

Inequality and the Marriage Gap

Nawid Siassi*

Universidad Carlos III de Madrid

January 23, 2013

Abstract

Marriage is one of the most important determinants of economic prosperity, yet most existing theories of inequality ignore the role of the family. This paper documents that the cross-sectional distributions of earnings and wealth display a high degree of concentration, even when disaggregated into single and married households. At the same time, there is a large marriage gap: married people earn on average 27 percent more income, and they hold 34 percent more net worth. To account for these empirical facts, I develop a stochastic OLG model with female and male agents, who (i) are randomly selected into single or married households at the beginning of their economic lives; (ii) face uninsurable labor market risk henceforth; (iii) and make Pareto-efficient decisions if married. In a calibrated version of the model, I show that matching patterns by educational attainment, an effective tax bonus for married couples and directed bequests are key to explaining the marriage gap in earnings and wealth. A policy experiment of moving from joint tax filing for married couples to separate filing yields output and welfare gains.

Keywords: Inequality; Wealth Distribution; Marriage Gap; Incomplete Markets.

JEL Classification Numbers: D13; D31; D91; E21.

*Address: Departamento de Economía, Calle Madrid, 126, 28903 Getafe (Madrid), Spain. E-Mail: Nawid.Siassi@uc3m.es. I wish to thank Arpad Abraham, Martin Flodén, Piero Gottardi, Francesc Obiols-Homs, my thesis advisor Salvador Ortigueira, Nicola Pavoni, Mike Mariathasan and seminar participants at the European University Institute and at Universidad Carlos III de Madrid for valuable comments.

1 Introduction

Marriage is one of the most important determinants of economic prosperity. Yet, perhaps surprisingly, most existing theories of inequality abstract from the role of the family: the standard framework for studying inequality treats all households as being comprised of a single decision-maker, without making the role of the marital status explicit. The main contribution of this paper is to fill this void and present a theory that can account for the observed inequality between single and married households.

The cross-sectional distributions of earnings, income and wealth in the United States display a large degree of concentration.¹ When disaggregated into married and single households, economic prosperity remain very unequally distributed *within* both subgroups. At the same time, there is a striking divergence *between* both subgroups: on average, married people have 49.4 percent higher labor earnings, they earn 26.8 percent more income, and they are 33.5 percent richer than singles. This disparity, the marriage gap, is not driven by extremely rich households as the corresponding ratios of medians look similar. In light of the empirical relevance of the family – in the year 2009, half of the adult population in the United States was married – reconciling the strong association between marital status and economic outcomes is a challenge that models of inequality must face.

In order to account for these stylized facts, I construct an extended version of the neoclassical growth model with incomplete markets and idiosyncratic risk, which is the standard framework for studying inequality (see Aiyagari (1994)). Following recent studies, I mix some desirable features of both life-cycle and dynasty models by assuming stochastic transitions from working age to retirement and eventually death. Throughout working age, individuals receive idiosyncratic labor efficiency shocks. They use buffer-stock savings in a riskless asset, subject to a borrowing constraint, to smooth consumption over time, and they save for retirement. My further modeling choices are motivated by the focus of this paper. I assume that there are equally many female and male individuals who are randomly selected into households of different sizes when entering the economy. Some households consist of one person (“single”), others consist of two persons (“married”). Two-person households, formed by a female and a male, pool their income and make Pareto-efficient decisions on individual consumptions, labor supplies and joint savings.

A calibrated version of the benchmark model is largely successful in accounting for the facts from the data. The model generates a significant degree of inequality within the subgroups of single and married households, and it predicts a positive marriage gap for earnings, income and wealth. Three factors are key for the success of the model. First, one of the novel features I propose in this paper relates to the explicit distinction between intentional and accidental bequests: if intergenerational

¹See, for example, Díaz-Giménez, Glover and Ríos-Rull (2007), Heathcote, Perri and Violante (2010), and Hintermaier and Königer (2010).

ties are tighter in families with descendants, they have a stronger incentive to transmit their estates to the next generation. Since married households tend to have more descendants, a dynastic saving motive adds to explaining the marriage gap in wealth. Secondly, I acknowledge that the U.S. tax code encourages married couples to file their income taxes jointly. Joint filing will often result in a more favorable tax bracket and, thus, raise permanent disposable income. Since precautionary saving is tightly associated with a target wealth-to-permanent-income ratio, married couples are led to save more. Thirdly, I introduce heterogeneity in permanent abilities and a marriage market that reflects empirically observed marriage patterns across different education groups. The interplay between a higher propensity to marry for college-educated people and the assortative matching component is a key ingredient for generating the marriage gap in earnings and wealth.

I quantitatively assess the relevance of these three factors by solving a more parsimonious version of the model. This restricted model can be considered as a version of the standard incomplete-markets life-cycle model that is simply augmented by two household types – single and married. The restricted model performs substantially worse in accounting for the marriage gap in earnings and income, and it counterfactually predicts *lower* per-capita wealth for married people than for singles. The reason is that the risk-sharing component of household saving attenuates the precautionary motive for families with multiple earners.² In a series of counterfactual experiments, I show that the three factors laid out above all contribute significantly to reconciling the model with the data. The benchmark model also does a good job of accounting for the empirical shares of liquidity-constrained single and married households: while roughly one sixth of all married households classify as liquidity constrained (using Zeldes’ (1989) definition), the corresponding fraction of single households is almost 13 percentage points higher. Consistent with the data, the benchmark model further predicts that the marriage gap in wealth is particularly large for poor and middle-class households, whereas the role of the marital status is somewhat less important for rich households. Finally, I employ the model to conduct a hypothetical policy reform that abolishes the possibility of joint tax filing for married couples. My results indicate that moving to separate tax filing would imply sizable welfare gains through an effective redistribution from married to single households.

This paper relates to two strands of literature. First, it builds upon earlier work that studies cross-sectional income and wealth inequality in general equilibrium frameworks with heterogeneous agents, e.g. Aiyagari (1994), Huggett (1996), Krusell and Smith (1998), Castañeda et al. (2003) and De Nardi (2004). All of these studies, however, abstract from modeling the marital status of a household. A second body of literature makes this distinction explicit by considering single and

²Ortigueira and Siassi (2012) study the role of intrahousehold risk sharing in an environment with infinitely-lived workers who are subject to uninsurable employment risk. They find strong excess saving effects for workers who do not have access to family insurance, in particular among wealth-poor households.

married households separately; examples include Aiyagari et al. (2000), Greenwood et al. (2003), Cubeddu and Ríos-Rull (2003) and Hong and Ríos-Rull (2003).

To the best of my best knowledge, there is no theoretical work on the role of the marital status and cross-sectional inequality in a joint context. The study most closely related is a working paper by Guner and Knowles (2004) who investigate the link between marriage and wealth in a dynamic OLG setting. In their model, single agents and married couples make decisions on consumption, hours worked and savings, and they decide whom to marry and when to divorce, in anticipation of future outcomes. The authors show that their model can generate a significantly positive wealth gap. The main difference to my framework is that Guner and Knowles (2004) model consumption within married households as a public good and calibrate it using estimates for adult equivalence scales, which may be a driving force for their results. Since their model only consists of three periods, it also neglects the strong intrahousehold insurance effect on precautionary saving and probably performs poorly when tested along the cross-sectional dimension.

The divergence in effective taxation between single and married households and the role of joint tax filing has been the subject of a recent study by Guner, Kaygusuz and Ventura (2008). These authors construct a life-cycle economy populated by single and married workers who differ according to their labor efficiency and age. At the heart of their analysis lies an exogenous utility cost of participating in the labor market. This assumption allows them to focus on the extensive margin of married female labor supply. The authors use their model to evaluate various tax reforms, *inter alia* the abolition of joint tax filing. Their results associate substantial welfare gains with such a reform and, thus, share a commonality with my own findings.

The role of a bequest motive to generate a lifetime saving profile consistent with the data has been recently put under examination. De Nardi (2004) shows that intentional bequests can explain the emergence of very large estates and, therefore, help to generate a high degree of concentration at the upper tail of the wealth distribution. Fuster et al. (2008) study the significance of intergenerational links for the impact of various tax reform proposals. They confront two polar frameworks: a pure life-cycle model, on the one hand, and a dynastic model with altruistic links, on the other hand. They find that tax reforms can have very different implications depending on whether individuals derive utility from bequeathing to their descendants or not. Laitner (2001) introduces the existence of intentional and accidental bequests in a common framework. In his model, a constant fraction λ of households care about their heirs; the remaining households only cares about their own utility. In comparison to his approach, a novel element in my model is to relate the existence of a bequest motive explicitly to the presence of a descendant.

The remainder of the paper is organized as follows. Section 2 documents the empirical facts motivating this study. In section 3, I present my benchmark model economy and define a stationary equilibrium. The calibration strategy is described in Section 4, and Section 5 contains my results.

Concluding remarks are offered in Section 6.

2 The Data

To motivate this study, this section documents a number of stylized facts for the distributions of earnings, income and wealth in the United States. Most of the data analysis is based on the 2007 wave of the Survey of Consumer Finances (SCF).³ One advantage of the SCF is that it provides information on all three variables of interest for this study, whereas e.g. the Current Population Survey (CPS) does not collect any data on household wealth. A second advantage is that the Survey of Consumer Finances explicitly oversamples wealthy households and employs appropriate weighting schemes to adjust for higher non-response rates among rich households. Therefore, the SCF provides a more accurate description of the upper tails of the various distributions, as distinguished from other U.S. household surveys such as the CPS or the PSID.

For the purpose of this study, I restrict the sample to comprise only households where the head is at least 25 years old. I make an additional adjustment by excluding the wealth-richest 1 percent of households for the following reason. In a previous study, Castañeda, Díaz-Giménez and Ríos-Rull (2003) find that matching the concentration at the very top of the wealth distribution requires a small-probability state of extremely high hourly wages. For instance, in their benchmark economy agents in the highest efficiency state are more than 100 times more productive than those in the second-highest state, and they are more than 1,000 times more productive than agents in the lowest state. In the model presented in the next section, agents draw their labor efficiency based on a stochastic earnings process that has been estimated from PSID data. Since the very rich households are neither present in the PSID nor in my model, I choose to abstract from them.⁴

The upper panel in Table 1 summarizes a selection of distributional statistics for labor earnings, total income and wealth across households in the U.S. economy. As is well known, all three variables are very unequally distributed, with wealth being by far the most concentrated one among them. For instance, households belonging to the bottom 40 percent of the respective distribution earn 11.7 percent of income and they hold only 1.8 percent of total wealth. The Gini coefficient – a more sensitive concentration measure for the upper tail of the distribution – exceeds 0.5 for all variables of interest and is particularly high for wealth (0.72). These facts indicate that the cross-sectional distributions of earnings, income and wealth are highly skewed to the right, with fat lower tails and a very thin upper tail.

The middle and lower panels in Table 1 display the same set of statistics when the sample is

³A detailed description of the data and variable definitions are provided in Appendix I.

⁴Recent studies by Heathcote, Storesletten and Violante (2010) and Hintermaier and Königer (2011) pursue a similar strategy.

Table 1. Summary statistics

	Mean (\$)	Median (\$)	Gini	Bottom 40%
All Households				
Labor earnings	57,076	37,364	0.59	4.72
Total income	73,868	48,745	0.50	11.71
Wealth	398,872	136,750	0.72	1.81
Married Households				
Labor earnings	78,919	56,560	0.53	8.79
Total income	99,768	68,489	0.48	13.51
Wealth	551,571	191,800	0.72	2.82
Single Households				
Labor earnings	26,416	15,426	0.63	0.54
Total income	39,348	29,617	0.44	14.15
Wealth	206,555	63,846	0.71	0.66

Source: 2007 wave of the Survey of Consumer Finances (SCF).

partitioned into married and single households. As can be seen, earnings, income and wealth remain very unequally distributed within the two subsamples. Moreover, married households earn significantly more income and they hold substantially more assets than single households, even when dividing by the number of household members (cf. Table 1, first two columns). To make this point more explicit, I define the *marriage gap* in variable x (e.g. mean labor earnings) as

$$\Delta(x) \equiv 100 \cdot \left(\frac{1}{2} x^{\mathcal{M}} / x^{\mathcal{S}} - 1 \right), \quad (1)$$

where I divide the value for married households, $x^{\mathcal{M}}$, by 2 in order to compute the per-capita value. The marriage gap $\Delta(x)$ is then obtained as the percentage deviation between married and single households. Table 2 reports the resulting values for this measure. As can be seen in the table, married individuals earn on average 49.4 percent more labor income than singles. The marriage gap in average total income amounts to 26.1 percent, and married people are on average 33.5 percent richer than singles. Put differently, while only about 60 percent of the population in the sample is married, they hold almost 80 percent of total wealth and they earn 82 percent of total labor income.

In the remaining part of this section, I will assess how the marriage gaps in earnings, income and wealth evolve over the life cycle and to what extent they are driven by extreme observations. One hypothesis is that many people get married later than at the age of 25, which would imply that they enter the sample of married households at a later point of their increasing life-cycle profile of earnings and wealth. To check for this possibility, I restrict the sample to households where

Table 2. The marriage gap

	Popul share (%)	Δ^{Mean}	Δ^{Median}
Labor earnings			
Full sample	100.0	+ 49.4	+ 83.3
Age 30+	92.1	+ 55.3	+ 97.8
Working age (25 – 59)	69.9	+ 25.6	+ 13.8
Retirement age (60+)	30.1	+ 140.0	–
Full sample (CPS)	100.0	+ 27.6	+ 52.5
Total income			
Full sample	100.0	+ 26.8	+ 15.6
Age 30+	92.1	+ 30.1	+ 20.4
Working age (25 – 59)	69.9	+ 17.5	+ 10.8
Retirement age (60+)	30.1	+ 43.7	+ 9.5
Full sample (CPS)	100.0	+ 12.8	+ 24.5
Wealth			
Full sample	100.0	+ 33.5	+ 50.2
Age 30+	92.1	+ 34.3	+ 45.6
Working age (25 – 59)	69.9	+ 38.4	+ 116.8
Retirement age (60+)	30.1	+ 50.8	+ 7.6
Retirement - no widows	21.4	+ 69.2	+ 30.3

Sources: 2007 wave of the Survey of Consumer Finances (SCF) and March 2010 wave of the Current Population Survey (CPS).

the head is at least 30 years old, an age by which most of the first marriages have been formed.⁵ Under the given hypothesis, one would expect the marriage gaps in earnings, income and wealth to be smaller than for the full sample. Instead, they are even larger (cf. Table 2). This suggests that the marriage gap at young age is comparatively small and then opens up over the life cycle.

To further investigate this point, I divide the sample into working-age and retirement-age subgroups. Labor earnings for working-age married individuals are on average 25.6 percent higher than for working-age singles. The corresponding statistic for retired individuals is with +140 percent much higher, but this result is mainly driven by the fact that there are more married people participating in the labor market at old age. As for the marriage gaps in income and wealth, the picture looks similar: married individuals earn on average 17.5 percent more income and they hold 38.4 percent more assets. During retirement age, the discrepancy jumps up to +43.7 percent (total income) and +50.8 percent (wealth) respectively.

⁵The U.S. Census Bureau reports that the median age at first marriage for women and men in the year 2009 was 26.5 and 28.4 respectively.

It should be noted that the sample of single households comprises widowed individuals, and that the share of widows is substantially larger for the retirement subsample. In order to evaluate the importance of this fact, the last line of Table 2 reports the marriage gap in wealth for people at retirement age if widowed singles are excluded. The discrepancy rises even further to +69.2 percent (compared to +50.8 percent when widows are not excluded). This suggests that widowed retired individuals are substantially better off than other singles, which is consistent with the previous findings under the assumption that the death of one spouse does not intrinsically elevate the survivor's economic well-being to a large extent.⁶

Are these results driven by extreme observation, e.g. by very rich households? To answer this question, the last column in Table 2 documents the respective values for the marriage gap in median values. The disparity in wealth between the median married household and the median single household is with +50.2 percent even more pronounced than the corresponding value for means (+33.5 percent), whereas the opposite is true for total income (+15.6 percent compared to +26.8 percent). The gap in median labor earnings for households of all age groups rises to +83.3 percent. A more meaningful comparison between working-age married and single households yields a gap of +13.8 percent. As a final exercise, I investigate whether my previous findings are in some way specific to the relatively small SCF sample by performing a similar analysis for the Current Population Survey (CPS) with its much larger sample size. As can be seen in Table 2, the marriage gaps in average earnings (+27.6 percent) and income (+12.8 percent) remain notably high. They are slightly smaller than in the SCF, which is perhaps not surprising as income-rich households are underrepresented in the CPS.⁷

To summarize, the preceding empirical analysis has uncovered two main findings: first, the cross-sectional distributions of earnings, income and wealth in the United States are highly concentrated and skewed to the right. This holds true for the sample of all households, and for the two subsamples of single and married households. Second, married people earn considerably higher income and they are richer than singles. This marriage gap – i.e. the difference in per-capita values between married and single individuals – is positive over the life cycle and robust to outliers. With the aim of constructing a theory that is consistent with these empirical facts, I now turn to presenting my benchmark model.

⁶In the model I abstract from widows by assuming that married couples decess jointly.

⁷There is no comparative value for wealth since the CPS does not collect any data on household net worth.

3 The Model

3.1 Preliminaries

Demographics. Consider an economy that is populated by a continuum of measure one of households. In each period $t = 0, 1, 2, \dots$, a cohort of new individuals enters the economy. Half of them are born as females, the other half as males. The life cycle of an individual consists of three phases: household formation, working age and retirement. The first phase takes place before an agent enters the economy (“time 0”) and determines whether individuals will commence their working lives as singles or couples. Once households are formed, they are either comprised of one single adult or two married adults, one female and one male. Households can be either in the working-age or the retirement stage. At the end of each period, working-age households face a constant exogenous probability of becoming retired, and retired households face a constant exogenous probability of dying. When a retired household dies, its members are replaced by an equal number of newborn agents. The deceased household’s financial wealth is liquidated and transmitted to the next generation.⁸ In addition, married couples face an exogenous probability of separation throughout the working-age and retirement stages (“divorce”). Divorced agents form single households for the rest of their lives.

Preferences. All agents enjoy the consumption of an aggregate good and of leisure time. Preferences for agents of gender $g \in \{f, m\}$ can be described by a per-period utility function $U^g(c_t, l_t)$, where c_t and l_t denote consumption and leisure in period t respectively, and a common discount factor $\beta \in (0, 1)$. I will assume that U^g is strictly increasing and strictly concave in each of its arguments, twice continuously differentiable and satisfies the Inada conditions. In addition, agents derive utility from bequeathing their estate to their descendants.

Employment opportunities. In each period, agents are endowed with one unit of disposable time and an individual level of labor productivity e that depends on their history of idiosyncratic shocks. Retired agents are not productive at all, i.e. $e = 0$. In the working-age phase, the labor productivity of individual i at time t is given by

$$e_t^i = \exp(\xi^i + z_t^i), \quad (2)$$

where ξ^i is a permanent component that is determined when an agent is born and may be interpreted as an ability shock. I assume that ξ^i is drawn from a finite set Ξ that contains zero as an element. The time-varying part of labor productivity, z_t^i , evolves according to an AR(1) process,

$$z_t^i = \rho z_{t-1}^i + \epsilon_t^i, \quad \text{with } \epsilon_t^i \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_\epsilon^2), \quad (3)$$

⁸The transmission of wealth is detailed below.

where ρ measures the longevity of temporary productivity shocks. To model transitions to retirement, at the end of each period, there is a probability ϕ^R that labor productivity is set to zero permanently, i.e. $e_t^i = 0$, $\forall t = t + 1, t + 2, \dots$. Agent i 's labor productivity in period t can then be summarized as $s_t^i \equiv (\xi^i, e_t^i)$, where $s_t^i \in S \equiv \Xi \times \mathbb{R}$. Note that s implicitly describes whether an agent is in working age, $e > 0$, or retired, $e = 0$.

Household formation. Before a new cohort of agents enters the working-age stage, it is determined whether they will start their economic lives in a one-person (“single”) or two-person (“married”) household. I assume that there is a marriage market that randomly matches two individuals i and j of opposite gender according to an exogenous probability q_{ξ^i, ξ^j}^g that potentially depends on their relative abilities. The latter assumption allows me to model the positive assortative matching component of couple formation, which implies the significant correlation between permanent wages that is observed in the data.⁹ Once two individuals are matched, they enter into a cooperative bargaining process that prescribes efficiency for the resulting allocation. They can fully commit to this outcome until their marriage is dissolved exogenously or they die together. The married household maximizes a weighted sum of its members’ utilities where relative weights are set upon matching and remain fixed thereafter. In this paper I adopt a unitary model of the household and treat utility weights as parameters. Individuals who are left unmatched remain singles and form one-person households. After individuals have been selected into households, it is furthermore determined whether the household has descendants ($d = 1$) or not ($d = 0$). The presence of descendants has an impact on the bequest motive and occurs according to an exogenous probability that depends on the marital status: married households are assigned descendants with probability π^M , and single-person households are assigned descendants with probability π^S .

Intergenerational transmission. Successive generations of individuals are linked through the transmission of assets in the following stylized way. I assume that every deceased married couple leaves their estate to *two* new entrants in equal shares, i.e. each of the two entrants inherits half of the assets. Moreover, I assume that for every deceased single agent there is another single agent of opposite gender who deceases at the end of the same period; upon death, their estates are pooled and left to two new entrants in equal shares. From the perspective of newborn individuals, initial asset holdings are determined when they enter the working age, i.e. after the household formation stage. Each new entrant inherits either half of the estate of a randomly selected deceased married couple, or half of the sum of assets of two randomly selected deceased single individuals of opposite gender.¹⁰

⁹See Hyslop (2001).

¹⁰Two clarifying remarks are in order. First, an underlying assumption is that there is a “veil of ignorance” between deceased agents and newborns, in the sense that neither side knows the other’s identity until the transmission of assets takes place. Second, a constant population size requires that each deceased individual has on average one descendant. For simplicity, I assume that pairs of agents with $d = 1$ bequeath only to two heirs – they could have

Government. The government levies taxes on income, collects payroll taxes and pays out benefits to retired individuals. Income taxation for single and married households is characterized by two functions, $\tau^S(y)$ and $\tau^M(y)$, where total household income y is composed of labor earnings, capital income and retirement benefits. Payroll taxes are levied on a flat-rate basis on labor earnings, where the tax rate is denoted by τ_p . Retirement benefits are allowed to depend on the gender and ability mix of all household members. The government cannot issue any debt and is thus required to balance its budget on a period-by-period basis.

Firms. Production of the aggregate good is conducted by a continuum of competitive firms. The representative firm operates a technology that can be represented by the Cobb-Douglas production function $F(K, L) = K^\alpha L^{1-\alpha}$, where K is the aggregate stock of capital, L is aggregate labor and $0 < \alpha < 1$ is the capital's share of income. Female and male labor are assumed to be perfect substitutes, $L \equiv \lambda L^m + (1 - \lambda)L^f$, where λ is a parameter that pins down relative productivities and can thus be used to model the gender gap in wages. The firm's maximization problem is static: given a rental price of capital r and gross wages for females and males \bar{w}^f and \bar{w}^m , respectively, first-order conditions are:

$$F_K(K, L) = r + \delta \quad (4)$$

$$\lambda F_L(K, L) = \bar{w}^m \quad (5)$$

$$(1 - \lambda)F_L(K, L) = \bar{w}^f, \quad (6)$$

where $\delta > 0$ denotes the depreciation rate of capital. Net wage rates for females and males will be denoted by $w^f = (1 - \tau_p) \bar{w}^f$ and $w^m = (1 - \tau_p) \bar{w}^m$, respectively.

Market structure. A crucial assumption for the model at hand is that there are no markets for state-contingent contracts in the economy; hence, workers cannot insure perfectly against idiosyncratic labor market uncertainty. Also, there is no annuity market to insure individual mortality risk. The only asset in the economy is physical capital, which pays out the risk-free interest rate r . Moreover, I assume that individuals in this economy are not allowed to borrow, which imposes a zero lower bound on their asset holdings. The latter assumption also implies that agents cannot die in debt.

3.2 The Problem of the Household

Single household. For a single agent of gender g the relevant state variables are current wealth a , a vector describing the agent's labor efficiency $s = (\xi, e)$, and whether there are descendants, more than two heirs – and that their other descendants inherit from pairs of agents with $d = 0$. Equivalently, one could assume that each individual has descendants and only those with $d = 1$ have a bequest motive.

$d \in \mathcal{D} = \{0, 1\}$.¹¹ The problem of a single household can be formulated recursively as

$$\begin{aligned} \mathcal{V}^g(a, s, d) &= \max_{c, l, a'} \left\{ U^g(c, l) + [1 - \phi(s)] \beta \mathbb{E} [\mathcal{V}^g(a', s', d) | s] + \phi(s) \beta \mathcal{Z}^g(a', d) \right\} \quad (7) \\ \text{s.t.} \quad & c + a' = y - \tau^S(y) + a \\ & y = b^g(s) + (1 - l) e w^g + ra \\ & c \geq 0, \quad 0 \leq l \leq 1, \quad a' \in \mathcal{A}, \quad \text{and (2), (3),} \end{aligned}$$

where $\mathcal{A} = [0, \bar{A}]$, and \bar{A} is an upper bound for asset holdings that is sufficiently large such that it never binds. Recall that when a household retires, its labor efficiency is permanently set to zero, $e = 0$. The function $\phi(s)$ describes the probability of dying at the end of the period and takes on a positive value only when a household is already retired. That is,

$$\phi(s) = \phi(\xi, e) = \begin{cases} \phi^D & , \text{ if } e = 0, \\ 0 & , \text{ if } e > 0. \end{cases}$$

Similarly, retirement benefits $b^g(s)$ are only paid out to retired households. The value of bequeathing remaining estates to descendants, $\mathcal{Z}^g(a', d)$, depends positively on a' and will be described in more detail later.¹²

Married household. Consider now the maximization problem faced by a married household. As explained above, the utility of each individual in the household carries a weight, reflecting the relative power of that individual in the household. Under full commitment, that is, when household members can commit to future intra-household allocations, individual weights are set when the household is formed and remain unchanged thereafter. Let $\mu \in [0, 1]$ be the Pareto weight on the female's utility. Denote by $\mathbf{s} = (s^f, s^m)$ the pair of states describing the labor productivity of both members in a married household, where $\mathbf{s} \in \mathbf{S} \equiv S \times S$. A married household solves

$$\begin{aligned} \mathcal{V}(a, \mathbf{s}, d) &= \max_{c^f, c^m, l^f, l^m, a'} \left\{ \mu U^f(c^f, l^f) + (1 - \mu) U^m(c^m, l^m) + \phi(\mathbf{s}) \beta \mathcal{Z}(a', d) \right. \quad (8) \\ &\quad \left. + [1 - \phi(\mathbf{s})] \beta \mathbb{E} \left[(1 - \psi) \mathcal{V}(a', \mathbf{s}', d) + \psi \mathcal{S}(a', \mathbf{s}', d) \mid \mathbf{s} \right] \right\} \\ \text{s.t.} \quad & c^f + c^m + a' = y - \tau^M(y) + a \\ & y = b(\mathbf{s}) + (1 - l^f) e^f w^f + (1 - l^m) e^m w^m + ra \\ & \mathcal{S}(a, \mathbf{s}, d) = \mu \mathcal{V}^f(a/2, s^f, d) + (1 - \mu) \mathcal{V}^m(a/2, s^m, d) \\ & c^f, c^m \geq 0, \quad 0 \leq l^f, l^m \leq 1, \quad a' \in \mathcal{A}, \quad \text{and (2), (3).} \end{aligned}$$

¹¹For notational convenience, I will suppress the dependence on interest rates and wages. Since this paper focuses on stationary equilibria only, prices will be constant over time.

¹²It would be straightforward to extend the model so as to allow for an intergenerational transmission of earnings ability (see De Nardi (2004)). Since the model is already fairly complex, I abstract from this possibility.

Married households that do not die at the end of the period face a constant probability ψ of divorcing. In case a divorce occurs, the joint continuation value \mathcal{S} can be constructed as the weighted sum of individual continuation values \mathcal{V}^f and \mathcal{V}^m , and all assets are split equally between the two household members.

Bequest motive. A retired individual who does not survive into the following period potentially derives utility from leaving his estate (“bequest motive”). A key feature of my model is that the taste to bequeath wealth depends not only on the size of estates left, but also on the presence of descendants. More specifically, I make the following two assumptions: (a) Individuals only have a bequest motive if they have descendants, i.e. $\mathcal{Z}^f(a', 0) = \mathcal{Z}^m(a', 0) = \mathcal{Z}(a', 0) = 0$; (b) Individuals with descendants are fully altruistic towards them, i.e. their bequest function is equal to the value function of the inheritor.¹³ Since the information set for all agents is restricted to the mere presence of descendants rather than their identity – i.e. whether they are female/male, married/single, college-educated or not etc. – the bequest utility corresponds to the expected value function for a generic newborn agent, which can be constructed as the weighted average of expected value functions for agents of both genders, education levels and marital statuses. As stated above, I assume that married couples with descendants leave their estate to two entrants in equal shares. Single agents pool their estates with a randomly selected single agent of opposite gender (second “parent”) and leave the pooled estate in equal shares to two entrants. Since single agents do not know the quantity of assets contributed by the other parent before dying, they form rational expectations based on the actual distribution of assets.¹⁴

3.3 Stationary Equilibrium

To keep notation as compact as possible, I will define the state space for all types of households as $\mathbf{X} \equiv \mathbf{S} \times \mathcal{A} \times \mathcal{D}$, where I arbitrarily impose $\mathbf{S} = S \times \{0\} \times \{0\}$ if the household is single.¹⁵ The Borel algebra generated by an appropriate family of subsets of \mathbf{X} is denoted by \mathcal{B} . Let $\nu(B)$ be a probability measure describing the mass of households in $B \in \mathbf{X}$, where $\nu(B)$ is defined on \mathcal{B} . Denote by $P(a, \mathbf{s}, d, B)$ the probability that a household at state (a, \mathbf{s}, d) will transit to a state that lies in $B \subset \mathcal{B}$ in the next period. The transition function P can be constructed as

$$P(a, \mathbf{s}, d, B) = \int_{\mathbf{s}' \in B_{\mathbf{s}}} \mathcal{I}_{a'(a, \mathbf{s}, d) \in B_a} \Omega(\mathbf{s}, d\mathbf{s}'),$$

where \mathcal{I} is an indicator function taking on a value of 1 if its argument is true and 0 otherwise, $\Omega(\mathbf{s}, \mathbf{s}' \in B_{\mathbf{s}})$ is the probability that the exogenous state next period belongs to $B_{\mathbf{s}} \subseteq \mathbf{S}$, and $B_{\mathbf{s}}$

¹³The assumption of full altruism has found some support in recent studies, e.g. Castañeda et al. (2003) and Fuster et al. (2008).

¹⁴See Appendix III for a formal derivation of the bequest function for single and married households.

¹⁵As a consequence, the state \mathbf{s} characterizes (i) the labor efficiency of all household members, (ii) whether the household is in working age or retirement; and (iii) whether the household is single or married.

and B_a are the projections of \mathcal{B} on \mathcal{S} and \mathcal{A} respectively.

Definition: A stationary recursive competitive equilibrium with incomplete markets in this economy is a list of functions $\{\mathcal{V}^f, \mathcal{V}^m, \mathcal{V}, c^f, c^m, l^f, l^m, a', K, L^f, L^m\}$, a measure of households ν , a set of prices $\{r, w^f, w^m\}$ and a government policy $\{\tau, \tau_p, b\}$ such that:

- 1) For given prices, taxes and benefits, $\mathcal{V}^f, \mathcal{V}^m$ and \mathcal{V} solve (7) – (8), and $c^f(a, \mathbf{s}, d), c^m(a, \mathbf{s}, d), l^f(a, \mathbf{s}, d), l^m(a, \mathbf{s}, d)$ and $a'(a, \mathbf{s}, d)$ are the associated policy functions.
- 2) For given prices, K, L^f and L^m satisfy the firm's first-order conditions (4) – (6).
- 3) Aggregate factor inputs are generated by the policy functions of the agents:

$$K = \int_{\mathbf{X}} a'(a, \mathbf{s}, d; \mu) d\nu, \quad (9)$$

$$L^f = \int_{\mathbf{X}} e^f [1 - l^f(a, \mathbf{s}, d)] d\nu, \quad (10)$$

$$L^m = \int_{\mathbf{X}} e^m [1 - l^m(a, \mathbf{s}, d)] d\nu. \quad (11)$$

- 4) The time-invariant stationary distribution ν is determined by the transition function P as

$$\nu(B) = \int_{\mathbf{X}} P(a, \mathbf{s}, d, B) d\nu \quad \text{for all } B \in \mathcal{B}. \quad (12)$$

- 5) The government budget is balanced:

$$\int_{\mathbf{X}} \tau(y) d\nu + \tau_p [\bar{w}^f L^f + \bar{w}^m L^m] = \int_{\mathbf{X}} b d\nu. \quad (13)$$

4 Parameterization and Calibration

The length of a period is set to one year. The model contains ten preference parameters, $\beta, \varphi_c^f, \varphi_c^m, \varphi_l^f, \varphi_l^m, \sigma^f, \sigma^m, \gamma^f, \gamma^m$ and μ , and five demographic parameters, $\phi^D, \phi^R, \psi, q_c^f, q_c^m$. There are seven technology parameters, $\alpha, \lambda, \delta, \rho, \sigma_\epsilon, \varrho, \xi_{coll}$, four matching probabilities, $q_{0,0}^f, q_{0,\xi_c}^f, q_{\xi_c,0}^f, q_{\xi_c,\xi_c}^f$, and two probabilities of having a descendant, π^S and π^M . Finally, I have to specify thirteen parameters describing the policy of the government, $\kappa_0^S, \kappa_1^S, \kappa_0^M, \kappa_1^M, b_0^f, b_{\xi_c}^f, b_0^m, b_{\xi_c}^m, b_{0,0}, b_{\xi_c,0}, b_{0,\xi_c}, b_{\xi_c,\xi_c}, \tau_p$. Overall, the model is characterized by 41 parameters.

I will set 22 parameters using a *a priori* information. These include all demographic and technology parameters, the four matching probabilities, the two probabilities of having a descendant and the preference parameters $\varphi_c^f, \sigma^f, \sigma^m$ and μ . The remaining parameter values are chosen to match an equal number of moment conditions in the model, i.e. they solve a system of 19 equations in 19 unknowns.

Table 3. Baseline parameters

Description	Parameter	Value	Description	Parameter	Value
Discount factor	β	0.966	Gender premium	λ	0.562
Female risk aversion	σ^f	1.75	College premium	ξ_{coll}	0.560
Male risk aversion	σ^m	1.5	Wage persistence	ρ	0.914
Utility weight (f)	φ_c^f	1	Wage volatility	σ_ϵ	0.206
Utility weight (m)	φ_c^m	2.20	Cross-spouse correlation	ϱ	0.15
Utility weight (f)	φ_l^f	3.91	Fraction with college (f)	q_c^f	0.396
Utility weight (m)	φ_l^m	1.40	Fraction with college (m)	q_c^m	0.394
Regulates Frisch elasticity	γ^f	1.8	Matching probability	$q_{0,0}^f$	0.566
Regulates Frisch elasticity	γ^m	4.0	Matching probability	q_{0,ξ_c}^f	0.201
Probability of retiring	ϕ^R	1/35	Matching probability	$q_{\xi_c,0}^f$	0.237
Probability of dying	ϕ^D	1/20	Matching probability	q_{ξ_c,ξ_c}^f	0.666
Probability of divorce	ψ	0.0092	Probability descendant	π^S	0.34
Capital share	α	0.36	Probability descendant	π^M	0.89
Capital depreciation rate	δ	0.1	Pareto weight	μ	0.5

Demographics. My strategy is to set demographic parameters and matching probabilities such that the resulting composition of the population in the model mimics the actual population in the United States. Life-cycle parameters are determined as follows: individuals enter the working-age stage when they are 25, and they retire and die stochastically. I target the expected durations of their working lives and retirement to be 35 years and 20 years respectively. The shares of newborn females and males with college education are set to their empirical counterparts of 39.6 percent and 39.4 percent respectively. Furthermore, I target the empirical population shares of married couples with all four combinations of education mixes and the share of marriages leading to a divorce. From CPS data for the year 2009, I estimate that 50 percent of all households are married. Across all married households, there are 32.1 percent where both spouses are college-educated, 42.2 percent where neither spouse has college education, 11.4 percent where only the husband has been college-educated and 14.2 percent where only the wife has college education. As for the divorce rate, I target a 40-percent probability that married couples divorce before dying.¹⁶ Given the demographic structure of the model, these targets uniquely pin down the four matching probabilities and the divorce probability, which implies that they can be calibrated externally (see Appendix II for a formal derivation). Finally, the probabilities of having descendants are set to match the population shares of single (69%) and married (89%) households with children.

¹⁶Divorce rates in the U.S. are typically estimated to be 40-50 percent with higher rates for teenage marriages, in particular in the first 5-10 years. Since agents enter my model at the age of 25, I choose a rate of 40 percent.

Preferences. Instantaneous utility functions for females and males are parameterized as follows,

$$U^g(c, l) = \varphi_c^g \frac{c^{1-\sigma^g} - 1}{1 - \sigma^g} + \varphi_l^g \frac{l^{1-\gamma^g} - 1}{1 - \gamma^g} \quad \text{for } g = f, m, \quad (14)$$

where φ_c^g and φ_l^g are parameters and σ^g is the coefficient of relative risk aversion of an individual of gender g . I normalize φ_c^f to 1, which is equivalent to dividing both instantaneous utility functions by this parameter. Non-gender-based estimates of the average coefficient of relative risk aversion between 1 and 3 are common. When gender is taken into account, females are found to be more risk-averse than males. I set individual preferences for risk to $\sigma^f = 1.75$ and $\sigma^m = 1.5$. Estimates for males' Frisch elasticity of labor supply in the presence of potentially binding borrowing constraints range from 0.2 to 0.6 (see Domeij and Flodén 2006). Blundell and MaCurdy (1999) find that for females this elasticity is 3-4 times larger than for males. I target values of 0.42 and 1.25 for males and females, respectively. Utility weights are fixed to align with estimates for the amount of time people spend on market work. Specifically, I target median hours worked by single females (34.3 %), married females (30.4 %) and married males (38.0 %) as fractions of their discretionary time.¹⁷ As is well-known in the literature, the subjective discount factor β can be used to match a capital-output ratio of 3. Finally, the female's Pareto weight in married households is set at $\mu = 0.5$.

Labor efficiency. The stochastic process for labor productivity is modeled as a mapping from observed distributions of hourly wages net of fixed heterogeneity. Denote by ω_t^i the log hourly wage of individual i at time t and specify its evolution as

$$\omega_t^i = \beta_0^i + \beta_1 x_t^i + \tilde{z}_t^i + \iota_t^i, \quad \text{with } \iota_t^i \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_\iota^2), \quad (15)$$

where β_0^i represents an unobserved fixed effect, x_t^i is a vector of observable characteristics such as age, gender and education, ι_t^i reflects measurement error and \tilde{z}_t^i is the time-varying component of an individual's log wage which corresponds to z_t^i in the model and evolves as posited in equation (3). Following Heathcote et al. (2010), I assume that women and men face the same stochastic process for \tilde{z} . This assumption mitigates the selection bias that may be caused by the unobservability of wages for non-working individuals: the reason is that the wage process specified above can be estimated using only data for men, for whom selection is a minor concern. The two parameter values characterizing the labor productivity process (3) are set as in Flodén and Lindé (2001). Using the same specification, these authors estimate $\rho = 0.914$ and $\sigma_\epsilon = 0.206$ from PSID data. In addition, I allow for a correlation structure of temporary shocks within married households. Following Heathcote et al. (2010), I target a cross-spouse correlation for temporary shocks of 0.15.¹⁸ As for the permanent component of labor productivity, I impose that ξ can take

¹⁷These estimates are based on the 2010 wave of the Current Population Survey. I make the assumption that the disposable daily time endowment is 15 hours.

¹⁸See Hyslop (2001).

Table 4. Fiscal policy parameters

Description	Parameter	Value	Description	Parameter	Value
Tax function	κ_0^S	0.2033	Retirement benefits	b_0^m	0.1267
Tax function	κ_1^S	0.0497	Retirement benefits	$b_{\xi_c}^m$	0.1585
Tax function	κ_0^M	0.1740	Retirement benefits	$b_{0,0}$	0.2143
Tax function	κ_1^M	0.0733	Retirement benefits	$b_{\xi_c,0}$	0.2452
Payroll tax rate	τ_p	0.0480	Retirement benefits	b_{0,ξ_c}	0.2386
Retirement benefits	b_0^f	0.1060	Retirement benefits	b_{ξ_c,ξ_c}	0.2675
Retirement benefits	$b_{\xi_c}^f$	0.1341			

on one of two values in $\Xi = \{0, \xi_c\}$, where ξ_c captures the skill wage premium of college-educated individuals. Using CPS data, I estimate a ratio between average wages of college-educated individuals and non-college educated individuals of 1.75. The value of ξ_{coll} is set to match this target.

Technology. Using estimates for the annual capital depreciation rate and the capital share of income, I set $\delta = 0.1$ and $\alpha = 0.36$, which are both standard values in the macro literature. The parameter λ , which characterizes the gender wage gap, is set to match a ratio between average female and male hourly wages of 0.78, as estimated from the CPS.

Government. Following Guner, Kaygusuz and Ventura (2008), I construct income tax functions for single and married households based on estimates for effective taxes paid as a function of reported income. Using IRS data, Guner, Kaygusuz and Ventura (2008) estimate mean income and the average tax rate corresponding to each tabulated income bracket, and then fit a non-linear equation to the data for single and married households separately. The respective income tax functions are parameterized as follows:

$$\tau^S(y) = [\kappa_0^S + \kappa_1^S \log(y)] y \quad (16)$$

$$\tau^M(y) = [\kappa_0^M + \kappa_1^M \log(y)] y. \quad (17)$$

The coefficients they estimate are obtained by normalizing average income in each income bracket by mean household income. Hence, I have to adjust them appropriately by taking into account mean household income in my benchmark economy. Retirement benefits are calibrated by implementing a version of the U.S. Social Security system into my model economy. An exact implementation would require keeping track of each household's lifetime earnings history, which is computationally expensive. Instead, I employ a simpler version where pensions are a function of average earnings during working life for each household type. For instance, retirement benefits for married households with two college-educated spouses are calculated on the basis of average labor earnings by wives and husbands in college/college-households. Since the current distribution of

pension payments partly reflects changes in labor supply patterns over the last decades, I make an additional adjustment by using actual Social Security income data from the 2010 Current Population Survey (see Section 3 in Appendix II for a detailed description of the calibration of retirement benefits). Finally, the linear payroll tax rate τ_p is simply set to balance the budget of the government.¹⁹

5 Results

5.1 The Benchmark Model

As a first step, I will investigate whether my calibrated benchmark model is consistent with the cross-sectional distributions of earnings, income and wealth in the U.S. economy. To this end, I compute a selected variety of distributional statistics in my benchmark economy and compare them to their empirical counterparts. As a second step, I will ask whether my model does a significantly better job of accounting for the data than a more parsimonious model, which lacks some of the features in the benchmark framework.

Table 5 displays a summary of my results. For each of the three variables of main interest – labor earnings, income and net worth –, I contrast five distributional statistics of my benchmark economy with their corresponding empirical values. The first three statistics pertain to the shape of the respective distribution as measures of cross-sectional inequality, whereas the final two statistics directly address the marriage gap in means and medians. Table 5 suggests that the benchmark model does a good job of accounting for the salient features of the data. Firstly, the model succeeds to generate a degree of dispersion that is in accordance with the U.S. distributions of earnings, income and wealth, perhaps with the exception of the respective upper tails where the model slightly understates the degree of concentration. Secondly, the benchmark model generates a positive marriage gap in means and medians for all three variable of interest. In the following, I will discuss these results in more detail.

A closer look at the first four rows in Table 5 reveals that the simulated economy does a good job of accounting for the cross-sectional distribution of labor earnings. At the upper tail, the 5% households with the highest income from labor earnings earn 21 percent of total labor income, which is fairly close to the empirical value of 27.6 percent. The 40% households with the lowest income from labor earnings earn only 0.3 percent of total labor income. The discrepancy to the empirical value of 4.7 percent partially reflects the assumption that households are not allowed

¹⁹The implied value for τ_p in my benchmark model is 4.8%, which is lower than the 2007 US payroll tax rate of 15.3%. An alternative calibration strategy, which is to set τ_p to its historical value and assume that the resulting budget surplus is used for government consumption, yields very similar results.

Table 5. Main results

	Bottom 40%	Top 5%	Gini	Δ^{Mean}	Δ^{Median}
Labor earnings					
Data	4.7	27.6	0.59	+ 49.4	+ 83.3
Benchmark model	0.3	21.4	0.60	+ 28.1	+ 85.9
Restricted model	0.3	19.9	0.59	+ 23.0	+ 82.7
Total income					
Data	11.7	27.0	0.50	+ 26.8	+ 15.6
Benchmark model	16.0	17.2	0.40	+ 13.5	+ 31.9
Restricted model	18.8	15.5	0.37	+ 6.7	+ 32.8
Wealth					
Data	1.8	43.6	0.72	+ 33.5	+ 50.2
Benchmark model	2.4	26.4	0.64	+ 20.7	+ 89.0
Restricted model	3.4	22.2	0.59	- 18.6	+ 5.8

Note: Δ^{Mean} and Δ^{Median} measure the marriage gap in means and medians, respectively. (cf. equation (1))

to work anymore once they are retired (36.4 percent of the population). As for total income, the model does a very good job of accounting for the bottom tail. The fact that it slightly underpredicts the degree of dispersion at the upper tail is related to the shape of the wealth distribution. The reason is that there are not sufficiently many households accumulating extreme levels of wealth and, hence, having large incomes from capital gains. By contrast, the model is consistent with the lower tail of the wealth distribution: in the benchmark economy, the poorest 40% households hold only 2.4 percent of total wealth, which comes very close to the empirical value of 1.8 percent.

To what extent can the benchmark economy account for the inequality between married and single individuals, as reflected by the marriage gap? In the model, married individuals earn on average 28.1 percent more labor income, their average total income is 13.5 percent higher, and they hold 20.7 percent more assets than singles. While the ordering across the three variables is the same as in the data, the model does not quite match the magnitude of the empirical values. For instance, the marriage gap in mean income predicted by the model is only about half as large as the corresponding empirical value. When looking at median values, the marriage gaps in the model are generally larger than their empirical counterparts. For instance, the median married individual in the model holds more than 89 percent more assets than the median single individual, as opposed to an empirical value of 50 percent. On the other hand, the marriage gap in median labor earnings in the model, 85.9 percent, is very close to the empirical value of 83.3 percent. The performance of the model along the life-cycle dimension and across the wealth distribution are

discussed in more detail in Sections 5.3 and 5.4.

The benchmark model generates a positive marriage gap in per-capita means and medians for all three variables of interest. How do these findings compare to the implications of a more parsimonious framework that lacks some features introduced in the benchmark model? Specifically, what is the role played by directed bequests, differential income taxation and permanent abilities? To answer this question, I deactivate these features and solve a *restricted* version of the model (a precise description will be provided in the next section). As can be seen in Table 5, the restricted model performs considerably worse at accounting for the data in almost all dimensions. Mean earnings and income display a marriage gap that is considerably lower than in the benchmark framework. More importantly, the restricted model fails entirely to generate a positive wealth gap – on the contrary, it predicts that married individuals are almost 19 percent poorer than singles. In other words, the restricted model is not consistent with the household saving behavior that we observe in the data.

It is important to note that the restricted model closely resembles standard general equilibrium theories of inequality in a life-cycle setting (e.g. Castañeda et al. (2003) and De Nardi (2004)). While some of these models do a better job of matching the bottom and top tails of the income and wealth distribution, most of them abstract from the distinction between singles and couples and, hence, cannot be tested along this key dimension. The virtue of the restricted model is to show that augmenting a standard model in the most straightforward way, i.e. by introducing two household types, leads to counterfactual predictions. The benchmark model does a significantly better job of explaining the emergence of a positive per-capita wealth and income gap. Hence, it is natural to ask: which of its additional features are crucial and how do they contribute to the overall performance of the model?

5.2 The Marriage Gap: A Decomposition Analysis

In an attempt to shed more light on this question, I simulate a series of alternative models and compare their predictions for the marriage gap in earnings, income and wealth. The main objective is to evaluate the relative importance of three factors that distinguish the benchmark model from the restricted model: (i) the distinction between intentional and accidental bequests; (ii) the tax bonus for married households as implied by the different income tax functions; and (iii) the role of permanent abilities in conjunction with positive assortative mating. Starting from the benchmark model, I shut these three channels down one by one and recalibrate each model appropriately. The restricted model then refers to the setup in which all three channels are deactivated.

To be more precise, the three alternative models are specified as follows.²⁰ First, I eliminate the

²⁰Appendix IV contains a full list of parameter values for all four model economies.




Figure 1: Effective tax functions

bequest channel by making the assumption that all bequests are accidental (M1). Since no household draws utility from leaving their estates to further generations, the presence of descendants becomes irrelevant. Next, I remove the tax bonus implied by the estimated effective tax functions for single and married households (M2). Figure 1 makes this point more explicit: the solid blue line plots average tax rates as a function of total household income for married couples using the estimated coefficients. By contrast, the bold red line plots average tax rates as a function of the aggregated income of two single households earning exactly the same income.²¹ For my second counterfactual model, I shift the tax curve for married couples upwards such that it crosses the curve for single households at median household income (dotted black line).²² Finally, I deactivate the role of education by removing any heterogeneity in permanent abilities, i.e. $\xi_{coll} = 0$ (M3).²³

Table 6 presents my results. All models predict a positive marriage gap for labor earnings and income, although the extent to which they are consistent with the respective empirical values varies considerably. Between the restricted model and the data, there is a gap of 26.4 percentage points

²¹For instance, a married couple earning 80,000 \$ annually pays an average tax rate of 10.8 percent, whereas two singles with 40,000 \$ income each pay 12.4 percent. Note that the tax function for two singles is a lower bound, because any combination of two singles with different individual incomes yields a higher effective tax rate.

²²While this assumption does not capture all benefits for married households implied by the U.S. tax code, it is instructive to measure the pure level effect of lower taxes. In section 5.4, I run a different experiment by eliminating joint tax filing for married couples altogether.

²³I have experimented with a series of alternative models where the three channels are deactivated in a different order; quantitatively, my results are very similar.

Table 6. The marriage gap: Comparison of alternative models

	Labor earnings		Total income		Wealth	
	Δ^{Mean}	Δ^{Median}	Δ^{Mean}	Δ^{Median}	Δ^{Mean}	Δ^{Median}
Data	+ 49.4	+ 83.3	+ 26.8	+ 15.6	+ 33.5	+ 50.2
Benchmark model	+ 28.1	+ 85.9	+ 13.5	+ 31.9	+ 20.7	+ 89.0
M1: No bequest motive	+ 28.8	+ 87.6	+ 12.6	+ 32.0	+ 4.4	+ 54.4
M2: No beq/No tax bonus	+ 29.9	+ 90.8	+ 12.6	+ 32.5	- 4.5	+ 36.5
M3: Restricted model	+ 23.0	+ 82.7	+ 6.7	+ 32.8	- 18.6	+ 5.8

for average labor earnings and a gap of 20.1 percentage points for average income. The benchmark model closes these gaps by roughly 5 and 7 percentage points respectively. The main channel driving this result is heterogeneity in permanent abilities as college-educated individuals exhibit a higher propensity to marry. It is important to note, however, that this result is not implied by my calibration strategy. Since labor earnings are composed of both wages and hours worked, individual and aggregate outcomes are determined by the decision problems of all household types and, more specifically, policy functions for labor supply. The model correctly accounts for the fact that college-educated individuals work less, but they still earn more income due to higher wages. The decomposition analysis also suggests that the tax bonus and the bequest motive seem to be less important determinants for the marriage gap in earnings and income.

As for per-capita wealth, there is a striking dissimilarity across the various models. The restricted model predicts that married agents are on average 18.6 percent poorer than singles, which amounts to a gap of 52 percentage points to the empirical value. The three additional features introduced in the benchmark model can account for more than 75% of this gap. Across the three channels, including stronger dynastic links in households with descendants contributes the largest share with roughly 40% (or 16 percentage points). The intuition is that most married households decumulate wealth at a slower pace late in life so as to leave more estate to the following generation. Factoring in the effective income tax bonus implied by the estimated tax functions for married and single households further contributes significantly to closing the gap (the relative contribution is 22% or 9 percentage points). The reason is that lower income taxes yield a higher permanent

disposable household income; this implies that married households accumulate more buffer-stock assets in order to reach their target wealth-to-permanent-disposable-income ratio.²⁴ Finally, the inclusion of permanent abilities into the model yields a relative contribution of roughly 38% (or 14 percentage points). College-educated individuals have a higher permanent income, so they will target a larger buffer stock of wealth to insure against labor income risk. Moreover, the US Social Security system implies that retirement benefits are a concave function of earnings. Hence, in order to ensure a large average replacement rate, college-educated people accumulate more private retirement saving during working age. Since college-educated people have a higher propensity to get married, both of these effects contribute to the role of education as the third key channel in the model.

Why does the restricted model yield counterfactual predictions for the marriage gap in wealth? The main driving force for households to accumulate assets in this model is the precautionary motive: since asset markets are incomplete, the presence of borrowing constraints implies that agents cannot fully insure against labor income risk. Single households use precautionary saving and precautionary labor supply to insure against this risk. Married individuals, additionally, have access to intrahousehold risk sharing, which enables them to share labor income risk efficiently through income pooling and spousal labor supply. As a consequence, they save less for precautionary reasons.²⁵ This mechanism yields a large negative wealth gap in the restricted model.

As for the marriage gap in medians, all four models yield similar predictions for labor earnings and total income. The results for labor earnings are slightly harder to interpret, because they are considerably influenced by the fact that a large share of the population has zero earnings. As discussed in Section 5.4, the benchmark model does a very good job in accounting for the median marriage gap in earnings when evaluated along the life-cycle dimension. By contrast, the median gap in wealth varies considerably across the four models: while in the benchmark model the median married individual holds almost 90 percent more assets than the median single individual, the restricted model predicts almost no gap. Note that decomposing the contribution of each of the three channels yields similar numbers as above. Compared to the empirical value, the benchmark model overpredicts the median wealth gap. In order to shed more light on the wealth gap between married and single households, I now turn to evaluating the performance of the model along the asset distribution.

²⁴See Carroll (1997).

²⁵The lower is the cross-spousal correlation of temporary shocks, the stronger is this effect. For instance, if the correlation was 0 instead of 0.15, the gap to the data would widen by another half a percentage point.

Table 7. Fraction of liquidity-constrained households

	Data	Model
All Households	20.1 %	25.7 %
Married Households	15.0 %	19.9 %
Single Households	27.7 %	31.2 %

5.3 The Marriage Gap in Wealth: Model vs. Data

A recurrent theme in this paper is that married households accumulate more wealth than single households. If asset markets are incomplete, an important saving motive is to build up a buffer stock of wealth so as to insure against adverse labor market shocks and smooth consumption intertemporally. In the previous section I have argued that the precautionary saving motive puts a strong downward pressure on the marriage gap in wealth since married households have access to intrahousehold risk sharing as an additional source of insurance. In light of this result, I now turn to the question whether the model can account for the fraction of liquidity-constrained single and married households in the US. Moreover, I will analyze how the marriage gap in wealth shapes along the asset distribution and compare the model’s prediction to the empirical pattern.

Following Zeldes (1989), I define a household as liquidity-constrained if its current non-housing net worth is lower than two months of income, where I assume that housing constitutes half of each household’s total net worth.²⁶ These assumptions imply a threshold capital-income ratio of 3, below which a household is classified as liquidity-constrained. Table 7 shows that, according to this indicator, about 20 percent of all households in the US are liquidity-constrained.²⁷ Moreover, there is a considerable difference between the sample of single and married households: while only 15 percent of married couples face liquidity constraints, the corresponding proportion of single households is almost twice as large (27.7 percent).

Can the benchmark economy account for these differences? The share of liquidity-constrained households predicted by the model, 25.7 percent, is slightly larger than the empirical value of 20.1 percent (Table 7, last column). The model is also very successful in replicating the distinct shares of single and married households: the proportion of liquidity-constrained singles is 11 percentage points greater than the proportion of liquidity-constrained married couples, which matches almost exactly the empirical difference. These findings suggest that married households

²⁶In principle, it would be possible to use SCF portfolio data to identify each household’s share of total wealth held in housing. Since I do not model housing as an asset in my benchmark economy, I abstract from this additional source of heterogeneity. Using 1998 SCF data, Díaz and Luengo-Prado (2010) find that the proportion of total wealth in housing represents between 96 percent for poor households and 27 percent for rich households. I pick 50 percent as the liquidable share. Choosing a higher value would simply scale up all numbers in Table 8 equally.

²⁷Jappelli, Pischke and Souleles (1998) and, more recently, Gorbachev and Dogra (2008) obtain similar values.

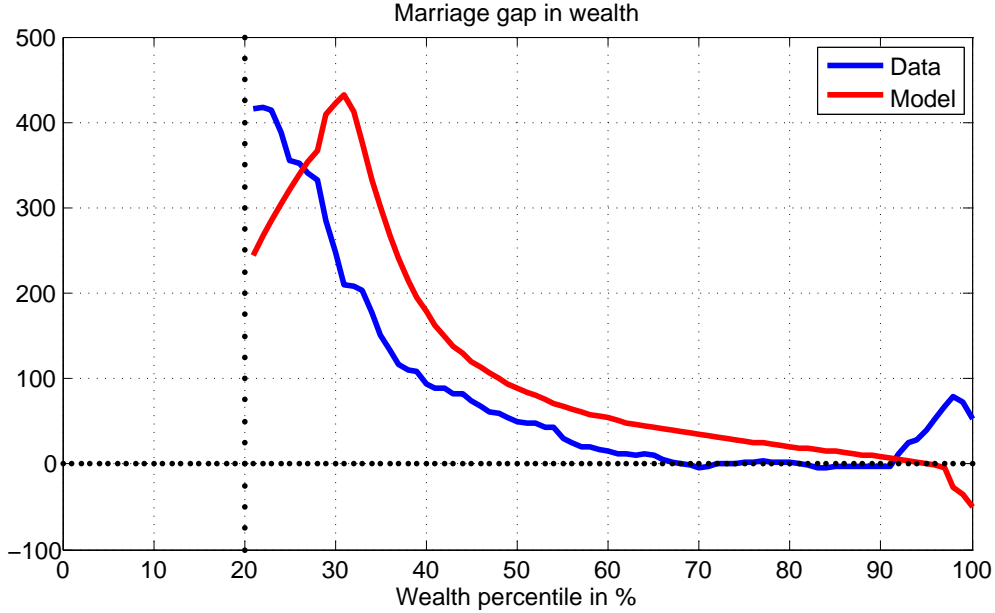


Figure 2: Marriage gap in wealth by percentile

find themselves in financial turmoil at a considerably lower frequency than single households. Since uninsurable income risk and the presence of borrowing constraints are widely acknowledged as crucial determinants of inequality, the ability to account for the fraction of liquidity-constrained single and married households is an important success of the model.

To further investigate this result, I now turn to discuss the shape of the marriage gap in wealth along the asset distribution. Specifically, I compute the percentage differential between married and single households for each percentile of the respective distribution in the data and in my model (Figure 2).²⁸ Focusing on the data, two features stand out. First, the marriage gap in wealth is positive along the entire distribution. Second, the gap is very large for households in the bottom half of the distribution and then declines sharply. For instance, married individuals at the 30th percentile of the wealth distribution of married households hold about three times more assets than single individuals at the 30th percentile of the wealth distribution of single households. By contrast, there is almost no discrepancy at the 70th percentile. Interestingly, there is a jump at the very top, which suggests that the majority of extremely rich people is married.

How well can the benchmark model account for the data? The curve computed from the model yields a fairly similar shape: the marriage gap in wealth is positive along the entire distribution –

²⁸By way of example, in the data the median married individual holds about 50 percent more assets than the median single individual (cf. Table 1, last column). I omit plotting the lower 20 percent of the distribution, because some households in the data hold very few or even negative assets. The vast majority of single and married households in this part of the distribution classify to be liquidity constrained.

Table 8. The marriage gap along the life cycle

	Δ^{Mean}			Δ^{Median}		
	Data	Benchmark	Restricted	Data	Benchmark	Restricted
Labor earnings						
All households	+ 49.4	+ 28.1	+ 22.9	+ 85.6	+ 85.9	+ 80.2
Working age	+ 25.6	+ 6.6	+ 2.0	+ 13.8	+ 13.1	+ 10.1
Total income						
All households	+ 26.8	+ 13.5	+ 6.5	+ 15.6	+ 31.9	+ 15.6
Working age	+ 17.5	+ 6.3	- 1.0	+ 10.8	+ 11.5	+ 6.3
Retirement age	+ 43.7	- 2.8	- 2.9	+ 9.5	- 2.3	- 10.8
Wealth						
All households	+ 33.5	+ 20.7	- 18.8	+ 50.2	+ 89.0	+ 8.5
Working age	+ 38.4	+ 2.0	- 29.7	+ 116	+ 16.1	- 25.8
Retirement age	+ 50.8	+ 47.9	- 10.2	+ 7.6	+ 173	+ 98.2

with the exception of a small group of very rich households – and, moreover, it is larger for poorer households. Note that there is an increasing section below the 30th wealth percentile. This part of the distribution is characterized by a large fraction of liquidity-constrained single households, whereas most married households are unconstrained. The curve reaches its peak around the 30th percentile and then starts to decline again, consistent with the data. Overall, these findings give rise to the notion that married individuals tend to hold more assets than singles, even when conditioned on their position in the wealth distribution. This disparity is particularly large for households in the bottom half of the distribution. This suggests that the marital status plays a stronger role for poor and middle-class households.²⁹

5.4 Along the life cycle

While the benchmark economy successfully generates a positive marriage gap for earnings, income and wealth, it cannot fully account for the respective empirical values (see Section 5.1). For instance, the marriage gap in average earnings predicted by the model is +28.1%, as opposed to an empirical value of +49.4%. Since labor earnings are constrained to be zero during in the retirement, perhaps it is more meaningful to look at the marriage gap during working age. For this reason, I now turn to analyzing the performance of the model along the life-cycle dimension.

Table 8 presents my results. Married individuals in my benchmark economy earn on average 6.6%

²⁹Focusing on employment risk, Ortigueira and Siassi (2012) find that intrahousehold risk sharing has its largest impact among low-wealth households.

Table 9. Policy reform: Joint vs. separate tax filing

	K	L	Y	Gini	Δ^{Earn}	Δ^{Inc}	Δ^{Wealth}	Welfare
Joint filing	1.39	0.250	0.463	0.64	+ 28.1	+ 13.5	+ 20.7	–
Separate filing	1.45	0.251	0.471	0.64	+ 29.6	+ 15.5	+ 43.7	+ 1.2%

Note: Δ denotes the average marriage gap as defined in equation (1). The Gini coefficient is computed for wealth.

more labor income during working age than single individuals. Hence, there is still a significant difference to the empirical value of 25.6%. On the other hand, the marriage gap in median earnings matches the data almost exactly (13.1% vs. 13.8%). This suggests that the disparity in averages between the model and the data may be partially driven by a group of married individuals with very high earnings realizations. This result carries over to the marriage gap in total income: while the benchmark economy underpredicts the average income gap along the life cycle, it does a considerably better job for the median income gap. As for wealth, the model underpredicts the average marriage gap during working age, and it comes very close to matching the data value for the retirement period. Perhaps the biggest failure of the model is the difference in medians during working age and retirement. This suggests that the benchmark model cannot fully account for the distribution of the marriage gap along the wealth distribution (see Section 5.3) and lacks other empirically relevant saving motives, e.g. uninsurable health shocks.

5.5 Policy experiment: Separate tax filing

My results so far indicate that married households benefit substantially from the bonus embodied in the U.S. tax code. The reason is that they face considerably lower effective tax rates than two singles with the same characteristics, in particular, if individual incomes are very different. The decomposition analysis carried out in Section 5.2 suggests that an upward-shift of effective tax rates for married households would reduce the average marriage gap in wealth by about 8 percentage points. This section addresses the same issue in a more direct way by asking: what are the implications of a policy reform that abolishes the possibility to file taxes jointly? I evaluate this question in the context of a hypothetical tax reform in my benchmark economy. The fundamental modification is that all agents, single or married, are now subject to the same effective tax schedule. Specifically, I assume that married individuals face the singles' tax function (τ^S), where income from capital gains is split equally between the two spouses. Additional tax revenues resulting from this reform are redistributed in the form of lump-sum transfers to all agents in the economy. For the purpose of this study, I will focus on long-run effects and leave transitional dynamics aside.³⁰

Table 9 contrasts a selection of aggregate variables under both tax regimes. In the new stationary

³⁰I also abstract from changes in marriage decisions, fertility patterns or educational choices that could potentially result from the reform. These issues are left for future research.

equilibrium after the reform, aggregate capital is 4.1% larger, aggregate labor increases by 0.5%, and total output grows by 1.8%. Perhaps surprisingly, abolishing joint tax filing does not reduce inequality between single and married households. By contrast, the marriage gap in earnings, income and wealth widens even more. This seems to stand in stark contrast to the decomposition analysis carried out in Section 5.2. What is the intuition behind this result? Figure 1 shows that an upward shift of the tax schedule for married households, as imposed in Section 5.2, yields a curve that still looks quite different from the tax function for singles. Importantly, the tax function for singles implies significantly higher tax rates for low income values. Once married households have to file their taxes separately, they face considerably higher rates when hit by adverse income shocks. As a result, they choose to rely more on precautionary savings.

The benchmark model predicts aggregate welfare gains under the new tax regime: expected welfare, measured in consumption equivalent units for a newborn individual, increases by 1.2% (1.0% for females and 1.3% for males). At the bottom of this result lies a fiscal redistribution from married to single households: the tax reform generates 9.2% higher income tax revenues which are redistributed through a lump-sum transfer worth 1.3% of mean per-capita income (roughly \$650 in 2007 dollars). The reform effectively provides more income insurance for single households by shifting the tax burden to married couples who are better insured against idiosyncratic risk.

6 Concluding Remarks

While there has been significant progress in developing macroeconomic models that can account for the cross-sectional distributions of earnings, income and wealth, most of the existing theories do not explicitly acknowledge the role of the marital status. The main contribution of this paper is to derive a number of empirical facts pertaining to inequality across single and married households, and to present a model that is consistent with the salient features of the data.

This paper should be considered as an intermediate step towards a refined understanding of the interaction between marriage and economic inequality. The model has various limitations that are worth exploring in future research. For instance, endogenizing marriage formation and destruction could help to account for the unexplained remainder of the marriage gap in wealth: if poor singles have a harder time to find a spouse than rich singles, and adverse labor market shocks lead to divorces, it is to be expected that the predicted wealth disparity increases even further. Another interesting extension would be to explicitly model economies of scale in two-person households, for instance, by means of a public good. A careful quantitative analysis would, however, be complicated by the fact that there is yet no empirical consensus on the magnitude of such scale economies within the household. Finally, this paper indicates important implications for policy design. It could be interesting to study policies that acknowledge the demand for redistribution by targeting single and married households in different ways.

References

- [1] Aiyagari, S.R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109(3), pp. 659-684.
- [2] Aiyagari, S.R., Greenwood, J. and Guner, N. (2000), “On the State of the Union,” *Journal of Political Economy*, 108(2), pp. 213-244.
- [3] Blundell, R. and MaCurdy, T.E. (1999), “Labor Supply: A Review of Alternative Approaches,” in: Ashenfelter, Q., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 3A.
- [4] Castañeda, A., Díaz-Giménez, J. and Ríos-Rull, J. (2003), “Accounting for the U.S. Earnings and Wealth Inequality,” *Journal of Political Economy*, 111(4), pp. 818-857.
- [5] Carroll, C. (1997), “Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112(1), pp. 1-55.
- [6] Cubeddu, L. and Ríos-Rull, V. (2003), “Families as Shocks,” *Journal of the European Economic Association*, 1(2-3), pp. 671-682.
- [7] De Nardi, M. (2004), “Wealth Inequality and Intergenerational Links,” *Review of Economic Studies*, 71, pp. 743-768.
- [8] Díaz, A. and Luengo-Prado, M. (2010), “The Wealth Distribution with Durable Goods,” *International Economic Review*, 51(1), pp. 143-170.
- [9] Díaz-Giménez, J., Glover, A. and Ríos-Rull, V. (2011), “Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 Update,” Federal Reserve Minneapolis.
- [10] Domeij, D. and Flodén, M. (2001), “The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased,” *Review of Economic Dynamics*, 9, pp. 242-262.
- [11] Flodén, M. and Lindé, J. (2001), “Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?,” *Review of Economic Dynamics*, 4, pp. 406-437.
- [12] Fuster, L., Imrohorglu, A. and Imrohorglu, S. (2008), “Altruism, Incomplete Markets, and Tax Reform,” *Journal of Monetary Economics*, 55(1), pp. 65-90.
- [13] Gorbachev, O. and Dogra, K. (2010), “Evolution of Consumption Volatility for the Liquidity Constrained Households over 1983-2004,” mimeo.
- [14] Greenwood, J., Guner, N. and Knowles, J. (2003), “More on Marriage, Fertility, and the Distribution of Income,” *International Economic Review*, 44(3), pp. 827-862.

- [15] Guner, N., Kaygusuz, R. and Ventura, G. (2008), “Taxation, Aggregates and the Household,” IZA Discussion Papers 3318, Institute for the Study of Labor (IZA).
- [16] Guner, N. and Knowles, J. (2004), “Marital Instability and the Distribution of Wealth,” Working Paper, Department of Economics, University of Pennsylvania.
- [17] Heathcote J., Perri, F. and Violante, G.L. (2010), “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006,” *Review of Economic Dynamics*, 13(1), pp. 15-51.
- [18] Heathcote J., Storesletten, K. and Violante, G.L. (2009), “The Macroeconomics of the ‘Quiet Revolution’: Understanding the Implications of the Rise in Women’s Participation for Economic Growth and Inequality,” mimeo.
- [19] Heathcote J., Storesletten, K. and Violante, G.L. (2010), “The Macroeconomic Implications of Rising Wage Inequality in the United States,” *Journal of Political Economy*, 118(4), pp. 681-722.
- [20] Hintermaier T. and Königer, W. (2011), “On the Evolution of the US Consumer Wealth Distribution,” *Review of Economic Dynamics*, 14(2), pp. 317-338.
- [21] Hong, J. and Ríos-Rull, V. (2007), “Social Security, Life Insurance and Annuities for Families,” *Journal of Monetary Economics*, 51(1), pp. 118-140.
- [22] Huggett, M. (1996), “Wealth Distribution in Life-Cycle Economies,” *Journal of Monetary Economics*, 38(3), pp. 469-494.
- [23] Hyslop, D. (2001), “Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings,” *American Economic Review*, 91(4), pp. 755-777.
- [24] Jappelli, T., Pischke, J. and Souleles, N. (1998), “Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources,” *The Review of Economics and Statistics*, 80(2), pp. 251-262.
- [25] Krusell, P. and Smith, A. (1998), “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), pp. 867-896.
- [26] Laitner, J. (2001), “Secular Changes in Wealth Inequality and Inheritance,” *Economic Journal*, 111(474), pp. 691-721.
- [27] Ortigueira, S. and Siassi, N. (2012), “How Important is Intra-Household Risk Sharing for Savings and Labor Supply,” Working Paper.
- [28] Zeldes, S. P. (1989), “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy*, 97(2), pp. 305-346.

7 Appendix I: Data Sources and Variable Definitions

7.1 Survey of Consumer Finances (SCF)

The analysis is based upon the 2007 wave of the Survey of Consumer Finances (SCF). The SCF is conducted in three-year intervals and gathers detailed information on the financial situation of families in the United States. The survey is designed to obtain a sufficiently large and unbiased sample of wealthy households and provides appropriate weighting schemes to adjust for nonrespondents. The primary unit of observation is the household. A household comprises either an economically dominant single individual or a couple (married or living as partners) as well as all other individuals in the household who are financially interdependent with that individual or couple. I classify a household as married if the head of the household is legally married (and not separated), where I follow the SCF convention of defining the head as the male in core couple households. The following two restrictions are applied on the basic sample. First, I exclude all households where the head is less than 25 years old. Second, I exclude the richest 1 percent of married households and the richest 1 percent of single households. The remaining sample consists of 3,580 households; 2,366 of them are classified as married and 1,232 of them are classified as single. I employ the following variable definitions:

Labor earnings. Sum of wages + salaries plus two 66 % of business + self-employment income.

Total income. Sum of all income sources before taxes, i.e. wage income, self-employment income, net asset income, and private and public transfers.

Wealth. Net worth of the household, i.e. value of real and financial assets net of liabilities.

College education. Individual has obtained a college degree (variables x5904 and x6104).

7.2 Current Population Survey (CPS)

The analysis is based upon the March 2010 Supplement of the Current Population Survey (CPS). The CPS is the primary source of labor force statistics in the United States. A representative sample of currently around 60,000 households is interviewed about a set of demographic and labor force questions at a monthly frequency. The Annual Social and Economic Supplement (or ‘March Supplement’) augments the basic survey by a set of more detailed questions on income and is extended by an additional sample of around 34,5000 households.

I choose the family as the basic unit of observation and henceforth refer to it as a household interchangeably.³¹ As a measure of preliminary data cleaning, I drop all households in which there is at least one individual whose hourly wage is lower than half of the federal minimum wage

³¹The primary unit of observation defined in the CPS is the “housing unit”, which may include multiple families.

in 2009. In addition, I exclude all households where the head is less than 25 years old. The remaining sample consists of 190,200 individuals organized in 77,368 households.

Next, I divide the main sample into two subsamples and call them married households and single households. A household is defined to be married if its head is currently married, and single otherwise. The definition of the household head can be described as follows. I first check whether there is an individual in the family that is of working age, i.e. between 25 and 60 years old. If so, the head is defined to be the oldest male person of working age or, alternatively, the oldest female person of working age in case there is no working-age male. If there are no individuals of working age in the household, I define the household head to be the oldest adult male (i.e. ≥ 18 years). If there are no adult males in the household, the head is defined to be the oldest adult female.

As a result, the subsample of married households consists of 126,165 individuals organized in 38,593 households. The subsample of single households consists of 64,035 individuals organized in 38,775 households. Labor earnings and total income are defined as for the SCF, where for both variables, respectively, I trim the data by excluding the bottom 0.5 percent of observations in order to deal with the problem of underreporting. Other variables are defined as follows:

Wage. To obtain an individual's hourly wage rate, I divide annual labor earnings by annual hours worked. Annual hours worked are computed as the product of the number of weeks worked and the number of hours worked per week (items *'wkwork'* and *'hrswk'*).

Gender premium. I consider all working-age individuals who participate in the labor force, i.e. who work at least 260 annual hours (this corresponds to an average of 1 hour per working day). This sample consists of 74,792 individuals. The gender premium is computed as the ratio between the average wage rate earned by females and the average wage rate earned by males.

College: Individuals are defined to be college-educated if they have obtained some college degree (value of 41+ in item *'a-hga'*). To compute the college premium, I consider the same sample that is used to compute the gender premium. Then I calculate the ratio between the average wage rate earned by college-educated individuals and the average wage rate earned by non-college educated individuals. To partly control for the gender gap, I do this separately for both genders and then take the average.

8 Appendix II

8.1 Population flows

Denote by n^M , n^f and n^m the respective measures of married, single female and single male households in the population, where the total population size is normalized to $n^M + n^f + n^m = 1$.

Recall that married households die together and that the probability of dying is equal across retired single households of both genders. These assumptions imply that $n^f = n^m$. Denote by n_W^i and n_R^i the measures of working-age and retired households for each household type $i \in \{\mathcal{M}, f, m\}$, where by definition $n_W^i + n_R^i = n^i$. Population flows between different household types can then be described as follows:

Retired married households	– Inflow:	$\phi^R (1 - \psi) n_W^{\mathcal{M}}$
Retired married households	– Outflow:	$\phi^D n_R^{\mathcal{M}} + (1 - \phi^D) \psi n_R^{\mathcal{M}}$
Working-age married households	– Inflow:	newborn married
Working-age married households	– Outflow:	$\phi^R n_W^{\mathcal{M}} + (1 - \phi^R) \psi n_W^{\mathcal{M}}$
Retired single households	– Inflow:	$\phi^R n_W^g + \psi (1 - \phi^D) n_R^{\mathcal{M}} + \psi \phi^R n_W^{\mathcal{M}}$
Retired single households	– Outflow:	$\phi^D n_R^g$
Working-age single households	– Inflow:	$(1 - \phi^R) \psi n_W^{\mathcal{M}} + \text{newborn singles}$
Working-age single households	– Outflow:	$\phi^R n_W^g$

where $g = f, m$. For example, the measure of married households flowing out of retirement after each period is equal to the sum of two masses: the measure of households who die at the end of the period, $\phi^D n_R^{\mathcal{M}}$, and the measure of households who survive but divorce at the end of the period, $(1 - \phi^D) \psi n_R^{\mathcal{M}}$. The measures of newborn single and married households will be described in more detail below.

It is important to note that population measures in a stationary equilibrium must be constant, i.e. inflows and outflows for each household type have to exactly offset each other. For instance, consider the measure of retired married households. Equalizing the first two expressions and using $n_W^{\mathcal{M}} = n^{\mathcal{M}} - n_R^{\mathcal{M}}$ yields

$$n_R^{\mathcal{M}} = \frac{\phi^R (1 - \psi)}{\phi^D + (1 - \phi^D) \psi + \phi^R (1 - \psi)} \cdot n^{\mathcal{M}}.$$

Similarly, one can find an expression for the measure of retired female households:

$$n_R^f = \frac{\phi^R n^f + (1 - \phi^D) \psi n_R^{\mathcal{M}} + \phi^R \psi n_W^{\mathcal{M}}}{\phi^D + \phi^R}.$$

To derive the measures of newborn households, note that the mass of households dying at the end of each period is equal to $\left(\frac{1}{\phi^R} + \frac{1}{\phi^D}\right)^{-1} \equiv \bar{n}$. In a stationary equilibrium, this implies that the total mass of newborn households is equal to \bar{n} . The share of newborn married households is determined by the weighted sum of individual matching probabilities during the household formation stage, $\sum_{i,j} q_{\xi^i, \xi^j}^g q_i^g \equiv \bar{q}$, where q_i^g is the probability that an individual of gender g is born with ability level i , and \bar{q} represents the unconditional probability of being matched. It is then easy to show that the measure of newborn married households in each period is equal to

$\bar{n}\bar{q}/(2-\bar{q})$. Equivalently, the measure of newborn single households is equal to $2\bar{n}(1-\bar{q})/(2-\bar{q})$, where half of them are single females and the other half single males.

8.2 Calibration of demographic parameters

The objective is to pin down values for $q_{0,0}^f, q_{0,\xi_c}^f, q_{\xi_c,0}^f, q_{\xi_c,\xi_c}^f$ and ψ given predetermined values for ϕ^R, ϕ^D, q_c^f and q_c^m and five empirical moments: the population shares of married couples with all four educational combinations and the share of marriages that lead to a divorce. First, note that the unconditional probability for a married working-age couple to be hit by a separation shock before dying is given by $(1-\psi)^{\frac{1}{\phi^R}+\frac{1}{\phi^D}}$. For any combination of ϕ^R and ϕ^D , the single-period divorce probability ψ then determines the targeted life-time divorce probability.

To pin down the matching probabilities, recall from the previous subsection that the measure of newborn married households flowing into working age is given by $\bar{n}\bar{q}/(2-\bar{q})$. In a stationary equilibrium, this number has to equal the corresponding outflow:

$$\bar{n}\bar{q}/(2-\bar{q}) = \left[\phi^R + (1-\phi^R)\psi \right] n_W^{\mathcal{M}} = \left[\phi^R + (1-\phi^R)\psi \right] \left[n^{\mathcal{M}} - n_R^{\mathcal{M}} \right]$$

Plug in the expression for $n_R^{\mathcal{M}}$ as derived above:

$$\begin{aligned} \bar{n}\bar{q}/(2-\bar{q}) &= \left[\phi^R + (1-\phi^R)\psi \right] \left[n^{\mathcal{M}} - \frac{\phi^R(1-\psi)}{\phi^D + (1-\phi^D)\psi + \phi^R(1-\psi)} \cdot n^{\mathcal{M}} \right] \\ \bar{n}\bar{q}/(2-\bar{q}) &= \left[\phi^R + (1-\phi^R)\psi \right] \left[\frac{\phi^D + (1-\phi^D)\psi}{\phi^D + (1-\phi^D)\psi + \phi^R(1-\psi)} \right] n^{\mathcal{M}}. \end{aligned}$$

Given the overall share of married households $n^{\mathcal{M}}$, which is simply the sum of the empirical population shares of married couples with all four educational combinations, one can solve this expression for \bar{q} , the unconditional probability of being matched. It is then straightforward to determine the four conditional matching probabilities $q_{0,0}^f, q_{0,\xi_c}^f, q_{\xi_c,0}^f$ and $q_{\xi_c}^f$ based on \bar{q}, q_c^f and the empirical frequencies of the four educational combinations across married households.

8.3 Calibration of retirement benefits

In order to set numerical values for pension benefits, I implement a version of the U.S. Social Security system into my model economy. In the United States, the Social Security Administration keeps track of a worker's earnings throughout his/her working career. Retirement benefits are then computed as a function of average indexed monthly earnings (AIME), which takes into account the highest 35 years of a worker's earnings. This function is piecewise-linear and concave in its argument, the AIME, where the thresholds are called bend points. Specifically, retirement benefits are computed as the sum of (i) 90 percent of the AIME up to the first bend point (\$656

in 2006), plus (ii) 32 percent of the AIME from the first bend point to the second bend point (\$3,955 in 2006), plus (iii) 15 percent of the AIME in excess of the second bend point. For married households there is an additional regulation, the spousal benefit rule, which states that a married individual is entitled to 50% of his/her spouse's benefits. Hence, married households' benefits can be computed as the maximum between (a) the sum of individual benefits and (b) 1.5 times the highest of the two benefits. In the benchmark model, the spousal benefit rule applies to married households where the husband is college-educated and the wife is not.

An exact implementation of the U.S. Social Security system would require keeping track of each individual's earning history, which is computationally very expensive. Instead, I employ a simplified version that relies on average earnings during working life for a given household type. For instance, retirement benefits for married households with two college-educated spouses are calculated on the basis of average labor earnings by wives and husbands in college/college-households. The bend points are computed as fractions of mean individual pre-tax earnings during working life (\$40,813 in 2006). Since average labor earnings are generally a function of benefits themselves, the calibration proceeds in an iterative procedure until pension benefits are consistent with the Social Security formula.

Since the current empirical distribution of pension payments partly reflects changes in labor supply patterns over the last decades, I make an additional adjustment by using actual Social Security income data from the 2010 Current Population Survey. I select the sample of households where both the head and the spouse – if present – are between 65 and 74 years old and, for all educational compositions, estimate the average replacement rate of Social Security income as a fraction of mean household income. Retirement benefits in the model are then obtained as a weighted average of the value implied by my implementation of the Social Security system as described above, and the empirical value, where the latter carries a relative weight equal to the population share of retired households.

9 Appendix III: The Bequest Function

The bequest motive for single and married agents is closely linked to the expected value function for a newborn individual, which I denote as \mathcal{W} . Based on the likelihood to be selected into gender, education and household type respectively, and given that the initial time-varying labor efficiency component z is drawn from a Normal distribution through ϵ , one can construct \mathcal{W} as a function

of initial assets:

$$\begin{aligned} \mathcal{W}(a) = & \frac{1}{2} \sum_{g=f,m} \sum_{\xi^i \in \Xi} q_{\xi^i}^g \left(\sum_{\xi^j \in \Xi} q_{\xi^i, \xi^j}^g \int \pi^{\mathcal{M}} \tilde{V}^g(a, \mathbf{s}, 1) + (1 - \pi^{\mathcal{M}}) \tilde{V}^g(a, \mathbf{s}, 0) d\epsilon \right. \\ & \left. + \left(1 - \sum_{\xi^j \in \Xi} q_{\xi^i, \xi^j}^g \right) \int \pi^{\mathcal{S}} V^g(a, \mathbf{s}, 1) + (1 - \pi^{\mathcal{S}}) V^g(a, \mathbf{s}, 0) d\epsilon \right), \end{aligned}$$

where $V^g(a, s, d)$ is the value function of a single agents of gender g as defined in (7), and $\tilde{V}^g(a, \mathbf{s}, d)$ is the value function of a married agent of gender g . Note that $\tilde{V}^g(a, \mathbf{s}, d)$ is implicitly defined by the married household's Pareto problem and can be computed accordingly.

The bequest function for married households with descendants, $\mathcal{Z}(a', 1)$, can now be derived as follows. Since there is no intergenerational transmission of abilities, the bequest motive does not depend on individual labor efficiencies and is, thus, identical for both agents in the household. As a result, the bequest function does not depend on relative Pareto weights and simply reads

$$\mathcal{Z}(a', 1) = \mathcal{W}(a'/2).$$

The bequest function for single agents is slightly more involved, because they pool their estates with a randomly selected single agent of opposite gender. Since they do not know the quantity of assets contributed by the other parent before dying, they form rational expectations based on the actual distribution of assets in the population. Denote by ν^g the measure of single households of gender $g \in \{f, m\}$. Then the bequest function for a single female household is given by

$$\mathcal{Z}^f(a', 1) = \int_S \int_A \phi(s) \mathcal{W}\left((a' + a^P(a, s^m, 1)) / 2\right) d\nu^m,$$

where I have renamed the policy function of single male agents to $a^P(a, s^m, 1)$ in order to avoid any confusion with a' . Similarly, the bequest function for a single male household is given by

$$\mathcal{Z}^m(a', 1) = \int_S \int_A \phi(s) \mathcal{W}\left((a' + a^P(a, s^f, 1)) / 2\right) d\nu^f.$$

10 Appendix IV - Parameter Values

Table A1. Parameter values

Description	Parameter	Benchmark	NoBeqMotive	NoBeq/NoTax	Restricted
Discount factor	β	0.966	0.971	0.972	0.975
Female risk aversion	σ^f	1.75	1.75	1.75	1.75
Male risk aversion	σ^m	1.50	1.50	1.50	1.50
Utility weight (f)	φ_c^f	1	1	1	1
Utility weight (m)	φ_c^m	2.20	2.14	2.03	2.12
Utility weight (f)	φ_l^f	3.91	3.77	3.68	4.25
Utility weight (m)	φ_l^m	1.40	1.34	1.31	1.54
Regulates Frisch elasticity	γ^f	1.8	1.8	1.8	1.8
Regulates Frisch elasticity	γ^m	4.0	4.0	4.0	4.0
Probability of retiring	ϕ^R	1/35	1/35	1/35	1/35
Probability of dying	ϕ^D	1/20	1/20	1/20	1/20
Probability of divorce	ψ	0.0092	0.0092	0.0092	0.0092
Capital share	α	0.36	0.36	0.36	0.36
Capital depreciation rate	δ	0.1	0.1	0.1	0.1
Gender premium	λ	0.562	0.562	0.562	0.562
College premium	ξ_{coll}	0.560	0.560	0.560	–
Wage persistence	ρ	0.914	0.914	0.914	0.914
Wage volatility	σ_ϵ	0.206	0.206	0.206	0.206
Cross-spouse correlation	ϱ	0.15	0.15	0.15	0.15
Fraction with college (f)	q_c^f	0.396	0.396	0.396	–
Fraction with college (m)	q_c^m	0.394	0.394	0.394	–
Matching probability	$q_{0,0}^f$	0.566	0.566	0.566	0.821
Matching probability	q_{0,ξ_c}^f	0.201	0.201	0.201	–
Matching probability	$q_{\xi_c,0}^f$	0.237	0.237	0.237	–
Matching probability	q_{ξ_c,ξ_c}^f	0.666	0.666	0.666	–
Probability descendant	π^S	0.34	0	0	0
Probability descendant	π^M	0.89	0	0	0
Pareto weight	μ	0.5	0.5	0.5	0.5
Tax function	κ_0^S	0.2033	0.2029	0.2015	0.2144
Tax function	κ_1^S	0.0497	0.0497	0.0497	0.0497
Tax function	κ_0^M	0.1740	0.1734	0.2073	0.2263
Tax function	κ_1^M	0.0733	0.0733	0.0733	0.0733
Payroll tax rate	τ_p	0.0480	0.0470	0.0160	0.0450
Retirement benefits	b_0^f	0.1060	0.1070	0.1100	0.0970
Retirement benefits	$b_{\xi_c}^f$	0.1341	0.1359	0.1397	–
Retirement benefits	b_0^m	0.1267	0.1279	0.1309	0.1198
Retirement benefits	$b_{\xi_c}^m$	0.1585	0.1600	0.1639	–
Retirement benefits	$b_{0,0}$	0.2143	0.2158	0.2219	0.2115
Retirement benefits	$b_{\xi_c,0}$	0.2452	0.2481	0.2550	–
Retirement benefits	b_{0,ξ_c}	0.2386	0.2414	0.2482	–
Retirement benefits	b_{ξ_c,ξ_c}	0.2675	0.2706	0.2781	–