



Universiteit
Leiden
The Netherlands

Financial decision-making, gender and social norms in Zambia: preliminary report on the quantitative data generation, analysis and results

Barr, A.; Dekker, M.; Mwansa, F.; Zuze, T.L.

Citation

Barr, A., Dekker, M., Mwansa, F., & Zuze, T. L. (2020). Financial decision-making, gender and social norms in Zambia: preliminary report on the quantitative data generation, analysis and results. *Asc Working Paper Series*. Retrieved from <https://hdl.handle.net/1887/92344>

Version: Not Applicable (or Unknown)

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/92344>

Note: To cite this publication please use the final published version (if applicable).

African Studies Centre Leiden
The Netherlands

Financial decision-making, gender and social norms in Zambia: Preliminary report on the quantitative data generation, analysis and results

Abigail Barr

(School of Economics, University of Nottingham)

Marleen Dekker

(African Studies Center, Leiden University)

Floyd Mwansa

(Financial Sector Deepening Zambia)

Tia Linda Zuze

(Financial Sector Deepening Zambia/ReSeP, Stellenbosch University)

ASC Working Paper 147 / 2020

African Studies Centre Leiden
P.O. Box 9555
2300 RB Leiden
The Netherlands

Telephone +31-71-5273372
Website www.ascleiden.nl
E-mail asc@asc.leidenuniv.nl

© Abigail Barr (School of Economics, University of Nottingham)
Marleen Dekker (African Studies Center, Leiden University)
Floyd Mwansa (Financial Sector Deepening Zambia)
Tia Linda Zuze (Financial Sector Deepening Zambia/ReSeP, Stellenbosch University)

Financial decision-making, gender and social norms in Zambia: Preliminary report on the quantitative data generation, analysis and results

Abigail Barr¹, Marleen Dekker², Floyd Mwansa³ and Tia Linda Zuze⁴



¹ School of Economics, University of Nottingham

² African Studies Center, Leiden University

³ Financial Sector Deepening Zambia

⁴ Financial Sector Deepening Zambia/ReSeP, Stellenbosch University

Acknowledgements

This study was funded by FSDZ (DFID project number 203000) and NWO Westerdijk Talent Scheme project number 014.041.029. We would like to thank Patricia Lowes for her assistance in the generation of the fieldwork materials and FSDZ staff for their warm reception in the office and useful feedback during our study design. The data collection for this study would not have been possible without the excellent guidance by Testify Shampongo and Kingsley Sambo (FSDZ), impeccable field management by Samuel Msanide and our team of excellent field researchers Berlynah Nkosi Ziba, Bernard Mwanza, Chibwe Lumpa, Chipasha Kafwimbi, Cliff Kabengele, Elizabeth Chulu, Ellimo Mwiinga, Emelia Banda, Grace Mutale, Jeanine Mbaya, Jonas Ngoma, Joseph Mwiinga, Julius Masiwa, Liseli Ndombo, Nicholas Hatontola, Nsangu Tembo, Racheal Lupiya, Solomon Dakah, Tapiwa Christine Janda and Victor Chambawilo. This study was approved by the ERES Converge Institutional Review Board Zambia (ref. 2019-Jun-012), the Nottingham School of Economics Research Ethics Committee and the ASCL Research Ethics Review Board.

Abstract

This document presents the preliminary findings from the quantitative data generation and analysis conducted as part of the project “Financial decision-making, gender and social norms in Zambia”. Using a series of specially designed behavioural experiments, we generated an extensive set of insights into the normative environment within which spouses in Eastern Province, Zambia, make decisions about individual money holding and saving. Here are some of those insights. Spouses in Eastern Province, Zambia, are willing to compromise household-level earnings in order to maintain individual control over money. Wives, but not husbands, are more likely to compromise household-level earnings in order to maintain individual control over money, when they can keep that money and their actions hidden from their spouses. Individually-held behavioural prescriptions, i.e., the “shoulds” and “oughts” that individuals have in mind and reference as guides for their own behaviour and as benchmarks against which to evaluate others’ behaviour, inform decision-making about maintaining individual control over money at a cost to the household. Further, when individuals know that their spouses will find out about their decisions regarding maintaining individual control over money (or not) at a cost to the household, the individuals take their spouses’ opinions about what they should do into account, i.e., they compromise. There is strong but not unequivocal evidence pointing to the existence of a social norm, i.e., a “should” or “ought” that is collectively held and enforced by members of a community, forbidding saving in secret from one’s spouse, with the secrecy not the saving being the problem. Assuming it exists, this social norm forbidding saving in secret from one’s spouse applies to both husbands and wives, and this is acknowledged by both husbands and wives. However, the extent to which violations of this norm are tolerated depends on who is doing the violating and who the evaluating. In patrilineal communities (as compared to matrilineal communities), both husbands and wives are especially intolerant of secret saving by husbands and in both patrilineal and matrilineal communities, wives are less tolerant than husbands of secret saving by husbands and more tolerant than husbands of secret saving by wives. This relative tolerance of secret saving by wives notwithstanding, just under one in three wives and one in six husbands think that a man is justified in beating his wife if he discovers that she is saving in an e-wallet or has joined a savings group without his knowledge and, as grounds for wife beating, saving in secret is on a par with neglecting the children, visiting friends or family in secret and refusing to have sex. For further insights, see the main text of the report.

I. Introduction

The aims of this project are to investigate the social norms that guide and constrain the financial decision-making of Zambian men and women, and to explore ways of incorporating knowledge about these norms into financial services, financial literacy programmes and policy design. In this report, we present the findings from the first part of the project, within which we used quantitative methods to address a pre-defined set of research questions. During the second part of the project (yet to be undertaken) we plan to initiate stakeholder discussions about how best to incorporate the knowledge into the designs of financial services, financial literacy programmes and other complementary interventions. Discussing the findings with communities such as those that participated in the data generating activities will be an important part of this, but so too will be presentations to and discussions with representatives from both for-profit and non-profit financial service providers, advocates for the expansion of financial inclusion, policy makers and regulators.

Using recently developed methods drawn from behavioural and experimental economics, especially adapted for this project, we planned to **identify** the content of a range of social norms relating to money holding, money use and money management and **quantify** the strength of those social norms.

By the **content of a social norm**, we mean the **socially prescribed behaviour** for an individual facing a particular set of behavioural choices. The **strength of a social norm** can be defined as the speed with which social approval declines as an individual's chosen behaviour moves away from the behavioural prescription. However, it can also be defined with reference to the likelihood of or severity with which an individual who deviates from the behavioural prescription is punished. Finally, it can be defined with reference to the strength of the correlation between the behavioural prescription and actual behaviour. Drawing on the conceptual framework proposed by Bicchieri (2006), we planned to investigate all three of these types of norm strength and how they relate one to another, while focusing specifically on financial-decision-making-relevant social norms.

Our specially adapted methods also facilitated the measurement of individually held behavioural prescriptions relating to financial decision-making and this turned out to be very important as the data we gathered indicates that:

- financial decision-making by Zambian men and women is guided and constrained by social norms to a lesser extent than we originally expected; and
- individually held behavioural prescriptions and differences in such prescriptions between spouses may have important implications for the decisions they make.

We also planned to investigate the extent to which financial-decision-making-relevant social norms are gender specific and whether the social norms that women acknowledge as guides to the way that they should behave differ from the social norms that men think should be guiding women's behaviour and vice-versa. Here, with regard to the one social norm that was clearly identified in our data, we found that the content of the norm was not gender specific, but that deviations from the norm by husbands were tolerated less than deviations by wives, especially by wives. With regard to individually held behavioural prescriptions, we found marked differences between the genders in terms of both what each thought that they themselves should do and what each thought the other should do.

Finally, we planned to investigate whether financial-decision-making-relevant social norms vary across ethnic groups (tribal groupings within the Zambian context). Here, we focused on the Ngoni, who are traditionally patrilineal, and the Chewa, who are traditionally matrilineal. We found no evidence of differences in the content of financial-decision-making-relevant social norms and individually held

behavioural prescriptions between the two ethnic groups. However, we did find some significant differences in the extent to which deviations from those norms and prescriptions are tolerated.

There are many reasons why an individual's behaviour might deviate from a social norm. Most notably it may be individually sub-optimal to adhere to a social norm. If this is the case, the individual might choose to deviate from the norm when they can hide their actions and adhere to the norm when they cannot. Further, they might be willing to compromise economic efficiency in order to generate and take advantage of opportunities to hide their actions, i.e., they may be willing to divert effort and resources away from more productive activities in order to increase the chances of keeping norm-violations hidden.

Within the context of a household, even in the absence of gender differentiated social norms, if there are differences in individually held behavioural prescriptions between husbands and wives, those husbands and wives might be willing to compromise household-level economic efficiency in order to generate and take advantage of opportunities to hide their actions. Specifically, if a wife (husband) knows what her husband (his wife) thinks she (he) should do, but disagrees, she (he) might be willing to compromise household-level economic efficiency in order to generate and take advantage of opportunities to do as she (he) wishes without her husband (his wife) knowing.

When, in the interests of economic development, we wish to facilitate behavioural change, it is important to know what is guiding and constraining individual behaviour. If social norms are guiding and constraining behaviour, social norms need to be changed. This can be difficult because social norms tend to be held in place by multiple mechanisms. An individual deviates from a social norm and, if he or she has internalised the norm, feels guilt. An individual deviates from a social norm and feels ashamed owing to social disapproval, either imagined or expressed by other members of her or his society. An individual deviates from a social norm and is punished, materially, physically or psychologically, by other members of her or his society. To change a social norm and, hence, remove it as a constraint on individual behaviour, one needs to address each of these mechanisms, and this involves working with the entire or a large part of the society to which the norm belongs. In contrast, while an individual who deviates from a behavioural prescription that they hold and that pertains to their own behaviour might feel guilt, he or she is not disapproved of or punished by other members of her or his society. So, only one mechanism needs to be addressed and there is no need to work with entire or large parts of societies. This renders individually held behavioural prescriptions easier to change. Finally, an individual who deviates from a behavioural prescription that an important other in her or his life, e.g., her or his spouse, holds might be disapproved of or punished by that other, but will not experience guilt. In this case, successful behaviour-changing interventions need to be directed at both those whose behaviour we wish to change and those whose individually held behavioural prescriptions need to be changed, but the diversity of views within both individuals' society (or societies) can be used to facilitate such change. Development practitioners in possession of knowledge about what is guiding and constraining individual behaviour can design services and interventions that support sustainable behaviour change and do not lead to unintended consequences, especially for the disempowered.

In this report, we set out our findings from the quantitative data generation conducted in July-August 2019 and the subsequent analysis of that data. The findings reported below are preliminary in the sense that there is more to be done that could generate further insights. This notwithstanding, the findings are comprehensive in the sense that they address all of the questions we set out in our proposal and have been subjected to an extensive range of scientific robustness checks.

The report is structured as follows: Section II presents the research agenda that we set out prior to collecting our data; Section III describes the methods we used to generate our data (III.i), our sample design (III.ii), and our analytical methods (III.iii); and Section IV presents our data analysis and results.

II. Quantitative research agenda

By engaging a sample of Zambian women and men in a series of incentivised tasks, i.e., tasks for which they would be paid, with the amount they were paid being dependent on the decisions that they and others made during the tasks, we planned to address the following questions:

1. Are spouses in Zambia willing to compromise household economic efficiency in order to maintain individual control over money?
2. Specifically, when given the opportunity to do so, do spouses hide money from each other even when hiding is costly?
3. Are wives and husbands differentially inclined to incur costs with the aim of hiding and, thereby, maintaining individual control over money?
4. Are the decisions made by spouses about money hiding, money holding, and money and financial service use informed by social norms about who within households should hold and make decisions about money?
5. Specifically, if a wife chooses to hold and make independent decisions about money and financial service use, is she violating a norm?
6. Do the social norms that women acknowledge as guides to the way they should behave differ from the social norms that men think should be guiding women's behaviour (and vice-versa)?
7. Do social norms about who within a household is supposed to hold and make decisions about money vary systematically depending on whether a community is matrilineal or patrilineal?

Based on the findings of previous studies undertaken outside Zambia (see, for example, Ashraf (2009), Iversen et al (2011), Kebede et al (2014), Jakiele and Ozier (2016), Almås et al (2018) and Barr et al (2019)), we expected the answers to questions 1 to 3 to be "yes". However, it was important to establish whether and to what extent these findings replicate in Zambia as they related directly to the conceptual foundation upon which we build.

Studies addressing questions 4 to 7 are extremely rare. Based on Afzal et al (2016), discussions taking place during webinars focusing on financial inclusion and women's empowerment around the time when we were formulating our research questions, and preliminary discussions with women in Zambia, we expected the answers to questions 4 to 7 to be "yes" as well. However, in the interests of scientific rigor, we allowed for the answers to these questions to be "no" and ensured that our research methods would allow us to investigate, not only social norms but also individually held behavioural prescriptions, whether and how they differ between women and men and matrilineal and patrilineal communities and how they impact on decision-making.

III. Quantitative research methods

III.i Data generation methods

The Allocation Task (AT): At the core of our quantitative research design was an adaptation of the incentivised task designed by Almås et al (2018). Within our version of the task wives and husbands had to make a series of eleven decisions about whether a sum of money should be placed in their own hands or in the hands of their spouse. Within each husband-wife couple taking part in the study, one

spouse was randomly picked to be the decision-maker (DM), while the other spouse was passive (PS). The DM was guided through the task under carefully controlled conditions. Most importantly, the DM and PS were not allowed to discuss the task before the DM made her/his decisions and the PS only found out about the task and, in one treatment, the decisions made by his/her spouse after the latter had completed the task.

Each DM started the task by choosing either K45 to self or K45 to spouse (referred to below as the 45:45 decision),⁵ and then worked through a further ten decisions designed to quantify her/his willingness to pay to have an amount of money in either her/his own or her/his spouse's possession. The further ten decisions were: K40 to self or K45 to spouse (the 40:45 decision); K35 to self or K45 to spouse (the 35:45 decision); K30 to self or K45 to spouse (the 30:45 decision); K25 to self or K45 to spouse (the 25:45 decision); K20 to self or K45 to spouse (the 20:45 decision); K15 to self or K45 to spouse (the 15:45 decision); K10 to self or K45 to spouse (the 10:45 decision); K5 to self or K45 to spouse (the 5:45 decision); zero to self or K45 to spouse (the 0:45 decision); and K45 to self or K40 to spouse (the 45:40 decision). Once a DM had completed all eleven decisions, they rolled a die to establish which, if any, would be used to determine their own and their spouses' earnings from the task.

Each DM was randomly assigned to one of two treatments, either the **Zero Deniability Treatment** or the **High Deniability Treatment**. Under the **Zero Deniability Treatment** each PS (later) learned about the decisions made by his/her DM spouse and the DM knew that this would happen before making his/her decisions. Under the **High Deniability Treatment**, the PSs received their payoff, if one was due, but the specifics of the treatment ensured that they learnt nothing (or, in a few cases, very little)⁶ about the decisions made by their DM spouses.

The decisions the DMs took in this task inform us about their willingness to compromise household efficiency in order to maintain individual control over money (see research question 1). Comparing decisions across treatment arms revealed their preferences regarding the use of opportunities to hide money (see research question 2) and comparing wives' and husbands' decisions under the two treatments allowed us to investigate gender differences in these preferences and behaviours (see research question 3).

The Evaluation Task Part 1 (ET1): Using survey questions and methods originally proposed by Burks and Krupka (2012) and Krupka and Weber (2013), we set out to identify the content and strength of the individually held behavioural prescriptions and social norms that pertain to the AT within each of the communities in which husbands and wives participated in the AT. It is important to note that we were not assuming that individuals and communities have pre-formed individual prescriptions and social norms relating specifically to the task. Rather, we were assuming that they would have individual prescriptions and social norms pertaining to money holding and use by wives and husbands and would be adept at spotting the relevance of and applying these prescriptions and norms to emergent decision-making contexts.

⁵ In the original design, the first AT decision was K50 to self vs. K50 to spouse. However, data from the technical pilot suggested the existence a preference discontinuity between K50 and K45. Starting the AT decision series with K45 to self vs. K45 to spouse, obviated any problems that this discontinuity might have caused at the analysis stage.

⁶ PSs who received K40 would have been able to infer that their DMs chose K40 to spouse over K45 to self in the 11th AT decision.

To elicit the **individually held behavioural prescriptions**, under carefully controlled conditions (most importantly, no talking between participants) we asked each PS and DM spouse to evaluate the appropriateness of:

- a DM wife choosing K45 for herself and zero for her husband in the AT;
- a DM wife choosing K45 for her husband and zero for herself in the AT;
- a DM husband choosing K45 for himself and zero for his wife in the AT; and
- a DM husband choosing K45 for his wife and zero for himself in the AT.

They indicated the level of appropriateness of each of these actions on the following four point scale:

- very appropriate (quantified as =1 in the analysis below)
- somewhat appropriate (quantified as =0.33 in the analysis below)
- somewhat inappropriate (quantified as =-0.33 in the analysis below)
- very inappropriate (quantified as =-1 in the analysis below).

The DM spouses undertook this evaluation task after they had completed the AT. The PS were talked through the first decision in the AT before making their evaluations.

We, then, asked the DM spouses to go through the same evaluation process when the AT decision was between K35 for self and K45 for spouse.⁷

To investigate whether the individually held behavioural prescriptions corresponded to or where distinct from any relevant **social norms** we asked the DM and PS spouses (again, under carefully controlled conditions) to repeat the evaluations. However, this time, prior to making their evaluations, they were told that, once they had made their evaluations (16 in the case of the DM spouses, 8 in the case of the PS spouses), one would be randomly selected and, if their own evaluation matched the most common evaluation given in this selected case within their *session* (more on this below) they would receive a payoff of K30.

Here, who is being incentivised to match evaluations with whom, i.e., who is in a *session* with whom, is critical. We ran four evaluation *sessions* in each of the communities included in our study, one attended by PS wives, one by PS husbands, one by DM wives and one by DM husbands, with the sessions involving DM spouses run after they had completed the AT. So, each participant was incentivised to match evaluations with a sample of others of the same gender, from the same community and with the same prior experience of the AT.

If a relevant social norm exists and is referenced during this task, when the incentive to match is in place, for a given DM type facing a given dichotomous choice, we would expect one option to be considered very or somewhat appropriate and the other to be considered very or somewhat inappropriate by the large majority or all of the participants in a session. In addition, we would expect the corresponding unincentivized individual evaluations to be more or similarly diverse. If instead, the unincentivised individual evaluations are less diverse than the incentivised ones, it indicates that the participants have no readily available social norm to help them match their evaluations and earn the K30 payoff.

⁷ In the original design, PSs were also going to undertake this second set of evaluations and both DMs and PSs were going to undertake one further set of evaluations. However, our technical pilot indicated that this would not be possible given the daily time constraint we faced. This time constraint was binding, as we needed to complete all the data-generating activities in any given village within a single day in order to ensure internal validity.

If a social norm is found to exist, then the incentivised evaluations can be used to address research questions 5 and 6 about whether the social norm is gender-differentiated and question 7 about whether the social norm varies between matrilineal and patrilineal ethnic groups. If the data indicates that no social norm exists, then the unincentivised individual evaluations can be used to investigate whether and how individually held behavioural prescriptions differ between men and women and between members of matrilineal and patrilineal ethnic groups. In either case, the evaluations can also be used to investigate whether and how prior experience of the AT impacts on the apparent content and strength of the prescriptions or norms.

The decisions made by the DM spouses during the AT combined with the ET1 evaluations allowed us to investigate whether and how decisions made by spouses about money hiding and money holding are informed by social norms (see research question 4) and/or individually held behavioural prescriptions.

The Evaluation Task Part 2 (ET2): Using the same methods as described above, we also identified the content, variability and strength of individually held behavioural prescriptions and social norms that pertain to a list of choices about money that husbands and wives might engage in during their everyday lives. These were presented to the participants in the form of a vignette and corresponding choice set. Participants were asked to evaluate the appropriateness of each of the choices presented using the four point scale described above. The vignette and choice set was as follows:

A wife (husband) finds that she (he) often has a bit of cash left over after all of the essentials have been bought. One day she (he) is thinking about what to do with this money and starts considering the following options:

- Give it to her husband (his wife)
- Spend it on treats for the family
- Put it in her (his) usual hiding place in the home and tell her husband (his wife) about the money, but not where it is hidden
- Put it in her (his) usual hiding place in the home without telling her husband (his wife)
- Put it in a mobile money e-wallet and tell her husband (his wife)
- Put it in a mobile money e-wallet without telling her husband (his wife)
- Join a savings group and tell her husband (his wife)
- Join a savings group without telling her husband (his wife).

Only the PS completed this task and,⁸ owing to time constraints, in each session, the participants focused either on a husband or a wife considering these options. Given the latter, we are dependent on cross-subject comparisons when investigating gender-differentiated behavioural prescriptions and norms.

The Evaluation Task Part 3 (ET3): Finally, we used an adaptation of the same method to evaluate individually held opinions and social norms relating to intimate partner violence. Here, we built on the widely used Demographics and Health Survey (DHS) instrument. To be consistent with the DHS, we asked about the justifiability (rather than the appropriateness) of wife beating within each of a set of scenarios. However, we rendered the task relevant to financial decision-making by adding to the list

⁸ In our original design, the DMs also engaged in this task. However, our technical pilot indicated that this would not be possible given the daily time constraint.

of scenarios.⁹ In addition, as in ET1 and ET2, we invited the participants to go through the task, first with, then, without the matching-to-others-responses incentives. Here is the list of scenarios we used:

- Do you think a man is justified in beating his wife if she neglects their children?
- Do you think a man is justified in beating his wife if she burns the food?
- Do you think a man is justified in beating his wife if she argues with him?
- Do you think a man is justified in beating his wife if she visits family or friends without his permission?
- Do you think a man is justified in beating his wife if he discovers that she is saving in a mobile money e-wallet?
- Do you think a man is justified in beating his wife if he discovers that she has joined a savings group?

The data generated using ET2 and ET3 allowed us to extend the investigation into the content and strength of individually held behavioural prescriptions and social norms impacting on choices made in everyday life relating to the use of financial services (see research questions 5, 6, 7).

Survey: All of the participants in the AT and evaluation tasks also responded to a short survey designed to elicit individual and household characteristics, data on individual financial service use and other relevant background measures. Thus far, we have made only limited use of the survey data. Our research design enabled us to use alternative and often more effective methods, specifically, randomisation and fixed effects estimation, to control for differences in individual and household characteristics when seeking the answers to our research questions. However, in further analysis, the use of the survey data may yield additional insights.

III.ii Sample design

We undertook the quantitative data generation in Eastern Province, which is representative of Zambia as a whole (excluding Lusaka and the Copperbelt) with respect to living standards and use of financial services. Within Eastern Province, we engaged 11 Chewa (matrilineal) communities and 11 Ngoni (patrilineal) communities that were comparable with respect to livelihoods and living standards.

Within each of these 22 communities, in accordance with our post-technical pilot research plan, we conducted:

- 1 workshop session involving 4 husbands and 4 wives who undertook the AT under the zero deniability treatment;¹⁰
- 1 workshop session involving 4 husbands and 4 wives who undertook the AT under the high deniability treatment;
- 1 workshop session involving 8 wives (not also involved as DMs in the AT sessions) who undertook the first part of ET1 followed by ET2, and ET3;
- 1 workshop session involving 8 husbands (not also involved as DMs in the AT sessions) who undertook the first part of ET1 followed by ET2, and ET3;
- 1 workshop session involving 8 wives (previously involved as DMs in the AT sessions) who undertook all of ET1;
- 1 workshop session involving 8 husbands (previously involved in as DMs the AT sessions) who undertook all of ET1.

⁹ We included the DHS scenarios in our list to use as a benchmark against which to compare those of greater relevance to our research aims.

¹⁰ Originally, we planned to work with 10 rather than 8 participants in each session. However, the technical pilot indicated that this would not be possible given the daily time constraint.

We ran these six sessions plus the survey in each community within a single day. This went a long way towards ensuring zero communication between participants about the tasks prior to them making their decisions and evaluations. However, it did not preclude the members of one community that had completed the tasks informing the members of another community who had yet to complete the task about the tasks so that the latter could discuss the tasks before engaging. In the few village clusters in which we thought this could be an issue, we conducted unstructured interviews with small samples of participants to find out what, if anything, they had heard about the tasks prior to their workshops. Thus, we learnt that some had heard that we could be trusted, i.e., that, to the extent that the participants could tell, we were paying people at the end of the workshops in accordance with what we were telling people during the workshops. However, none of those interviewed felt that they had gained sufficient prior understanding of the tasks to have been able to collude with their fellow community members.

Had it been implemented to the letter; this sample design would have yielded:

- 1936 AT decisions made by 88 husbands and 88 wives under the zero deniability treatment;
- 1936 AT decisions made by 88 husbands and 88 wives under the high deniability treatment;
- 1408 ET1 series 1 evaluations made by 176 husbands, 88 of whom had previously engaged in the AT, and 176 wives, 88 of whom had previously engaged in the AT;
- 704 ET1 series 2 evaluations made by 88 husbands and 88 wives, all of whom had previously engaged in the AT;
- 2,816 ET2 evaluations made by 88 husbands and 88 wives, none of whom had previously engaged in the AT; and
- 2,464 ET3 evaluations made by 88 husbands and 88 wives, none of whom had previously engaged in the AT.

However, one wife had to leave a workshop in order to take her baby to a clinic after completing the AT but before engaging in ET1 or the survey and one man registered for a workshop with a woman who, during the survey, we discovered was not the man's wife. The AT decisions of the wife who left and all the evaluations made by her husband are included in the analyses presented below. All of the data pertaining to the unmarried man and woman who registered as a couple were deleted from the dataset. These two attendance issues account for all the apparent inconsistencies between the by-design sample sizes listed above and the actual sample sizes reported below.

We conducted all 22 data-generating workshops during the second half of August 2019.

III.iii Analytical methods

Most of the findings presented below have been derived using statistical regression analysis. However, with the aim of making our findings accessible to a diverse audience, wherever possible, in the main text of this report, we present the analysis in graphical form. The regression model specifications and the estimations upon which the graphs are based and some of the additional estimations run to check the robustness of results are presented in the Analytical Appendix.

IV Analysis and results

This section is organised by research question. For each question we present the relevant data in graphical form and then the results of the relevant regression analyses. As we proceed, the conclusions we draw are clearly labelled and presented in bold text).

1. Are spouses in Zambia willing to compromise household economic efficiency in order to maintain individual control over money?


We can answer this question using the data from the AT. Figure 1 (below) plots the proportion of spouses allocating to self (vertical axis) for each of the 11 AT decisions (horizontal axis) under the high deniability treatment (green dots) and zero deniability treatment (orange dots). The figure reveals considerable willingness on the part of spouses to compromise household-level economic efficiency in order to maintain individual control over money.¹¹

In total, 42% of the Allocation Task decisions were to self, despite the fact that allocating to self was the inefficient option in 82% of the decisions. Across the 9 decisions in which allocating to self was inefficient, on average the spouses were willing to incur a cost of K7.94 (17.6%) to be the one holding the money at the end of the workshop. Put another way, on average when making these decisions, the spouses preferred to have K37.06 in their own hand rather than K45 in the hand of their spouse.

Around 30% (107) of the AT decision-making spouses chose to allocate to spouse in the 40:45 decision, i.e., the one decision in which allocating to self was the efficient option. Of these 43% were wives.

Conclusion: Spouses in Zambia are willing to compromise household economic efficiency in order to maintain individual control over money.

Figure 1: Allocations to self in the Allocation Task



Source: authors' data.

¹¹ It is worth noting that over 50% of the DMs flick-flacked between self and spouse across these decisions. This flick-flacking is inconsistent with any simple economic model of decision-making. In about half of these cases it looks like the DM switched from self to spouse early, experienced regret as they saw the to-self payoff decline from one decision to the next, and so switched back for a while to secure a higher chance of some payoff for themselves. However, we do not know, for sure that this was their reasoning and, in the other half of cases, there is no easy way to rationalise the data. It is this flick-flacking that is behind the relatively flat sections of the plots between decision 35:45 and 10:45. Throughout our regression analysis, we treat these inconsistencies as "noise", i.e., as contributors to regression model errors.

2. Specifically, when given the opportunity to do so, do spouses hide money from each other even when hiding is costly?


Figure 1 shows that, for 7 out of the 11 AT decisions, a greater proportion of spouses allocated to self under the high deniability treatment (green) than under the zero deniability treatment (orange) and, for the remaining 4 decisions, the proportions were similar. This suggests that, when they can hide their compromising of household efficiency in order to maintain individual control over money, spouses compromise it more. However, when we analyse the husbands' and wives' decisions together, the cross-treatment difference is not statistically significant (see Table A1, Model 1, in the Analytical Appendix).

Before drawing a conclusion here, we should consider the analysis relating to question 3.

3. Are wives and husbands differentially inclined to incur costs with the aim of hiding and, thereby, maintaining individual control over money?

In Figure 2 (below) we separate the husbands' and wives' AT decisions. The panel on the right indicates that, when wives can hide their compromising of household efficiency in order to maintain individual control over money, they compromise it more. In contrast, the panel on the left indicates that husbands' choices are indistinguishable across the treatments. Regression analysis lends support to these indications (see Table A1, Model 2, in the Analytical Appendix). The treatment effect is statistically significant for wives ($p=0.022$) but not for husbands and, under the high deniability treatment, wives compromise household efficiency more than husbands ($p=0.003$), while under the low deniability treatment, there is no difference between husbands' and wives' behaviour.

Figure 2: Allocations to self by wives and husbands in the Allocation Task



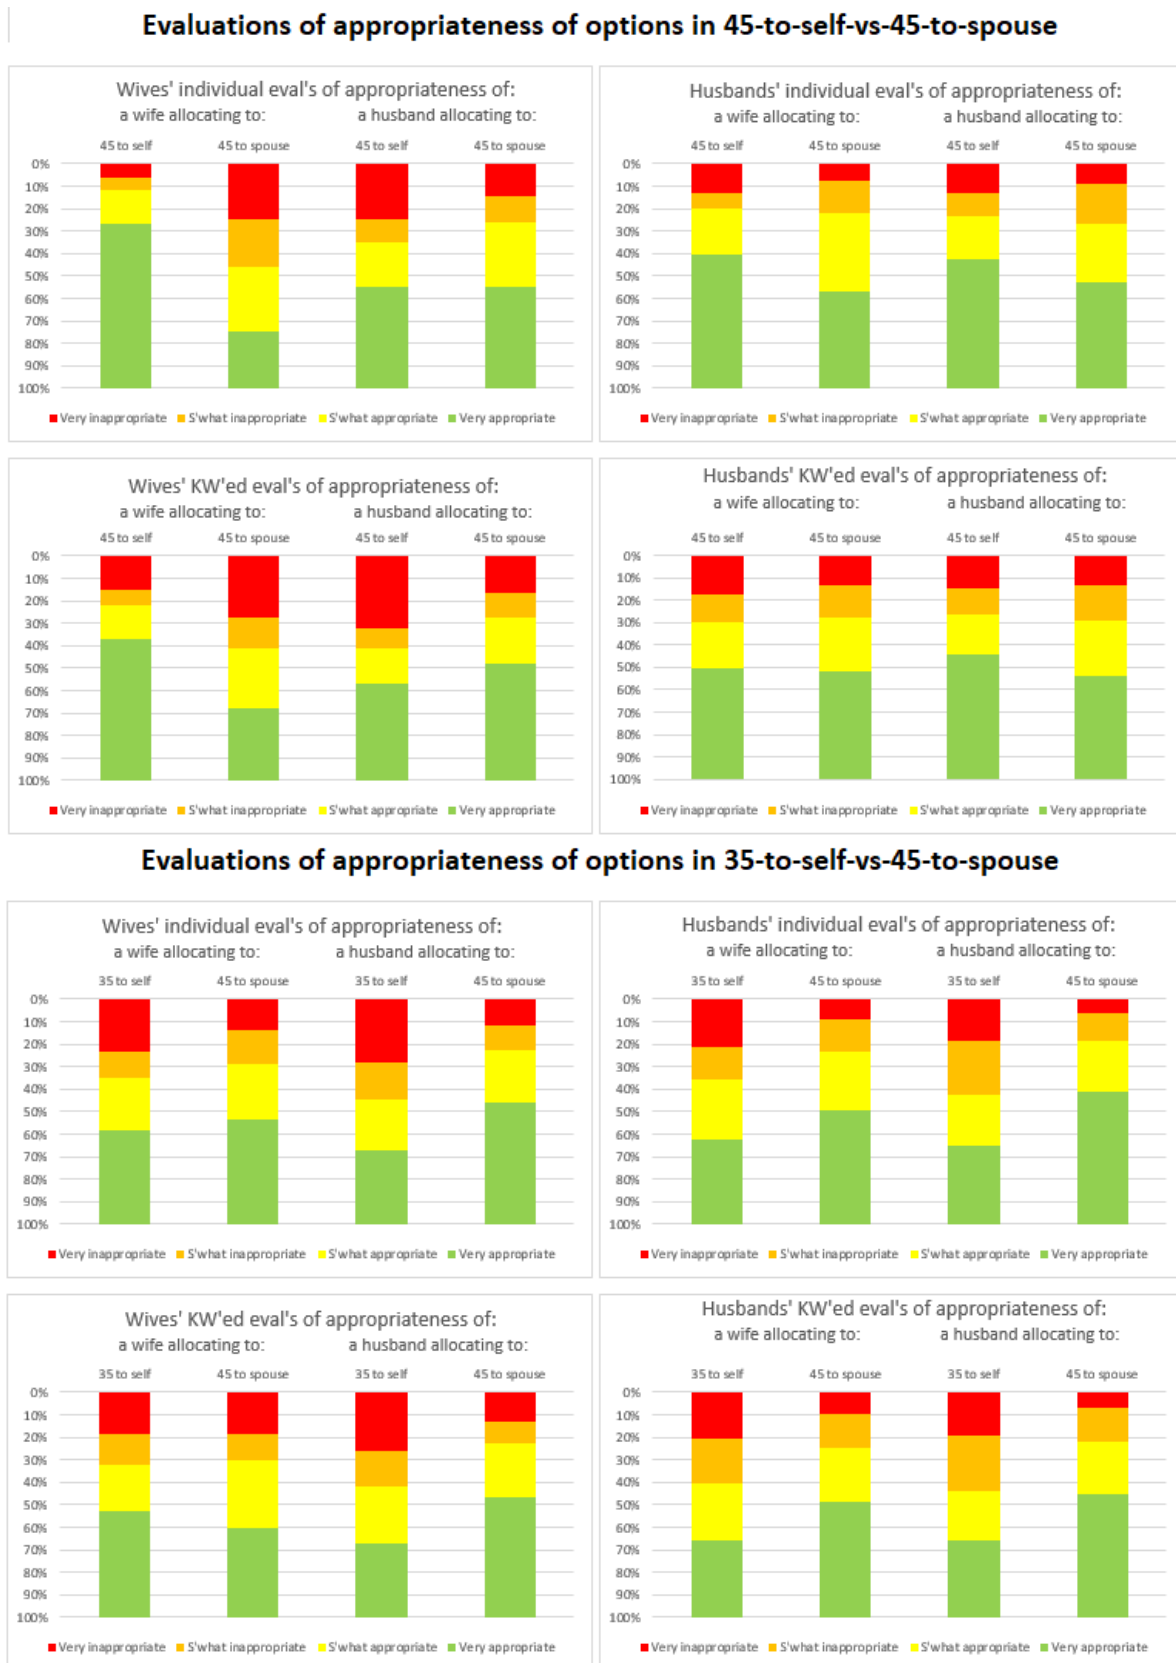
Source: authors' data

Conclusion: Wives are more likely to compromise household-level efficiency in order to maintain individual control over money when they can keep that money and their actions hidden from their husbands. This is not the case for husbands.

4. Are the decisions made by spouses about money hiding, money holding, and money and financial service use informed by social norms about who within households should hold and make decisions about money?

We start this section by looking at the ET1 data which relates to individually held behavioural prescriptions and social norms about money holding that may have affected decision-making in the AT. The ET1 data are presented in Figure 3. The evaluations made by wives are presented on the lefthand side of the figure. The evaluations made by husbands are presented on the right. The evaluations relating to the 45:45 AT decision are presented in the top half of the figure. The evaluations relating to the 35:45 AT decision are presented in the bottom half. Within the top half of the figure, the upper two panels present the evaluations made in the absence of the incentive to match and the lower two panels present the evaluations made when the incentive to match was in place. Within each panel, the first column presents the evaluations of a wife allocating to herself, the second column presents the evaluations of a wife allocating to her husband, the third column presents the evaluations of a husband allocating to himself, and the fourth column presents the evaluations of a husband allocating to his wife. Within each column, the green part indicates the proportion evaluating the action as very appropriate, the yellow part indicates the proportion evaluating the action as somewhat appropriate, the orange part indicates the proportion evaluating the action as somewhat inappropriate, and the red part indicates the proportion evaluating the action as very inappropriate.

Figure 3: Evaluations of appropriateness of options in the Allocation Task (ET1)



Source: authors' data

In order to inform behaviour, a social norm first needs to exist. Currently, in the behavioural experimental literature, there is no agreed method for formally testing a null hypothesis of the form “there is no social norm pertaining to decisions in context ...”. However, there are several characteristics of the ET1 data that are inconsistent with the existence of a social norm informing decisions in the AT. Two of these characteristics are apparent from Figure 3. First, for each evaluator-type-DM-type-action case, there was quite a lot of variation in evaluations, even when the incentive to match was in place. Second, where we see less variation in individual evaluations – most notably, over 70% of wives indicated that they individually thought it was very appropriate for a wife to allocate to herself in the 45:45 decision – there was greater variation when the incentive to match was in place than when it was absent. This is inconsistent with husbands or wives having shared views about what husbands and wives should do that they can reference when trying to match, i.e., it is inconsistent with the existence of social norms.

However, these characteristics of the dataset could be misleading. It could be that the variation within each column in Figure 3 is owing not to the absence of social norms but to differences in social norms across communities and, hence, across the sessions within which the evaluations were made and the incentive to match was relevant. To investigate this possibility, we conducted an evaluation-by-session-level analysis. Across 88 sessions, a total of 528 evaluations of possible actions in the AT were made, while the incentive to match was in place. Among these 528 evaluations:¹²

- perfect matching, i.e., all participants making the same evaluation when the incentive to match was in place, was achieved in only 6 (1.1%);
- partial matching, i.e., all stating that the action was either very or somewhat appropriate or all stating that the action was either very or somewhat inappropriate, was achieved in only a further 54 (10.2%);
- there was a single modal evaluation in 440 (83.3%); but
- within these 440, on average, fewer than 5 out of the 8 participants’ evaluations matched the mode.

Further, statistical regression analysis reveals that there was significantly more within session variation in the evaluations when the incentive to match was in place than when it was not (see Table A2 in the Analytical Appendix). So, even when we allow for the possibility that social norms relevant to the AT vary across communities, we find statistically significant patterns in our data that are inconsistent with the existence of any such social norm.

Conclusion: The decisions made by spouses in the AT were not informed by social norms about who within households should hold money.

However, this does not rule out the possibility that individually held behavioural prescriptions and inter-spousal differences in such prescriptions informed the AT decisions. To investigate this, we define three new variables. The first variable, which we shall call the *relative evaluation*, is the relative appropriateness of allocating to self as the evaluation made by a specific evaluator of a specific DM-type facing a specific decision choosing to allocate to self minus the evaluation made by the same evaluator of the same DM-type facing the same decision choosing to allocate to spouse. Given the way in which the action-specific evaluations are quantified, the maximum value that a relative evaluation can take is 2, indicating that allocating to self is very appropriate and allocating to spouse

¹² In debriefs with the field researchers, we heard that many spouses indicated that K30 was sufficient incentive to try and match others’ evaluations, but struggled to imagine what those others’ evaluations would be.

is very inappropriate, the minimum value it can take is -2, indicating that allocating to self is very inappropriate and allocating to spouse is very appropriate, and a relative evaluation of zero can be interpreted as indifference between the two possible actions. The second variable, which we shall call the *self-relevant prescription*, is the relative evaluation made by evaluator i of a DM of the same type as i facing AT decision 45:45. In the analysis below, this variable captures what a DM thinks he or she should do in the AT. The more positive (negative) this variable the surer the DM is that he or she should allocate to self (spouse). On average, the DMs' *self-relevant prescription* was 0.261 indicating that, on average, they thought they should allocate to self. However, it varied markedly across DMs (see Figure A1, Analytical Appendix). The third variable, which we shall call the *inter-spousal prescription difference*, is the relative evaluation made by evaluator i 's spouse of a DM of the same type as i facing AT decision 45:45 minus i 's *self-relevant prescription*. In the analysis below, this variable captures the difference between what a DM's spouse thinks the DM should do in the AT and what the DM him- or herself thinks he or she should do.¹³ If this variable is negative (positive) it indicates that a DM's spouse is less (more) inclined to think that the DM should allocate to self than the DM him- or herself. On average, the DMs' *inter-spousal prescription difference* was -0.354 indicating that, on average, the PSs were less inclined than their DM spouses to think that the DMs should allocate to self. However, this variable also varied markedly across DMs (see Figure A1, Analytical Appendix).

We would expect that a DM's own individually held behavioural prescriptions would inform his or her own behaviour. So, our *ex ante* expectation was that the higher a DM's *self-relevant prescription* the more likely they would be to allocate to self in the AT.

The foundations for our *ex ante* expectations relating to the *inter-spousal prescription difference* were a little more complicated. First, for this second variable to affect what a DM does in the AT, the DM needs to have some sense of what his or her spouse thinks. The DMs and their PS spouses had no opportunity to discuss the AT before each made their ET1 evaluations so, for this to be the case, the DM's spouse would need to successfully extrapolate from what they know about their spouses in everyday life. Let us assume that this is the case. Second, we would expect the *inter-spousal prescription difference* to impact on a DM's decision-making only when the DM knows that their spouse will, at some point, find out what they chose to do in the AT, i.e., if the DM undertook the AT under the zero deniability treatment. Under the high deniability treatment, the DM is free to do as he or she thinks appropriate, which might incorporate some aspect of what they think their spouse thinks. They do not need to consider further behavioural adjustments to accommodate aspects of what they think their spouse thinks that they have not incorporated into their own thinking because their spouse is not going to find out what they did. So, our *ex ante* expectations were that the higher the *inter-spousal prescription difference* the more likely a DM would be to allocate to self under the low deniability treatment in the AT, but that the *inter-spousal prescription difference* would have no effect on decisions made under the high deniability treatment.

Introducing these two new variables into the regression analysis of AT decisions (see Table A3 in the Analytical Appendix) yields the following results:

- as expected, the higher a DM's *self-relevant prescription* the more likely they were to allocate to self in the AT irrespective of the treatment to which they were assigned ($p=0.030$ and 0.003 for the low and high deniability treatments respectively);

¹³ We focus on the 45:45 AT decision options because, unlike the 35:45 decision options, these were evaluated by both the DMs and their PS spouses.

- as expected, the higher the *inter-spousal prescription difference* the more likely a DM was to allocate to self under the low deniability treatment in the AT; and
- as expected, the *inter-spousal prescription difference* had no effect on decisions made under the high deniability treatment.

Before concluding, it is important to mention that, while these findings are consistent with individually held behavioural prescriptions informing AT decision-making, the analysis above does not provide unequivocal support for this as a conclusion. This is because the DMs' own individually held behavioural prescriptions were elicited after they had made the AT decisions and we know that making such decisions affected the ET1 evaluations that were subsequently made (see Table A11 and Figure A2 in the Analytical Appendix) and this leads to what is commonly referred to as an endogeneity problem.

Conclusions: Individually held behavioural prescriptions informed decision-making in the AT. Specifically, the more appropriate individuals thought it was for them to allocate money to themselves rather than their spouses, the more likely they were to do precisely that. Further, when the individuals knew their spouses would find out what they had done in the AT, they took their spouses' opinions about what they should do into account when making their decisions, i.e., they compromised.

The findings about the relationship between own and spouses' individually held behavioural prescriptions and own decisions are important not only because they support the conclusions above, but also because they speak to the validity of our experimental designs. Had we been unable to identify a positive relationship between what DMs thought they should do in the AT and what they actually did in the AT, it would have cast doubt on the DMs' understanding of the AT and ET1 and the link between the two and, hence, on all of the findings generated using the AT and ET1 data that are presented above. Finding no relationship between the inter-spousal prescription difference and DMs' decisions in the AT would have been less damning. However, the fact that we do find such a relationship and that the relationship that we find is consistent with our *ex ante* expectations indicates that the spouses who participated in our research not only had a reasonable understanding of the AT and ET1 and the link between the two, but were also able to draw on experiences from their everyday lives when trying to imagine what their spouses were thinking about the tasks. This is tremendously important as it signals that our findings have external validity, i.e., can be considered relevant when thinking about husbands' and wives' behaviour in everyday life. The logic behind this is that, if the spouses consider everyday life to be of relevance when deciding what to do in the experiment, then we should consider what the spouses do in the experiment as informative about what they do and why they do what they do in everyday life.

We now turn to the ET2 data. In ET2, we applied the same methods as in ET1 but to the elicitation of behavioural prescriptions relating to everyday life decision-making scenarios. Because the methods are the same and the decision-making scenarios have more in common with everyday life, it is reasonable to assume that ET2 will have at least as much if not more external validity.

The ET2 data are presented in Figure 4. The evaluations by wives are presented in the top half of the figure. The evaluations by husbands are presented in the bottom half. The evaluations of a wife taking each of the options she has regarding what to do with a bit of spare cash are presented on the left of the figure. The evaluations of a husband taking each of the options are presented on the right. Within the top half of the figure, the upper two panels present the evaluations made in the absence of the

incentive to match and the lower two panels present the evaluations made when the incentive to match was in place. This same applies to the bottom half of the figure. Within each panel, each column presents the evaluations of a wife or husband choosing a specific option and the columns are ordered from left to right as follows:

1. give the cash to her or his spouse
2. spend the cash on treats for the family
3. hide the cash at home and tell spouse about the cash but not where it is hidden
4. hide the cash at home and don't tell spouse about the cash
5. put the money in an e-wallet and tell spouse
6. put the money in an e-wallet and don't tell spouse
7. join a savings group and tell spouse
8. join a savings group and don't tell spouse.

Within each column, the green part indicates the proportion evaluating the action as very appropriate, the yellow part indicates the proportion evaluating the action as somewhat appropriate, the orange part indicates the proportion evaluating the action as somewhat inappropriate, and the red part indicates the proportion evaluating the action as very inappropriate.

Here, unlike in the corresponding figure for ET1, for some actions we see very strong modal evaluations. Most notably, across all the panels, saving, by hiding cash at home, putting it in an e-wallet or joining a savings group, and keeping it secret from one's spouse is thought most often to be very inappropriate. In contrast, saving and telling one's spouse is thought most often to be very appropriate. Regression analyses indicate that these apparent patterns in the data are statistically significant. Saving and keeping it secret from one's spouse is significantly less appropriate than any of the other options on average and for each evaluator-type-decision-maker-type combination ($p < 0.001$ in all case, see Table A4, Analytical Appendix). Saving and telling one's spouse is significantly more appropriate than any of the other options on average and for each evaluator-type-decision-maker-type combination ($p < 0.003$ in all case, see Table A5, Analytical Appendix). Among the savings options, joining a savings group and telling one's spouse is considered most appropriate irrespective of who is evaluating whom ($p < 0.001$ on average, but not for all evaluator-type-decision-maker-type combinations, see Table A6, Analytical Appendix). Saving in an e-wallet and telling one's spouse and hiding the cash at home and telling one's spouse take second and third place respectively (see Table A6, Analytical Appendix).

These findings are indicative of the existence of a social norm forbidding saving in secret from one's spouse. Taken at face value, the data could also be signalling the existence of a social norm prescribing saving and telling one's spouse. However, given how difficult saving is for people like those participating this study, it is probably better to describe this as a collective or commonly held aspiration rather than as a social norm.

Figure 4: Evaluations of appropriateness of options regarding what to do with “spare cash” (ET2)

Evaluations by wives



Evaluations by husbands



Source: authors' data

To further our investigation into whether a social norm forbidding saving in secret from one's spouse exists, we conducted an evaluation-by-session-level analysis similar to the presented above for ET1. Across 88 sessions, a total of 352 ET2 evaluations were made, while the incentive to match was in place. Among these 352 evaluations:

- perfect matching, i.e., all participants making the same evaluation when the incentive to match was in place, was achieved in 50 (14.2%)
- partial matching, i.e., all stating that the action was either very or somewhat appropriate or all stating that the action was either very or somewhat inappropriate, was achieved in a further 72 (20.5%)
- there was a single modal evaluation in 336 (95.5%); and
- within these 336, on average, more than 6 out of the 8 participants' evaluations matched the mode.

Further, statistical regression analyses revealed that there was no more within session variation in the evaluations when the incentive to match was in place than when it was not (see Table A7 in the Analytical Appendix) and that the evaluations made in the absence and in the presence of the incentive to match are very highly correlated ($p < 0.0001$) with 63% of the variation in each being explained by the other.

The analysis above indicates that the spouses were more successful with regard to matching their evaluations in ET2 as compared to ET1. In particular, that there was a single mode in over 95% of the evaluations and that, within these, on average, more than 6 out of the 8 participants' evaluations matched that mode is consistent with the existence of a collectively held and collectively recognised behavioural prescription, i.e., with the existence of a social norm. That the evaluations made in the absence and in the presence of the incentive to match are highly correlated is consistent with the social norm having been internalised by many of the participants in our research study. However, this strong correlation combined with the finding that there was no less within session variation in the evaluations when the incentive to match was in place means that we cannot unequivocally conclude that a social norm exists.

There is one more element in our dataset that is relevant to the issue of whether a social norm forbidding saving and keeping it secret from one's spouse exists. Social norms are often enforced through the punishment of norm violators and ET3 was specifically designed to investigate whether a wife secretly belonging to a savings group or secretly saving in an e-wallet were considered justifiable grounds for wife-beating.


Figure 5 presents the ET3 data. The upper panel presents the evaluations made in the absence of the incentive to match. The lower panel presents the evaluations made in the presence of the incentive to match. In each panel, there is a pair of bars relating to each of the questions asked during the task. The order of the questions, left to right, is as follows:

1. Do you think a man is justified in beating his wife if she neglects their children?
2. Do you think a man is justified in beating his wife if she burns the food?
3. Do you think a man is justified in beating his wife if she argues with him?
4. Do you think a man is justified in beating his wife if she visits family or friends without his permission?
5. Do you think a man is justified in beating his wife if he discovers that she is saving in a mobile money e-wallet?
6. Do you think a man is justified in beating his wife if he discovers that she has joined a savings group?

Figure 5: Evaluations of justifiability of wife beating (ET3)


Incentive to match absent

Wife beating justifiable if she...



Incentive to match in place

Wife beating justifiable if she...



Source: authors' data

Within a pair of bars, the height of the left-hand dark pink bar indicates the proportion of wives responding that the man is justified in beating his wife and the height of the left-hand dark pink bar indicates the proportion of husbands responding that the man is justified in beating his wife.

The most notable feature in the ET3 data is that wives are more inclined than husbands to think wife-beating is justifiable. A similar finding was reported in the 2018 Zambian DHS, where 46% of women versus 26% of men believed that it was justifiable for wives to be beaten by their husband for any one of the following: burning food; arguing; going out without telling him; neglecting children; making major household purchases without telling him; and refusing sex. However, here, the issue under investigation is not whether and how evaluations differ depending on the gender of the evaluator, but on whether and to what extent saving in secret is considered punishable.

We focus on the evaluations made when the incentive to match was absent because, as in ET1, introducing the incentive to match increased the within session variance in evaluations, particularly for husbands (see Table A8, Analytical Appendix).

When the incentive to match was absent, 30% of wives and 15% of husbands indicated that a man beating his wife if he discovered that she was saving in an e-wallet or had joined a savings group was justifiable. After accounting for the marked gender difference, the patterns of relative justifiability across the seven scenarios are similar for husbands and wives (see notes below Table A9, Analytical Appendix). Beating a wife who is discovered to be saving in an e-wallet or to have joined a savings group is as justifiable as beating a wife for neglecting children, visiting friends or family in secret and refusing to have sex, and more justifiable than beating wife for burning the food or arguing (see Table A9, Analytical Appendix). These findings are also consistent with the existence of a social norm forbidding saving and keeping it secret from one's spouse. However, they too do not provide unequivocal evidence of its existence.

Conclusions: There is strong but not unequivocal evidence pointing to the existence of a social norm forbidding saving in secret from ones spouse. If such a social norm exists, it has been internalised by many spouses in Eastern Province. If such a norm does not exist, then an individually held behavioural prescription forbidding saving in secret from ones spouse is held by the very large majority of spouses in Eastern Province. Just under one in three wives and one in six husbands think that a man is justified in beating his wife if he discovers that she is saving in an e-wallet or has joined a savings group. Thus, as a grounds for wife beating, saving in secret is on a par with neglecting the children, visiting friends or family in secret and refusing to have sex. There is also evidence pointing to the existence of a collective or commonly held aspiration to save.

5. Specifically, if a wife chooses to hold and make independent decisions about money and financial service use, is she violating a norm?

To address this question, we draw on the analysis and findings presented above in relation to research question 4.

Conclusions: A wife choosing to hold money in the AT, i.e., choosing to allocate money to herself rather than her husband in the AT, is not violating a social norm. However, it may be inconsistent with an individual behavioural prescription that either she or her husband or both hold. There is strong but not unequivocal evidence that a wife choosing to save in secret from her husband is violating a social norm and some think that her husband would be justified in beating her under such circumstances. However, here it is important to note that it is the secrecy and not the saving that is the violation.

6. Do the social norms that women acknowledge as guides to the way they should behave differ from the social norms that men think should be guiding women's behaviour (and vice-versa)?

To address this question, we draw and build on the analysis and findings presented above in relation to research question 4.

The analysis of the ET2 data presented above, indicates the existence of a social norm forbidding saving in secret from one's spouse exists. Figure 4 and the corresponding regression analyses go on to indicate that both husbands and wives think that this behavioural prescription applies to both husbands and wives; saving and keeping it secret from one's spouse was found to be considerably and significantly less appropriate than any of the other ET2 options both on average and for each evaluator-type-decision-maker-type combination ($p < 0.001$ in all cases, see Table A4, Analytical Appendix).

However, Figure 4 and the regressions also indicate that the extent to which violations of this norm are tolerated varies depending on whether the violator is a husband or wife and on who is doing the evaluation. Wives are significantly less tolerant than husbands of secret saving by husbands ($p < 0.001$, see Table A4, Analytical Appendix) and significantly more tolerant than husbands of secret saving by wives ($p < 0.001$, see Table A4, Analytical Appendix). Correspondingly, wives think that husbands saving and telling their wives is significantly more appropriate than do the husbands themselves ($p < 0.001$, see Table A5, Analytical Appendix).

Conclusion: Assuming it exists, the social norm forbidding saving in secret from one's spouse applies to both husbands and wives and this is acknowledged by both husbands and wives. However, the extent to which norm violation is tolerated depends on who is doing the violating and who the evaluating; wives are less tolerant than husbands of secret saving by husbands and more tolerant than husbands of secret saving by wives.

There is one other aspect of Figure 4 that is worth investigating in relation to this question. Irrespective of whether one focuses on the evaluations made in the presence or absence of the incentive to match, the figure indicates that the very large majority (<85%) of both husbands and wives thought it very appropriate for husbands to give their spare cash to their spouses, i.e., their wives. The majority of husbands also thought it very appropriate for wives to give their spare cash to their spouses. However, the husbands were, on average, more tolerant of a wife deviating from this behavioural prescription than a husband deviating ($p = 0.008$, see Table A10, Analytical Appendix). In contrast, the wives had much more varied opinions about the appropriateness of a wife giving her spare cash to her husband and, on average, were even more tolerant than the husbands of a wife deviating from this behavioural prescription ($p = 0.067$, see Table A10, Analytical Appendix).


Conclusion: A social norm prescribing that husbands should give their spare cash to their wives may exist and be acknowledged by both husbands and wives. However, wives do not acknowledge the existence of a social norm prescribing that wives should give their spare cash to their husbands and, if husbands acknowledge such a norm, they are far more tolerant of wives deviating from the norm than they are of husbands deviating.

Recall that the analysis of the AT and ET1 data indicated that, while the DMs' decisions in the AT were not informed by social norms, they were informed by the DMs' own and their spouses' individually held behavioural prescriptions. This being the case, it is interesting to investigate whether the

individually held behavioural prescriptions that guided the wives' behaviour in the AT differ from those that guided the husbands' behaviour.

Figure 6 (below) presents the average relative evaluations by each evaluator-type for each DM-type-decision-type combination. Recall that the maximum value that a relative evaluation can take is 2, indicating that allocating to self is very appropriate and allocating to spouse is very inappropriate, the minimum value it can take is -2, indicating that allocating to self is very inappropriate and allocating to spouse is very appropriate, and a relative evaluation of zero can be interpreted as indifference between the two possible actions. The heights of the bars indicate the size of the average relative evaluations. The p-value reported above or below each bar relates to the null hypothesis that the mean relative evaluation is zero, i.e., that the evaluator considers allocating to self neither more nor less appropriate than allocating to spouse. The whiskers on each bar define the 95% confidence interval of the corresponding average relative evaluation. The p-values and confidence intervals are derived from a regression (see Table A11, Model 1, in the Analytical Appendix).

Figure 6: Relative evaluations of actions in the AT by types of evaluator and decision-maker



Source: authors' data

The most important and notable findings from this analysis are that:

- On average, husbands are indifferent to the allocating actions of DM husbands and close to indifferent to the actions of DM wives facing the 45:45 decision (p=0.387 and 0.071 respectively);
- On average, wives are close to indifferent to the allocating actions of DM husbands facing the 45:45 decision (p=0.090), but consider wives allocating to themselves as relatively appropriate (p<0.001);
- Wives think it significantly more appropriate for wives to allocate to themselves in the 45:45 decision than do husbands (p<0.001);
- Wives think it significantly more appropriate for wives to allocate to themselves in the 45:45 decision than for husbands to allocate to themselves (p<0.001);

- Wives think it somewhat less appropriate for husbands to allocate to themselves in the 45:45 decision than do husbands ($p=0.066$);
- In general, both wives and husbands think it less appropriate to allocate to self when doing so is costly to the household, i.e., when facing the 35:45 decision ($p=0.023$ or lower).
- However, on average, wives are indifferent to the allocating actions of DM wives facing the 35:45 decision ($p=0.132$).

Finally, it is worth noting that when we separate out the evaluations of the DM and PS spouses, we find that the PS wives' evaluations of husbands' actions are more extreme than the DM wives' evaluation. Specifically, compared to the DM wives, the PS wives think it even more appropriate that wives allocate to themselves when there is no cost of doing so and think it significantly inappropriate for husbands to do likewise (see Table A11, Model 2, and Figure A2 in the Analytical Appendix).

Conclusions: Individually held behavioural prescriptions pertaining to the AT vary depending on who is doing the prescribing, to whom the prescription applies and on the specifics of the decision being faced. Wives have much stronger opinions than husbands about what wives and husbands should do. There is strong evidence that wives think that wives should allocate to themselves when doing so imposes no cost on the household and weak evidence that they think that husbands should allocate to their wives. In general, both wives and husbands think it is inappropriate for spouses to allocate to themselves when doing so imposes a cost on the household. However, wives think that this cost is offset when it is a wife who is allocating to herself.

5. Do social norms about who within a household is supposed to hold and make decisions about money vary systematically depending on whether a community is matrilineal or patrilineal?

The analysis of the ET2 data revealed a social norm forbidding saving in secret from one's spouse (see text above and Table A4, Analytical Appendix). When we build on the analysis of the ET2 data by distinguishing between spouses living in matrilineal and patrilineal communities, we find no evidence of differences in social norm content, but several differences in the strength of the behavioural prescriptions as they apply to wives and husbands that are consistent with prior expectations.

Specifically, with regard to the social norm forbidding saving in secret from one's spouse we find:

- violations of the norm by husbands are tolerated less by both wives and husbands in patrilineal communities; and
- wives in patrilineal communities have especially varied opinions about and are, on average, close to indifference about violations of the norm by wives (see Table A12, Analytical Appendix).

With regard to the appropriateness of saving and telling one's spouse, when we distinguish between spouses living in matrilineal and patrilineal communities, we find:

- while wives in both patrilineal and matrilineal communities consider it very appropriate for husbands to save and tell their wives, the former consider it significantly more appropriate than the latter;
- wives in patrilineal communities have especially varied opinions about and are, on average, close to indifference about whether wives should save and tell their spouses (see Table A13, Analytical Appendix).

Finally, with regard to the appropriateness of giving spare cash to one's spouse, when we distinguish between spouses living in matrilineal and patrilineal communities, we find that:

- while husbands in both patrilineal and matrilineal communities consider it very appropriate for husbands to give their spare cash to their wives, the former consider it significantly more appropriate than the latter (see Table A14, Analytical Appendix).

Each of these findings is consistent with husbands in patrilineal communities being seen as having greater responsibility towards their wives and their children compared to husbands in matrilineal communities.

Conclusion: Social norms about who within a household is supposed to hold and make decisions about money do not vary systematically depending on whether a community is matrilineal or patrilineal. However, the extent to which deviations from the norms and prescriptions by wives and husbands are tolerated does vary systematically depending on whether a community is matrilineal or patrilineal and the variations are consistent with husbands in patrilineal communities being seen as having greater responsibility towards their wives and their children.

References

- Afzal, U., d'Adda, G., Fafchamps, M. and Said, F. (2016) "Gender and Agency within the Household: Experimental Evidence from Pakistan, mimeo.
- Almås, I., Armand, A., Attanasio, O., Carneiro, P. (2018) "Measuring and Changing Control: Women's Empowerment and Targeted Transfers", *The Economic Journal*, 128(612), pp. F609–F639.
- Ashraf, N. (2009), 'Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines', *American Economic Review*, 99(4), 1245-1277.
- Barr, A., M. Dekker, W. Janssens, B. Kebede, B. Kramer 'Cooperation in Polygynous Households', *American Journal of Economics: Applied Economics*, 11(2), 266-83, April 2019.
- Bicchieri, C., (2006) *The grammar of Society: The nature and dynamics of social norms*, Cambridge University Press.
- Burks, S. V. and Krupka, E. L. (2012) "A multimethod approach to identifying norms and normative expectations within a corporate hierarchy: Evidence from the financial services industry", *Management Science*, 58(1), pp. 203-217.
- Iversen, V., C. Jackson, B. Kebede, A. Munro, and A. Verschoor. 2011. "Do Spouses Realise Cooperative Gains? Experimental Evidence from Rural Uganda." *World Development* 39 (4): 569–578.
- Jakiela, P. and O. Ozier (2016), 'Does Africa need a rotten kin theorem? Experimental evidence from village economies', *The Review of Economic Studies*, 83(1), 231-268.
- Kebede, B., M. Tarazona, A. Munro, and A. Verschoor. 2014. "Intra-Household Efficiency: An Experimental Study from Ethiopia." *Journal of African Economies*, 23, 105-150.
- Krupka, E. L. and Weber, R. A. (2013), "Identifying social norms using coordination games: Why does dictator game sharing vary?" *Journal of the European Economic Association*, 11(3), pp. 495-524.

Analytical Appendix

Table A1: Regression analysis of Allocation Task decisions

Dependent variable = 1 if allocated to self in Allocation Task

	Model 1	Model 2			
		Husbands	Wives	Low deniability	High deniability
Constant		0.590*** (0.080)	0.609*** (0.080)	0.590*** (0.080)	0.570*** (0.084)
High deniability [#]	0.036 (0.025)	-0.020 (0.040)	0.091** (0.039)		
DM=Wife (dmw) [#]				0.019 (0.044)	0.130*** (0.042)
AT decision:					
40:45	-0.234*** (0.035)		-0.234*** (0.035)		
35:45	-0.217*** (0.036)		-0.217*** (0.036)		
30:45	-0.299*** (0.032)		-0.299*** (0.032)		
25:45	-0.285*** (0.038)		-0.285*** (0.038)		
20:45	-0.305*** (0.034)		-0.305*** (0.034)		
15:45	-0.342*** (0.032)		-0.342*** (0.032)		
10:45	-0.339*** (0.036)		-0.339*** (0.036)		
5:45	-0.385*** (0.036)		-0.385*** (0.036)		
0:45	-0.570*** (0.038)		-0.570*** (0.038)		
45:40	0.003 (0.025)		0.003 (0.025)		
Village fes	yes		Yes		
Script reader fes	yes		Yes		
Observations	3,861		3,861		
R-squared	0.163		0.172		

Notes: Marginal effects from two LPM estimations presented; an observation is a decision in the Allocation Task; each participant made 11 decisions; AT decision basis for comparison is 45:45; standard errors adjusted for non-independence within sessions (44) in parentheses; *** p<0.01, ** p<0.05, * p<0.1; # p-value on difference in effect across genders/treatments is 0.74.

Source: authors' data

Table A2: Regression analysis of within-session variation in ET1 evaluations

Dependent variable = SD = the within session standard deviation in an evaluation of a specific evaluator-type-DM-type-action

	Model 1	Model 2			
		Husbands evaluating DM husbands	Husbands evaluating DM wives	Wives evaluating DM husbands	Wives evaluating DM wives
Constant	0.653*** (0.014)	0.615*** (0.038)	0.605*** (0.036)	0.648*** (0.040)	0.659*** (0.046)
Incentive to match (IM)	0.046*** (0.011)	0.024 (0.014)	0.047*** (0.018)	0.027 (0.024)	0.086*** (0.026)
Patrilineal (PL)		-0.046 (0.038)	-0.027 (0.037)	0.007 (0.042)	-0.033 (0.040)
Allocate to self (AS)		0.044 (0.035)	0.048* (0.028)	0.110*** (0.035)	-0.151*** (0.037)
35:45 decision (D35:45)		-0.008 (0.029)	0.022 (0.032)	0.004 (0.034)	0.030 (0.034)
AS x D35:45		0.045 (0.045)	0.067* (0.040)	-0.012 (0.045)	0.207*** (0.046)
Experimental controls	no			Yes	
Observations	1,056			1,056	
R-squared	0.011			0.113	

Notes: An observation is an evaluation by all the participants in a session; Basis for comparison = allocate to spouse in 45:45 decision; Experimental controls = a dummy indicating that 7 (not 8) spouses participated in the session, a dummy indicating that the participants had not previously engaged in the AT and the latter interacted with the genders of the evaluator and the decision-maker; robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Source: authors' data


The specification of Model 2 is as follows:

$$\begin{aligned}
 SD_j = & \alpha_0 + \alpha_1 IM_j + \alpha_2 PL_j + \alpha_3 AS_j + \alpha_4 D35:45_j + \alpha_5 (AS \times D35:45)_j \\
 & + \beta_0 wife_j + \beta_1 (wife \times IM)_j + \beta_2 (wife \times PL)_j + \beta_3 (wife \times AS)_j + \beta_4 (wife \times D35:45)_j \\
 & + \beta_5 (wife \times AS \times D35:45)_j + \gamma_0 dmw_j + \gamma_1 (dmw \times IM)_j + \gamma_2 (dmw \times PL)_j + \gamma_3 (dmw \times AS)_j \\
 & + \gamma_4 (dmw \times D35:45)_j + \gamma_5 (dmw \times AS \times D35:45)_j + \delta_0 (dmw \times wife)_j \\
 & + \delta_1 (dmw \times wife \times IM)_j + \delta_2 (dmw \times wife \times PL)_j + \delta_3 (dmw \times wife \times AS)_j \\
 & + \delta_4 (dmw \times wife \times D35:45)_j + \delta_5 (dmw \times wife \times AS \times D35:45)_j + \mathbf{C}\boldsymbol{\theta} + u_j
 \end{aligned}$$

where: $wife_j = 1$ if evaluators were wives and zero otherwise; $dmw_j = 1$ if decision-maker in scenario being evaluated was a wife and zero otherwise; \mathbf{C} is a vector of experimental controls (see table note for details); u_j is the error term; and the other variables are as defined in Table A2.

The effect of the incentive to match is positive and significantly different from zero on average and for both husband evaluators and wife evaluators when evaluating wives' actions after controlling for the characteristics of the action being evaluated and whether the evaluator lives in a patrilineal community. The effect of the incentive to match is positive but not significantly different from zero for both husband evaluators and wife evaluators when evaluating husbands' actions after controlling for the characteristics of the action being evaluated and whether the evaluator lives in a patrilineal community. Adding village fixed effects leaves the results reported above qualitatively unchanged.

Figure A1: The distributions of self-relevant prescriptions and inter-spousal prescription differences across DMs



Source: authors' data

Table A3: Regression analysis of Allocation Task decisions including behavioural prescriptions

Dependent variable = 1 if allocated to self in Allocation Task

	Model 2 (from Table A1)		Model 3	
	Low deniability	High deniability	Low deniability	High deniability
Constant	0.590*** (0.080)	0.570*** (0.084)	0.595*** (0.068)	0.568*** (0.073)
DM=wife (dmw)	0.019 (0.044)	0.130*** (0.042)	-0.002 (0.041)	0.093** (0.041)
<i>Self-relevant prescription</i>			0.077** (0.034)	0.082*** (0.026)
<i>Inter-spousal prescription difference</i>			0.038** (0.018)	2.5e ⁻⁴ (0.016)
DM=wife (dmw)	0.019 (0.044)	0.130*** (0.042)	-0.002 (0.041)	0.093** (0.041)
AT decision fes	yes		Yes	
Village fes	yes		Yes	
Script reader fes	yes		Yes	
Observations	3,861		3,850	
R-squared	0.172		0.195	

Notes: Marginal effects from two LPM estimations presented; an observation is a decision in the Allocation Task; each participant made 11 decisions; 11 observations are lost between Models 2 and 3 owing to the departure of one wife after she had engaged in the AT but before she had engaged in ET1; standard errors adjusted for non-independence within sessions (44) in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Source: authors' data

Table A4: Regression analysis of ET2: Is saving and keeping it secret from spouse less appropriate than other options?

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Incentive to match absent					
Constant	0.600*** (0.039)	0.586*** (0.041)	0.545*** (0.079)	0.805*** (0.025)	0.467*** (0.097)
Save and keep secret	-1.111*** (0.080)	-1.172*** ^a (0.064)	-0.985*** ^b (0.137)	-1.563*** ^{ac} (0.076)	-0.729*** ^{bc} (0.200)
Observations	2,808			2,808	
R-squared	0.362			0.390	
Incentive to match in place					
Constant	0.583*** (0.038)	0.615*** (0.031)	0.465*** (0.078)	0.799*** (0.030)	0.455*** (0.031)
Save and keep secret	-1.132*** (0.085)	-1.188*** ^a (0.071)	-0.947*** ^b (0.146)	-1.633*** ^{ac} (0.053)	-1.132*** ^{bc} (0.218)
Observations	2,808			2,808	
R-squared	0.367			0.399	

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; coefficients from linear regressions presented; basis for comparison = non-saving and save and tell spouse options; coefficients from linear regressions presented; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; ^a diff. in marginal effect between husbands evaluating husbands and wives evaluating husbands sig at $p < 0.001$; ^b diff. in marginal effect between husbands evaluating wives and wives evaluating wives sig at $p < 0.001$; ^c diff. in marginal effect between wives evaluating husbands and wives evaluating wives sig at $p < 0.001$.

Source: authors' data

The specification of Model 1 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 SaveSecret_v + u_{ijv}$$

where:

- $Eval_{ijv}$ = i 's evaluation of DM j taking option v ;
 $SaveSecret_v$ = 1 if j hides cash at home and does not tell spouse, saves in an e-wallet and does not tell spouse, or joins a savings group and does not tell spouse; = 0 if j takes any of the other five options; and
 u_{ijv} = the error term.

The specification of Model 2 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 SaveSecret_v + \beta_2 wife_i + \beta_3 dmw_j + \beta_5 (wife * SaveSecret)_{iv} + \beta_5 (wife * dmw)_{ij} + \beta_6 (dmw * SaveSecret)_{jv} + \beta_7 (dmw * wife * SaveSecret)_{ijv} + u_{ijv}$$

where:

- $wife_i$ = 1 if evaluator i is a wife; = 0 if evaluator i is a husband; and
 dmw_j = 1 if decision-maker j , whose actions are being evaluated, is a wife; = 0 if j is a husband.

Table A5: Regression analysis of ET2: Is saving and telling spouse more appropriate than other options?

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Incentive to match absent					
Constant	-0.066** (0.032)	-0.041 (0.045)	-0.044 (0.066)	-0.159*** (0.044)	-0.023 (0.082)
Save and tell spouse	0.667*** (0.063)	0.501*** ^a (0.050)	0.587*** (0.112)	1.008*** ^a (0.070)	0.576*** (0.173)
Observations	2,808	2,808			
R-squared	0.131	0.143			
Incentive to match in place					
Constant	-0.099*** (0.036)	-0.024 (0.054)	-0.080 (0.056)	-0.218*** (0.040)	-0.074 (0.105)
Save and tell spouse	0.686*** (0.070)	0.517*** ^a (0.062)	0.507*** (0.122)	1.080*** ^a (0.050)	0.645*** (0.192)
Observations	2,808	2,808			
R-squared	0.135	0.152			

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; coefficients from linear regressions presented; basis for comparison = non-saving and save and keep secret options; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; ^a diff. in marginal effect between husbands evaluating husbands and wives evaluating husbands sig at $p < 0.001$.

Source: authors' data

The specification of Model 1 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 SaveSecret_v + u_{ijv}$$

where:

- $Eval_{ijv}$ = i 's evaluation of DM j taking option v ;
 $SaveSecret_v$ = 1 if j hides cash at home and does not tell spouse, saves in an e-wallet and does not tell spouse, or joins a savings group and does not tell spouse; = 0 if j takes any of the other five options; and
 u_{ijv} = the error term.

The specification of Model 2 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 SaveTell_v + \beta_2 wife_i + \beta_3 dmw_j + \beta_5 (wife * SaveTell)_{iv} + \beta_5 (wife * dmw)_{ij} + \beta_6 (dmw * SaveTell)_{jv} + \beta_7 (dmw * wife * SaveTell)_{ijv} + u_{ijv}$$

where:

- $wife_i$ = 1 if evaluator i is a wife; = 0 if evaluator i is a husband; and
 dmw_j = 1 if decision-maker j , whose actions are being evaluated, is a wife; = 0 if j is a husband.

Table A6: Regression analysis of ET2: Appropriateness of different types of saving

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Incentive to match absent					
Constant	0.299*** (0.079)	-0.015 (0.121)	0.174 (0.154)	0.747*** (0.070)	0.295* (0.164)
Ewallet	0.353*** (0.076)	0.523** (0.204)	0.432*** (0.136)	0.130* (0.076)	0.326** (0.134)
Savings group	0.551*** (0.072)	0.902*** (0.112)	0.674*** (0.139)	0.176** (0.079)	0.447*** (0.135)
H ₀ : coeff on Ewallet = coeff on Savings group	p<0.001	p=0.003	p=0.076	p=0.009	p=0.057
Observations	1,053	1,053			
R-squared	0.105	0.173			
Incentive to match in place					
Constant	0.345*** (0.078)	0.136 (0.121)	0.030 (0.155)	0.808*** (0.063)	0.409*** (0.146)
Ewallet	0.251*** (0.077)	0.348 (0.209)	0.379** (0.177)	0.077 (0.060)	0.197* (0.104)
Savings group	0.477*** (0.075)	0.720*** (0.127)	0.811*** (0.145)	0.084 (0.059)	0.288** (0.126)
H ₀ : coeff on Ewallet = coeff on Savings group	p<0.001	p=0.006	p=0.007	p=0.710	p=0.112
Observations	1,053	1,053			
R-squared	0.076	0.162			

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; sample restricted to saving and telling spouse options; coefficients from linear regressions presented; basis for comparison = hide cash at home; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Source: authors' data

The specification of Model 1 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 Ewallet_v + \beta_2 SavGrp_v + u_{ijv}$$

where:

$Eval_{ijv}$ = i 's evaluation of DM j taking option v ;
 $Ewallet_v$ = 1 if j saves in an e-wallet; = 0 if j hides cash at home or joins a savings group;
 $SavGrp_v$ = 1 if j joins a savings group; = 0 if j hides cash at home or saves in an e-wallet; and
 u_{ijv} = the error term.

The specification of Model 2 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 Ewallet_v + \beta_2 SavGrp_v + \beta_3 wife_i + \beta_4 dmw_j + \beta_5 (wife * Ewallet)_{iv} + \beta_6 (wife * SavGrp)_{iv} + \beta_7 (wife * dmw)_{ij} + \beta_8 (dmw * Ewallet)_{jv} + \beta_9 (dmw * SavGrp)_{jv} + \beta_{10} (dmw * wife * Ewallet)_{ijv} + \beta_{11} (dmw * wife * SavGrp)_{ijv} + u_{ijv}$$

where $wife_i$ and dmw_j are defined as before.

Table A7: Analysis of within-session variation in ET2 evaluations

Dependent variable = SD = the within session standard deviation in an evaluation of a specific option

	Model 1	Model 2			
		Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Constant	0.524*** (0.026)	0.457*** (0.080)	0.591*** (0.069)	0.595*** (0.078)	0.658*** (0.083)
Incentive to match (IM)	0.007 (0.015)	0.013 (0.029)	-0.002 (0.035)	0.003 (0.032)	0.014 (0.031)
Patrilineal (PL)		-0.059 (0.054)	-0.059 (0.054)	-0.031 (0.035)	-0.031 (0.035)
To spouse (TS)		-0.115 (0.090)	-0.116 (0.102)	-0.426*** (0.098)	0.103 (0.070)
Hide at home (HH)		0.330*** (0.088)	0.053 (0.064)	-0.245** (0.096)	0.000 (0.065)
Mobile money (MM)		0.151 (0.162)	-0.096 (0.125)	-0.387*** (0.096)	-0.146 (0.100)
Savings group (SG)		-0.137 (0.104)	-0.266** (0.116)	-0.430*** (0.092)	-0.240* (0.135)
Keep saving secret (S)		-0.118 (0.078)	-0.047 (0.118)	-0.078 (0.131)	0.053 (0.111)
MM x S		0.161 (0.114)	0.212* (0.124)	0.141** (0.067)	0.130 (0.090)
SG x S		0.436*** (0.057)	0.389*** (0.144)	0.263** (0.119)	0.270** (0.100)
Experimental controls	no		Yes		
Observations	704		704		
R-squared	<0.001		0.314		

Notes: An observation is an evaluation by all the participants in a session; Basis for comparison = treats for family; Experimental controls = a dummy indicating that 7 (not 8) spouses in the session, a dummy indicating that the order of the possible actions was reversed and the latter interacted with gender of evaluator and evaluated; robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Source: authors' data

The specification of model 2 is as follows:

$$\begin{aligned}
SD_j = & \alpha_0 + \alpha_1 IM_j + \alpha_2 PL_j + \alpha_3 TS_j + \alpha_4 FT_j + \alpha_5 MM_j + \alpha_6 SG_j + \alpha_7 S_j + \alpha_8 (MM \times S)_j + \alpha_9 (SG \times S)_j \\
& + \beta_0 wife_j + \beta_1 (wife \times IM)_j + \beta_2 (wife \times PL)_j + \beta_3 (wife \times TS)_j + \beta_4 (wife \times FT)_j \\
& + \beta_5 (wife \times MM)_j + \beta_6 (wife \times SG)_j + \beta_7 (wife \times S)_j + \beta_8 (wife \times MM \times S)_j \\
& + \beta_9 (wife \times SG \times S)_j + \gamma_0 dmw_j + \gamma_1 (dmw \times IM)_j + \gamma_2 (dmw \times PL)_j + \gamma_3 (dmw \times TS)_j \\
& + \gamma_4 (dmw \times FT)_j + \gamma_5 (dmw \times MM)_j + \gamma_6 (dmw \times SG)_j + \gamma_7 (dmw \times S)_j + \gamma_8 (dmw \times MM \times S)_j \\
& + \gamma_9 (dmw \times SG \times S)_j + \delta_0 (dmw \times wife)_j + \delta_1 (dmw \times wife \times IM)_j + \delta_2 (dmw \times wife \times PL)_j \\
& + \delta_3 (dmw \times wife \times TS)_j + \delta_4 (dmw \times wife \times FT)_j + \delta_5 (dmw \times wife \times MM)_j \\
& + \delta_6 (dmw \times wife \times SG)_j + \delta_7 (dmw \times wife \times S)_j + \delta_8 (dmw \times wife \times MM \times S)_j \\
& + \delta_9 (dmw \times wife \times SG \times S)_j + \mathbf{C}\boldsymbol{\theta} + u_j
\end{aligned}$$

where: $wife_j = 1$ if evaluator is a wife; $dmw_j = 1$ if a wife is being evaluated; \mathbf{C} is a vector of experimental controls (see note for Table ?? for details); u_j is the error term; and the other variables are as defined in Table ??.

The effect of the incentive to coordinate is insignificantly different from zero on average and for each possible evaluator-type-decision-maker-type combination both before (not shown) and after controlling for the characteristics of the evaluated option and whether evaluator lives in a patrilineal or matrilineal community. Adding village fixed effects leaves results qualitatively unchanged.

Table A8: Analysis of within-session variation in ET3 evaluations

Dependent variable = SD = the within session standard deviation in an evaluation of a specific option

VARIABLES	Model 1	Model 2	
		Husbands evaluating	Wives evaluating
Constant	0.341*** (0.019)	0.329*** ^a (0.039)	0.426*** ^a (0.028)
Incentive to match (IM)	0.045*** (0.014)	0.074*** ^a (0.023)	0.016 ^a (0.012)
Patrilineal (PL)	-	-0.075* ^a (0.043)	0.047 ^a (0.038)
BurnFood	-	-0.041** (0.019)	-0.056** (0.028)
Argue	-	-0.025 ^a (0.033)	-0.124*** ^a (0.034)
VisitFamFrnd	-	0.028 (0.033)	0.006 (0.022)
RefuseSex	-	-0.065** (0.027)	-0.048 (0.034)
EWallet	-	0.005 (0.035)	-0.051 (0.036)
SavGrp	-	0.015 (0.041)	-0.073* (0.037)
Observations	616	616	
R-squared	0.016	0.168	

Notes: An observation is an evaluation by all the participants in a session; Basis for comparison = Wife neglects the children; Experimental controls = a dummy indicating that 7 (not 8) spouses participated in the session; robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Source: authors' data

The specification of Model 2 is as follows:

$$SD_j = \alpha_0 + \alpha_1 IM_j + \alpha_2 PL_j + \alpha_3 BurnFood_j + \alpha_4 Argue_j + \alpha_5 VisitFamFrnd_j + \alpha_6 RefuseSex_j + \alpha_7 EWallet_j + \alpha_8 SavGrp_j + \beta_0 wife_j + \beta_1 (wife \times IM)_j + \beta_2 (wife \times PL)_j + \beta_3 (wife \times BurnFood)_j + \beta_4 (wife \times Argue)_j + \beta_5 (wife \times VisitFamFrnd)_j + \beta_6 (wife \times RefuseSex)_j + \beta_7 (wife \times EWallet)_j + \beta_8 (wife \times SavGrp)_j + \mathbf{C}\theta + u_j$$

where: $wife_j = 1$ if evaluators were wives and zero otherwise; \mathbf{C} is a vector of experimental controls (see table note for details); u_j is the error term; and the other variables are as defined in the Table.

Table A9: Analysis of ET3 evaluations made in the absence of the incentive to match

Dependent variable = Just = 1 if evaluator responded that wife beating was justified

	Model (1)	H0: coeff on variable in row = coeff on...	
		Ewallet (2)	SavGrp (3)
Constant	0.160*** (0.027)		
Wife	0.148*** (0.028)		
Patrilineal	-0.021 (0.028)		
BrnFood	-0.048** (0.024)	0.049**	0.084*
Argue	-0.060** (0.028)	0.046**	0.048**
VisitFamFrnd	0.054** (0.025)	0.172	0.140
RefuseSex	-0.031 (0.024)	0.176	0.221
Ewallet	0.011 (0.031)		0.803
SavGrp	0.006 (0.030)	0.803	
Individual fes	yes		
Observations	2,457		
R-squared	0.043		

Notes: An observation is an evaluator i responding to one of the ET3 questions; sample restricted to responses given in absence of incentive to match; coefficients from one linear probability presented in column (1); standard errors adjusted for non-independence within sessions (88) in parentheses; basis for comparison = wife neglects children; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; column (2) presents results (p-values) of tests of equivalence between coefficient on Ewallet and coefficients on each of the other scenario identifiers; column (3) presents results (p-values) of tests of equivalence between coefficient on SavGrp and coefficients on each of the other scenario identifiers.

Source: authors' data

The specification of the model in column (1) is as follows:

$$Just_{ij} = \alpha_0 + \alpha_1 wife_i + \alpha_2 patrilineal_i + \beta_1 BrnFood_j + \beta_2 Argue_j + \beta_3 VisitFamFrnd_j + \beta_4 RefuseSex_j + \beta_5 EWallet_j + \alpha_6 SavGrp_j + u_{ij}$$

where: $wife_i = 1$ if evaluator was a wife and zero otherwise; $patrilineal_i = 1$ if evaluator lived in a patrilineal community and zero otherwise; u_{ij} is the error term; and the other variables identify the scenario j .

When added to this model, a set of six interaction terms, each involving $wife$ and one of the scenario identifiers, is jointly insignificant ($p=0.276$).

Table A10: Regression analysis of ET2: For whom is giving cash to the spouse more appropriate than other options?

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Incentive to match absent					
Constant	0.124*** (0.025)	0.047 (0.042)	0.132** (0.059)	0.119*** (0.031)	0.198*** (0.052)
Give cash to spouse	0.479*** (0.077)	0.802*** a (0.046)	0.353** ab (0.154)	0.805*** c (0.058)	-0.039 bc (0.141)
Observations	2,808	2,808			
R-squared	0.031	0.049			
Incentive to match in place					
Constant	0.102*** (0.024)	0.071 (0.044)	0.058 (0.046)	0.087*** (0.030)	0.192*** (0.053)
Give cash to spouse	0.450*** (0.084)	0.785*** a' (0.059)	0.411** a'b' (0.166)	0.798*** c' (0.031)	-0.192 b'c' (0.147)
Observations	2,808	2,808			
R-squared	0.027	0.050			

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; coefficients from linear regressions presented; basis for comparison = non-saving and save and keep secret options; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; ^a diff. in marginal effect between husbands evaluating husbands and husbands evaluating wives sig at $p = 0.008$ (^{a'} diff. sig at $p = 0.04$); ^b diff. in marginal effect between husbands evaluating wives and wives evaluating wives sig at $p = 0.067$ (^{b'} diff. sig at $p = 0.009$); ^c diff. in marginal effect between wives evaluating husbands and wives evaluating wives sig at $p < 0.001$ (^{c'} diff. sig at $p < 0.10$).

Source: authors' data

The specification of Model 1 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 ToSpouse_v + u_{ijv}$$

where:

$Eval_{ijv}$ = i 's evaluation of DM j taking option v ;

$toSpouse_v$ = 1 if j gives cash to spouse; = 0 if j takes any of the other seven options; and

u_{ijv} = the error term.

The specification of Model 2 is as follows:

$$Eval_{ijv} = \beta_0 + \beta_1 ToSpouse_v + \beta_2 wife_i + \beta_3 dmw_j + \beta_5 (wife * ToSpouse)_{iv} + \beta_5 (wife * dmw)_{ij} + \beta_6 (dmw * ToSpouse)_{jv} + \beta_7 (dmw * wife * ToSpouse)_{ijv} + u_{ijv}$$

where $wife_i$ and dmw_j are defined as above.

Table A11: Regression analysis of relative evaluations of allocating to self in the AT when the incentive to match is not in place

Dependent variable = evaluation of to-self - evaluation of to-spouse

	Model 1	Model 2	
		Engaged in AT prior to ET1	Did not engage in AT
Constant	0.059 (0.068)	0.057 (0.082)	0.061 (0.108)
Wife	-0.190* (0.102)	-0.015 (0.119)	-0.365** (0.158)
Dmw	0.027 (0.080)	0.008 (0.095)	0.045 (0.129)
wife*dmw	0.779*** (0.145)	0.415*** (0.143)	1.143*** (0.224)
D35:45	-0.455*** (0.102)	-0.453*** (0.080)	-0.453*** (0.080)
wife*D35:45	0.196 (0.151)	0.021 (0.138)	0.021 (0.138)
dmw*D35:45	0.121 (0.125)	0.140 (0.114)	0.140 (0.114)
wife*dmw* D35:45	-0.673*** (0.185)	0.057 (0.082)	0.061 (0.108)
Observations	2,100	2,100	
R-squared	0.083	0.093	

Notes: An observation is an evaluator i evaluating the options faced by a decision-maker of type j when facing decision d ; coefficients from linear regressions presented; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Source: authors' data

The specification of Model 1 is as follows:

$$Reval_{ijd} = \beta_0 + \beta_1 wife_i + \beta_2 dmw_j + \beta_3 (wife * dmw)_{ij} + \beta_4 D35:45_d + \beta_5 (wife * D35:45)_{id} + \beta_6 (dmw * D35:45)_{jd} + \beta_7 (wife * dmw * D35:45)_{ijd} + u_{ijd}$$

where:

- $Reval_{ijd}$ = i 's relative evaluation of DM j allocating to self instead of to spouse in decision d ;
 $wife_i$ = 1 if evaluator i is a wife; = 0 if evaluator i is a husband;
 dmw_j = 1 if decision-maker j , whose actions are being evaluated, is a wife; = 0 if j is a husband;
 $D35:45_d$ = 1 if j is facing AT decision 35:45; = 0 if j is facing AT decision 45:45; and
 u_{ijd} = the error term.


Including evaluator fixed effects in the estimation of Model 1 yields almost identical estimates of β_2 to β_7 and almost identical corresponding standard error estimates. Estimating Model 1 using the evaluations made when the incentive to match was in place yields qualitatively similar results.

The specification of Model 2 is as follows:

$$Reval_{ijd} = \beta_0 + \beta_1 wife_i + \beta_2 dmw_j + \beta_3 (wife * dmw)_{ij} + \beta_4 D35:45_d + \beta_5 (wife * D35:45)_{id} + \beta_6 (dmw * D35:45)_{jd} + \beta_7 (wife * dmw * D35:45)_{ijd} + \gamma_0 AT_i + \gamma_1 (AT * wife)_i + \gamma_2 (AT * dmw)_{ij} + \gamma_3 (AT * wife * dmw)_{ij} + u_{ijd}$$

where AT_i = 1 if evaluator i engaged in the AT prior to ET1 and = 0 if i did not engage in the AT.

Figure A2: Relative evaluations of actions in AT decision 45:45, distinguishing between DM and PS spouses



Note: This figure is derived from Model 2, Table A11.

Source: authors' data

Table A12: Regression analysis of ET2: Differential appropriateness of saving and keeping it secret in Matrilineal and Patrilineal communities

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Matrilineal					
Constant	0.644*** (0.031)	0.656*** (0.035)	0.617*** (0.035)	0.764*** (0.037)	0.513*** (0.081)
Save and keep secret	-1.216*** (0.069)	-1.068*** ^a (0.093)	-1.122*** (0.161)	-1.560*** ^b (0.082)	-1.074*** (0.055)
Patrilineal					
Constant	0.522*** (0.067)	0.567*** (0.045)	0.339*** (0.117)	0.843*** (0.042)	0.406*** (0.144)
Save and keep secret	-1.048*** (0.154)	-1.333*** ^a (0.065)	-0.802*** (0.214)	-1.723*** ^b (0.033)	-0.507 (0.365)
Observations	2,808	2,808			
R-squared	0.371	0.413			

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; coefficients from linear regressions presented; basis for comparison = non-saving and save and tell spouse options; coefficients from linear regressions presented; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; ^a diff. in marginal effect between evaluations made by members of matrilineal and patrilineal communities sig at $p = 0.024$; ^b diff. in marginal effect between evaluations made by members of matrilineal and patrilineal communities sig at $p = 0.072$.

Source: authors' data

Table A13: Regression analysis of ET2: Differential appropriateness of saving and telling spouse in Matrilineal and Patrilineal communities

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Matrilineal					
Constant	-0.092* (0.047)	0.078 (0.069)	-0.050 (0.097)	-0.197*** (0.059)	-0.210** (0.094)
Save and tell spouse	0.746*** (0.066)	0.473*** (0.087)	0.656*** (0.151)	1.003*** ^a (0.071)	0.854*** (0.077)
Patrilineal					
Constant	-0.106* (0.054)	-0.147*** (0.042)	-0.106 (0.063)	-0.244*** (0.050)	0.039 (0.162)
Save and tell spouse	0.626*** (0.122)	0.569*** (0.082)	0.383** (0.169)	1.176*** ^a (0.036)	0.470 (0.330)
Observations	2,808			2,808	
R-squared	0.137			0.162	

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; coefficients from linear regressions presented; basis for comparison = non-saving and save and tell spouse options; coefficients from linear regressions presented; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; ^a diff. in marginal effect between evaluations made by members of matrilineal and patrilineal communities sig at $p = 0.035$.

Source: authors' data

Table A14: Regression analysis of ET2: Appropriateness of giving cash to spouse in Matrilineal and Patrilineal communities

Dependent variable = evaluation (very appropriate (1.00), somewhat appropriate (0.33), somewhat inappropriate (-0.33), very inappropriate (-1.00))

	Model 1	Model 2			
	All evaluations	Husbands evaluating husbands	Husbands evaluating wives	Wives evaluating husbands	Wives evaluating wives
Matrilineal					
Constant	0.119*** (0.030)	0.169*** (0.046)	0.117 (0.075)	0.081* (0.044)	0.107 (0.073)
Give cash to spouse	0.551*** (0.082)	0.692*** ^a (0.091)	0.633*** (0.192)	0.780*** (0.052)	0.026 (0.052)
Patrilineal					
Constant	0.085** (0.036)	-0.045 (0.037)	0.010 (0.048)	0.094** (0.039)	0.262*** (0.064)
Give cash to spouse	0.347** (0.144)	0.895*** ^a (0.059)	0.226 (0.166)	0.821*** (0.031)	-0.373 (0.147)
Observations	2,808		2,808		
R-squared	0.030		0.058		

Notes: An observation is an evaluator i evaluating an options faced by a decision-maker of type j ; coefficients from linear regressions presented; basis for comparison = non-saving and save and keep secret options; standard errors adjusted for non-independence within sessions (88) in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$; ^a diff. in marginal effects between matrilineal and patrilineal communities sig at $p = 0.035$.

Source: authors' data