Journal of Public Economics 239 (2024) 105213

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Short communication

Listen to her: Gender differences in information diffusion within the household ☆

Dietmar Fehr [a,b], Johanna Mollerstrom [c], Ricardo Perez-Truglia [d,*]

[a] University of Stuttgart, Germany
[b] CESifo, Germany
[c] George Mason University and Research Institute of Industrial Economics (IFN), United States of America
[d] University of California, Los Angeles, United States of America

ARTICLE INFO

JEL classification:
C93
D13
D83
J16

Keywords:
Household economics
Information diffusion
Gender
Survey
Experiment

ABSTRACT

We study how economic information diffuses within the household, leveraging an information-provision experiment with a representative sample of households from Germany. A random sample of household members received information about their household's position in the income distribution. When provided with information directly, there are no gender differences in how individuals update their beliefs. However, we observe significant gender disparities in the diffusion of information within the household. Specifically, when only the husband receives the information, it influences the wife's beliefs; however, when only the wife receives the information, it does not affect the husband's beliefs.

1. Introduction

About half a century ago, the United Kingdom (UK) government changed the allocation of subsidies for families with children, directing them towards mothers instead of fathers. According to the economic models of the time, such a policy should have no impact on families' behavior. The central tenet of these models was that the household functions as an integrated unit in which preferences are aligned and information is available to all members of the household (Samuelson, 1956; Becker, 1981). However, these basic assumptions have been criticized as unrealistic. Empirical evidence supports this criticism: In the UK, providing child allowance to mothers rather than fathers led to spending patterns more in line with the intention of the policy to cover necessities, such as clothing, for the family's children (Lundberg and Pollak, 1996; Lundberg et al., 1997; Ward-Batts, 2008). Subsequent results have corroborated the view that households do not necessarily function as an integrated unit with common preferences over monetary resources and that, relative to men's choices, women's spending choices are deemed to be more beneficial to the family's children (e.g., Dizon-Ross and Jayachandran, 2023). This is a prominent reason why cash transfer programs to the poor often target women as beneficiaries (Duflo, 2003; Almås et al., 2018; Armand et al., 2020; Field et al., 2021).

There has been a growing interest in economic research aimed at understanding how households function in the real world (e.g., Lundberg and Pollak, 1996; Ashraf, 2009; Chiappori and Mazzocco, 2017). The focus has been on how households manage resources such as goods and money. A highly relevant question that has received comparably little attention so far is how households manage *information*. Arguably, information is as crucial a resource as money, given that limitations on information accessible to spouses can impact their decision-making. The importance of information in intra-household decision-making has long been emphasized in the sociological literature (e.g., Dwyer and Bruce, 1988; Zelizer, 2005), while the common assumption in economics is that household members pool their information, in particular

☆ We thank the co-editor and two anonymous reviewers for their excellent feedback as well as Gautam Rao, Frank Schilbach, Basit Zafar, and seminar audiences in Berlin, Munich, Konstanz, and Tucson for valuable discussions. Petr Novak and Paul Schlowak provided superb research assistance. We are grateful to Bettina Zweck (Kantar Public Germany), David Richter (DIW Berlin), and Carsten Schroeder (DIW Berlin) for their support in implementing the project. This project received financial support from the German Research Foundation (DFG) through individual grant FE 1452/3-1 (Fehr) and from the German Institute for Economic Research (DIW Berlin, Mollerstrom). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

* Corresponding author.
E-mail addresses: dietmar.fehr@ivr.uni-stuttgart.de (D. Fehr), jmollers@gmu.edu (J. Mollerstrom), ricardo.truglia@anderson.ucla.edu (R. Perez-Truglia).

https://doi.org/10.1016/j.jpubeco.2024.105213
Received 30 July 2023; Received in revised form 31 May 2024; Accepted 26 August 2024
0047-2727/© 2024 Elsevier B.V. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

when interests align (Chiappori, 1992; Lundberg and Pollak, 1996). In this paper, we challenge this assumption and provide novel evidence on gender differences in how economic information diffuses within the household.

Studying diffusion of information within the household presents some empirical challenges: we need a setting in which we can both observe the two spouses independently and repeatedly in their natural environment and also manipulate decision-relevant information exogenously. For this purpose, we leverage existing data from a two-year survey experiment with a representative sample of Germans (Fehr et al., 2022).[1] Our survey revolves around perceptions about the relative position of the household in the income distribution. This setting is well suited for studying learning due to widely documented evidence that individuals have significant misperceptions about their relative income (Cruces et al., 2013; Engelhardt and Wagener, 2017; Karadja et al., 2017; Fehr et al., 2022) and because perceptions about relative income are important for households in natural settings. For example, perceived relative income has been shown to affect preferences for redistribution (Cruces et al., 2013; Engelhardt and Wagener, 2017; Karadja et al., 2017; Fehr et al., 2022), subjective well-being (Perez-Truglia, 2020), and a wide range of decisions such as where to live (Bottan and Perez-Truglia, 2022) and whether to change employers (Card et al., 2012; Cullen and Perez-Truglia, 2022).

In our baseline survey, we first elicited respondents' beliefs about their household rank on the national and global income scale in an incentivized way. All adult household members were interviewed by professional interviewers in private, without the possibility of communicating with each other, so respondents could not share any information during the baseline survey even if they wanted to. After eliciting the prior beliefs, half of the respondents received accurate information about their household's income rank. We randomized this information provision at the individual level to create variation within households. Thus, this resulted in households where both spouses, only the wife or husband, or nobody received the information, enabling us to explore how respondents acquire knowledge through direct information provision and indirectly through the diffusion of information within the household.

A year later, we conducted a follow-up survey with the same respondents, where we again asked incentivized questions about the household's income rank. Although there was no opportunity for spouses to communicate during the interviews, they had ample opportunity to talk about the household's income rank in the year that passed between the two survey waves, if they chose to do so. Importantly, we did not provide explicit incentives to share the information with other household members and did not inform the respondents that we would ask questions about relative income again a year later. As a result, information sharing evolved endogenously and naturally, with respondents freely choosing to share information with other household members or to refrain from doing so.

We start by documenting how individuals learn from information directly (i.e., when they receive it themselves). When respondents directly receive information on their true income ranks, the information has a significant and persistent effect on beliefs even after a year has passed. More importantly, men and women seem to incorporate the information to a similar degree when it is given directly to them. After one year, the learning rate is around 0.2 and does not differ statistically between women and men. More precisely, for each percentage point shock in the information given directly to a respondent, the perceived income rank as measured a year later is higher by about 0.22 percentage points for women and 0.16 for men.

In contrast, we find stark differences by gender in how information diffuses within the household, with a substantially lower pass-through of information from wives to husbands than vice versa. If husbands receive information about the true income rank directly, whereas their wives do not, we observe a pass-through to their wives' belief that is about as strong as if the wives received the information directly. However, if a wife receives the information directly but not their husband, we see no effect on her husband's belief. The gender difference in indirect learning rates (0.19 for women vs −0.01 for men) is large and statistically significant.

Our findings further indicate that this phenomenon is specific to the household context, as men do not disregard information received from women in general. We show that men are equally likely to incorporate information given to them directly by male versus female interviewers. However, we find little evidence that the observed gendered information flow within households is due to asymmetries in financial knowledge and experiences, different communication and information acquisition patterns of women and men, or gender differences in the interest in information about relative income. Instead, our results indicate a difference in the way husbands and wives update the information provided by each other.

We contribute to the emerging literature on information flows within households. The bulk of this literature is concerned with decision situations in which incentives are *non-aligned* and preferences differ, such as fertility decisions (Ashraf et al., 2014; Apedo-Amah et al., 2020; Ashraf et al., 2022).[2] The evidence from these experiments shows that information in such settings only sometimes flows freely and that information barriers can result in inefficient behavior (e.g., Ashraf, 2009; Ashraf et al., 2014, 2022). For example, Ashraf et al. (2022) conducted an information intervention in which they informed husbands or wives about the risks to maternal health. Consistent with our findings, they find that the information spills over from husbands to wives but not in the other direction. We contribute to this literature by studying a real-world situation in which incentives are *aligned*, which is arguably one of the more common settings in practice, yet one that has received little attention. A notable exception is a study by Conlon et al. (2022) that focuses, as we do, on a situation with *aligned* preferences. In their laboratory experiment with 400 married couples from Chennai, India, the husband or wife receives signals about the number of differently colored balls in an urn. They can pass on this information to their spouse, and the spouse can subsequently use it to make an optimal guess about the color of the ball that is drawn next. Despite explicit incentives to share this information and, consistent with our own findings, Conlon et al. (2022) document pronounced gender differences in the diffusion of information: Although wives took the information discovered by their husbands into full consideration, husbands did not do the same with the information revealed to their wives.

We complement the work of Conlon et al. (2022) in several important ways. First, unlike their stylized setting (participants received information by drawing balls from an urn), our setting is one of endogenous and naturally occurring information diffusion. Our subjects could naturally share the information in their daily lives over the span of a whole year, but we did not provide explicit incentives to do so. Second, rather than studying beliefs about an abstract object (the colors of balls from an urn), we study a belief that households arguably care about above and beyond the context of our experiment: their relative income. Third, the observed gender differences in information diffusion hold in very different cultural and economic contexts. For example, gender norms differ substantially between Germany and India. According to the World Values Survey, 52% of Indians agree with the statement

[1] In the original study, we measure how beliefs about relative income affect preferences for redistribution (Fehr et al., 2022). In this follow-up work, we further analyze the data to explore gender differences in information diffusion.

[2] More generally, there are some studies exploring gender differences in how information flows outside households (e.g., Beaman and Dillon, 2018; BenYishay et al., 2020; Cullen and Perez-Truglia, 2023).

that men should have more rights to a job than women if jobs are scarce, while only 15% of Germans agree with the same statement. In summary, while Conlon et al. (2022) have a more controlled setting, it is also more artificial. On the other hand, our work is set in a natural field setting, which comes with somewhat less control but arguably higher external validity. Taken together, these two studies paint a consistent picture that even in environments with aligned interests, gender barriers to information flow exist and are robust across different cultural and economic contexts.

2. Research design and data

We implemented two tailor-made survey modules in the Innovation Sample of the German Socio-Economic Panel (SOEP-IS). The SOEP-IS is a comprehensive longitudinal study that, once a year, surveys a representative sample of the German population on a wide range of topics. It is the ideal test-bed for our research question and offers several advantages over other survey modes: First, all household members over 16 years of age are interviewed by professional interviewers in computer-assisted interviews conducted in person. Second, we can follow up with little attrition a year later. Third, face-to-face interviews provide significant control, minimize non-response, and allow us to clarify misunderstandings instantly. Important for our purposes, they also prevent the search for information and communication between household members *during* and *between* the interviews *within a wave* because the interviews were conducted privately with each member of a household. Fourth, through the SOEP, we have access to a rich set of measures of socioeconomic indicators. Fifth, the SOEP team implements various safeguards to ensure high data quality, such as pre-testing new items and performing plausibility and consistency checks after data collection (for more details, see Goebel et al., 2019).

Baseline Survey: At the beginning of the baseline survey, each respondent stated their household income before taxes and the number of household members.[3] We used this information to explain and inform each respondent about their per-capita, pre-tax household income.[4] Subsequently, each respondent assessed their rank in the national (i.e., German) and the global income distribution based on their per-capita pre-tax income on a scale from 0 (poorest percentile) to 100 (richest percentile) in randomized order. We incentivized both income rank assessments to ensure that it was optimal for respondents to answer truthfully, and each assessment that was correct to the closest percentile was rewarded with €20. To prevent communication within the household and to avoid social desirability bias potentially impacting answers, respondents stated their ranks in private (i.e., without other household members or the interviewer seeing the tablet screen).

Approximately 10–15 min later, after respondents had answered several questions unrelated to our research, we randomized half of the respondents to a treatment, providing them with accurate information about their household's true income rank in the national and global income distributions. The information briefly explained the source of the information and then, based on the respondents' stated per capita, pre-tax household income, revealed the share of people who are poorer at the national and global levels.[5] The information was read out by the interviewer, who additionally visualized the information with customized graphs to facilitate understanding (see Appendix Figure A1 for a screenshot). The other half of the respondents received no information. Randomization was implemented at the individual level through the survey software and each respondent had an equal chance of receiving the information or not.

Follow-up Survey: One year later, we implemented our second survey module with the same sample of respondents. The setup of the follow-up survey was closely modeled after the baseline survey. That is, we first collected information on household income and the number of household members and explained the concept of per-capita household income. We then asked respondents to state their rank in the national and global income distributions in private. Again, we rewarded accurate predictions (this time, we paid €10 for each accurate prediction). The main difference to the baseline survey was that we did not provide information on the true income rank in either context in the follow-up survey. Instead, we elicited respondents' willingness to pay (WTP) for information about their true rank in the national and global income distributions using a list-price version of the Becker-DeGroot-Marschak method (Becker et al., 1964).[6] Finally, we asked *treated* respondents whether they had shared the information on the true income rank that they received in the baseline survey with anyone in the household during the past year, and we asked *all* respondents whether they, during this time, had looked for information about the distribution of national and global income.

Data: Our data contains the two survey modules that we implemented in the 2017 and 2018 waves of the SOEP-IS. A total of 1392 respondents took part in the baseline survey, while 1144 participated in the second survey (82% of the 1392 respondents in the baseline survey). We focus our analysis on single-member and two-person, mixed-gender households as explained in Section 3 below. This restriction results in a sample of 1164 respondents in the baseline survey and 989 respondents in the follow-up survey (85% of the 1164 respondents in the baseline survey). One potential concern is that the experiment's information provision could have affected the decision to participate in the follow-up survey. However, there is no significant difference in the attrition rates between the control and treatment neither in the *full sample* (17% vs. 19%, p-value = 0.392 for t-test of proportions) nor in the *restricted sample* (14% vs. 17%, p-value = 0.289 for t-test of proportions). In Appendix Tables A1–A6, we present several specifications showing that treatment status does not predict participation in the follow-up survey (for the restricted and full sample). Moreover, and as expected, the observable pre-treatment characteristics are balanced across treatment and control groups. Appendix Tables A7 present the results for the full and restricted sample and also split the samples by gender (for more details, see Section A.4 in the Appendix).

3. Empirical strategy

We want to estimate the direct and indirect impact of information provision on beliefs about income ranks one year later. We define T_i^{direct} as a treatment indicator variable, taking the value 1 if a respondent received direct information on their household income rank in the baseline survey and 0 otherwise. Similarly, $T_i^{indirect}$ is an indicator variable that takes the value 1 if the respondent did not receive the information directly, but another member of their household did, and

[3] Note that spouses should be well-informed about household income because the basis of income taxes of couples in Germany is their joint income. Moreover, evidence suggests that more than 75 percent of German couples pool their financial resources (Lott, 2017). We observe a similar share of couples who pool their financial resources, see Appendix A.7 for details.

[4] Note that estimates of the global income distribution are only available at the per-capita, pre-tax level.

[5] Consequently, the information provided could differ somewhat between members of the same household. We discuss this and the potential consequences thereof in Section 4.1.

[6] For both pieces of information, we presented five scenarios in which respondents had to decide between receiving information about their true rank in the income distribution and receiving a monetary reward that progressively increased from 10 cents to 10 euros. Respondents made their decision in private, and we informed them that one randomly selected decision for each piece of information (national and global) would be implemented. Possible payments and information provisions were made at the end of the survey.

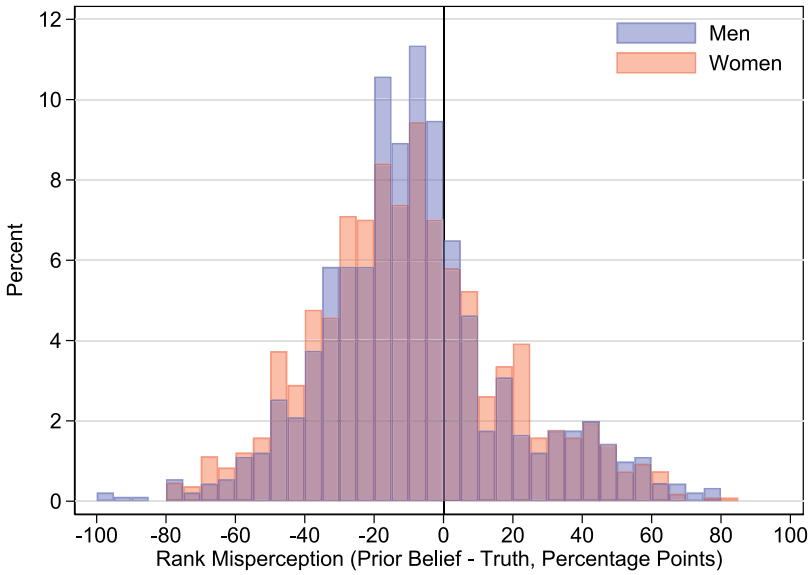

Fig. 1. Misperceptions of Income Ranks, by Gender.
Notes: Distribution of misperceptions about income rank in the baseline survey for female (red) and male respondents (blue). Misperceptions are calculated as the difference between prior beliefs about income rank and true income rank. Positive (negative) differences correspond to overestimation (underestimation) of own income rank. Data from baseline, i.e., before the respondent (or their spouse) actually received any information ($n = 1978$).

0 otherwise.[7] Let r_i^{prior} denote the perceived income rank in the baseline survey (i.e., the prior belief before receiving information) and r_i^{info} denote the information about the income rank that could be shown to the respondent. Consequently, $r_i^{info} - r_i^{prior}$ is the potential treatment: i.e., the misperception about the income rank. A positive difference indicates an underestimation, and a negative difference indicates an overestimation of the income rank. The direct information shock is given by $(r_i^{info} - r_i^{prior}) \cdot T_i^{direct}$, while the indirect information shock is given by $(r_i^{info} - r_i^{prior}) \cdot T_i^{indirect}$. Let $r_i^{posterior}$ be the posterior belief about the income rank in the follow-up survey. We then use the following specification to estimate direct and indirect learning rates.

$$r_i^{posterior} = \alpha^{direct}(r_i^{info} - r_i^{prior}) \cdot T_i^{direct} + \alpha^{indirect}(r_i^{info} - r_i^{prior}) \cdot T_i^{indirect} + \beta_1(r_i^{info} - r_i^{prior}) + \beta_2 X_i + \epsilon_i \quad (1)$$

The coefficients α^{direct} and $\alpha^{indirect}$ tell us how correcting misperceptions — directly or indirectly through information provided to the spouse — affect beliefs one year later. The parameter α^{direct} measures the direct learning rate, i.e., the effect of an additional percentage point of information shock given directly to individual i on their posterior belief. The parameter $\alpha^{indirect}$ measures the indirect rate of learning, i.e., the rate of pass-through between the information provided to respondent i's spouse and respondent i's belief one year later. X_i is a vector of control variables that include the demographic characteristics of the respondent and the household. We estimate Eq. (1) separately for female and male respondents and cluster standard errors at the household level.

For our baseline specification, we restrict our sample to single-member households and households consisting of two adult partners, that is, husbands and wives ($n = 989$). We include single-member households to strengthen statistical power in the analysis of direct learning. We exclude households in which other adult household members in addition to the spouses were interviewed, to avoid dealing with cases in which information can be transmitted from multiple household members (e.g., adult children, grandparents). We further restrict the sample to mixed-gender partners — same-sex households are a negligible share of the sample, and thus we do not have enough data to study them separately. Finally, we observe beliefs about each respondent's income rank at the national and global levels. In the analysis, we pool these two responses, as differentiating between the two belief statements is inessential for our purposes. This gives us two income-rank observations for each respondent, resulting in a total of $n = 1,978$ observations. In Appendix Section A.6, we show that our results are not sensitive to any of the specification choices listed above.

4. Results

4.1. Misperceptions about income ranks

Misperceptions about their own household income rank are common among both women and men. Fig. 1 shows the distribution of misperceptions (measured as perceived minus the actual percentile) at baseline, separated by gender; the difference in the distribution of misperceptions between women and men is statistically insignificant (Kolmogorov–Smirnov test, p-value = 0.126). For example, women underestimate their rank by approximately 9 percentage points, on average, and men by approximately 10 percentage points, a difference that is small and statistically insignificant (p-value = 0.411).

Next, we compare perceptions within two-person households, i.e., between husbands and wives. Panel (a) of Fig. 2 shows a binned scatterplot of misperceptions about income rank, with the wives on the y-axis and their husbands on the x-axis. If husbands and wives have similar levels of misperception regarding their households' income rank, their misperceptions would align along the 45-degree line. However, misperceptions do not align perfectly in this way, suggesting a significant disagreement about the income ranks between spouses. Although rank misperceptions within a household are correlated, the correlation is far from perfect ($\rho = 0.55$). In other words, husbands and wives tend to harbor rather different misperceptions.

A potential concern is that differences in misperceptions about *relative* income are a mechanic result of disagreements about *absolute* income. To address this concern, panel (b) of Fig. 2 shows a binned scatterplot of the stated household income for the wives (y-axis) and

[7] The implicit assumption here is that for directly informed respondents it should not matter whether another household member received information or not. We provide a direct test of this assumption in Appendix A.5.

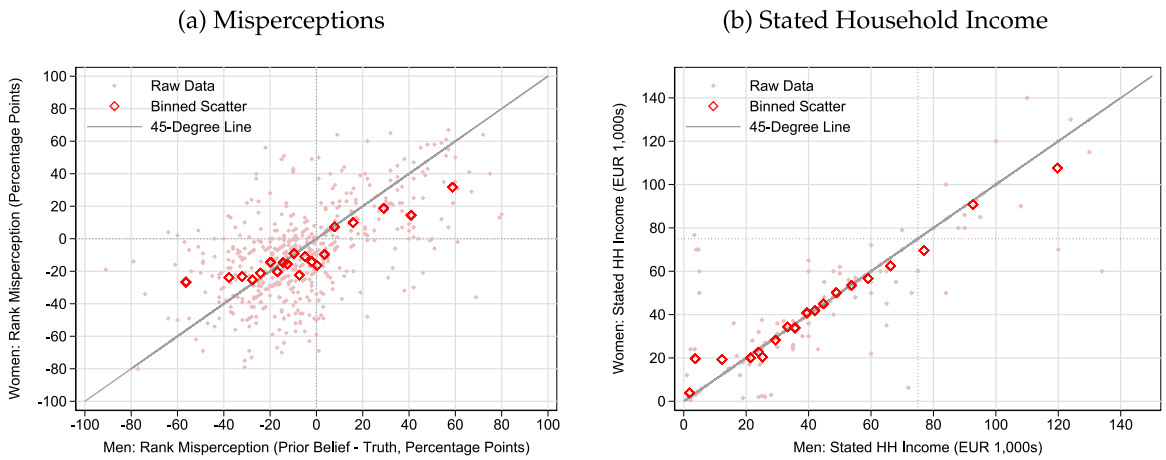

Fig. 2. Misperceptions of Income Ranks and Stated Household Income within Households.
Notes: **Panel (a)** shows the correlation between misperceptions about the income rank of husbands and wives (within the household), and **panel (b)** shows the correlation of the stated household income of husbands and wives (within the household). Misperceptions are calculated as the difference between prior beliefs about income rank and true income rank. Stated household income is the yearly gross household income measured in 1000 euros. Both figures show scatter plots of the raw data (light red) and binned scatterplots (red diamonds). For the binned scatterplot, we group the variables on the x-axis into 20 equally sized bins and calculate the mean of the x and y variables within each bin. Both figures use data from the baseline survey, and we restrict the sample to two-person, mixed-gender households ($n = 536$).

their husbands (x-axis). In contrast to misperceptions of relative income, the stated household incomes line up almost perfectly on the 45-degree line, with a correlation coefficient of $\rho = 0.95$.[8] This suggests that spouses largely agree about their absolute household income and that misperceptions about relative income cannot be attributed to disagreement about absolute income.

4.2. Direct and indirect effects of information on posterior beliefs

Fig. 3 presents coefficient plots of our main result (for the corresponding regression results in table form, see Appendix Table A10). The effect of indirect information diffusion is entirely driven by women whose husbands directly received information about the true income ranks of the household. Panel (a) of Fig. 3 reveals that for each additional percentage point in the direct information shock, men updated their posterior belief by 0.16 and women by 0.22 percentage points. Importantly, the difference between these two estimates (0.16 and 0.22) is not only small but also statistically insignificant (p-value = 0.391).[9] The observed direct learning rates are sizable, considering that we measure the posteriors about a year later. Generally, the learning rate should be lower than the perfect pass-through rate (i.e., $\alpha < 1$), even if measured immediately after the information provision. First, from a Bayesian perspective, respondents form posterior beliefs by taking a weighted average between the signal provided to them and their prior beliefs. Thus, if respondents find the information untrustworthy or feel very confident about their prior beliefs, they should update only partially. Second, when the posterior beliefs are elicited months later, the effects of information can be diluted because subjects forget the information provided in the experiment or incorporate new information. In fact, evidence shows that the effect of information on beliefs can decrease substantially even over the course of a few months (e.g., Cavallo et al., 2017; Bottan and Perez-Truglia, 2022). Therefore, a substantial dilution would be expected a full year after the information was provided.

Although there is no difference in how men and women treat information that was revealed to them directly, the information provided to their spouses generates a different picture. When a wife received information about the actual income ranks through her husband, the effect on her posterior belief one year later was substantial (0.19 percentage points, p-value = 0.010) and almost as strong as if the information were provided directly. In stark contrast, when a husband was not directly informed about the true household income rank but his wife was, he did not adjust his beliefs one year later (−0.01 percentage points, p-value = 0.906). The difference in indirect learning rates between wives and husbands is both sizable (0.19 vs. −0.01) and statistically significant (p-value = 0.040).

Panel (b) of Fig. 3 presents a falsification test to probe the robustness of these results. We measure the effect of direct and indirect information provision on prior beliefs about household income rank. Given that we elicited these beliefs before the information experiment, we expect to observe no effect of the information on these prior beliefs. This is exactly what we find: the direct and indirect placebo learning rate is close to zero, statistically insignificant, and precisely estimated in all specifications.[10] In Appendix Section A.6, we further show that our results are robust to not pooling the beliefs about national and global income ranks (Appendix Table A12), using the full sample (Appendix Table A13), and focusing only on two-person households (Appendix Table A14).

The evidence presented consistently points to pronounced gender differences in information diffusion. Our preferred interpretation of these findings is that wives are more likely to incorporate the information shared by their husbands than husbands are to incorporate the information shared by their wives. This interpretation is consistent with the findings of Conlon et al. (2022), who designed a laboratory experiment to unravel this mechanism. In our field setting, controlling and observing how household members share information is more difficult, so it is more challenging to rule out alternative stories. However, we next provide evidence against some of these alternative explanations.

We start by examining whether the observed gender differences occur beyond the household context. To do so, we leverage the assignment of female or male interviewers to households over which households have no control. The interviewers read out the information on income ranks and showed the respondent a customized graph on a

[8] About 62 percent of couples perfectly agree on their household income, and for 78 percent, the disagreement is less than 5,000 euros.

[9] To test for the difference in learning rates across gender, we present estimates from interacting all relevant variables with gender in Appendix Table A11.

[10] The results from this falsification test are also presented in Appendix Table A10. The top and bottom panels correspond to the same regression, but while the posterior beliefs are the dependent variable in the top panel, the prior beliefs are the dependent variable in the bottom panel.

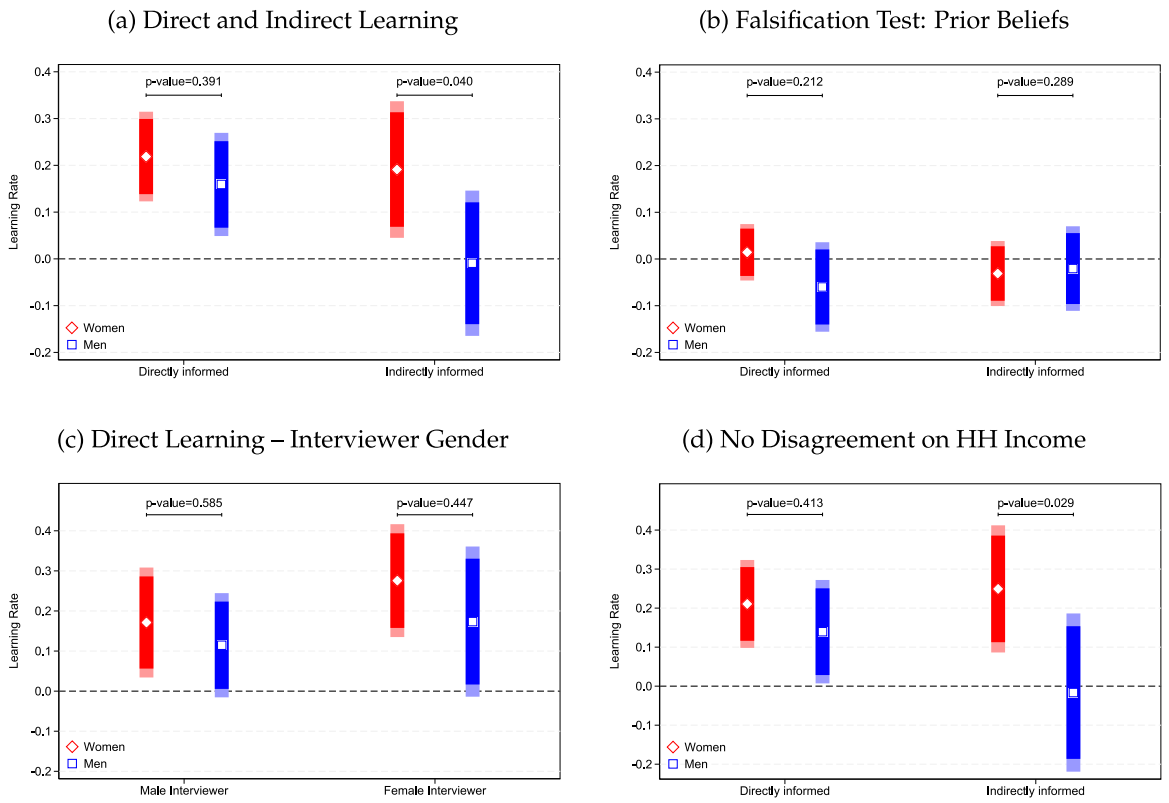

Fig. 3. Direct and Indirect Learning from the Information Shocks.
Notes: Coefficient plots of learning rates from OLS regressions estimating the effect of information provision on beliefs about income rank as outlined in Eq. (1) in Section 3. The sample is restricted to single-member and two-person, mixed-gender households, and standard errors are clustered at the household level. The bands around the coefficient estimates indicate 90% (light color) and 95% (intense color) confidence intervals. **Panel (a)** shows the effect of providing direct information to a respondent (α^{direct}) or indirect information through a respondent's partner ($\alpha^{indirect}$) on this respondent's beliefs about income rank one year after the intervention (posteriors). We estimate (α^{direct}) and ($\alpha^{indirect}$) separately for women, shown in red, and men, displayed in blue. **Panel (b)** shows a falsification test from estimating Eq. (1) for women and men using beliefs about income rank in the same year (prior beliefs). **Panel (c)** shows the effect of providing direct information to a woman or man (α^{direct}) on their beliefs about the income rank one year after the intervention (posterior) by the gender of the interviewer. **Panel (d)** replicates **Panel (a)** but restricts two-person, mixed-gender households to households that state exactly the same household income.

tablet visualizing the information treatment (as discussed in Section 2 above).[11] Panel (c) of Fig. 3 provides suggestive evidence that there is no difference in the reaction of men: They update their beliefs in a similar fashion regardless of whether the interviewer is female or male (0.17 vs. 0.11). The difference between the two coefficients is small and statistically insignificant (p-value = 0.674; see also Appendix Table A15). This suggests that our findings are more likely the result of within-household dynamics than of a more general phenomenon where men neglect to incorporate information they receive from any woman.

Next, we explore whether observed gender differences are specific to spousal relationship patterns. First, men typically contribute more to household income than women, particularly in our setting.[12] A common perception in this context is that the spouse who contributes more is better informed about household income and therefore their views on financial matters carry more weight. Although it is possible that this mechanism is at work in our setting, it is unlikely to fully explain our results. We observe similar results when we focus our analysis on samples in which this mechanism should not operate. First, we see little disagreement between spouses about absolute household income (see Fig. 2b). More importantly, if we focus only on households where both partners report exactly the same absolute household income, we still see the same gender difference in indirect learning. Panel (d) of Fig. 3 shows that the indirect learning rate of women, in this case, is 0.25 (relative to −0.02 for men). Second, linking our survey module to a previous SOEP-IS wave with questions on financial decision making among couples provides further suggestive evidence against asymmetric knowledge, financial views, and financial decision-making power within households. Restricting the analysis to couples where (i) both say that they decide equally on financial matters (Table A17) and (ii) both say that they pool their income (Table A18), we again see that women, but not men, learn indirectly from their spouses.

Second, another natural channel that could help explain our findings is gender differences in communication patterns. For example, if a wife does not communicate the information, her husband would be unable to learn from her, or if men are more interested in the topic of income ranks or, more generally, in financial matters, they may be more likely to share the information with their spouses. Although we cannot completely exclude these potential differences in communication patterns, there is evidence against their significance. The most direct evidence on communication patterns uses a follow-up survey question about whether directly informed respondents shared the income-rank information with other household members after the baseline survey. These data are, of course, merely a proxy for information sharing. Importantly, responses are likely to be subject to substantial recall bias. When individuals are asked whether they did something a year ago, their ability to recall these events is far from perfect. Thus, we expect these responses to systematically underestimate the share of individuals who respond affirmatively (see e.g.,

[11] The gender composition of interviewers is roughly balanced — 55% male vs. 45% female interviewers.

[12] In our sample, 83% of employed men work full time, while the corresponding share for women is only 47%.

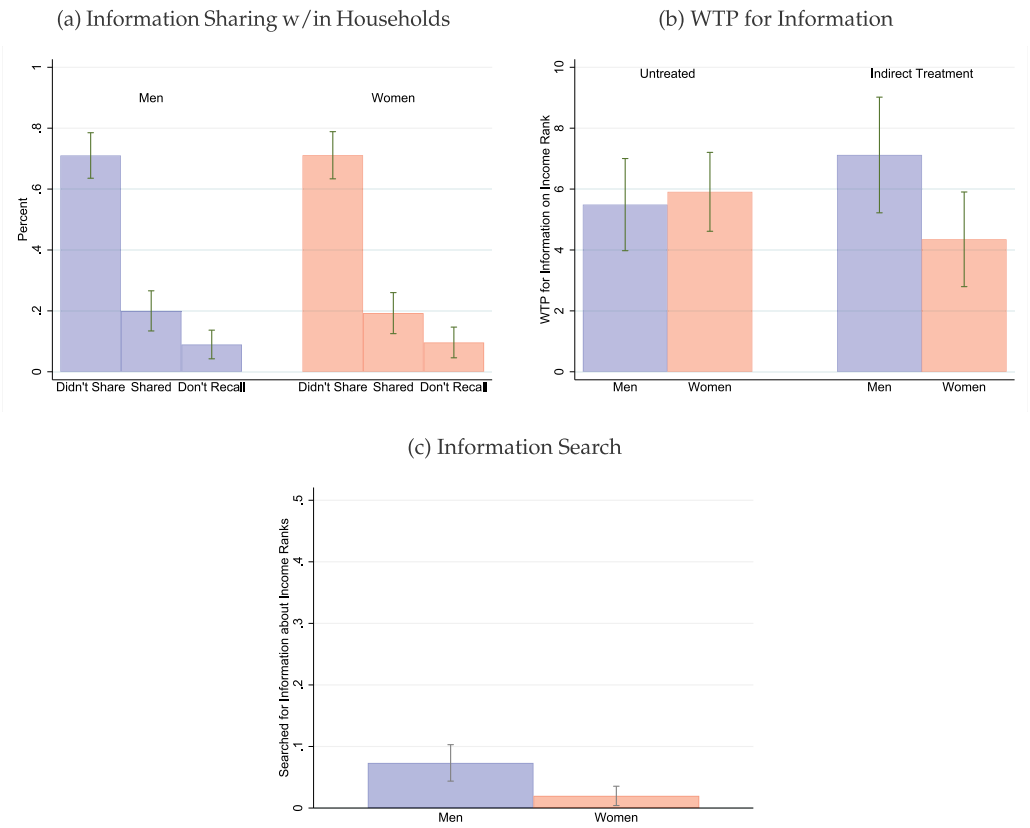

Fig. 4. Information Search, Information Sharing and Willingness to Pay for Information.
Notes: **Panel (a)** shows the fraction of treated women and men in two-person households who say that they (i) did not share their income rank information within their household after the baseline survey, (ii) shared this information, and (iii) did not recall information sharing ($n = 280$, Fisher's exact test, p-value = 1.000). **Panel (b)** shows the average WTP for rank information for untreated women and men in two-person households ($n = 290$, p-value = 0.607) and indirectly treated women and men ($n = 300$, p-value = 0.027). **Panel (c)** shows the share of women and men in two-person households who said they searched for information about income ranks after the baseline survey ($n = 604$).

Schacter, 1999; Bound et al., 2001).[13] However, we note that sharing the experimental information – and being aware that one has done so – is not a necessary condition for social learning as spouses may discuss the issue more generally and, for example, share their (updated) beliefs. With those caveats in mind, panel (a) of Fig. 4 shows that a non-negligible share of respondents said they shared the information within the household. Most importantly, we do not find any evidence that wives and husbands differ in the propensity to share information: 21% vs. 22% (p-value = 0.899, test of proportions).

We also provide two pieces of evidence that men and women have similar levels of interest in relative income information. First, the data on prior beliefs documents a small and insignificant gender difference in the average misperception (husbands 21.2 percentage points vs. wives 22.4 percentage points, p-value = 0.239). This evidence suggests that, prior to our baseline survey, husbands and wives had acquired similar levels of information on income ranks. Second, using data from the follow-up survey on the willingness to pay for information indicates that there is no gender difference in the interest in information about relative income. The average WTP for information was 6 euros for the national and global ranks each. This is substantial given that the maximum WTP is 10 euros and is also high compared to other studies that elicit WTP for other types of information (e.g., Khattak et al., 2003; Angulo et al., 2005; Allcott and Kessler, 2019; Fuster et al., 2022). Although it is possible that part of this demand for information is introduced artificially through our experiment, there are reasons to believe that the respondents can be genuinely interested in the topic due to its instrumental value.[14] We find no evidence that men are more interested than women in information on relative income. Panel (b) of Fig. 4 compares the average WTP for information between men and women. Looking at respondents in untreated households (i.e., households in which no information was received), we see that the WTP for information does not differ much between women and men (5.9 euros vs. 5.5 euros, p-value = 0.607). If we look at uninformed respondents in households with an informed member, we see that uninformed women have a significantly lower WTP than uninformed men (4.4 euros vs. 7.1 euros, p-value = 0.027). This supports our main result: women likely have a lower WTP because they already received it from their husbands, whereas men pay substantial amounts for information that they could have learned from their wives.

Finally, given the long time span between the baseline and follow-up surveys, one potential concern is that the respondents may have

[13] In addition to the recall bias, respondents may have been reluctant to admit sharing the information because the interview protocol was rather strict in preventing communication *during* the interview, so the respondents may have worried that they were also not supposed to share the information after the interview.

[14] In both the baseline and follow-up surveys, we incentivized the assessment of income ranks. Although we did not tell respondents that we would elicit this information in the following year, some respondents may nevertheless expect this opportunity and thus express interest in the information. On the other hand, people may be genuinely interested in this information, for example, when it becomes publicly available due to transparency policies (Perez-Truglia, 2020). Other evidence suggests that employees are interested in learning about the salaries of their peers and that this information has a significant impact, e.g., on whether to stay with a company (Card et al., 2012; Cullen and Perez-Truglia, 2022, 2023).

obtained information about income ranks from sources other than their spouses. To address this concern, we draw on a follow-up survey question on information search about relative income in-between the baseline and follow-up surveys. Panel (c) of Fig. 4 shows that only a small share of respondents in two-person households reported having searched for rank information on their own (2% of women and 7% of men). Focusing on indirectly treated respondents, we observe, however, no significant differences (Fisher's exact test, p-value = 0.114). Thus, seeking information from other sources is unlikely to significantly explain the gender differences in information diffusion.

In summary, these exercises suggest that the observed gender difference in information diffusion is unlikely to result from asymmetric financial knowledge and experience, differences in information-sharing and acquisition patterns between women and men, or differential interest in information about income ranks. Instead, our preferred interpretation is that, relative to men, women are more prone to incorporate information from their spouses.

5. Conclusions

Our study documents gender-specific barriers to information flow within households in a naturally occurring setting in a representative sample of Germans. We shed light on the boundaries and underlying reasons for these barriers, but some important questions remain open. First, gender stereotypes, such as that "men have to earn more than their wives" (e.g., Kamenica et al., 2015), can also play a role in our context. A natural starting point to address this possible issue would be to examine whether the results are similar in more female-dominated domains. Second, our study focuses on one important aspect of household decision-making – perceived relative income – but extending the examination beyond these beliefs, for example, to inflation expectations, effectiveness and safety of vaccines, etc., and exploring other contexts, including different developed and developing countries, is necessary to get a more complete picture of information diffusion within households. Finally, it is possible that women are ineffective in communicating their knowledge to men (see e.g., Bjorkman Nyqvist et al., 2024), so it would be fruitful to explore the communication patterns between women and men more thoroughly.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All the code has been posted online, and the data can be easily accessed online as well. See the online supplementary material for more details.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jpubeco.2024.105213.

References

Allcott, H., Kessler, J.B., 2019. The welfare effects of nudges: A case study of energy use social comparisons. Am. Econ. J. Appl. Econ. 11 (1), 236–276.

Almås, I., Armand, A., Attanasio, O., Carneiro, P., 2018. Measuring and changing control: Women's empowerment and targeted transfers. Econ. J. 128 (612), F609–F639.

Angulo, A.M., Gil, J.M., Tamburo, L., 2005. Food safety and consumers' willingness to pay for labelled beef in Spain. J. Food Prod. Mark. 11 (3), 89–105.

Apedo-Amah, M., Djebbari, H., Ziparo, R., 2020. Gender, information and the efficiency of household production decisions: An experiment in rural Togo. IZA Discussion Paper.

Armand, A., Attanasio, O., Carnerio, P., Lechene, V., 2020. The effect of gender-targeted conditional cash transfers on household expenditures: Evidence from a randomized experiment. Econom. J. 130 (631), 1875–1897.

Ashraf, N., 2009. Spousal control and intra–household decision making: An experimental study in the Philippines. Amer. Econ. Rev. 99 (4), 1245–1277.

Ashraf, N., Field, E., Lee, J., 2014. Household bargaining and excess fertility: An experimental study in Zambia. Amer. Econ. Rev. 104 (7), 2210–2237.

Ashraf, N., Field, E., Voena, A., Ziparo, R., 2022. Gendered spheres of learning and household decision making over fertility. Working Paper 28220.

Beaman, L., Dillon, A., 2018. Diffusion of agricultural information within social networks: Evidence on gender inequalities from Mali. J. Dev. Econ. 133, 147–161.

Becker, G.S., 1981. A Treatise on the Family. Harvard University Press, Cambridge, MA.

Becker, G.M., DeGroot, M.H., Marschak, J., 1964. Measuring Utility by a Single-Response Sequential Method. Behav. Sci. 9 (3), 226–232.

BenYishay, A., Jones, M., Kondylis, F., Mobarak, A.M., 2020. Gender gaps in technology diffusion. J. Dev. Econ. 143, 102380.

Bjorkman Nyqvist, M., Jayachandran, S., Zipfel, C., 2024. A mother's voice: Impacts of spousal communication training on child health investments. J. Dev. Econ. 168, 103263.

Bottan, N., Perez-Truglia, R., 2022. Choosing your pond: Location choices and relative income. Rev. Econ. Stat. 104 (5), 1010–1027.

Bound, J., Brown, C., Mathiowetz, N., 2001. In: Heckman, J.J., Leamer, E. (Eds.), Chapter 59 - Measurement Error in Survey Data. In: Handbook of Econometrics, vol. 5, Elsevier, pp. 3705–3843.

Card, D., Mas, A., Moretti, E., Saez, E., 2012. Inequality at work: The effect of peer salaries on job satisfaction. Amer. Econ. Rev. 102 (6), 2981–3003.

Cavallo, A., Cruces, G., Perez-Truglia, R., 2017. Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. Am. Econ. J.: Macroecon. 9 (3), 1–35.

Chiappori, P.-A., 1992. Collective labor supply and welfare. J. Polit. Econ. 100 (3), 437–467.

Chiappori, P.-A., Mazzocco, M., 2017. Static and intertemporal household decisions. J. Econ. Lit. 55 (3), 985–1045.

Conlon, J., Mani, M., Rao, G., Ridley, M.W., Schilbach, F., 2022. Learning in the household. Working Paper.

Cruces, G., Perez-Truglia, R., Tetaz, M., 2013. Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. J. Public Econ. 98, 100–112.

Cullen, Z., Perez-Truglia, R., 2022. How much does your boss make? The effects of salary comparisons. J. Polit. Econ. 130, 766–822.

Cullen, Z., Perez-Truglia, R., 2023. The salary taboo: Privacy norms and the diffusion of information. J. Public Econ. 222, 104890, URL https://www.sciencedirect.com/science/article/pii/S0047272723000725.

Dizon-Ross, R., Jayachandran, S., 2023. Detecting mother-father differences in spending on children: A new approach using willingness-to-pay elicitation. Am. Econ. Rev. Insights 5, 445–459.

Duflo, E., 2003. Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in South Africa. World Bank Econ. Rev. 17 (1), 1–25.

Dwyer, D., Bruce, J., 1988. A Home Divided: Women and Income in the Third World. Stanford University Press, Palo Alto, CA.

Engelhardt, C., Wagener, A., 2017. What do germans think and know about income inequality? A survey experiment. Socio-Econ. Rev. 16, 743–767.

Fehr, D., Mollerstrom, J., Perez-Truglia, R., 2022. Your place in the world – Relative income and global inequality. Am. Econ. J. Econ. Policy 14 (4), 232–268.

Field, E., Pande, R., Rigol, N., Schaner, S., Troyer Moore, C., 2021. On her own account: How strengthening women's financial control impacts labor supply and gender norms. Amer. Econ. Rev. 111 (7), 2342–2375.

Fuster, A., Perez-Truglia, R., Wiederholt, M., Zafar, B., 2022. Expectations with endogenous information acquisition: An experimental investigation. Rev. Econ. Stat. 104 (5), 1059–1078.

Goebel, J., Grabka, M., Liebig, S., Kroh, M., Richter, D., Schroeder, C., Schupp, J., 2019. The German socio-economic panel (SOEP). Jahrbücher für Nationalökonomie und Statistik 239 (2), 345–360.

Kamenica, E., Bertrand, M., Pan, J., 2015. Gender identity and relative income within households. Q. J. Econ. 130 (2), 571–614.

Karadja, M., Mollerstrom, J., Seim, D., 2017. Richer (and holier) than thou? The effect of relative income improvements on demand for redistribution. Rev. Econ. Stat. 99, 201–212.

Khattak, A.J., Yim, Y., Prokopy, L.S., 2003. Willingness to pay for travel information. Transp. Res. C 11 (2), 137–159.

Lott, Y., 2017. When my money becomes our money: Changes in couples' money management. Soc. Policy Soc. 16 (2), 199–218.

Lundberg, S.J., Pollak, R.A., 1996. Bargaining and distribution in marriage. J. Econ. Perspect. 10 (4), 139–158.

Lundberg, S.J., Pollak, R.A., Wales, T.J., 1997. Do husbands and wives pool their resources? Evidence from the UK. Child benefit. J. Hum. Resour. 32 (3), 463–480.

Perez-Truglia, R., 2020. The effects of income transparency on well-being: Evidence from a natural experiment. Amer. Econ. Rev. 110 (4), 1019–1054.
Samuelson, P.A., 1956. Social indifference curves. Q. J. Econ. 70 (1), 1–22.
Schacter, D.L., 1999. The seven sins of memory: insights from psychology and cognitive neuroscience. Am. Psychol. 54 (3), 182.
Ward-Batts, J., 2008. Out of the wallet and into the purse: Modeling family expenditures to test income pooling. J. Hum. Resour. 43 (2), 325–351.
Zelizer, V.A., 2005. The Purchase of Intimacy. Princeton University Press, Princeton, NJ.