

DISCUSSION PAPER SERIES

DP16591

The Insurance Role of the Family

Hans Fehr and Fabian Kindermann

LABOUR ECONOMICS
MACROECONOMICS AND GROWTH
PUBLIC ECONOMICS

CEPR

The Insurance Role of the Family

Hans Fehr and Fabian Kindermann

Discussion Paper DP16591
Published 30 September 2021
Submitted 29 September 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Labour Economics
- Macroeconomics and Growth
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Hans Fehr and Fabian Kindermann

The Insurance Role of the Family

Abstract

Besides love and affection the family also provides economic benefits. Beyond gains from specialization and economies of scale, it serves as a provider of insurance against various risks individuals face throughout their life. This insurance role of the family has changed during past decades owing to several factors: a fundamental transition in the gender wage gap and female labor force participation, the legal framework, and the dynamics of household formation over the life cycle. The present chapter reviews recent studies that quantify the importance of family insurance and studies its interplay with social insurance as well as the private insurance market.

JEL Classification: D10, D15, D19, J22

Keywords: family insurance, Household decision making, Added Worker Effect

Hans Fehr - hans.fehr@uni-wuerzburg.de
University of Würzburg

Fabian Kindermann - fabian.kindermann@ur.de
University of Regensburg and CEPR

The Insurance Role of the Family

Hans Fehr
University of Würzburg,
CESifo, and Netspar

Fabian Kindermann
University of Regensburg,
CEPR, HCEO, and Netspar

September 29, 2021

Abstract

Besides love and affection the family also provides economic benefits. Beyond gains from specialization and economies of scale, it serves as a provider of insurance against various risks individuals face throughout their life. This insurance role of the family has changed during past decades owing to several factors: a fundamental transition in the gender wage gap and female labor force participation, the legal framework, and the dynamics of household formation over the life cycle. The present chapter reviews recent studies that quantify the importance of family insurance and studies its interplay with social insurance as well as the private insurance market.

JEL Classification: D10, D15, D19, J22

Keywords: family insurance, household decision making, added worker effect

Contact Information:

Fehr: hans.fehr@uni-wuerzburg.de, Kindermann: fabian.kindermann@ur.de

1 Introduction

Family decisions are central determinants of macroeconomic outcomes: marital behavior and assortative mating shape the income and wealth distribution; fertility choices and parental investment in the education of children can explain a large part of the Western industrialisation process; and labor supply decisions are predominantly made by two potential earners in a household context. Not surprisingly, recent surveys by Doepke and Tertilt (2016) or Greenwood et al. (2017) document that intra-family decision making and its consequences for the macroeconomy have increasingly received attention in the economic literature.

Over and above these aggregate outcomes, the joint behavior of family members and the sharing of resources within a family may also provide insurance against shocks in an uncertain world. Since information barriers are typically fairly low within a family, such an insurance could even outperform private or public insurance schemes, which may suffer from typical market failure problems. There are various risks individuals face throughout their life cycle that can be insured within the family context. When the primary earner in the family loses her or his job, a secondary earner might step in to absorb the looming loss of family income. This so-called “added-worker effect” can shape the optimal design of unemployment or disability insurance, but can also impact economic dynamics along the business cycle. A surviving spouse might inherit the entire family wealth and use it to deal with longevity risk. Children may help out their parents when they become old, providing informal care or company. In turn, when a government competes with such family insurance arrangements by providing social insurance, for example through social security, unemployment insurance or social assistance schemes, this can have a significant impact on family ties, family decision making and also family formation.

The idea that the family acts as a risk-sharing institution is not new. In their seminal paper, Kotlikoff and Spivak (1981) already have argued that many family practices that deal with job loss, bankruptcy or longevity risk can be explained as implicit insurance contracts written *ex ante* between selfish family members. In a quantitative analysis, they show that a marriage at age 30 can substitute roughly 45 to 50 percent of a fairly priced annuity. Brown and Poterba (2000) compare the valuation of fair annuities for singles and married couples under various specifications for preferences, pre-existing annuities and survivor benefit ratios. Owing to extensive risk sharing possibilities within the couple, they argue that the utility gains from joint life annuities are small. In some circumstances they do not even compensate actual loads on annuity products, which may to some degree explain the so-called annuity puzzle. Di Tella and MacCulloch (2002) argue that publicly provided unemployment benefits may crowd out intra-family transfers by more than one-for-one, since defection from informal family insurance arrangements becomes more attractive and employed family members have to pay taxes. When family ties are strong (so that the family provides a high level of insurance), the optimal size of the welfare state is zero. But when family ties are weak,

the government should be the only provider of insurance. Leroux and Pestieau (2014) come to a very similar conclusion with respect to the pension system. Of course, this view also implies that the absence or weakness of private and public insurance provision (for example in developing countries) may be an important incentive for marriage. Boldrin and Jones (2002) develop an overlapping generation model with fertility choice, in which parents' old-age consumption directly enters the utility function of their children. This generates an endogenous transfer from children to parents and helps to explain the demographic transition prior to the industrialization. Boldrin et al. (2015) also highlight the negative effects of government provided pensions and increased access to capital markets on fertility. An altruistic motive of children towards their parents may also alter the optimal design of long-term care (LTC) policy for dependent parents, see Pestieau and Sato (2008).

While many of these studies influenced our view of the role of the family as provider of insurance, this chapter focuses on two aspects that have shaped the more recent empirical and quantitative literature. First, the emergence of micro data sets that combine detailed socio-demographic and labor market information allows a deeper and more sophisticated empirical analysis of the interaction between human capital accumulation, family formation and labor supply decisions. As noted by Goldin (2020), changes in social norms, educational achievements and career opportunities nowadays allow women to follow a “career and family” goal, which in turn strengthens the insurance role of the family. Consequently, the risk-sharing consequences of labor market search and participation decisions of married couples are the focus of Sections 2 and 3.¹ Second, the rapid advances in computational speed and software allowed researchers to develop quantitative partial and general equilibrium models with various sources of risk as well as family structures, which allow to isolate the insurance role of families in a dynamic perspective. Section 4 therefore discusses results from such approaches with respect to individual behavior, while Section 5 concentrates on the interplay between public insurance and the formation of families. Finally, we offer an alternative view on families as both a provider insurance and a source of risk in Section 6. Section 7 offers some concluding remarks and points to potential avenues for future research.

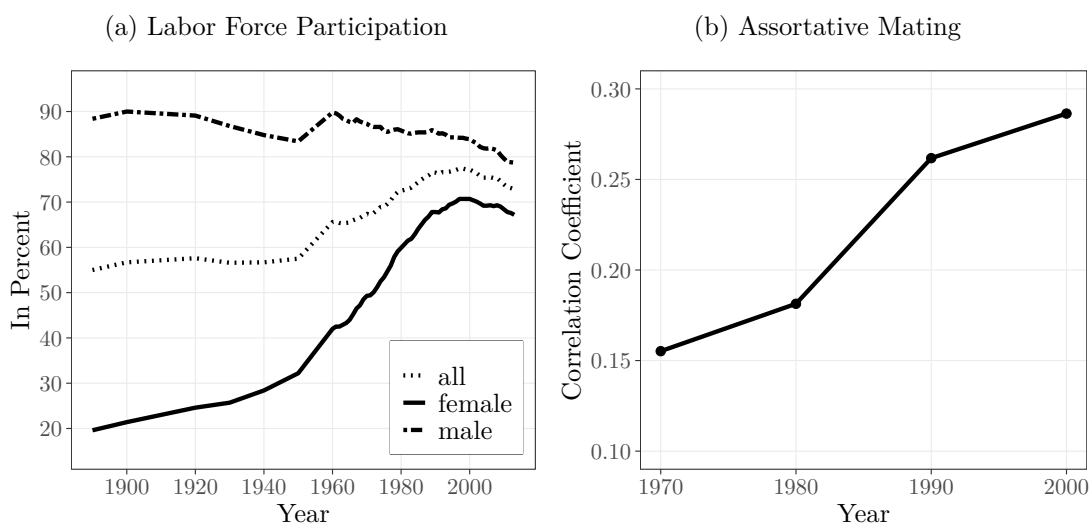
2 Family insurance and the labor market

In the twentieth century, the family underwent a substantial transition. Families became smaller, educational attainment rose and there were fewer children. In the context of labor market performance, there have been two trends that stand out very prominently. A substantial rise in female labor force participation on the one hand, and an increasing degree of “economic homogamy” on the other. The left panel of Figure 1 shows the labor force participation of men and women in the U.S. over the course of the twentieth and the early twenty-first century.

¹We thereby deliberately exclude a huge body of literature that deals with fertility decisions, see Doepke et al. (2022) for a recent survey.

While around the year 1900, about 20 percent of all women participated in the labor force, this share rose to a substantial 70 percent in 2000 and since then has approximately stagnated. The labor force participation of men, on the other hand, has been constant at 90 percent until the mid twentieth century (with the exception of World War II), and since then has declined slightly, but remained at a high level. The trends that accompany the rise in female labor force participation are manifold, amongst them a substantial decline in fertility (Doepke et al., 2022), a rise in the return to female labor supply (Galor and Weil, 1996), the ability to better combine families and careers (Goldin, 2020), and a technological revolution in the household (Greenwood et al., 2005).²

Figure 1: Changes in Labor Force Participation and Assortative Mating in the US



Sources: (a) Figure reconstructed from Doepke and Tertilt (2016).

(b) Figure reconstructed from Bredemeier and Juessen (2013).

The evolution in economic homogamy relates to the observation that educational attainment, wages as well as the earnings of married partners have become more alike over time. A rising assortative mating in terms of educational attainment has been documented in, for example, Schwartz and Mare (2005), Blossfeld (2009) or Fernández et al. (2005). Bredemeier and Juessen (2013) measure the correlation coefficient of husbands' and wives' decile positions in the wage distribution at different points in time using data from the Current Population Survey (CPS). They find that this correlation coefficient has almost doubled in between the 1970s and 2000s, see the right panel of Figure 1. At the same time, they show that weekly working hours have been growing the most among women who are married to higher wage men in the same time span. A quantitative assessment in a standard model of household labor supply reveals that this change in female labor supply behavior can be attributed to a large part to assortative mating.

²Greenwood et al. (2016) provide a quantitative assessment of the importance of many of these channels in a unified model of marriage, divorce, educational attainment, and married female labor-force participation.

While labor earnings from female spouses have always been an important source of income within wage poorer households, the fact that high wage women are nowadays married more frequently to high wage men has led to a reallocation of paid labor within richer couples such that partner’s earnings become more alike, see also Gonalons-Pons and Schwartz (2017). There may be several reasons for the rise in assortative mating. Blossfeld and Timm (2003) argue that colleges and universities act as marriage markets and that the expansion of higher education (especially for women) has contributed substantially to assortative mating. Calvo et al. (2021) attribute sorting in the marriage market and sorting in the labor market to the nature of home production. When home production inputs of two partners are complements, then there is positive sorting in the marriage and the labor market in equilibrium, with the result that partners share home tasks more equally and provide similar labor market hours.

2.1 The added worker effect

The transition in family labor supply described above has important implications for the degree of labor market insurance a couple can provide to themselves. When both adult members of a family have a strong attachment to the labor force and a similar earnings capacity, the distribution of earnings within the family will be more equal than under a one-breadwinner model. But most importantly, one partner can compensate for potential transitory earnings losses of the other partner. Lundberg (1985) is among the first to investigate this so-called “added worker effect” both conceptually and empirically. She argues that, if one wants to measure the added worker effect consistently in the data, one has to pay tribute to its nature as an insurance device. As such, the added worker effect should be understood as a temporary change in a wife’s labor force participation upon a temporary (an most likely exogenous) lay-off of the husband. To consistently estimate this effect, Lundberg (1985) calculates joint monthly labor force transition rates between employment, unemployment and non-participation of husbands and wives. Her estimates show that, at least for for a white woman, the probability of entering the labor force is about 25 percent higher if her husband is unemployed than when he is employed. In addition, her probability of leaving the labor force is 33 percent lower. For Black and Hispanic women, the evidence of an added worker effect is less clear. These estimates seem large at first sight. However, one has to bear in mind that monthly transition rates are generally very small. To put these numbers into perspective, Lundberg (1985) simulates impulse response functions of female participation and employment upon the lay-off of a husband. Her simulations reveal that if 100 men were to become unemployed, this would lead to an additional participation of three wives and an employment of two wives.

The small quantitative magnitude of the added worker effect as measured by Lundberg (1985) can be rationalized in several ways. First, changing from non-participation to participation is a quite costly task for every worker. Hence, a couple will only fall back on such a strategy at last resort, when all other options like intensive search of the primary earner and running down liquid savings are not

an option (Heckman and MaCurdy, 1980). A second contributing factor may be the time span of 1969 to 1973 in which the data was collected, a time where female labor force participation was not at its peak yet and where the correlation of wages between wives and husbands was still rather low. Mankart and Oikonomou (2016) support this latter idea. Using data from the CPS, they show that the strength of the added worker effect has in fact increased in the U.S. over the time period from 1980 to 2000. While in the 1980s the likelihood to enter the labor force for a wife with an unemployed husband was about 5 percentage points higher than for a wife with an employed husband, this likelihood has doubled to 10 percentage points in the 2000s. Within a standard search model with two household earners, they show that this increase in the strength of the added worker effect in the U.S. can be rationalized by a decline in the gender wage gap, changes in search frictions and changes in the labor force participation costs of women.

By how much a couple relies on the added worker effect to self-insure unemployment or wage shocks certainly depends on the economic environment they are living in. Cullen and Gruber (2000) point to the fact that government provided unemployment insurance (UI) can not only crowd-out individual search effort, it might also impact on the labor supply of other family members during periods of unemployment. Using data from the Survey of Income and Program Participants (SIPP), they quantify the crowd-out effect of UI for husbands on the labor supply of wives and find it to be substantial. They report that “for each dollar of unemployment insurance receipt wives earn 73 cents less”. Choi and Valladares-Esteban (2020) elaborate on this idea in a quantitative incomplete markets model with one- and two-earner households who can decide to work, save and consume. Their findings indicate that married couples value unemployment insurance substantially less than singles, which has important implications for the design of public policy. Summing up, the design of public insurance and redistribution is important in understanding the extent of family-provided insurance. This also helps us rationalize that the empirical findings related to the added worker effect vary substantially across countries with different welfare regimes. It is therefore not surprising that the evidence for an active family-base unemployment insurance through an added worker effect is much weaker in Austria (Halla et al., 2020), Germany (Illing et al., 2021), the Netherlands (De Nardi et al., 2021), and Norway (Blundell et al., 2015).

2.2 From earnings to consumption insurance

Despite the rather weak evidence for an admittedly quite narrow definition of the added worker effect, two-earner couples can still provide insurance. The mere pooling of resources for consumption can act as insurance device against individual fluctuations in consumption and therefore utility. To see the driving forces of this type of insurance, let us assume that a couple consists of two partners who each receive some earnings w_1 and w_2 , respectively. Earnings are risky and follow a joint distribution with individual variances $\sigma^2(w_1)$ and $\sigma^2(w_2)$ as well as a correlation coefficient $\rho(w_1, w_2)$. If each of the partners consumed their individual earnings

entirely on their own, i.e. $c_i = w_i$, then the variances of individual consumption would just be

$$\text{Var}(c_1) = \sigma^2(w_1) \quad \text{and} \quad \text{Var}(c_2) = \sigma^2(w_2).$$

If instead the couple entirely pooled their resources and each partner always got half of the family income to consume, then the variance of individual consumption would read

$$\text{Var}(c_1) = \text{Var}(c_2) = 0.25[\sigma^2(w_1) + \sigma^2(w_2)] + \varrho(w_1, w_2)\sigma(w_1)\sigma(w_2).$$

In the rather special but illustrative case in which both variances were identical, $\sigma^2(w_1) = \sigma^2(w_2)$, and the earnings of the two partners were uncorrelated $\varrho(w_1, w_2) = 0$, income pooling would obviously halve the variance of consumption of each partner.

While of course very simple and stylized, the previous considerations already point to the major themes and potential obstacles of family consumption insurance:

- In order to provide such insurance, both partners need to work and generate earnings. Couples facing higher risk – either because they don't have other means of smoothing consumption, like liquid savings, or because wage risk generally increases – should hence be characterized by a higher labor force participation or longer labor hours of the secondary earner.
- In turn, when a couple has to rely on specialization in the family, for example due to the presence of children, this might impede family insurance.
- Furthermore, when shocks between partners are strongly correlated, for example when both partners work in the same occupation or even the same firm, then consumption smoothing opportunities might be small.
- Lastly, when partners don't share their income equally, for example because of different bargaining positions or limited commitment to consumption plans, then intra-family insurance might be weakened.

In the following we will point to several studies that investigate these issues further.

Blundell et al. (2008) were among the first to empirically study the extent of consumption insurance against permanent and transitory income shocks. They construct a new dataset based on information from both the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX) to investigate how increases in earnings risk over the 1970s to 1990s materialized in consumption inequality. They report a divergence between the consumption and the income distribution over this time period. More specifically, while the variance of earnings and consumption increased almost in parallel in the early 1980s, the two series

decoupled around 1985. The variance of earnings then increased further throughout the 1990s, but the variance of consumption stagnated pointing to additional insurance possibilities for households. Blundell et al. (2008) attribute this finding to the nature of earnings shocks. The increase in earnings risk in the early 1980s was due to an increase in the variance of permanent shocks, which was then replaced by an increase in the variance of transitory shocks in the late 1980s and early 1990s. Since it is much easier for households to insure transitory shocks, the pass-through of earnings risk into consumption risk weakened. Their results also indicate that, next to progressive taxes and government transfers, family labor supply might have played a role for households in smoothing earnings shocks.

The extent of intra-family insurance against transitory shocks is at the heart of Ortigueira and Siassi (2013). They study a general equilibrium model with incomplete markets, in which households consisting of a female and a male worker are subject to idiosyncratic unemployment shocks. Unemployment shocks are transitory and can be correlated across partners. Households decide about how much to consume, work and save in each period and can use both precautionary savings as well as intra-family reallocations of labor supply to smooth shocks over time. Ortigueira and Siassi (2013) point to the fact that intra-family insurance through labor supply is mostly prevalent among low-wealth, liquidity constrained households. In line with the discussion in the previous section, households with sufficient wealth can smooth shocks by running down their buffer stock savings. Low-wealth households, on the other hand, provide insurance primarily through the added worker effect. To quantify the extent of intra-family insurance, the authors compute the pass-through of income losses from an unemployment spell to consumption expenditure. While about 35 percent of a liquidity constrained single household's income loss directly transmits into consumption, this number only amounts to 17 percent for a comparable household with access to family insurance. Furthermore, they point to the fact that the availability of family insurance might dampen the accumulation of precautionary savings.

Blundell et al. (2016) make a clear distinction between the insurance role of the family against transitory and permanent wage shocks. They exploit detailed data on family labor earnings, wealth and consumption from the PSID to estimate a structural life-cycle model of joint family decision making. Their model features non-separabilities in both consumption and labor supply as well as partners' leisure consumption, the latter reflecting the idea that couples might enjoy spending time together. There are three potential channels of insurance against wage shocks: (i) self-insurance through precautionary savings, (ii) family insurance through joint labor supply responses to shocks, and (iii) governmental insurance through progressive taxes. Not surprisingly, Blundell et al. (2016) find that both partners in a family are exposed to transitory and permanent wage shocks and that these shocks tend to be somewhat positively correlated. While the evidence on the correlation coefficient is not very strong from a statistical perspective, their findings at least provide some indication for the view that couples tend to work in similar jobs, which might be a result of increased assortative mating (Juhn

and Potter, 2007). By looking at transitory and permanent shocks separately, the authors can provide estimates for two distinct but equally important sets of elasticities. Transitory wage shocks, as affecting household income in only one year, can be interpreted as leaving household wealth approximately constant. The resulting elasticity estimates to such shocks hence resemble the Frisch labor supply and consumption elasticities. Permanent wage shocks, on the other hand, affect household wealth substantially, and hence should be interpreted as Marshallian elasticities.

The findings in Blundell et al. (2016) indicate that a couple's response to transitory and permanent after-tax wage changes is quite distinct. While the Frisch elasticity responses to own wage shocks is quite large – about 0.7 for men and 1.0 for women – and in line with previous finding in the literature, the Marshallian elasticities are much smaller – about -0.1 for men and 0.4 for women. Hence, wealth effects apparently are important for understanding households' labor supply choices. More important from the point of view of this chapter, however, is the finding that the cross-elasticities are substantially different for transitory and permanent wage shocks. If one partner's wage falls transitorily, then both members of a couple will reduce their labor hours indicating that (from a Frisch perspective) labor hours of the partners are complements. At the same time, consumption expenditure of the family increases, which has two economic interpretations: First, leisure and monetary consumption of a couple are complements in the short-run and, second, the typical couple has enough resources to fully insure a transitory wage shock by resorting to buffer stock wealth, see again Blundell et al. (2008) or Kaplan and Violante (2010). The picture looks different upon the arrival of a permanent wage shock, which is much harder to insure by means of precautionary savings. Once the wage of a partner falls permanently, labor supply of the other partner increases in order to counteract the decline in wage income. Consumption of the family falls as a result of the pronounced wealth effect. Note that the compensating labor supply response to negative permanent wage changes is most pronounced for wives when the husband's wage falls. In the other direction, the evidence of family insurance is much weaker, which might be a result of the fact that the selected sample only includes stable married couples with permanently working male spouses. Nevertheless, these results show that family insurance is a powerful means to insure permanent wage changes in particular. To summarize their findings, they “calculate that, on average, of the total amount of consumption ‘insured’ against permanent shocks to the male's wage through behavioral responses, about 63 percent comes from family labor supply and only about 17 percent comes from self-insurance through savings, with the rest explained by taxes and transfers.” (Blundell et al., 2016).

Wu and Krueger (2021) provide a more structural decomposition of family insurance. They replicate the permanent and transitory wage shocks and the corresponding consumption insurance found in Blundell et al. (2016) in a partial equilibrium life-cycle model with endogenous household labor supply. Their model allows to isolate the mechanisms by which a wage shock to the male earner in

the family is mitigated until it ends up in household consumption. In the case of permanent wage shocks, consumption insurance increases by almost 20 percentage points when moving from a one-earner model with exogenous labor supply, in which shocks can only be insured by precautionary savings, to a two-earner household model with endogenous family labor supply. While allowing for a labor supply response of the primary earner raises consumption insurance by 5.9 percentage points, the presence of a second earner contributes an additional 13.8 percentage points of insurance. With all insurance mechanisms present, male labor supply actually falls in response to a negative wage shock which highlights the importance of female earnings and labor supply adjustments. The insurance provided by a spouse is roughly constant over the life cycle, while the insurance provided by savings and social security rises with age so that older households predominantly rely on savings to smooth consumption. These findings are confirmed by welfare calculations which show that labor supply responses of the secondary earner reduce the overall welfare cost of wage shocks for the primary earner by roughly 40 percent.

Extending their own structural model by allowing for the presence of children and time spent on home production, Blundell et al. (2018) show that complementarities in leisure consumption as well as the substitutability of time input into home production of two partners are both important in understanding the responses to wage shocks within the family. In their structural life-cycle model, the arrival of children is exogenous, but once children are there, parents have to provide child care services. Within this framework, there are two important observations to make. First, and unlike in Blundell et al. (2016), the labor supply response of a mother with respect to an increase in her husband's wage is negative. If there were only complementarities in leisure consumption, one would expect the response to be of opposite sign, as increased work hours for the husband should also lead to higher labor supply of the wife. However, in the presence of home produced child care, an expansion of work effort of the husband leads to a decline in his hours of child care services, which needs to be compensated by more hours spent with children by the wife. Second, the Frisch elasticity of a mother with respect to her own wage is large, as a higher wage leads the cost of home produced child care to increase.

The previous discussion has mostly centered around consumption insurance in the cross section, assuming wage risk to be constant over time. Another strand of literature tries to identify how wage risk has changed over time and how this might have affected risk sharing within the family. Attanasio et al. (2005) were one of the first to develop a quantitative life-cycle model to study the role of female labor supply as an insurance device against income risk. They consider a family household consisting of two opposite-sex members who optimally choose consumption, savings, and female labor supply. The male spouse always works and his human capital evolves exogenously over the life-cycle, but human capital of the female spouse is endogenous to prior labor market experience. Female labor force participation, in turn, decreases with the relative cost of purchasing child

care services. The model calibration matches the main life cycle participation behavior of U.S. women born in the 1940s. In a counterfactual exercise, Attanasio et al. (2005) then double the size of permanent income shocks and compute the resulting changes to savings and female labor supply. As one would expect, higher uncertainty increases female labor market participation especially when children have left the household. This labor supply effect is much stronger when the household is borrowing constrained or is not able to self-insure via savings. Therefore, the welfare cost of higher income uncertainty can be reduced by adjusting female labor supply accordingly.

Heathcote et al. (2010) bring the analysis of rising wage inequality to a fully-fledged general equilibrium model with overlapping generations. Men and women in their model first optimally choose between two levels of education. Based on their educational choice they are then matched to form a household and enter the labor market as married couples. The household optimizes consumption, savings, and individual labor supply facing idiosyncratic labor productivity risk. Within this model, the authors strive to quantify the impact of changes in the U.S. wage distribution from the mid 1960s onwards on the macroeconomy and household welfare. Changes in the wage distribution are modeled by shifts in structural parameters like the demand for skilled labor, the demand for female labor, the variance of wage shocks, total factor productivity and the costs of education. The calibration broadly reproduces the empirical trends of the cross-sectional distribution over working hours, earnings and consumption after 1965. Heathcote et al. (2010) finally use the model to isolate the effects of the single drivers of rising wage inequality. On average, changes in the wage structure led to an increase in household welfare, but some poorer household types are hit by adverse demand shifts in the 1980s. The average welfare gains mainly emerge from increased educational participation and a reallocation of time within the household as reactions to changes in the college premium, the gender wage gap and the volatility of wage shocks. While the model accounts for roughly three-quarters of the increase in female labor hours over the considered time period, this increase is mainly driven by a narrowing gender wage gap and not by the rise in uncertainty.

The welfare costs of a rising volatility of wage risk is also under investigation in Park and Shin (2020), who compare the U.S. in the early 1970s to the U.S. in the early 2000s. Instead of modeling an education choice and a matching stage, they consider heterogeneous risk preferences and allow individuals to choose the risk profile of their job. Furthermore, they point to the fact that, in the time period they consider, the volatility of male earnings has increased but the volatility of female earnings declined. This suggests that gender differences in the trends of wage volatility might be important for understanding the welfare costs of increased wage risk. In fact, they find that their model exhibits a much smaller welfare cost of rising wage risk as compared to the standard family labor supply model with homogeneous risk aversion, without self-selection into jobs and with uniform wage risk across genders. As in previous studies, Park and Shin (2020)

find that households are effective in smoothing transitory wage shocks through precautionary savings, but they can also share permanent wage risk through the added worker effect.

2.3 Limitations of the consumption insurance model

All studies discussed so far have taken a unitary view on household decision making process. They assumed that the household has a unique joint utility function, that household choices are made in the best interest of everyone, and that resources are pooled entirely to guarantee a maximum amount of consumption insurance. An alternative view to the household decision making process is that of a model where each household member has their own preferences about consumption and labor supply and household decisions are the consequence of some form of agreement. There are multiple approaches to formulating such household decision models:

- non-cooperative behavior, in which the ultimate outcome for each household member is determined by Nash equilibrium;
- bargaining, in which each household member has an outside option and partners bargain about the surplus of cooperation;
- collective models, which assume that all family outcomes are Pareto optimal, but non-labor income is shared between partners according to some sharing rule that itself depends on individual characteristics.

Chiappori and Mazzocco (2017) provide an excellent survey of different modeling assumptions, their testable implications and their validity when cross-checked with household level data. What all these approaches have in common and what separates them from the unitary model is that the allocation within the household – in the simplest case consumption and labor supply of each spouse – typically depend on additional characteristics or “individual decisions powers”. The decision power of an individual can then typically be proxied by some observable variables like individual wages, human capital, age, or more broadly by so-called distribution factors.

Returning to the issue of consumption insurance, the non-unitary household world offers a particular issue that might seriously impede consumption risk sharing between household members: the issue of commitment. In the ideal world, partners were to make state contingent plans about consumption and labor supply allocations for each potential state of the world they might experience in future periods. When there is perfect commitment, such a state contingent plan would entail perfect consumption insurance against all components of risk that are unpredictable at the point in time at which the plan is made. If, however, there is limited commitment, then each household member might ask themselves at each point in time and each state of the world, whether it is still worth accepting an allocation

that was negotiated in the past or whether it would be better to leave the household and choose the best possible outside option. The mere threat of leaving the household might then lead to a renegotiation of the allocation between partners, or might even lead to the dissolution of a household.³

The literature on limited commitment in intra-family consumption and labor supply choices is still thin. Mazzocco (2007) was among the first to formalize and test commitment issues in intertemporal household decision making. In particular, he writes down two models of dynamic household decisions: First, a model of full commitment in which household choices are always on the ex-ante Pareto frontier. Second, a model incorporating the assumption that households can't fully commit to plans and might therefore renegotiate allocations at later points in time when new information about important economic variables has arrived. He uses data from the CEX to show that the full-commitment (or unitary) model is strongly rejected by the data and that intertemporal household behavior should better be described by limited commitment models.

Lise and Yamada (2019) use detailed panel data from Japan to study the allocation of resources within the household both in the cross section and within households over time. Their dataset comprises information on the private consumption expenditures of husbands and wives as well as expenditure for the entire household. In addition, they can observe hours of market work, home production and leisure for each household member, some measures of market wages and other household characteristics. This detailed information on intra-household allocations allows the authors to document a series of striking facts. They find that the majority of consumption expenditure is made for the entire household. Only 21 percent of expenditure is purely private, and out of these private spendings women consume on average about 30 percent. Wives contribute about 30 percent of market hours and 86 percent of home production hours. However, Lise and Yamada (2019) also document a substantial amount of cross-sectional heterogeneity. To this end, they construct a dynamic collective household choice model, in which both wife and husband have their own preferences over private consumption, a home produced public good and leisure. Allocations within the household are determined by some Pareto weights on the individual utility functions, which can depend on current household characteristics but also forecastable future elements, like expected wage growth. These weights can be revised by the couple upon the arrival of new information, like a persistent wage shock to some partner, reflecting the idea of Mazzocco (2007) that planned consumption allocations can be due to commitment problems.

The empirical findings from their estimated model are consistent with the idea

³Note that, in economic theory, it is not easy to describe situations in which the household would actually break up. As long as there is some surplus to be shared between partners and as long as the costs of renegotiating plans are not extremely high, there is typically an allocation that makes both partners better off as compared to the outside option. One avenue to go in order to allow for endogenous dissolution of households would be to consider information asymmetries as in Friedberg and Stern (2014).

that the intra-family distribution of resources is achieved by some form of bargaining or collective decision mechanism in which the two partners have different bargaining or decision powers. In particular, they document several facts: First, the allocation of resources at the time of marriage reflects expectations about future wage growth or career perspectives. Wives who either have a higher wage or a high expected wage growth receive a larger Pareto weight within the intra-family allocation. Furthermore, changes in wages can trigger changes in Pareto weights, reflecting the idea of an increased decision power for the partner who receives a positive wage shock. Weights are, however, only revised upon the arrival of major shocks. Small wage shocks typically leave the intra-household resource allocations unchanged. This points to a true limited commitment story as an obstacle to consumption insurance within the family, in the sense that substantial changes in the outside option of one partner can affect the amount of private consumption each partner in a marriage can realize. Last but not least, utility weights are in general more favorable to men, and especially so when the wife is not in employment. In the latter case, weights are also more stable over time.

Summing up this section, commitment issues can hinder consumption insurance within the family. There are other studies that point to commitment problems in intra-family decision making, like for example Rasul (2008) or Doepke and Kindermann (2019) in the context of fertility decisions and Voena (2015) who points to the importance of divorce law and therefore threat points in shaping intra-family allocations. The observation that commitment issues can give rise to imperfect consumption insurance is not new, see for example Kocherlakota (1996) for a theoretical analysis. Yet, especially in the intra-family context, the empirical evidence is still scarce. This is mostly due to a lack of detailed data on private consumption expenditures in the family. In addition, future work might also shed more light on the role of information asymmetries in shaping family outcomes, like for example in Friedberg and Stern (2014).

3 Families and aggregate fluctuations

In the previous section, we pointed to the role of the modern two-earner family in insuring idiosyncratic wage, employment and consumption risk. However, the economic consequences of the added worker effect and intra-family insurance possibilities are much broader. One might suspect, for example, that families also play a role in mitigating business cycle risks and that they can potentially shape the business cycle itself. This notion is not only supported by the evidence on risk-sharing possibilities, but can also be grounded in the fact that women often work in jobs that are less exposed to cyclical wage and unemployment risk than men, see for example Hoynes et al. (2012), Doepke and Tertilt (2016), and Vandenbroucke and Zhu (2018).⁴

The literature in macroeconomics and finance has used countercyclical variations

⁴One exception from this rule is the recent Covid-19 recession, in which women were disproportionately affected by job loss, see Alon et al. (2021).

in labor income risk as driver for business cycle models for a while. However, empirical investigations of the nature of such variations over the cycle has remained scarce, mostly owing to data limitations. Using very rich data from the U.S. Social Security Administration, Guvenen et al. (2014) have pointed to the fact that the *variance* of idiosyncratic income shocks is actually *flat* over the cycle. What is strongly *procyclical*, however, is the *skewness* of shocks. In recessions, workers more frequently experience drops in (real) wages or even extended periods of unemployment. Positive shocks, on the other hand, like a salary raise within the firm or a change to a better-paid job become inherently less likely. This asymmetry pulls down the entire wage distribution in recessions leading income shocks to be more left-skewed.

In the context of family insurance, the evidence of whether two-earner families are able to better cope with variations in the skewness of shocks over the business cycle is, so far, inconclusive. Busch et al. (2021) use panel data from the U.S., Germany, Sweden and France to investigate insurance possibilities against left-skewed income shocks in recessions. While they find some role for the tax and transfer system in ensuring large negative income shocks, they also report that within-family reallocations in labor supply are not effective in mitigating fluctuations in skewness over the cycle. Yet, Pruitt and Turner (2020) come to a different conclusion. They study administrative data of millions of households from the Internal Revenue Service of the United States and find that intra-household insurance actually is a powerful instrument in dealing with much of the risk facing primary earners. In their dataset, which comprises the years 2000 to 2014, they document both a procyclical skewness and also a countercyclical variance of labor income shocks for male earners. Moving from individual to household earnings, they find that the addition of a second earner reduced the risk a household faces by a substantial amount. This means that both the cyclicity of the variance and the skewness is less pronounced for two-earner households and that such households face less tail risk in recessions. Using a very stylized household model with a risk aversion parameter of 1.25, they calculate that the certainty equivalent earnings of two-earner households are 19 percent higher than that of males alone.

Mankart and Oikonomou (2017) provide a more structural approach to investigating family risk-sharing possibilities over the business cycle. Their analysis starts from the stylized fact that unemployment usually spikes in recessions and, hence, exhibits countercyclical behavior. Aggregate labor force participation, however, is entirely flat over the cycle. A standard search and matching model is not consistent with such an observation. If the economy enters a recession where more households get laid off and the job-finding probability falls, households that have the option of buffering the adverse effects from the recession with private wealth would withdraw entirely from the labor force, as searching for a job becomes less attractive. Those households would return to the market once the economy recovers and the chances of finding a job have increased. Hence, participation would be strongly procyclical. However, Mankart and Oikonomou (2017) argue that in a model with two-earner families, there is a counteracting effect to this argument.

When the economy enters a recession and the likelihood for the primary earner to face a job loss increases, it becomes more attractive for the secondary earner to search for a job in order to insure the household against potential future income losses. This increased search effort and therefore labor force participation of secondary earners – still predominantly women – offsets the procyclical participation behavior of men and renders the aggregate participation rate acyclical. The data support this view of the labor market, as the authors report a countercyclical participation rate but a procyclical employment rate for married women.

Using a novel methodology to calculate the added worker effect and studying data from the CPS, Guner et al. (2020) provide additional evidence for the aforementioned arguments. They find that, when shutting down any considerations regarding family insurance against idiosyncratic income shocks, women’s employment would look much more like that of men, exhibiting a strong negative skewness in recessions. Hence, when secondary earners would search under the same economic trade-offs as primary earners, their participation would be procyclical and employment strongly countercyclical as well. The considerations underlying the added worker effect, however, are fundamentally different as already argued by Mankart and Oikonomou (2017). As a result, women are exposed to less cyclical unemployment and wage risk. Birinci (2021), on the other hand, comes to a different conclusion. He finds only small spousal earnings responses to a job displacement of the family head in the PSID and argues, similar to Cullen and Gruber (2000), that generous unemployment benefits crowd out family insurance. As those benefits are most generous during recessions, it is not surprising that he finds the spousal earnings response to be even lower in economic downturns, when family insurance might otherwise be most valuable.

The increase in labor force participation of women over the course of the last century certainly had an impact on aggregate economic performance as well. Albanesi (2019) documents that the growth in women’s labor supply is responsible for a substantial part of TFP growth in the U.S. in the 1980s. In addition, the fact that female employment typically correlates less with the business cycle has contributed to a decline in the cyclicalities of aggregate labor hours in the 1980s, the so-called the great moderation. Finally, taking a deeper look at recessions, the strong growth in female employment over the course of the 20th century has led employment recoveries to be extremely fast until around 1990. From that time onwards, female labor force participation suddenly stagnates and aggregate female employment behaves much more like that of men, meaning that employment takes an extended period of time after a recession to recover. The result are so-called jobless recoveries, meaning a slow recovery of the labor market even when GDP has returned to trend already after a recession. Olsson (2020) and Fukui et al. (2021) confirm this view. The latter also provide a sufficient statistics approach to measuring the impact of female labor force participation on aggregate employment. Finally, Bardoczy (2020) argues that the added worker effect can act as an automatic stabilizer in recessions. When families provide consumption insurance through strategic labor supply behavior over the cycle, consumption expenditure

declines less in a recession as compared to a situation with only single-earner households. As a result, aggregate demand exhibits less cyclical which calms the consequences of shocks to aggregate economic performance.

When family insurance is important for households, one might suspect that recessions also have an impact on the formation and dissolution of families. In fact, Schaller (2013) as well as Hellerstein et al. (2013) provide supportive evidence for this idea by showing that divorce is procyclical. This is consistent with the idea that the benefits of being a two-earner family may increase in recessions. Consequently, partners who are on the edge of breaking up will stay together for a bit longer in order to overcome the economic burdens of a downturn. The opposite is, however, true for the marriage rate which is procyclical as well, see Schaller (2013) and Bellido and Marcén (2021). Finally, Dyrda et al. (2012) show that during recessions, individuals tend to live in larger households. This is mostly driven by young individuals who either move back to their parents or do not leave the household as they would otherwise do. Such an extended family risk sharing mechanism can have an impact on the labor supply response of young individuals to business cycle shocks. As a result, the macroeconomic Frisch elasticity of labor supply with respect to wage changes increases.

4 Social insurance and family insurance

In the previous sections, we extensively discussed the role of families and particularly family labor supply in insuring individual risk. However, family insurance mechanisms are only partial, as perfect insurance exploits the law of large numbers. This is obviously impossible for two-earner households. In addition, families are not the only insurance mechanism in an economy. Governments usually provide some form of social safety net or run public insurance systems and private insurance firms exist as well. This section discusses the interplay between insurance provided by the government and/or the market and family structure. The interplay between public and family insurance has already been sketched in previous sections, where we argued that public unemployment insurance can crowd out the added worker effect. Here, we particularly focus on a dynamic life-cycle perspective, where families either consist of two-earner households or of parents and children who pool risks and self-insure via savings. In the first part, we study the impact of family-oriented welfare and social security regulations on the labor supply of married and single women. Next, we focus on informal care provision of families as an insurance device against various health risks before we finally analyze portfolio choice and life insurance demand of families.

4.1 Social security and spousal labor supply

Standard life-cycle models of labor supply and savings typically assume a lump-sum pension benefit at retirement and only model survival risk at the household level. Consequently, such models entirely neglect widows and widowers. Furthermore, in most countries public pension benefits are related to former earnings and include so-called “auxiliary benefits” for surviving and/or living spouses. For ex-

ample, the U.S. social security system allows married individuals to collect either one's own pension benefits or half of the spouse's entitlement. When a spouse dies, the survivor receives either the own or the deceased spouse's entitlement. Since married women typically are secondary earners and tend to outlive their husbands, these provisions can have a significant effect on married female labor supply during working years.⁵

A good starting point for the analysis of such social security regulations is Kaygusuz (2015). In his quantitative simulation model, men and women enter economic life either as a married couple or as singles and they remain in this status until death. Couples are matched on the basis of five different educational categories (i.e. 25 potential combinations) and they decide about consumption, labor supply and participation of the female spouse. Eliminating spousal and survivor benefits increases the labor force participation of married women by 4.7 percent and aggregate labor supply by 0.8 percent in the long run. Married households also increase their savings, leading to a higher capital stock and output. While single-earner couples are predominantly hurt by the reform, a majority of households experiences welfare gains.

This result provides a useful benchmark for further discussion. However, since Kaygusuz (2015) abstracts from income risk during the working phase, he does not capture the interplay between family insurance and social insurance. This issue is taken up by all following studies. Sanchez-Marcos and Bethencourt (2018) pay particular attention to endogenous human capital formation, a variable family size due to the presence of children as well as detailed eligibility rules and early retirement provisions for spousal and survivor benefits. The subsequent removal of these benefits increases the female employment rate on average by 4 and 10 percentage points, respectively. Owing to self-insurance possibilities within the family, the economy-wide consumption inequality increases only slightly. But this last result mainly stems from the long-run perspective taken by Sanchez-Marcos and Bethencourt (2018). Nishiyama (2019) simulates the transition path resulting from a reform of pensions benefits rules. He shows that the elimination of auxiliary benefits hurts the majority of married households in the short run, because they can only partially adjust their self-insurance provision to a sudden social security reform.

The studies discussed so far assumed stable marriages and abstracted from both divorce risk and the timing of marriage. Recently this gap between modeling structure and reality has been filled by allowing for marital transitions over the life cycle. Groneck and Wallenius (2020) consider a heterogeneous household structure by including a labor supply choice between part-time and full-time work and endogenous human capital formation of women, a male labor supply choice at the retirement margin and a socio-economic gradient to survival risk. The length of a marriage determines the eligibility for auxiliary benefits and the joint retirement of spouses creates a utility surplus. Household decisions are assumed to follow a

⁵They also have distributional consequences, but we will not discuss this any further.

collective household model, so that the allocation of resources between husbands and wives depends on welfare weights that in turn relate to the relative earnings potential of partners. The model is calibrated to match the employment rates of married women over the life cycle and the retirement decision of married men for three educational groups of U.S. cohorts born 1950-54. Eliminating auxiliary benefits induces earlier retirement of married men and increases labor supply of married women by more than 12 percent, the strongest effect in all studies considered so far. All in all, aggregate labor hours increase by roughly 2 percentage points. Finally, Borella et al. (2019) find even stronger employment effects in their structural life-cycle model that also considers marriage-related taxes, health risk during retirement and labor supply at the extensive and intensive margin.

While the distortions of a couples' labor supply decision implied by survivor benefits clearly deteriorate economic efficiency, a complete welfare evaluation of social security also has to take into account the provision of longevity insurance, which could differ substantially between singles and couples. For this reason, Fehr et al. (2017) quantify the efficiency effects of an elimination of social security for different household types. They focus on the German case. The German pension system, unlike that of the U.S., is less generous for spouses, hardly distorts family labor supply and mainly insures longevity risk. Table 1 reports the simulated efficiency effects in a series of models with different family structures. When there

Table 1: Efficiency effects of eliminating pensions with different family structures*

	Overall	Single	Married
No families	-1.20		
Constant status	-0.88	-1.34	-0.35
Martial Risk	-0.69		

Source: Fehr et al. (2017).

are only single households (“No families”), eliminating social security comes at a substantial loss of economic efficiency of about 1.2 percent of initial equilibrium resources. This efficiency loss mainly reflects the value of longevity insurance provided by the German pension system. If instead households enter the economy as either singles or married couples and this state remains unchanged for a household’s entire life, the overall efficiency loss of shutting down the public pension system is much smaller at around 0.9 percent. The reason for this becomes clear when we measure the efficiency effect separately by family type. While singles still experience substantial welfare losses, married couples value the longevity insurance from social security much less. The possibility to insure longevity risk on a private basis within the family causes the efficiency loss for married couples to amount to only 0.35 percent, where singles lose an additional 1 percent. Finally, in an economy with unstable families, meaning with divorce risk and remarriage, the overall efficiency effect roughly equals that of the economy with stable family

structures. One can conclude from this entire analysis that, when the numbers of singles steadily increases as it is the case in many Western societies, insurance provision by social security becomes more important.

This last point is in some way confirmed by Haan and Prowse (2020), who quantify the optimal policy mix for the German unemployment insurance and social assistance system in a structurally estimated life-cycle model. In the benchmark model that includes both singles and families, the optimal replacement rate for the level of social assistance is close to the status quo in Germany. In a model with only single household – i.e. where marriage probabilities are set to zero – the optimal social assistance level would be roughly 66 percent higher, which mainly reflects the missing family insurance from income pooling of couples. Quite surprisingly, Haan and Prowse (2020) find no role for unemployment benefits in Germany in either case. However, this may be related to the assumptions about the wage process. De Nardi et al. (2020) apply a quite similar approach to analyzing the optimal structure of welfare benefits in the UK. Comparing in-work benefits (such as the Working Tax Credit) and income support programs, they show that the calibration of the wage processes for men and women dramatically affects the optimal benefit structure. While the canonical wage process favors in-work benefits relative to income support, a more carefully estimated wage process generates a higher optimal income floor.

4.2 Health risk and informal care

Since the bulk of longevity risk is covered by old-age pension systems, health risks in the form of health shocks, disability or long-term care and the associated medical costs are considered the most important sources of risk for the elderly. Hence, one can expect health risk to have a significant impact on precautionary savings behavior. Traditionally, self-insurance against these types of risks has been analyzed in models of individual households, where the role of informal assistance by the family is completely neglected, see De Nardi et al. (2010). Only recently, simulation models have been developed that allow to quantify the central role the family can take in insuring against significant health shocks. Dobrescu (2015) estimates a structural life-cycle model in which the elderly are exposed to health shocks (not covered by public insurance) for three European regions. Individuals can either insure formally by purchasing insurance on the market or informally via transfers from spouses and children. When formal care provision is weak and social cohesion is strong, households keep inheritable wealth in order to trigger descendants to provide informal care in late life. Interestingly, the study finds a positive association between social cohesion and life expectancy.

While Dobrescu (2015) analyzes the role of health shocks and informal care for retired households, families are equally important in caring for individuals who experience health shocks already during working life. As younger households have less time to accumulate assets for self-insurance, family insurance in the form of spousal labor supply or help with childcare might become even more important to cover such risks. Ball and Low (2014) analyze consumption expenditure data of

men and women aged 25 to 60 from the British Household Panel Survey (BHPS) between the years 1991 and 2004. Individuals are categorized according to their work limitations. They find a decline in food expenditure of about 8 to 11 percent for individuals who receive disability insurance (DI) benefits, but have no other self-insurance available. The presence of a work-active spouse mitigates this number by 3.3 percentage points, indicating a significant degree of family insurance. The importance of spousal labor supply as an insurance mechanism against disability risk is also highlighted by Autor et al. (2019) who estimate a structural life-cycle model using detailed Norwegian register data. Their aim is to quantify the willingness to pay for DI for different household types. As it turns out, families have a significantly lower willingness to pay than singles and this difference almost completely disappears when spousal labor supply is not allowed to adjust to disability shocks. Furthermore, self-insurance via savings and borrowing only has a minor effect on the willingness to pay for DI. When the spouse has died already, family assistance for DI recipients could be provided by children. Surprisingly, Rennane (2020) shows that child support acts as a complement (and not a substitute) for income provided by DI in the U.S. Since recipients of higher DI benefits can compensate their children for their help, it is easier for these children to reduce their market labor supply.

In contrast to disability risk, long-term care risk typically materializes at older ages, so that family insurance is mainly provided in the form of caregiving from children. Informal care is then clearly a substitute to formal care.⁶ In the U.S., the latter is usually provided by medicaid and private nursing homes, but according to Barczyk and Kredler (2018) this only accounts for about one-third of total care hours. The remaining hours are provided informally. They calibrate a life-cycle model in which parents and children make non-cooperative savings decisions and bargain on intra-family transfers in the form of money and time (care). Their model features the main publicly provided care arrangements in the U.S. and replicates realistic shares of formal and informal care hours. Implementing formal and informal care subsidies – similar to those paid in Germany – generates large welfare gains, even when the expansion of subsidies is combined with a reduction in the size of the Medicaid program. The reason is an improved targeting of benefits. Informal care subsidies mainly appeal to low productivity children who then leave the labor force. High productivity children exploit the (higher) formal care subsidy and stay on the labor market.

Long-term care risk is also an important motivation for self-insurance in other countries. Imrohoroglu and Zhao (2018) develop an overlapping generations model in which parents and children form a household and derive utility from two-sided altruism. They strive to explain the increase in the Chinese savings rate since 1980. While the risk facing parents in old age were usually absorbed by their children in the past, the strictly enforced one-child policy has deteriorated this insurance channel and induced families to increase savings instead. This explains

⁶Klimaviciute and Pestieau (2018) provide a survey of the theoretical literature on formal and informal care and the role of the family.

the long-term upward trend in the Chinese savings rate. Short term fluctuations, on the other hand, are explained by changes in productivity growth.

Finally, Braun et al. (2017) do not directly address the implications of family insurance, but clearly document the benefits of public insurance in a family model. Their overlapping generation model of the U.S. economy distinguishes between singles and married households, allows for a wide range of income and health risks and accounts for the U.S. social security and means-tested social insurance programs. Eliminating social insurance generates a significant ex-ante welfare loss that varies considerable across household types.

4.3 Marital status, portfolio choice and life insurance demand

Kotlikoff and Spivak (1981) as well as Brown and Poterba (2000) already pointed to the fact that families can provide a similar degree of longevity insurance as annuity contracts. In the same way, market provided life insurance could substitute survivor benefits of social security, see Hong and Ríos-Rull (2012). One would therefore expect that both family insurance as well as the presence of private markets for annuities and life insurances would erode the insurance value of social security. In order to quantify the value of public insurance provision, Hong and Ríos-Rull (2007) develop an overlapping generation model in which households face marriage and divorce risks over the life cycle. The insurance value of social security is then isolated by comparing the long-run welfare effects of introducing social security (with survivor benefits) into economies with and without private insurance markets. As it turns out, the long-run welfare effects of social security are hardly affected by the market setting. Hong and Ríos-Rull (2007) therefore conclude that it provides only little longevity and survivor's risk insurance over and above what the market and family insurance can achieve.

In order to shed more light on this issue, recent studies provide more realistic models of household savings choices with different insurance products and investment risk in old age. A good starting point for such an analysis is Hubener et al. (2014), who completely abstract from social security and consider a retired couple facing an uncertain lifespan that chooses among risk-free bonds, risky stocks, term life insurances and (single or joint) annuities to manage retirement income. Households have a bequest motive and private insurance companies charge loadings which may, amongst others, result from asymmetric mortality beliefs between insurer and insured. If a household only lives from liquid wealth, there is no need to buy a life insurance and roughly 20 percent of financial wealth are invested in stocks. However, if the household's endowment also includes retirement income, the husband will buy life insurance in order to counteract the reduction in income of the wife in case of his death. Interestingly, the presence of children hardly affects the demand for life insurance. This indicates that the motivation for buying such insurance mainly comes from income provision and is not related to the bequest motive. Furthermore, under the presence of retirement income, the couple holds a much larger share of their financial wealth in stocks. Finally, instead of buying joint annuities early in retirement, the couple prefers to wait until one spouse dies

and the surviving spouse can purchase a single annuity at cheaper prices.

Whereas Hubener et al. (2014) only look at couples upon their entry into retirement, Wang (2019) jointly studies life insurance demand and labor supply over the entire life cycle. In a model that features realistic wage shocks for men and women, she shows that the gender gap in life insurance demand between men and women observed in the data is mainly driven by the gender income gap. She also highlights the role of children in household choices and the risk sharing effects of income growth and income correlation within the household.

While focusing on life insurance demand, Wang (2019) abstracts from risky asset choices and changes in family status due to marriages and divorce. The interaction between marital risk and portfolio choice is analyzed in Love (2010), who allows for uncertain medical expenses during retirement. Divorce induces men and women to adjust their portfolios differently. While men increase the riskiness of their portfolios, women respond by choosing a safer asset allocation. Remarriages after a divorce then trigger portfolio changes in the opposite direction. Similarly, the death of a spouse implies a move towards a safer asset allocation as well, but this effect is much more pronounced among women than among men. In all cases, these adjustments can be explained by the loss (or increase) in family insurance.

If at all, the studies by Hubener et al. (2014), Wang (2019) and Love (2010) only allow for very rudimentary pension benefit at retirement. Hence, they do not discuss the interaction between the pension system and individual portfolio choice. These issues are taken up by Hubener et al. (2016) and Li (2018), who analyze optimal financial behavior in sophisticated models with labor market, marital and mortality risk as well as an advanced public pension system. The former focus on the optimal benefit-claiming option, which offers an alternative to altering the (financial) portfolio structure for balancing economic shocks. The latter provides a general equilibrium model in which social security provides survivor benefits to the spouse and to dependent children. There is a negative mortality-income gradient and adverse selection in the life insurance market. The counterfactual experiments analyze the reduction of survivor benefits under alternative specifications of private insurance pricing rules. As it turns out, the reduction of survival benefits for dependent children can be welfare reducing owing to the inefficiency of the private insurance market.

5 Welfare system and family formation

The previous section has discussed various channels through which families provide insurance to their members and how this insurance provision interacts with the public insurance system. Up to this point, however, we focused on existing family structures abstracting from explicit marriage and divorce decisions. Besides love and affection, these decisions yet often depend on a number of economic considerations.⁷ If household risk sharing is important and acts as a substitute to

⁷Hess (2004) provides an analysis how the presence of love interacts with risk sharing in the decision to get married and divorced.

public insurance, we would expect a link between public welfare systems and the formation and stability of families.

The link between the public safety net and family structures is sometimes very direct and obvious. Guner and Knowles (2009) compare U.S. and Canadian welfare policies within a model of household formation and dissolution, endogenous fertility and human capital investments in children. While U.S. programs such as the Aid to Families with Dependent Children (AFDC) penalize women for marriage and rewarded them for out of wedlock fertility, Canadian welfare programs were more generous and less biased. Guner and Knowles (2009) show that these differences in policies could (at least partly) explain a higher single-parenthood and marital instability in the U.S. compared to Canada.

In 1996 the AFDC program in the U.S. was substituted by the Temporary Assistance to Needy Families (TANF) program which restricted public transfers for families to a maximum of 5 years. This reform impaired mainly single mothers with low earnings potential who had no support from a husband. Low et al. (2020) document that welfare utilization declined dramatically for this group, although immediate eligibility hardly changed. While benefits were being “banked”, the employment of single women increased and the divorce rate fell. Simulating this reform in a structural life-cycle model with endogenous marriage and divorce, they find that the anticipation of benefit exhaustion led to higher marital stability and raised the employment rate of single mothers. The simulations also show that the decline in welfare use of single mothers would have been significantly more severe, if there was no possibility of marriage, and vice versa for married women. This highlights the importance of family insurance as a substitute for public insurance in the U.S. Overall, however, it is not clear, whether the current U.S. tax and transfer system impedes or encourages marriage. A recent study by Ortigueira and Siassi (2021) argues that some anti-poverty provisions still induce single mothers to cohabit rather than to marry their partner.

Outside of the U.S., Persson (2020) analyzes the elimination of survivor insurance from future marriage contracts in Sweden in 1989. As in Guner and Knowles (2009) and Low et al. (2020), she highlights the role of expectations of future benefits for the marriage decision. Since the reform introduced special transitory provisions for non-married couples, Persson (2020) identifies a marriage boom right after the reform, which resulted in a lower match quality and finally increased the divorce rate. In a long-run perspective, the reform affected marriage formation decades before expected payouts and raised the degree of assortative mating in the marriage market.

Finally, Schulz and Siuda (2020) exploit a reform of the German unemployment insurance in 2003 to document the role of within-household insurance for marital behavior. With this reform, Germany tightened the means testing of unemployment assistance against a partner’s income. Consequently, new marriages with a partner who has a high unemployment risk became less attractive. Schulz and Siuda (2020) first document a positive correlation between foreign nationality and

unemployment risk. They then show that the reform resulted in fewer but more stable inter-ethnic marriages. The latter effect is due to positive selection.

Summing up, there is clear evidence that tying social insurance provision to marriage – either directly or indirectly – can have unintended but far reaching economic consequences through behavioral changes in the marriage market. Welfare reforms that reduce marriage rates and/or increase divorce probabilities may ultimately result in lower welfare owing to the counteracting effects on family insurance provision.

6 Alternative view: Families as a source of risk

In the previous sections, we pointed to the family as provider of insurance against various sources of risk. However, the formation and dissolution of a family can be interpreted as a source of risk as well. When a suitable partner arrives, the resources of two singles suddenly need to be shared. On the contrary, divorce leads to a division of formerly joint wealth and potentially of claims to social insurance systems. In reality, of course, changes in marital status are not entirely exogenous but result from individual choices. However, for a single to find a suitable partner or for married persons to experience new outside options still entails a serious amount of exogeneity. Hence, as a first-order approximation, treating family formation and dissolution as shocks can deliver valuable insights.

Cubeddu and Ríos-Rull (2003) were among the first to model changes in family status as a stochastic process. Their work focuses on the impact of changes in household structure on aggregate savings in different settings. They find that marital risk has a significant impact on household savings. However, the exact direction and quantity fundamentally depends on the decision process within the household, divorce rules and (re-)marriage patterns after divorce. Table 1 derived from Fehr et al. (2017) showed that incorporating marital risk into the household decision framework does not necessarily alter the insurance value of families. This might in part be owing to the fact that family members take such risks into account and self-insure using reallocations of labor supply and savings.

Low marriage rates and unstable family structures are a phenomenon of modern times. Between the 1970s and the 2000s, there has been a significant decline of crude marriage rates and a rise in divorce rates in almost all European countries.⁸ At the same time, female labor supply strongly increased and particularly so for married women. Driven by this observation, Fernandez and Wong (2014) compare the behavior of two U.S. cohorts born in 1935 and 1955 using a dynamic life-cycle model with labor market, longevity and marital risk (in the spirit of Cubeddu and Ríos-Rull, 2003) under incomplete markets. Their model accounts for the changes in assortative mating, in the age at first marriage, in fertility patterns, in marriage and divorce probabilities and for the decline in the gender wage gap. They find the increase in divorce risk to have the largest impact on married

⁸See for example Fehr et al. (2016). Stevenson and Wolfers (2007) provide a discussion of the driving forces of this development especially in the U.S.

women's labor force participation. Fehr et al. (2016) come to a slightly different conclusion. They simulate a rich overlapping generations model calibrated to match the changes in household structures in Germany. When the probability of marriage declines, singles build up more precautionary savings and increase their labor supply especially early in life. A higher divorce risk has a similar effect for married couples, but at a much lower scale. As a result and similar to Heathcote et al. (2010), the decline in the gender wage gap plays a more important role than in Fernandez and Wong (2014).

Of course, the divorce risk of married women can be insured at a much earlier stage. Since divorced women are typically worse off than their spouses owing to a lower earnings capacity, they may self-insure by investing in their human capital early on. Guvenen and Rendall (2015) study the interplay between education decisions early in economic life, marriage and divorce decisions throughout the life cycle as well as the time allocation between market work and home work in an overlapping generation economy. They start from a situation in which a consent divorce law regime determines education, marital and labor supply behavior. They then show that the expectation of a reform towards unilateral divorce (and the corresponding increase in divorce probabilities) induces currently unmarried women to invest more in their education. Changes in divorce law and their impact on self-insurance and family insurance are also at the heart of Voena (2015). In addition to consensual and unilateral divorce decisions, she also considers different rules for the division of property after divorce. Not surprisingly, she confirms that changes in divorce risk affect a couples' savings and labor supply behavior. In addition, she shows that risk-sharing in the family is higher under mutual consent divorce law as compared to unilateral divorce law (because of a higher divorce risk). Alternative financial arrangements after divorce hardly affect risk-sharing under mutual consent, but may play a role under unilateral law.

7 Summary

This chapter pointed to the various dimensions of individual and aggregate risk that can be insured through the family. Family insurance arrangements are of course imperfect. Still, they may be preferred by individuals to avoid information barriers or high loading factors in private insurance markets or the redistributive features embedded in some social insurance schemes. In turn, the expansion of social insurance through fiscal systems has eroded some of the gains of marriage. The data support the idea that individuals understand and value the insurance provision of the family, and that family insurance influences marital decisions at least to some degree. The transition in family structures from early and stable marriages towards less family stability over the course of the last century, on the other hand, clearly affected economic behavior by inducing individuals to seek other options for self-insurance.

Despite these valuable insights from the studies cited in this survey, there are still many issues in the context of family insurance that are not yet fully understood.

What is an optimal design of social insurance policies when marital decisions endogenously respond to fiscal policy design? How much redistribution does social insurance imply between singles and married couples when the latter can provide partial insurance by themselves? And is this redistribution intended? Should insurance policies (be they private or public) differentiate much more across family types? Are information problems really small in couples, and if not, how do they impede family risk sharing? What are means for partners to commit to future risk-sharing options in a world where family structures are less stable? Do social norms against female employment hold back family insurance in some countries? And how does family insurance (or the lack thereof) impact the business cycle in the future? Does the ongoing emancipation in the modern family lead to different aggregate outcomes when men and women form expectations differently? These and many more questions could and should be addressed in future research.

Of course, there are various other aspects related to family decision making that we did not capture in this chapter, but that are of equal importance for economists. Doepke et al. (2022) provide an extensive survey of current research on fertility decisions in the modern family. Inspired by the “career and family” view of Goldin (2020), they show that a balance of power and an equal sharing of the cost of having children between partners is a prerequisite for high fertility rates in modern economies. Bau and Fernández (2022) point to the role of culture in shaping family structures and family institutions all across the world. Baudin et al. (2021) argue that economists mostly study the so-called “nuclear family”, but that other types of family arrangements exist and that they may deliver different economic outcomes. Studying different family types and their economic behavior might therefore be important for understanding development processes or public policies. The mere fact that all of these recent surveys pay so much attention to the various roles of the family and to the economics of intra-family decision making demonstrates the importance of this topic in modern economics and should encourage researchers to study the economics of the family.

8 Bibliography

- Albanesi, S. (2019). Changing business cycles: The role of women’s employment. Working Paper 25655, National Bureau of Economic Research.
- Alon, T., S. Coskun, M. Doepke, D. Koll, and M. Tertilt (2021). From mancession to shecession: Women’s employment in regular and pandemic recessions. In M. S. Eichenbaum and E. Hurst (Eds.), *NBER Macroeconomics Annual 2021, volume 36*. University of Chicago Press, forthcoming.
- Attanasio, O., H. Low, and V. Sánchez-Marcos (2005). Female labor supply as insurance against idiosyncratic risk. *Journal of the European Economic Association* 3(2-3), 755–764.
- Autor, D., A. Kostol, M. Mogstad, and B. Setzler (2019). Disability benefits, consumption insurance, and household labor supply. *American Economic Review* 109(7), 2613–2654.
- Ball, S. and H. Low (2014). Do self-insurance and disability insurance prevent consumption loss on disability? *Economica* 81, 468–490.
- Barczyk, D. and M. Kredler (2018). Evaluating long-term-care options, taking the family seriously. *Review of Economic Studies* 85(2), 766–809.
- Bardoczy, B. (2020). Spousal insurance and the amplification of business cycles. Working paper, Northwestern University.
- Bau, N. and R. Fernández (2022). The family as a social institution. In S. Lundberg and A. Voena (Eds.), *Handbook of Family Economics*. North Holland, forthcoming.
- Baudin, T., B. De Rock, and P. Gobbi (2021). Economics and family structures. CEPR Discussion Papers 16516, Center for Economics Policy Research.
- Bellido, H. and M. Marcén (2021). Will you marry me? It depends (on the business cycle). *Review of Economics of the Household* 19(2), 551–579.
- Birinci, S. (2021). Spousal labor supply response to job displacement and implications for optimal transfers. Working paper, Federal Reserve Bank of St. Louis.
- Blossfeld, H.-P. (2009). Educational assortative marriage in comparative perspective. *Annual Review of Sociology* 35(1), 513–530.
- Blossfeld, H.-P. and A. Timm (2003). *Who marries whom? Educational systems as marriage markets in modern societies*, Volume 12 of *European Studies of Population*. Springer.

- Blundell, R., M. Graber, and M. Mogstad (2015). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics* 127, 58–73.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy* 126(S1), S73–S115.
- Boldrin, M., M. De Nardi, and L. E. Jones (2015). Fertility and social security. *Journal of Demographic Economics* 81(3), 261–299.
- Boldrin, M. and L. E. Jones (2002). Mortality, fertility, and saving in a Malthusian economy. *Review of Economic Dynamics* 5(4), 775–814.
- Borella, M., M. De Nardi, and F. Yang (2019). Are marriage-related taxes and social security benefits holding back female labor supply? Working Paper 26097, National Bureau of Economic Research.
- Braun, A. R., K. A. Kopecky, and T. Koreshkova (2017). Old, sick, alone, and poor: A welfare analysis of old-age social insurance programmes. *Review of Economic Studies* 84(2), 580–612.
- Bredemeier, C. and F. Juessen (2013). Assortative mating and female labor supply. *Journal of Labor Economics* 31(3), 603–631.
- Brown, J. R. and J. M. Poterba (2000). Joint life annuities and annuity demand by married couples. *Journal of Risk and Insurance* 67(4), 527–553.
- Busch, C., D. Domeij, F. Guvenen, and R. Madera (2021). Skewed idiosyncratic income risk over the business cycle: Sources and insurance. *American Economic Journal: Macroeconomics* forthcoming.
- Calvo, P. A., I. Lindenlaub, and A. Reynoso (2021). Marriage market and labor market sorting. Working Paper 28883, National Bureau of Economic Research.
- Chiappori, P.-A. and M. Mazzocco (2017). Static and intertemporal household decisions. *Journal of Economic Literature* 55(3), 985–1045.
- Choi, S. and A. Valladares-Esteban (2020). On households and unemployment insurance. *Quantitative Economics* 11(1), 437–469.
- Cubeddu, L. and J.-V. Ríos-Rull (2003). Families as shocks. *Journal of the European Economic Association* 1(2-3), 671–682.

- Cullen, J. B. and J. Gruber (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics* 18(3), 546–572.
- De Nardi, C., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy* 118(1), 39–75.
- De Nardi, M., G. Fella, M. Knoef, G. Paz-Pardo, and R. Van Ooijen (2021). Family and government insurance: Wage, earnings, and income risks in the Netherlands and the U.S. *Journal of Public Economics* 193, 104327.
- De Nardi, M., G. Fella, and G. Paz-Pardo (2020). Wage risk and government and spousal insurance. Working Paper 28294, National Bureau of Economic Research.
- Di Tella, R. and R. MacCulloch (2002). Informal family insurance and the design of the welfare state. *Economic Journal* 112, 481–503.
- Dobrescu, L. I. (2015). To love or to pay - savings and health care in older age. *Journal of Human Resources* 50(1), 254–299.
- Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2022). The economics of fertility: A new era. In S. Lundberg and A. Voena (Eds.), *Handbook of Family Economics*. North Holland, forthcoming.
- Doepke, M. and F. Kindermann (2019). Bargaining over babies: Theory, evidence, and policy implications. *American Economic Review* 109(9), 3264–3306.
- Doepke, M. and M. Tertilt (2016). Families in macroeconomics. In J. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics, Vol. 2*, pp. 1789–1889. North Holland.
- Dyrda, S., G. Kaplan, and J.-V. Ríos-Rull (2012). Business cycles and household formation: The micro vs the macro labor elasticity. Working Paper 17880, National Bureau of Economic Research.
- Fehr, H., M. Kallweit, and F. Kindermann (2016). Household formation, female labor supply, and savings. *Scandinavian Journal of Economics* 118(4), 868–911.
- Fehr, H., M. Kallweit, and F. Kindermann (2017). Families and social security. *European Economic Review* 91, 30–56.
- Fernandez, R. and J. C. Wong (2014). Divorce risk, wages and working wives: A quantitative life-cycle analysis of female labour force participation. *Economic Journal* 124(576), 319–358.
- Fernández, R., N. Guner, and J. Knowles (2005). Love and money: A theoretical and empirical analysis of household sorting and inequality. *Quarterly Journal of Economics* 120(1), 273–344.

- Friedberg, L. and S. Stern (2014). Marriage, divorce, and asymmetric information. *International Economic Review* 55(4), 1155–1199.
- Fukui, M., E. Nakamura, and J. Steinsson (2021). Women, wealth effects, and slow recoveries. *American Economic Journal: Macroeconomics* forthcoming.
- Galor, O. and D. N. Weil (1996). The gender gap, fertility, and growth. *American Economic Review* 86(3), 374–387.
- Goldin, C. (2020). Journey across a century of women – the 2020 Martin S. Feldstein lecture. *NBER Reporter* 3, 1–7.
- Gonalons-Pons, P. and C. R. Schwartz (2017). Trends in economic homogamy: Changes in assortative mating or the division of labor in marriage? *Demography* 54(3), 985–1005.
- Greenwood, J., N. Guner, G. Kocharkov, and C. Santos (2016). Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation. *American Economic Journal: Macroeconomics* 8(1), 1–41.
- Greenwood, J., N. Guner, and G. Vandenbroucke (2017). Family economics writ large. *Journal of Economic Literature* 55(4), 1346–1434.
- Greenwood, J., A. Seshadri, and M. Yorukoglu (2005). Engines of liberation. *Review of Economic Studies* 72(1), 109–133.
- Groneck, M. and J. Wallenius (2020). It sucks to be single! marital status and redistribution of social security. *Economic Journal* 131(633), 327–371.
- Guner, N. and J. Knowles (2009). Why is the rate of single-parenthood lower in Canada than in the U.S.? A dynamic equilibrium analysis of welfare policies. *Canadian Journal of Economics/Revue canadienne d'économique* 42(1), 56–89.
- Guner, N., Y. Kulikova, and A. Valladares Esteban (2020). Does the added worker effect matter? Discussion Paper 14346, Centre for Economic Policy Research.
- Güvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy* 122(3), 621–660.
- Güvenen, F. and M. Rendall (2015). Women's emancipation through education: A macroeconomic analysis. *Review of Economic Dynamics* 18(4), 931–956.
- Haan, P. and V. Prowse (2020). Optimal social assistance and unemployment insurance in a life-cycle model of family labor supply and savings. Working paper, DIW Berlin.
- Halla, M., J. Schmieder, and A. Weber (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics* 12(4), 253–87.

- Heathcote, J., K. Storesletten, and G. L. Violante (2010). The macroeconomic implications of rising wage inequality in the United States. *Journal of Political Economy* 118(4), 681–722.
- Heckman, J. J. and T. E. MaCurdy (1980). A life cycle model of female labour supply. *Review of Economic Studies* 47(1), 47–74.
- Hellerstein, J. K., M. S. Morrill, and B. Zou (2013). Business cycles and divorce: Evidence from microdata. *Economics Letters* 118(1), 68–70.
- Hess, G. D. (2004). Marriage and consumption insurance: What’s love got to do with it? *Journal of Political Economy* 112(2), 290–318.
- Hong, J. H. and J.-V. Ríos-Rull (2007). Social security, life insurance and annuities for families. *Journal of Monetary Economics* 54(1), 118–140.
- Hong, J. H. and J.-V. Ríos-Rull (2012). Life insurance and household consumption. *American Economic Review* 102(7), 3701–3730.
- Hoynes, H., D. L. Miller, and J. Schaller (2012). Who suffers during recessions? *Journal of Economic Perspectives* 26(3), 27–48.
- Hubener, A., R. Maurer, and O. S. Mitchell (2016). How family status and social security claiming options shape optimal life cycle portfolios. *Review of Financial Studies* 29(4), 937–978.
- Hubener, A., R. Maurer, and R. Rogalla (2014). Optimal portfolio choice with annuities and life insurance for retired couples. *Review of Finance* 18(1), 147–188.
- Illing, H., J. F. Schmieder, and S. Trenkle (2021). The gender gap in earnings losses after job displacement. Working Paper 29251, National Bureau of Economic Research.
- Imrohoroğlu, A. and K. Zhao (2018). The chinese saving rate: Long-term care risks, family insurance, and demographics. *Journal of Monetary Economics* 96(1), 33–52.
- Juhn, C. and S. M. Potter (2007). Is there still an added-worker effect? Staff Reports 310, Federal Reserve Bank of New York.
- Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2(4), 53–87.
- Kaygusuz, R. (2015). Social security and two-earner households. *Journal of Economic Dynamics and Control* 59, 163–178.
- Klimaviciute, J. and P. Pestieau (2018). The public economics of long-term care: A survey of recent contributions. *Annals of Public and Cooperative Economics* 89(1), 49–64.

- Kocherlakota, N. R. (1996). Implications of efficient risk sharing without commitment. *Review of Economic Studies* 63(4), 595–609.
- Kotlikoff, L. J. and A. Spivak (1981). The family as an incomplete annuities market. *Journal of Political Economy* 89(2), 372–391.
- Leroux, M. L. and P. Pestieau (2014). Social security and family support. *Canadian Journal of Economics* 47(1), 115–143.
- Li, Y. (2018). Economic analysis of social security survivor insurance. *International Economic Review* 59(4), 2043–2073.
- Lise, J. and K. Yamada (2019). Household sharing and commitment: Evidence from panel data on individual expenditures and time use. *Review of Economic Studies* 86(5), 2184–2219.
- Love, D. A. (2010). The effects of marital status and children on savings and portfolio choice. *The Review of Financial Studies* 23(1), 385–432.
- Low, H., C. Meghir, L. Pistaferri, and A. Voena (2020). Marriage, labor supply and the dynamics of the social safety net. Working Paper 24356, National Bureau of Economic Research.
- Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics* 3(1), 11–37.
- Mankart, J. and R. Oikonomou (2016). The rise of the added worker effect. *Economics Letters* 143, 48–51.
- Mankart, J. and R. Oikonomou (2017). Household search and the aggregate labour market. *Review of Economic Studies* 84(4), 1735–1788.
- Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *Review of Economic Studies* 74(3), 857–895.
- Nishiyama, S. (2019). The joint labor supply decision of married couples and the U.S. social security pension system. *Review of Economic Dynamics* 31, 277–304.
- Olsson, J. (2020). Structural transformation of the labor market and the aggregate economy. Working paper, University of Edinburgh.
- Ortigueira, S. and N. Siassi (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics* 60(6), 650–666.
- Ortigueira, S. and N. Siassi (2021). The u.s. tax-transfer system and low-income households: Savings, labor supply, and household formation. *Review of Economic Dynamics* forthcoming.

- Park, S. and D. Shin (2020). Welfare consequences of rising wage risk in the united states: Self selection into risky jobs and family labor supply adjustments. Working Paper 6/2020, Victoria University of Wellington.
- Persson, P. (2020). Social insurance and the marriage market. *Journal of Political Economy* 128(1), 252–300.
- Pestieau, P. and M. Sato (2008). Long-term care: The state, the market and the family. *Economica* 75, 435–454.
- Pruitt, S. and N. Turner (2020). Earnings risk in the household: Evidence from millions of US tax returns. *American Economic Review: Insights* 2(2), 237–54.
- Rasul, I. (2008). Household bargaining over fertility: Theory and evidence from Malaysia. *Journal of Development Economics* 86(2), 215–241.
- Rennane, S. (2020). A double safty net? Understanding interactions between disability benefits, formal assistance, and family support. *Journal of Health Economics* 69, 102248.
- Sanchez-Marcos, V. and C. Bethencourt (2018). The effect of public pensions on women’s labor market participation over a full life cycle. *Quantitative Economics* 9(2), 707–733.
- Schaller, J. (2013). For richer, if not for poorer? Marriage and divorce over the business cycle. *Journal of Population Economics* 26(3), 1007–1033.
- Schulz, B. and F. Siuda (2020). Marriage and divorce: The role of labor market institutions. Working Paper 8508, CESifo Munich.
- Schwartz, C. R. and R. D. Mare (2005). Trends in educational assortative marriage from 1940 to 2003. *Demography* 42(4), 621–646.
- Stevenson, B. and J. Wolfers (2007). Marriage and divorce: Changes and their driving forces. *Journal of Economic Perspectives* 21(2), 27–52.
- Vandenbroucke, G. and H. Zhu (2018). Unequal pink slips? Gender and the risk of unemployment. *The Regional Economist* 26(2).
- Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review* 105(8), 2295–2332.
- Wang, N. (2019). The demand for life insurance in a heterogeneous-agent life cycle economy with joint decisions. *Geneva Risk and Insurance Review* 44(2), 176–206.
- Wu, C. and D. Krueger (2021). Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. *American Economic Journal: Macroeconomics* 13(1), 79–113.