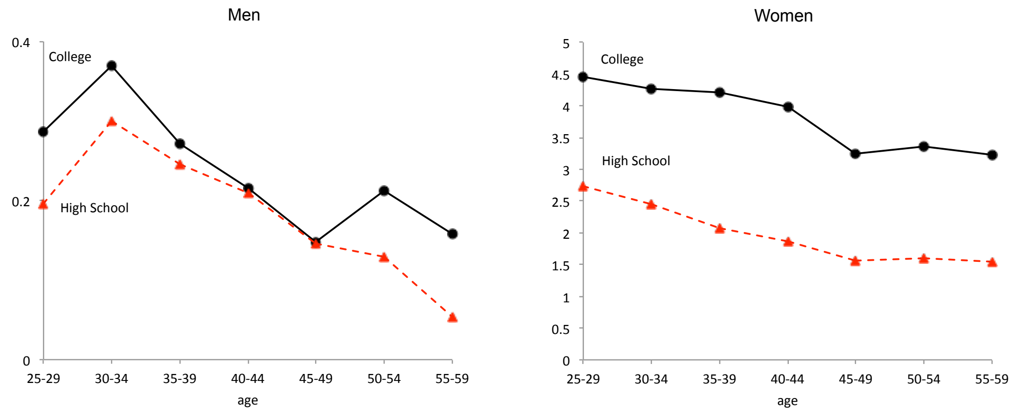

Table A.1: Transition probabilities from single to first marriage

	age	Men College	Men High School	Women College	Women High School
1935 Cohort	30	51.24	44.19	18.14	41.01
	35	26.06	18.90	19.52	3.84
	40	22.65	15.91	23.64	5.63
	45	17.32	6.18	1.63	1.00
1955 Cohort	30	41.89	32.27	41.26	21.84
	35	33.17	11.42	24.39	14.99
	40	21.78	21.28	18.91	15.52
	45	24.68	9.48	14.35	2.28

Note: CPS 1962-2008. Probabilities are calculated using the evolution of the proportion of people who are "never married" between the age shown and 5 years before. These probabilities are conditional on being single.

Table A.2: Proportion of women who are mothers during the ages of 25-29, by cohort, education and marital status

		Single	Married	Divorced
1935 Cohort	High School	8.22	90.81	62.46
	College	2.07	90.98	43.14
1955 Cohort	High School	10.39	81.25	52.37
	College	3.70	59.87	32.59

Note: CPS 1962-2008. Proportion of women between the ages of 25-29 who report having at least one own child in the household.

Table A.3: Implied log wage intercepts for men and women, by education

age	College Men	HS Men	age	College Women	HS Women
25-29	0.61	0.09	25-29	0.61	0.10
30-34	0.69	0.19	30-34	0.63	0.11
35-39	0.59	0.14	35-39	0.61	0.10
40-44	0.54	0.10	40-44	0.75	0.12
45-49	0.47	0.04	45-49	0.79	0.14
50-54	0.54	0.02	50-54	0.80	0.16
55-59	0.48	-0.06	55-59	0.84	0.18

Table A.4: Implied log wage intercepts for men and women, by education - Linear model

	College Men	HS Men		College Women	HS Women
25-29	0.29	0.20	25-29	4.45	2.73
30-34	0.37	0.30	30-34	4.27	2.46
35-39	0.27	0.25	35-39	4.21	2.07
40-44	0.22	0.21	40-44	3.98	1.86
45-49	0.15	0.15	45-49	3.24	1.57
50-54	0.21	0.13	50-54	3.35	1.60
55-59	0.16	0.05	55-59	3.23	1.55

Table A.5: Calibration targets linear model, model vs. data

	Married LFP				Wages	
	College		HS		Gender Wage Ratio	
Ages	Data	Model	Data	Model	Data	Model
25-29	34.59	36.02	26.40	28.06	0.75	0.72
30-34	32.74	35.53	34.22	35.33	0.65	0.64
35-39	47.10	45.40	44.81	43.37	0.60	0.60
40-44	62.34	64.07	54.72	56.54	0.55	0.59
45-49	69.97	68.14	57.30	56.72	0.55	0.62
50-54	68.85	68.66	56.69	58.61	0.59	0.60
55-59	62.50	60.59	48.22	48.67	0.63	0.65
	Divorced LFP				Wages, Women	
	College		HS		Skill Premium	
Ages	Data	Model	Data	Model	Data	Model
25-29	60.90	59.80	60.90	59.80	1.40	1.42
30-34	66.05	65.00	66.05	65.00	1.44	1.42
35-39	73.95	74.52	73.95	74.52	1.39	1.40
40-44	73.87	76.49	73.87	76.49	1.34	1.36
45-49	76.00	76.28	76.00	76.28	1.39	1.42
50-54	73.38	75.02	73.38	75.02	1.48	1.41
55-59	65.13	65.24	65.13	65.24	1.53	1.44
			Skill Premium, men:		1.43	1.43

Note: The total distance between moments and targets is computed as a weighted average of the difference between each moment and target. The moments and targets for the gender wage ratio and skill premium are both multiplied by 100. The weights are such that the 28 moments for LFP receive half of the total weight and the 15 moments for wages receive the other half of the weight. Within each group, the moments are equally weighted.

B Appendix: Data

B.1 LFP

We use the Current Population Survey from 1962- 2010 to compute labor force participation rates. We construct two synthetic cohorts of women born between 1934-1936 and 1954-1956 and compute their LFP as the proportion of women during each age bracket (e.g. 25-29) who report being in the labor force.

Data is available for the most part of the life-cycles of these women. However, note that the women who are aged 26-29 in 1962 correspond directly to the women from our cohort. However, we do not have any information about these women's behavior when they are 25. Thus, in order to compute the LFP for the first period of the women from the 1935 cohort, we also add to the sample women who are 25 years old in 1962 (these women would have been born in 1933). The opposite problem happens for the 1955 cohort. The data ends in 2010 when the median woman of that cohort is 55 years old. The labor force participation behavior for those women during their last period of work (aged 55-59) is computed from the LFP behavior of the women who are aged between 55-59 in 2010 (some of which belong to our cohort - those aged 55 and 56 - and some who were born between 1951-1953).

B.2 Wages

Hourly wages used for the skill premium and the gender wage ratios are computed from the CPS using the individuals' reported labor income and hours and weeks worked last year. We use the sample of white men and women who do not live in group quarters. Prior to 1977, for hours per week, we use the variable which reports the hours worked in the previous week, by intervals; we use the midpoint of the interval. From 1977 onwards, we use the variables for "usual hours worked per week" (last year) and the continuous variable for number of weeks worked last year. Whenever we compute lifetime averages for a variable, we first compute the average of the variable over each year and then average across years. Sample weights are used throughout (PERWT). Concerning top-coded observations, we follow the procedure in Katz and Autor (1999). We mutiply all top-coded observations until 1996 by 1.5. After 1996, top-coded observations in the CPS correspond to the average value of all top-coded observations, thus we do not impose further treatment. We compute the gender wage ratio as the ratio of the average wage of women versus men. The skill premium is computed analogously, using the average wage of college versus high school.

B.3 Income Process

$$\ln y_{egt} = \tau_{eg,t} + \gamma_{eg1} x_t + \gamma_{eg2} x_t^2 + \lambda_e \ln \theta - \delta(1 - P_{t-1}) + w_{et}$$

$$w_{et} = z_{et} + \eta_{et}, \qquad \eta \sim N(0, \sigma^2_{\eta,e})$$

$$z_{et} = \rho_e z_{e,t-1} + \epsilon_{et}, \qquad \epsilon \sim N(0, \sigma^2_{\epsilon,e})$$

Age Profiles:

We use the pooled sample of PSID for the years 1968-2009, restricted to white males who are heads of households. We exclude individuals in the Latino, SEO and immigrant samples. We also drop observations from people younger than 25 and people older than 65 years old and those who report being self-employed. We choose only individuals with at least 8 (not necessarily consecutive) observations. Furthermore, we drop individuals with missing, top-coded and zero earnings those with zero, missing or more than 5840 annual work hours. Individuals with changes in log earnings greater than 4 or less than -2 are also eliminated from the sample. This leaves us with 1645 individuals in the "low" education group and 1261 in the "high" education group.[54]

First, we construct data on hourly wages ($y_{em,t}$) for men of education level e using data on earnings and total hours worked. We then run the following regression in order to estimate the parameters $\hat{\gamma}_{em1}, \hat{\gamma}_{em2}$ of the second degree polynomial on age:

$$\ln y_{em,t} = \mathbb{D}_{emt} + \gamma_{em1} age_t + \gamma_{em2} age_t^2 + w_{em,t}$$

where $\mathbb{D}_{emt}$ is a set of year dummies.

Given the residuals from the regression above, we estimate the parameters for the persistent and transitory shocks using the Minimum Distance Estimator (Chamberlain (1984)). The methods of estimating this process are standard in the literature (see e.g. Heathcote, Storesletten and Violante (2004) for a detailed description of the method). Note that we allow for time-varying $\sigma^2_{\epsilon,e}$ and $\sigma^2_{\eta,e}$ during the estimation process. In the model, we use as inputs the average value across the sample and this is what we report.

In order to estimate the time-varying age intercepts for men $\tau_{em,t}$, although we use all the waves from 1968 to 2009 of the PSID, these years still do not fully span the life-cycles of our two cohorts. Namely, our 1935 cohort was 33 years old in 1968 and our 1955 cohort was 54 years old in 2009. We assume that the wage intercept for the first period of the 1935 cohort (ages 25-29) is the same as the one for the second period (ages 30-34), whereas the last intercept for the 1955 cohort (ages 55-59) is the same as the period before (ages 50-54). As can be seen in the Figures A.1 in the Online Appendix these are reasonable assumptions. We use $\mathbb{D}_{emt}$ to compute the estimates for the time-varying intercepts.

In order to obtain a value for each $\tau_{em,t}$ in the model, we average the $D_{em,s}$ values obtained over

[54] The Panel Study of Income Dynamics (PSID) is the longest panel survey conducted in the US, starting in 1968. Interviews were conducted on an annual basis until 1997, and from then onwards, biennially.

the 5 years corresponding to the $\tau_{em,t}$ interval. Recalling that our earliest data is from 1968 (when our 1935 cohort is 33 years old), our assumptions imply that $\tau_{em,1} = \mathbb{D}_{em,1968}$. We calibrate $\tau_{em,1}$ internally. For the 1955 cohort, on the other hand, we are missing the last intercept $\tau_{em,5}$, since the cohort is only observed up to the median age of 55 in 2009. We thus assume that $\tau_{em,5} = \tau_{em,4}$.

B.4 Divorce and Remarriage Rates

To compute divorce rates ideally one would keep track of each agent's marriage duration conditional upon year of marriage (age). Doing so, however, would add significantly to the computational complexity by increasing the state space. Moreover, it is difficult to estimate with precision the probabilities of divorce/remarriage conditional on both year and duration of marriage/divorce due to small sample size in each year-of-marriage bin. Given these considerations, we choose a simple alternative. Recall that in our model the first period corresponds to the ages of 25-29, the second period, to the ages of 30-34, and so on. Conditional on gender and education, we compute the proportion of people who begin age 25 as married but who undergo at least one divorce between the ages of 25-29. This gives us the probabilities of divorce between periods 1 and 2 in the model. We proceed analogously for the remaining periods, deriving 4 age-varying divorce probabilities by education and gender. Thus, our model reproduces, period-by-period, the correct proportion who divorced in each age bracket.

For remarriage rates, due to sample size restrictions, we proceed in a different way.[55] We calculate, conditional on gender and education, the proportion of people who remarry before the twentieth anniversary of their last divorce and assume a uniform remarriage rate over this twenty year interval.[56]

[55] The number of people in the SIPP who entered ages 25 and 30 as divorced is too small – around 30 for men and 15 for women – to employ the same procedure.

[56] We augment our birth cohorts by 2 years due to sample size (i.e. the 1935 cohort is defined as all people born between 1933-1937). To calculate the remarriage rate we proceed in the following manner: individuals who divorced after the age of 30 contribute to the remarriage rate if they remarried within 20 years of their divorce and prior to the

B.5 Fertility

The timing of fertility shocks is as follows. For the 1935 cohort, single women and college women of all marital status receive a fertility shock in the first period that takes the value of zero or one. Divorced women and married high-school women receive a fertility shock that takes the value of zero or two. The probability of receiving a non-zero value is calibrated so as to match the initial proportions (age 25-29) in the data by marital status and education as reported in Table A.2. In period 2, divorced and single women are not hit by fertility shocks whereas all women who are married in period 2 are assumed to have an additional child. Lastly, all women who were married in both periods 1 and 2 receive an additional fertility shock in the second period that can take the value of zero or two for high-school women, and zero or one for college women. The frequency of shocks is calibrated to generate 2.54 children per college woman, 3.20 for high school women and an overall average of 3.00 children per woman for the 1935 cohort.[57]

For the 1955 cohort the structure and values of fertility shocks are the same as those for the 1935 cohort. The frequencies in the initial period are adjusted so as to match the proportions in the data for the 1955 cohort between the ages of 25-29 as reported in Table A.2. In the second period, once again, no divorced or single women receive any additional children. All women who were not married in period zero and got married at the end of period one receive a child. Lastly, once again, all women who were married in both periods are hit by a fertility shock that can take the value of zero or one.[58] This generates the following numbers of children over their lifespan:

age of 60, if not, they are considered divorced. For those individuals who enter the age of 30 already divorced, the 20 year window in which to count as "remarried" commences as of age 30. This is the correct procedure since the initial proportion of individuals across marital states already accounts for divorces and remarriages that took place before the age of 30.

[57]The numbers for average number of children were computed using the PSID by calculating the average number of children ever born to women from each of our cohorts by the time they reached aged 40. Due to sample size constraints, we define the 1935 cohort as women born in 1933-1937 and the 1955 cohort those born in 1953-1957. In order to generate the correct number of average children, we assign a 98% probability of a non-zero fertility outcome during the second period for women who have been married for two periods if they have college education and 43% probability if they have high school.

[58]Once again, to generate the correct number of average children seen in the data, we assign the probabilities of receiving an extra child during the second period to women who have been married for 2 periods of 73% if they are

college women have 1.74 children on average, high school women have 2.11 and, overall there are 1.91 children per woman for the 1955 cohort.

B.6 Consumption Deflator:

We use an altered McClements scale ($\mathbf{e}(k_t; s)$) in order to deflate household consumption. Table B.6 reproduces the original McClements scale in normalized for one adult.

Table B.6: McClements Scale

			+1 child, by age:						
1 adult	2 adults	+ 1 adult	0 - 1	2- 4	5-7	8-10	11-12	13-15	16-18
1	1.64	+0.75	+ 0.148	+0.295	+ 0.344	+0.377	+0.41	+0.443	+0.59

Since we have 5 year periods, and our children are aged 0-4, 5-9, 10-14, 15-20, we weigh the scale accordingly. For example, a child aged 0-4 will add: $0.4(0.148) + 0.6(0.295) = 0.2362$.

The scale $\mathbf{e}(k_t; s)$ is constructed using the number of adults in the household (1 if $s = S, D$ and 2 if $s = M$) and the number of children and their respective ages (k_t).

B.7 Pensions:

To compute retirement benefits for a model household, we modify the approach used in Heathcote, Storesletten and Violante (2010) in order to avoid keeping track of an individual's average earnings over their lifecycle. More specifically, we take each individual's last observed earnings y_T and compute social security benefits as follows: 90% of y_T up to a first threshold equal to $0.38\bar{y}_T$, where $\bar{y}_T$ is the average observed earnings in the economy during the last period, plus 32% of y_T from this bendpoint to a higher bendpoint equal to $1.59\bar{y}_T$, plus 15% of the remaining y_T exceeding this last bendpoint. For married households, this process is done for both the husband and the wife; the household total benefits are either the sum of their benefits or 1.5 times the husband's benefits,

college and 52% if they are high school.

whichever one is highest.

C Appendix: Model Solution

In our model, households have a finite horizon, so the dynamic problem is solved numerically by backwards recursion from the last period of life using value function iteration. At each age, the households solve for their consumption savings rule and participation decisions taking as given their state variables that period and next period's value function. In addition to assets, our model has two other potential continuous state variables: the persistent components of earnings w_t^{em} and of the wife, w_t^{ef}. Handling more than one continuous state variable is possible but computationally costly. We choose to discretize these two variables, leaving assets as our only continuous state. We use 25 nodes whose locations are age-dependent for each of the earnings components.

During the working stages of the lifecycle, our model combines a discrete decision (whether the woman participates in the market) and a continuous decision (the amount of savings). This combination may lead to non-concavities in the value function. Furthermore, the existence of transitions across marital states also requires some attention. For all periods $t > 4$, since there are no longer any transitions across different marital status, the maximization problem for the single men and women and divorced men is a straightforward consumption-savings problem. The problem of the married couple and of the divorced woman combines the discrete participation choice of the wife together with the continuous choice of assets. The combination implies that concavity of the value function is not guaranteed even if one controls for the participation decision that period. Given enough uncertainty the value function conditional on today's participation is concave We follow Attanasio et al. (2008) and impose (and check) a unique level of reservation assets a_t^* at which, given the values of all other state variables, the conditional value functions (working vs not working) intersect only once and thus the woman's participation decision switches at that point from not working to working. Thus conditional on all other states, for all values $a_t < a_t^*$, the woman works and for all values $a_t > a_t^*$ the woman does not work. We numerically check both the concavity of the conditional value functions and the uniqueness of the reservation asset level.

The optimization problem of the household who enters a period $t \leq 4$ as a married couple must take into account the continuation values of the husband and wife, which are different. It is also important to note that the solution of the optimization problem for both divorced and single agents involves the calculation of a fixed point. Recall that each agent optimally chooses her/his asset accumulation and labor force participation/experience (for women), taking into account her/his expectations about their potential spouses optimal decision paths while the potential spouse does the same. This generates a computationally iterative process whereupon the optimal (asset and labor market experience) decisions for each potential couple are given by solutions consistent the optimal decision paths taken as given by their counterpart, e.g. in order for $a^{\mathbf{d,s}}_{m,t+1}, P^*_t, a^{\mathbf{d,s}}_{f,t+1}$ to be a solution in a given period, then it must be the case that the potential wife chose $a^{\mathbf{d,s}}_{f,t+1}$ and P^*_t taking as given her expectations about her potential husbands optimal decision path while her potential husband chose $a^{\mathbf{d,s}}_{m,t+1}$ taking as given his expectations over his potential wifes path of optimal decisions. In this sense, although a marital type is assigned exogenously to each agent, the state space of their potential match in each period is endogenously determined in the model and each agents forms the correct expectations over the endogenous path of optimal decisions of their potential spouse. Thus the algorithm is as follows:

- Given $a^{\mathbf{d}}_{m,t+1,0}$ and all the other state variables, the ex-wife chooses $a^{\mathbf{d}}_{f,t+1,0}$ and P_t.
- The ex-husband takes $a^{\mathbf{d}}_{f,t+1,0}$ chosen by the wife and chooses $a^{\mathbf{d}}_{m,t+1,1}$.
- If $a^{\mathbf{d}}_{m,t+1,1} = a^{\mathbf{d}}_{m,t+1,0}$ then this decision point is done. If not, then the wife will take$a^{\mathbf{d}}_{m,t+1,1}$ as given and solve for a new value of $a^{\mathbf{d}}_{f,t+1,1}$ and P_t.

- This process is iterated until convergence, defined as $|a^{\mathbf{d}}_{m,t+1,j} - a^{\mathbf{d}}_{m,t+1,j+1}| < 1^{-10}$

If the process above converges after D iterations, then the asset levels at time $t+1$ for the divorced female and male are given by $a^{\mathbf{d}}_{f,t+1} = a^{\mathbf{d}}_{f,t+1,D}$ and $a^{\mathbf{d}}_{m,t+1} = a^{\mathbf{d}}_{m,t+1,D}$