

How Important is Intra-household Risk Sharing for Savings and Labor Supply?*

Salvador Ortigueira[†]
Universidad Carlos III de Madrid

Nawid Siassi[‡]
European University Institute

September 23, 2011

Abstract

While it is recognized that the family is primarily an institution for risk sharing, little is known about the quantitative effects of this informal source of insurance on savings and labor supply. In this paper, we present a model where workers (females and males) are subject to idiosyncratic employment risk and where capital markets are incomplete. A household is formed by a female and a male, who make collective decisions on consumption, savings and labor supplies. We find that intra-household risk sharing has its largest impact among wealth-poor households. While the wealth-rich use mainly savings to smooth consumption across unemployment spells, wealth-poor households rely on spousal labor supply. For instance, for low-wealth households, average hours worked by wives of unemployed husbands are 8% higher than those worked by wives of employed husbands. This response in wives' hours makes up 9% of lost family income. We also study the crowding out effects of public unemployment insurance on other sources of private insurance, and consumption losses upon an unemployment spell.

Keywords: Intra-household risk sharing; Collective household model; Idiosyncratic unemployment risk; Incomplete markets; Precautionary motive.

JEL Classification Numbers: D13, D91, E21.

*We thank Claudio Michelacci, Josep Pijoan-Mas, Gianluca Violante, an anonymous referee and an Associate Editor for this journal for helpful comments and suggestions. Salvador Ortigueira gratefully acknowledges financial support from the Spanish Ministerio de Ciencia e Innovación under grant 2011/00049/001. Nawid Siassi thanks the Economics Department at New York University for its hospitality while working on this project.

[†]Address: Economics Department, Calle Madrid 126, 28903 Getafe (Madrid), Spain. E-Mail: Salvador.Ortigueira@uc3m.es.

[‡]Address: Economics Department, Via della Piazzuola 43, 50133 Florence, Italy. E-Mail: Nawid.Siassi@eui.eu

1 Introduction

The lack of a formal, private insurance market against employment risk makes this type of risk different from most of others faced by individuals. Even though public, compulsory unemployment insurance schemes are present in many countries, they typically fall short of providing full insurance and workers must rely on self-insurance and on informal insurance mechanisms in order to smooth consumption across unemployment spells. Precautionary savings is the main instrument individuals can use as self-insurance against employment risk. The family, on the other hand, is the main informal insurance mechanism available to individuals, as information and payment enforceability are better within than between households.¹

In this paper, we present an incomplete markets economy with idiosyncratic employment risk and assess quantitatively the role of the family as provider of insurance. Intra-household risk sharing, more than any other informal insurance mechanism, has important behavioral implications that affect not only the demand for self insurance, but also how this is crowded out by public insurance programs. Indeed, recent empirical evidence on patterns of insurance against employment risk sheds light on these crowding out effects. More specifically, using a large panel of U.S. households, Cullen and Gruber (2000) and Engen and Gruber (2001) estimate the response in two forms of insurance —accumulation of financial assets and spousal labor supply— to changes in the generosity of public unemployment insurance (UI) and find significant crowding out effects on both. The extent to which public insurance crowds out other forms of (private) insurance is of paramount importance for public policy assessment.²

Our model economy consists of a large number of two-person households, each pooling risks and making *collective* decisions on individual consumptions, labor supplies and joint savings in a risk-free asset, subject to a borrowing constraint. Risk sharing within the two-person household is assumed to be efficient, with individual weights in the household's utility function determined, among other variables, by their relative earning ability. There is a firms sector producing a homogeneous good with capital and labor services, and a government providing UI.

In order to assess the consequences of within-household risk sharing, the equilibrium in this economy is compared to that arising in an economy where individuals lack access to insurance from the family and are left with self-insurance and UI as their only instruments to cope with employment risk. This latter framework corresponds to a standard Aiyagari-Huggett economy augmented with a labor-leisure choice, which has been studied by, e.g., Flodén and Lindé (2001),

¹Blundell *et al.* (2008) estimate the degree of consumption insurance from U.S. data and find evidence that the family plays an important insurance role. See also Shore (2010) for an empirical analysis of family risk sharing over the business cycle.

²See, e.g., Attanasio and Rios-Rull (2000), Golosov and Tsyvinski (2007) and Chetty and Saez (2010) for analyses on the optimal level of social insurance when other forms of private insurance are also available.

Marcet *et al.* (2007) and Pijoan-Mas (2006), among others. In addition to serving as a reference against which our collective household economy can be compared, the Aiyagari-Huggett model is shown to largely overestimate the measure of precautionary savings against employment risk found by Engen and Gruber (2001). This is an indication that this latter model embodies far less insurance opportunities than those available to U.S. households.

The equilibrium of our model economy contains a distribution of households over financial assets and spouses' employment status, and we can thus assess the effects of intra-household risk sharing for different groups of households. Since intra-household risk sharing has its largest impact among low-wealth households, special attention is devoted to households with financial assets below two months of average income. This is the group that Zeldes (1989) termed as liquidity-constrained households, and that represent almost 20% of total U.S. households in 2001.

Our first question addresses the extent to which our model with family insurance can account for patterns of insurance against employment risk observed across U.S. households. The first exercise mimics the empirical work conducted by Engen and Gruber (2001), which regresses households' assets on UI and finds an elasticity of the assets-to-income ratio with respect to the generosity of UI equal to -0.28 . This elasticity provides a measure of the precautionary savings motive, and also of the crowding out effects of UI on households' financial wealth. In a calibration version of our model with family insurance this elasticity is -0.38 , in contrast to an elasticity of -0.64 we find in the Aiyagari-Huggett model. This suggests that abstracting from risk sharing at the level of the household introduces an important bias in this elasticity.

We then inquire on the use of spousal labor supply as an insurance mechanism, and on the magnitude of consumption losses upon an unemployment spell. Our first finding is that while the wealth rich use savings to smooth consumption across unemployment spells, wealth-poor households rely on spousal labor supply. For instance, in the group of liquidity-constrained households (using Zeldes' (1989) definition), average hours worked by wives of unemployed husbands are 8% higher than those worked by wives of employed husbands. This response in wives' hours makes up 9% of lost family income. The implications of family insurance for consumption losses upon unemployment are sizable, especially for low-wealth households. For instance, we find that a liquidity-constrained male lacking insurance from the family suffers a consumption loss of 30%, while the same male suffers a loss of 8% when given access to family insurance. When we compute the fraction of household income loss that translates into household consumption loss for liquidity-constrained households, we find 35% under no family insurance, and 17% with family insurance.

We include two applications and one extension of our model. First, we assess the value of intra-household risk sharing. As suggested by our results above, this value is higher for individuals in

low-wealth households. As an example, the value of family insurance to an unemployed individual with an employed spouse and no assets represents more than 5% of per period consumption of a similar individual with no family insurance. This value decreases with the level of wealth. Second, we study the consequences of family insurance for optimal UI and find that it is one of its key determinants. If we compute optimal UI for households with different wealth levels, instead of focusing on ex-ante utility, we find that family insurance creates a wedge in optimal replacement rates that decreases with wealth. For instance, the optimal replacement rate for the average two-person household with no assets is 15%, while this rate increases to 60% for a similar household lacking family insurance. This wedge closes at wealth levels equal to six months of average income. As an extension, we introduce marital shocks in our benchmark model so that married and single households co-exist. Within this framework we study the robustness of our results on intra-household risk sharing, the implications of marital risk on households' savings, and compute consumption losses upon employment and/or marital shocks.

There is a vast literature —too large to be discussed here— assessing the effects of idiosyncratic income risk on consumption, labor supply and savings. Most of this literature adopts the bachelor household formulation to study individual responses to income shocks, and the extent of self-insurance. A recent example of this type of exercise is the paper by Low *et al.* (2008). These authors assume that individuals (they focus only on males) are subject to a rich array of idiosyncratic shocks, including productivity and employment shocks. These shocks are assumed to differ in their available insurance opportunities (employment shocks are partially insured by public UI while productivity shocks are not). The authors then use a bachelor household model to assess the effects of these shocks and the individual willingness to pay to avoid them. Since they consider endogenous mobility choices, their paper extends previous results in the literature by adding a new channel from shocks to individual responses to shocks.

Kotlikoff and Spivak (1981) is one of the first papers in economics to study the family as a provider of insurance to its members. In their model unexpected longevity is the only risk faced by individuals. They show that within-household efficient risk-sharing closes much of the utility gap between no annuities and complete annuities. For example, the utility gain of marriage at age 30 is about 50% of the utility gain of an annuities market.

A more recent exception to the use of the bachelor household formulation is the work of Attanasio *et al.* (2005), who present a partial equilibrium model with a two-person, *unitary* household to assess the response of female labor market participation (extensive margin) to idiosyncratic earnings risk within the family. While male participation is exogenous, female participation is endogenous and assumed to affect her human capital formation. The authors find that the higher the uncertainty, the higher female participation. They also find that the welfare cost of uncertainty is lower when households can adjust female labor market participation.

Heathcote *et al.* (2010) also use a two-person, unitary household model to study the welfare implications of the observed changes in the U.S. wage structure. In particular, they present an incomplete-markets, life-cycle model to quantify the effects of the rising college premium, the narrowing wage gender gap and the increasing wage volatility. Their model allows for an endogenous education choice and for a process matching females and males into households. Even though the welfare consequences of the above-mentioned changes in wages are highly heterogeneous across different types of households, they find that, on average, recent cohorts of households enjoy welfare gains, as the new structure of wages translates into higher educational attainment.

The remaining of the paper is organized as follows. Section 2 describes the economic environment, defines the steady-state equilibrium, and presents some properties of decision rules for collective households. It also presents a parameterization and calibration of this economy, and shows results on the role played by intra-household risk sharing. Section 3 contains our results on how family insurance shapes the effects of changes in public insurance. Section 4 presents some extensions and applications, and Section 5 concludes.

2 The Economic Environment

Consumers The economy is populated by a continuum of measure two of infinitely-lived consumers. Half of this population is referred to as females (f), and the other half as males (m). They supply time to work in the firms sector and face idiosyncratic risk in the form of employment shocks. Employment shocks, s , take on values in $S \equiv \{0, 1\}$ and follow a Markov chain with transition matrix Π^i , for $i = f, m$. Thus, $\pi_{s'|s}^i$ is the probability for an agent of gender i to receive employment shock s' tomorrow conditional on employment shock s today. These probabilities satisfy $\sum_{s'} \pi_{s'|s}^i = 1$, $\pi_{s'|s}^i > 0$, and $\pi_{1|1}^i \geq \pi_{1|0}^i$ for $i = f, m$. The long-run probabilities of the two employment shocks in S are denoted by q_0^i and q_1^i . There are no others shocks in the economy.

Markets are incomplete. Households can save in a non-state-contingent asset, a , that pays the risk-free interest rate r . There is a borrowing constraint represented by $a \geq \underline{a}$.

Lifetime preferences for an agent of gender i over stochastic consumption and leisure streams are

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U^i(c_t, l_t), \quad \text{for } i = f, m, \quad (2.1)$$

where c_t denotes consumption and l_t is leisure. We make the following assumptions on U^i : **A1)** Utility $U^i(c, l) : \mathbb{R}_+ \times [0, 1] \rightarrow \mathbb{R}$ is bounded, continuous and twice continuously differentiable in its interior. **A2)** U^i is separable in consumption and leisure. **A3)** U^i is strictly increasing and strictly concave in each of its arguments. Moreover, $\lim_{c \rightarrow 0} U_c^i(c, l) = +\infty$, and $\lim_{l \rightarrow 0} U_l^i(c, l) = +\infty$.

Firms The aggregate good is produced by competitive firms with neoclassical production function

$F(K, L)$, where K is the aggregate stock of capital and L is aggregate labor. Capital depreciates at rate $\delta > 0$. Aggregate labor is defined as $L \equiv \lambda L^m + (1 - \lambda)L^f$, where $0 < \lambda < 1$ is a parameter. Given r and gross wage rates \bar{w}^f and \bar{w}^m , the firm's first-order conditions are:

$$F_K(K, L) = r + \delta \quad (2.2)$$

$$\lambda F_L(K, L) = \bar{w}^m \quad (2.3)$$

$$(1 - \lambda)F_L(K, L) = \bar{w}^f. \quad (2.4)$$

Public Insurance A public unemployment insurance scheme (UI) pays out b^i , for $i = f, m$, to workers hit by the unemployment shock. UI is financed on a period-by-period basis with taxes on labor income.

2.1 The Bachelor versus the Collective Household Model

We consider two different risk-sharing arrangements, each of them defining in turn a different type of household. We start out by presenting the problem of the bachelor household. This is the definition of household that has dominated not only the literature on precautionary savings, but also most of the macroeconomic literature. A single breadwinner chooses sequences of consumption, leisure and asset holdings in order to maximize his/her own lifetime utility. In most studies adopting this framework, the income process is estimated using data on males. The second type of household we study is a dynamic version of the collective household model pioneered by Chiappori (1988). A household is formed by two individuals who make collective decisions on consumptions, labor supplies and savings. In order to understand the consequences of intra-household risk sharing we compare the allocations generated by these two household arrangements.

Bachelor Households A household formed by a single agent of gender i solves

$$v^i(s, a; w^i, r) = \max_{c, l, a'} \left\{ U^i(c, l) + \beta \sum_{s'} \pi_{s'|s}^i v^i(s', a'; w^i, r) \right\} \quad (2.5)$$

$$s.t. \quad c + a' = w^i(1 - l)s + (1 - s)b^i + (1 + r)a \quad (2.6)$$

$$c \geq 0, 0 \leq l \leq 1, \text{ and } a' \in [\underline{a}^i, \bar{a}], \quad (2.7)$$

where $\pi_{s'|s}^i$ are the elements of Π^i and $w^i = (1 - \tau^i)\bar{w}^i$ denotes after-tax wage rates. A version of this model, where there is a measure one of same-gender workers, is the workhorse model in the literature of uninsurable idiosyncratic risk, precautionary savings and labor supply.³ By construction, the bachelor household does not engage in informal insurance arrangements with other workers. The only sources of insurance available to this type of household are public insurance, own savings and own labor supply.

³See, e.g., Flodén and Lindé (2001), Marcet *et al.* (2007) and Pijoan-Mas (2006).

Collective Households We now consider two-person, collective households formed by an egotistical female and an egotistical male who share labor market risk and reach efficient intra-household allocations. Following the literature of collective households (see Chiappori and Donni 2010 for a recent survey), the utility of each individual in the household carries a weight, reflecting her/his relative power in the household. Individual weights are assumed to depend on variables such as the population sex ratio, relative earnings and public insurance. Under full commitment to future intra-household allocations, individual weights are set when the household is formed and remain unchanged thereafter. Thus, transitory shocks, which are small relative to lifetime income, are assumed to have no effect on individual weights. Only variables known or predicted at the time of household formation can affect those weights.⁴ We write the Pareto weight on females' utility as $\mu(x, \mathbf{z}) \in (0, 1)$, where function μ is assumed to be differentiable with respect to its first argument. Variable x is a measure of the relative earning ability of the two spouses, which we write as,

$$x \equiv \frac{q_1^f(1 - \tau^f)\bar{w}^f + q_0^f b^f}{q_1^m(1 - \tau^m)\bar{w}^m + q_0^m b^m}, \quad (2.8)$$

where q_i^j is the long-run probability of employment state i for an agent of gender j . I.e., there are four sources of female-male earnings differences that affect relative Pareto weights: 1) Different gross wages; 2) Different tax rates; 3) Different levels of UI; and, 4) Finally, females and males may be subject to different employment and unemployment spells. Vector \mathbf{z} includes variables such as the population sex ratio, the initial contribution to household wealth, etc., which we do not model explicitly in this paper. In our model the Pareto weight function, $\mu(x, \mathbf{z})$, is not obtained as the outcome of an explicit bargaining process between females and males.⁵ Instead, we will use estimates by Browning *et al.* (1994) to pin down the derivative of μ with respect to x .

Household-level state variables for the two-person, collective household are the vector of employment shocks, $\mathbf{s} = (s^f, s^m)$, and the level of asset holdings, a . The state space of a household is $X = S \times S \times [\underline{a}, \bar{a}]$. The transition matrix for \mathbf{s} is denoted by Π . In the case of uncorrelated employment shocks within the household matrix Π is simply obtained as $\Pi \equiv \Pi^m \otimes \Pi^f$. Otherwise, a typical element of Π is written as $\pi_{s^f s^m}^f [(1 - \varsigma)\pi_{s^m' | s^m}^m + \varsigma \mathcal{I}_{\{s^f = s^m'\}}]$, where \mathcal{I} is an indicator function and parameter $\varsigma \in [0, 1]$ pins down the extent of positive correlation in employment shocks within the household. The vector of after-tax wages for the household, (w^f, w^m) , is denoted by \mathbf{w} .

⁴Mazzocco (2007) uses CEX data to test, and reject, the hypothesis of intra-household commitment. Since our model abstracts from permanent shocks and assumes only transitory shocks to labor income, we will initially retain, for the sake of analytical tractability, the assumption of commitment. We discuss below the implications of this assumption for our results.

⁵In a recent paper, Heathcote *et al.* (2009) endogenize the Pareto weight as the solution to a symmetric Nash bargaining problem within the household.

The maximization problem of a collective household with female's Pareto weight $\mu(x, \mathbf{z})$ is

$$V(\mathbf{s}, a; x, \mathbf{z}, r) = \max_{c^f, c^m, l^f, l^m, a'} \left\{ \mu(x, \mathbf{z})U^f(c^f, l^f) + [1 - \mu(x, \mathbf{z})]U^m(c^m, l^m) + \beta \sum_{\mathbf{s}'} \pi_{\mathbf{s}'|\mathbf{s}} V(\mathbf{s}', a'; x, \mathbf{z}, r) \right\} \quad (2.9)$$

s.t.

$$c^f + c^m + a' = \sum_{i=f,m} w^i(1 - l^i)s^i + \sum_{i=f,m} (1 - s^i)b^i + (1 + r)a \quad (2.10)$$

$$c^f, c^m \geq 0, \quad 0 \leq l^f, l^m \leq 1, \quad \text{and} \quad a' \in [\underline{a}, \bar{a}], \quad (2.11)$$

where $\pi_{\mathbf{s}'|\mathbf{s}}$ are the elements of Π . In our model, \mathbf{z} is the only source of variation in Pareto weights across households. We represent the distribution of these weights in the population of households by $G(\mu)$, with support $M \equiv (0, 1)$.

The utility function of the collective household depends, *via* the Pareto weight, on wages and policy variables, yielding household demands that fail to meet the Slutsky conditions. This failure is the defining feature of the collective model.⁶ Household decisions depend not only on total household income, but also on who receives the income. The dependency of the household's utility function on prices and policy must be acknowledged when setting the Frisch elasticities of labor supply for females and males. In particular, these elasticities are functions of the derivative of the Pareto weight with respect to wages (see the Appendix for a derivation of Frisch elasticities). Likewise, the household's attitude towards risk depends both on individual preferences and on the Pareto weight [a derivation of the household's risk attitude under CRRA utility functions is presented in the Appendix. For a two-period, collective model with uncertainty see Mazzocco (2004)].

The assumption of egoistical preferences is not crucial. Actually, Browning *et al.* (2006) show that under caring preferences of the form $U^i(c^i, l^i) + \nu^i U^j(c^j, l^j)$, where $0 < \nu^i \leq 1$, is agent i 's caring parameter, the utility function of the household can be written down as for the case of egotistical preferences, after a re-definition of Pareto weights.

We now present the first-order conditions to the maximization problem (2.9)-(2.11). Risk-sharing within the household implies that the ratio of marginal utilities of consumption equals relative Pareto weights, and is thus independent of the realized vector of employment shocks. That is,

$$\mu U_c^f = (1 - \mu)U_c^m. \quad (2.12)$$

This equation defines individual risk-sharing rules, which, for a given level of household consumption, specify how much is consumed by each of its members. It is straightforward to show that the first-order derivatives of the risk-sharing rules are positive and given by the product of the

⁶See Browning *et al.* (2006) for a formal definition of collective *versus* unitary models of the household.

household's coefficient of absolute risk aversion and the individual's coefficient of absolute risk tolerance.⁷ Therefore, the member of the household showing higher risk tolerance will be the one absorbing most of the variation in total household consumption. (In the Appendix we present the derivatives of the risk-sharing rules for the case of CRRA utility functions.)

First-order conditions to female and male labor supply are, respectively,

$$\frac{U_l^f}{U_c^f} \geq w^f s^f \quad \text{with inequality if } l^f = 1 \quad (2.13)$$

$$\frac{U_l^m}{U_c^m} \geq w^m s^m \quad \text{with inequality if } l^m = 1. \quad (2.14)$$

Moreover, if the labor supply decision is interior for both household members then

$$\frac{U_l^f}{w^f s^f} = \frac{1 - \mu}{\mu} \frac{U_l^m}{w^m s^m}. \quad (2.15)$$

The first-order condition to savings is,

$$U_c^f = \beta(1+r) \sum_{s'} \pi_{s'|s} U_c^{f'} \quad \text{if } a' > \bar{a} \quad (2.16)$$

$$U_c^f \geq \beta(1+r) \sum_{s'} \pi_{s'|s} U_c^{f'} \quad \text{if } a' = \underline{a}. \quad (2.17)$$

Proposition 1 below presents some properties of the value function and decision rules for a household with Pareto weight $\mu \in M$. (Our proof in the Appendix is for the case $b^f = b^m = 0$, but it is straightforward to show that the results hold for the case of positive UI, provided $w^i l^i \geq b^i$ for $i = f, m$, i.e., earnings are higher than UI.)

Proposition 1. Assume **A1** – **A3**, $w > 0$, $(1+r) > 0$, $\beta(1+r) \leq 1$. Then:

- (a) $V(\mathbf{s}, a, \mu)$ is strictly increasing and strictly concave in a . Decision rules $c^f(\mathbf{s}, a; \mu)$, $c^m(\mathbf{s}, a; \mu)$, $l^f(\mathbf{s}, a; \mu)$, $l^m(\mathbf{s}, a; \mu)$ and $a'(\mathbf{s}, a; \mu)$ are continuous in a and strictly positive.
- (b) Decision rules for consumption, $c^f(\mathbf{s}, a; \mu)$ and $c^m(\mathbf{s}, a; \mu)$, are strictly increasing in a . Decision rules for savings, $a'(\mathbf{s}, a; \mu)$, and leisure, $l^f(s^f = 1, s^m, a; \mu)$, $l^m(s^m = 1, s^f, a; \mu)$, are increasing in a .
- (c) Decision rules for consumption are increasing in the own employment shock: $c^j(s^j = 1, s^i, a; \mu) \geq c^j(s^j = 0, s^i, a; \mu)$.
- (d) Decision rules for leisure are increasing in the spouse's employment shock: $l^j(s^j = 1, s^i = 1, a) \geq l^j(s^j = 1, s^i = 0, a; \mu)$.
- (e) If $\beta(1+r) \leq 1$, then for all $a \in [\underline{a}, \bar{a}]$, $a'(s^f = 0, s^m = 0, a; \mu) \leq a$ (with strict inequality if $\underline{a} < a < \bar{a}$ and $\beta(1+r) < 1$).

⁷Risk tolerance is defined as the reciprocal of risk aversion.

We now present some results on the asymptotic properties of the consumption program, savings and labor supply of a household with Pareto weight μ , for different values of wages, (w^f, w^m) , and of the interest rate, r . We extend results by Marcet *et al.* (2007) for the bachelor household to our two-person, collective household model. We also extend the results to non-homogeneous utility functions. With this aim, let us denote by $\tilde{a}(\mu)$ the minimum level of asset holdings for which both spouses within a household with Pareto weight μ will stop supplying labor. The value of $\tilde{a}(\mu)$ and the proof of Proposition 2 below are presented in the Appendix.

Proposition 2: Assume **A1** – **A3**, $\bar{a} > \tilde{a}(\mu)$, $w > 0$ and $(1 + r) > 0$. Then:

- (a) If $\beta(1 + r) \leq 1$, for any $a \leq \tilde{a}(\mu)$, $a'(\mathbf{s}, a; \mu) \leq \tilde{a}(\mu)$.
- (b) If $\beta(1 + r) = 1$, for any $a \geq \tilde{a}(\mu)$ and any \mathbf{s} we have $a'(\mathbf{s}, a; \mu) = a$, $l^f(\mathbf{s}, a; \mu) = 1$, $l^m(\mathbf{s}, a; \mu) = 1$ and $c^f(\mathbf{s}, a; \mu) + c^m(\mathbf{s}, a; \mu) = ar$ such that $\mu U_c^f = (1 - \mu) U_c^m$.
- (c) If $\beta(1 + r) = 1$ and $a \leq \tilde{a}(\mu)$, then $a_t \xrightarrow{a.s.} \tilde{a}(\mu)$, $c_t^i \xrightarrow{a.s.} \tilde{c}^i(\mu)$, $l_t^i \xrightarrow{a.s.} 1$, $i = f, m$.

In the case $\beta(1 + r) < 1$, the household can reach any value of asset holdings from any initial condition in finite time, and a stationary distribution arises in the long run. Moreover, in the case $\beta(1 + r) = 1$ and $a \leq \tilde{a}(\mu)$, capital accumulation in the long run is bounded and it converges asymptotically to $\tilde{a}(\mu)$. This is in contrast to the case of inelastic labor supply where savings asymptotically grow to infinity if $\beta(1 + r) = 1$. As it should be apparent from these results, the endogenous labor-leisure decision changes the asymptotic behavior of consumption and wealth by removing income uncertainty. That is, when household wealth is high enough, labor supply equals zero and, thus, employment shocks no longer affect household income. Hence, under non-stochastic income, unbounded asset accumulation is no longer optimal when $\beta(1 + r) = 1$.

If we set $\bar{a} > \max_{\mu \in M} \tilde{a}(\mu)$ and choose initial capital holdings for all households with relative Pareto weight μ such that $a_0(\mu) \leq \tilde{a}(\mu)$, then the upper bound on asset holdings, which was imposed to guarantee existence and uniqueness of the value function, is never binding.

2.2 The Steady-State Equilibrium

A stationary equilibrium in the collective household economy is defined as follows. Let $\psi(B; \mu)$ be a probability measure, defined on the Borel sigma algebra \mathcal{B} , describing the mass of households with fixed Pareto weight μ at each $B \in \mathcal{B}$. Denote by $P(\mathbf{s}, a, B; \mu)$ the probability that a household with Pareto weight μ at state (\mathbf{s}, a) will transit to a state that lies in $B \in \mathcal{B}$ in the next period. The transition function P can be constructed as,

$$P(\mathbf{s}, a, B; \mu) = \sum_{\mathbf{s}' \in B_{\mathbf{s}}} \Pi_{\mathbf{s}'|\mathbf{s}} \mathcal{I}_{\{a'(\mathbf{s}, a; \mu) \in B_a\}},$$

where \mathcal{I} is an indicator function and $B_{\mathcal{S}}$ and B_a are the projections of B on $S \times S$ and $[\underline{a}, \bar{a}]$, respectively. We are now ready to define the stationary equilibrium.

Definition: A stationary recursive competitive equilibrium with incomplete markets in the economy with collective households is a list of functions $\{V, c^f, c^m, l^f, l^m, a', K, L^f, L^m\}$, a measure of households ψ and a set of prices $\{r, \bar{w}^f, \bar{w}^m\}$, taxes $\{\tau^f, \tau^m\}$ and UI $\{b^f, b^m\}$ such that:

- 1) For given prices, taxes and UI, V is the solution to (2.9) – (2.11), and $c^f(\mathbf{s}, a; \mu)$, $c^m(\mathbf{s}, a; \mu)$, $l^f(\mathbf{s}, a; \mu)$, $l^m(\mathbf{s}, a; \mu)$ and $a'(\mathbf{s}, a; \mu)$ are the associated optimal policy functions.
- 2) For given prices, K , L^f and L^m satisfy the firm's first-order conditions (2.2) – (2.4).
- 3) Aggregate factor inputs are generated by the policy functions of the agents:

$$K = \int_M \int_X a'(\mathbf{s}, a; \mu) d\psi dG, \quad (2.18)$$

$$L^f = \int_M \int_X s^f [1 - l^f(\mathbf{s}, a; \mu)] d\psi dG, \quad (2.19)$$

$$L^m = \int_M \int_X s^m [1 - l^m(\mathbf{s}, a; \mu)] d\psi dG. \quad (2.20)$$

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

$$\psi(B; \mu) = \int_X P(\mathbf{s}, a, B; \mu) d\psi \quad \text{for all } B \in \mathcal{B}. \quad (2.21)$$

- 5) The government budget is balanced: $q_0^f b^f + q_0^m b^m = \tau^f \bar{w}^f L^f + \tau^m \bar{w}^m L^m$.

Under assumptions **A1** – **A3** the interest rate in the stationary equilibrium must be such that $\beta(1+r) < 1$.

2.3 Parameterization and Calibration

We assume identical preferences for females and males and parameterize their common instantaneous utility function as follows,⁸

$$U(c, l) = \frac{c^{1-\sigma} - 1}{1-\sigma} + \varphi \frac{l^{1-\gamma} - 1}{1-\gamma}. \quad (2.22)$$

Since our exercise is on intra-household insurance, and since insurance is tightly connected to wealth, it is important that we match the proportion of low-wealth households in the U.S. Thus, and following ideas in Krusell and Smith (1998), we assume that households are heterogeneous in

⁸In a previous version of this paper we assumed different preferences for females and males. Our main results are however unaffected by this assumption.

terms of the discount factor, β . Namely, there is a fraction, say χ , of impatient households with discount factor β_L . Patient households use discount factor $\beta_H > \beta_L$.

The production technology is the standard Cobb-Douglas function, $F(K, L) = K^\alpha L^{1-\alpha}$, where labor is $L \equiv \lambda L^m + (1 - \lambda)L^f$. Parameter α is the capital's share of income and λ pins down relative gross wages, since $\bar{w}^f/\bar{w}^m = (1 - \lambda)/\lambda$.

We make the simplifying assumption of a degenerate distribution over vector \mathbf{z} , implying a common relative Pareto weight across households. Both μ and its derivative with respect to x , $\mu_x(x, \mathbf{z})$, are pinned down below.

Parameter Values We set values to five preference parameters: $\beta_L, \beta_H, \sigma, \varphi, \gamma$; three technology parameters: α, λ and δ ; and the four parameters in the two transition matrices Π^f and Π^m . The public insurance program contains four parameters: τ^f, τ^m, b^f and b^m . Finally, we have the borrowing limit, \underline{a} , and the two values μ and μ_x .

One period is set to one quarter. We choose $\underline{a} = 0$ so that households face a no-borrowing constraint. We set $\sigma = 2$, which is a standard value in the macro literature. The depreciation rate of capital is set at $\delta = 0.025$, and the capital's share of income, α , equal to 0.36. We impose equal tax rates for females and males, $\tau^f = \tau^m$, which allows us to pin down the value of λ from *a priori* information on the gender wage gap. We set λ equal to 0.575, corresponding to a wage gap of 0.74, the value reported by Heathcote *et al.* (2010) for the 2004 U.S. economy.

Transition probabilities for idiosyncratic employment shocks are assumed to be identical for females and males, as the average difference between female and male unemployment rates over the period 1980-2009 is practically zero. We set $\pi_{1|1}^i = 0.09$ and $\pi_{1|0}^i = 0.06$ for $i = f, m$, which match an average employment rate of 93%, after normalizing with the participation rate.⁹ In our benchmark economy within-household employment shocks are uncorrelated, an assumption that is supported by SIPP data (Survey of Income Program Participation). Indeed, within-household unemployment correlation for households where husband and wife report different occupation is 0.05. For households reporting same occupation this correlation is 0.23. However, only 3.2% of the households report same occupation for husband and wife. (For a detailed explanation on the calculation of these correlations, see Shore and Sinai 2010.) We will study the sensitivity of our results to positive correlation in employment shocks within the household.

Remaining parameters are set to match the following targets: 1) Married working females' and males' average hours of work represent 28% and 40% of their discretionary time, respectively.¹⁰

⁹These transition probabilities are taken from Imrohoroglu (1989), Krusell and Smith (1998) and Marcet *et al.* (2007).

¹⁰Mazzocco *et al.* (2008) use PSID data and obtain mean annual hours worked by working married females and males equal to 1660 and 2312, respectively. We assume 16 hours of daily discretionary time.

2) Estimates for males' Frisch elasticity in the presence of borrowing constraints range from 0.2 to 0.6 (see Domeij and Flodén 2006). For females, estimates are typically higher (see Blundell and MaCurdy 1999). We will target values of 0.50 and 0.85 for males and females, respectively. 3) The capital-to-output ratio is 10. 4) The average net UI replacement rate in the U.S. is roughly 30 percent (see OECD 2010). We set b^f and b^m to match this target as fractions of the average wage income, both for females and males. Labor income tax rates are set to balance the budget of the public insurance program. 5) The derivative of the Pareto weight function with respect to the expected income differential, μ_x , is set to match sharing rule estimates presented in Browning *et al.* (1994).¹¹ 6) The fraction of liquidity-constrained households in the U.S. (using Zeldes' 1989 definition of liquidity constrained as holding non-housing wealth below two months of average income) averages 17% for the period 1983-2004, as reported by Gorbachev and Dogra (2010). 7) Finally, we target the ratio of wealth held by the bottom quintile in the wealth distribution over total wealth to be less than one percent. Table 1 presents our benchmark economy.

[Insert Table 1 here]

Aggregate values at the steady-state equilibrium are: $Y = 1.2722$, $K = 12.6814$, $L = 0.3490$, and $r = 0.0111$. This steady state matches relatively well the bottom tail of the wealth distribution, but it does less well at matching the upper tail. We could improve on the upper tail by introducing three different discount factors, as in Krusell and Smith (1998), instead of only two. However, since our focus is on the effects of intra-household risk sharing at the level of the household, and since these effects are small for households with high levels of wealth, it is not crucial that we match the upper part of the wealth distribution. Hence, our model contains the structure needed to conducting the study on intra-household risk sharing.

2.3.1 Policy Functions: A First Look at the Effects of Intra-household Risk Sharing

We start by presenting policy functions in our collective economy and then assess the effects of intra-household risk sharing on households' savings and hours worked. The left panel of Figure 1 presents savings policy functions (for convenience we plot the net change in asset holdings $a' - a$), for households with low discount factor (top chart on left panel) and with high discount factor (bottom chart on left panel). Among impatient households, positive net savings are observed

¹¹These authors use data on couples with no children to estimate the parameters of the sharing rule: they find that increasing the wife's contribution to household income from 25% to 75% (holding total expenditure constant) raises her share in total expenditure by about 2.3%. We use this estimate to obtain μ_x as follows. Starting from the parameters that conform our benchmark economy, we increase the value of x — e.g. by raising w^f/w^m — and then compute the new Pareto weight, say $\tilde{\mu}$, such that the implied increase in the wife's relative contribution to household income yields an increase in the wife's share of total expenditure, $c^f/(c^f + c^m)$, that matches the one implied by the sharing rule as estimated in Browning *et al.* (1994).

only when the two spouses are employed and hold low levels of assets. Households with one or two unemployed spouses choose non-positive net savings. Among patient households, those with two employed spouses choose positive net savings at all values in the support of the equilibrium distribution of assets. Patient households where at least one of the spouses is unemployed choose negative net savings at a large set of low asset holdings. Negative net savings are larger in households where the male is unemployed.

Policy functions for hours are shown in the right panel of Figure 1 (top chart for impatient households and bottom chart for patient ones). For both groups, hours decrease with household wealth. As asset holdings approach the borrowing limit, policy functions for hours bend upwards, capturing the fact that asset-poor households use labor supply to smooth consumption more intensively. Hours increase if the spouse is unemployed, both for females and males, and the increase is especially marked for females in asset-poor households. For example, a female in a household with no assets will supply almost half of her available time to work if the spouse is unemployed, as opposed to 0.36 when the spouse is employed, which represents a decline of 28%.

The savings and hours effects of intra-household risk sharing are shown in Figures 2 and 3. Figure 2 plots excess savings of two bachelors (each with wealth $a/2$) over a two-person collective household (with wealth a). The two top charts show excess savings for impatient and patient households across employment shocks. Clearly, although risk sharing affects the savings decisions of all households across the wealth distribution, its effects are strongest among wealth-poor households. The bottom charts of Figure 2 show average excess savings.

The top chart of Figure 3 plots excess hours worked by two bachelors (each with wealth $a/2$) over hours worked by a two-person collective household (with wealth a). (Since there is not much difference in excess hours between patient and impatient households we plot the average of the two.) For all households where only the male is employed, intra-household risk sharing increases household hours. For households where the female is employed, with the exception of low-wealth households with the male unemployed, intra-household risk sharing decreases household hours. The bottom chart of the Figure shows the average of excess hours across households along the employment distribution. As it is apparent, the effects of intra-household risk sharing on hours are strongest among wealth-poor households.

3 Intra-household Risk Sharing and Public Insurance

The use of savings and labor supply as insurance mechanisms depends on the generosity of public insurance. The extent to which the ability to share risks within the household shapes the crowding out effects of public insurance is explored in this section.

3.1 Household Financial Assets and the Generosity of UI

An implication of our model, as of any model with uninsurable income risk, is that household asset holdings increase with income uncertainty. Since the level of UI is directly correlated with household income risk, Engen and Gruber (2001) exploit the variation in generosity of unemployment insurance schedules across U.S. states to test this implication and to estimate the extent of the precautionary savings motive. These authors use SIPP data—which follows a cross section of households over a period of 2.5 years—in combination with data on UI available to these households under their state/date insurance system. They regress household financial assets (normalized by average household income) on the generosity of UI, controlling for a vector of demographic and economic characteristics of the household. They find an elasticity of the average household’s financial assets-to-income ratio with respect to UI equal to -0.28 . That is, reducing the replacement rate of UI by 50% would increase the household’s assets-to-income ratio by 14%.

We use our model economy to compute the elasticity of the average assets-to-income ratio with respect to UI. The purpose of this exercise is twofold. On the one hand, we use it as a test for our model to match this estimated measure of the precautionary savings motive. On the other hand, we also compute this elasticity using the bachelor household model and assess by how much it overestimates the precautionary motive. In this latter model there is no intra-household risk sharing and, therefore, variation in UI amounts to larger changes in household income risk and, consequently, to larger effects on savings.

In order to mimic the empirical exercise conducted by Engen and Gruber (2001) we proceed as follows. Since these authors rely on the exogenous variation in UI for workers living in different states in the U.S., we interpret the level of UI in our benchmark economy as the average value across all states. Then, we vary the level and compute asset-to-income ratios holding equilibrium prices unchanged, a strategy in accordance with the existence of a unique financial and labor market across states. However, Pareto weights and the distribution of households over asset holdings are let to change with UI. Thus, our exercise compares the differential asset-to-income ratio of households across states that provide these households with differing levels of UI, which is exactly what Engen and Gruber (2001) do in their empirical work. The results of this exercise are presented in Table 2. Our collective household model yields an elasticity of the assets-to-income ratio with respect to UI equal to -0.38 , accounting thus fairly well for the empirical elasticity estimated by Engen and Gruber (2001). On the contrary, the bachelor household model overpredicts this elasticity by a factor of more than two. This shows that intra-household risk sharing plays a crucial role in the determination of the elasticity of the assets-to-income ratio and, therefore, that this informal source of insurance is key when assessing the crowding out effects of public insurance.

The relative success of our collective model at matching this empirical elasticity lends support to the view that the two-person household embeds the most relevant informal insurance arrangements available to individuals. Indeed, some authors have emphasized that the extended family, friends and other social networks play only a negligible insurance role (see, e.g., Blundell *et al.* 2008).

[Insert Table 2 here]

The results in this section are robust to the introduction of correlated employment shocks within the household. The elasticity of the assets-to-income ratio with respect to UI in the collective economy increases only to -0.39 when the correlation in employment shocks is set to 0.3.

3.2 Spousal Labor Supply as Insurance

As shown above, spousal labor supply is a potentially important source of household self-insurance. The change in a household member’s labor supply induced by an unemployment spell of another household member —the added worker effect— has been largely studied in the empirical literature. Most of this literature has focused on the labor supply response of married women to their husband’s unemployment spells.¹²

Early literature on the added worker effect (see Cullen and Gruber 2000 for a short review) has singled out liquidity constraints as one of the main reasons married women increase hours worked during their husband’s unemployment spells. Empirical estimates have, however, produced mixed results, failing to find strong support for this effect.¹³ Cullen and Gruber (2000), using SIPP data for married couples aged between 25 and 54 years old, report means for wives’ monthly hours worked during husbands’ spells of employment and unemployment, respectively. Conditional on working women, these authors find that the average amount of work per month by wives of unemployed husbands is 149 hours, as opposed to 132.4 by wives of employed husbands. When non-working wives are included, average hours are 98.2 and 97.9, respectively.

In this section we use our collective households model to study the response of female labor supply to males’ unemployment spells in two groups of households. In order to highlight the role of liquidity constraints for wives’ labor supply responses, we follow Zeldes (1989) in defining a household as liquidity constrained if its non-housing wealth is less than two months of average income. Table 3 below reports the added worker effect in our model economy. For the group

¹²The main argument in favor of restricting the attention to labor supply of women is that they are the secondary wage earners in most households (according to Cullen and Gruber 2000, in 87% of married couples in the U.S. the husband earns more and in 73% the husband works more hours).

¹³Stephens (2002) estimates the added worker effect taking into account not only the current period of the husband’s job loss but also the periods before and after a job loss. This author finds small pre-displacement effects but large, persistent post-displacement effects.

of liquidity-constrained households, average hours worked by wives of unemployed husbands are 8.6% higher than those worked by wives of employed husbands, an increase comparable to that found by Cullen and Gruber (2000) in their sample of working women. When all households are taken into account the increase in hours is only 2.7%. That is, spousal labor supply is an important insurance mechanism for wealth-poor households but not for the wealth rich.

[Insert Table 3 here]

How effective is wives' labor supply as insurance against income fluctuations due to husbands' unemployment? To answer this question we compute, for each level of assets, a , the fraction of lost family income that is made up by the wife's response to the husband's unemployment spell,

$$\frac{[h^f(0, 1, a) - h^f(1, 1, a)]w^f}{h^m(1, 1, a)w^m - b^m},$$

where $h^f(0, 1, a)$ denotes hours worked by a female with an unemployed husband and $h^f(1, 1, a)$ is female hours if the husband is employed. The denominator represents lost income due to husband's unemployment. The numerator is the increase in income due to the wife's response in hours. For the group of liquidity-constrained households, the wives' response makes up about 9.6% of lost family income, while this number is only 2.5% when we consider all households. Wealth-rich households use savings to smooth consumption upon husband's unemployment. Liquidity-constrained households must rely, however, on spousal labor supply.

Spousal Labor Supply and the Generosity of UI According to some authors, the moderate to nil added worker effect found in the data may be partially explained by the presence of public insurance. That is, unemployment payments during the husband's unemployment spell crowd out wife's labor supply. Cullen and Gruber (2000) estimate this effect and find that a 50% reduction in potential UI of the husband (75 USD per week) would imply an increase in monthly hours worked by the wife (conditional on working) of 13.42 hours, which amounts to an increase of about 9%. They also find a differentially larger response of wives' labor supply among liquidity-constrained households.

We use our model economy to compute the crowding out of UI on wives' labor supply. Table 4 below presents the results of this exercise. A 50% reduction in UI received by the husband increases wife's hours by 8.71% in the group of liquidity-constrained households. This increase is only 1.53% when all households are considered. The relatively higher sensitivity of spousal labor supply to UI among liquidity-constrained households found in our model is in line with the finding of Cullen and Gruber (2000).

[Insert Table 4 here]

Even though a direct comparison of our results with those in Cullen and Gruber (2000) is not

straightforward, our findings in Table 6 are remarkably close to empirical estimates. It should be noted, however, that their estimated increase in hours of work by the wife after a 50% reduction in UI received by the unemployed husband is not statistically significant, thus hindering the assessment of our model’s predictions.

3.3 Consumption Loss Upon Unemployment

Under imperfect capital markets, the loss of the job implies a reduction in the level of individual consumption. The degree of transmission of unemployment shocks to consumption depends on factors such as the generosity of UI, on the level of accumulated wealth and on whether risks are shared within the household.

In this section we use our benchmark economy to assess the contribution of intra-household risk sharing to individual consumption insurance, as measured by the degree of transmission of unemployment shocks to consumption. We do so by comparing individual consumption losses upon unemployment in the collective household model to those in the bachelor model. We compute the percentage change in consumption upon unemployment, $\Delta c/c$, at all asset levels in the support of the corresponding equilibrium distribution. In the collective economy, individual consumption losses for females and males, both with an employed spouse and with an unemployed spouse, are computed as,

$$\frac{c^j(s^j = 0, s^i, a) - c^j(s^j = 1, s^i, a)}{c^j(s^j = 1, s^i, a)},$$

for $j = f, m, i = f, m$ and $i \neq j$, both for $s^i = 1$ and $s^i = 0$. For the bachelor economy, individual consumption losses upon unemployment are computed as, $(c^j(0, a) - c^j(1, a))/c^j(1, a)$ for $j = f, m$.

In panel (a) of Table 5 we report average individual consumption losses, both for the group of liquidity-constrained individuals and for all individuals. We use the respective equilibrium asset and employment distributions to average out individual consumption losses. The results show that intra-household risk sharing provides important consumption smoothing opportunities, especially for liquidity-constrained individuals. Thus, the average consumption loss for a liquidity-constrained female in the bachelor economy is -23.23% , against -4.5% in the collective economy. For a liquidity-constrained male, intra-household risk sharing reduces his consumption loss from -30.09% to -8.74% . These numbers imply that the family is an important provider of consumption insurance for a significant fraction of individuals. These results are robust to correlation in employment shocks within the household. For instance, a correlation of 0.3 implies consumption losses of -4.85% and -9.05% for liquidity-constrained females and males, respectively.

[Insert Table 5 here]

We also study household insurance by computing the fraction of household income loss (due

to an unemployment shock) that translates into household consumption loss. To do this we compute income and consumption losses upon an unemployment shock for each household across the asset and employment distributions. Then, we average out the percentage of income loss that is transmitted to consumption loss across all households. Panel (b) of Table 5 presents our results for the two household arrangements.

The fraction of income loss that transits to consumption loss in the group of liquidity-constrained households is non-negligible, even in the collective household economy. For an average household in this group, 17.31% of the household income lost due to an unemployment shock is absorbed by consumption. This result is consistent with the empirical finding of Blundell *et al.* (2008) about the degree of insurability of transitory income shocks. They find that the impact of these shocks on consumption is small when estimated from all households in their sample, but it is found to be larger, about 0.2, in the subsample of wealth-poor households (these authors define a household as wealth poor if its wealth is in the bottom 20 percent of the distribution of initial wealth).

The fraction of income loss that transmits to consumption loss under a 0.3 correlation of employment shocks within the household increases from 17.31% to 18.39% for liquidity-constrained households.

Consumption Loss and the Generosity of UI We now turn to the sensitivity of household consumption losses upon unemployment with respect to the generosity of UI. Browning and Crossley (2001), using a Canadian panel, estimate this sensitivity exploiting legislative changes introduced in 1993 and 1994 that reduced the replacement rate by about five percentage points. In total, 19,000 individuals who had experienced a job separation either before or after the policy reform were interviewed several times after the job loss. Browning and Crossley (2001) obtain two main results. First, the level of UI has small average effects on household consumption loss upon unemployment. In particular, a 10 percentage-point reduction in UI leads to an average fall in consumption of 0.8%.¹⁴ Second, the consumption effects of UI are not homogeneous across households, being substantially larger within the group of liquidity-constrained households at the time of job separation. (These authors also follow Zeldes' (1989) definition of liquidity-constrained households.) These results show the importance of UI as a consumption smoothing instrument for wealth-poor households.

Table 6 below presents the elasticities of consumption loss with respect to UI in our model economies with and without intra-household risk sharing. Our quantitative exercise explicitly acknowledges the panel dimension of the empirical exercise conducted by Browning and Crossley (2001). We compute the relative change in consumption from the period prior to the unemploy-

¹⁴Gruber (1997) uses U.S. data on food consumption from the Panel Study of Income Dynamics (PSID) and finds a larger mean effect of UI on consumption losses upon unemployment. This author estimates that a 10 percentage-point increase in UI reduces the fall in consumption by 2.65%.

ment shock to the period in which the job separation is realized. Saving decisions in the pre-unemployment period are used for the computation of consumption upon unemployment. That is, consumption of an individual of gender j in the period before unemployment is $c^j(s^j = 1, s^i, a)$ and consumption at the time of the job loss is $c^j(s^j = 0, s^i, a')$, where $a' = a'(s^j = 1, s^i, a)$. We then weigh consumption levels in both periods using the stationary distribution of employment shocks for the spouse (in the collective economy). Our collective economy accounts well for the elasticity of consumption loss with respect to UI in the group of liquidity-constrained households. It however underestimates this elasticity for the whole sample of households.

[Insert Table 6 here]

It is important to note that the empirical exercise in Browning and Crossley (2001) uses Canadian data, while our baseline parameter values have been chosen to match U.S. stylized facts. Then, rather than trying to account for this elasticity, our exercise in this section aims at shedding further light on the role of intra-household risk sharing. Thus, the elasticity predicted by the bachelor economy, -0.2262 , is more than two times the elasticity under collective households.

4 Applications and Extensions

The Value of Intra-household Risk Sharing In order to assess the value of intra-household insurance, we first remove from our benchmark economy all intra-household transfers which are not related to risk sharing. This is accomplished by setting the gender wage gap to zero and the relative Pareto weight to one, so that any transfer within the household is driven by risk pooling. We then compute the increase in bachelor consumption that would leave an individual in a collective household indifferent between remaining in the household and splitting up with half of the household's assets to remain bachelor thereafter. I.e., we compute the value of ζ that solves,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U^i(c_t^{col}, l_t^{col}) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U^i((1 + \zeta)c_t^{bach}, l_t^{bach}), \quad (4.1)$$

where a superscript *col* refers to allocations in the collective economy from initial conditions (s^i, s^j) and a_0 , and a superscript *bach* refers to allocations attained by the deviating individual from initial conditions s^i and $a_0/2$. The results vary widely depending upon own and spouse's employment status, discount factor and asset holdings. For instance, the value of intra-household insurance for an impatient, unemployed individual with an employed spouse and no assets represents 5.16% of per period consumption of a similar bachelor individual. This number goes down to 0.75% if the individual is employed. The value of intra-household insurance for the unemployed decreases if the spouse is also unemployed. Also, the value of family insurance is lower to patient than to impatient

individuals. When aggregating over employment shocks and discount factors, the welfare gain within the group of liquidity-constrained households averages 0.2387%. To conclude, our model suggests that the value of intra-household insurance is significant to wealth-poor, unemployed individuals and that it decreases with wealth, especially among patient households. Finally, note that the value of family insurance to unemployed individuals with no assets is higher than the average value of removing aggregate business cycle fluctuations, as estimated in the literature.

Optimal Unemployment Insurance We carry out two exercises to ascertain how access to insurance from the family shapes the welfare effects of UI, and how it affects the optimal provision of public insurance. As noted by Krusell *et al.* (2010), measuring the welfare effects of UI by comparing ex ante, average utilities may be misleading, and they advocate instead for looking at welfare effects for different groups of agents. We follow this line and compute optimal UI, and the welfare effects of changes in UI, for the average household with wealth level a .

Our first results show that family insurance is a key determinant of optimal public insurance, especially for wealth-poor households. The optimal replacement rate for an average collective household with no assets is about 15%, while this rate increases to 60% for the average bachelor household with no assets, i.e., a difference of 45 percentage points. If we look now at households with levels of wealth equal to two months of average income, the differential in optimal replacement rates for the collective and the bachelor household goes down to 30 percentage points. As was to be expected, the higher the wealth level, the lower the differential in optimal replacement rates implied by family insurance. Actually, at wealth levels equal to six months of average income, this differential becomes zero.

Second, we compute the welfare effects (in consumption equivalent units) of reducing the replacement rate by 10 percentage points from the value in our baseline economy. For the average collective households with no assets this decrease in UI amounts to an increase in welfare of 0.015%, while the average bachelor households experiences a decrease in welfare of -0.164% . Among the group of liquidity-constrained households, the reduction in UI increases welfare of collective households by 0.034%, and decreases welfare of bachelor households by -0.035% .

Marital Shocks As shown by Mazzocco (2007), the assumption of commitment might not be supported by empirical evidence. In general, lack of commitment leads to non-stationary Pareto weights and, in certain cases, to marriage dissolution. In order to understand how the possibility of marriage dissolution alters the saving behavior of our two-person, collective households, we now introduce exogenous marital shocks into our benchmark economy. It must be noted that this extension does not embody all the implications of lack of commitment, but it does include the most relevant one to our purposes, i.e., the end of the marriage and hence of family insurance. Marital risk is bound to affect household savings, consumption and labor supplies, and thus the

role played by intra-household insurance.

Two-person households face now an exogenous probability of dissolution, say $\phi_d > 0$. When the separation shock hits, household members become single, each with half of the household assets. We assume no divorce costs. Single individuals face a probability of being matched with another single of opposite gender, say $\phi_m > 0$, and form a two-person household. Matching is segmented in the sense that prospective spouses for a single agent with wealth a are assumed to have wealth in a neighborhood centered around a . Upon matching, a two-person household is immediately formed and all assets are shared.¹⁵ We choose ϕ_d so that half of the population lives as married couples (according to the 2001 U.S. census, 54.8% and 44% of men and women over 25 years of age were married, respectively) and that the average duration of marriage is 7 years. Note that in our dynastic model if 50 percent of the population lives as married households we must have that $\phi_d = \phi_m$.

We now assess the effects of risk of divorce on savings, on the extent of the precautionary motive and on consumption losses. The savings effects of marital risk have received some attention in the literature. An important research question in this literature is whether an increase in the risk of marriage dissolution leads to higher household savings.¹⁶ The exercise we conduct in this section consists in computing the change in savings that would follow from the removal of marital risk, both for liquidity-constrained households and for all households in the economy. That is, we start with $\phi_d = 0.024$ and then set this probability to zero. We find that savings go down by 4.2584% within the group of liquidity-constrained households, and 0.0292% when all households are considered. The removal of marital risk eliminates a need for precautionary savings and, therefore, decreases households savings. This result is consistent with the findings in González and Özcan (2008).

We now compute the elasticity of the assets-to-income ratio with respect to UI in this economy with marital risk, which is a measure of the precautionary motive. This elasticity is -0.48 , against the -0.38 we found under no marital risk. As expected, the extra risk brought about by marital shocks increases the responsiveness of household asset holdings to changes in UI.

Our final exercise computes consumption losses upon unemployment and/or divorce shocks. That is, we take an agent who is married and employed in period $t - 1$, and who transits to period t either as: (i) married and unemployed, (ii) single and employed, or (iii) single and unemployed. We compute the average loss in consumption of this generic agent within the group of liquidity-

¹⁵Even though different divorce rules and/or an endogenous marriage decision may have important quantitative effects, we believe our model provides a useful benchmark to assessing the effects of marital transitions.

¹⁶See González and Özcan (2008) for a recent empirical study using the legalization of divorce in Ireland in 1996. They find results suggesting that the risk of divorce brought about by the law was followed by an increase in the propensity to save of married couples.

constrained households. The consumption loss upon the unemployment shock, (i), is -6.42% . The marriage shock, (ii), implies a loss of consumption of -2.13% , and a combined unemployment and marriage shock, (iii), brings about a loss of consumption of -34.32% . In our framework without public goods within the household and without divorce costs, the unemployment shock has larger consumption effects than a marriage shock. The two shocks combined amplify consumption losses.

Three remarks concerning the scope of this version of the model, where married and single households co-exist, are in order. First, as indicated above, an implication of limited commitment is non stationarity in Pareto weights, a feature our model is silent about. It is conceivable that upon being hit by the unemployment shock, an agent sees her(his) Pareto weight reduced in favor of his (her) employed spouse. This shift towards the employed spouse may weaken the savings and labor supply effects of intra-household insurance, especially among low-wealth households. Second, if we interpret cohabitation as a form of marriage without family insurance, and assume that singles in our model are cohabiting couples (i.e., two singles under the same roof), then our model accounts for the differential in savings between cohabiting and married couples observed in the U.S. Negrasa and Oreffice (2010) use the 2000 U.S. Census and show that cohabiting heterosexual couples save more than their married counterparts. The explanation for this differential provided by our model is higher precautionary savings. Third, if we interpret singles in our model as non-cohabiting adults, then our model cannot account for the higher savings rate of married versus single households. The explanation of this failure is that our model abstracts from important drivers of household savings. Two such omissions are life-cycle considerations and households' housing investment.

5 Concluding Remarks

In this paper we assess quantitatively the effects of intra-household risk sharing on savings and labor supply within a model of idiosyncratic unemployment risk. With this purpose, we present a model economy where households are formed by a female and a male, who make collective decisions and lack access to a complete capital market. Our model is a dynamic version of the standard collective model of the household developed by Chiappori and co-authors since the 1980's, which assumes efficient risk sharing within the household. Equipped with this model, we then ask about the quantitative effects of this informal insurance arrangement on households' savings and labor supplies, on the extent of the precautionary motive, on the crowding out effects of public insurance and on consumption losses upon unemployment. In light of our results, we conclude that intra-household risk sharing has large quantitative effects on all these margins explored. Importantly, we find that our model economy accounts relatively well for key elasticities of savings and spousal labor supply with respect to UI, as estimated by Engen and Gruber (2001) and Cullen and Gruber

(2000), respectively. We also show that standard models, which abstract from intra-household risk sharing, fail to match those elasticities. A conclusion we draw from the exercise in this paper is that ignoring risk sharing at the level of the household introduces an important bias not only on the extent of the precautionary motive but also on the distortionary effects of public insurance programs.

The model presented in this paper can be used to address a number of related questions. In particular, we plan to use versions of this model to shed further light on a recent debate about gender-based taxation. A number of scholars have argued in favor of taxing females and males differently on the grounds of their different elasticities of labor supply. The interplay of income tax rates with Pareto weights within the household is bound to introduce tradeoffs that have been so far overlooked in this debate. An important extension that is worth pursuing is the consideration of permanent income shocks under no commitment to future household allocations. As mentioned above, under no commitment Pareto weights are non stationary, a feature that is likely to affect the role played by intra-household risk sharing. Finally, in light of recent results on the degree of household insurability against different types of shocks and the evolution of consumption and income inequality (see, e.g., Blundell *et al.* 2008), we need models that can account for the observed ability of households to insure different kinds of risks. Models with two-person households and perfect risk sharing within the household are a first step in this direction.

References

- [1] Aiyagari, S.R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109. pp. 659-684.
- [2] Attanasio, O. and Rios-Rull, J.V. (2000), “Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?” *European Economic Review*, 44, pp. 1225-1258.
- [3] Attanasio, O., Low, H. and Sanchez-Marcos, V. (2005), “Female Labor Supply as Insurance Against Idiosyncratic Risk,” *Journal of the European Economic Association*,
- [4] Blundell, R. and MaCurdy, T.E. (1999), “Labor Supply: A Review of Alternative Approaches,” in: Ashenfelter, Q., Card, D. (Eds.), *Handbook of Labor Economics*, vol. 3A. Elsevier, Amsterdam.
- [5] Blundell, R., Pistaferri, L. and Preston, I. (2008), “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98(5), pp. 1887-1921.
- [6] Browning, M., Bourguignon, F., Chiappori, P.A. and Lechene, V. (1994), “Income and Outcomes: A Structural Model of Intrahousehold Allocation,” *Journal of Political Economy*, 102(6), pp. 1067-1096.

- [7] Browning, M., Chiappori, P.A. and Lechene, V. (2006), "Collective and Unitary Models: A Clarification," *Review of Economics of the Household*, 4, pp. 5-14.
- [8] Browning, M. and Crossley, T.F. (2001), "Unemployment Insurance Benefit Levels and Consumption Changes," *Journal of Public Economics*, 80, pp. 1-23.
- [9] Chetty, R. and Saez, E. (2010), "Optimal Taxation and Social Insurance with Endogenous Private Insurance," *American Economic Journal: Economic Policy*, 2, pp. 85-114.
- [10] Chiappori, P.-A. (1988), "Rational Household Labor Supply," *Econometrica*, 56, pp. 63-89.
- [11] Chiappori, P.-A. and Donni, O. (2010), "Non-unitary Models of Household Behavior: A Survey of the Literature," in: Molina, A. (eds.) *Household Economic Behaviors*, Berlin: Springer.
- [12] Cullen, J.B. and Gruber, J. (2000), "Does Unemployment Insurance Crowd out Spousal Labor Supply?," *Journal of Labor Economics*, 18, pp. 546-572.
- [13] Domeij, D. and Floden, M. (2001), "The labor-supply elasticity and borrowing constraints: why estimates are biased," *Review of Economic Dynamics*, 9, pp. 242-262.
- [14] Engen, E.M. and Gruber, J. (2001), "Unemployment Insurance and Precautionary Savings," *Journal of Monetary Economics*, 47, pp. 545-579.
- [15] Flodén, M. and Lindé, J. (2001), "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, 4, pp. 406-437.
- [16] Golosov, M. and Tsyvinski, A. (2007), "Optimal Taxation with Endogenous Insurance Markets," *Quarterly Journal of Economics*, 122, pp. 487-534.
- [17] González, L., and Özcan, B. (2008), "The Risk of Divorce and Household Saving Behavior," IZA DP No. 3726.
- [18] Gorbachev, O. and Dogra, K. (2010), "Evolution of Consumption Volatility for the Liquidity Constrained Households over 1983-2004," mimeo.
- [19] Gruber, J. (1997) "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87, pp. 192-205.
- [20] Heathcote J., Storesletten, K. and Violante, G.L. (2010), "The Macroeconomic Implications of Rising Wage Inequality in the U.S.," *Journal of Political Economy*, 118(4), pp. 681-722.
- [21] Heathcote J., Storesletten, K. and Violante, G.L. (2009), "The Macroeconomics of the 'Quiet Revolution'," mimeo.
- [22] Huggett, M. (1993), "The Risk-free Rate in Heterogenous-agent Incomplete Insurance Economies," *Journal of Economic Dynamics and Control*, 17, pp. 953-969.

- [23] Imrohoroglu, A. (1989), "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 97(6), pp. 1364-1383.
- [24] Kotlikoff, L.J. and Spivak, A. (1981), "The Family as an Incomplete Annuities Market," *Journal of Political Economy*, 89, pp. 372-391.
- [25] Krusell, P., Mukoyama, T. and Şahin, A. (2010), "Labor-Market Matching with Precautionary Savings and Aggregate Fluctuations," *Review of Economic Studies*, 77(4), pp. 1477-1507.
- [26] Krusell, P. and Smith Jr., A. (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, pp. 867-896.
- [27] Low, H., Meghir, C. and Pistaferri, L. (2008), "Wage Risk and Employment Risk over the Life Cycle," WP08/06 The Institute For Fiscal Studies.
- [28] Marcet, A., Obiols-Homs, F. and Weil, P. (2007), "Incomplete Markets, Labor Supply and Capital Accumulation", *Journal of Monetary Economics*, 54(8), pp. 2621-2635.
- [29] Mazzocco, M. (2004), "Savings, Risk Sharing and Preferences for Risk," *American Economic Review*, 94(4), pp. 1169-1182.
- [30] Mazzocco, M. (2007), "Household Intertemporal Behavior: A Collective Characterization and a Test of Commitment," *Review of Economic Studies*, 74(3), pp. 857-895.
- [31] Mazzocco, M., Ruiz, C. and Yamaguchi, S. (2008), "Labor Supply, Wealth Dynamics and Marriage Decisions," California Center for Population Research (CCPR).
- [32] Negrasa, B., Oreffice, S. (2010), "Sexual Orientation and Household Savings: Do Homosexual Couples Save More?," IZA DP No. 4961.
- [33] Pijoan-Mas, J. (2006), "Precautionary Savings or Working Longer Hours?," *Review of Economic Dynamics*, 9, pp. 326-352.
- [34] Shore, S.H. (2010), "For Better, for Worse: Intra-household Risk-sharing over the Business Cycle," *Review of Economics and Statistics*, 92(3), pp. 536-548.
- [35] Shore, S.H. and Sinai, T. (2010), "Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?" *Review of Economics and Statistics*, 92(2), pp. 408-424.
- [36] Stephens Jr., M. (2002), "Worker Displacement and the Added Worker Effect," *Journal of Labor Economics*, 20(3), pp. 504-537.
- [37] Zeldes, S. P. (1989), "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97(2), pp. 305-346.

Table 1. Baseline Parameters

Description	Parameter	Value	Description	Parameter	Value
Relative risk aversion	σ	2	Discount factor	β_H	0.989
Regulates Frisch elasticity	γ	3	Discount factor	β_L	0.982
Utility weight	φ	1.27	Share of impatient households	χ	0.20
Pareto weight	μ	0.575	Unemployment insurance	b^f	0.083
Derivative Pareto weight	μ_1	0.038	Unemployment insurance	b^m	0.161
Capital share	α	0.36	Relative wages	λ	0.575
Capital depreciation rate	δ	0.025			

Table 2. Unemployment Insurance and Financial Assets

	Elasticity of average assets-to-income ratio w.r.t. UI replacement rate
Data (Engen and Gruber 2001)	-0.28
Collective Household Economy	-0.38
Bachelor Household Economy	-0.64

Notes: This table shows how household asset holdings respond to the generosity of UI.

Table 3. Female Labor Supply and Male Employment Status

	Liquidity-constrained households	All households
Employed Husband	166.0	145.0
Unemployed Husband	180.2	148.8

Notes: This table shows average monthly hours of work by working females in households with employed and unemployed males in our baseline economy with collective households.

Table 4. UI and Female Labor Supply During Male's Unemployment Spells

	Liquidity-constrained	
	households	All households
10% reduction in b^m	+0.86%	+0.28%
50% reduction in b^m	+8.71%	+1.53%

Notes: This table shows the percentage increase in female labor supply upon a male's unemployment spell yielded by 10% and 50% reductions in UI in our economy with collective households.

Table 5. Individual and household consumption loss upon unemployment

(a) Individual consumption loss upon unemployment:				
	Collective Model		Bachelor Model	
	Liquidity-constrained individuals	All individuals	Liquidity-constrained individuals	All individuals
Females, $\Delta c^f / c^f$	-4.5%	-0.92%	-23.23%	-2.46%
Males, $\Delta c^m / c^m$	-8.74%	-1.78%	-30.09%	-2.44%

(b) Fraction of household income loss that transmits to household consumption loss:				
	Collective Model		Bachelor Model	
	Liquidity-constrained households	All households	Liquidity-constrained households	All households
$\Delta c / \Delta y$	17.31%	3.55%	35.66%	3.34%


Notes: Panel (a) of this table presents individual insurance as measured by the percentage of consumption lost upon an unemployment shock. Panel (b) presents household insurance as measured by the degree of transmission of income loss to consumption upon an unemployment shock.

Table 6. Elasticity of Household Consumption Loss to UI

	Liquidity-constrained	
	households	All households
Data (Browning and Crossley 2001)	-0.0922	-0.05
Collective Household Economy	-0.0954	-0.0175
Bachelor Household Economy	-0.2262	-0.0212

Notes: Sensitivity of household consumption loss upon unemployment with respect to the generosity of UI.

(a) Savings Policy Functions



(b) Policy Functions for Hours Worked







Figure 1: Policy Functions in the Collective Household Economy

(a) Excess Savings (Low Beta)



(b) Excess Savings (High Beta)






Figure 2: Savings Effect of Intra-household Risk Sharing

Excess Hours Worked





Figure 3: Hours Effect of Intra-household Risk Sharing