The Role of Marital Status for the Evaluation of Bankruptcy Regimes*

Tsung-Hsien Li Jan Sun

Abstract

The consumer finance literature has emphasized the importance of income and expense risk for the evaluation of bankruptcy regimes. Single and married households differ in the risks they face. In this paper, we build the first quantitative consumer default model that explicitly models singles and couples. We calibrate our model to the United States in 2019 and estimate expense shocks separately for single and married individuals. Our calibrated model generates large differences in bankruptcy rates across marital status as in the data. We examine how the preferred degree of bankruptcy leniency differs between singles and couples. There are several channels at work: Differences on the income side between singles and couples cause couples to prefer a stricter bankruptcy regime due to the intra-household insurance channel. However, increased risk for couples due to divorce and on the expense side outweighs the first channel. The net effect is that couples prefer more lenient bankruptcy than singles. Our findings suggest that marital status is important to take into account for the evaluation of bankruptcy regimes.

Keywords: Consumer Credit, Bankruptcy, Default, Bankruptcy Regulation, Marital Status

JEL Classifications: D13, D14, D15, E21, E49, G18, G51, J12, K35

*We are extremely grateful to our main advisor Michèle Tertilt for her guidance and support. We are grateful to Klaus Adam for in-depth discussions and to seminar participants at the Mannheim Macro Seminar for many helpful comments. Funding by the "Stiftung Geld und Währung" (Li) and the German Research Foundation (DFG) through CRC TR 224 (project A3, Sun) is gratefully acknowledged. We acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation through grant INST 35/1134-1 FUGG. **Li**: Assistant Research Fellow of Economics at Academia Sinica (email: tsunghsien1124@gmail.com; website: https://tsunghsien1124.github.io/). **Sun**: Senior Economist at Oxford Economics (email: jsun@oxfordeconomics.com; website: https://sites.google.com/view/jansun/startseite).

1 Introduction

Household bankruptcy rates in the United States differ significantly across marital status. This fact has been documented by many empirical studies. Sullivan, Warren, and Westbrook (2000) show that single individuals are over-represented among bankruptcy filers while married are under-represented in the U.S. in 1991.[1] In particular, divorced individuals made up 23% of all bankruptcy filers while only representing 9.7% of the U.S. population at that time. More recently, Fisher (2019) documents that being divorced is highly correlated with bankruptcy conditional on a host of socio-economic characteristics, such as age, race, education, income, employment, home-ownership, etc. In a similar vein, Agarwal, Chomsisengphet, and Liu (2011) estimate that married credit card holders are 32% less likely to declare bankruptcy compared to non-married ones, again conditional on a large range of socio-economic controls. Fay, Hurst, and White (2002) highlight the importance of divorce events as their results suggest that the probability of bankruptcy increases by 86% in the year following a divorce.

However, until now the quantitative consumer default literature has exclusively employed models that do not differentiate between single and couple households.[2] All households are modeled as a single entity. Prior work, such as Livshits, MacGee, and Tertilt (2007), emphasized the importance of income and expense risk for the welfare evaluation of bankruptcy regimes. Couple households differ from single households in these aspects: First, couples have potential access to two earners. A spouse can adjust his/her labor supply in response to the partner's earnings shocks, the *intra-HH insurance channel*. Previous literature has emphasized this channel both empirically (see e.g. Blundell, Pistaferri, and Preston (2008)) as well as structurally (see e.g. Ortigueira and Siassi (2013)). Thus, couples face different income risks than singles. Second, couples by nature face different expense risks compared to singles. Couples consist of two individuals who both can suffer expense shocks. In addition, couples are also at risk of divorce, an event that is costly and often goes hand in hand with bankruptcy as mentioned earlier. Finally, couples also benefit from economies of scale in consumption whereas singles do not.[3]

[1] See Figure 6.1 in Sullivan et al. (2000).

[2] In the following, we will use the terms couple and married household interchangeably.

[3] Clearly, there are other dimensions of heterogeneity between couples and singles, such as childbirth and childcare. As our paper is the first that models couples and singles explicitly in a consumer default setting, we abstract from these differences in this paper as a first step.

Our contribution is building a consumer default model that explicitly differentiates between single and couple households for the first time and examining the implications for the welfare evaluation of bankruptcy regimes. We build upon the life-cycle Huggett-type model with consumer default in Livshits et al. (2007). Individuals derive utility from consumption and leisure. They are subject to wage and expense shocks. In response to these shocks, individuals choose how much to work. Households also choose how much to save or borrow as well as whether to repay their loans or to default. Default is costly and a fraction of the filing household's wage is garnished. Furthermore, couples can become divorced. Financial services are provided by competitive financial intermediaries. These intermediaries price loans according to the individual default risk of the borrower. There is no asymmetric information in our model.

We calibrate our model to the U.S. population in 2019. Standard parameters are taken from the literature. Other parameters are exogenously calibrated to data whenever possible. We pay particular attention to the estimation of expense shocks. Previous literature has shown that medical expenses are a main component of these shocks and that they are key for welfare implications.[4] We estimate out-of-pocket medical expenses separately for single and married individuals using data from the Medical Expenditure Panel Survey (MEPS) for 2018 and 2019.[5] The remaining parameters, including the discount factor, the utility weight on consumption, the wage garnishment rate, and the annual transaction cost of lending are internally calibrated to match the several aggregate data moments. We also evaluate our model on a set of untargeted moments.

Using our calibrated model, we first compute the ex-ante welfare of newborns across different garnishment rates for single and couple households. There exists a trade-off when considering the optimal garnishment rate: On the one hand, a lower garnishment rate is a more lenient bankruptcy regime and thus allows households to default more cheaply in response to adverse shocks. On the other hand, a stricter bankruptcy regime via a higher garnishment rate reduces default rates, resulting in lower default premia in equilibrium and as a result cheaper loans for households.[6] We find that single and married households prefer different degrees of bankruptcy leniency. In our model, couples

[4] For example, Livshits et al. (2007).

[5] Out-of-pocket, not total, medical expenses are the relevant type of expense for this model. This is because out-of-pocket expenses have to be paid by the individuals themselves and are thus relevant to the bankruptcy decision. Expenses paid by insurance do not affect a household's bankruptcy choice.

[6] In the literature, the first channel is referred to as smoothing consumption *across states*. The second one is described as smoothing *over time*.

prefer a more lenient regime than singles. We then decompose which sources of heterogeneity between singles and couples drive this finding. We find that differences on the income side of households cause couples to prefer a **stricter** regime than singles. This is because couples have access to the intra-household insurance channel: Spouses can adjust their own labor supply in response to their partners' wage shocks. This additional source of insurance for couples means that they have less need for the insurance provided by bankruptcy. However, this channel is outweighed by the additional risk faced by couples through divorce and because couples are hit more often by expense shocks (due to consisting of two individuals instead of one). These two factors make couples value bankruptcy more than singles. The net effect is that couples prefer a more lenient regime. Overall, our results suggest that marital status is an important source of heterogeneity when evaluating the welfare implications of different bankruptcy regimes.

The remainder of the paper is structured as follows: Section 2 gives an overview of the related literature. In Section 3, we summarize how household bankruptcy works in the U.S. with a focus on differences between single and married filers. Section 4 presents the model framework. In Section 5, we detail the calibration of our model and the estimation of medical expense shocks from data. Section 6 presents the welfare result of our baseline model as well as a decomposition of the different channels at work. Section 7 contains a robustness check. Section 8 concludes with some potential further avenues for research.

2 Related Literature

In this section, we discuss the literature most closely related to our paper. We begin with papers from the consumer default literature. Afterward, we summarize some papers which emphasize the importance of modeling singles and couples separately. Two important ingredients in our model are the intra-household insurance channel and expense shocks. We conclude this section by presenting a brief overview of some papers that have looked at these two elements.

Our paper is closely related to the structural literature on consumer default. Livshits et al. (2007) is one of the workhorse models in this literature. The authors analyze two different bankruptcy regimes: "fresh start (FS)" and "no fresh start (NFS)". In a FS economy households are allowed to discharge debt by defaulting. In contrast, in a NFS economy debt cannot be discharged but must instead be repaid under a repayment plan. They build a heterogeneous agent lifecycle model and examine the welfare consequences of the two regimes. They find that the nature of expense and income uncertainty is crucial for the welfare assessment of the two regimes. Livshits, MacGee, and Tertilt (2010) build on the same model to investigate the rise in consumer bankruptcy filings between 1970 to 2002. They argue that the most important drivers behind this development are lower lending transaction costs and decreasing costs of bankruptcy. Chatterjee, Corbae, Nakajima, and Ríos-Rull (2007), another workhorse model in this literature, build a consumer default model with infinitely-lived households. They study the welfare implications of a means test introduced in the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005.[7] The authors show that this policy change leads to welfare gains in their model. Herkenhoff (2019) examines the relationship between changes in consumer credit access and business cycles. To do so, he builds a consumer default model with a labor market subject to search and matching frictions. Chatterjee, Corbae, Dempsey, and Ríos-Rull (2020) build a consumer default model with asymmetric information between lenders and borrowers. In their model lenders compute a credit score for each borrower to better assess their creditworthiness. A paper that models endogenous labor supply in a consumer default framework like mine is Exler (2019). In contrast to our paper, none of the studies above model single and couple households separately.

Our paper is also related to a growing literature emphasizing the importance of explicitly modeling singles and couples for understanding household behavior and for analyzing policy implications. De Nardi, French, Jones, and McGee (2021) try to understand the savings behavior of retired households using a structural approach. In order to do so, they show that it is important to model marital status along with medical expenses and bequest motives. Borella, De Nardi, and Yang (2021) look at how the dependence of taxes and retirement benefits on marital status affects female labor supply. Their results suggest that an elimination of marriage-related provisions would result in increased labor force participation of women and lead to large welfare gains for most households. Guner, Kaygusuz, and Ventura (2012) study the effects of two tax reforms while explicitly modeling couples: a proportional income tax and a reform allowing separate filing for married individuals. They illustrate that both reforms lead to a large increase in the labor

[7] The means test is a test that Chapter 7 bankruptcy applicants have to pass before filing. It prevents households with income above certain thresholds to file.

supply of married females. Bacher (2021) examines how marital status affects the investment choices of households. She finds that divorce risk lowers the demand for illiquid assets. Failure to explicitly account for marital status thus results in overestimation of the attractiveness of housing and housing-related policies. In contrast to these studies, our paper is the first to examine the importance of distinguishing between single and couple households in a consumer default setting.

One important mechanism in our model is the intra-household insurance channel. There exists a large literature that examines how spousal labor supply acts as insurance against risk. On the one hand, empirical studies have investigated this channel. Blundell et al. (2008) use the PSID to study changes in income and consumption inequality. Among other results, they document that family labor supply plays an important role for insuring against permanent income shocks. Also using the PSID, Shore (2010) examines the cyclical properties of the intra-household insurance channel. On the other hand, there exists a structural literature which studies this mechanism. Kotlikoff and Spivak (1981) is an early paper that studies the insurance provided by a family. Attanasio, Low, and Sánchez-Marcos (2005) use a lifecycle model with a unitary family to study the role of spousal labor supply as insurance against earnings risk. They show that the welfare cost of increased uncertainty is higher if female labor supply cannot be adjusted. Ortigueira and Siassi (2013) use an Aiyagari-Huggett setup with couples in order to quantify the effects of this channel on household behavior. In their model couples suffer a much weaker consumption drop upon unemployment relative to single individuals.[8] Compared to these papers, we embed the intra-household insurance channel in a consumer default framework and investigate its effects on bankruptcy behavior.

Another important ingredient in our model are expense shocks. There is a sizeable literature that tries to quantify medical expenses and look at their effect on different household dynamics. French and Jones (2004) use data from the Health and Retirement Survey and the Assets and Health Dynamics of the Oldest Old to estimate a stochastic process for health care costs of old people. Hubbard, Skinner, and Zeldes (1994) build a lifecycle model including uncertain longevity and medical expenses. They show that these features help match wealth and consumption dynamics in the U.S. Palumbo (1999) builds a structural model in which elderly individuals suffer from medical expense shocks. He

[8] Their results also suggest that wealth-poor households rely more on spousal labor supply as an insurance mechanism.

demonstrates that uncertain medical expenses help explain the dissaving behavior of retired people. De Nardi, French, Jones, McGee, and Rodgers (2020) estimate medical expenses for retired individuals using the Health and Retirement Study. They document that medical expenses in addition to bequests can explain asset changes around an individual's death. In contrast to these studies, we focus on medical expenses for working age people and on the effects of these expenses on household default behavior. Another paper in the consumer default literature that highlights the importance of expense shocks is Livshits et al. (2007). They find that expense uncertainty is crucial for the comparison of different bankruptcy regimes and show that with larger expense uncertainty households prefer a more lenient regime.

3 Institutional Details

Households in the United States have access to two types of bankruptcy: Chapter 7 and Chapter 13. Chapter 7 bankruptcy (also called liquidation) allows households to discharge their unsecured debt.[9] In return, assets above a certain exemption level are liquidated to repay the creditors. However, not everyone is allowed to file for Chapter 7. The Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 introduced a so-called means test. This means test prevents individuals above certain income thresholds from filing for Chapter 7.[10] Under Chapter 13 bankruptcy (also called reorganization) debtors can propose a repayment plan. Compared to Chapter 7, Chapter 13 allows debtors to keep their assets in general. As Chapter 7 bankruptcy is the most prevalent form of household bankruptcy, we follow most of the consumer default literature in focusing on this type in our paper.

Compared to filing for Chapter 7 as a single household, filing as a married couple is more complex. When a married couple decides to default, it has three choices: (1) File jointly for bankruptcy, (2) one spouse files whereas the other does not, and (3) both spouses file but separately. When a couple files jointly, the debts of both spouses get discharged. In return, assets of both spouses are eligible to be liquidated. Depending on the state, a joint filing may allow a couple to double certain asset exemptions and thus

[9] Some notable exceptions include student loans, childcare, and alimony.

[10] For more details, see e.g. https://www.uscourts.gov/services-forms/bankruptcy/bankruptcy-basics/chapter-7-bankruptcy-basics.

let couples keep more of their assets.[11] Filing jointly also has the advantage of lower per capita legal fees, as lawyer and court costs only have to be paid once.

When one spouse chooses to file for bankruptcy and the other does not, which debts are discharged and which assets can be liquidated depends on state laws. In general, U.S. states can be divided into those that follow ”common law” and those that follow ”community property law”.[12] In common law states assets acquired during marriage belong to the acquiring spouse only (unless the asset was acquired in the names of both). In these states only the filer's own debts are discharged. The other spouse remains liable for any of her/his own debt as well as for debts that belong to both spouses. Assets that belong to the non-filing spouse cannot be liquidated however, only assets that belong to the filer. This is in contrast to community property states where assets purchased during marriage belong to both spouses by default. Certain community property states allow the discharge of joint debts even if only one spouse files. At the same time, in some community property states joint assets may be liquidated even if one spouse did not file. As our framework only allows joint assets for couples, we see our model as representative of community property states.

4 The Model

Each model period lasts three years. Our setup uses a life-cycle model where households start life at age 20. Households are identical ex-ante. They live for 16 periods and die at age 68 with probability one. Households maximize their discounted lifetime utility from consumption and leisure.

There are two different household types: singles and couples. Singles can be either female or male. Couples consist of one female and one male individual. They always start life and die together. There is no marriage in the model. Singles always stay single. Couples are subject to an exogenous, random divorce shock. The probability of this divorce shock is the same for all couples. After divorce, there is no possibility of re-marriage. Couples are modeled in a unitary framework and they jointly decide on their actions.

[11] For an overview of the exemption regulations in different U.S. states, see e.g. `https://www.nolo.com/legal-encyclopedia/bankruptcy-exemptions-state`.

[12] Most U.S. states follow common law. Those that follow community property law are Arizona, California, Idaho, Louisiana, Nevada, New Mexico, Texas, Washington, and Wisconsin. Many European countries also follow community property law.

Individuals in the economy can choose how much time to work. They are subject to idiosyncratic uncertainty including wage shocks as well as expense shocks. The wage process follows a persistent autoregressive process and differs between female and male individuals. Expense shocks are i.i.d. and vary between singles and couples as well as across ages. In the following sections, we suppress the dependence of the expense shock on marital status and age for better readability. In addition, households are subject to a deterministic life-cycle productivity profile. There is no aggregate uncertainty.

Households can save or borrow to smooth consumption. Asset markets are incomplete. Households have access to only one-period non-contingent bonds. Couples have access to a joint asset.[13] Importantly, households may declare bankruptcy in order to insure themselves against shocks. In case of default, a fraction of the household's wage is garnished.[14] In return all debts and expense shocks are discharged. A household may also default when it has savings but is subject to an expense shock. When this happens, all savings are lost. Couples are only allowed to jointly file for bankruptcy.[15]

We assume that all individuals are born with zero assets and face no expense shock in their first period of life. There is however heterogeneity regarding the starting wage. For simplicity, we assume that it is drawn from a uniform distribution over all possible wage realizations.

Loans and saving services are extended by a perfectly competitive financial intermediary sector. This sector takes as given the exogenous risk-free rate. Loans are priced such that in expectation the financial intermediary makes zero profit on every loan it extends. Households take as given this loan pricing schedule. There is no asymmetric information.

The timing within each period is as follows: First, households realize their wage and expense shocks. In response to these shocks, they choose how much to work and whether to default or not. If a household does not default, it also chooses how much to save or borrow. Couples realize their divorce shock at the beginning of each period. If a couple becomes divorced, the two individuals separate and make their own choices as divorced individuals for the period. Assets are split equally upon divorce.

[13] For tractability we assume that spouses in a couple cannot have separate asset holdings.

[14] In practice, Chapter 7 bankruptcy does not entail wage garnishments. However, one may interpret the garnishment as an honest effort of the borrower to repay his debts as required by U.S. bankruptcy law under the good faith requirement.

[15] As couples only have access to a joint asset, it does not make sense to model separate bankruptcy filings for them within this model setup.

In the rest of this section, we formally introduce the problem of households in Section 4.1. Section 4.2 describes the problem of the financial intermediaries. Finally, in Section 4.3 we define the equilibrium.

4.1 Households

In this section, we describe the problem of households. First, we start with the problem facing single households. Afterward, we lay out the problem of couples before turning to the divorced' problem.

4.1.1 Singles

Each period, singles choose whether to default or repay:[16]

$$V_{S,g,j}(a, z_g, \kappa) = \max\{V^R_{S,g,j}(a, z_g, \kappa), V^D_{S,g,j}(z_g)\} \tag{1}$$

where $V_{S,g,j}(a, z, \kappa)$ is the value function of a single (S), with gender g and age j. The value depends on the individual's asset position a, wage shock z_g drawn from a persistent AR(1) process, and i.i.d. expense shock κ. $a > 0$ denotes savings, whereas $a < 0$ denotes borrowing. $V^R_{S,g,j}$ denotes the value of repayment, while $V^D_{S,g,j}$ is the value of defaulting.[17]

The repayment value function for singles is given by:

$$V^R_{S,g,j}(a, z_g, \kappa) = \max_{(a',n)} u\,(c, l) + \beta \cdot \mathbb{E}_{z'_g,\kappa'} \left\{ V_{S,g,j+1}(a', z'_g, \kappa') \right\} \tag{2}$$

s.t.

$$c + q^{(a')}_{S,g,j}(z_g) \cdot a' \leq e_j \cdot z_g \cdot n + a - \kappa \tag{3}$$

$$l = T - n \tag{4}$$

In Equation (2) singles choose their next period asset position a' and labor supply n. They maximize utility from consumption c and leisure l and the expected next period

[16] Note again that we suppress the dependence of the expense shock κ on marital status and age.

[17] The value of defaulting does not depend on assets a or expense shock κ. This is because, in case of default, all assets and expense shocks are discharged.

value discounted by β. The expectation is taken over the next-period realization of the wage shock z'_g and expense shock κ'.

Equation (3) shows the budget constraint. The resources available to the household are shown on the right-hand side. Income is determined by a life-cycle productivity component e_j, the persistent wage shock z_g (which can differ across gender g), and the labor choice n. a is the asset position that individuals entered the period with. Available resources are reduced by the expense shock κ. On the left-hand side, the individual can choose consumption c and next period's asset position a'. $q^{(a')}_{S,g,j}(z_g)$ denotes the discount pricing schedule for loans. Note that the price for a loan depends on its size a' as individuals are more likely to default on a larger loan ceteris paribus. The price also depends on marital status S, gender g, age j, and productivity z_g as all these variables influence the repayment probability of the loan. Equation (4) shows the time constraint with total time endowment given by T.

Similarly, the value function for the default case for singles is given by:

$$V^D_{S,g,j}(z_g) = \max_n u\,(c,l) + \beta \cdot \mathbb{E}_{z'_g,\kappa'} \left\{ V_{S,g,j+1}(0, z'_g, \kappa') \right\} \tag{5}$$

s.t.

$$\begin{aligned} c &\leq (e_j \cdot z_g \cdot n) \cdot (1-\phi) \\ l &= T - n \end{aligned} \tag{6}$$

Note that in case of default, there is no asset choice to be made: In the period of default, no borrowing or saving is allowed ($a' = 0$). Equation (6) shows the budget constraint for the default case. All debts (or savings) a and expense shocks κ are discharged. In return, a fraction ϕ of the individual's wage is garnished.

4.1.2 Couples

Couples in our framework are modeled using a unitary approach. All choices are made jointly by the two individuals in a couple.[18]

Analogous to the case for singles, couples choose whether to jointly repay or default

[18] A similar framework is used in Borella et al. (2021) for example.

each period:

$$V_{C,j}(a,z_f,z_m,\kappa_f,\kappa_m) = \max\{V^R_{C,j}(a,z_f,z_m,\kappa_f,\kappa_m), V^D_{C,j}(z_f,z_m)\} \tag{7}$$

Here, the subscripts denote the female f or male m in a couple.

The repayment value function for couples is given by:

$$\begin{aligned} V^R_{C,j}(a,z_f,z_m,\kappa_f,\kappa_m) &= \max_{(a',n_f,n_m)} u\left(\frac{c}{\eta}, l_f\right) + u\left(\frac{c}{\eta}, l_m\right) \\ &+ \beta \cdot \Bigg((1-\psi) \cdot \mathbb{E}_{z'_f,z'_m,\kappa'_f,\kappa'_m} \left\{ V_{C,j+1}(a',z'_f,z'_m,\kappa'_f,\kappa'_m) \right\} \\ &+ \psi \cdot \left(\mathbb{E}_{z'_f,\kappa'_f} \left\{ V_{Div,f,j+1}(\frac{a}{2}, z'_f, \kappa'_f) \right\} + \mathbb{E}_{z'_m,\kappa'_m} \left\{ V_{Div,m,j+1}(\frac{a}{2}, z'_m, \kappa'_m) \right\} \right) \Bigg) \end{aligned} \tag{8}$$

s.t.

$$\begin{aligned} c + q^{(a')}_{C,j}(z_f,z_m) \cdot a' &\leq e_j \cdot z_f \cdot n_f + e_j \cdot z_m \cdot n_m + a - \kappa_f - \kappa_m \qquad (9) \\ l_f &= T - n_f \\ l_m &= T - n_m \end{aligned}$$

In Equation (8) couples jointly maximize the sum of their individual utilities and their expected continuation value. Consumption in couples is adjusted by an equivalence scale η. This scale captures economies of scale in consumption within couples. With $\eta < 2$ couples can consume more than what they could consume if they were living separately. With probability ψ a couple gets hit by an exogenous divorce shock next period. Thus, with probability $(1-\psi)$ the relevant continuation value is that of couples and with probability ψ it is the sum of the two divorced continuation values denoted by V_{Div}. We assume that in our model assets (or debts) get split 50-50 in the event of divorce.

Equation (9) shows the budget constraint. Note that couples are only allowed to save/borrow in one joint asset a'. Furthermore, each individual is subject to its own idiosyncratic productivity (z_f, z_m) and expense (κ_f, κ_m) shocks. Couples face their own loan pricing schedule q_C which is different from the one for singles q_S.

For the case of default, the value function is given by:

$$V^{D}_{C,j}(z_f, z_m) = \max_{(n_f,n_m)} u\left(\frac{c}{\eta}, l_f\right) + u\left(\frac{c}{\eta}, l_m\right)$$
$$+ \beta \cdot \Bigg((1-\psi) \cdot \mathbb{E}_{z'_f,z'_m,\kappa'_f,\kappa'_m}\left\{V_{C,j+1}(0, z'_f, z'_m, \kappa'_f, \kappa'_m)\right\}$$
$$+ \psi \cdot \left(\mathbb{E}_{z'_f,\kappa'_f}\left\{V_{Div,f,j+1}(0, z'_f, \kappa'_f)\right\} + \mathbb{E}_{z'_m,\kappa'_m}\left\{V_{Div,m,j+1}(0, z'_m, \kappa'_m)\right\}\right)\Bigg) \tag{10}$$

s.t.

$$c \leq (e_j \cdot z_f \cdot n_f + e_j \cdot z_m \cdot n_m) \cdot (1-\phi) \tag{11}$$
$$l_f = T - n_f$$
$$l_m = T - n_m$$

The interpretation of Equation (10) is analogous to Equation (8). Note in Equation (11) that in case of default a fraction ϕ of both spouses' income is garnished.

4.1.3 Divorced

The decision problem for divorced individuals is identical to the one for singles, except in the period of divorce in which an additional divorce cost has to be paid.

Divorced choose whether to repay or default each period:

$$V_{Div,g,j}(a, z_g, \kappa) = \max\{V^{R}_{Div,g,j}(a, z_g, \kappa), V^{D}_{Div,g,j}(z_g)\} \tag{12}$$

The repayment value function for divorced is given by:

$$V^{R}_{Div,g,j}(a, z_g, \kappa) = \max_{(a',n)} u(c, l) + \beta \cdot \mathbb{E}_{z'_g,\kappa'}\left\{V_{S,g,j+1}(a', z'_g, \kappa')\right\} \tag{13}$$

s.t.

$$c + q^{(a')}_{S,g,j}(z_g) \cdot a' \leq e_j \cdot z_g \cdot n + a - \kappa - \kappa_{Div} \tag{14}$$
$$l = T - n$$

Note that the continuation value in Equation (13) is the one for singles as the divorced problem is identical to the singles' one after the first period of divorce. For this reason the relevant pricing schedule in Equation (14) is also the one for singles. The divorce cost is captured by κ_{Div} and represents monetary costs from divorce such as lawyer fees.

The default value function for divorced is given by:

$$V^{D}_{Div,g,j}(z_g) = \max_{n} u\left(c,l\right) + \beta \cdot \mathbb{E}_{z'_g,\kappa'}\left\{V_{S,g,j+1}(0,z'_g,\kappa')\right\} \tag{15}$$

s.t.

$$\begin{aligned} c &\leq (e_j \cdot z_g \cdot n) \cdot (1-\phi) \\ l &= T - n \end{aligned} \tag{16}$$

Equation (16) shows that default also discharges the cost from divorce κ_{Div} in addition to other expense shocks.

4.2 Financial Intermediaries

Banks have access to funding at the exogenous, risk-free rate r_f. They operate in a perfectly competitive environment and every loan is priced such that it yields zero profit in expectation. Households differ in their default risk depending on their marital status, gender, age, and persistent wage. Furthermore, the loan size also plays a role for the default risk. Households are more likely to default on larger loans ceteris paribus. As a result, banks condition on all these variables when pricing their loans.

The bond price of a loan with size a' for a single S of gender g and age j with wage z is given by

$$q^{(a')}_{S,g,j}(z_g) = \\ = \begin{cases} \left(\mathbb{P}^{(a')}_{S,g,j}(z_g) \cdot 1 + (1-\mathbb{P}^{(a')}_{S,g,j}(z_g)) \cdot \mathbb{E}\left(\frac{\Gamma}{a'+\kappa'}|d'=1\right)\right) \cdot \frac{1}{1+r_f+\tau} & \text{if } a' < 0 \\ \frac{1}{1+r_f} & \text{if } a' \geq 0 \end{cases} \tag{17}$$

where τ is a borrowing wedge capturing the transaction cost of making loans. $\mathbb{P}$ denotes the repayment probability next period. If the household defaults, the bank will garnish a fraction ϕ of the wage. This leads to an expected recovery of $\mathbb{E}\left(\frac{\Gamma}{a'+\kappa'}|d'=1\right)$, where $d'=1$ indicates default.[19] The recovery Γ in case of default is given by

$$\Gamma = (e_{j+1} \cdot z'_g \cdot n') \cdot \phi$$

The loan bond price for a couple C of age j is given by

$$q^{(a')}_{C,j}(z_f, z_m) = \\ = \begin{cases} \left((1-\psi) \cdot \mathbb{E}\{\hat{\mathbb{P}}_C\} + \psi \cdot (\frac{a}{2} \cdot \mathbb{E}\{\hat{\mathbb{P}}_{Div,f}\} + \frac{a}{2} \cdot \mathbb{E}\{\hat{\mathbb{P}}_{Div,m}\})\right) \cdot \frac{1}{1+r_f+\tau} & \text{if } a' < 0 \\ \frac{1}{1+r_f} & \text{if } a' \geq 0 \end{cases} \tag{18}$$

where $\mathbb{E}\{\hat{\mathbb{P}}_C\}$ and $\mathbb{E}\{\hat{\mathbb{P}}_{Div,g}\}$ denote expected repayment and recovery amounts from couples and divorced respectively. Recall that ψ denotes the probability of divorce. The expected repayment and recovery amounts are defined as below:

$$\mathbb{E}\{\hat{\mathbb{P}}_C\} = \overbrace{\mathbb{P}^{(a')}_{C,j}(z_f, z_m) \cdot 1}^{\text{repayment}} + \overbrace{(1 - \mathbb{P}^{(a')}_{C,j}(z_f, z_m)) \cdot \mathbb{E}\left(\frac{\Gamma_C}{a' + \kappa'_f + \kappa'_m}|d'=1\right)}^{\text{recovery}}$$
$$\mathbb{E}\{\hat{\mathbb{P}}_{Div,g}\} = \mathbb{P}^{(a')}_{Div,g,j}(z_g) \cdot 1 + (1 - \mathbb{P}^{(a')}_{Div,g,j}(z_g)) \cdot \mathbb{E}\left(\frac{\Gamma_{Div,g}}{a' + \kappa' + \kappa_{Div}}|d'=1\right)$$

where $\mathbb{P}$ again denotes the repayment probability next period. The recovery amounts in case of default are given by:

$$\Gamma_C = (e_{j+1} \cdot z'_f \cdot n'_f + e_{j+1} \cdot z'_m \cdot n'_m) \cdot \phi$$
$$\Gamma_{Div,g} = (e_{j+1} \cdot z'_g \cdot n'_g) \cdot \phi$$

4.3 Equilibrium

Given a risk-free rate r_s a recursive competitive equilibrium is given by a set of value functions $(V^R_S, V^D_S, V^R_C, V^D_C, V^R_{Div}, V^D_{Div})$, a set of policy functions $(c_S, c_C, c_{Div}, a'_S,$ $a'_C, a'_{Div}, d'_S, d'_C, d'_{Div}, n_S, n_{C,f}, n_{C,m}, n_{Div})$, and a set of bond pricing functions (q_S, q_C) such that:

1. The value functions satisfy Equations (1), (2), (5), (7), (8), (10), (12), (13), (15).

[19] Note here that the garnished wage will be used to proportionally repay incurred expense shocks on top of the debt.

2. The policy functions are the associated optimal policy rules.

3. The bond price schedules satisfy the zero profit conditions (17) and (18).

We compute the equilibrium value functions, policy functions, and bond price schedules by backward induction starting at the final age period. Further computational details are given in Appendix A.1.

5 Calibration

We choose the baseline calibration year as 2019 and calibrate the model to the U.S. population. There are three sets of parameters: (1) Those that are standard and which we take from the literature, (2) parameters that we exogenously calibrate to direct empirical counterparts, and (3) parameters that we internally calibrate to have the model match certain data moments. Table 1 contains an overview of all exogenously chosen parameters, whereas Table 2 lists the internally calibrated parameters. Table 3 summarizes the targeted moments used for the calibration.

We assume that individuals may choose to work full-time ($n = 1$), part-time ($n = 0.5$) or not at all ($n = 0$). The utility function is given by

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

where γ is the risk aversion coefficient and which we set to 2 which is a standard value in the macro literature. ω denotes the utility weight of consumption and is important for the labor supply choice of individuals. We thus internally calibrate this parameters to match the average hours worked of singles at age 50 as estimated by Borella, De Nardi, and Yang (2018). This yields a value of $\omega = 0.56$. The annual discount factor is calibrated to match the fraction of households in the Survey of Consumer Finances (SCF) 2019 with negative net worth.[20] We restrict the SCF sample to households with a head aged between 20 and 68 in line with our model. This results in an annual value of about 0.973 (and thus $\beta = 0.973^3 = 0.92$).

[20] The definition of net worth follows Herkenhoff (2019). It is computed as the difference between a household's liquid assets, such as checking and savings accounts, and credit card debt. We prefer this measure of net worth as our model does not include illiquid assets like housing.

Parameter	Symbol	Value	Source
Income processes			
Persistence, women	ρ_f	0.963	Borella et al. (2018)
Persistence, men	ρ_m	0.973	Borella et al. (2018)
Variance, women	$\sigma^2_{\epsilon,f}$	0.014	Borella et al. (2018)
Variance, men	$\sigma^2_{\epsilon,m}$	0.016	Borella et al. (2018)
Lifecycle productivity	e_j		Livshits et al. (2010)
Expense shocks	κ		Own (MEPS data)
Annual savings rate	r_f	3.44%	Gourinchas and Parker (2002)
Total weekly time endowment	T	60 hours	Alon et al. (2020)
Risk aversion coefficient	γ	2	Standard
Annual probability for divorce	$\frac{\psi}{3}$	1.1%	American Community Survey (2019)
Divorce cost	κ_{Div}	\$11,300	Martindale-Nolo Research (2019)
Equivalence scale in couples	η	1.64	Voena (2015)

Table 1: Exogenously Chosen Parameters

The persistent wage process is taken from Borella et al. (2018). We use their estimated wage process because they estimate wage processes separately for men and women using PSID data. The authors assume an AR(1) process in log wages: $\ln z^i_{g,t+1} = \rho_g \ln z^i_{g,t} + \epsilon^i_{g,t}, \epsilon^i_{g,t} \sim N(0, \sigma^2_{\epsilon,g})$ for gender $g \in \{f, m\}$. The estimated process for women shows slightly lower persistence than the one for men ($\rho_f = 0.963$ vs. $\rho_m = 0.973$) and a smaller shock variance ($\sigma^2_{\epsilon,f} = 0.014$ vs. $\sigma^2_{\epsilon,m} = 0.016$). We convert their annual estimates to triennial values and then discretize them into two five-state Markov processes using the Rouwenhorst method. The lifecycle productivity profile is taken from Gourinchas and Parker (2002).[21]

We set the risk-free savings rate to 3.44% following Gourinchas and Parker (2002).[22] This implies a three-year risk-free rate on savings of 10.68%. The transaction cost of lending τ is calibrated internally to match the average interest rate on credit cards in the 2019 SCF. We again restrict the sample to household heads aged between 20 and 68 and also exclude households that report no credit card debt or a non-positive interest rate.[23] This results in an annual value for the transaction cost of lending of 0.93%. Together with the risk-free savings rate, this implies a three year risk-free lending rate of around 13.7%.[24]

[21] This profile is also used in Livshits et al. (2007) and Livshits et al. (2010).

[22] This is the value used in Livshits et al. (2010). Voena (2015) uses a similar value of 3%.

[23] We exclude observations with no credit card debt as these households use credit cards for transactional, and not borrowing, purposes. We leave out observations with non-positive interest rates as these are usually temporary promotional rates.

[24] $(1 + 0.0344 + 0.0093)^3 \approx 1.1369$

Parameter	Value	Data Target
Annual Discount factor	0.97	Frac. of HH with neg. net worth
Consumption weight	0.56	Avg. hours of singles at 50
Wage garnishment rate	0.395	Ch. 7 bankruptcies per HH
Ann. transaction cost of lending	0.93%	Avg. credit card interest rate

Table 2: Internally Calibrated Parameters

Notes: Source for fraction of HH with neg. net worth: Survey of Consumer Finances 2019. Source for avg. hours of singles at 50: Borella et al. (2018). Source for Ch. 7 bankruptcies per HH: American Bankruptcy Institute. Source for avg. credit card interest rate: Survey of Consumer Finances 2019.

The wage garnishment rate is crucial for the amount of default in the economy. We set it to $\phi = 0.395$ to match the number of Ch. 7 bankruptcies per household in the U.S. in 2019 as reported by the American Bankruptcy Institute. Total time endowment $T = 1.5$ is taken from Alon, Doepke, Olmstead-Rumsey, and Tertilt (2020). For a full-time job of $n = 1$ corresponding to 40 hours, this value for T implies a total weekly time endowment of 60 hours.

The parameters governing the divorce shock are calibrated as follows: The probability of a divorce shock is pinned down by the divorce rate in the U.S. using data from the 2019 American Community Survey. This yields an annual divorce rate of 1.1% resulting in a three-year rate of 3.3%. The costs of divorce are taken from a survey conducted by Martindale and Nolo Research in 2019.[25] Their survey gives an estimate of average legal costs of divorce in 2019 of $11,300.

The equivalence scale for consumption in couples is taken from Voena (2015). She calibrates the degree of economies of scale for couples to 1.4023, which implies an equivalence scale of 1.64.[26]

As it turns out, expense shocks are crucial for welfare implications. Thus, we estimate these shocks myself in order to allow for heterogeneity in expenses across marital status and age. The next subsection details our procedure.

[25] See also https://www.nolo.com/legal-encyclopedia/ctp/cost-of-divorce.html.

[26] More specifically, she models spousal consumption as $x = ((c^H)^\rho + (c^W)^\rho)^{\frac{1}{\rho}}$, where x are household expenditures and c^H (c^W) is the consumption of the husband (wife). Assuming that consumption is split equally between husband and wife (as is the case in our model), $\rho = 1.4023$ implies an equivalence scale of 1.64.

Moment (in %)	Data	Model
Default rate (aggregate)	1.1	1.1
Fraction of borrowers (aggregate)	20.7	20.7
Avg. interest rate (aggregate)	16.3	16.3
Avg. hours of singles at 50	1786	1795

Table 3: Targeted Moments

5.1 Expense Shocks

The previous literature has highlighted several sources of unexpected expenses which are important to consider for the bankruptcy decision of households. Among these sources are out-of-pocket medical, divorce, and childcare expenses resulting from unplanned pregnancies, see e.g. Livshits et al. (2007). As our framework explicitly models divorce, we include divorce expenses in the divorce shock instead of the expense shock. Furthermore, our model abstracts from children and childcare. As such, we also abstract from childcare expenses when estimating the expense shock to feed into our model.

To estimate medical expenses we use data from the Medical Expenditure Panel Survey (MEPS) from the years 2018 and 2019. The MEPS features an overlapping cohorts design and follows each cohort for two years. It collects detailed information on medical expenditures of households and includes many demographic attributes. In particular, MEPS also collects the source of payment. This is important because medical expenses relevant for the bankruptcy decision of households are those that have to be paid by themselves (out-of-pocket).[27]

To estimate these shocks we largely follow the approach laid out in Livshits, MacGee, and Tertilt (2003). We focus on Panel 23 which covers the years 2018 and 2019. There are two issues with using out-of-pocket spending reported in MEPS out of the box for our estimation.

The first issue is that MEPS underreports out-of-pocket medical spending compared to aggregate sources. Average per capita out-of-pocket spending in MEPS for 2018 (2019) in Panel 23 was \$826.45 (\$834.01). Using National Health Expenditure Data the same figure for 2018 (2019) was \$1184.39 (\$1233.06).[28] Under the assumption that the factor of

[27] As opposed to payments covered by insurance.

[28] Total out-of-pocket medical expenditures in 2018 (2019) were \$386.5 billion (\$403.7B billion). The total U.S. civilian non-institutionalized population in 2018 (2019) was 326 million (327 million).

underreporting is constant across the population, we adjust MEPS out-of-pocket expense numbers in 2018 (2019) by a factor of 1.43 (1.48).

The second issue is that out-of-pocket medical spending in MEPS does not include bad debts, i.e. medical bills unpaid by households. However, these bills are part of the medical expenses faced by households and influence households' default behavior. As such we construct a measure of medical expenses by adding bad debt to the out-of-pocket spending reported in MEPS. The American Hospital Assocation (2020) reports total U.S. uncompensated hospital care cost in 2018 (2019) of \$41.3 billion (\$41.61 billion).[29] This corresponds to 3.68% (3.49%) of total U.S. hospital spending in 2018 (2019).[30] Assuming that this ratio also holds for the total medical sector, we get an estimate of bad debt in the U.S. medical sector of \$111.2 billion (\$110.81 billion) for 2018 (2019).[31] We allocate this sum to all individuals in the MEPS data who were not insured in at least one month of 2018 (2019) proportional to the difference between their charges and expenditures.

More specifically, we compute adjusted out-of-pocket medical expenses $\widetilde{OOP}_i^Y$ facing an individual i in year Y in the following way:

$$\widetilde{OOP}_i^Y = a^Y \cdot OOP_i^Y + b^Y \cdot we_i^Y \cdot (charge_i^Y - exp_i^Y) \quad (19)$$

where $a^{2018} = 1.43$ ($a^{2019} = 1.48$) is the adjustment factor from before. OOP are the out-of-pocket expenses recorded in MEPS, b^Y is a factor to allocate the previously estimated aggregate bad debt to individuals, and we is an indicator of whether an individual was uninsured for at least one month. $charge_i^Y$ are the total medical charges facing an individual i and exp_i^Y are the total medical expenses paid by any source. The difference (interacted with the insurance status) is thus a measure of individual bad debt.

A period in our model lasts three years. As such, we want to estimate medical expenses over a three-year period. However, MEPS only follows each panel for two years. To construct medical expenses in the third year while taking into account potential persistence

[29] Uncompensated care is a measure of care for which the hospital received no payment from patient or insurer. It includes bad debts and financial assistance provided by the hospital.

[30] National Health Expenditure Data reports total U.S. hospital spending in 2018 (2019) of \$1122.6 billion (\$1193.7 billion).

[31] National Health Expenditure Data reports total personal health care expenditures in 2018 (2019) of \$3021.8 billion (\$3175.2 billion).

Figure 1: Average Per Capita 3-Year Out-of-Pocket Medical Expenses

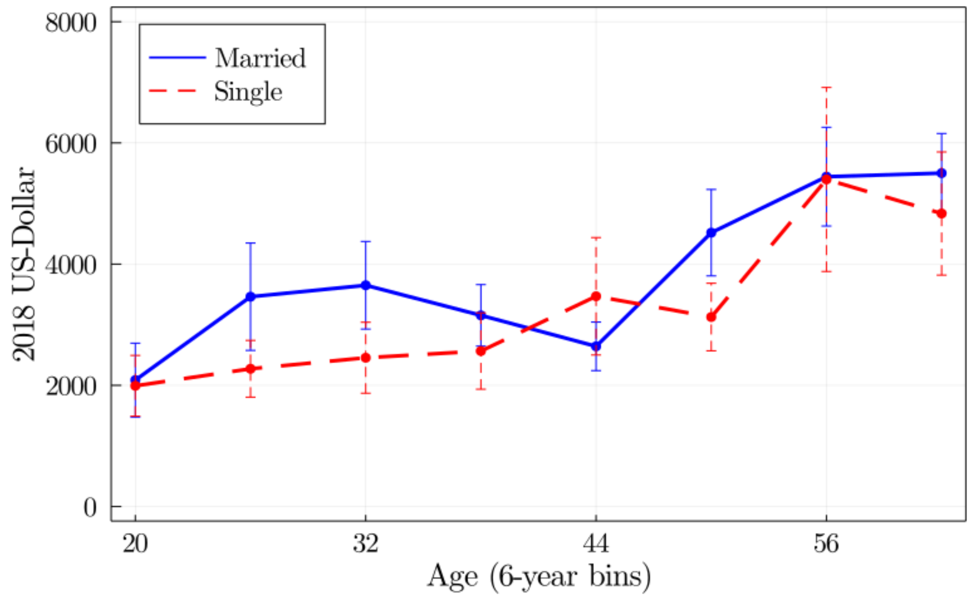

Notes: Confidence bands are 2 SE. Out-of-pocket medical expenses are costs that are not covered by insurance and have to be paid by the patients themselves.

in these expenses we first estimate the following regression:

$$\widetilde{OOP}_i^{Y2} = \alpha + \beta \cdot \widetilde{OOP}_i^{Y1} + \epsilon_i \tag{20}$$

We find estimated values of $\hat{\alpha} = 654.56$ and $\hat{\beta} = 0.47$.

Then, we estimate the third year expenses using these estimated parameters:

$$\widehat{\widetilde{OOP}}_i^{Y3} = \hat{\alpha} + \hat{\beta} \cdot \widetilde{OOP}_i^{Y2} + e_i$$

where e_i is drawn from the residual distribution of (20). The final 3-year expense is the sum of the expenses in year one, two, and three.

Figure 1 plots the estimated average per capita 3-year out-of-pocket medical expenses across six-year age bins. We can see that these expenses amount to thousands of US-Dollars. Unsurprisingly, medical expenses also tend to increase as people get older. Married individuals seem to have slightly higher mean expenses in their late 20s and early 30s compared to singles. Figure 2 offers a more disaggregated view. In the top row, we can see that the higher medical expenses for married individuals are primarily driven by married females. One plausible explanation is that these reflect higher expenses due to pregnancy

Figure 2: Avg. Per Capita 3-Year OOP Medical Expenses - Detailed View

Notes: Confidence bands are 2 SE. Out-of-pocket medical expenses are costs that are not covered by insurance and have to be paid by the patients themselves.

and childbirth.[32]

In our model, expense shocks depend on marital status as well as age and can have three realizations: $\kappa_{status}^{age} = \{0, \kappa_{1,status}^{age}, \kappa_{2,status}^{age}\}$. To translate the estimated medical expenses into our model, we first subset the data across marital status and age. we categorize individuals into married, single (including divorced), and others, as well as six-year age groups. Note that we subset the sample into six-year age bins due to sample size limitations. As a model period lasts three years, two consecutive age groups in the model face the same expense shocks.[33] In order to compute the kind of medical expenses that can trigger default, we focus on the largest estimated expenses and compute the 95th and 98th percentile for each subset. The large shock κ_2 is pinned down by the mean expense of the top 2%. The smaller shock κ_3 is determined by the mean expense of the next 3%. The corresponding shock probabilities in our model are thus:

[32] For additional results showing the distribution of expenses across singles and couples, see Appendix A.3.
[33] For example, 20-22 and 23-25 year old individuals face the same expense shocks in the model.

Figure 3: Estimated Expense Shock Magnitudes

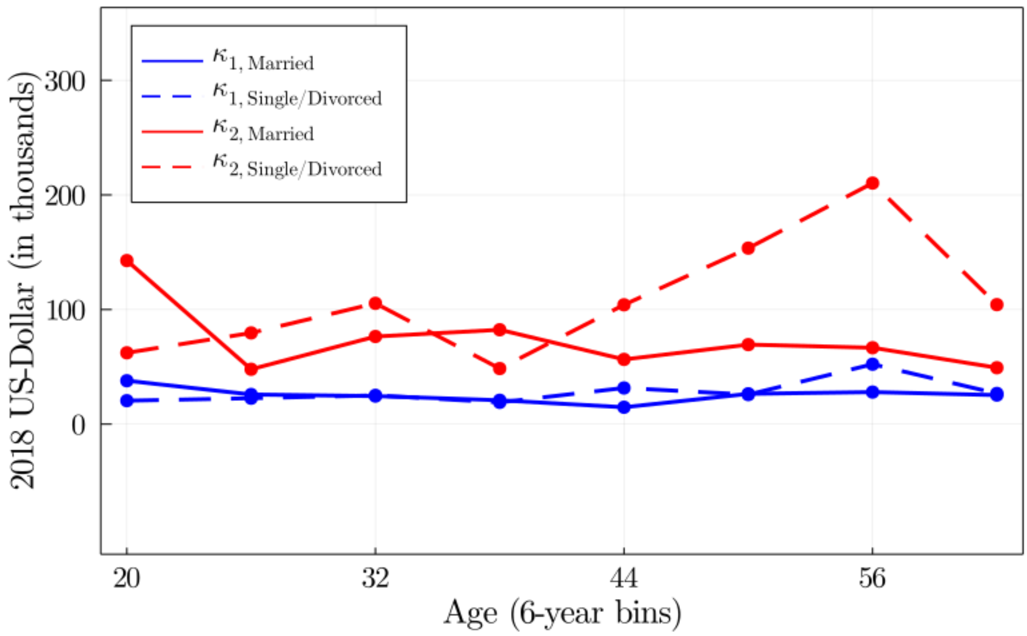

Notes: Source: MEPS (2019). κ_1 is computed as the mean out-of-pocket expenses among the 95th to 98th percentile. κ_2 is computed as the mean out-of-pocket expenses among the 98th to 100th percentile.

$\pi_\kappa = \{0.95, 0.03, 0.02\}$. Figure 3 illustrates the estimated magnitudes of the expense shock process $\kappa^{age}_{status} = \{0, \kappa^{age}_{1,status}, \kappa^{age}_{2,status}\}$. One can see that these large medical expenses can amount to several hundred thousand US-Dollars. In addition, single and divorced individuals have larger expenses compared to married at later ages.[34]

5.2 Model Validation

We evaluate our model fit on a set of untargeted moments that are commonly used in the literature. The results are summarized in Table 4.

Note that while we target the aggregate default rate, fraction of borrowers, and average interest rate in our calibration, the moments for the various subgroups (singles, couples, and divorced) were not used. Regarding default rates, we use data from the American Bankruptcy Institute for the aggregate default rate. However, the institute does not publish default rates across marital status. Instead, we use the Survey of Consumer Finances 2019 to compute default rates across marital status. One can see that our model manages to generate default rates that are higher for singles and in particular divorced households relative to couples as in the data. For the fraction of borrowers, our model overpredicts

[34]For additional results regarding the estimation of expenses, see Appendix A.2.

Moment (in %)	Data	Model
Default rate*		
Default rate, Singles	0.85	1.56
Default rate, Couples	0.84	0.67
Default rate, Divorced	1.36	1.89
Fraction of borrowers		
Fraction of borrowers, Singles	20.2	24.2
Fraction of borrowers, Couples	20.4	18.5
Fraction of borrowers, Divorced	23.3	21.7
Avg. interest rate		
Avg. interest rate, Singles	16.2	17
Avg. interest rate, Couples	16.1	15.7
Avg. interest rate, Divorced	16.8	16.7
Debt-to-Income Ratio (cond. on borrowing)	15.2	25.7
DTI, Singles	14.8	34.5
DTI, Couples	13.6	20.3
DTI, Divorced	20.9	26.8

Table 4: Untargeted Moments

Notes: * The American Bankruptcy Institute does not publish default rates across marital status. Instead, the default rates for the subgroups are computed using the Survey of Consumer Finances 2019.

borrowing by singles compared to the data but replicates the higher need for borrowing among divorced relative to couples. Similarly, our model generates a too high interest rate for singles but reflects how divorced households face more expensive loans than couples.

The debt-to-income (DTI) ratio conditional on borrowing in the data is computed using the SCF 2019. For our measure of debt, we use the same net worth definition following Herkenhoff (2019) described earlier. First, we compute for each household with negative net worth the ratio of net worth and wage income.[35] Then, we winsorize the top 1% of the resulting distribution following Herkenhoff (2019) as we consider these observations as outliers with outsized influence on the mean. Finally, we compute the average ratio for all households with heads aged between 20 and 68. The DTI ratio in the model is similarly derived by computing the ratio of debt to income for all households with negative assets and then averaging. Overall, our model overpredicts indebtedness for all subgroups.

Figure 4 depicts bankruptcy filing rates across age for singles, couples, and divorced

[35] We focus on wage income as our measure of income in the data as this is most in line with our model setup.

Figure 4: Bankruptcy Filing Rates (Normalized) Across Age

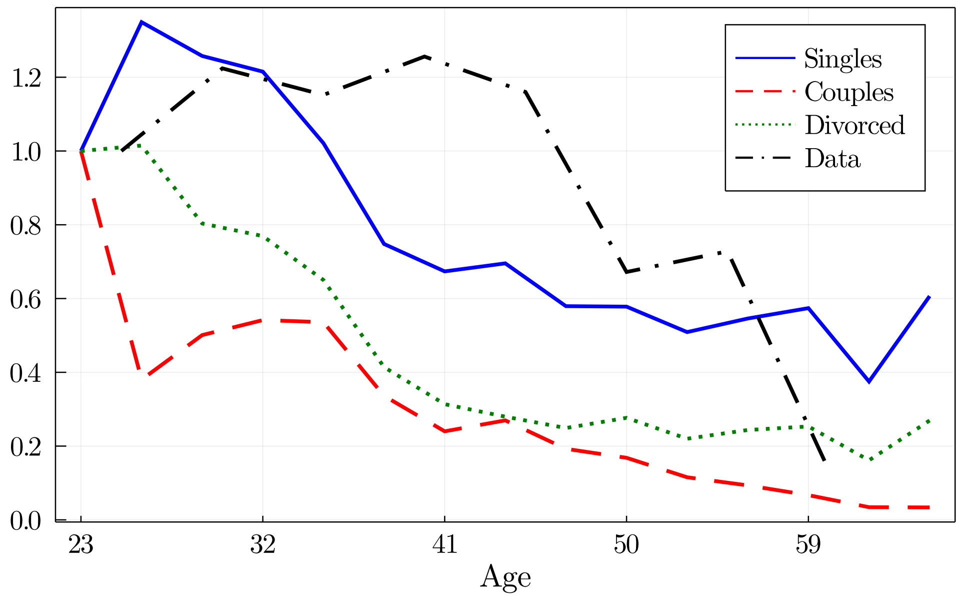

Notes: The series are normalized such that each series starts at one. "Singles", "Couples", and "Divorced" denote the bankruptcy filing rates for the respective household groups in our model after normalization. The data series is taken from Sullivan et al. (2000).

in our model (normalized so that every series starts at one). The data numbers are taken from Sullivan et al. (2000).[36] Note that our model is calibrated to the year 2019, whereas the data come from a much earlier time period (pre 2000). Our model matches the overall decline in bankruptcy rates over the lifetime but overpredicts filings rates for younger households. This is driven by the fact that in the model young households borrow heavily to smooth consumption over their lifetime. This causes an increase in defaults in turn.

6 Welfare Analysis

In this section, we analyze the welfare implications of our model. In particular, we are interested in whether single households and couples differ in their preferred leniency of the bankruptcy regime.

In theory, the welfare-maximizing bankruptcy leniency is unclear ex-ante. There are two opposing forces at work: On the one hand, a more lenient bankruptcy regime (corresponding to a lower garnishment rate in our model) makes bankruptcy less costly for

[36] Unfortunately, it is not possible to plot bankruptcy filing rates across age and marital status using the Survey of Consumer Finances. Bankruptcy is too rare as an event and the sample size of the survey is not large enough.

households. Thus, households can default more cheaply when hit by a bad shock in order to smooth their consumption. In the literature, this channel is commonly referred to as smoothing consumption *over states*. On the other hand, a more lenient bankruptcy regime makes it more likely for households to default c.p. Financial intermediaries anticipate this and will require a higher default premium. As a result, lower default costs lead to higher interest rates on loans and make smoothing consumption *over time* more difficult.[37]

To measure welfare across different garnishment rates we use the ex-ante well-being of single women and men as well as of couple households. That is, we use the following welfare criterion:

$$\begin{aligned} W_S &= \mathbb{E}_{z_g}\left\{V_{S,g,j=0}(a=0, z_g, \kappa=0)\right\} \\ W_C &= \mathbb{E}_{z_f, z_m}\left\{V_{C,j=0}(a=0, z_f, z_m, \kappa_f=0, \kappa_m=0)\right\} \end{aligned} \tag{21}$$

Recall that all newborns ($j = 0$) start life with no assets ($a = 0$) and no expense shock ($\kappa = 0$), but that there is heterogeneity regarding the starting wage (z_g).[38]

6.1 Baseline

In this section, we examine the welfare implications of varying the leniency of the bankruptcy regime by changing the garnishment rate ϕ between 0.1 and 0.9. Figure 5 illustrates the resulting welfare curves for single female and male as well as couple households. We can see that there are sharp differences regarding the preferred bankruptcy leniency between single and couple households. Whereas single households prefer intermediate garnishment rates in the range between 0.3 and 0.5, couples prefer a more lenient regime with a garnishment rate of 0.1.[39]

6.2 Decomposition

What are the channels that drive the heterogeneous welfare implications for single versus couple households in the previous section? Single and couple households differ in a

[37] See also Zame (1993).

[38] As we are only interested in the shape of the welfare curve (the location of the maximum in particular), there is no need to convert this welfare measure into consumption equivalent variation.

[39] The differences between single women and men are driven by heterogeneity between their wage processes.

Figure 5: Welfare - Baseline

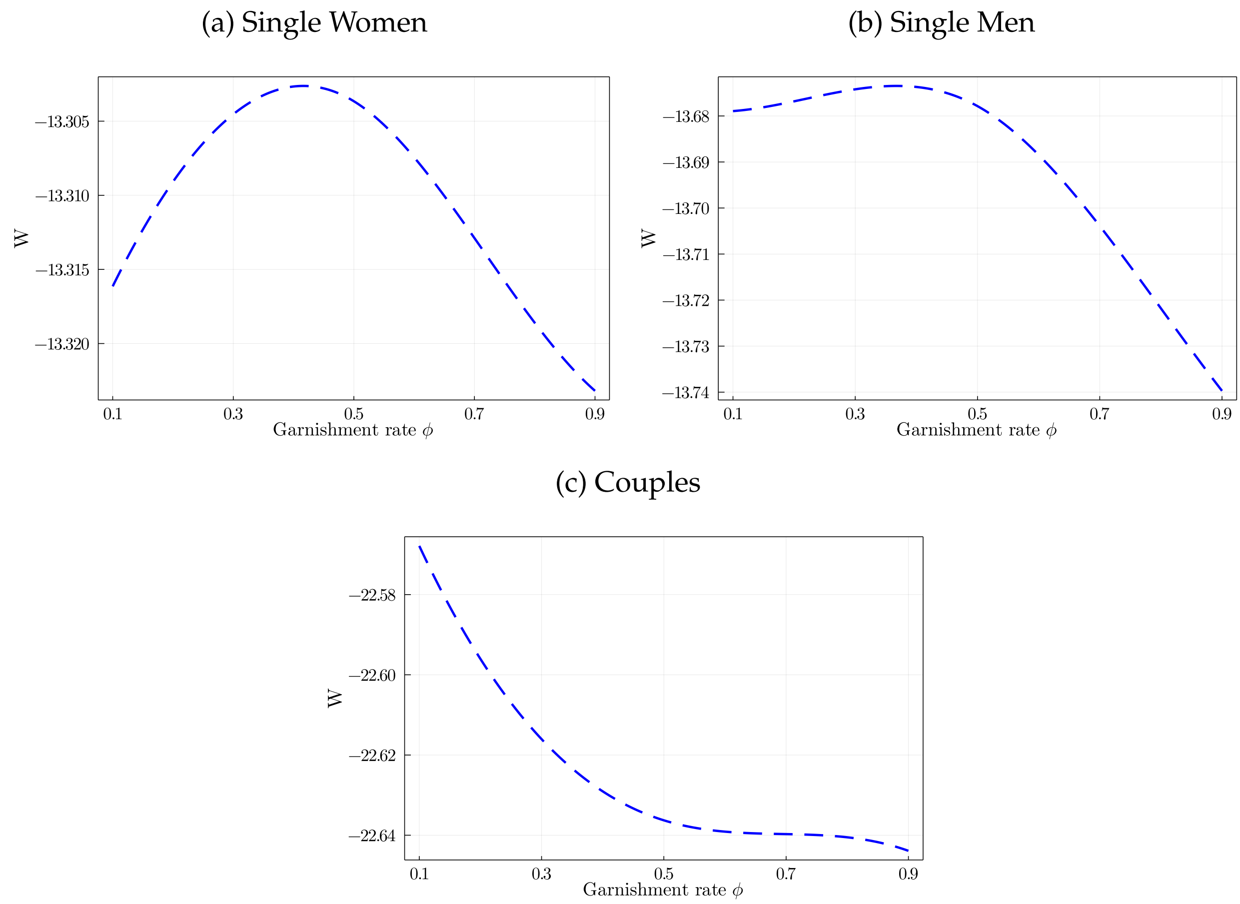

Notes: The welfare measure W is defined in Equation (21).

number of ways:

1. **Income side**: In couple households, each individual is subject to their own wage process. In addition, each spouse can choose his/her labor supply. Single households by nature are only subject to one wage process.

2. **Expense side**: Both spouses in couple households are subject to their own expense shock process. This implies that couple households are hit by at least one expense shock more often than single households.

3. **Divorce channel**: Couple households are subject to a divorce shock, singles are not. The divorce itself is costly.

4. **Economies of scale**: Couples households enjoy economies of scale in consumption, singles do not.

In the following of this section, we separately look at each of these channels in turn

Figure 6: Welfare - Only Income Differences between Singles and Couples

(a) Single Women

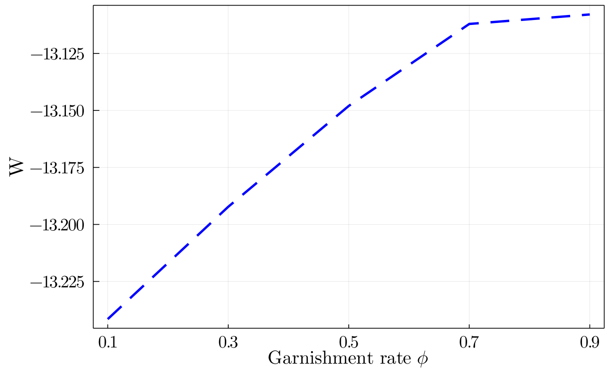

(b) Single Men

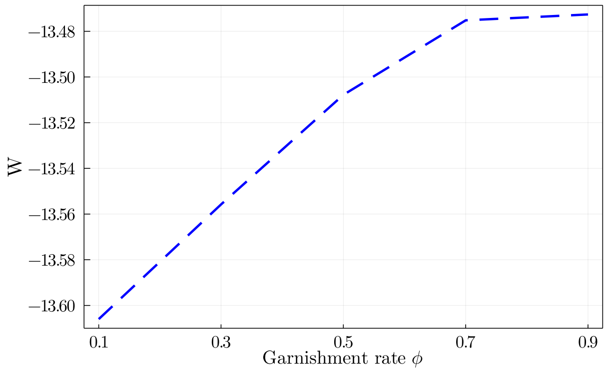

(c) Couples

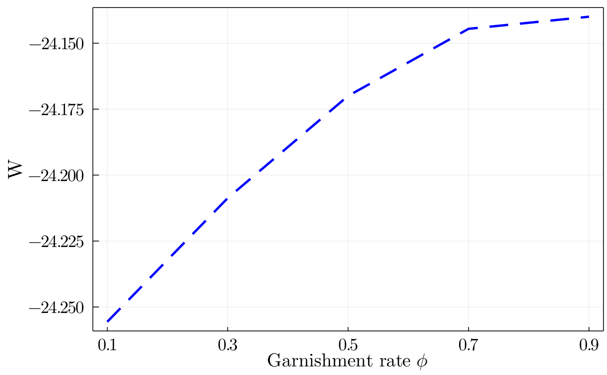

Notes: The welfare measure W is defined in Equation (21). In this experiment, all expense shocks, divorce shocks, and economies of scale in consumption are turned off. The only remaining differences between singles and couples are on the income side.

and examine to what extent they drive the baseline welfare results in Section 6.1.

6.2.1 Income Side

To start off, we first look at the effects of differences on the income side between singles and couples. Recall from Section 4.1 that we model wage processes at an individual level. Thus, singles are subject to one wage process while each individual in a couple is subject to its own wage process.

In order to isolate the income side, we turn off the divorce shock and economies of scale for couples as well as all expense shocks. Figure 6 shows the resulting welfare curves across different garnishment rates for singles and couples. We can see that in this counterfactual all households prefer a very strict bankruptcy regime with a garnishment rate of 0.9. This result is primarily driven by the removal of expense shocks. Without expense shocks, households no longer have any need for bankruptcy to discharge large expenses. As a result, the benefit of a lenient bankruptcy regime is reduced as smoothing over states becomes less important.

There is a large literature that examines risk-sharing within families.[40] One result from this literature is the existence of the so-called *intra-household insurance channel*. Within a family, spousal labor supply acts as an additional insurance mechanism against labor shocks. A spouse can adjust his or her labor supply in response to the partner's wage realization.

This channel is also active in our model. One way to illustrate it is by examining the degree of consumption insurance available to singles versus couples. To do so, we use our calibrated model with the benchmark garnishment rate, where all expense shocks, divorce shocks, and economies of scale remain turned off, to simulate a household panel of marital status, consumption and labor productivity. we then use the simulated data to run the following regression separately for single or married individuals i:[41]

$$\Delta \log(c_{it}) = \delta + \mu \cdot \Delta \log(z_{it}) + \nu_1 \cdot age_{it} + \nu_2 \cdot age_{it}^2 + \epsilon_{it}$$

where we regress log changes in productivity z_{it} on log changes in consumption c_{it} while controlling for age and the square of age. The coefficient of interest is μ which measures the degree to which changes in labor productivity pass-through to changes in consumption. A higher value for μ indicates a higher pass-through and thus a lower degree of consumption insurance. we find that in our model couples enjoy a stronger degree of consumption insurance: The estimated coefficient for singles is $\hat{\mu}_S = 0.88$ whereas the one for couples is $\hat{\mu}_C = 0.43$.

Intuitively, the existence of an additional mechanism to smooth consumption over states for couples should mean that they rely less on default to smooth over states relative to singles. As a result, couples should prefer a stricter bankruptcy regime compared to singles. Figure 7 shows that this is indeed the case. In order to understand the effect of the intra-household insurance channel on the preferred bankruptcy regime it is necessary to get an interior optimum for one of the household types (either singles or couples). To

[40] See Section 2 for an overview.

[41] This regression has been used in the literature to measure the degree of consumption insurance in data by for example Blundell et al. (2008) and in simulated data by Voena (2015) among others.

Figure 7: Welfare - Effect of Intra-Household Insurance Channel

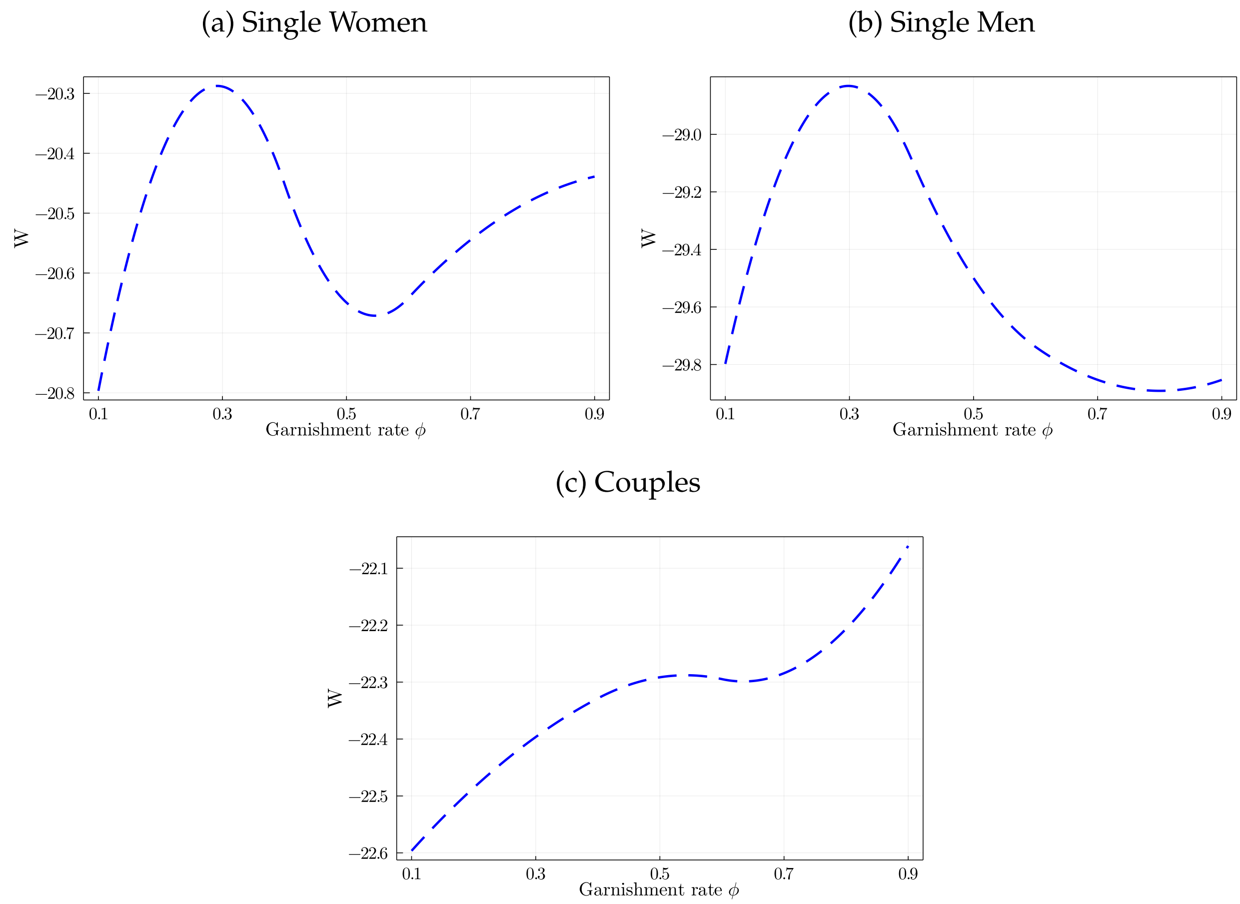

Notes: The welfare measure W is defined in Equation (21). In this experiment, all expense shocks, divorce shocks, and economies of scale in consumption are turned off. The only remaining differences between singles and couples are on the income side. In addition, the standard deviations of the income processes are artificially inflated by a factor of fifteen in order to generate an interior optimum for at least one group.

do so, we artificially multiply the standard deviations of the income processes by fifteen times in Figure 7 compared to Figure 6.[42] We can see that now couples prefer a higher garnishment rate compared to singles, in line with the intuition.

The conclusion from these experiments is that the income side cannot explain why couples prefer a more lenient bankruptcy regime than singles in the baseline. In fact, as Figure 7 shows, differences on the income side should make couples prefer a **stricter** regime relative to singles. As a result, there must be a counteracting force among one (or several) of the remaining channels.

Figure 8: Welfare of Couples - Experiments

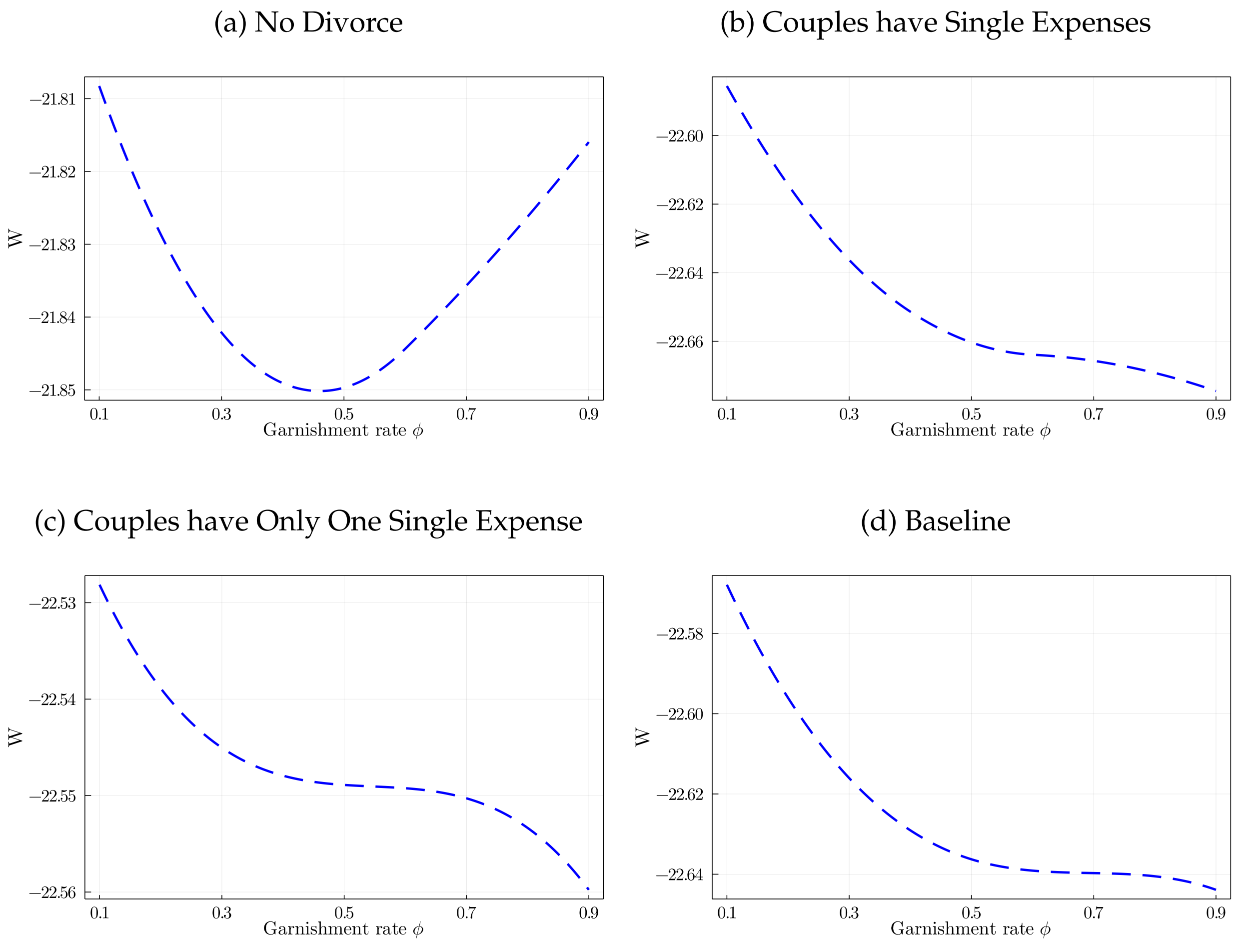

Notes: The welfare measure W is defined in Equation (21). "No divorce" describes the experiment where the divorce shock for couples is turned off. "Couples have single expenses" refers to the experiment where the calibrated expense process for married individuals is replaced by the one for singles. "Couples have only one single expense" refers to the experiment where additionally, couples are subject to only one expense process instead of two.

6.2.2 Expense Side and Divorce

In this section, we investigate to what extent the expense side and the divorce shock can explain the welfare findings from Section 6.1. To do so, we start with the baseline model and eliminate the differences between singles and couples for each channel one at a time. The results are depicted in Figure 8.

To examine the influence of the divorce channel, we turn off the divorce shock for couples by setting the probability of the shock to 0. Figure 8a shows that even after turning off the divorce shock couples still prefer the most lenient bankruptcy regime. Figure 8d depicts the baseline welfare results of couples for comparison. We can see that compared

[42] The exact multiplier is chosen arbitrarily. It only needs to be large enough to generate an interior optimum.

to the baseline case higher garnishment rates become relatively more appealing as the welfare curve displays a U-shape. Thus, while the divorce channel alone cannot explain why couple households prefer a lower garnishment rate than singles, it seems to be one important driver.

The expense side is somewhat more complicated. First, singles differ from couples in terms of their calibrated expense shock process. In particular, the expense shock sizes of a single individual differ from those of a married individual.[43] Moreover, we model expense shocks at the individual level. Single households are thus subject to one expense shock process whereas couples are subject to two. In order to disentangle these two channels, we first replace the expense shock process for married individuals using the one for singles. We call this experiment "Couples have single expenses." Figure 8b shows that the welfare curve in this case looks similar to the baseline case. We thus conclude that differences in the calibration of the expense shock processes are not the main driver behind the welfare findings.

Next, we assume that in addition couples are only subject to one expense shock process instead of two. This means that the expense side of singles is now identical to the one of couples. We name this experiment "Couples have only one single expense." Figure 8c illustrates that couples still prefer the most lenient bankruptcy regime. However, compared to the baseline the welfare curve now is much flatter (notice the different y-axis scaling).[44]

6.2.3 Economies of Scale

One final difference between singles and couples is that couples benefit from economies of scale in consumption. Again starting from the baseline model, we turn off economies of scale in couples' consumption by setting $\eta = 2$. Figure 9 illustrates that this barely affects the shape of the welfare curve across garnishment rates compared to the baseline in Figure 8d. As a result, we conclude that economies of scale are not the feature that drive the baseline welfare results.

[43] See also Section 5.1.

[44] In Appendix A.3 we show the welfare result when we assume that the expense shock processes within couples are perfectly correlated.

Figure 9: Welfare of Couples - No Economies of Scale

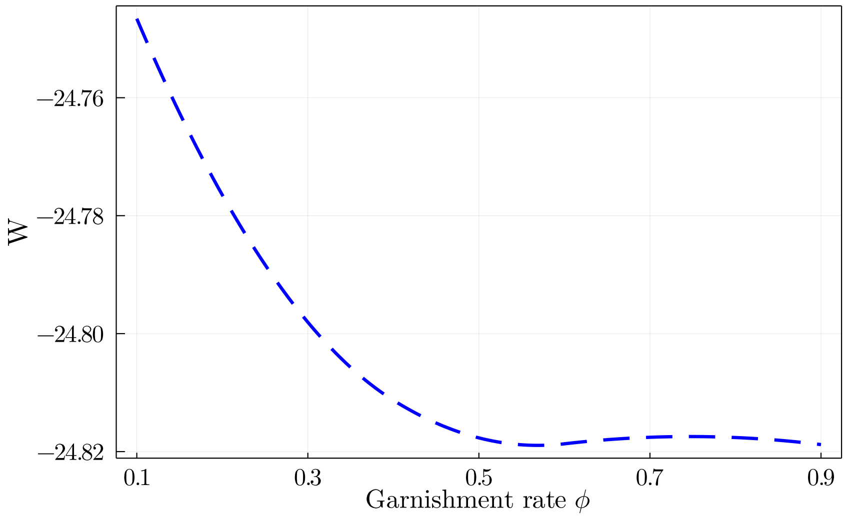

Notes: The welfare measure W is defined in Equation (21). "No economies of scale" refers to the experiment where couples do not benefit from economies of scale in consumption.

6.2.4 Combinations of Different Channels

Putting these results together, we can see that no channel by itself can explain the finding that couples prefer the most lenient bankruptcy regime. Thus, it must be a combination of channels. Figure 10 depicts different combinations of the previous experiments. For example, in Figure 10a we turn off both the the divorce shock as well as the economies of scale in consumption for couples. We can see in Figure 10c that turning off the divorce shock and making the expense side of couples identical to the one for singles is key to explaining why couples prefer the most lenient regime in the baseline.

Table 5 summarizes the previous experiments. The baseline welfare result, that couples prefer more lenient bankruptcy, is driven by two factors: 1) Default is an important insurance mechanism against divorce for couples. This channel does not exist for singles. 2) A more lenient bankruptcy makes it cheaper to default in response to expense shocks. Couples benefit in particular from this channel because they are hit by expense shocks more often.

Regarding the first point, it is unsurprising to see that divorce is a factor that makes couples prefer a more lenient bankruptcy regime. First, a divorce shock itself is costly and can be seen as a type of expense shock. Second, a divorce transforms one couple into two single households. Singles do not have access to intra-household insurance increasing their risk from wage fluctuations. Figure 11 shows how default behavior among house-

Figure 10: Welfare of Couples - Experiments - Combinations

(a) No Divorce and No Economies of Scale

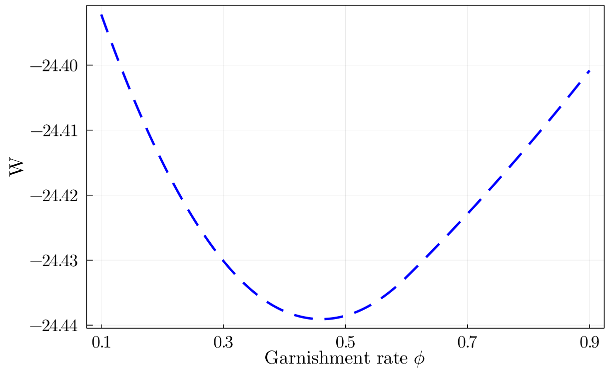

(b) No Scale, Only One Single Expense

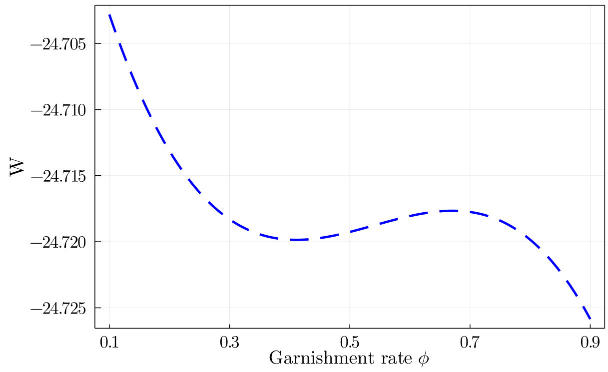

(c) No Divorce, Only One Single Expense

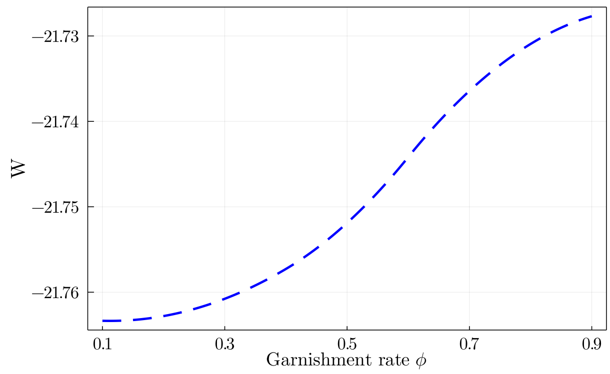

Notes: The welfare measure W is defined in Equation (21). ”Only one single expense” refers to the experiment where the expense side of couples is identical to the one of singles.

holds evolves around divorce. We can see that the default rate quadruples at the time of divorce (event time equal to 0). Naturally, a higher garnishment rate makes default after divorce more costly.

For the second point, we compute the fraction of households that default conditional on receiving any expense shock in the baseline. Figure 12 summarizes the results. In the left figure we depict the fraction of singles that default after receiving any expense shock (blue, solid line), a small expense shock (red, dashed line) or a large shock (green, dash-dotted line). The right figure shows the fraction of couples that default after receiving any expense shock (blue, solid line), one small shock (red, solid), two small shocks (red, dashed), one large shock (green, solid), one large and one small shock (black, dash-dotted), or two large shocks (green, dashed).

For couples the fraction of households that default after receiving one small shock (red,

Experiment	Couples prefer stricter bankruptcy than singles?
No divorce	No
No economies of scale	No
No divorce, no economies of scale	No
Couples have single expenses	No
Couples only one single expense	No
No scale, only one single expense	No
No divorce, only one single expense	Yes

Table 5: Decomposition Exercises

Notes: "Couples have single expenses" refers to the experiment where the calibrated expense process for married individuals is replaced by the one for singles. "Couples only one single expense" refers to the experiment where additionally couples are subject to only one expense process instead of two.

Figure 11: Fraction of Defaulters around Divorce Event

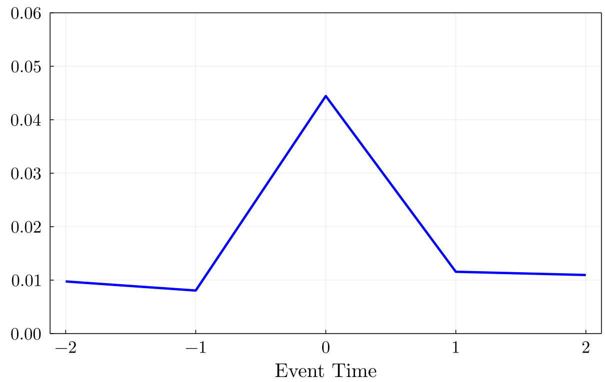

Notes: Event time denotes the time relative to the divorce event. The divorce event happens at time 0.

solid line in right figure) drops strongly when garnishment rates increase from 0.1 to 0.2 (a drop of around 50%) or from 0.1 to 0.3 (a drop of around 75%). Compared to this, the fraction of single households that default after suffering a small expense shock (red, dashed line in left figure) only drops by around 18% when the garnishment rate increases from 0.1 to 0.2 or by around 29% when the rate increases from 0.1 to 0.3.[45] These results suggest that for couples garnishment rates higher than 0.1 quickly limit the usefulness of default to insure against smaller expense shock realizations. As couples get hit by at least one expense shock more often than singles, this makes lenient bankruptcy regimes relatively more attractive to couples than singles.

[45]Similarly, going from the baseline to a garnishment rate of 0.1 the default rate for singles increases by a factor of 3.8, while it increases by a factor of 8.1 for couples.

Figure 12: Fraction of Defaulters cond. on Receiving Expense Shock

(a) Singles (b) Couples

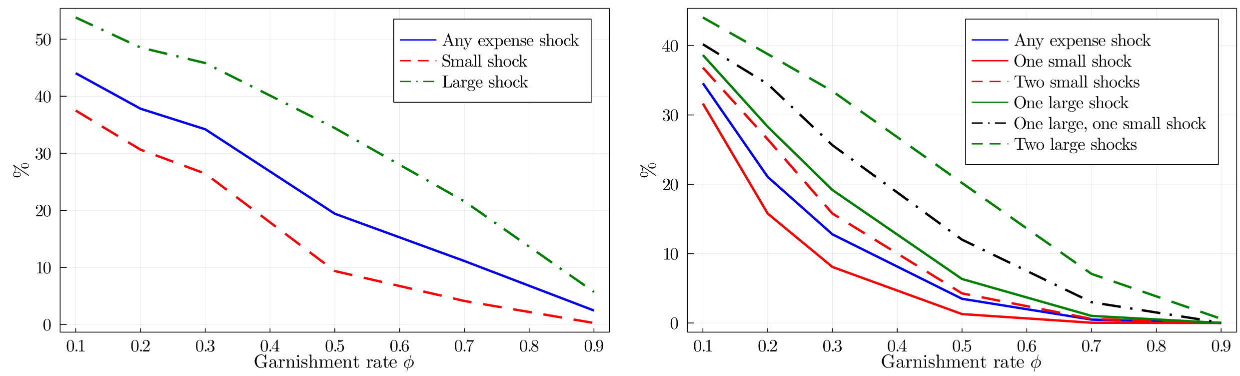

Notes: This figure illustrates the fraction of households that default after receiving a certain combination of expense shocks.

7 Further Robustness Checks

One concern might be that the welfare result in Figure 5 is driven by how we model the cost of bankruptcy. Recall from Section 4.1 that in case of default a fraction ϕ of a household's wage is garnished. In particular, for couples the labor income of both spouses is garnished. Could couples prefer a lower garnishment rate than singles simply due to larger income losses for a given rate?

To examine this possibility, we change the bankruptcy cost from wage garnishment to a fixed cost. To be precise, the budget constraint in case of default now looks as follows for singles:

$$c \leq (e_j \cdot z_g \cdot n) - \phi$$

And for couples:

$$c \leq (e_j \cdot z_f \cdot n_f + e_j \cdot z_m \cdot n_m) - \phi$$

This means that the absolute costs of bankruptcy are now identical for single and couple households. We then vary the bankruptcy cost ϕ from 0.01 to 0.3.[46]

[46] Note that for a fixed bankruptcy cost, it is only possible to solve the model for ϕ up to around 0.3 in our calibration. Beyond that empty budget sets start to appear: In certain states it is impossible for households to either repay or default, as both choices will lead to non-positive consumption. With proportional wage garnishment this can never happen as it is always possible to default.

Figure 13: Welfare - Fixed Bankruptcy Cost

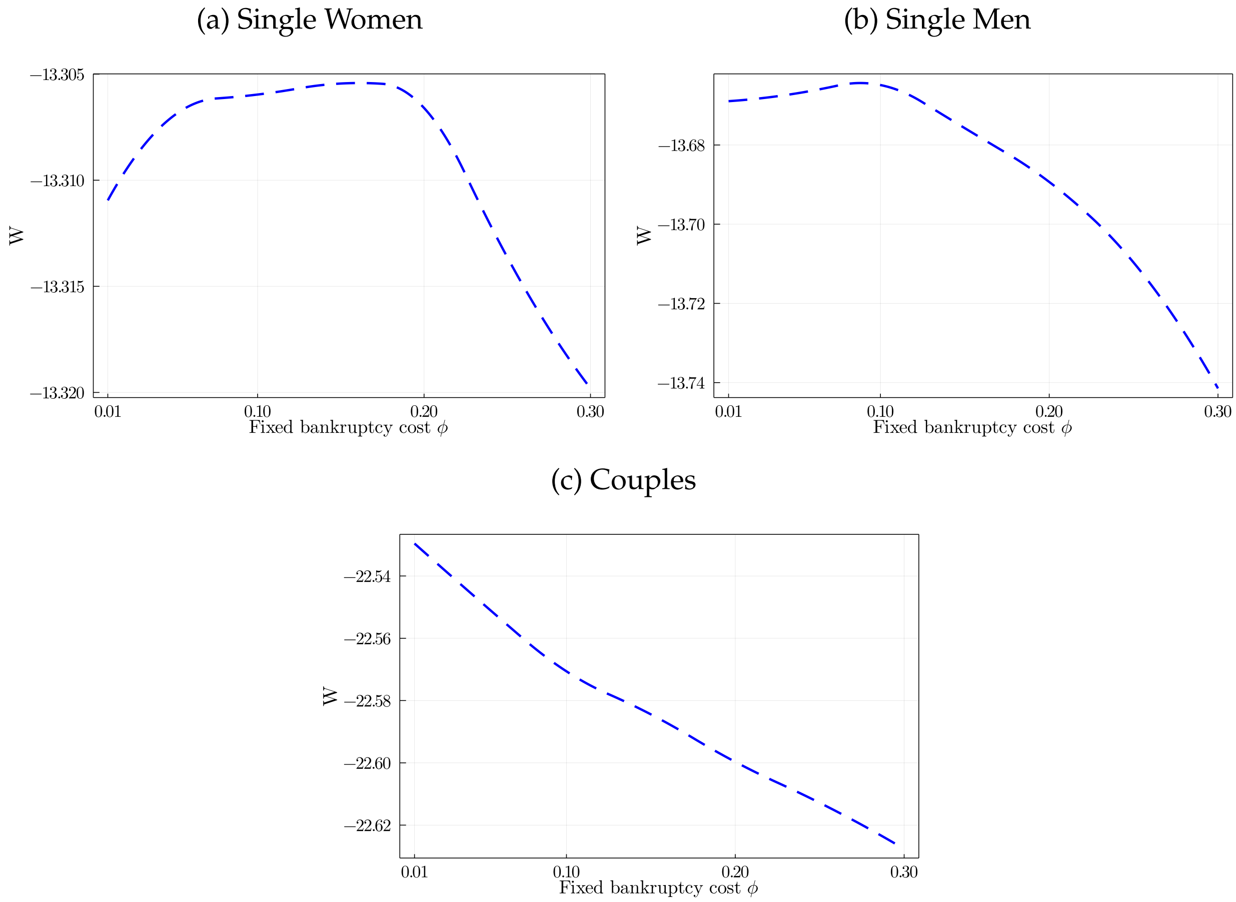

Notes: The welfare measure W is defined in Equation (21). In this experiment the proportional wage garnishment cost of bankruptcy is replaced by a fixed cost that is the same for singles and couples.

Figure 13 illustrates that the previous welfare results still hold in this case. Again, couples prefer a more lenient bankruptcy regime compared to singles.

8 Conclusion

Bankruptcy rates in the U.S. differ strongly across marital status. In particular divorce has been shown to be an important driver of household bankruptcies. However, until now the quantitative consumer default literature has ignored differences in marital status. Work in this literature models all households as a single entity. In this paper, we address this gap and investigate how household marital status affects the welfare implications of bankruptcy regulation. To do so, we build a consumer default model that is the first to explicitly model both single and couple households. In addition, our model also allows for couples to divorce.

Using a calibrated version of our model we examine the welfare effects of different degrees of bankruptcy leniency for singles and couples separately. We find that there are large differences between these two types of households in our model: Couples prefer a more lenient bankruptcy regime compared to singles. We show that this finding is driven mainly by differences between couples and singles on the income and expense sides as well as divorces. In terms of the income side the main distinction is that couples have access to intra-household insurance whereas singles do not. We show that this difference in fact makes couple prefer a stricter bankruptcy regime than singles. However, in contrast to singles couples can also be divorced. Furthermore, couples suffer expenses more often as there are two individuals. These two factors make default more valuable to couples relative to singles and their influence outweighs the income side. The net effect is then that couples prefer a more lenient bankruptcy regime than singles. To summarize, our results suggest that ignoring household heterogeneity across marital status in consumer default models may not be an innocuous choice for welfare analysis and thus policy experiments.

One natural extension of the model would be to endogenize the divorce decision of couples.[47] In the current model, a couple may get divorced even when one spouse would be badly off after divorce. On the one hand, in a model with endogenous divorce such a spouse may instead choose to compensate his/her partner by adjusting the allocation of resources within marriage. This could lessen the need for bankruptcy after divorce and make couples prefer a stricter bankruptcy regime. On the other hand, endogenous divorce may increase precautionary savings of couples as illustrated in Doepke and Tertilt (2016). This increase would make the higher interest rates for borrowing associated with more lenient bankruptcy regimes less costly.

Another promising avenue for future research could be to allow for separate asset holdings in couple households. This extension would enable researchers to model separate bankruptcy filings within couples.[48] Separate asset holdings could lead to interesting situations where couples allocate assets and debt strategically. For instance, couples may allocate less debt/more savings to the spouse who would be worse off in divorce.[49] This allocation could make bankruptcy after divorce less important.

[47] Endogenous divorce can be modeled in a limited commitment and endogenous bargaining framework such as Voena (2015).

[48] As a result, common law states can be sensibly modeled in such a framework. See also Section 3.

[49] Separate assets would require the researcher to specify how assets are divided upon divorce. One possibility is title-based distribution under which assets are divided according to the title of ownership. Another option is equitable distribution.

It could also be interesting to include marriage in the model. This addition could help the model to better match the data. Adding an exogenous marriage shock would be relatively straightforward. Such a change could make singles prefer a more lenient bankruptcy regime than in the current baseline, as they anticipate turning into a couple household later in life. At the same time, the prospect of marriage may also change the behavior of singles. They may want to borrow more in younger ages as they count on benefiting from economies of scale in consumption later in life as a couple. More borrowing could make singles prefer a stricter bankruptcy regime with lower interest rates.

A further interesting angle could be to examine this model through the lens of gender equality. The empirical literature has highlighted that it is often women with children who end up in precarious financial situations after divorce. Extending the model in this paper by adding children and childcare would allow researchers to model this problem as well as search for policy interventions in the bankruptcy law space that could alleviate this issue.

References

AGARWAL, S., S. CHOMSISENGPHET, AND C. LIU (2011): "Consumer bankruptcy and default: The role of individual social capital," *Journal of Economic Psychology*, 32, 632—-650.

ALON, T., M. DOEPKE, J. OLMSTEAD-RUMSEY, AND M. TERTILT (2020): "This Time It's Different: The Role of Women's Employment in a Pandemic Recession," CEPR Discussion Paper No. 15149.

AMERICAN HOSPITAL ASSOCATION (2020): "Fact Sheet: Uncompensated Hospital Care Cost," Available at: https://www.aha.org/fact-sheets/2020-01-06-fact-sheet-uncompensated-hospital-care-cost (Accessed: 04.04.2022).

ATTANASIO, O., H. LOW, AND V. SÁNCHEZ-MARCOS (2005): "Female Labor Supply as Insurance against Idiosyncratic Risk," *Journal of the European Economic Association*, 3, 755—-764.

BACHER, A. (2021): "Housing and Savings Behavior across Family Types," Available at: https://annikabacher.github.io/Bacher_JMP.pdf (Accessed: 14.03.2022).

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921.

BORELLA, M., M. DE NARDI, AND F. YANG (2018): "The aggregate implications of gender and marriage," *The Journal of the Economics of Ageing*, 11, 6–26.

——— (2021): "Are Marriage-related taxes and social security benefits holding back female labor supply?" NBER Working Paper No. 26097.

CHATTERJEE, S., D. CORBAE, K. P. DEMPSEY, AND J.-V. RÍOS-RULL (2020): "A Quantitative Theory of the Credit Score," NBER Working Paper No. 27671.

CHATTERJEE, S., D. CORBAE, M. NAKAJIMA, AND J.-V. RÍOS-RULL (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75, 1525–1589.

DE NARDI, M., E. FRENCH, J. B. JONES, AND R. MCGEE (2021): "Why Do Couples and Singles Save During Retirement?" NBER Working Paper No. 28828.

De Nardi, M., E. French, J. B. Jones, R. McGee, and R. Rodgers (2020): "Medical Spending, Bequests, and Asset Dynamics Around the Time of Death," NBER Working Paper No. 26879.

Doepke, M. and M. Tertilt (2016): "Families in Macroeconomics," *Handbook of Macroeconomics*, 2B, 1789—-1891.

Exler, F. (2019): "Personal Bankruptcy and Wage Garnishment," Available at: https://sites.google.com/site/floexler/research (Accessed: 07.11.2021).

Fay, S., E. Hurst, and M. J. White (2002): "The Household Bankruptcy Decision," *American Economic Review*, 92, 706—-718.

Fisher, J. D. (2019): "Who Files for Personal Bankruptcy in the United States?" *The Journal of Consumer Affairs*, 53, 2003—-2026.

French, E. and J. B. Jones (2004): "On the Distribution and Dynamics of Health Care Costs," *Journal of Applied Econometrics*, 19, 705—-721.

Gourinchas, P.-O. and J. A. Parker (2002): "Consumption Over the Lifecycle," *Econometrica*, 70, 47—-89.

Guner, N., R. Kaygusuz, and G. Ventura (2012): "Taxation and Household Labour Supply," *Review of Economic Studies*, 79, 1113–1149.

Herkenhoff, K. F. (2019): "The Impact of Consumer Credit Access on Unemployment," *Review of Economic Studies*, 86(6), 2605–2642.

Hubbard, G., J. Skinner, and S. P. Zeldes (1994): "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy*, 40, 59—-125.

Kotlikoff, L. J. and A. Spivak (1981): "The Family as an Incomplete Annuities Market," *Journal of Political Economy*, 89, 372—-391.

Livshits, I., J. MacGee, and M. Tertilt (2003): "Consumer Bankruptcy: A Fresh Start," Federal Reserve Bank of Minneapolis Working Paper No. 617.

——— (2007): "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 97, 402–418.

——— (2010): "Accounting for the Rise in Consumer Bankruptcies," *American Economic Journal: Macroeconomics*, 2, 165–193.

ORTIGUEIRA, S. AND N. SIASSI (2013): "How important is intra-household risk sharing for savings and labor supply?" *Journal of Monetary Economics*, 60, 650–666.

PALUMBO, M. G. (1999): "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle," *Review of Economic Studies*, 66, 395—-421.

SHORE, S. H. (2010): "For Better, For Worse: Intrahousehold Risk-Sharing over the Business Cycle," *The Review of Economics and Statistics*, 92, 536—-548.

SULLIVAN, T. A., E. WARREN, AND J. L. WESTBROOK (2000): *The Fragile Middle Class: Americans in Debt*, New Haven and London: Yale University Press.

VOENA, A. (2015): "Yours, Mine and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review*, 105, 2295–2332.

ZAME, W. R. (1993): "Efficiency and the Role of Default When Security Markets are Incomplete," *American Economic Review*, 83, 1142–1164.

A Appendix

A.1 Computational Details

Variable	Symbol	# Points	Range
Wage processes			
Women	z_f	5	$\{0.42, 0.64, 1.0, 1.55, 2.41\}$
Men	z_m	5	$\{0.33, 0.58, 1.0, 1.73, 2.99\}$
Lifecycle productivity profile	e_j	16	$\{0.77, 0.82, ..., 1.01, 0.95\}$
Initial distribution of newborn wage		5	$\{0.2, 0.2, 0.2, 0.2, 0.2\}$
Expense shocks	κ	$3 \times 2 \times 16$	
Assets	a	240	$[-5.0, 13.0]$
Labor choice	n	3	$\{0, 0.5, 1\}$

Table 6: Grids Used for Model Computation

Notes: Expense shocks differ across marital status (single vs. married) and age (16). In each subgroup there are three possible expense shock realizations.

Table 6 summarizes the computational grids used in the solution of the model. The grids for the wage processes are discretized into two 5-state Markov processes using the Rouwenhorst method. The asset grid is equally spaced with 100 grid points in the negative space and 140 points in the positive one. The asset grid is chosen sufficiently wide such that no household will ever pick a point on the boundary of the grid.

We solve the model backward. Starting in the final period of life, we first compute the value functions for singles, divorced, and couples. Note that households are not allowed to die with debt. Thus, we do not need the loan pricing functions to solve for the value functions in the final period. Using the computed value functions we can derive the repayment choice for every state. We can use these choices to compute repayment probabilities which in turn give me the loan pricing schedules in the second-to-last period. Using these pricing functions we can again compute the value functions in this period. We iterate this procedure backward until we reach the first period of life.

A.2 Expense Shock Estimation Details

Figure 14 depicts histograms of annual out-of-pocket medical expenses as reported in MEPS 2019 for single and married individuals across different age groups.[50] We winsorize the top 5% of observations in each subgroup.

Figure 15 depicts the distribution of bad medical debt across ages for singles/divorced and married. More specifically, we plot the difference in charges versus expenditures for 2019: $(charge_i^{2019} - exp_i^{2019})$ in Equation (19). One can see that single individuals have more bad debt than married ones in the age groups 44-49 and 56-61. Note that we winsorize the top 5% of observations in each subgroup.

A.3 Additional Results

Figure 16 depicts lifecycle profiles of average per capita consumption, labor supply, income and assets in the baseline model. Figure 17 shows interest rates across age and marital status in the baseline model as well as in the data.

In Figure 18 we show the welfare results for the counterfactual in which, starting from the baseline, we assume that the expense shock processes within couples are perfectly correlated. We can see that couples prefer the most lenient bankruptcy regime in this case.

In Figure 19 we plot the welfare results for married women and men individually across garnishment rates ϕ. We compute the value of individuals within couples by taking the optimal consumption, labor, asset and default policy functions derived from the couples' problem as given. Let $c^* \equiv c^*(a, z_f, z_m, \kappa_f, \kappa_m)$, $a^* \equiv a^*(a, z_f, z_m, \kappa_f, \kappa_m)$, and $l_g^* \equiv l_g^*(a, z_f, z_m, \kappa_f, \kappa_m)$ denote these policy functions. Use $V_{C,j}^g$ to denote the value of a married individual with gender g and age j.

The value is then given by:

$$V_{C,j}^g(a, z_f, z_m, \kappa_f, \kappa_m) = u\left(\frac{c^*}{\eta}, l_g^*\right)$$
$$+ \beta \cdot \Bigg((1-\psi) \cdot \mathbb{E}_{z'_f, z'_m, \kappa'_f, \kappa'_m} \left\{ V_{C,j+1}^g(a^*, z'_f, z'_m, \kappa'_f, \kappa'_m) \right\}$$

[50] Note that these expenses are the out-of-pocket expenses as recorded in MEPS 2019 (adjusted for consistency with aggregate data). They do not include estimated bad debt.

$$+ \psi \cdot \mathbb{E}_{z'_g, \kappa'_g} \left\{ V_{Div,g,j+1}\left(\frac{a^*}{2}, z'_g, \kappa'_g\right) \right\} \Bigg)$$

To measure the welfare of individuals within couples, we again use ex-ante well-being:

$$W_C^g = \mathbb{E}_{z_f, z_m} \left\{ V_{C,j=0}^g(a = 0, z_f, z_m, \kappa_f = 0, \kappa_m = 0) \right\} \tag{22}$$

where W_C^g denotes the ex-ante value of a married individual with gender g.

In Figures 19a and 19b we can see that married women prefer a more lenient bankruptcy regime than married men. This result is driven by differences in female and male labor supply across garnishment rates as shown in Figures 19c and 19d.

Figure 14: Histograms of Out-of-Pocket Expenses (in 2018 US-Dollars)

(a) Singles Aged 32-37

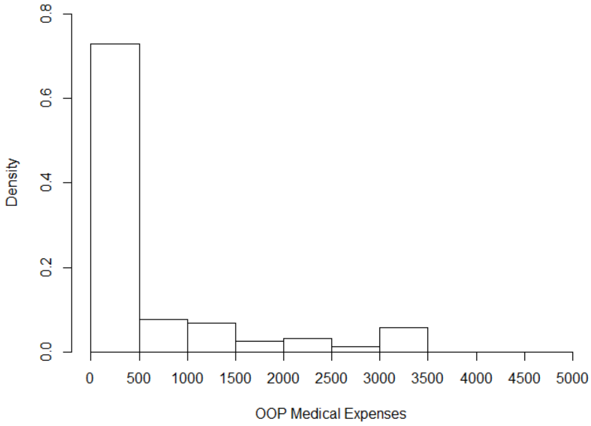

(b) Married Aged 32-37

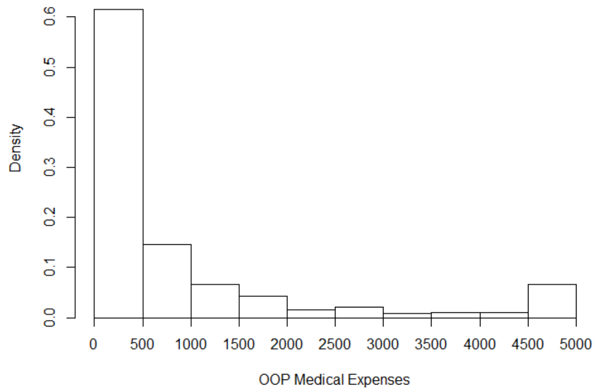

(c) Singles Aged 44-49

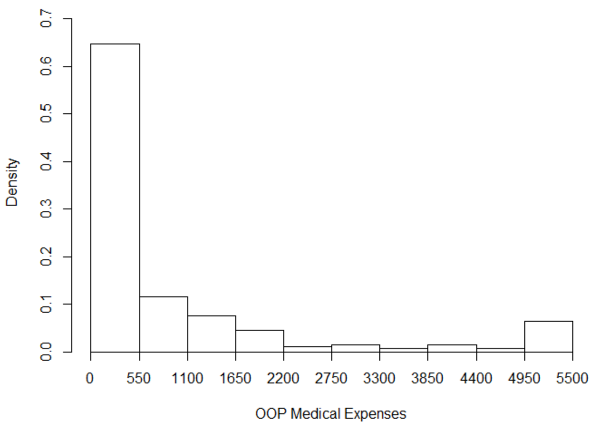

(d) Married Aged 44-49

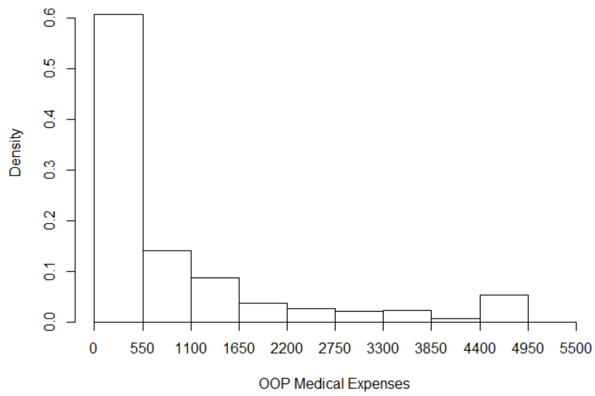

(e) Singles Aged 56-61

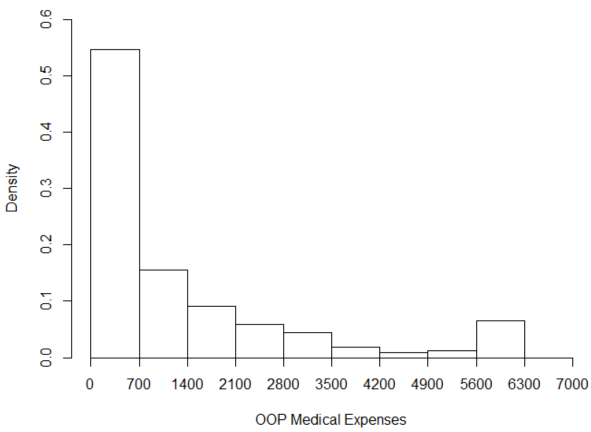

(f) Married Aged 56-61

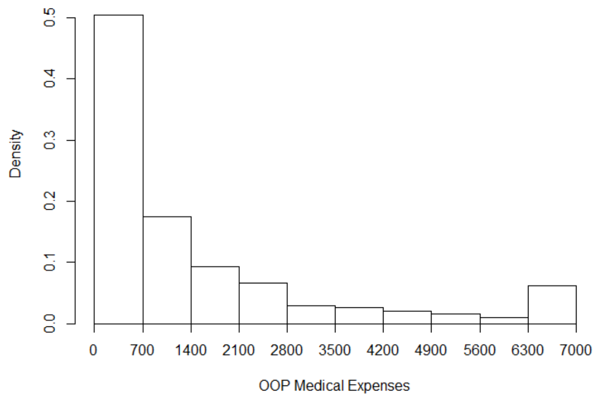

Notes: Source: MEPS (2019). The top 5% of observations in each subgroup are winsorized.

Figure 15: Histograms of Individual Bad Debt

(a) Singles Aged 32-37

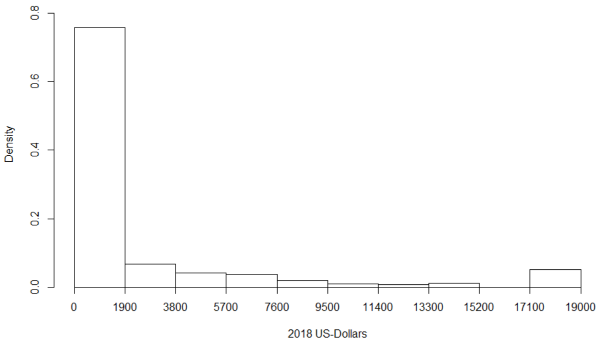

(b) Married Aged 32-37

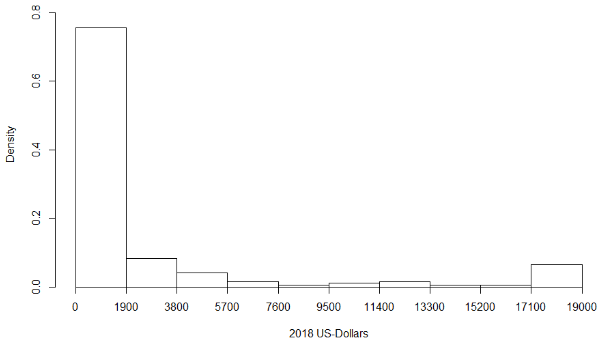

(c) Singles Aged 44-49

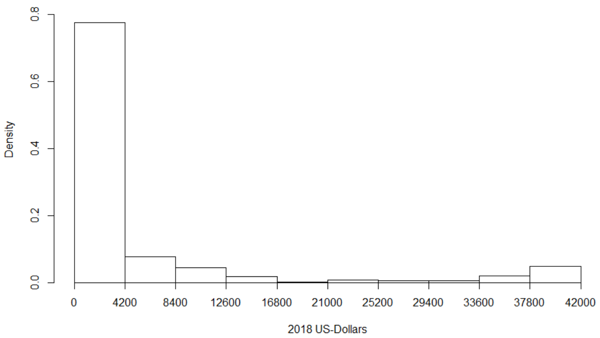

(d) Married Aged 44-49

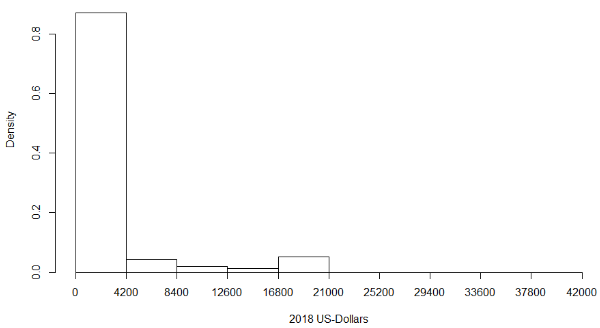

(e) Singles Aged 56-61

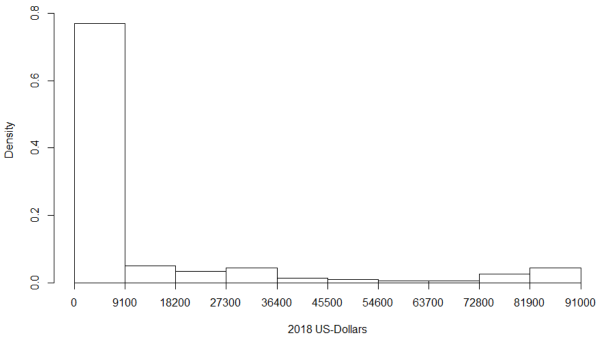

(f) Married Aged 56-61

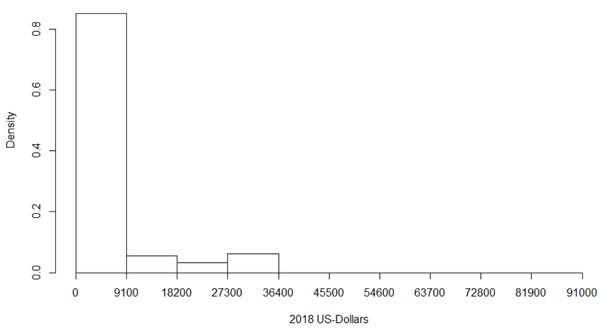

Notes: Source: MEPS (2019). Bad debt is measured as the difference between total charges and total expenditures. See also Equation (19). The top 5% of observations in each subgroup are winsorized.

Figure 16: Model Lifecycle Profiles

(a) Average Per Capita Consumption

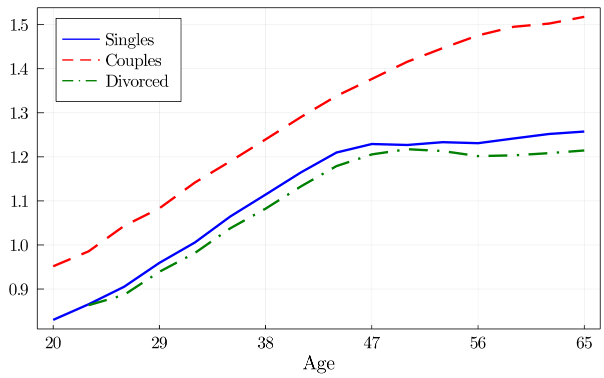

(b) Average Per Capita Labor Supply

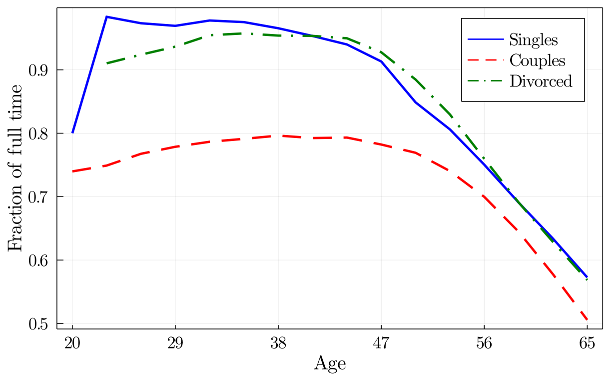

(c) Average Per Capita Labor Income

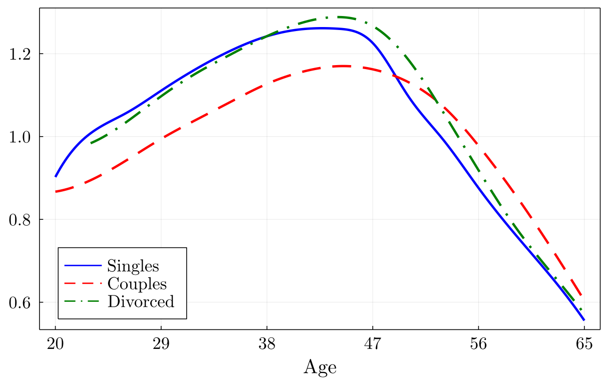

(d) Average Per Capita Assets

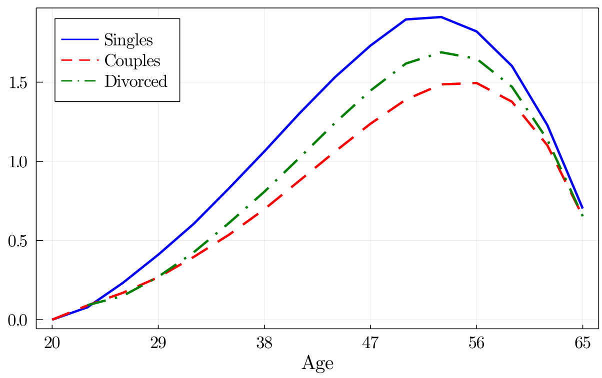

Notes: Since our model has a period length of three years, the x-axis denotes 3-year age brackets.

Figure 17: Interest Rates Across Age

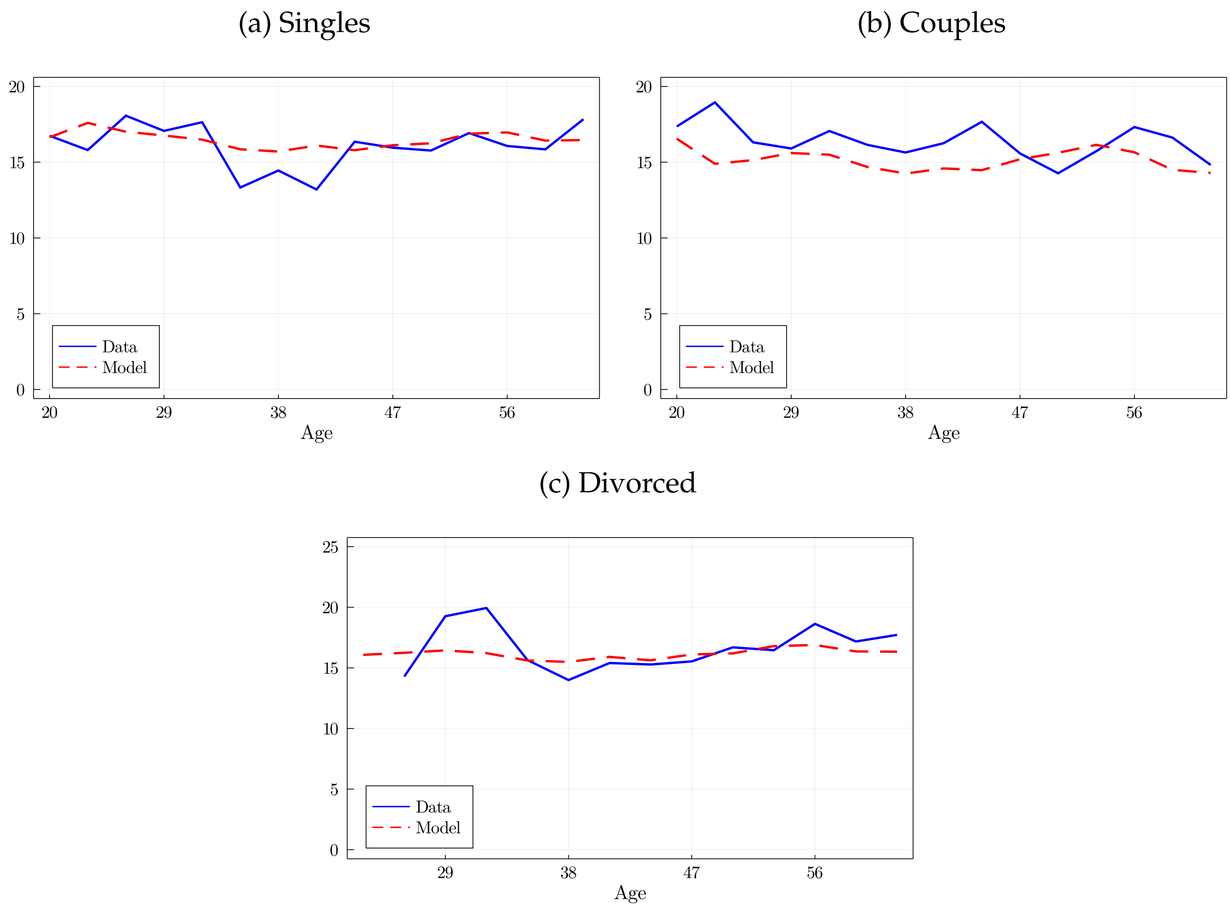

Notes: Source for data: SCF 2019. Since our model has a period length of three years, the x-axis denotes 3-year age brackets.

Figure 18: Welfare of Couples - Perfectly Correlated Expenses

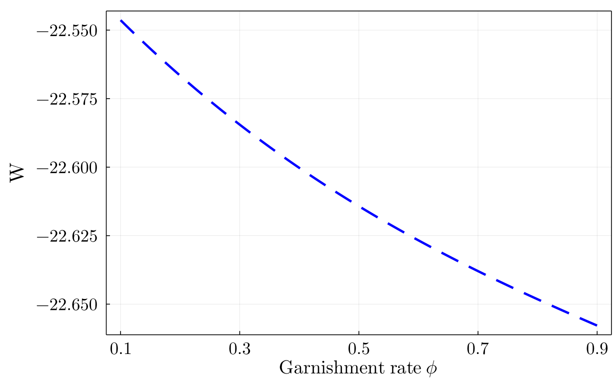

Notes: The welfare measure W is defined in Equation (21). In this experiment the expense shock processes within couples are assumed to be perfectly correlated.

Figure 19: Within Couple Results

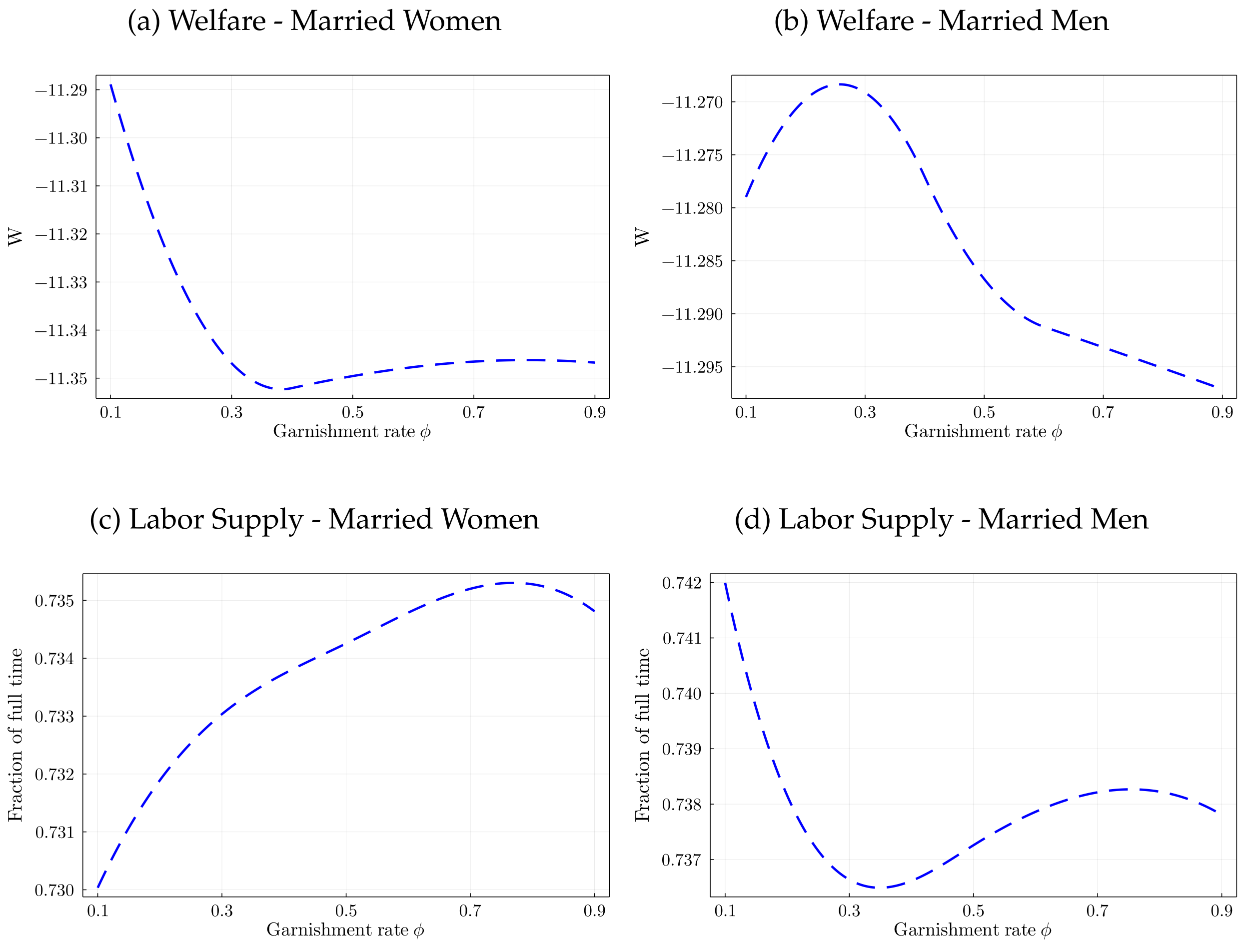

Notes: The welfare measure W is defined in Equation (22). These plots show the results for individual welfare and labor supply of married women and men.