

Bargaining Breakdown: Intra-Household Decision-Making and Women's Employment

Matt Lowe and Madeline McKelway*

October 23, 2017

Abstract

Women in the developing world often lack the power to make key household decisions. This comes at a cost – myriad evidence suggests that the preferences of women are more aligned with development goals than those of men. We use a field experiment to test whether the lack of decision-making power of wives in India is due to a lack of information, or a lack of communication with husbands. We partnered with India's largest carpet manufacturer to offer employment opportunities to 495 married women. Gender differences in preferences meant there was an intra-household tension: women were often interested in working outside of the home, while their husbands opposed the idea. We experimentally varied how the job opportunity was presented to couples. To test for the effects of information, and the incentives of husbands to withhold it, we randomized whether enrollment tickets and job information were given to the women or to their husbands. For the non-targeted spouse, we cross-randomized whether they were informed about the job opportunity, giving variation in whether husbands had plausible deniability. To test for the importance of communication, some couples received the ticket and information together, with a chance to discuss the job. Overall, enrollment was low at 17%. Information was not a barrier to enrollment – providing women with information about the opportunity had no effect because husbands did not strategically withhold information, despite having plausible deniability. Surprisingly, we find that having couples discuss the opportunity together decreased enrollment, by 6 to 9 percentage points. We conclude that policymakers should tread with care: intra-household communication may not be easily manipulated without unintended consequences for decision-making.

*Department of Economics, MIT. E-mail: mlope@mit.edu, mckelway@mit.edu. Many thanks to Daron Acemoglu, Abhijit Banerjee, Esther Duflo and Frank Schilbach for helpful comments and suggestions. Thanks also to Bakul Chugh, Maddy Dutta and Aditya Madhusudan for outstanding research assistance. We are grateful for financial support from the Weiss Family Fund. This RCT was pre-registered in the AEA registry with ID AEARCTR-0001678.

1 Introduction

Policymakers are often faced with the challenge of convincing households to enroll in programs that would achieve development goals (e.g. education, health care, or financial services). But in doing so, policymakers often confront frictions within the household. In many cases, husbands and wives have different preferences for these programs, the wife has less bargaining power, and the wife's preference is more aligned with development goals than the husband's. An implication is that policymakers should present opportunities in a way that would encourage women's involvement in the household take-up decision. But little is known about how exactly policymakers could do so. Would simply informing women about the opportunity in addition to or instead of husbands allow women to take a larger role in decisions? Should policymakers explicitly encourage or facilitate joint deliberation between husbands and wives? Could gestures that symbolically place the decision in women's hands make the decision actually be in women's hands? And more generally, are household take-up decisions even susceptible to such small nudges from policy?

We investigate these questions in the context of one important household decision in rural India: female labor supply. By partnering with a large carpet export firm, we are able to offer full time weaving employment for women. Evidence suggests employment for women would improve household well-being on a number of dimensions and help in the achievement of many development goals ([Heath and Jayachandran \(2016\)](#)). However, female labor force participation (FLFP) in this setting represents a decision in which spouses' preferences strongly and systematically diverge. In particular, often wives are interested in working outside of the home, but their husbands are opposed to them doing so. Given the low bargaining power of women in this setting, intra-household frictions often bar women from working. Our setting is therefore one in which policymakers would want to encourage women's involvement in the household's take-up decision. By orchestrating all recruitment activities for six of the firm's new women weaving centers, we are able to experimentally vary how the employment opportunity is presented to households. In doing so, we test the effects of several policies to encourage greater involvement of women in household decision-making.

Women's employment is a high-stakes household decision that is of independent interest. Understanding household decisions regarding women's employment in India is particularly important, given that the country's FLFP is puzzlingly low¹ and, by some measures, even falling.² Further, certain features of this household decision make it generalizable to other decisions of interest. In particular, employment

¹Figure 1 uses country-level 2014 data from the World Bank to plot the relationship between FLFP and log GDP per capita. India has low FLFP by global standards, and low FLFP even conditional on its GDP per capita.

²See "[Why Aren't India's Women Working?](#)" where Pande and Moore report that from 2004 to 2011, India's FLFP fell from 31 to 24 percent, despite its economy growing at an average annual rate of 7 percent.

decisions are public within the household; they cannot be made by bypassing one's spouse and therefore must be made collectively. The public nature of this decision closely approximates many important decisions households make (e.g. education of children, migration). While such decisions are certainly desirable to study, our context is rare in offering the ability to randomize aspects of this high-stakes decision. This gives us the ability to cleanly identify treatment effects and to test mechanisms underlying household behavior.

The first step of our experiment was to print unique tickets that enabled specific women to enroll in the job with our partner firm; no woman could enroll without her unique ticket, and no ticket could be used to enable enrollment of any other woman but the woman for whom it was printed. Importantly, we randomly excluded some women from receiving tickets, and this was common knowledge. As will be seen, this gave certain ticket recipients plausible deniability: whenever the ticket recipients' spouse would not be informed about the ticket, the recipient could plausibly deny having received it.

We then randomized whether the ticket was given to the woman or to her husband, and cross-randomized which of three information sets the non-ticketed spouse was assigned to. Those in the first group received no information, enabling their ticketed spouses to withhold information about the wife's eligibility to enroll in the job. A second group received full information but received it separately from the ticketed spouse, eliminating the possibility of information withholding but still allowing the possibility of conversation avoidance. The final group received full information at the same time as the ticketed spouse and then the two were given an opportunity to discuss the decision in the presence of a surveyor. In all cases, the ticketed spouse was aware of which information group their spouse would be in.

This design allows us to test the effects of three ways policymakers might present programs to encourage women's involvement in the household take-up decision. The first policy would inform women about the opportunity instead of or in addition to husbands. In many developing country settings, and particularly in ours, women rarely leave the home; this makes it much easier to inform men about an opportunity than women. But when husbands are opposed to take-up of a particular opportunity, they have an incentive to withhold information about the opportunity from their wives so that the household does not enroll. In this case, there may be gains from informing women directly about the opportunity. However, incentives to withhold information interact crucially with bargaining power. Husbands may not need to withhold information if they know final decisions will reflect their preferences regardless of their wives' knowledge of the opportunity. If so, there would be no gains from informing women.

Next, we test a policy that facilitates joint deliberation between husbands and wives. Even when there is full information about an opportunity, it is possible that husbands will de facto make decisions unilaterally. If so, there may be gains to kickstarting a decision-making process that is joint. To do so, policymakers could bring husbands and wives together to discuss an opportunity in a moderated

environment. Such an environment would foster a cordial conversation about a particular opportunity. Even if true opinions are not conveyed during the discussion, this could help individuals realize their spouses have opinions about the decision, encourage them to seek opinions of their spouse later on, and thereby kickstart a more joint decision-making process. On the other hand, encouraging joint decision-making may backfire: husbands may resent suggestions of equality in discussion or forced discussion could produce arguments. Further, joint decisions may sacrifice efficiency of unilateral decisions (e.g. by requiring negotiation costs) and make it more difficult for a decision to be made at all.

Finally, we test the effect of policies that encourage women's involvement in decision-making by symbolically putting the decision in their hands. Many opportunities require that households have completed physical forms or obtained physical vouchers before households can enroll. In such cases, policymakers have the option of symbolically empowering women in a decision by initially distributing forms or vouchers to wives. The predicted effect of such gestures is, however, ambiguous. They may raise take-up by nudging households to give women more say in decision-making. But on the other hand, such gestures may backfire if they cause husbands to take offense or erode husbands' perceptions of ownership over take-up decisions.

Overall enrollment was low at 17%, and none of these policies raised enrollment rates. First, there is no evidence that husbands strategically withhold information from their wives; they fully disclose information about the employment opportunity and there are no gains to providing the information to women in addition to or instead of husbands. We argue this is because substantial bargaining power allows husbands to make decisions without needing to withhold information.

Second, bringing women to the decision-making table by encouraging discussion lowered enrollment by 6 to 9 percentage points. To explain this, we argue that the discussion forces wives to express early interest in the job, when usually they would test the waters, gradually persuading the husband of their preferences over time. By expressing immediate interest, wives signal to their husbands that they are "bad" types – those that break gender norms – and this leads husbands to veto enrollment altogether.

Finally, we find that symbolically empowering women in decision-making by giving wives enrollment tickets backfires in the short-run, dramatically raising dropout rates. We present evidence that this is due to men being more supportive and accommodating of a decision they feel was their own. Taken together, these results suggest that the intra-household decision-making process is a fragile one. Policymakers should take caution when considering how to present programs in a way that would encourage women's involvement in the household take-up decision.

Our findings relate most closely to literature on intra-household decision-making. A large literature suggests that increasing women's bargaining power improves household welfare along a number of dimensions such as family health, children's educational attainment, food consumption, and female labor

supply (Duflo (2003), Duflo and Udry (2004), Field et al. (2016), Qian (2008), and Thomas (1990)). An implication is that policymakers may want to present opportunities to households in a way that lead women to have greater involvement in the take-up decision. However, literature on how they might do so is sparse. We contribute by providing novel evidence on three policies in particular: (i) informing women about an opportunity instead of or in addition to their husbands; (ii) encouraging discussion about the opportunity in a moderated environment so as to kickstart subsequent joint decision-making; and (iii) using symbolic gestures to place decisions in women’s hands. Results in Ashraf et al. (2014) speak to the effects of (i). In particular, Ashraf et al. (2014) provide evidence that women may strategically withhold contraception decisions from their husbands when given the opportunity to do so. Importantly, the authors document strategic information withholding on the part of the spouse with less bargaining power; it is unclear whether spouses with more bargaining power withhold information and therefore unclear whether household decisions would change as a result of policies that simply provide women with information about an opportunity. Indeed, combining our results with results of Ashraf et al. (2014) provides evidence of an important interaction between strategic withholding and bargaining power. In another related paper, Ashraf (2009) uses an intervention akin to (ii), varying whether spouses are together or separate when making observable savings decisions. Results suggest that allowing discussion can indeed shift household decisions. But there is an important distinction between our intervention and Ashraf’s. Participants in Ashraf’s study must make their decision immediately after discussing, whereas ours make their decision several days later. Our intervention seeks to use an initial discussion to influence the nature of subsequent decision-making, and therefore serves a very different purpose than Ashraf’s.

In sum, we find null or negative effects of the policy nudges we consider. Though discouraging, these results complement the literature, which typically finds gains from empowering women. In particular, gains from shifting women’s intra-household bargaining power over extended periods (as in Duflo (2003), Duflo and Udry (2004), Field et al. (2016), Qian (2008), and Thomas (1990)) are not easily replicated in the short-run by increasing women’s role in a particular decision. Attempting to present opportunities in a way that explicitly encourages women’s involvement in take-up decisions deserves careful thinking and caution.

The remainder of this paper is organized as follows. The subsequent section provides relevant information on our context and setting. We then detail our experimental timeline and interventions in Section 3. Section 4 summarizes our data and sample, while Section 5 details the results. We conclude in Section 6.

2 Context

2.1 Gender Norms in Uttar Pradesh

Our study takes place in rural villages in eastern Uttar Pradesh, India. Uttar Pradesh is one of India's poorest states; the median husband in our sample earns \$68 per month.³ This is also a setting with strong gender norms. Purdah (the practice of women veiling their faces and staying out of sight of men and strangers) is an important feature of local culture. In our pilot data, 82% of women said they practiced purdah and 86% said the practice is very or rather important. Arranged marriage and patrilocality is commonplace: at the time of marriage, women leave their native villages and become a part of a previously unknown family in a new village. Married women are typically confined to their homes and responsible for all household chores, including childrearing, cooking, tending to livestock, and household cultivation. Whereas 82% of husbands in our sample had done activities to earn income in the previous 3 months, only 13% of their wives had done so. This norm is strongly ingrained in cultural mindsets, with 87% of men and 85% of women in our sample believing that husbands should earn more income than their wives.⁴ These factors, along with ample anecdotal evidence, suggest women have very little bargaining power within the household.

2.2 Partner Firm and Women's Job Opportunity

We partner with a large carpet manufacturing and exporting firm working in this region. Carpet weaving has been a major occupation in the region since the time of the British rule,⁵ when the British set up the industry to take advantage of low labor costs. Weaving is generally considered a low caste occupation and, as with most formal-sector employees in this setting, weavers are predominantly male.⁶

The firm imports wool, usually from Rajasthan, India, to a factory in Uttar Pradesh, where it is converted to yarn. The firm supplies yarn to thousands of loom owners located in villages across Uttar Pradesh. These loom owners operate small looms in their villages (see Figure 2) that employ local men to weave carpets by hand. The firm purchases woven carpets from the loom owners, packages them, and ships them to buyers globally. The firm now hopes to establish looms in villages that are owned and staffed entirely by women. They have several motivations for doing so. First, a larger pool of labor would allow the firm to more easily take advantage of periods of high export demand for carpets. Second, men increasingly migrate to cities for work, resulting in local labor shortages. Third, the firm has a long

³Converted using the market exchange rate during the study period of Rs. 66 per \$1.

⁴Though below we will show that other work-related preferences between wives and husbands systematically differ.

⁵The region is in fact known as the "Carpet Belt" of northern India.

⁶74% of weavers in Uttar Pradesh-based respondents to the Indian Human Development Survey (2005) are male.

history of promoting development in the region through corporate social responsibility projects (e.g. they have sponsored the construction of many sanitary latrines in local villages). They view staffing the new looms with women as a means of empowering women to promote development of the region.

The firm has now arranged for the construction of several new village looms that are owned by women. They recruit women to staff each loom from the village surrounding the loom. As women usually have no prior experience in weaving, the job begins with a four-month training period. The firm pays the women during this time.⁷ Overall, this job offers impoverished households a respectable and stable wage.⁸ By many objective measures, the job is desirable: it is near potential employees' homes, involves safe and comfortable work, requires reasonable hours, and demands no prior training. To respect gender norms, the firm takes many steps to ensure that women would interact only with other women while at work. Qualitative evidence suggests participants also see the job as desirable: 88% say workers in this job would be completely safe, and only 11% say the job is low status. Nevertheless, the firm faces very low enrollment and retention rates. Our experiment finds average enrollment rates of 17% and average dropout rates (conditional on enrollment) of 48%.

2.3 Gender Differences in Preferences Regarding Women's Employment

Intra-household issues are at the heart of low take-up. Ample anecdotal evidence suggests that women themselves are very interested in taking the job but their husbands will not allow. Figure 3 visualizes reasons for not enrolling provided in our sample of married couples, and Figure 4 visualizes reasons for dropping out provided by women in many looms (both in and out of our experimental sample). Both husbands and wives report that a key reason for the wife not enrolling is that her husband or family (i.e. the husband's parents and siblings) would not allow: this is the third most important reason husbands provide and the second most important reason wives provide. It is also the fourth most common reason for dropping out. Indeed, our partner firm believes opposition from husbands is a key reason for why women do not take and stay in the job.

The most common reason provided for both not enrolling and for dropping out is that there is no one else to do household chores. It is likely that this reason also reflects the bargaining position of the husband and the husband's lack of support for the job. In particular, had the husband been willing to take on more household chores or ask his family to do so, the wife may have been able to work.

More generally, data from our baseline survey (detailed below) provide quantitative evidence of sys-

⁷Initially, the women were to be paid monthly in cash. However, India's demonetization occurred in the middle of our study, resulting in a rushed transition toward paying directly into bank accounts.

⁸Even during the training period, the daily wage offered to women is only slightly below that which a skilled male weaver would earn.

tematic differences in preferences towards women’s employment within the household. We asked both husbands and wives, separately, how appropriate they thought it would be for men or women in their household to hold a full-time job outside of the house in three different occupations: construction, weaving and teaching. The former and latter were chosen to reflect the most and least, respectively, male-dominated occupations in this area,⁹ whilst weaving was chosen to match the actual job opportunity we offer in the experiment. Respondents give their answer on a 0 to 2 scale: inappropriate, somewhat appropriate, or completely appropriate. By asking a question about appropriateness (rather than actual interest in a job) before we advertised the actual weaving position, we hoped to elicit a reasonably accurate measure of the gender work norms held in this area.

To investigate the extent to which spouses disagree about women’s employment, we first estimate how closely correlated a married couple’s responses are. We report in Table 1 regressions of the wife’s work norm answers on her husband’s answers. These simple bivariate regressions show that there is a significant positive correlation across spouses, consistent with either positive assortative marriage matching (arranged marriage matching progressive men with progressive women) or ex-post social influence within the marriage (spousal peer effects). That said, the correlation is quite small: it is around 0.23 for women weavers. It follows that in this context there is reasonable variation in how much spouses agree about work norms.

We next investigate whether the nature of this disagreement is that husbands are systematically more opposed to women working than women themselves are. The first six columns of Table 2 compare husbands and wives responses to these questions, by regressing responses on an indicator for the respondent being a wife along with couple fixed effects. In the first three columns, we see that there is no systematic disagreement by gender in perceived appropriateness of men working in the three occupations.¹⁰ In contrast, the next three columns show that there is a systematic divergence of preferences by gender in perceived appropriateness of women working. Women feel it would be significantly more appropriate for women in their homes to work in all three jobs than their husbands do.¹¹ The divergence is somewhat smaller for teaching, consistent with it being a less male-dominated profession. The divergence is highest for weaving, our occupation of interest.

The final column considers a variable measured on the baseline survey after information was given about the weaving job opportunity for women with our partner firm. Wives are asked how interested they

⁹Amongst Uttar Pradesh-based respondents to the Indian Human Development Survey (2005), 95% of construction workers and 72% of teachers are male.

¹⁰The omitted mean is high: both wives and husbands on average think it is between somewhat and completely appropriate for men in their households to work in either occupation. The mean is especially high for teaching, presumably reflecting the occupation’s high pay and prestige.

¹¹The omitted mean is in each case notably lower, making the perceived appropriateness of women working on average lower than that for men.

are in the job, whilst their husbands are asked how interested they are in their wives taking the job. Again, answers are on a 0 to 2 scale: not interested, somewhat interested, or very interested. Consistent with the gender differences in perceived appropriateness, wives report significantly more interest in the actual job than their husbands. In addition to this, the level of interest is reasonably high, with 57% of women reporting being very interested and 20% being somewhat interested. Taken together, these results suggest that (i) spouses disagree about women’s employment in general and for our specific job opportunity; (ii) the nature of this disagreement is that husbands are more opposed to the opportunity than are wives;¹² and (iii) women’s interest in our specific job opportunity is high.

2.4 Intra-household Decision-Making and Women’s Employment

Our study investigates intra-household decision-making regarding women’s employment in rural Uttar Pradesh. Evidence presented in this section suggests this is an ideal setting for studying intra-household decision-making. This is a cultural context that features strong gender norms and highly unequal intra-household bargaining power. There are also clear differences in spousal preferences regarding women’s employment decisions, with wives being systematically more supportive of themselves or other women joining the labor force than are their husbands. Finally, women’s employment represents a very real-world, high-stakes decision a household must make. Further, the setting is particularly well suited to answer the specific policy questions outlined in the introduction; not only do policy objectives align more closely with wives’ preferences than husbands’, but we also have the ability to experimentally test policies designed to give women more involvement in the household decision.

3 Experimental Design

3.1 Timeline

We conducted the experiment in conjunction with the opening of six new weaving centers, each of which had slots for 20 women weavers. The firm gave permission for us to orchestrate all recruitment activities for these centers. Recruitment for and opening of each center occurred sequentially between September 2016 and January 2017. We now describe the 3 to 4 week process used for each center’s recruitment.

¹²Note we take as a premise that this difference exists but our experiment cannot to speak to why it exists. It could be because the earnings of women are more of a threat to the gender identity of the husbands than the women themselves (Bertrand et al. (2015)), or because household chores would have to be divided more equally should a woman work. Given that married couples live in the husband’s native village, it could also be that men differentially bear the social cost of women violating gender norms regarding work.

3.1.1 Census

The first step was to conduct a census of the catchment area for the center. The catchment area was defined by the loom owner as the area from which the firm would have recruited women in the absence of the study. These areas typically consisted of the entire village surrounding the loom, but excluded hamlets inhabited by high castes.¹³ We visited each home in the catchment area and surveyed the household head, asking him or her to list all adults in the household along with their genders, ages, marital status, and contact information. A catchment area's census typically took 4 to 7 days. Using the census data, we identified all women in the firm's required age range for job eligibility (ages 18 to 30) along with a "pair" for the woman.¹⁴ If the woman was married, the pair was her husband; otherwise it was her household head. Note that while we included eligible unmarried women in recruitment, our analysis focuses on decision-making in married couples and excludes unmarried pairs. We then randomly assigned treatment at the couple level, stratifying by village, hamlet, and caste.

3.1.2 Baseline and Intervention

The second step was a baseline survey which contained our experimental intervention. This lasted for 6 to 8 days and occurred 1 to 6 weeks after the census ended. It involved individual surveys of all eligible women, and separate individual surveys of their pairs. These surveys began with a surveyor meeting the participant at his or her home. Importantly, the surveyor and the participant then moved to a completely private place in the yard or fields near the home where they could not be overheard. To ensure households' approval of female participants being alone with a surveyor, all surveyors were female. Each surveyor was randomly assigned a group of participants in a randomly ordered list, subject to the two members of each pair being assigned to different surveyors. On the evening before each day of baseline surveying, surveyors would schedule surveys for the following day by calling their assigned participants in the assigned order. If the participant was not available the following day, the surveyor scheduled the appointment for the earliest possible time during the baseline period. If the participant could not be reached over the phone, the surveyor would visit the home the following day to either take the survey or schedule the survey at the soonest available time. Participants that could not be found or were not available during this baseline period were considered to have attrited from the study.¹⁵ The

¹³High castes tend to not see weaving as a job appropriate for their class and are particularly opposed to women working outside of the home.

¹⁴In addition, pairs were dropped if either the woman or her pair was not available for surveys in the next month.

¹⁵This procedure was slightly different for the first center. In this case, we followed the same ordering and surveyor assignment procedure but appointments were not set in advance. We adopted the appointment-setting starting with the second center to reduce attrition.

random ordering of participants prioritized married couples so that if surveyors ran out of time it was unmarried participants that were excluded.¹⁶

The baseline survey itself had two parts. The first part was a questionnaire that asked about demographics, employment, and attitudes towards women's employment. The second part of the survey was our experimental intervention. The intervention recruited women for the job according to our experimental design (detailed below).

3.1.3 Enrollment Day

On the day after the baseline survey ended, we hosted an enrollment day at the weaving center. Any woman wishing to enroll in the job was required to come with her pair to the weaving center between 7am and 7pm on that day. Those that enrolled were also required to present unique enrollment tickets given to one member of the pair during the baseline survey (detailed below). The requirement for women to attend with their pairs is important, for both practical and conceptual reasons. Practically, it eliminated scenarios in which a woman would enroll without her husband's permission, thereby reducing dropout and subsequent intra-household discord. Conceptually, it means we can interpret enrollment as a decision made jointly by the household.

3.1.4 Endline

Finally, we conducted an endline survey in the 3 to 5 days following enrollment. The purpose of this survey was to help us understand how enrollment decisions had been made. To this end, it included a quiz about job information to assess participants' knowledge of the job¹⁷ along with questions on the pair's decision-making process. Endline surveys were generally done over the phone. From the third center onwards, we were able to conduct endline surveys in person for any participants not reachable over the phone. After about 40% of our data was collected, it was clear that treatments did not have the effects on enrollment that were initially expected. To best understand these patterns, we added several questions to the endline survey at this point. We amended our AEA pre-registration to reflect this change. As a result, some of the endline responses presented below are missing for the first 40% of couples surveyed.

¹⁶Our schedule was restricted by the firm's need to open centers by a particular date. This meant that we could not guarantee canvassing of all those eligible within the time allotted for a center's recruitment activities.

¹⁷Surveyors did not tell participants whether answers were correct or incorrect.

3.2 Treatments

3.2.1 Plausible Deniability

Our intervention involved randomly varying the style of recruitment each pair received. We began by printing enrollment tickets. Each ticket had the names of a particular pair written on it along with a unique identification number. No woman could enroll without her unique ticket and no-one but the woman for whom a particular ticket was printed could enroll using the ticket. Crucially, tickets were only printed for 90% of randomly chosen women. The remaining 10% of pairs received baseline surveys but following the questionnaire, a randomly chosen member of the pair was told that the eligible woman in the pair had not received a ticket. Any participants in the 90% assigned to receive information about the ticket (via randomizations detailed below) were also told that some participants would be surveyed and not receive tickets. It was therefore common knowledge that we had not printed tickets for all eligible women and that being surveyed did not signal a ticket had been printed. This meant no one knew if a particular woman could enroll unless told by a surveyor or by someone else who was told by a surveyor and chose to share the information. This system ensures plausible deniability: whenever only one spouse was informed about the ticket (determined by treatment assignments detailed below), he or she could plausibly deny having received the ticket. Non-ticketed pairs exist only for this purpose and our analysis focuses on pairs that actually received tickets.

3.2.2 Assigning a Ticketed Spouse

The delivery of the ticket and job information to married couples varied according to two, cross-randomized treatments.¹⁸ The first treatment determined which spouse would receive the ticket: 50% of the time the ticket was given to the wife, and 50% of the time to the husband. Anyone who received the ticket received full information about the job details, the enrollment process, and what his or her pair would be told by a surveyor.

3.2.3 Information Given to the Non-Ticketed Spouse

The second, cross-randomized treatment determined what information the non-ticketed spouse was told by a surveyor about the ticket and job opportunity. In one-third of couples, the non-ticketed spouse was told nothing, and his or her survey ended immediately after the baseline questionnaire. We refer

¹⁸We followed a simpler procedure for unmarried couples as they were to be excluded from analysis. In a randomly chosen 50% of pairs, the eligible woman was notified of whether a ticket had been printed for her. If so, she also received the ticket and job details. In the other 50%, the household head received this information. The non-ticketed member of the pair received a baseline survey but no further information.

to couples in this treatment as *NoInfo*, meaning the non-ticketed spouse received no information. Importantly, ticketed spouses in *NoInfo* could withhold information about job eligibility if they desired; a surveyor would never tell their non-ticketed spouses that a ticket had been given and the ticketed spouse could plausibly deny having received a ticket because 10% of women did not have tickets printed. Because tickets were required to enroll, such withholding was a means by which one spouse could prevent enrollment without ever having to argue his or her case.

In another one-third of couples, the non-ticketed spouse was told that his or her pair had or would receive an enrollment ticket along with details of the job and enrollment process. We refer to couples in this treatment as *Info*. A priori, we expected any effect of this group would be driven by provision of information about the ticket rather than job and enrollment details as ticket information is specific to an individual couple, whereas job and enrollment details could be spread across households in a village. However, we used an additional randomization to allow for the possibility that providing job and enrollment details would have a treatment effect beyond the effect of providing information about the ticket alone. In particular, we split the *Info* treatment in two: in 50% of *Info* couples (one-sixth of the full sample) the non-ticketed spouse was told that his or her spouse had or would receive an enrollment ticket but did not receive job and enrollment details, while the rest of non-ticketed spouses in *Info* couples received this information plus details. While there are some significant differences in endline knowledge of job details between these two groups, there are no significant differences in enrollment or decision-making outcomes. We therefore combine these two subgroups into a single *Info* group for analysis. In contrast to *NoInfo* couples, both spouses in *Info* couples knew that the wife could enroll in the job and no information could be withheld. That said, spouses in *Info* were told about the opportunity separately and the intervention did not affect how they interacted with one another.

The final one-third of couples were assigned to the *Discuss* group. Here the non-ticketed spouse was present while the ticket and job details were given to the ticketed spouse. The survey then paused for a full three minutes and the couple was encouraged to discuss the opportunity together. Data from surveyor evaluations find 80% of *Discuss* couples did indeed discuss the opportunity during this time.¹⁹ Two surveyors were present during the discussion but remained silent and did not provide any additional information about the job. At the end of the discussion, the surveyor handed the enrollment ticket to the ticketed spouse. The goal of the discussion was not to provide additional information about the opportunity or to make spouses share their true preferences with one another; surveyors remained silent during the discussion and their presence may have made couples less likely to share their true opinions. Rather, the goal of the discussion treatment was to kickstart a joint decision-making process that could later continue in private. Whilst couples in both *Info* and *Discuss* had full information about the opportunity, only

¹⁹The other 20% either remained silent or discussed something else.

the *Discuss* treatment directly affected how spouses interacted with one another about the opportunity.

A key logistical difference between this treatment and the others is that husbands and wives had to be together when job information was given, but separate when taking the baseline questionnaire. Initially, we scheduled husbands and wives surveys simultaneously, separated the spouses to individually take the baseline questionnaire, and rejoined for the job information and discussion. However, we became concerned that this procedure could introduce selective attrition as surveyors would reveal treatment status when setting appointments. Balance tables (see next section) do not suggest this introduced selective attrition but, nevertheless, we modified the procedure roughly 25% of the way through the sample to eliminate the possibility of selective attrition. Under the new procedure, spouses were contacted individually to complete the baseline questionnaire in the same manner in which individuals in all other treatment groups were contacted. After the questionnaire was complete, the participant was told that surveyors had limited time to complete all surveys and therefore wanted to complete the second half of the survey with the participant's spouse present once the spouse had taken the individual questionnaire. 100% of couples that completed individual baseline questionnaires were able and willing to schedule a second joint appointment. The only remaining concern is that postponing of job information meant that, on average, *Discuss* couples received job information closer to the enrollment date. We show below that results are unaffected by the addition of fixed effects for days between enrollment and ticket receipt.

In total we have six treatments. These are visualized in Figure 5.

4 Data and Sample

4.1 Sample Size and Characteristics

Our census identified 817 married women that were eligible for the job. Tickets were printed for 732 (90%) of them. The 817 women and their husbands formed the sample to be approached for the baseline survey and intervention, but the remainder of the paper focuses only on couples with tickets. Couples were considered to have attrited from the study if either spouse could not take the baseline survey. The requirement that both spouses be surveyed raised attrition rates but was enforced because our treatments would be difficult to interpret if some couples had only one spouse approached. Overall, 32% of the 732 couples attrited from the study. As visualized in Figure 6, the vast majority (75%) of attrition was due to one or both members of the couple being out of town, away all day for work, or otherwise unavailable during the 6 to 8 day baseline period. The two other leading causes for attrition were surveyors running

out of time to complete baseline surveys for one or both members of the couple (11% of attrition),²⁰ or one or both members of the couple not consenting to take the survey (9% of attrition). This produced a sample of 495 married couples (990 adults) for analysis. The first column of Table 3 finds rates of attrition do not differ significantly by treatment group,²¹ which is consistent with participants not knowing what treatment group they belonged to prior to taking the survey. This is reassuring in light of concerns mentioned above regarding the possibility of selective attrition in 25% of the *Discuss* sample. Of the 990 married adults in the sample, 830 (84%) completed endline surveys. Figure 6 shows the vast majority (87%) of endline attrition was due to the adult being unreachable over the phone or in person. 10% of attrition was due to an adult not consenting to participate in the endline survey, and 4% was due to the participant not having a phone.²² The first column of Table 4 tests for differential rates of endline attrition by treatment group. There is slightly higher attrition from the endline sample in *Discuss*, *Wife Tick* and *Info, Husb. Tick* groups but differences are not jointly significant.

Table 3 also includes characteristics of the 495 couples in the analysis sample, separately by treatment group. The average wife is 26 years old and her husband 29 years. About 39% of women and 13% of men have no education. These couples have 2.3 children on average, with 40% of them having been married for over 10 years, 32% for 5 to 10 years, and 28% for less than 5. As we might expect, far fewer women in the sample were employed than men (13% vs. 82%).²³ The final two columns summarize stated appropriateness of women weavers on a 0 to 2 scale.²⁴ The final row of the table provides a p-value for the joint test of equality of the six means. There is some imbalance in wife's age, marriage of over 10 years, wife's employment, and husband's stated appropriateness for women weavers (4 of 12 tests reject at the 10% level). Participants did not know their treatment status before taking baseline surveys so could not have selectively attrited. The only exception is the 25% of *Discuss* couples mentioned above. *Discuss* couples do have slightly higher levels of husbands' stated approval of female weavers but otherwise imbalances are not driven by *Discuss* groups.²⁵ We view these imbalances as unfortunate, but a result of

²⁰Enrollment dates were set in advance and therefore imposed a hard deadline on when baseline surveys had to be completed. 96% of couples excluded from the study because surveyors ran out of time came from the first center, before we had precise estimates of how long baseline surveys would require and before we had implemented the system of appointment setting. As mentioned above, baseline surveys were conducted in a random order in all villages so attrition for this reason is not selective.

²¹The p-value for the joint test of equality of attrition levels across treatment groups (conditional on strata) is 0.6.

²²This was only a reason for attrition for the first two centers as starting from the third we conducted endline surveys in person for those that did not have phones.

²³Where employment is defined as having done activities in the last 3 months to earn income.

²⁴Note this is the same variable used in Tables 1 and 2.

²⁵Note that imbalance on women's employment is driven by the *Discuss, Wife Tick* group. However, women's employment in the *Discuss, Husb. Tick* group is similar to all other groups. Because there could only have been selection into *Discuss* and not into *Discuss, Wife Tick* or *Discuss, Husb. Tick* in particular, we view this as unfortunate random sampling variation and not driven by selection.

random sampling variation. Nevertheless, we estimate all regressions both with and without the full set of covariates presented in this table, and show our results to be similar in either case. Table 4 presents the corresponding tests in the sample of adults that completed the endline survey. The imbalances in this table are similar to those in the baseline balance check, suggesting attrition out of endline surveys was not selective.

4.2 Outcome Variables

The main outcome variables are enrollment and retention in the job. The first outcome we consider is whether an eligible woman enrolled in the job on enrollment day. However, 9 married couples that received tickets contacted our team after enrollment day to say they wanted to enroll in the job but could not attend enrollment day.²⁶ Most of these couples contacted the team within one week of the enrollment day and all within one month of enrollment day. We therefore create an alternative outcome variable for enrolling on or within a month of enrollment day. In February 2017, we assessed one-month retention rates of all enrollees. This data provide two additional outcome variables: dropped out within one month of starting work (observed only for those than enrolled on or within one month of enrollment day), and enrolled on or within one month of enrollment day and stayed in the job for at least one month.

To test directly for information withholding we use a dummy for the non-ticketed spouse knowing at endline that a ticket had been given to his or her spouse. Note this question was asked of adults in couples that did and did not receive tickets, so was phrased in a way that would not give away the correct answer. In particular, the surveyor first asked whether the participant him or herself was given a ticket when surveyed a few days prior, and then asked whether the participant's spouse had received a ticket when surveyed a few days prior. To address the concern that respondents might not have believed they had plausible deniability, we also look at whether the ticketed spouse in *NoInfo* or *Info* groups answered a baseline comprehension question about what his or her spouse would know correctly. As mentioned above, all ticketed spouses were fully informed about what their non-ticketed spouse would learn from a surveyor. To gauge and enhance understanding, the surveyor also asked comprehension questions during this section and told the participant whether the answer was correct or incorrect. Answers to this question can be seen as a proxy for whether ticketed spouses' initial reaction was to believe what the surveyor said the non-ticketed spouse would know.²⁷

²⁶The reasons provided were: husband was working on enrollment day (2 cases), one or both were not in town on enrollment day (2 cases), child was sick on enrollment day (2 cases), and the couple confused the date of or procedure for enrollment (3 cases).

²⁷Note this variable is only a proxy for initial reaction, not final beliefs, as the participant was told whether the answer was correct or incorrect immediately after answering.

A next class of outcome variables measure intra-household decision-making with data from the endline survey. We use: the number of discussions the married couple had about the opportunity after the first discussion,²⁸ who (husband, wife, or both) initiated these discussions,²⁹ who (husband, wife, or both) had ultimate power to make the enrollment decision, how much spouses disagreed about whether the wife should enroll (on a 0 to 2 scale), and how inconsiderate one’s spouse was of one’s own opinion (on a 0 to 2 scale).³⁰ Whenever we consider one of these decision-making outcomes, we use the average response provided by the husband and wife. If one member of the couple did not take an endline survey, we use the answer from the spouse that did take the survey. We define decision-making outcomes at the couple level in this way for several reasons. First, decision-making outcomes pertain to the decision process of a couple so are inherently couple-level outcomes. Second, this method avoids underweighting couples in which one spouse’s answer is missing, either because the spouse did not take an endline survey or because the spouse provided an answer on the survey that prevented the question of interest from being asked.³¹ Lastly, self reports of decision-making are noisy; relying on spousal cross-reports by averaging spouses’ outcomes whenever possible reduces this noise.

Finally, we consider perceptions of overall job desirability at endline (measured on a 0 to 2 scale). We consider perceived desirability of husbands and wives separately given gender differences in preferences about the job.

4.3 Empirical Specification

Our main empirical specification takes the form:

$$y_i = \alpha_s + \beta T_i + \gamma X_i + \varepsilon_i$$

²⁸Number of discussions is measured by asking for the number of times a participant discussed the opportunity with his or her spouse. If the participant was in the *Discuss* group, we asked for the number of times the two discussed the opportunity outside of the discussion during the survey, and then added one to this number if the surveyor reported that the couple discussed the opportunity during their survey. We then subtract one from the number of discussions to compute the number after the first. We do so because couples in the *Discuss* group were encouraged to have a conversation in front of the surveyor and this discussion is not necessarily comparable to a first discussion in other treatment groups.

²⁹For couples in *Discuss*, this referred to initiating conversation outside of the conversation with the surveyor

³⁰The data on ultimate power and considerateness were only collected for the latter three centers, as explained in Section 3.1.4.

³¹Adults that said they had no discussions with their spouse were not asked who initiated discussion. Adults outside of the *Discuss* treatment who said they had no discussions were not asked how considerate their spouse was. Non-ticketed adults in *NoInfo* that did not believe they had received a ticket were not asked who had ultimate power over the enrollment decision, how much the couple disagreed about the enrollment decision, or who currently had the ticket.

where i denotes a married couple, T_i denotes a vector of treatment indicators, and α_s are strata fixed effects.³² X_i includes the 12 covariates presented in the balance checks (Table 3 and 4). Every regression is estimated with and without X_i to show that our results are not likely driven by the aforementioned balance issues. y_i is one of the outcome variables discussed in the previous section, and robust standard errors are used unless otherwise stated.

5 Results

5.1 Information Withholding

We first test whether spouses strategically withhold information to manipulate enrollment decisions. Since husbands are generally more opposed to their wives enrolling than wives themselves are, it is plausible that husbands would strategically withhold job information from their wives in order to prevent enrollment. On the other hand, husbands in this setting might have such strong bargaining power that whether they share or hide the information is irrelevant; they can fully disclose information about the opportunity and still veto enrollment. We first look at enrollment effects to test for the information withholding hypothesis.

Recall that ticketed spouses in *NoInfo* can withhold information about the wife’s ability to enroll in the job but ticketed spouses in *Info* cannot. If information withholding is strategic, it should occur more when men receive tickets. Our test of strategic information withholding is therefore a difference in difference, comparing changes in enrollment rates between *Info* and *NoInfo* when husbands versus wives receive tickets. We include *Discuss* couples in this regression to improve power but postpone discussion of the *Discuss* treatment effect to the next section. In particular, we use the specification

$$y_i = \alpha_s + \beta_1 Info_i + \beta_2 Husb.Tick_i + \beta_3 (Info_i \times Husb.Tick_i) + \beta_4 Discuss_i + \beta_5 (Discuss_i \times Husb.Tick_i) + \gamma X_i + \epsilon_i$$

where y_i is an enrollment outcome at the couple-level, $Info_i/Discuss_i = 1$ if the couple is assigned to the *Info/Discuss* treatment group, and $Husb.Tick_i = 1$ if the husband was given the job ticket. If husbands use information withholding as a way to prevent their wives from enrolling in the job, we would expect

³²Randomization was stratified by village, hamlet, and caste to create 36 total strata. After allowing for baseline attrition, some strata do not have all six treatment groups represented. To make use of all observations in identifying treatment effects, we pooled strata that did not have one member of each treatment group. Whenever required, we pooled two strata in the same village and caste but different hamlets. Pooling across different hamlets rather than across castes or villages keeps members of a strata as similar to one another as possible. After pooling, we have 25 strata.

to find $\beta_2 < 0$, $\beta_3 > 0$.

The test is shown in the first four columns of Table 5. On both measures of enrollment, with and without covariates, β_2 and β_3 are small and insignificant. There is no evidence here that husbands strategically withhold information to prevent their wives from enrolling. The next two columns test withholding more directly by using as an outcome a dummy for the non-ticketed spouse knowing about the ticket. Non-ticketed spouses in the *Info* group are more likely to know about the ticket, but this knowledge does not differ based on the recipient of the ticket. These results firmly reject the hypothesis that husbands hide ticket receipt from their wives to manipulate enrollment decisions.

Though the opportunity to withhold information was stressed to ticketed spouses in *NoInfo* groups (i.e. they were told their spouse would be surveyed, that their spouse would not receive information about the ticket during this survey, and that many women in the village did not have tickets printed for them), one possibility is that they were doubtful that the secret would be kept by surveyors. The final two columns of Table 5 test this indirectly by considering responses to a comprehension question. After we told ticketed spouses what we would tell their spouse, we asked a comprehension question about what their spouse would be told by surveyors about the ticket and job opportunity. We find that roughly 90% answered this question correctly in *NoInfo* groups, and that confusion was actually more likely in the *Info* groups, where the question was answered correctly far less often.³³ There is therefore not evidence that ticketed participants in *NoInfo* were inherently doubtful that they could plausibly deny having received the ticket should they choose to do so.

In this setting, husbands do not strategically withhold information from their wives. We believe this to be consistent with husbands having such strong bargaining power in the household that they do not need to manipulate information to have a decision go in their favor. A policy corollary of this is when take-up requires a joint decision, informing wives about opportunities (instead of or in addition to husbands) may not move take-up.

5.2 Discussion

Next we investigate the effect of the *Discuss* treatment. Recall this treatment kickstarts collective decision-making by providing job information to both spouses together and giving the two three minutes to discuss the opportunity. Table 6 provides effects on enrollment. Given insignificant differences between the *NoInfo* and *Info* treatment groups, we pool these into a single omitted category. We find large and significant negative effects: kickstarting collective decision-making lowers enrollment 6 to 9 p.p., a sta-

³³In all cases, after answering the comprehension question the surveyor would give the actual correct answer. So even in the *Info* groups respondents should understand the exact extent of information sets.

tistically significant effect with either enrollment measure, and with or without covariates. Given low overall levels of enrollment, this amounts to an enormous 35 to 47% gutting of enrollment rates.

The remaining columns of Table 6 test whether the effect of *Discuss* differs based on which spouse received the ticket. The point estimate on *Discuss* is negative regardless of who got the ticket, though the point estimate is 7 or 8 p.p. more negative if the husband received the ticket. However, this effect is not statistically significant. Subsequent tables therefore investigate why the discussion treatment lowers enrollment, without distinguishing between couples in which the husband versus wife was given the ticket.

5.2.1 Explaining the Discussion Effect

To understand why the *Discuss* group backfires, first note that outside of the *Discuss* treatment, the vast majority (86%) of couples report having discussed the opportunity with their spouse; it cannot then be that the discussion treatment started a discussion when there would have been none, but rather changed the nature of that discussion. We propose instead the following explanation: given strong gender norms in the region, when wives show too much interest in the job, they signal to their husbands that they are a “bad” type of woman – in the most extreme case, a woman that seeks relations with other men. Consistent with this, anecdotal evidence suggests that the safety concerns expressed by husbands are at least partially a mask for their fears of their wives meeting other men. If this signalling exists, women interested in the job have to convey this interest carefully to their husbands – they should test the waters, rather than declare outright interest. The more interest they show, the more they signal they are a bad type, leading the husband to veto enrollment altogether. By testing the waters, they might bring the husband to believe that it is his own idea, and ultimately have him allow her to enroll.

In this model, the *Discuss* group forces the wife to reveal her type – by encouraging immediate discussion, wives who are interested signal to their husbands immediately that they are bad types, lowering enrollment. But why can’t wives test the water even in the *Discuss* group? One argument is that the presence of the surveyors change the dynamics of the discussion. The husband expresses more support for the job opportunity since the surveyors are listening, and this causes the wife to feel able to express her interest. In the long-run, this reduces enrollment.

In general, spouses have strategies of how to successfully persuade each other of their preferences, and the *Discuss* group prevents one type of strategy that the wife might use. This story is one plausible explanation for the Discussion effect. In ongoing fieldwork, we are using vignette experiments to test this mechanism directly.

5.2.2 Competing Explanations

Tables 7 and 8 rule out competing explanations. One potential explanation is that couples in *Discuss* on average learned about the enrollment opportunity later, which affected their enrollment behavior. Table 7 replicates the first four columns of Table 6 but includes fixed effects for days between enrollment and ticket receipt. This leaves point estimates and results essentially unchanged.

A second alternative explanation is that the *Discuss* treatment forces discussion, produces a more argumentative discussion than the discussion that would have occurred naturally, and sees lower enrollment as a result. However, the first four columns of Table 8 provide evidence that contradicts this explanation. We find no evidence that spouses were more likely to disagree about the opportunity in the discussion group. Adults in the discussion group do not feel their spouse was more inconsiderate of their opinions; to the contrary, adults feel their spouses were significantly more considerate of their opinions in the discussion group.

Another possibility is that the presence of a surveyor in the discussion distorts conversation in a way that leads to less enrollment. One potential story is that the husband wants to prove to an outsider that he has control within the household and therefore he asserts opinions that are particularly against enrollment. This, however, is not at all consistent with qualitative reports from surveyors; if anything, couples tried to be more cordial in the presence of a surveyor with no surveyors saying the discussion was argumentative and 47% of surveyors saying the husband and the wife seemed equally interested in the job. This is also consistent with the explanation we propose above, which relies on husbands expressing support for the job in front of surveyors. Further, even if spouses were not fully genuine in the presence of a surveyor, they could always discuss the opportunity later to convey their true preferences. The final columns in Table 8 suggest that couples in the *Discuss* treatment had significantly more conversations after the first, which suggests they could undo any disingenuous opinions shared in the presence of a surveyor in discussions on their own.

A final idea is that couples in the *Discuss* treatment make a more informed decision and that the most informed decision is to not enroll. First note that couples in *Discuss* were given the same information by surveyors as couples in *Info* so a more informed decision could only come from having made a more carefully reasoned decision. If this were the case, the couples in the *Discuss* group that did enroll should have lower dropout rates due to making a better decision. However, data on retention presented below is not consistent with this.

5.3 Gender Effects on Enrollment

Finally, we present tests for a treatment effect of the gender of the ticket recipient on enrollment in Table 9. There is no effect of the gender of the ticket recipient on enrollment rates, consistent with the earlier result that there is no evidence for strategic information withholding.

5.4 Effects on Retention

Next we examine treatment effects on dropout and retention rates. Given the low enrollment (17%), our regressions conditional on enrollment have relatively few observations, giving us limited power to study the drivers of dropout. In light of low power and guided by results above, we pool treatment groups, regressing only on *Husb.Tick_i* and *Discuss_i*. Given there are 25 strata and only 81 couples with dropout status, we also exclude strata fixed effects from dropout regressions.

Results are presented in Table 10. We first consider the effects of *Discuss*. The third and fourth columns find discussion has a negative and significant effect on retention rates. This effect is large, lowering retention by 5 to 6 p.p. from a base of 8%. The first two columns find no effect of *Discuss* on dropout, which suggests that the effects on retention are driven by the effects on enrollment presented and explored above.

Next we consider the effects of *Husb.Tick*. The first two columns find that couples in which the husband received the ticket are significantly less likely to dropout. The effect is very large, lowering dropout rates 21 to 24 p.p. The next two columns summarize the enrollment and dropout effects by considering retention as an outcome. Though only marginally significant, the point estimates on *Husb.Tick* are large relative to the control mean; they suggest *Husb.Tick* raises retention by 3 to 4 p.p. from a low base of 8 p.p. This is despite the small negative effects of *Husb.Tick* on enrollment (Table 9). It follows that giving the husband the ticket improves retention by reducing dropout rates.

One explanation for the large effect of *Husb.Tick* on dropout is selection; even if *Husb.Tick* does not have an average effect on enrollment rates (Table 9), it may have an effect on the characteristics of those who enrolled, and these characteristics might predict dropout. This selection story is not supported in the data however, as column (2) of Table 10 shows the effect of *Husb.Tick* remains even conditional on covariates.

A more plausible explanation is as follows. Enrollment in the job opportunity produces drastic changes in household lifestyles. Women were previously responsible for household chores but now have less time to do them. If she is to continue working, the woman needs support from her household to ensure the chores are still completed. Many households are not willing or able to adjust to the change, which leads dropout rates to be high (indeed, the modal reason for dropout is that there is no-one else

to do the household chores). We suspect that when the woman’s continuing work required adjustments within the household, the husband was more likely to support these adjustments (either by helping with chores himself or persuading his family to help with them) if he had received the ticket and felt the enrollment decision was his idea.

We test this explanation directly by using data on the reasons for dropout. In column (1) of Table 11, we restrict the sample to only those that dropped out, and regress whether the reason given was household chores on *Husb. Tick*. There is a large negative and significant effect of 41 p.p. That is, those that dropped out when the husband got the ticket are far less likely to state household chores as the reason, consistent with husbands helping accommodate the wife’s work when given the ticket.

In the remaining columns we regress husbands’ and wives’ perceived job desirability at endline on *Husb. Tick*. Before considering results, note what this test cannot tell us. Because this variable was measured at endline, it was measured before the job began and cannot reflect experience with the job. Further, the measured relationship cannot tell us anything about information on job details as *Husb. Tick* is orthogonal to job details given to each spouse. Any relationship cannot be driven by adults believing that ticketing the husband is a signal of job quality because they likely knew other village couples in which the wife was ticketed. Nevertheless, we find that ticketed husbands view the job as significantly more desirable than non-ticketed husbands. We interpret this as evidence that husbands like the job more when they feel ownership over the enrollment decision, which leads them to be more supportive of the wife staying in the job. Note also that there is no relationship between the wife’s perceived job desirability and the wife receiving the ticket; increasing the husband’s like for the job does not come at the expense of the wife’s like for the job.

5.5 Summary Effects

Table 12 and Figure 7 present a summary of results on enrollment and retention. In particular, they present pairwise comparisons between the *NoInfo*, *Husb. Tick* treatment group and the remaining five groups. These summary results are consistent with the three main results provided above: (i) there is no evidence of information withholding; (ii) discussion depresses enrollment; and (iii) ticketing the husband has no effect on enrollment rates but significantly reduces dropout rates and therefore improves retention.

6 Conclusion and Discussion

In many cases, it is thought that women’s preferences more closely align with development goals than men’s. A natural goal of policy is therefore to encourage greater involvement of women in decision-

making. In this paper, we consider whether and how policymakers might do so by altering the way in which opportunities are presented to households. We consider one particular opportunity: a job for women. By orchestrating the recruitment process of a large carpet manufacturing firm in India, we test specific strategies policymakers might use to encourage greater involvement of women in a particular decision.

None of the three policies we tried were effective and two even backfired. We first find that informing women about an opportunity when husbands have an incentive to withhold information did not work. Husbands have such strong bargaining power within the household that they do not need to manipulate information to ensure decisions to reflect their preferences; there is no gain from providing additional information to women. Second, facilitating an initial discussion between spouses so as to kickstart a joint decision-making process dramatically reduced enrollment. This is contrary to idea that ensuring at least one discussion between the husband and wife is joint might ensure that the wife's preferences are not marginalized in subsequent decision-making. Third, placing decisions in wives' hands through symbolic gestures may backfire; we find dropout rates are much lower when the husband was given the enrollment ticket. This result is consistent with the idea that husbands might be more likely to support and accommodate a decision they feel ownership over.

Given these surprising results, it is worth placing our findings in the broader context of intra-household literature. A number of studies (e.g. [Duflo \(2003\)](#), [Duflo and Udry \(2004\)](#), [Field et al. \(2016\)](#), [Qian \(2008\)](#), and [Thomas \(1990\)](#)) find gains from giving women a greater role in household decisions by shifting women's intra-household bargaining power over extended periods of time. Our results suggest that these gains are not easily emulated in the very short-run through policies that increase women's role in a particular decision. Recent findings (e.g. [Ashraf et al. \(2014\)](#)) suggest policymakers could empower women in a particular decision by offering them the opportunity to bypass their husbands altogether. However, labor supply and many important household decisions cannot be made by bypassing one's spouse. The challenge of increasing women's role in such decisions remains.

We propose two avenues for future research. First, the exact bargaining process regarding decisions like labor supply could be studied with more precision, with attention paid to measuring beliefs, affect, and precise details on spouses' debates. Such fine-grained analysis might shed insight into exactly why treatments like our discussion treatment would backfire, and how similar treatments might be made effective. Second, researchers should innovate; we must design and test alternative interventions that might empower women in decision-making.

References

- ASHRAF, N. (2009): “Spousal control and intra-household decision making: An experimental study in the Philippines,” *The American Economic Review*, 99, 1245–1277.
- ASHRAF, N., E. FIELD, AND J. LEE (2014): “Household bargaining and excess fertility: an experimental study in Zambia,” *The American Economic Review*, 104, 2210–2237.
- BERTRAND, M., E. KAMENICA, AND J. PAN (2015): “Gender identity and relative income within households,” *The Quarterly Journal of Economics*, 130, 571–614.
- DUFLO, E. (2003): “Grandmothers and granddaughters: old-age pensions and intrahousehold allocation in South Africa,” *The World Bank Economic Review*, 17, 1–25.
- DUFLO, E. AND C. UDRY (2004): “Intrahousehold resource allocation in Cote d’Ivoire: Social norms, separate accounts and consumption choices,” Tech. rep., National Bureau of Economic Research.
- FIELD, E., R. PANDE, N. RIGOL, S. SCHANER, AND C. T. MOORE (2016): “On Her Account: Can Strengthening Women’s Financial Control Boost Female Labor Supply?” .
- HEATH, R. AND S. JAYACHANDRAN (2016): “The Causes and Consequences of Increased Female Education and Labor Force Participation in Developing Countries,” Tech. rep., National Bureau of Economic Research.
- QIAN, N. (2008): “Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance,” *The Quarterly Journal of Economics*, 123, 1251–1285.
- THOMAS, D. (1990): “Intra-household resource allocation: An inferential approach,” *Journal of human resources*, 635–664.

Tables and Figures

Figure 1: Female Labor Force Participation: India vs. The World




Figure 2: A Hand-Knotted Loom



Figure 3: Reasons for Non-Enrollment

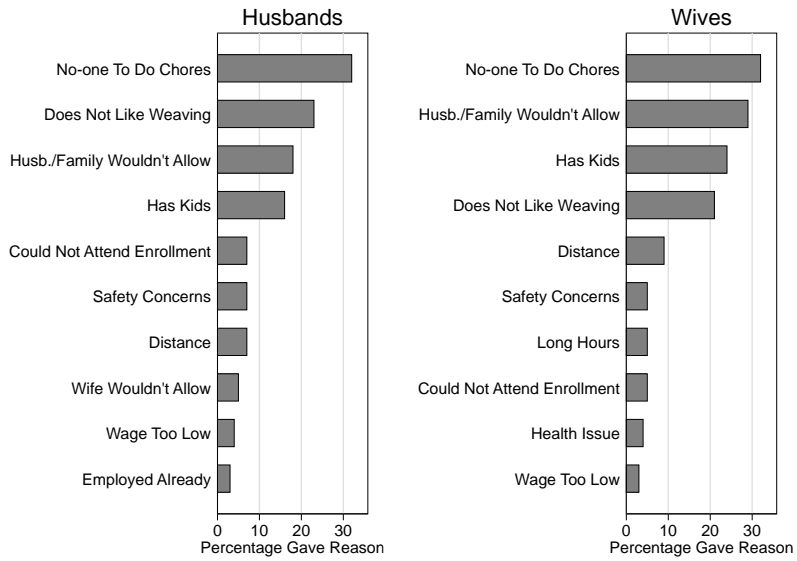


Figure 4: Reasons for Dropout

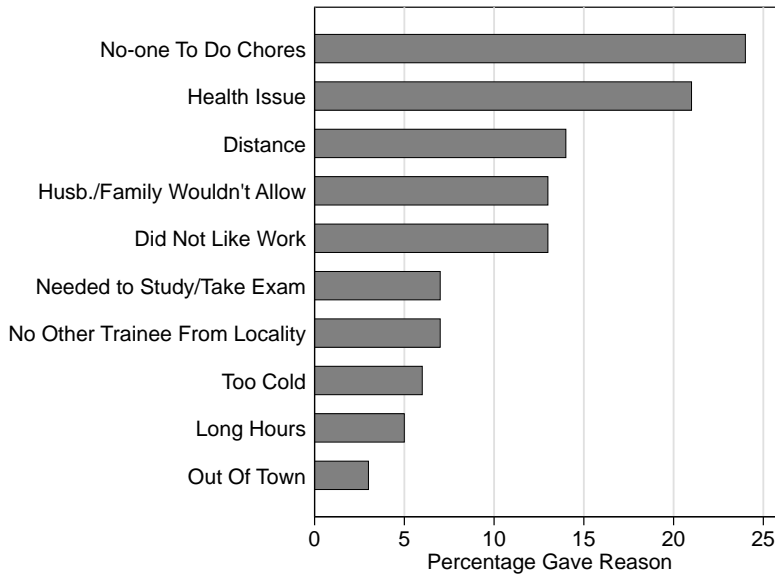


Figure 5: Experimental Design

		<i>Ticketed Spouse:</i>	
<i>Non-Ticketed:</i>		Husband	Wife
NoInfo			
Info			
Discuss			

Figure 6: Reasons for Attrition

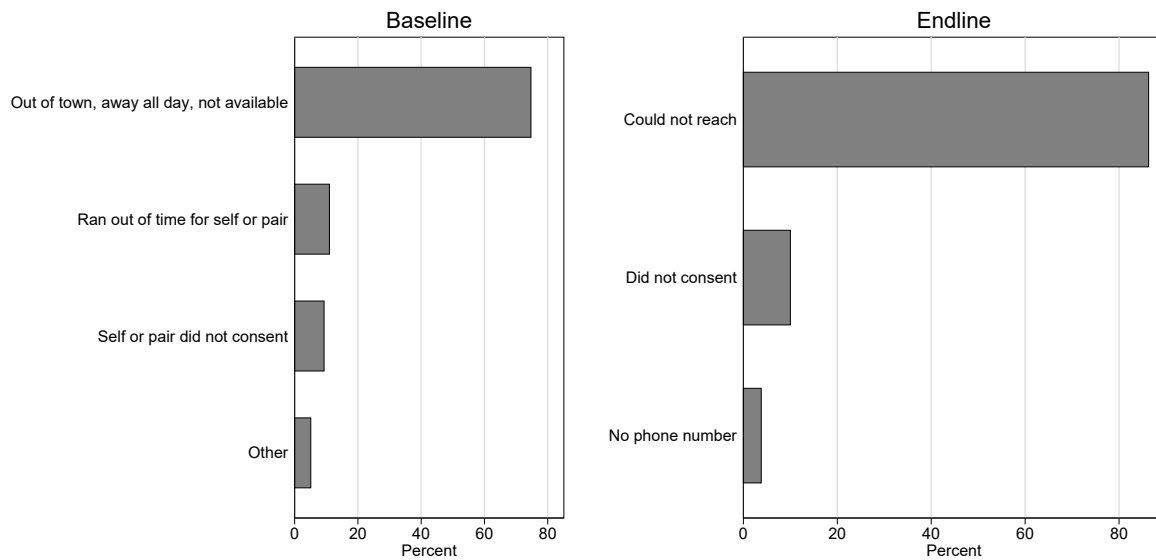


Figure 7: Summary of Results

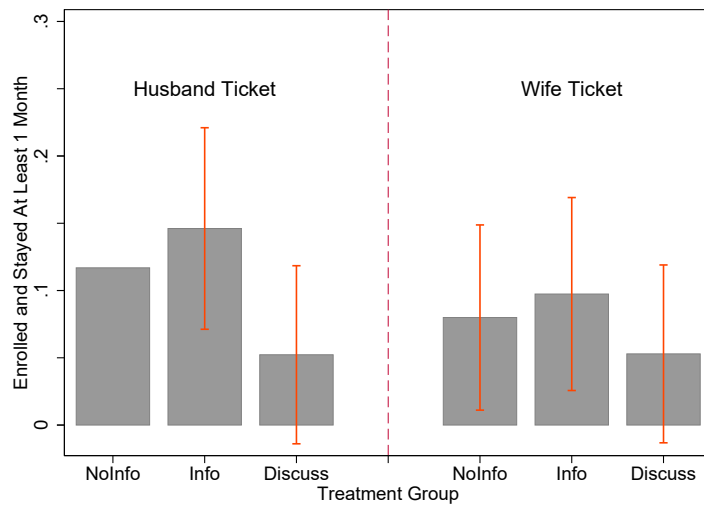
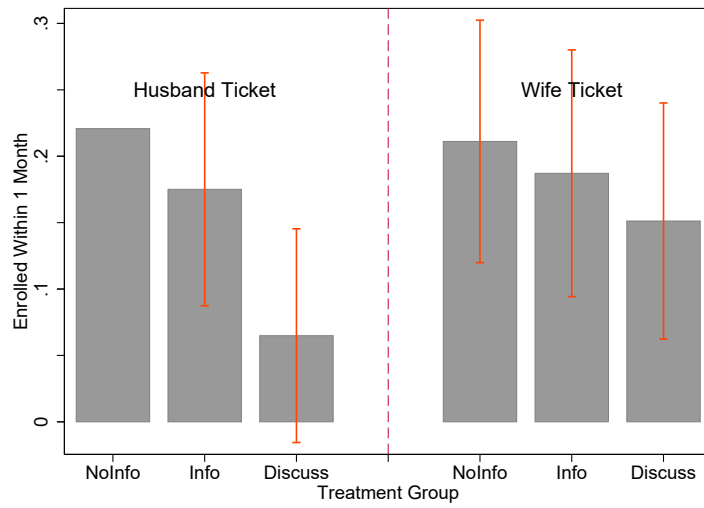



Table 1: Are Spouses' Employment Preferences Correlated?

	Wife Answer					
	<i>Appropriate for Man</i>			<i>Appropriate for Woman</i>		
	Construction (1)	Weaver (2)	Teacher (3)	Construction (4)	Weaver (5)	Teacher (6)
Husband Answer	0.18*** (0.05)	0.21*** (0.05)	0.01 (0.04)	0.21*** (0.05)	0.23*** (0.04)	0.07* (0.04)
N	482	481	481	479	479	477
R^2	0.03	0.05	0.00	0.04	0.05	0.01

Robust standard errors in parentheses. Appropriateness outcomes take values: 0 = Inappropriate, 1 = Somewhat appropriate, 2 = Completely appropriate. Constant not shown.

Table 2: Gender Preferences for Women's Employment

	<i>Appropriate for Man</i>			<i>Appropriate for Woman</i>			Job Interest (7)
	Construction (1)	Weaver (2)	Teacher (3)	Construction (4)	Weaver (5)	Teacher (6)	
Wife	-0.04 (0.05)	0.05 (0.05)	-0.02 (0.02)	0.29*** (0.05)	0.31*** (0.05)	0.11*** (0.04)	0.14** (0.06)
Husband Mean	1.36	1.40	1.91	0.57	0.81	1.66	1.16
N	976	974	975	973	973	971	500
R^2	0.59	0.61	0.51	0.61	0.63	0.55	0.71
Couple FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses. Appropriateness outcomes take values: 0 = Inappropriate, 1 = Somewhat appropriate, 2 = Completely appropriate. Interest in Job takes values: 0 = Not at all/not very interested, 1 = Somewhat interested, 2 = Very interested. The final column restricts to only those couples who were *both* asked about their interested during baseline.

Table 3: Baseline Balance Check

	Attrit:		Age of		No Education		Fertility		Years Married		Employed		Woman Weaver Pref.	
	Baseline (1)	Wife (2)	Husb. (3)	Wife (4)	Husb. (5)	10+ (7)	5-10 (8)	0-5 (9)	Wife (10)	Husb. (11)	Wife (12)	Husb. (13)		
No Info, Husb. Tick.	0.35 (0.04)	26.65 (0.59)	29.32 (0.64)	0.39 (0.06)	0.13 (0.04)	0.44 (0.06)	0.27 (0.05)	0.29 (0.05)	0.13 (0.04)	0.81 (0.05)	1.19 (0.10)	0.79 (0.10)		
Info, Husb. Tick.	0.30 (0.04)	25.39 (0.50)	27.81 (0.58)	0.39 (0.05)	0.22 (0.05)	0.34 (0.05)	0.35 (0.05)	0.30 (0.05)	0.06 (0.03)	0.80 (0.04)	1.00 (0.10)	0.69 (0.09)		
Discuss, Husb. Tick.	0.29 (0.04)	26.51 (0.47)	29.29 (0.48)	0.34 (0.05)	0.12 (0.03)	0.46 (0.05)	0.35 (0.05)	0.19 (0.04)	0.13 (0.04)	0.85 (0.04)	1.05 (0.10)	1.00 (0.10)		
No Info, Wife Tick.	0.35 (0.04)	25.09 (0.47)	28.33 (0.53)	0.33 (0.05)	0.07 (0.03)	0.33 (0.05)	0.35 (0.05)	0.32 (0.05)	0.06 (0.03)	0.80 (0.04)	1.20 (0.09)	0.67 (0.10)		
Info, Wife Tick.	0.29 (0.04)	27.32 (0.69)	30.19 (0.71)	0.47 (0.05)	0.13 (0.04)	0.49 (0.05)	0.26 (0.05)	0.26 (0.05)	0.15 (0.04)	0.84 (0.04)	1.18 (0.09)	0.78 (0.09)		
Discuss, Wife Tick.	0.35 (0.04)	26.87 (0.73)	29.27 (0.76)	0.42 (0.06)	0.14 (0.04)	0.31 (0.05)	0.35 (0.06)	0.35 (0.06)	0.25 (0.05)	0.83 (0.04)	1.09 (0.10)	0.91 (0.10)		
N	732	490	493	490	493	485	485	485	490	493	490	493		
p-value	.6	.03	.13	.58	.13	.06	.55	.26	.01	.71	.6	.07		

Robust standard errors in parentheses, unit of observation is couple-level. The p-value is from joint test for equality of means across the six treatment groups, conditional on strata FE.

Table 4: Endline Balance Check

	Attrit:		Age of		No Education		Fertility		Years Married		Employed		Woman Weaver Pref.												
	Endline	(1)	Wife	(2)	Husb.	(3)	Wife	(4)	Husb.	(5)	Fertility	(6)	10+	(7)	5-10	(8)	0-5	(9)	Wife	(10)	Husb.	(11)	Wife	(12)	Husb.
No Info, Husb. Tick.	0.15 (0.03)	26.78 (0.62)	29.40 (0.64)	0.41 (0.06)	0.12 (0.04)	2.25 (0.19)	0.48 (0.06)	0.26 (0.05)	0.26 (0.05)	0.13 (0.04)	0.80 (0.05)	1.17 (0.10)	0.80 (0.10)												
Info, Husb. Tick.	0.19 (0.03)	25.39 (0.56)	27.86 (0.64)	0.37 (0.06)	0.21 (0.05)	2.18 (0.20)	0.35 (0.06)	0.34 (0.06)	0.31 (0.05)	0.06 (0.03)	0.76 (0.05)	1.01 (0.10)	0.74 (0.10)												
Discuss, Husb. Tick.	0.14 (0.03)	26.44 (0.50)	29.09 (0.50)	0.34 (0.05)	0.10 (0.03)	2.58 (0.18)	0.43 (0.06)	0.38 (0.06)	0.19 (0.04)	0.12 (0.04)	0.87 (0.04)	1.05 (0.10)	0.99 (0.10)												
No Info, Wife Tick.	0.15 (0.03)	24.86 (0.47)	28.21 (0.56)	0.34 (0.06)	0.06 (0.03)	2.11 (0.16)	0.31 (0.05)	0.36 (0.06)	0.32 (0.05)	0.05 (0.02)	0.78 (0.05)	1.20 (0.10)	0.66 (0.10)												
Info, Wife Tick.	0.13 (0.03)	26.80 (0.60)	29.72 (0.65)	0.44 (0.06)	0.12 (0.04)	2.39 (0.19)	0.49 (0.06)	0.24 (0.05)	0.27 (0.05)	0.15 (0.04)	0.84 (0.04)	1.12 (0.10)	0.74 (0.09)												
Discuss, Wife Tick.	0.23 (0.04)	27.03 (0.76)	29.39 (0.84)	0.41 (0.06)	0.15 (0.04)	2.24 (0.18)	0.32 (0.06)	0.37 (0.06)	0.32 (0.06)	0.25 (0.05)	0.82 (0.05)	1.19 (0.10)	0.97 (0.11)												
N	990	826	828	826	828	826	820	820	820	826	828	826	828												
p-value	.29	.02	.24	.71	.05	.43	.06	.29	.39	.01	.33	.62	.07												

Standard errors clustered at couple-level in parentheses, unit of observation is adult-level throughout. The p-value is from joint test for equality of means across the six treatment groups, conditional on strata FE.

Table 5: Information Withholding

	Enrolled (=1)		Enrolled Within 1 Month (=1)		Non-Ticketed Knows About Ticket (=1)		Ticketed Comp. Q Correct (=1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Info (=1)	-0.03 (0.05)	-0.04 (0.05)	-0.02 (0.06)	-0.02 (0.06)	0.13* (0.07)	0.12* (0.07)	-0.85*** (0.04)	-0.84*** (0.04)
Husb. Tick. (=1)	0.02 (0.06)	0.00 (0.05)	0.02 (0.06)	0.01 (0.06)	0.03 (0.07)	0.03 (0.08)	-0.04 (0.04)	-0.04 (0.04)
Info X Husb. Tick. (=1)	0.01 (0.07)	0.03 (0.07)	-0.03 (0.08)	-0.02 (0.08)	-0.04 (0.10)	-0.03 (0.11)	0.07 (0.07)	0.05 (0.06)
Discussion Treatments	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Strata	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	Yes	No	Yes	No	Yes
Omitted Group Avg.	0.18	0.18	0.21	0.21	0.72	0.72	0.93	0.93
Obs. Unit	Couple	Couple	Couple	Couple	Couple	Couple	Couple	Couple
N	495	484	495	484	414	410	330	325

Robust standard errors in parentheses. Enrolled is an indicator for enrolling on enrollment day. Enrolled Within 1 Month is an indicator for enrolling within 1 month of enrollment day. Non-Ticketed Knows About Ticket is an indicator for the non-ticketed spouse knowing his/her spouse received an enrollment ticket. Ticketed Comprehension Question Correct is an indicator for the ticketed spouse answering the baseline comprehension about what his/her spouse would learn from a surveyor correctly. The final two columns exclude the discussion group. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Discussion Effect on Enrollment

	Enrolled (=1)		Enrolled Within 1 Month (=1)		Enrolled (=1)		Enrolled Within 1 Month (=1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Discuss (=1)	-0.06*	-0.08***	-0.06**	-0.09***	-0.02	-0.04	-0.03	-0.05
	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
Husb. Tick. (=1)					0.02	0.02	0.01	-0.00
					(0.04)	(0.04)	(0.04)	(0.04)
Discuss X Husb. Tick. (=1)					-0.07	-0.07	-0.07	-0.08
					(0.06)	(0.06)	(0.06)	(0.06)
Strata	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	Yes	No	Yes	No	Yes
Omitted Group Avg.	0.17	0.17	0.19	0.19	0.16	0.16	0.19	0.19
Obs. Unit	Couple	Couple	Couple	Couple	Couple	Couple	Couple	Couple
N	495	484	495	484	495	484	495	484

Robust standard errors in parentheses. Enrolled is an indicator for enrolling on enrollment day. Enrolled within 1 month is an indicator for enrolling within 1 month of enrollment day. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Discussion: Alternative Explanations I

	Enrolled (=1)		Enrolled Within 1 Month (=1)	
	(1)	(2)	(3)	(4)
Discuss (=1)	-0.05 (0.03)	-0.07*** (0.03)	-0.06* (0.03)	-0.09*** (0.03)
Day from Enrollment FE	Yes	Yes	Yes	Yes
Strata	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	Yes
Omitted Group Avg.	0.17	0.17	0.19	0.19
Obs. Unit	Couple	Couple	Couple	Couple
Num. Obs.	495.00	484.00	495.00	484.00

Standard errors clustered by couple in parentheses. Enrolled is an indicator for enrolling on enrollment day. Enrolled within 1 month is an indicator for enrolling within 1 month of enrollment day. Day from Enrollment FE are fixed effects for the number of days between enrollment and when the couple's ticket was given. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Discussion: Alternative Explanations II

	Spouses Disagreed, Couple Avg.		Spouse Inconsiderate, Couple Avg.		N. Discussions After First, Couple Avg.	
	(1)	(2)	(3)	(4)	(5)	(6)
Discuss (=1)	0.03 (0.08)	0.07 (0.08)	-0.16** (0.07)	-0.14** (0.07)	0.67** (0.34)	0.71** (0.36)
Strata	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	Yes	No	Yes
Omitted Group Avg.	1.20	1.20	0.62	0.62	2.42	2.42
Obs. Unit	Couple	Couple	Couple	Couple	Couple	Couple
Num. Obs.	271	268	262	259	446	438

Robust standard errors in parentheses. Spouses Disagreed is the level of disagreement between spouses about the job opportunity, measured on a 0-2 scale. Spouse Inconsiderate is the extent to which one's spouse was inconsiderate of one's own opinion about the job opportunity, measured on a 0-2 scale. Number Discussions After First is the number of discussions spouses had after their first discussion. In each regression, the outcome used is the average of the husband and wife responses. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Ticket Recipient and Enrollment

	Enrolled (=1)		Enrolled Within 1 Month (=1)	
	(1)	(2)	(3)	(4)
Husb. Tick. (=1)	-0.00 (0.03)	-0.01 (0.03)	-0.02 (0.03)	-0.03 (0.03)
Strata	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	Yes
Omitted Group Avg.	0.15	0.15	0.18	0.18
Obs. Unit	Couple	Couple	Couple	Couple
N	495	484	495	484

Robust standard errors in parentheses. Enrolled is an indicator for enrolling on enrollment day. Enrolled within 1 month is an indicator for enrolling within 1 month of enrollment day. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Treatment Effects on Dropout

	Dropped Out Within First Month (=1)		Enrolled and Stayed At Least 1 Month (=1)	
	(1)	(2)	(3)	(4)
Discuss (=1)	0.08 (0.14)	-0.05 (0.16)	-0.05** (0.02)	-0.06** (0.02)
Husb. Tick. (=1)	-0.24** (0.11)	-0.21* (0.12)	0.04 (0.02)	0.03 (0.02)
Strata	No	No	Yes	Yes
Covariates	No	Yes	No	Yes
Omitted Group Avg.	0.60	0.60	0.08	0.08
Obs. Unit	Couple	Couple	Couple	Couple
N	81	80	491	480

Robust standard errors in parentheses. Dropped Out Within First Month is an indicator for dropping out of the program within the first month of the center's opening. It is only observed for those that enrolled on or within one month of enrollment day. Enrolled and Stayed At Least 1 Month for the wife enrolling and staying at least one month into the program. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Husband Ticket and Dropout: Channels

	Dropped Out Due to Household Chores (=1)		Job Desirability		
	(1)	(2)	(3)	(4)	(5)
Husb. Tick. (=1)	-0.41*** (0.12)	0.20** (0.10)	0.16 (0.10)	-0.01 (0.10)	0.03 (0.10)
Strata	No	Yes	Yes	Yes	Yes
Covariates	No	No	Yes	No	Yes
Omitted Group Avg.	0.48	1.03	1.03	1.06	1.06
Obs. Unit	Drop Out Women	Husband	Husband	Wife	Wife
N	39	244	241	234	233

Robust standard errors in parentheses. Dropped out due to household chores is a dummy for dropping out due to demands of household chores. Job desirability is the perceived desirability of the job, measured on a 0-2 scale. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12: Summary of Results

	Enrolled (=1)		Enrolled Within 1 Month (=1)		Enrolled and Stayed At Least 1 Month (=1)	
	(1)	(2)	(3)	(4)	(5)	(6)
Info, Husb. Tick. (=1)	-0.02 (0.05)	-0.02 (0.05)	-0.05 (0.06)	-0.05 (0.05)	0.03 (0.05)	0.03 (0.05)
Discuss, Husb. Tick. (=1)	-0.10** (0.05)	-0.12*** (0.05)	-0.12** (0.05)	-0.16*** (0.05)	-0.05 (0.04)	-0.06 (0.04)
NoInfo, Wife Tick. (=1)	-0.02 (0.06)	-0.00 (0.05)	-0.02 (0.06)	-0.01 (0.06)	-0.04 (0.04)	-0.04 (0.04)
Info, Wife Tick. (=1)	-0.05 (0.05)	-0.05 (0.05)	-0.04 (0.06)	-0.03 (0.06)	-0.03 (0.04)	-0.02 (0.04)
Discuss, Wife Tick. (=1)	-0.05 (0.05)	-0.07 (0.05)	-0.06 (0.06)	-0.07 (0.05)	-0.06 (0.04)	-0.06 (0.04)
Strata	Yes	Yes	Yes	Yes	Yes	Yes
Covariates	No	Yes	No	Yes	No	Yes
Omitted Group Avg.	0.19	0.19	0.22	0.22	0.12	0.12
Obs. Unit	Couple	Couple	Couple	Couple	Couple	Couple
N	495	484	495	484	491	480

Robust standard errors in parentheses. Enrolled is an indicator for enrolling on enrollment day. Enrolled within 1 month is an indicator for enrolling within 1 month of enrollment day. Enrolled and Stayed At Least 1 Month for the wife enrolling and staying at least one month into the program. Covariates include wife's age, husband's age, wife uneducated, husband uneducated, number of children, time since marriage, wife employed, husband employed, wife's approval of women weavers, and husband's approval of women weavers.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$