

NBER WORKING PAPER SERIES

LISTEN TO HER: GENDER DIFFERENCES IN
INFORMATION DIFFUSION WITHIN THE HOUSEHOLD

Dietmar Fehr
Johanna Mollerstrom
Ricardo Perez-Truglia

Working Paper 30513
<http://www.nber.org/papers/w30513>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2022

We thank Petr Novak and Paul Schlowak for 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Dietmar Fehr, Johanna Mollerstrom, and Ricardo Perez-Truglia. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Listen to Her: Gender Differences in Information Diffusion within the Household
Dietmar Fehr, Johanna Mollerstrom, and Ricardo Perez-Truglia
NBER Working Paper No. 30513
September 2022
JEL No. D1,D83

ABSTRACT

Efficient diffusion of economic information plays a critical role in the functioning of society, and, more specifically, households. We study information diffusion between spouses in a representative sample of the German population. We focus on an important economic belief: the household's perceived rank in the income distribution. Our survey experiment consists of two waves. During each wave, all adult members of a household are interviewed separately with no possibility to communicate with each other. In the first wave, we randomly select a subset of respondents to receive accurate information about their household's income rank. By chance, some members of a household, but not others, receive the information. A year later, we re-survey all members of the same households in the second wave, with the aim of measuring whether the information provided in the experiment had a long-lasting effect on their beliefs. We find that receiving information directly persistently affected household members' beliefs. This direct learning worked similar for men and women. By contrast, for household members who did not receive information directly, we find striking gender differences in indirect learning. When information was provided to the husband, it affected his wife's beliefs to a similar degree as if she had directly received the information herself. By contrast, when the information was provided to the wife, it did not affect her husband's beliefs.

Dietmar Fehr
University of Heidelberg
Bergheimer Str 58
Heidelberg 69115
Germany
d7fehr@gmail.com

Ricardo Perez-Truglia
Haas School of Business
University of California, Berkeley
545 Student Services Building #1900
Berkeley, CA 94720-1900
and NBER
ricardotruglia@berkeley.edu

Johanna Mollerstrom
Vernon Smith Hall 5028
George Mason University
3434 Washington Blvd
Arlington, VA 22201
jmollers@gmu.edu

The online appendix and replication material can be found in the following URL:
<http://doi.org/10.17605/OSF.IO/3A9C5>

Introduction

The social sciences have a longstanding interest in understanding how information diffuses in society. The reason is that efficient information flow helps reduce the likelihood that economic agents will harbor misperceptions of economic realities when deciding for example which economic opportunities to pursue. The household is an area of particular interest for information flows, as it is an arena where many imperative economic decisions are made [1]. In simplified economic models of the household, it is generally assumed that the household functions as an integrated unit, meaning preferences are commonly aligned, and resources—both monetary and informational—that are made available to one spouse are in practice available to both [1-3]. This leads to the convenient prediction that it does not matter which spouse receives money or information, as the subsequent actions of the household will not change.

This assumption was challenged relatively early on with respect to monetary resources [4-7]. For example, exploiting a policy change implemented in the 1970s in the United Kingdom (UK), researchers documented how the child allowance was spent very differently when the money was directed to the mother, rather than the father [1,8-9]. This indicates that households, with respect to monetary resources, do not function as an integrated unit with common preferences. The findings from the UK suggested that providing the child allowance to mothers rather than fathers helped keep spending more in line with the policy's intention of covering necessities, like clothing, for the family's children. Subsequent results have corroborated this, and it is one of the reasons that aid projects in developing countries are increasingly being designed with women as the target recipients [10-11].

Arguably, information is at least as important a resource as money. However, it took some time for the idea of perfect information flow (and resulting uniform beliefs) within households to become a subject of experimental scrutiny. Experiments in this area that have been conducted have generally focused on settings where incentives are not aligned, *e.g.*, where one spouse has a reason to keep certain information hidden from their partner. It has been shown that information in such settings does not always flow freely, and that this can result in inefficient behavior [12-14].

We investigate information flows between spouses in circumstances where incentives are aligned. While this question is important due to the fact that many household decisions have a common objective, empirical evidence in this setting is sparse. In a social learning experiment conducted with spouses in Chennai, India, individuals were either informed directly about facts needed to maximize earnings or had to learn from their spouse what they discovered [15-16]. In this highly stylized setting (participants received information through drawing balls from an urn), participants had an explicit incentive to share information. The findings indicate that women treat their own and their husband's information as equally valuable. That is, women use both pieces of information to

help maximize earnings. However, this is not true for husbands; their beliefs respond less to information that was discovered by their wives than information they revealed themselves. This does not seem to result from their wives' lack of attempted communication, as even when husbands had perfect information about what their wives learned, they discounted it heavily.

We study a representative sample of a Western population—Germany—in a setting where the decision to share information (or not) is natural, rather than an explicit part of the experiment. Here, the information is not abstract but related to a feature that is central to the household's everyday economic behavior: the own household's rank in the income distribution. In a separate paper, we use these data to investigate how beliefs about income ranks relate to preferences for national and global redistribution [17]. The data can also be used to investigate gender differences in information diffusion within the household, which we do here.

In our first survey (baseline), we elicited respondents' beliefs about their household's rank on the national and global income scale (respondents were incentivized to answer correctly). Household income rank is believed to be a key determinant behind a person's preferences for redistribution in society [18-19], but is often misperceived by individuals [17, 20-21]. Later in the first survey, half of the respondents were randomly selected to receive truthful information about their household's actual rank in both income distributions.

A year later, we conducted a second survey where we again asked incentivized questions about the household's income rank at the national and global level. For both surveys, all members of a household over the age of 16 were interviewed, *i.e.*, in the case of a married couple, and for couples living together in a marriage-like relationship, both spouses responded individually to the survey. During the interview for the first survey, household members could not communicate with each other (this was strictly enforced by the trained interviewers who conducted the interview in person). Also, the treatment (information on the household's actual income rank) was randomly allocated at the individual level, so some household members were directly informed and others were not. However, in the year that passed between the two surveys, spouses had ample opportunity to discuss the information (importantly, however, they were neither encouraged to, nor dissuaded from, doing so.). By measuring the beliefs of household income rank in the second survey, we can therefore investigate how the information has diffused (or not) throughout the household during the course of a year. Our setting is thus one of endogenously occurring information sharing, without incentives attached—a setting which is arguably very common in everyday life for most households.

Results

We surveyed $n=1,144$ people in $N=776$ households, but restrict the sample analyzed here to households with two adult partners (regardless of whether they are married or not) and single households. The resulting sample consists of $n=989$ respondents and $n=1,978$ person-income rank

observations. We include single households in the baseline specification to maximize statistical power: while indirect learning is not possible for single households, these still help us identify the degree of direct learning. Nevertheless, as discussed in the Supplementary Information (SI), our results are robust to excluding single households.

We start by observing that misperceptions of household income rank are common. Figure 1 shows the distribution of misperceptions (measured as perceived minus actual percentile) at baseline, by gender. Both females and males harbor misperceptions, and these do not vary significantly between gender in sign or size ($M_{female} = -0.100, SE = 0.008, M_{male} = -0.089, SE = 0.009$, t-test of differences in means: $p = 0.41$). We can thus conclude that the information provided by our treatment has the potential to move respondents' beliefs about their household income rank closer to the true rank. Figure 2 shows binned scatterplots of misperceptions about income rank (Panel A) and stated absolute household income within households (Panel B). While the rank misperceptions within a household are correlated, the correlation is far from perfect ($\rho = 0.52$) and substantially lower than the correlation of stated household income ($\rho = 0.96$). This is, in itself, an indication that beliefs about income rank are not homogenous in the household. Rather, there is, on average, some disagreement between spouses about their household's income rank. This is not the case for stated household income.

How much does new information about income ranks diffuse within households? All members of the surveyed households were re-surveyed a year later, and each household member was again interviewed separately without the ability to communicate with their partner. We find that informing a respondent directly about their household's true income rank does move perceptions and increase accuracy. The learning rate for direct information is 0.164 ($SE = 0.034, p < 0.01$), which indicates that for each percentage point shock in information given directly to the respondent, the perceived income rank a year later is higher by about 0.16 percentage points.

Panel A of Figure 3 illustrates that the coefficient estimate for the direct learning rate is similar across gender. Both men and women updated their posterior beliefs by about 0.16-0.17 percentage points ($p = 0.93$) when they were themselves directly informed about the true household income rank. While there is no difference in how men and women treat information that was revealed to them directly, information provided to their partners shows a very different picture. When a woman received the information indirectly, the effect on her belief about income rank one year later was at least as strong as if she was directly informed (0.19 percentage points, $p = 0.01$). In contrast, when a man was indirectly informed about the true household income rank through his female partner, he did not adjust his beliefs one year later (-0.01 percentage points, $p = 0.91$). The difference in indirect learning rates between women and men is also statistically significant ($p = 0.04$). We can therefore conclude that female household members tend to adjust their beliefs in response to direct

and indirect information, while male household members tend to ignore the information unless they are directly informed, in which case they react as strongly as directly informed women.

Panel B of Figure 3 tests the robustness of these results through a falsification test. We measure the effect of direct and indirect information provision on beliefs about household income rank at baseline. Given that we elicited these beliefs before any information revelation on true household income rank, we expect to see no effect on these prior beliefs. The figure confirms this expectation, as the placebo learning rate is close to zero, statistically insignificant, and precisely estimated.

We cannot say exactly why women react to information revealed to their spouse whereas men do not; however, we find, similarly to the Chennai-study, that gender differences in intended communications do not appear to be the reason. We asked treated respondents in the second survey wave whether they could recall having shared information that they learned a year earlier. Male and female household members did not differ in the propensity to report having done so (test of gender difference in propensity to say that information was shared, $p = 0.32$).

Discussion


The fact that people underweight information discovered by others can lead to inefficient and faulty beliefs. This is true not only in society at large, but also in a context where information diffusion should meet few barriers: one's own household. Using a representative sample of Germans, we show that this is indeed the case. We document that while providing information to men influences their female partners' belief formation as much as if the information had been provided directly to the female, the opposite is not true. In fact, men's beliefs do not react at all to information provided to their female partners. This indicates the existence of asymmetries in information sharing in the household in a naturally-occurring setting. Future research should extend the setting explored here to other types of beliefs (*e.g.*, inflation expectations, effectiveness and safety of vaccines) and other contexts (*e.g.*, other developed and developing countries).

References

- [1] Lundberg S, R A Pollak 1996 Bargaining and Distribution in Marriage. *J Econ Persp* 10 (4) 139-158
- [2] Becker, G. S 1981 *A treatise on the family*. Cambridge, MA: Harvard University Press
- [3] Samuelson P A 1956 Social Indifference Curves, *Quart J Econ* 70(1) 1-22
- [4] Zelizer V A 2005 *The Purchase of Intimacy*. Princeton, NJ: Princeton University Press
- [5] Dwyer D, J Bruce 1988 *A Home Divided: Women and Income in the Third World*. Palo Alto, CA: Stanford University Press
- [6] Bittman M, P England, L Sayer, N Folbre, G Matheson 2003 When Does Gender Trump Money? Bargaining and Time in Household Work. *Am J Soc* 109 186-214
- [7] Pahl J M 1983 The Allocation of Money and the Structuring of Inequality Within Marriage. *Soc Rev* 31(2) 237-62
- [8] Lundberg S J, R A Pollak, T J Wales 1997 Do Husbands and Wives Pool Their Resources? Evidence from the U.K. Child Benefit. *J Hum Resource* 32(3) 463–80
- [9] Ward-Batts, J 2008 Out of the wallet and into the purse: Modeling family expenditures to test income pooling. *J Hum Resource* 43(2) 325–51
- [10] Duflo, E 2003 Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. *World Bank Econ Rev* 17(1) 1-25
- [11] Armand A, O Attanasio, P Carnerio, V Lechene 2020 The Effect of Gender-Targeted Conditional Cash Transfers on Household Expenditures: Evidence from a Randomized Experiment. *Econ J*, 130(631) 1875-1897
- [12] Ashraf, Nava, "Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines," *American Economic Review*, 2009, 99 (4), 1245–77
- [13] Ashraf N, E Field, and J Lee 2014 Household Bargaining and Excess Fertility: An Experimental Study in Zambia. *Am Econ Rev* 104(7) 2210– 37
- [14] Ashraf N, E Field, A Voena, R Ziparo 2020 Maternal Mortality Risk and Spousal Differences in the Demand for Children. NBER WP 28220

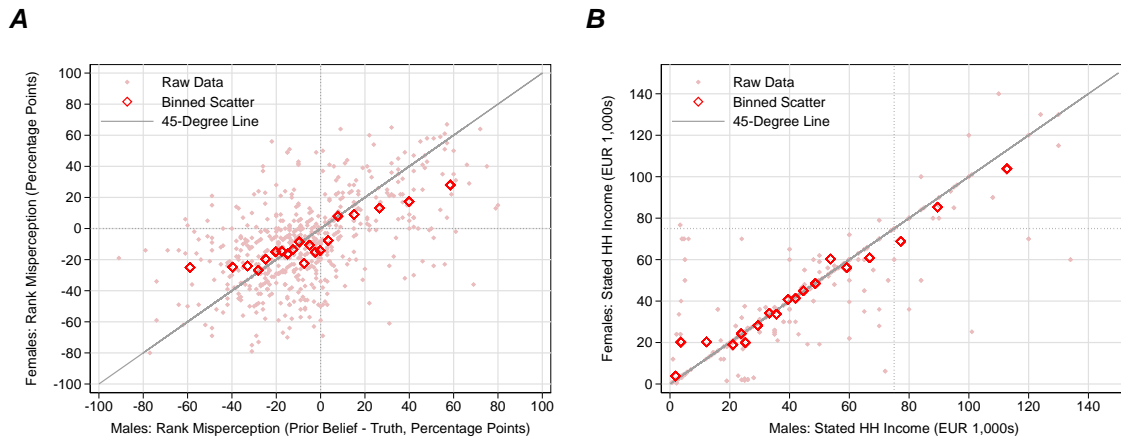
- [15] Conlon J, M Mani, G Rao, M W Ridley, F Schilbach 2021 Learning in the Household. NBER WP 28844
- [16] Conlon J, M Mani, G Rao, M W Ridley, F Schilbach 2022 Not Learning from Others. NBER WP 30378
- [17] Fehr D, J Mollerstrom, R Perez-Truglia 2022 Your Place in the World – Relative Income and Global Inequality. *Am Econ J – Econ Pol* 14(4) 1-37
- [18] Meltzer A, S Richard 1981 A rational theory of the size of government *J Pol Econ* 89 914–927
- [19] Alesina A, P Giuliano 2011 Preferences for redistribution. In *Handbook of Social Economics*, Elsevier, vol 1 93–131
- [20] Cruces G, R Perez-Truglia, M Tetaz 2013 Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *J Pub Econ* 98 100–112
- [21] Karadja M, J Mollerstrom, D Seim 2017 Richer (and holier) than thou? The effect of relative income improvements on demand for redistribution. *Rev Econ Stat* 99 201–212
- [22] Goebel, J., Grabka, M. M., Liebig, S., Kroh, M., Richter, D., Schröder, C., and Schupp, J. (2019). The German socio-economic panel (SOEP). *Jahrbücher für Nationalökonomie und Statistik*, 239(2), 345-360.

Figure 1. Both Genders Misperceive Their Household Income Rank



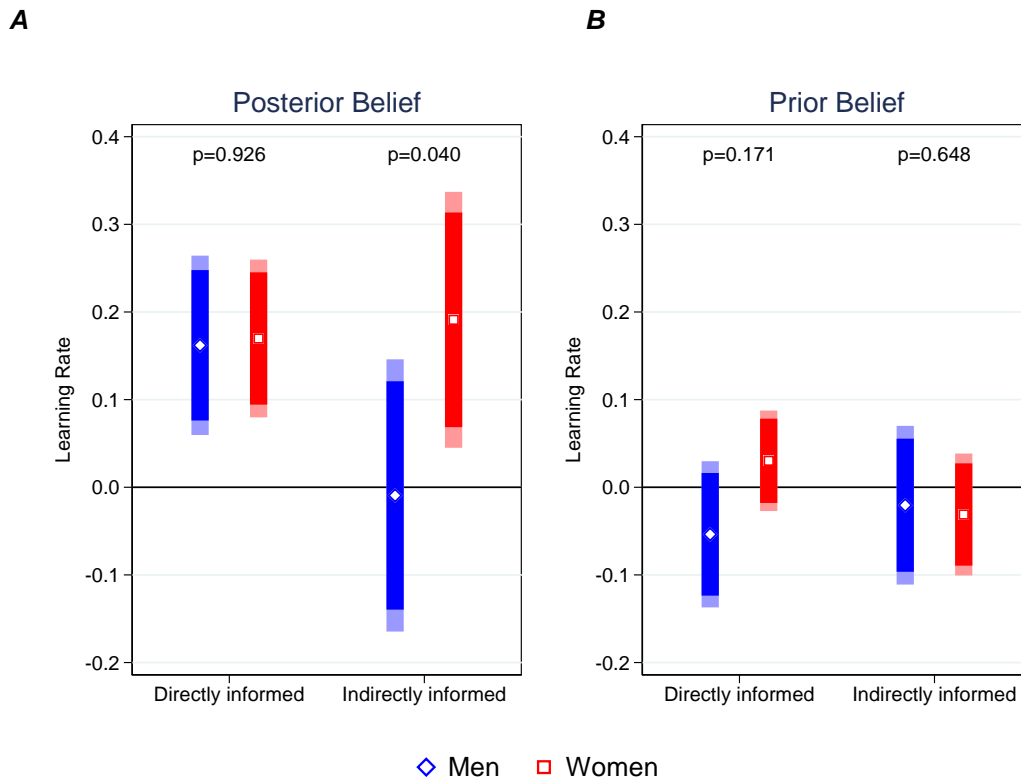
Notes: Distribution of misperceptions about income rank in the baseline survey for female (red) and male respondents (blue). The mean misperceptions is 0.1 for females and 0.09 for males and the distributions do not differ statistically (Kolmogorov-Smirnov test, $p > 0.12$). Misperceptions are calculated 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 ($n=1,978$).

Figure 2. Misperceptions Between Partners Are Correlated, But Far From Perfectly So



Notes: Panel A: correlation between misperceptions about income rank of females and males (within household). Panel B: correlation of stated household income of females and males (within household). Misperceptions are calculated as the difference between prior beliefs about income rank and true income rank. Stated household income is measured in 1,000 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 variable within each bin. Both figures use data from the baseline survey and we restrict the sample to 2-person, mixed-gender households ($n=1,210$).

Figure 3. Women Listen To Their Partner, Men Do Not



Notes: Coefficient plots of learning rates from OLS regressions estimating the effect of information provision on beliefs about income rank one year later (posterior beliefs) with standard errors clustered at the household level (along with 90% (light color) and 95% (intense color) confidence bands). The learning rate corresponds to the pass-through of information on true income rank within household, *i.e.*, the effect of providing direct information to a respondent or indirect information through a respondent's partner on this respondent's beliefs about income rank one year after the intervention. Panel A shows the main result: there is no evidence that the learning rate differs across men and women if they receive the information directly (Directly informed), whereas the learning rate for women is substantially higher than for men if they are indirectly informed (*i.e.*, if women receive the information through their male partner vs. men receiving information through their female partner). Panel B shows a falsification test: there is no effect of direct and indirect information transmission on beliefs about income rank in the same year (prior beliefs) as beliefs are elicited *before* the information provision. Regressions control for respondent's income, the number of household members, the prior belief about the income rank, the change in the true income rank in between the two surveys, and the following demographic characteristics: age and indicator variables for education, disability, unemployment, retirement, self-employment, political party and household location (East/West Germany). The sample is restricted to single and two-person, mixed-gender households (with or without children under 16 living at home). The results are robust to using the full sample and restricting the sample to two-person, mixed gender households only. The p-values for testing for gender differences (on top of the coefficient plots) are derived from fully saturated specifications (*i.e.*, interacting all relevant variables with an indicator for gender). For more details see the Supplementary Information (SI).

Supplementary Information

Extended Methods

Survey implementation: We implemented a tailor-made survey module in two waves of the Innovation Sample of the German Socio-Economic Panel (SOEP-IS) [17]. The SOEP-IS is a longitudinal study that surveys a representative sample of the German population on a wide range of topics once a year. The surveys are conducted by trained professional interviewers in computer-assisted face-to-face interviews. The SOEP-IS offer several advantages over other survey modes. First, the SOEP team implements various safeguards to ensure high data quality, such as pre-testing new items, and conducting plausibility and consistency checks after data collection [22]. Second, the face-to-face interviews provide significant control. For example, due to the fact that interviews were conducted in private with each member of a household, information look-up and communication between household members during and between the interviews within a wave was not possible. Third, while we only designed a module of the survey, we have access to responses to questions in all modules, including a rich set of measures of socio-economic indicators. Moreover, due to the longitudinal character of the SOEP, we can track outcomes in years before and after the baseline survey.

Design – Baseline Survey: In the baseline survey, we asked respondents to assess their household rank in the income distribution and then provided truthful and accurate information about their income rank to a randomly selected subset of respondents. Specifically, we asked respondents to state their perceived rank in the national (German) and global income distribution (in a randomized order). Due to the fact that estimates of the global income distribution are only available on per-capita pre-tax level, we explained and informed all respondents about their per-capita pre-tax household income based on their stated absolute household income. They then state their rank in the national and global income distributions on a scale from 0 (poorest percentile) to 100 (richest percentile) in private to avoid a social desirability bias (*i.e.*, without the interviewer seeing the tablet screen). We incentivized both assessments of income rank for accuracy and respondents received €20 for each assessment that was correct to the closest percentile (ensuring that it was optimal for them to answer in a way that elicited the true mode of their beliefs). About 10-15 minutes later, after answering several SOEP-IS questions unrelated to our research, we randomized half of the respondents into a treatment providing them information about their true income rank in the national and global income distributions. The information revealed how many percentiles are poorer at the national and global level, based on their stated pre-tax per-capita household income, and we additionally visualized this information using customized graphs to make it easier to understand and digest. The other half of respondents received no information.

Design – Second Survey: The setup of the second survey closely followed the setup of the baseline survey. That is, we first collected information on household income and the number of household members and explained what per-capita household income is. We then asked respondents to state their rank in the national and global income distributions. Again, we rewarded accurate predictions in the same way as in the baseline survey, but this time we only 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 this survey. The second survey allows us to investigate (i) whether the information provided to the survey participants in the baseline survey had persistent effects a year later and (ii) how this information diffused within households.

Data: Our analysis builds on our two survey modules that we implemented in the 2017 and 2018 waves of the SOEP-IS. A total of 1,392 respondents took part in the baseline survey, while 1,144 participated in the second survey (82 percent of the 1,392 respondents in the baseline survey). The lower participation rate in the second survey may raise a concern that providing information on the household income rank in the baseline survey could have affected the decision to participate in the second survey. However, this is not the case as there is not significant difference in the attrition rates between individuals who were in the control group (17 percent attrition), and individuals who were in the treatment group in the baseline survey (19 percent attrition, p-value=0.432 for t-test of proportions). Moreover, as expected, the observable pre-treatment characteristics are balanced across all treatment groups (for more detail, see also additional analysis in the replication package).

Regression Specification: The main result presented in Figure 3 is based on the following regression specification. 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 of the household if the individual was in the treatment group. Consequently, $r_i^{info} - r_i^{prior}$ is the misperception about the income rank (*i.e.*, a positive number indicates underestimation and a negative number indicates overestimation of the income rank). T_i is a treatment indicator variable taking on the value 1 if the respondent received direct information on their households' income rank in the baseline survey and 0 otherwise. T_i^{peer} takes value 1 if the respondent did not receive the information, but their spouse did, and 0 otherwise. Recall that any sharing of information among spouses must take place *after* the baseline survey and thus *after* they received information about the income rank, as each interview was conducted in private and communication between family members during the interview was not permitted. We then use the following specification to estimate the direct and indirect effect of providing information on the household's income rank:

$$r_i^{posterior} = \alpha \cdot (r_i^{info} - r_i^{prior}) \cdot T_i + \alpha^{peer} \cdot (r_i^{info} - r_i^{prior}) \cdot T_i^{peer} \quad (1)$$

$$+\beta_1(r_i^{info} - r_i^{prior}) + \beta_2 X_i + \epsilon_i$$

The dependent variable, $r_i^{posterior}$, is the perceived income rank in the second survey, and X_i is a vector of control variables that include the demographic characteristics of the respondent and the household. The coefficient α is the direct effect of information on perceived income rank, *i.e.*, it is the rate of pass-through between the information given, and subsequent beliefs. For example, a coefficient of 0.1 would indicate that for each percentage point shock in information given, the posterior belief a year later is higher by 0.1 percentage points. The coefficient of interest is α^{peer} , which is the indirect effect of information on perceived income rank, *i.e.*, the rate of pass-through between the information we gave to the partner of a respondent and their own beliefs one year later. Note that we should not expect a perfect pass-through rate in both cases (*i.e.*, $\alpha = 1$ and/or $\alpha^{peer} = 1$). In theory, Bayesian respondents would form posterior beliefs by taking a weighted average between the signal provided to them and their prior beliefs. We estimate (1) separately for women and men and cluster standard errors at the household level.

Analysis: We restrict our sample to single households and households with two mixed-gender adult partners ($n=989$). Note that we only consider two-person households to avoid cases in which information travels from children or other household members to parents or adults. While we have two-person households in our sample who never received information about their true income rank from us, we cannot be certain that one member obtains such information (*e.g.*, through active search) and discusses this with their partner. For this reason, we include single households as a “pure” control group (*i.e.*, any information on income rank can only have a direct effect on beliefs one year later). Finally, we observe beliefs about the income rank at the national and global level for each respondent. In the analysis, we pool these two observations, as differentiating between the two belief statements is inessential for our purposes, resulting in $n=1,978$ person-income rank observations. In the replication package we present the regression results underlying our main result presented in Figure 3. We show that this result is robust to: (i) using no restriction on the sample (*i.e.*, using the full sample, $n=2,259$); (ii) using only partners ($n= 1,203$); (iii) not splitting the sample along the gender dimension (*i.e.*, using a fully saturated model with interactions); and (iv) not pooling the beliefs about national and global income ranks.