

Household gender disparity and portfolio choices

Giulia Bettin¹

Emma Sarno²

Md Taslim Uddin³

Alberto Zazzaro⁴

May 22, 2024

¹Università Politecnica delle Marche and MoFiR

²Università di Napoli L'Orientale

³Asian University for Women

⁴Università di Napoli Federico II, CSEF and MoFiR

Abstract

This paper examines how women's economic empowerment within households affects financial portfolio choices. Using data from the Bank of Italy's Survey on Household Income and Wealth for 2004–2020, we measure empowerment via the female income share and examine its association with participation in risky assets and the share of financial wealth invested in equities and private bonds. The results show that a higher female income share is, on average, associated with lower exposure to risky assets. However, this aggregate effect conceals a key asymmetry: in male-headed households, a greater female income share reduces risky investment, whereas in female-headed households it increases it. Tests based on imputed individual risk preferences provide little support for a bargaining model in which income shares reweight partners' risk attitudes. Instead, the evidence supports a unitary decision-making framework à la Becker, in which the identity and confidence of the household financial decision-maker are central. Financial-sector employment of the household head substantially attenuates these effects, whereas similarity in partners' risk preferences does not, pointing to self-confidence rather than altruistic preference aggregation as the primary mechanism underlying household portfolio choices.

Keywords: Household finance, Gender disparity, Asset allocation, Household unitary decision-making, Household bargaining, Risk aversion.

JEL Codes: D13; D14; G11; J16.

1 Introduction

A large body of research in household finance and behavioral economics documents systematic gender differences in financial behavior, linking these differences to variations in risk aversion, self-confidence, financial literacy, expectations, and gender norms (Barber and Odean, 2001; Lusardi and Mitchell, 2008; Croson and Gneezy, 2009; Ke, 2021; Guiso and Zaccaria, 2023). These findings are central to our understanding of individual portfolio choice. Yet most financial decisions are not made by isolated individuals, but within households, where preferences, resources, information, and decision rights interact.

This paper asks whether women’s economic empowerment within the household affects household financial investment decisions. More specifically, we study whether the share of household income earned by women and men affects both the probability that the household holds risky assets and the fraction of financial wealth invested in such assets. We also investigate the mechanism by which this effect operates. Does the relative contribution to household income affect portfolio choices by shifting bargaining power between spouses or partners? Or, are financial decisions delegated to a single household member (male or female)? In this case, does the relative contribution of women and men to household income affect investment choices through its impact on the decision-maker’s self-confidence, or through its influence on the degree of altruism with which the preferences of the partner are taken into account?

Addressing these questions is important because household financial decisions differ fundamentally from individual ones. When individuals form a household, investment choices are no longer determined by a single set of preferences, but instead reflect intra-household interactions shaped by marital status, differences in economic autonomy, and the allocation of decision-making power between partners (Love, 2010; ?; Bertocchi et al., 2014; Gomes et al., 2021). In this context, a partner’s contribution to household income may play a dual role. First, it can shift the balance of decision-making power within the household, thereby giving greater weight to his/her preferences and greater responsibility in household financial decisions. Second, it may affect the behavior of those who effectively manage the household portfolio, for instance, by increasing or reducing their confidence in their own financial skills and decisions, or by altering the extent to which they accommodate the partner’s preferences.

The distinction is theoretically relevant and empirically consequential. A bargaining view of the household implies that financial choices result from a weighted aggregation of

individual preferences, with weights depending on each partner’s economic resources or outside options (McElroy and Horney, 1981). Under this approach, the effect of women’s income share on risky investment should depend on the relative risk aversion of women and men within the household. If women are more risk-averse than men, a higher female income share should reduce risky investment; if they are less risk-averse, it should increase it; and if risk preferences are similar, it should have little or no effect. Importantly, these effects should hold regardless of whether the financial decision-maker is male or female. By contrast, a unitary or dictatorship view of the household emphasizes who controls financial decisions (Becker, 1981b). In this case, the same increase in women’s income share may have opposite effects depending on whether the household portfolio is managed by a woman or by a man.

We study these issues using the Bank of Italy’s Survey on Household Income and Wealth (SHIW) over the period 2004–2020. The SHIW is particularly well-suited for our purposes because it provides detailed information on household income, wealth, demographic characteristics, financial assets, individual income and education, and the identity of the household member responsible for financial decisions, as well as her/his financial literacy and risk propensity. We measure women’s economic empowerment using the female income share, defined as the ratio of total female income to total household income. Our main outcomes are two standard measures of risky financial investment: an indicator for holding risky assets and the share of financial wealth invested in risky assets. In the empirical analysis, risky assets are defined as the sum of equities and private bonds.

Before turning to mixed-gender households, the Italian data provide a useful benchmark on the individual gender gap in portfolio choice. Among single-person households, women invest substantially less in risky assets than men (see Table 1). Female singles hold equities in only 1.91 percent of cases, compared with 6.32 percent for male singles. The difference is even more striking in portfolio shares: equities account for 0.09% of financial wealth among single women and 2.7% among single men. A similar pattern emerges for the broader risky-asset measure, with holding rates of 7.84 percent for single women and 13.92 percent for single men, and portfolio shares of 5.99 and 9.65 percent, respectively. These descriptive figures confirm that the standard gender gap documented in the international literature is clearly present in the Italian SHIW data.

The relevance of the Italian setting for studying the role of women’s economic empowerment in household financial decisions is supported by the regression results reported in Table 2,

Table 1: Portfolio composition by gender

	Equity		Risky assets		Obs.
	<i> Holding </i>	<i> Share </i>	<i> Holding </i>	<i> Share </i>	
<i>Female</i>	1.91%	0.09%	7.84%	5.99%	10,426
<i>Male</i>	6.32%	2.7%	13.92%	9.65%	6,011

Source. Bank of Italy, SHIW (2004-2020). *Female* and *Male* denote single-person female and male households, respectively. Risky assets are defined as the sum of equities and private bonds. *Holding* is a dummy variable equal to 1 if the household holds equities or, alternatively, risky assets. *Share* denotes the proportion of the household’s financial wealth invested in equities or, alternatively, in risky assets.

which are based on the full SHIW sample and control for household income and time effects. First, both pooled OLS estimates of asset holdings and pooled Tobit estimates of portfolio shares indicate that single-person households exhibit portfolio choices that differ significantly from those of multi-person households and are, on average, more risk-taking. Second, this effect is far from homogeneous across genders. The coefficient associated with being a single male is systematically larger than that associated with being a single female. The probability that a single male holds equities is 4.3 percentage points higher than for multi-person households, whereas for a single female it is only 1.3 percentage points higher. For the broader category of risky assets, the corresponding figures are 6.8 and 3.6 percentage points for single men and single women, respectively. The same gender gap is evident in portfolio shares, which are significantly higher for men than for women, both for equities and for risky assets more broadly.

These preliminary, descriptive results serve two purposes. First, they validate in our empirical setting the basic premise that women and men differ in their financial investment behavior. Second, they show that household composition matters for portfolio choice. The fact that single men and single women behave differently suggests that individual gender traits, preferences, confidence, or social norms are relevant. At the same time, because our main analysis focuses on mixed-gender households, it is necessary to examine how these individual characteristics are combined, weighted, or overridden when financial decisions are made jointly or delegated to a single household member.

Our empirical analysis proceeds in three steps. First, we estimate the average relationship between the female income share and household risky investment in the full sample of mixed-

Table 2: Portfolio choice by household composition

	Equity			
	<i> Holding </i>		<i> Share </i>	
<i>Single</i>	0.024*** (0.003)		0.170*** (0.020)	
<i>Single female</i>		0.013*** (0.002)		0.037* (0.022)
<i>Single male</i>		0.043*** (0.004)		0.287*** (0.020)
Observations	61,508	61,508	49,830	49,830
	Risky assets			
	<i> Holding </i>		<i> Share </i>	
<i>Single</i>	0.048*** (0.004)		0.235*** (0.019)	
<i>Single female</i>		0.036*** (0.004)		0.175*** (0.024)
<i>Single male</i>		0.068*** (0.006)		0.306*** (0.025)
Observations	61,508	61,508	49,830	49,830

Source: Bank of Italy, SHIW (2004–2020). The table reports pooled OLS estimates for *Holding* , a dummy variable equal to 1 if the household holds equities or risky assets, and pooled Tobit estimates for *Share* , defined as the proportion of the household’s financial wealth invested in equities or in risky assets. The baseline specification is $y_{it} = \alpha + \beta Single_{it} + \gamma \log(income_{it}) + \lambda_t + \epsilon_{it}$, where *Single* is a dummy equal to 1 if the household is single-person and *income* denotes household income. The extended specification distinguishes single-person households by the head of household’s gender.

gender households. We use pooled linear probability models for risky asset holding and Tobit models for risky asset shares. The specifications control for household income, wealth, demographic characteristics, employment status, risk attitudes, marital status, homeownership, and region-by-year fixed effects. The baseline results show that a higher female income share is associated with lower investment in risky assets. This finding indicates that women’s contribution to household income is not neutral for household portfolio allocation.

Second, we test whether this average effect is consistent with a bargaining model of household financial decision-making. To do so, we examine whether the effect of female income share varies with the relative risk aversion of women and men within the household. Since the SHIW directly observes risk attitudes only for the household head, we impute risk

propensity for all adult household members using a random forest procedure. This allows us to classify households according to whether adult women are, on average, more risk-averse than men, less risk-averse than men, or equally risk-averse. The results do not support the bargaining interpretation. Female income share remains negatively related to risky investment even when women and men in the household display similar risk aversion. This is difficult to reconcile with a framework in which the only channel is the reweighting of male and female risk preferences.

Third, we examine whether the evidence is instead consistent with a unitary model in which the financial decision-maker's identity is central. We distinguish households by whether the person in charge of financial decisions is male or female and estimate selection-corrected models of household headship. The results reveal a striking asymmetry. In male-headed households, a higher female income share reduces both the probability of holding risky assets and the share of risky assets. In female-headed households, the effect is positive. This pattern is consistent with the idea that women's income share changes the self-confidence, authority, or perceived competence of the household financial decision-maker, rather than simply shifting the average weight of male and female preferences.

We then investigate the mechanism behind this unitary pattern. We consider two channels. The first is altruism: as the household head's contribution to household income increases, they may assign greater or lesser weight to the partner's preferences, depending on whether altruistic or self-interested motives prevail. If this were the main channel, the effect of female income share should disappear when women and men have similar risk preferences. The data do not support this prediction. The second mechanism is self-confidence in asset evaluation. In this interpretation, income contributions affect the decision-maker's confidence in assessing financial opportunities and risks. A higher female income share may strengthen confidence in female-headed households and weaken confidence or authority among men in male-headed households. Consistent with this view, the effect of female income share is attenuated when the household head is financially sophisticated, proxied by employment in the financial sector.

The paper contributes to the household finance literature in three ways. First, it connects the literature on gender differences in financial behavior to that on intra-household decision-making. We show that women's economic contribution within the household is a key determinant of risky investment, but that its effect cannot be understood without considering who manages household finances.

Second, the paper provides a direct empirical test of competing models of household portfolio choice. While bargaining models predict that the effect of women’s income share should depend on relative preferences, our evidence points to the importance of decision rights and headship. This finding is closely related to evidence that intra-household financial decision-making is shaped by the identity of the person who controls resources and makes financial choices (Bertocchi et al., 2014; Ke, 2021). It also suggests that income shares influence household finance not only through formal bargaining power, but also through psychological and organizational channels.

Third, the paper identifies self-confidence as a relevant mechanism through which women’s economic empowerment affects financial investment. This channel is consistent with the broader behavioral finance literature emphasizing overconfidence, underconfidence, and gender differences in financial decision-making (Barber and Odean, 2001; Lusardi and Mitchell, 2008). Our results suggest that confidence is not only an individual trait, but may also be shaped within the household by relative income contributions and by the gender of the household financial decision-maker.

Overall, our results show that women’s income share matters for household portfolio allocation, but not in the simple way predicted by a standard bargaining model. Women’s economic empowerment reduces risky investment on average, yet this average masks opposite effects across male- and female-headed households. The findings, therefore, highlight the importance of opening the black box of household decision-making. In household finance, who earns income matters; but who controls financial decisions matters just as much.

The remainder of the paper proceeds as follows. Section 2 places the paper within the related literature on gender differences in financial behavior, household finance, intra-household allocation, and social norms. Section 3 presents the SHIW data, the construction of the key variables, and the baseline empirical framework. Section 4 develops the theoretical framework and derives the empirical predictions of bargaining and unitary models of household portfolio choice, and presents the corresponding empirical evidence. Section 5 examines the mechanisms underlying the main results, distinguishing between altruism and self-confidence, and provides empirical evidence. Section 6 concludes.

2 Literature review

Four strands of literature are most closely related to this paper: gender differences in financial behavior; the effects of family structure on portfolio choice; models of intra-household portfolio allocation and decision-making; and the role of social norms, confidence, and psychological factors in household financial choices.

2.1 Gender differences in financial behavior

A large body of literature documents systematic gender differences in financial behavior and portfolio allocation. Compared with men, women are generally less likely to participate in financial markets, invest a smaller share of their wealth in risky assets, and adopt more conservative investment strategies (Jianakoplos and Bernasek, 1998; Sunden and Surette, 1998). Subsequent research links these differences to overconfidence, with Barber and Odean (2001) showing that men trade more aggressively and ultimately obtain lower net returns, as well as to lower financial literacy, more cautious return expectations, and reduced willingness to engage in risky financial environments (Dominitz and Manski, 2007; Lusardi and Mitchell, 2008; Croson and Gneezy, 2009; Charness and Gneezy, 2012). At the household level, however, these individual-level differences interact with bargaining dynamics, resource allocation, and decision authority in ways that individual-level analyses cannot fully capture.

2.2 Household finance: Family structure and portfolio choice

This insight motivates the growing literature on household finance, which emphasizes that portfolio allocation depends not only on individual preferences and characteristics, but also on marital status, family structure, and intra-household interactions (Gomes et al., 2021). Love (2010) shows that marital status transitions generate large and asymmetric portfolio adjustments by gender: divorce pushes men toward riskier allocations while moving women toward safer ones, with children further moderating these effects. Bertocchi et al. (2011) argue that marriage acts as an insurance mechanism for women, reducing background income risk and increasing participation in risky financial markets, though this effect weakens as female labor market participation rises.

2.3 Intra-household decision-making: Bargaining vs unitary models

A central issue concerns how household preferences are aggregated into financial decisions. The standard unitary model (Becker, 1981b) treats the household as a single decision-making unit maximizing a common utility function subject to pooled resources, while bargaining models (McElroy and Horney, 1981) predict that portfolio allocation shifts with changes in spouses' relative resources and outside options. Jianakoplos and Bernasek (2008) test these alternative frameworks directly and find that neither relative earnings nor other proxies for bargaining power significantly affect the risky asset share, a result more consistent with unitary decision-making, though the authors note that this is difficult to interpret without individual-level measures of risk aversion. Arano et al. (2010) qualify this finding: gender differences in retirement asset allocation disappear at the individual level after standard controls, but prove significant in the context of joint household decision-making. More structural approaches confirm that both partners' preferences matter, but in asymmetric ways. Addoum et al. (2016) and Addoum (2017) show that shocks to spouses' relative income generate changes in household risk aversion and risky asset holdings, with effects particularly pronounced in households where wives are more risk averse than husbands. Yilmazer and Lich (2015) document a starker asymmetry: only the risk tolerance of the spouse with greater bargaining power influences portfolio outcomes, while that of the less powerful partner is largely irrelevant. Gu et al. (2024) structurally estimate the distribution of bargaining power from household portfolio data across three countries, finding a significant gender gap that is only partially explained by observable characteristics such as income and employment, with the residual linked to perceived gender norms. Maura (2022) similarly finds that the collective model fits portfolio data substantially better than the unitary benchmark, and that household risk preferences shape the optimal risky share conditional on stock market participation.

2.4 Decision authority, financial headship, and social norms

An alternative perspective emphasizes who effectively controls financial decisions rather than how preferences are aggregated. Bertocchi et al. (2014), using the same SHIW data employed in this paper, show that the wife's probability of assuming financial responsibility rises with her relative age, education, and income, consistent with bargaining, but falls when she is employed, suggesting a specialization pattern that assigns management to the spouse with more available time. These patterns, however, are not fully explained by economic

characteristics alone. A growing body of evidence points to gender norms as an independent determinant of who controls financial decisions and with what consequences. Guiso and Zaccaria (2023), also using SHIW data, exploit variation across Italian cohorts and regions in the gender of the household head to show that more egalitarian norms raise market participation, equity holdings, diversification, and returns, with patriarchal norms that began receding in the early 1990s, when a pension reform raised the cost of adhering to traditional gender roles, a shift directly relevant to our sample period. Ke (2021), using data from U.S. household surveys, shows that the financial knowledge of the husband translates into stock market participation more readily than equivalent knowledge held by the wife, a pattern traced to gender identity norms that suppress women’s effective voice in financial decisions, as confirmed by a randomized experiment.

3 Data and Baseline Model

3.1 Data Sources

We utilize data from the Bank of Italy’s Survey on Household Income and Wealth (SHIW), which is conducted every other year on a representative sample of the Italian resident population. The SHIW provides a rotating panel comprising approximately 8,000 households per wave, offering detailed information on household demographics, income, wealth, and financial assets. Our period of analysis spans from 2004 to 2020 and covers eight survey waves. We restrict our sample to mixed-gender households containing both male and female adults. Within this sample, 69% are male-headed households, and 31% are female-headed households.

3.2 Empirical Specification

The baseline empirical specification takes the following form:

$$y_{hrt} = \alpha + \beta FIS_{hrt} + \sum_{z=1}^Z \gamma_z Control_{z,hrt} + \delta_{rt} + \epsilon_t \tag{1}$$

where h , r , and t denote the household, NUTS-2 region of residence, and time (wave), respectively. The key independent variable is the **Female Income Share** (FIS), which serves as our measure of women’s economic empowerment within the household. It is defined as the share of household income earned by women in the household relative to total household

income:

$$FIS = \frac{\text{Total Female Income}}{\text{Total Household Income}} \quad (2)$$

The dependent variable, y_{hrt} , measures the household's allocation to risky assets, captured in two ways. First, we define the dummy variable *Risk Holding*, which takes the value of 1 if the household holds equities and/or private bonds, and 0 otherwise. Second, we restrict attention to households investing in financial assets and define *FIS* as the ratio of the value of equities and private bonds held in the portfolio to the total value of the household's financial assets.

The vector $Control_z$ includes a comprehensive set of household and individual-level variables: household income and its square, wealth quartiles, number of adults, number of children, average risk aversion, average age of adult household members and its square, average education of adult household members, marital status (married, divorced), homeownership, and employment, self-employment and retired status (The complete list of variables used in the empirical analysis, together with their definitions, descriptive statistics, and mean differences between male- and female-headed households, is reported in Tables A1 and A2 in the Appendix). We also include region \times year fixed effects (δ_{rt}) to absorb local macroeconomic conditions and time-varying social norms, in the spirit of Guiso and Zaccaria (2023). We estimate pooled linear probability models and Tobit models, with standard errors clustered at the household level.

3.3 Baseline Results

Table 3 presents the baseline estimation results on the full sample of mixed-gender households. The coefficient on *FIS* is negative and highly significant in both risk-holding and risk-share specifications. Specifically, an increase in the female income share is associated with a 5.2-5.4 percentage-point decrease in the probability of holding risky assets and a substantial decrease in the overall portfolio share allocated to risky assets.

Table 3: Baseline results: FIS and asset Allocation

<i>FIS</i>	<i>Risk Holding</i>		<i>Risk Share</i>	
	-0.052*** [0.006]	-0.054*** [0.007]	-0.218*** [0.0291]	-0.222*** [0.030]
Controls	inc/wealth	all	inc/wealth	all
Region×Year FEs	yes	yes	yes	yes
Observations	41,916	41,916	34,732	34,732
R^2 /Pseudo R^2	0.210	0.225	0.197	0.216

Notes. Controls include: household income; household income squared; dummies for household wealth quartiles; average age of adults; average age of adults squared, average education of adults, HH employed, HH self-employed, HH retired, Partner employed, Partner self-employed; average risk aversion of adults, married status, divorced status, number of adults, number of children, homeownership. *** $p < 0.01$.

4 Household Decision-Making: Theory and Empirical Evidence

This section studies how women’s relative contribution to household income enters household financial decisions. We contrast two canonical views of intra-household decision-making. The first is a bargaining model, in the spirit of McElroy and Horney (1981), in which household choices aggregate the preferences of different household members, weighted by their relative bargaining power. The second is a unitary or dictatorship model, following the tradition of Becker (1974, 1981b), in which a single household member effectively controls financial decisions. The purpose of the exercise is not to provide a complete structural model of household behavior, but rather to derive sharp empirical implications that allow us to distinguish between these two views in the context of risky portfolio choice.

Consider a two-adult, mixed-gender household endowed with one unit of financial wealth. The household can allocate wealth between a risk-free asset with zero return and a risky asset s , whose return is distributed as $R_s \sim N(\mu_s, \sigma_s^2)$. Let α denote the share of wealth invested in the risky asset. Each adult household member $j \in \{F, M\}$, where F denotes the female member and M the male member, is characterized by a mean-variance utility function:

$$U_j(\alpha) = \alpha\mu_s - \frac{\gamma_j}{2}\alpha^2\sigma_s^2 - C \cdot \mathbb{1}_{\alpha>0}, \quad (3)$$

where γ_j captures individual risk aversion. Participation in the risky asset market also entails a fixed non-monetary cost C , so that investing in the risky asset requires paying $C \cdot \mathbb{1}_{\alpha>0}$.

4.1 The bargaining approach

Under the bargaining model, household utility is a weighted average of individual utilities:

$$V_H(\alpha) = \omega(FIS)U_F(\alpha) + [1 - \omega(FIS)]U_M(\alpha) - C \cdot \mathbb{1}_{\alpha>0}, \quad (4)$$

where FIS denotes the female income share. The bargaining weight of the female member increases with her income contribution, and, for expositional simplicity, we set $\omega(FIS) = FIS$. Substituting individual utilities into the household objective function gives:

$$V_H(\alpha) = \alpha\mu_s - \frac{\Gamma_H}{2}\alpha^2\sigma_s^2 - C \cdot \mathbb{1}_{\alpha>0}, \quad (5)$$

where household risk aversion is:

$$\Gamma_H = \gamma_M + FIS(\gamma_F - \gamma_M). \quad (6)$$

The optimal risky asset share is therefore:

$$\alpha^* = \begin{cases} 0 & \text{if } \Gamma_H > \tilde{\Gamma} = \frac{\mu_s^2}{2C\sigma_s^2}, \\ \frac{1}{\Gamma_H} \frac{\mu_s}{\sigma_s^2} & \text{if } \Gamma_H \leq \tilde{\Gamma} = \frac{\mu_s^2}{2C\sigma_s^2}. \end{cases} \quad (7)$$

Differentiating (7) with respect to FIS , it is straightforward to verify that

$$\frac{\partial Prob(\alpha_H^* > 0)}{\partial FIS} \begin{matrix} \leq \\ \geq \end{matrix} 0 \quad \text{and} \quad \frac{\partial \alpha_H^*}{\partial FIS} \begin{matrix} \leq \\ \geq \end{matrix} 0 \iff \gamma_F \begin{matrix} \geq \\ \leq \end{matrix} \gamma_M \quad (8)$$

Therefore, this framework delivers a simple testable implication. If household financial decisions result from a bargaining process in which members seek to have their preferences prevail, then when the female household member is more risk-averse than the male member, an increase in FIS raises household risk aversion and reduces both the probability of holding risky assets and the share of risky assets. If the female member is less risk-averse than the male member, the opposite should occur. If the two members have the same level of risk aversion, FIS should have no effect on household portfolio choices. Importantly, under this interpretation, the sign of the effect of FIS should depend on the relative risk aversion of women and men within the household, not on the gender of the person formally or informally

in charge of financial decisions.

4.1.1 Measuring risk preferences

Testing this prediction requires information on the risk preferences of all adult household members. The SHIW provides information on risk attitudes only for the survey respondent, who is typically the household head and the individual responsible for financial decisions. In particular, risk aversion is elicited by a survey question that asks respondents to rank their preferences among alternative investment profiles, ranging from very high returns with a high risk of capital loss to low returns with no risk of loss. Responses are recorded on a four-point scale from 1 (high risk tolerance) to 4 (strong risk aversion).¹ By survey design, this question was administered only to the household head, leaving risk preferences systematically unobserved for all other adult members, including spouses and partners whose attitudes are central to our intra-household bargaining framework. Excluding these observations would preclude construction of the gender-differentiated, household-level risk measures our analysis requires.

We address this through MissForest (Stekhoven and Bühlmann, 2012), a machine learning imputation algorithm built on Random Forests (Breiman, 2001). Unlike mean substitution, which compresses variance and biases estimates toward zero (Little and Rubin, 2002)—or Multiple Imputation by Chained Equations, which requires a separately specified parametric model for each incomplete variable and risks misspecification (van Buuren and Groothuis-Oudshoorn, 2011), MissForest imposes no distributional assumptions, handles continuous, binary, and categorical variables within a single unified framework, and learns non-linearities and interactions directly from the data. These properties make it particularly well-suited to our setting, where the determinants of risk tolerance likely interact in complex ways across gender, age, and socioeconomic background.

4.1.2 The MissForest algorithm

The algorithm works as follows. A Random Forest builds a large ensemble of decision trees, each trained on a bootstrap sample of observations with a random subset of predictors

¹Specifically, survey question C31 asks: “In managing your financial investments, would you say you have a preference for investments that offer: (a) very high returns, but with a high risk of losing part of the capital (1); (b) a good return, but also a fair degree of protection for the invested capital (2); (c) a fair return, with a good degree of protection for the invested capital (3); (d) low returns, with no risk of losing the invested capital (4)?”

considered at every split. For *risfin*, which takes $K = 4$ ordered values, the ensemble prediction is determined by majority vote across all B trees:

$$\hat{f}(\mathbf{x}^*) = \arg \max_{k \in \{1, \dots, K\}} \sum_{b=1}^B \mathbf{1}\{\hat{f}_b(\mathbf{x}^*) = k\}. \quad (9)$$

Node splits maximise the reduction in Gini impurity, $\text{Gini}(N) = 1 - \sum_{k=1}^K \hat{p}_k^2$, where \hat{p}_k is the proportion of observations in node N belonging to class k . Aggregating across many decorrelated trees substantially reduces prediction variance, and the out-of-bag (OOB) observations—those excluded from each tree’s bootstrap sample—provide an unbiased estimate of predictive accuracy without requiring a separate validation set.

MissForest applies this framework iteratively. Missing *risfin* entries are first initialized with the modal category observed among household heads. The algorithm then trains a classification forest on the observed portion of the data, predicts values for the missing entries, and replaces the placeholders with updated estimates. This process repeats, with each cycle using the most recently imputed values for all variables, until the proportion of category labels that change between successive iterations,

$$\Delta_{\text{cat}}^{(t)} = \frac{1}{|M_{\text{cat}}|} \sum_{(i,j) \in M_{\text{cat}}} \mathbf{1}\{\hat{X}_{ij}^{(t)} \neq \hat{X}_{ij}^{(t-1)}\}, \quad (10)$$

falls below $\tau = 10^{-5}$, subject to a maximum of ten iterations.

The set of variables used for the imputation includes age, gender, education, individual income, marital status, household role, occupational status, employment sector, and two binary indicators: *fin work*, which identifies whether the household member is employed in the financial services sector (defined as occupations coded as *aspett* $\in \{6, 7\}$, according to the SHIW classification), and *fin_study* (a finance-related educational qualification, *tipolau* $\in \{1, 4, 5, 6\}$), which are strong predictors of financial risk attitudes (Christiansen et al., 2008; Calvet and Sodini, 2014). Before estimation, we correct an important SHIW data convention: the survey replicates the household head’s *risfin* value across all household members by default. We reassign *risfin* to missing for all non-head individuals so that the algorithm trains exclusively on genuine survey responses. The procedure is implemented with 100 trees and a fixed random seed (`seed = 123`) for full replicability, and is applied separately to each of the seven survey waves spanning 2006 to 2020, ensuring that imputed values reflect the distributional properties and covariate structure prevailing in each period rather than a pooled

average that could obscure meaningful variation over time.

We assess the reliability of the imputed values using two complementary diagnostics applied wave by wave. First, we report the OOB classification accuracy produced by MissForest, an unbiased internal estimate of predictive performance that requires no separate holdout set. Second, we compute two-sample Kolmogorov-Smirnov (KS) statistics comparing the empirical distribution of observed *risfin* values (household heads with directly elicited responses) against the full distribution of imputed values across all adult members, as a practical measure of distributional distortion introduced by the imputation. Results are presented in Table 4, and the corresponding kernel density overlays for each wave are shown in Figure 1.

Table 4: MissForest imputation diagnostics by survey wave

Wave	OOB accuracy	KS statistic	Assessment
2006	0.650	0.142	Good
2008	0.671	0.170	Good
2010	0.563	0.117	Acceptable
2012	0.733	0.108	Strong
2014	0.699	0.196	Good
2016	0.656	0.151	Good
2020	0.583	0.114	Acceptable

Notes: OOB accuracy is $1 - \text{PFC}$, where PFC is the proportion of falsely classified observations from MissForest’s internal out-of-bag estimation; the random-guessing benchmark for a four-category variable is 0.25. The KS statistic is the two-sample Kolmogorov-Smirnov statistic comparing the empirical distribution of observed values for household heads and directly elicited responses with the distribution of imputed values across all adult household members. *Strong* ≥ 0.70 ; *Good* 0.62–0.69; *Acceptable* 0.55–0.61.

OOB accuracy ranges from 0.563 in 2010 to 0.733 in 2012, comfortably exceeding the 0.25 benchmark of random classification across four categories in every wave. Five of the seven waves achieve accuracy of at least 0.650, confirming that the observed individual and household characteristics carry substantial predictive content for financial risk propensity. The two lower-performing waves, 2010 and 2020, coincide with periods of exceptional macroeconomic disruption, namely the aftermath of the global financial crisis and the COVID-19 pandemic, respectively, when the relationship between socioeconomic characteristics and risk preferences may have been temporarily less stable. Even in these waves, however, accuracy remains more than twice the random-guessing baseline. KS statistics range from 0.108 in 2012 to 0.196 in 2014, indicating that the cumulative distributions of observed and imputed *risfin* values

differ by at most 20 percentage points across all waves. The associated p-values are all zero, a pattern that reflects the well-documented tendency of the KS test to flag even negligible differences as statistically significant with large samples, rather than as evidence of meaningful distributional distortion.

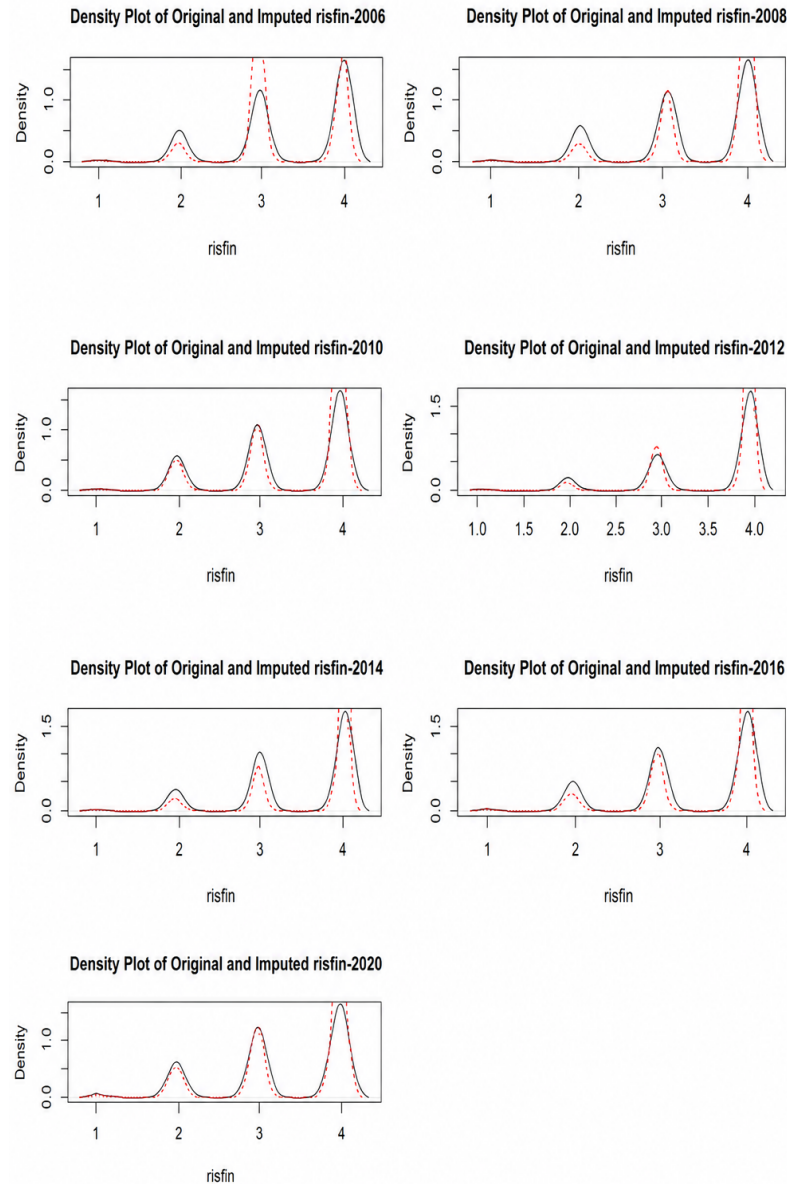


Figure 1: Kernel density plots of observed and imputed *risfin* values by survey wave

Notes. Each panel plots the kernel density of observed *risfin* values (solid black line, household heads with directly elicited responses) against the kernel density of imputed *risfin* values (dashed red line, all adult household members including non-heads) for the indicated survey wave. The four-point *risfin* scale runs from 1 (high risk tolerance) to 4 (strong risk aversion). The close correspondence between the two density curves in each wave indicates that the MissForest imputation preserves the marginal distribution of financial risk propensity without introducing systematic bias toward any single risk category.

The kernel density plots in Figure 1 provide more informative visual evidence: across

all seven waves, the imputed density (dashed red line) closely follows the observed density (solid black line), reproducing both the bimodal structure of the distribution-, concentrated at categories 3 and 4, consistent with the broad risk aversion documented in the household finance literature, and the relative magnitudes of the two peaks. No wave exhibits systematic compression toward a single category or meaningful displacement of distributional mass, suggesting the imputation preserves the shape of the underlying risk preference distribution with good fidelity. Taken together, the OOB accuracy figures and the density comparisons support the reliability of the imputed *risfin* values across the full sample period.

Following imputation, we construct the partner’s risk propensity, the gender-disaggregated average risk-preference for adult male and female members, and the overall household-level risk-preference measures used in the regression analysis. This allows us to incorporate the risk preferences of all household members into our analysis and aggregate them across genders, rather than relying exclusively on risk preferences reported by the household head. It also enables us to assess whether household financial investment decisions are consistent with a bargaining process among members with heterogeneous risk attitudes (McElroy and Horney, 1981; Browning and Chiappori, 1998; Mazzocco, 2004), or with a unitary framework in which within-household differences in risk preferences influence the choices of a benevolent household head (Becker, 1981a).

4.1.3 Evidence

To test whether the bargaining model can explain the observed pattern of household financial decisions, we classify households, using both observed and imputed measures of risk aversion, into three categories: adult women are, on average, more risk-averse than adult men, less risk-averse, or equally risk-averse. Formally, let $\Delta_{FM} = Avg(\gamma_F) - Avg(\gamma_M)$ denote the difference between the average risk aversion of female and male household members. Therefore, we estimate the following model:

$$y_{hrt} = \alpha + \beta_1 FIS \times \mathbb{1}_{\Delta_{FM} > 0} + \beta_2 FIS \times \mathbb{1}_{\Delta_{FM} < 0} + \beta_3 FIS \times \mathbb{1}_{\Delta_{FM} = 0} + \sum_{z=1}^Z \gamma_z Control_{z,hrt} + \delta_{rt} + \epsilon_t \quad (11)$$

where $\mathbb{1}_{\Delta_{FM} \lesseqgtr 0}$ denotes three indicator variables that take the value of one when Δ_{FM} is less than, greater than, or equal to zero, respectively. The comparative statics in (8) imply

that, under the bargaining model, $\beta_1 < 0$, $\beta_2 > 0$, and $\beta_3 = 0$, corresponding to cases in which $\Delta_{FM} > 0$, $\Delta_{FM} < 0$, and $\Delta_{FM} = 0$, respectively.

Table 5 reports the corresponding estimates. The evidence does not support the bargaining prediction. The coefficient of FIS is negative when women are more risk-averse than men, and it remains negative, albeit smaller in magnitude and less statistically significant, even when women are less risk-averse. More importantly, it remains negative and highly significant even when adult women and men within the household have the same level of risk aversion. In this latter case, when the preferences of household members coincide, the bargaining model predicts no effect of *FIS* on either risky-asset participation or portfolio shares. The persistence of a sizeable effect in this group is therefore difficult to reconcile with a model in which women’s income share matters only through changes in the relative weight attached to female and male risk preferences.

Table 5: Bargaining model: Female income share and relative risk aversion

	<i>Risk Holding</i>	<i>Risk Share</i>
$FIS \times \mathbb{1}_{\Delta_{FM} > 0}$	-0.099*** (0.010)	-0.342*** (0.042)
$FIS \times \mathbb{1}_{\Delta_{FM} < 0}$	-0.018* (0.011)	-0.056 (0.045)
$FIS \times \mathbb{1}_{\Delta_{FM} = 0}$	-0.043*** (0.008)	-0.200*** (0.037)
Controls	Yes	Yes
Region \times year fixed effects	Yes	Yes
Observations	41,916	34,732
R^2 /Pseudo- R^2	0.227	0.218

Notes: The table reports pooled linear probability estimates for *Risk Holding* and pooled Tobit estimates for *Risk Share*. $\Delta_{FM} = Avg(\gamma_F) - Avg(\gamma_M)$ denotes the average female–male risk-aversion differential within the household. $\mathbb{1}_{\Delta_{FM} > 0}$ identifies households in which adult women are, on average, more risk averse than adult men; $\mathbb{1}_{\Delta_{FM} < 0}$ identifies households in which adult women are less risk averse; $\mathbb{1}_{\Delta_{FM} = 0}$ identifies households in which adult women and men have the same average risk aversion. Standard errors are reported in parentheses. Controls are the same as in Table 3. In addition, we control for the average risk aversion of adult females and average risk aversion of adult males. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.2 The unitary household approach

We next consider a unitary model of household decision-making. In this framework, portfolio choices are not the outcome of a weighted average of individual preferences but rather reflect the preferences, beliefs, and confidence of the household member (the head of the household) who effectively manages financial investments. This view is consistent with a Beckerian representation of the household as if it maximized a single objective function, possibly that of an altruistic, benevolent household head (Becker, 1981b,a). In the present context, the key implication is that the effect of FIS should depend on who controls financial decisions. If the household head is female, a higher female income share may reinforce the decision-maker's confidence and authority in financial matters or influence her level of altruism. If the household head is male, an increase in the female income share (FIS), which corresponds to a decline in the male income share $1 - FIS$, may instead weaken the perceived dominance of the male decision-maker, reduce his confidence in portfolio choices, or alter his willingness to accommodate the preferences of other household members. Regardless of the specific channel through which FIS operates, the key implication is that, within a unitary household framework, the effects on investment choices should exhibit opposite signs depending on whether the household head is male or female.

We formalize this idea by allowing the household objective to differ according to the gender of the financial decision-maker:

$$V_{FH} = f_F(U_1, \dots, U_n, FIS) \quad (12)$$

$$V_{MH} = f_M(U_1, \dots, U_n, 1 - FIS) \quad (13)$$

where V_{FH} and V_{MH} denote the objective functions of female-headed and male-headed households, respectively. Unlike in the bargaining model, the sign of the effect of FIS is not pinned down by the relative risk aversion of women and men. Instead, it is expected to vary with the identity of the household head. However, a central prediction of the unitary model is that the effect of women's income share on risky investment should be different, and potentially opposite, in male- and female-headed households

We estimate Heckman selection models for household headship and portfolio outcomes, using widow or widower status as an exclusion restriction. Consistent with a unitary framework, the results reported in Table 6 show a sharp asymmetry. First, the probability that the

household head is a woman increases with women’s contribution to total household income (for brevity, the selection equation estimates are not reported in the table). Turning to the outcome equations, in male-headed households, *FIS* has a negative and statistically significant effect on both risky asset holding and the share of risky assets. In female-headed households, by contrast, the effect is positive and statistically significant. Quantitatively, an increase in the female income share reduces the probability of holding risky assets by 11 percentage points in male-headed households and increases it by 8.1 percentage points in female-headed households. The same sign reversal is observed for the share of risky assets.

Table 6: Unitary model: Female income share by gender of household head

	<i>Male headship</i>		<i>Female headship</i>	
	<i>Risk Holding</i>	<i>Risk Share</i>	<i>Risk Holding</i>	<i>Risk Share</i>
<i>FIS</i>	−0.110*** (0.020)	−0.085*** (0.016)	0.081*** (0.023)	0.058*** (0.022)
λ	0.008 (0.013)	0.017* (0.009)	0.025* (0.014)	0.003 (0.012)
Controls	Yes	Yes	Yes	Yes
Region×year fixed effects	Yes	Yes	Yes	Yes
Observations	28,876	24,335	13,040	10,397

Notes: The table reports Heckman selection estimates for risky asset holding and risky asset shares by gender of the household head. Male- and female-headed households are households in which the person in charge of financial decisions is male or female, respectively. Controls are the same as in Table 5. Widow/widower status is used as an exclusion restriction. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The differential impact of women’s economic empowerment on household investment choices is difficult to reconcile with a pure intra-household bargaining mechanism, but consistent with a unitary model in which the identity of the financial decision-maker plays a central role. The evidence suggests that women’s contribution to household income matters not only through the weight their preferences carry in shaping household investment decisions, but also through the effect on the allocation of financial decision rights. The evidence suggests that women’s contribution to household income matters not only through the weight their preferences carry in shaping household investment decisions, but also through its effect on the allocation of authority in financial decision-making. In other words, the female income share influences portfolio choices through the household decision-making structure: the same change in income composition has opposite implications depending on whether financial management is effectively entrusted to a man or to a woman.

5 Mechanisms

In this section, we examine the mechanisms through which the female income share may affect household investment decisions in a unitary household decision-making framework. We focus on two channels. The first is an *altruism channel*. Under this mechanism, the female income share affects portfolio choices by altering the extent to which the household head internalizes the utility of the other adult household member. The second is a *self-confidence channel*. Under this mechanism, a greater female contribution to household income may affect the head of household's self-confidence. A larger income contribution may increase the perceived authority and confidence of the household head when the head is a woman, while it may weaken the confidence or relative authority of a male head. These two channels generate distinct empirical implications, allowing us to distinguish between them.

5.1 The altruism channel

5.1.1 Theory

The altruism mechanism builds on a unitary representation of household financial decisions, like in equations (??) and (13), in which the household head chooses the portfolio but may assign positive weight to the partner's utility. In particular, assuming that household members are characterized by mean-variance utility functions, in female-headed households, the female head maximizes

$$\begin{aligned} V_{FH} &= U_F + a_F(FIS)U_M - C \cdot \mathbb{1}_{\alpha>0} \\ &= [1 + \alpha(FIS)]\mu_s + \frac{\gamma_F + a_F(FIS)\gamma_M}{2}\alpha^2\sigma_s^2 - C \cdot \mathbb{1}_{\alpha>0} \end{aligned} \quad (14)$$

while in male-headed households, the male head maximizes

$$\begin{aligned} V_{MH} &= U_M + a_M(1 - FIS)U_F - C \cdot \mathbb{1}_{\alpha>0} \\ &= [1 + \alpha(1 - FIS)]\mu_s + \frac{\gamma_M + a_M(1 - FIS)\gamma_F}{2}\alpha^2\sigma_s^2 - C \cdot \mathbb{1}_{\alpha>0}, \end{aligned} \quad (15)$$

where $a_F(FIS)$ and $a_F(1 - FIS)$ capture the degree of altruism of the female and male head, respectively. The optimal risky asset allocation in female- and male-headed households are

$$\alpha_{FH}^* = \begin{cases} 0 & \text{if } \Gamma_{FH} > \tilde{\Gamma}_{FH} = \frac{[(1+a_F(FIS))\mu_s]^2}{2C\sigma_s^2} \\ \frac{1+a_F(FIS)}{\Gamma_{MH}} \frac{\mu_s}{\sigma_s^2} & \text{if } \Gamma_{FH} \leq \tilde{\Gamma}_{FH} = \frac{[(1+a_F(FIS))\mu_s]^2}{2C\sigma_s^2} \end{cases} \quad (16)$$

$$\alpha_{MH}^* = \begin{cases} 0 & \text{if } \Gamma_{MH} > \tilde{\Gamma}_{MH} = \frac{[(1+a_M(1-FIS))\mu_s]^2}{2C\sigma_s^2} \\ \frac{1+a_M(1-FIS)}{\Gamma_{MH}} \frac{\mu_s}{\sigma_s^2} & \text{if } \Gamma_{MH} \leq \tilde{\Gamma}_{MH} = \frac{[(1+a_M(1-FIS))\mu_s]^2}{2C\sigma_s^2} \end{cases} \quad (17)$$

where $\Gamma_{FH} = \gamma_F + a_F(FIS)\gamma_M$ and $\Gamma_{MH} = \gamma_M + a_M(1 - FIS)\gamma_F$ denote the effective risk-aversion parameters. Differentiating (16) and (17) with respect to FIS , it is straightforward to verify that the effect of the share of income produced by females on risky asset holding and risky asset shares depends on the interaction between altruism and the gender gap in risk aversion:

$$sign \frac{\partial Prob(\alpha_{FH} > 0)}{\partial FIS} = sign \frac{\partial \alpha_{FH}^*}{\partial FIS} = sign \frac{\partial a_F}{\partial FIS} (\gamma_F - \gamma_M) \quad (18)$$

$$sign \frac{\partial Prob(\alpha_{MH} > 0)}{\partial FIS} = sign \frac{\partial \alpha_{MH}^*}{\partial FIS} = -sign \frac{\partial a_M}{\partial FIS} (\gamma_F - \gamma_M) \quad (19)$$

Although the female income share (FIS) may exert an ambiguous effect on household asset allocation, depending on whether altruism increases or decreases with it, the altruism channel delivers a clear, testable implication: if adult females and males are equally risk averse, FIS should have no effect on the asset allocation of either female- or male-headed households. Accordingly, if the female income share affects household portfolios primarily through the formation of altruistic preferences, its effect should vanish in the subsample of households in which adult women and men exhibit the same average level of risk aversion.

5.1.2 Evidence

To assess the relevance of the altruism channel in explaining the opposite effects of FIS on the financial investment choices of male- and female-headed households, we estimate a Heckman selection model on the restricted sample of households in which adult women and men exhibit the same average level of risk aversion. In this setting, the coefficient on FIS

should be statistically indistinguishable from zero for both risky-asset holdings and risky-asset shares. Table 7 reports the results. In male-headed households, the coefficient on *FIS* remains negative and statistically significant. An increase in the female income share reduces both the likelihood of holding risky assets and the share of risky assets, even when adult women and men exhibit the same average level of risk aversion. Similarly, in female-headed households, the coefficient on *FIS* remains positive and statistically significant for both outcomes.

Table 7: Altruism mechanism: households with equal average risk aversion

	<i>Male headship</i>		<i>Female headship</i>	
	<i>Risk holding</i>	<i>Risk share</i>	<i>Risk holding</i>	<i>Risk share</i>
FIS	-0.103*** (0.030)	-0.078*** (0.024)	0.089*** (0.033)	0.091*** (0.033)
λ	0.005 (0.019)	0.012 (0.014)	0.042** (0.020)	0.029* (0.017)
Controls	Yes	Yes	Yes	Yes
Region \times year FEs	Yes	Yes	Yes	Yes
Observations	11,702	9,630	5,418	4,241

Notes. Heckman selection models estimated on the subsample of households in which adult women and adult men have the same average risk aversion. Control and exclusion restriction are the same as in Table 6. Robust standard errors in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The persistence of significant effects in this restricted sample is inconsistent with the central prediction of the altruism channel. If the female income share operated only by changing the weight attached to the partner’s utility, it should not affect risky investment when male and female adults have identical average risk preferences. Instead, the coefficient on FIS remains large, statistically significant, and opposite in sign across male- and female-headed households. This pattern suggests that the female income share affects portfolio choices through a mechanism that goes beyond reweighting male and female risk preferences within the household objective function.

5.2 The self-confidence channel

5.2.1 Theory

The second mechanism links the female income share to the household head’s confidence in their own assessment of risky assets. Portfolio choices depend not only on “true” risk preferences, but also on the decision-maker’s perceived ability to evaluate the distribution of

risky returns. A more confident household head may perceive the investment environment as less uncertain, process financial information more effectively, or feel more entitled to take responsibility for risky financial decisions without taking into account preferences and opinions that differ from their own.

To parsimoniously capture this mechanism within our mean–variance framework, we assume that a lack of confidence in one’s own authority and judgment amplifies the negative weight assigned to investment risk in the household head’s utility function

$$V_{FH} = \alpha\mu_s - \frac{\gamma_F}{2} \frac{1}{c_F(\phi, FIS)} \alpha^2 \sigma_s^2 \quad (20)$$

$$V_{MH} = \alpha\mu_s - \frac{\gamma_M}{2} \frac{1}{c_M(\phi, 1 - FIS)} \alpha^2 \sigma_s^2 \quad (21)$$

where $c_F(\phi, FIS)$ denotes the confidence of a female household head in evaluating risky asset returns, and $c_M(\phi, 1 - FIS)$ the corresponding confidence of a male household head, both of which increase with their contribution to household income and their financial expertise, such as employment in the financial sector.

In this case, the optimal portfolio allocation for female- and male-headed M-Households is

$$\alpha_{FH}^* = \begin{cases} 0 & \text{if } \gamma_F > \tilde{\gamma}_F = \frac{c_F(\phi, FIS)\mu_s^2}{2C\sigma_s^2} \\ \frac{c_F(\phi, FIS)\mu_s}{\gamma_F} \frac{\mu_s}{\sigma_s^2} & \text{if } \gamma_F \leq \tilde{\gamma}_F = \frac{c_F(\phi, FIS)\mu_s^2}{2C\sigma_s^2} \end{cases} \quad (22)$$

$$\alpha_{MH}^* = \begin{cases} 0 & \text{if } \gamma_M > \tilde{\gamma}_M = \frac{c_M(\phi, 1-FIS)\mu_s^2}{C\sigma_s^2} \\ \frac{c_M(\phi, 1-FIS)\mu_s}{\gamma_M} \frac{\mu_s}{\sigma_s^2} & \text{if } \gamma_M \leq \tilde{\gamma}_M = \frac{c_M(\phi, 1-FIS)\mu_s^2}{C\sigma_s^2} \end{cases} \quad (23)$$

From (22) and (23), and consistent with the evidence reported in Table 6, an increase in the female share of household income raises both the likelihood that a female household head invests in risky assets and the share of such assets held, while exerting a negative effect on the risky investment choices of a male household head. Moreover, depending on whether the contribution to household income acts as a substitute for or a complement to financial experience in shaping the household head’s confidence, the impact of FIS may either amplify or attenuate the positive effect of the household head’s level of financial sophistication on risk-taking. In either case, its effect should be opposite for male and female household heads.

5.2.2 Evidence

We test this prediction by augmenting the Heckman selection model with an interaction between *FIS* and an indicator variable, *HFIN*, which equals 1 if the household head is employed in the financial sector. Employment in finance is used as a proxy for financial sophistication and confidence in asset evaluation.

Table 8 reports the results. In male-headed households, the coefficient on *FIS* is negative and statistically significant, while its interaction with *HFIN* is positive and statistically significant for risky asset holding and positive, though not statistically significant, for the risky asset share. In female-headed households, the pattern is reversed: the coefficient on *FIS* is positive and statistically significant, whereas the interaction between *FIS* and *HFIN* is negative, particularly strong, and statistically significant for the risky asset share. These findings are consistent with a self-confidence mechanism underlying the household head's unitary investment decisions. Employment in the financial sector and a higher share of household income both increase the likelihood of investing in risky assets and the portfolio share allocated to such assets. Moreover, the opposite signs of the interaction between *FIS* and *HFIN* across male- and female-headed households are consistent with the two acting as substitutes: financial sophistication reduces the marginal effect of income shares on the head of household's self-confidence. Accordingly, when the household head is already highly confident in their ability to accurately assess the risks of financial assets, changes in relative income contributions exert a more limited effect on portfolio choices.

6 Conclusions

This paper examines whether and how women's contributions to household income affect household financial investment decisions. The evidence shows that the female income share is an important determinant of household portfolio allocation. On average, a higher female income share is associated with lower exposure to risky assets. However, this average effect masks substantial heterogeneity by the gender of the household head responsible for financial decisions.

The main results indicate that the effect of the female income share does not vary systematically with the relative risk aversion of women and men within the household, which is inconsistent with the predictions of a bargaining model. Instead, it varies sharply with the

Table 8: Self-confidence mechanism: interaction with financial-sector employment

	Male headship		Female headship	
	Risk holding	Risk share	Risk holding	Risk share
FIS	-0.110*** [0.027]	-0.102*** [0.020]	0.099*** [0.026]	0.053** [0.026]
<i>HFIN</i>	0.156*** [0.021]	0.099*** [0.021]	0.305*** [0.073]	0.228*** [0.055]
FIS \times <i>HFIN</i>	0.140* [0.078]	0.037 [0.055]	-0.168 [0.119]	-0.214** [0.090]
λ	0.027 [0.019]	0.035*** [0.012]	0.047*** [0.016]	0.018 [0.013]
Controls	Yes	Yes	Yes	Yes
Region \times year FEs	Yes	Yes	Yes	Yes
Observations	28,876	24,335	13,040	10,397

Notes. Heckman selection models augmented with the interaction between the female income share and *HFIN*, an indicator equal to one when the household head is employed in the financial sector. Control and exclusion restriction are the same as in Table 6. Robust standard errors in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

gender of the financial decision-maker. In male-headed households, a higher female income share reduces both the probability of holding risky assets and the share of risky assets in the household portfolio. In female-headed households, the same increase has a positive effect on both outcomes. This pattern is consistent with a unitary model of household financial behavior in which the identity of the person who holds the purse strings is central.

The mechanism analysis further shows that the effect of the female income share is unlikely to operate through altruism. When the sample is restricted to households in which adult women and men have the same average risk aversion, the effect of the female income share remains statistically significant and has the opposite sign across male- and female-headed households. By contrast, the evidence supports a self-confidence channel. The effect of the female income share is attenuated when the household head is employed in the financial sector, suggesting that income shares matter less when the decision-maker already possesses financial expertise and confidence in asset evaluation.

Overall, women's economic empowerment affects household financial decisions not only by changing the distribution of resources within the household but also by interacting with the allocation of financial decision rights. The same increase in women's contribution to household income can lead to more or less financial risk-taking depending on who manages the household portfolio.

References

- Addoum, J. M. (2017). Household portfolio choice and retirement. *Review of Finance*, 21(4):1641–1683.
- Addoum, J. M., Kurnaz, M., and Veldkamp, L. (2016). Household portfolio diversification. Working Paper 22416, National Bureau of Economic Research.
- Arano, K., Parker, C., and Terry, R. (2010). Gender differences in financial risk tolerance. *Financial Services Review*, 19(2):141–155.
- Barber, B. M. and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics*, 116(1):261–292.
- Becker, G. S. (1974). A theory of marriage: Part ii. *Journal of Political Economy*, 82(2, Part 2):S11–S26.
- Becker, G. S. (1981a). Altruism in the family and selfishness in the market place. *Economica*, 48(1):1–15.
- Becker, G. S. (1981b). *A Treatise on the Family*. Harvard University Press.
- Bertocchi, G., Brunetti, M., and Torricelli, C. (2011). Marriage and stock market participation: Who call the shots? *Journal of Economic Behavior & Organization*, 80(3):562–578.
- Bertocchi, G., Brunetti, M., and Torricelli, C. (2014). Who holds the purse strings within the household? the determinants of intra-household financial decision-making. *Review of Economics of the Household*, 12:651–678.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.
- Browning, M. and Chiappori, P.-A. (1998). Efficient intra-household allocations: A general characterisation and empirical tests. *Econometrica*, 66(6):1241–1278.
- Calvet, L. E. and Sodini, P. (2014). Twin picks: Disentangling the determinants of risk-taking in household portfolios. *Journal of Finance*, 69(2):867–906.
- Charness, G. and Gneezy, U. (2012). Strong evidence for gender differences in risk taking. *Journal of Economic Behavior & Organization*, 83(1):50–58.

- Christiansen, C., Joensen, J. S., and Rangvid, J. (2008). Are economists more likely to hold stocks? *Review of Finance*, 12(3):465–496.
- Croson, R. and Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, 47(2):448–474.
- Dominitz, J. and Manski, C. F. (2007). Expected equity returns and portfolio decisions of older americans. *Journal of the European Economic Association*, 5(2-3):552–560.
- Gomes, F., Haliassos, M., and Ramadorai, T. (2021). Household finance. *Journal of Economic Literature*, 59(3):919–1000.
- Gu, R., Peng, C., and Zhang, W. (2024). The gender gap in household bargaining power: A revealed-preference approach. *The Review of Financial Studies*, page hhae039.
- Guiso, L. and Zaccaria, T. (2023). Gender, social norms, and household portfolio allocation. Discussion paper, CEPR.
- Jianakoplos, N. A. and Bernasek, A. (1998). Are women more risk averse? *Economic Inquiry*, 36(4):620–630.
- Jianakoplos, N. A. and Bernasek, A. (2008). Family financial risk taking when the wife earns more. *Journal of Family and Economic Issues*, 29(2):289–306.
- Ke, D. (2021). Who wears the pants? gender identity norms and household financial decisions. *Journal of Finance*, 76(3):1389–1425.
- Little, R. J. A. and Rubin, D. B. (2002). *Statistical Analysis with Missing Data*. Wiley, Hoboken, NJ, 2nd edition.
- Love, D. A. (2010). The effects of marital status and children on savings and portfolio choice. *Review of Financial Studies*, 23(1):385–432.
- Lusardi, A. and Mitchell, O. S. (2008). Planning and financial literacy: How do women fare? *American Economic Review*, 98(2):413–417.
- Maura, G. (2022). Intra-household bargaining and household financial decisions. Technical report, Working Paper.

- Mazzocco, M. (2004). Savings, risk sharing, and preferences for risk. *American Economic Review*, 94(4):1169–1182.
- McElroy, M. B. and Horney, M. J. (1981). Nash-bargained household decisions: Toward a generalization of the theory of demand. *International Economic Review*, 22(2):333–349.
- Stekhoven, D. J. and Bühlmann, P. (2012). MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118.
- Sunden, A. E. and Surette, B. J. (1998). Gender differences in the allocation of assets in retirement savings plans. *American Economic Review*, 88(2):207–211.
- van Buuren, S. and Groothuis-Oudshoorn, K. (2011). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, 45(3):1–67.
- Yilmazer, T. and Lich, S. (2015). Household portfolio choice and management. *Review of Economics of the Household*, 13(2):275–312.

Appendix: additional tables

Table A1: Variable definitions and statistics

Variable	Definition	N	Mean (SD)
Risk holding	Binary indicator = 1 if the household holds any risky financial assets (equities or private bonds), 0 otherwise.	61675	0.14 (0.35)
Risk share	Share of risky assets in total financial assets. Continuous variable in [0, 1].	49851	0.10 (0.24)
Core economic covariates			
FIS	Female income share in household: $yindex2 = fem_y / (fem_y + mal_y)$, where fem_y is total income of all female members and mal_y is total income of all male members. Ranges from 0 to 1; 0.5 indicates income parity. Proxies the relative economic bargaining power of women within the household.	61513	0.41 (0.38)
logy	Natural logarithm of total household income: $logy = \log(y)$.	61508	10.19 (0.72)
logy2	Squared natural log of household income	61508	104.3 (14.3)
Labor market status			
employed_2	Binary indicator = 1 if the household head is a wage or salaried employee, 0 otherwise.	61675	0.32 (0.47)
selfemployed_2	Binary indicator = 1 if the household head is self-employed, 0 otherwise.	61675	0.10 (0.30)
retired_2	Binary indicator = 1 if the household head is retired, 0 otherwise.	61675	0.36 (0.48)
partner_employed_1	Binary indicator = 1 if the household head's partner is a wage or salaried employee, 0 otherwise: $partner_employed_1 = 1$ if $partner_employ_status = 1$.	61675	0.19 (0.40)
partner_selfemployed	Binary indicator = 1 if the household head's partner is self-employed, 0 otherwise: $partner_selfemployed = 1$ if $partner_employ_status = 2$.	61675	0.05 (0.21)
Education and financial literacy			
avg_imp_u_risfin_adult	Average risk propensity of all adult members	61675	3.44 (0.58)
avg_edu_all_adults	Average education of all adult members (age ≥ 18)	61675	3.77 (1.55)
Wealth controls			
quantile2wealth	Binary indicator = 1 if the household belongs to the second quartile of the net wealth distribution ($w_quartile = 2$), 0 otherwise.	61675	0.25 (0.43)
quantile3wealth	Binary indicator = 1 if the household belongs to the third wealth quartile ($w_quartile = 3$), 0 otherwise.	61675	0.25 (0.43)
quantile4wealth	Binary indicator = 1 if the household belongs to the fourth (top) wealth quartile ($w_quartile = 4$), 0 otherwise.	61675	0.25 (0.43)
Household structure			
female	Binary indicator = 1 if the household head is female, 0 otherwise.	61,675	0.42 (0.49)
num_adult	Number of adult household members	61675	2.09 (0.96)
num_kid	Number of children (age < 18) in household	61675	0.35 (0.73)
married	Binary indicator = 1 if the household head is legally married, 0 otherwise.	61675	0.61 (0.49)
divorced	Binary indicator = 1 if the household head is divorced, 0 otherwise.	61675	0.08 (0.27)
homeowner	Binary indicator = 1 household owns home, 0 otherwise.	61675	0.72 (0.45)
Life Cycle			
avg_age_all_adults	Average age (in years) of all adult members (age ≥ 18) in household	61675	55.72 (16.4)
avg_age_all_adults_2	Squared average age of all adults	61675	3374 (1913)

Notes. The sample consists of all household-year observations from the SHIW waves 2006–2020. Standard deviations are in parentheses. *risk_share* has a smaller N because it is defined only for households with positive total financial assets. Wealth quartile dummies sum to one; their means equal 0.25 by construction. The omitted reference categories are the first wealth quartile, not employed (labour market status), and single (marital status).

Table A2: Univariate analysis.

	Male	Female	Diff.	<i>t</i> -stat.
Risk holding	0.18	0.10	0.08***	(28.36)
Risk share	0.11	0.07	0.04***	(20.06)
FIS	0.17	0.73	-0.56***	(-249.36)
logy	10.32	10.01	0.31***	(53.68)
logy2	106.95	100.65	6.30***	(55.57)
employed_2	0.36	0.27	0.09***	(23.74)
selfemployed_2	0.14	0.05	0.09***	(37.82)
retired_2	0.43	0.26	0.17***	(44.84)
partner_employed_1	0.21	0.17	0.04***	(13.09)
partner_selfemployed	0.04	0.05	-0.01***	(-7.24)
avg_imp_risfin_adult	3.38	3.50	-0.12***	(-25.05)
avg_edu_all_adults	3.93	3.56	0.37***	(29.27)
quantile2wealth	0.23	0.28	-0.05***	(-13.94)
quantile3wealth	0.25	0.24	0.01*	(2.47)
quantile4wealth	0.30	0.18	0.12***	(35.31)
num_adult	2.27	1.86	0.40***	(52.60)
num_kid	0.38	0.31	0.07***	(11.49)
married	0.77	0.39	0.37***	(98.94)
divorced	0.05	0.12	-0.06***	(-27.71)
homeowner	0.74	0.68	0.06***	(17.33)
avg_age_all_adults	54.24	57.74	-3.50***	(-25.79)
avg_age_all_adults_2	3180.73	3639.60	-458.87***	(-28.88)
Observations	35,647	26,028	61,675	

Notes: The sample comprises all household-year observations from the SHIW waves 2006–2020. The table reports means for male-headed and female-headed households separately, the difference between them, and the two-sample *t*-statistic (unequal variances) in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.