

MEASURING AND CHANGING CONTROL: WOMEN'S EMPOWERMENT AND TARGETED TRANSFERS*

Ingvild Almås, Alex Armand, Orazio Attanasio and Pedro Carneiro

This article uses a novel identification strategy to measure power in the household. Our strategy is to elicit women's willingness to pay to receive a cash transfer instead of their spouse receiving it. We selected participants from a sample of women who had already participated in a policy intervention in Macedonia offering poor households cash transfers conditional on having their children attending secondary school. The programme randomised transfers at the municipality level to either household heads (generally a male) or mothers. We show that women who were offered the transfer on average have stronger measured empowerment. Here, IV estimation confirms this result.

Most conditional cash transfer programmes around the world select a woman in the household to be the recipient of the transfer (Fiszbein *et al.*, 2009). The argument frequently used in support of targeting transfers to women is not only that such transfers promote gender equality and empower women but also that, through the empowerment of women, they benefit children as well.

Policy interventions that change the relative income of women *versus* men within households, such as the 1979 UK Reform of Child benefits analysed by Lundberg *et al.* (1997) and Ward-Batts (2008), the Mexican PROGRESA (Attanasio and Lechene, 2002, 2014), or the reform of the South African social pension programme studied by Duflo (2003), have been shown to have an effect on different family decisions (for a literature review, see Duflo, 2012). Consistent with this evidence, models of intra-household resource allocation that depart from the unitary framework (according to which the household can be considered as a decision unit with a well defined and fixed objective function) by and large imply that changing the control of resources results in different outcomes (Browning and Chiappori, 1998; for empirical applications, Thomas, 1990; Hoddinott and Haddad, 1995; Lundberg *et al.*, 1997; Doss, 2006; Ward-Batts, 2008).

* Corresponding author: Ingvild Almås, Institute for International Economic Studies, Stockholm University, SE-106 91 Stockholm, Sweden. E-mail: ingvild.almås@iies.su.se.

We thank Nava Ashraf, Richard Blundell, Martin Browning, Pierre-André Chiappori, Pascaline Dupas, Tore Ellingsen, Ernst Fehr, Eliana La Ferrara, Johannes Haushofer, Seema Jayachandran, Jeremy Lise, Karen Macours, Muriel Niederle, Fabien Postel-Vinay, Imran Rasul, Sylvie Lambert, Iliana Reggio, Kjell Salvanes, Michèle Tertilt, Bertil Tungodden, Alessandra Voena and seminar participants at IIES, Stockholm University, Tilburg University, Paris School of Economics, the Experimental and Behavioral Economics Seminar at Zurich University, the EDePo IFS conference in London, the Families in Macroeconomics conference, Edesheim, 8th COSME Gender Economics Workshop, University of Alicante, EEA Annual Meeting, the Normac conference, NOVAfrica and University of Essex for helpful comments. We gratefully acknowledge the financial support from 3ie International Initiative for Impact Evaluation (grant reference OW4-1022), The Choice Lab NHH, the Economic and Social Research Council (ESRC Professorial Fellowship ES/K010700/1), the European Research Council (Advanced Grants 249612), and the Research Council of Norway (project no 250415). When carrying out part of this research, Almås was associated with the Centre of Equality, Social Organization, and Performance (ESOP) at the Department of Economics at the University of Oslo. ESOP is supported by the Research Council of Norway.

However, there is no clear consensus on the precise mechanism through which households make decisions and allocate consumption when receiving a cash transfer, and there is limited evidence on the exact mechanism linking money transfers targeted to women and empowerment within the household. Moreover, measures of the relative decision power or bargaining strength within a household are rarely available, posing some difficult empirical challenges.

This article suggests and uses a novel method to measure the relative bargaining strength, or decision power, of women within the household. Rather than relying on traditional survey questions about who makes certain decisions regarding resource allocation within the household, we directly measure women's willingness to pay to gain control over income through an experiment. This experiment was implemented in urban areas of Macedonia. The women selected to participate were a sequence of choices between an amount A_k for themselves or an amount B_k for their husband (where A_k is usually smaller than B_k). The sequence of choices is designed to identify the value that makes the participants indifferent between receiving A_k and their husband receiving B_k . The experiment therefore elicits the participant's willingness to pay to become the recipient of a cash transfer offered to the household. We argue below that the measure from the experiment identifies bargaining power and does so in a more effective way than traditional survey-based measures.

The measure of empowerment from the experiment can be matched with data from a study of a nationwide cash transfer programme. All participants in the economic experiment are women in households eligible for the Macedonian 'Conditional Cash Transfer (CCT) for Secondary School Education', which provides cash transfers to poor households conditional on having their children enrolled in secondary school. In the three years prior to our laboratory experiment, the operation of the programme was not uniform across municipalities. In one group of randomly selected municipalities, the CCT was paid to the mother in the household, while in the remaining municipalities, the transfer was paid to the household head (generally a man). The random assignment of the programme modality across municipalities provides exogenous variation in the amount of resources potentially controlled by each household member. All women participating in the experiment also participated in at least one round of a household survey, containing detailed information on demographics, consumption, income and living conditions.

Taken together, the laboratory, field and survey data, constitute a unique data set which allows the identification of the empowerment effect, i.e. the effect of targeted transfers on women's bargaining position in the household. Our findings indicate that targeted transfers have a significant effect on female empowerment.

The rest of the article is organised as follows. Section 1 describes the economic experiment. Section 2 sketches a theoretical framework that allows us to justify the specific measure we have designed. Section 3 describes the data in more detail and gives descriptive statistics and Section 4 discusses the empirical analysis. Section 5 concludes.

1. Measuring Control: A Laboratory Experiment in the Field

In the past few decades, many surveys have included batteries of questions aimed at measuring the extent to which women are empowered within the family. A typical set

of questions, used in many different contexts, asks respondents to identify who is in charge of certain decisions, determining, for example, expenditures on different household consumption items, schooling, or various investments. Possible answers to these questions are that the wife is in charge, the husband is in charge, or spouses decide jointly. In many data sets, answers to these questions are bunched on 'both' categories, and a very limited variation is obtained.

In the context of conditional cash transfers, for instance, the PROGRESA evaluation survey included several of these questions. This CCT did not seem to have shifted the answers to these questions (see for instance Adato *et al.*, 2000). Therefore, if one were to interpret those results literally, one would conclude that the transfer programme, despite offering significant transfers to women, did not empower them. Yet, many studies claim that PROGRESA and other similar programmes did shift the position of women in the family in a substantive fashion, because of their impacts on the portfolio of household expenditures (Attanasio and Lechene, 2014).

One possible explanation of the failure to observe direct impacts of programmes such as PROGRESA on measured empowerment could be that empowerment is poorly measured in surveys. Adato *et al.* (2000), for instance, write: 'Women's status is difficult to quantify in the context of large household surveys like the ENCASEH and ENCELS. These surveys have several questions which attempt to tease out various aspects of women's status and bargaining power, such as attitudes towards women's roles, questions on who within the household takes a major responsibility for certain household decisions, questions on the disposition of women's income, and questions on women's mobility and freedom of movement. Nevertheless, household surveys are blunt instruments with which to examine intra-household relations, because the context of such decisions is often unstated, and without adequate understanding of the sociocultural context, survey results can easily be misinterpreted'.

One contribution of this article is to propose an alternative quantitative measure of women's empowerment within the household.¹ This measure is based on observing women's behaviour and choices in a laboratory setting. In particular, in the laboratory experiment that we describe in detail below, we offer the participants a cash transfer. The participant can choose that her husband becomes the recipient of the transfer, and in this case she does not have to pay anything. But she can also choose that she becomes the recipient, and in this case she may have to pay some of the offered amount. Through varying the amount, she has to pay in order to receive the transfer, we are able to elicit her willingness to pay for becoming the recipient of a cash transfer to the household.

The reason why this is a measure of female empowerment is simple. In a unitary setting, women should not be willing to pay anything in order to receive the transfer themselves, and should instead try to maximise the transfer amount. On the other hand, in a non-unitary model, the weaker the position of the woman in the household (the lower her control of resources), the more she should be willing to pay to obtain control of that transfer. In the next Section, we make this intuition precise within the

¹ An alternative could be to change the way in which the 'decision' questions are asked. Ashraf (2009), for instance, reports a much lower bunching at 'both' the answer for a survey in the Philippines. Interestingly, in that context, additional questions about who is the 'tie-breaker' are also asked.

framework of the collective model of Chiappori (1992), although a similar reasoning applies to other models of intra-household allocation. The remainder of this Section describes in detail the experiment we conducted to measure empowerment.

The experiment consists of a sequence of choices between two alternatives, which the participant is asked to make in sequential rounds. The participant is not informed about the algorithm determining the alternatives in each round. In the first round, the participant is asked to choose between an amount, A_1 , paid to her, or another amount B paid to her husband. B is kept constant across rounds, and we therefore refer to it as the 'stake' of the experiment. If the participant in the first round chooses A_1 (choosing herself to be the recipient), the amount A_2 in the following round is reduced by 75%. If, in the second round, the participant again chooses herself to be the recipient, the already reduced amount A_2 is again reduced by 75%. If the participant continues to choose herself to be the recipient, we keep reducing the amount by 75%, until the resulting amount is smaller than 20 MKD.²

If, however, in the first round, the participant chooses B (her husband to be the recipient) over A_1 , the amount A_2 is increased by 50 MKD in the next round. If the participant again chooses B , the amount A is once again increased by 50 MKD. If the participant chooses her husband to be the recipient a third time, the experiment stops.³

Consider now the case where the participant changes her choice when the value of A_k changes. If, in one round, the participant switches from A_k (choosing herself to be the recipient) to B (choosing her husband to be the recipient) or *vice versa*, the amount A_{k+1} offered in the next round is set to be the average between two amounts:

- (i) the amount offered in the current round; and
- (ii) the amount offered in the previous round where the participant made a different decision from the current one.

Therefore, the amount offered to participants increases or decreases depending on whether they switch from either receiving the transfer to letting their husband receive it, or *vice-versa*. This procedure continues as long as the difference between two consecutive amounts A_k is larger than 20 MKD. The decision to introduce a stopping rule of 20 MKD, rather than a smaller amount, is to avoid asking consecutive questions on amounts that are very similar in terms of monetary value, which would not be distinguishable to the participant in a meaningful way.

We piloted this algorithm and its details extensively. The experiment was run on a computer which, as the experiment progressed, automatically performed all the computations described above. Participants faced a new round of alternatives immediately after the decision in a given round was made and confirmed. Screenshots taken from the software are presented in online Appendix A.1. The algorithm was

² The amounts are expressed in Macedonian Denars (MKD). The exchange rate prevalent at the time of the experiment was 0.0215 US\$/MKD. This is computed as the average exchange rate with United States Dollars for the year 2014 (*Source*. National Bank of the Republic of Macedonia).

³ The reason why we stop the amount A from increasing further is to avoid situations in which the amount would become too high. At the same time, we do allow the amount A to become larger than the amount B in order to collect information about women that present a negative willingness to pay, e.g. collect information about those that would be willing to pay in order to ensure that their husband becomes the recipient of the transfer.

Always to Spouse		Switcher 1		Switcher 2		Always to Self	
Participant	Spouse	Participant	Spouse	Participant	Spouse	Participant	Spouse
550	600	550	600	550	600	550	600
600	600	138	600	138	600	138	600
650	600	344	600	35	600	35	600
	STOP	447	600	86	600		STOP
		396	600	61	600		
		421	600		STOP		
			STOP				

Fig. 1. Algorithm for the Determination of Offered Amounts: Some Examples

Notes. The Figure presents four possible combinations of a participant’s answers. From the left, the first shows a case in which the participant always chooses the option where her husband is the recipient. The second and the third show cases in which the participant switches after the first question. The fourth shows instead a case in which the participant decides that she will be the recipient independent of the cost. Colour Figure can be viewed at wileyonlinelibrary.com.

designed in order to identify the participant’s indifference point between the two alternatives of her being the recipient and her husband being the recipient. Some examples of the mechanics of the algorithm are presented in Figure 1.⁴

1.1. Protocol

As we discuss in more detail below, participants in the experiment were drawn from the sample of respondents to the household survey that was collected for the evaluation of the Conditional Cash Transfer programme in Macedonia. Women were invited, either through the phone or through visits to their homes, to participate in the experiment using contact information collected in that survey. The invitation stated that they were being asked to participate in an interview about the needs of women in their social stratum, and that they would be financially compensated for their time.

The experiment was carried out in a room where only the participant and a female assistant were present. Participation in the experiment did not involve any monetary costs for the participants. A driver picked up participants at their dwelling or at an agreed meeting point, and took them to the office where the experiment was carried out.⁵ Participants were given 300 MKD as a show-up fee. Participants were also told that one round of choices and its associated decision would be randomly selected to determine the actual payment at the end of the game.

⁴ It is important to note that we have no reason to believe that the participants understood the algorithm that we used, and that we saw no indication of strategic behaviour in the response, neither from the in-depth interviews with the participants in the focus groups after the pilot, nor from the data collected in the pilot and in the full experiment.

⁵ The decision to cover transportation costs was made to avoid that heterogeneity in this cost would influence the outcome of the experiment. In addition, we selected only participants living 10 kilometres or less away from the closest urban settlement, such that distance would not be an important factor in participants’ answers.

The starting value for B (the stake) was randomised among nine amounts, ranging from 400 MKD (8.60 US\$) to 800 MKD (17.20 US\$).⁶ The starting value for A was then defined to be $(B - 50 \text{ MKD})$, allowing the experiment to start from an initial choice for which it is costly for the participant to become the receiver. The different starting points and their value in US\$ are presented in Table 1. We report further details on the experiment and its protocol in online Appendix A.

In terms of size of the incentive, stakes ranged from 62% to 123% of average total daily household expenditure, which is roughly 650 MKD (16.98 US\$) in this sample. Relative to wages, the minimum offered stake corresponds to roughly 52% of the daily net wages for women who have completed secondary school, and 68% for women who has completed primary school only. The maximum offered stake corresponds to 103% and 136% of these wages, respectively.⁷

Following the incentivised version of the experiment, participants were asked to answer a non-incentivised version of the experiment, where the amounts were 10 times larger compared to the incentivised version. In this case, we asked the participants to choose across alternatives, thinking of the hypothetical situation they faced as if it corresponded to a real life scenario. Similarly to the incentivised version, initial values for B were randomised between nine amounts, ranging from 4,000 MKD (86.00 US\$) to 8,000 MKD (172.00 US\$). The starting value for A was then defined by $B - 500 \text{ MKD}$.

The actual payment from the experiment was given in envelopes, either to the participants or one for the participant and one for the spouse. The drivers drove the participants back to the dwellings and handed over the envelope to the spouse in those cases the spouse was due a payment. A main difference between our setting and that of e.g. Ashraf (2009), Ashraf *et al.* (2014) is that we in principle provided full information about earnings from the experiment to participants and spouses. We did, however, on purpose not give the spouses information about the choices the participants made and we included enough randomness in our design so that it was not possible to infer behaviour in the laboratory from the earnings. As we did not allow for hiding of resources but rather gave full information about earnings, the experiment is not designed to study how the possibility of hiding affects the results.⁸

⁶ Online Appendix C.2 presents a test for whether the willingness to pay is a function of the stakes. Results show that for both incentivised and non-incentivised cases, the willingness to pay is uncorrelated with the stakes, both when testing the coefficients individually and jointly.

⁷ Daily net wages for different educational levels are estimated using data provided by the Macedonian State Statistical Office. Gross wages by educational level were available for October 2010 (*Source*. 2010 Structure of Earnings of Employees) and net/gross wages were available for October 2010 and July 2014 (*Sources*. Average monthly net wage paid per employee, Average monthly gross wage paid per employee). We made use of the net/gross wage ratio in October 2010 to build net wages by educational level. We then computed wages by educational level in 2014, by using the nominal growth rate of wages from October 2010 to July 2014. At the time of the interview, the net daily wage for a woman who has completed primary school is estimated to be 590 MKD, while for a woman who has completed secondary school it is estimated to be 770 MKD.

⁸ A small study with about 60 subjects has been conducted (Almås *et al.*, 2016; Gudevold and Kjørholt, 2016) testing how choices in our setting may change with the possibility of hiding, and neither the choices in the experiment, nor the focus groups revealed strong effects of hiding but the sample size is small and the effect of information in our setting is thus still an avenue for future research.

Table 1
Starting Points

Incentivised		Non-incentivised	
Participant	Husband	Participant	Husband
350 (7.53)	400 (8.60)	3,500 (75.25)	4,000 (86.00)
400 (8.60)	450 (9.68)	4,000 (86.00)	4,500 (96.75)
450 (9.68)	500 (10.75)	4,500 (96.75)	5,000 (107.50)
500 (10.75)	550 (11.83)	5,000 (107.50)	5,500 (118.25)
550 (11.83)	600 (12.90)	5,500 (118.25)	6,000 (129.00)
600 (12.90)	650 (13.98)	6,000 (129.00)	6,500 (139.75)
650 (13.98)	700 (15.05)	6,500 (139.75)	7,000 (150.50)
700 (15.05)	750 (16.13)	7,000 (150.50)	7,500 (161.25)
750 (16.13)	800 (17.20)	7,500 (161.25)	8,000 (172.00)

Notes. Main units are expressed in Macedonian Denars (MKD). In parenthesis, we report the correspondence with US\$ using the prevalent exchange rate at the time of the experiment, 0.0215 US\$/MKD. This is computed as the average exchange rate with US\$ for the year 2014. *Source.* National Bank of the Republic of Macedonia.

2. Interpreting the Measures: A Theoretical Framework

The measurement tool we propose identifies how much a woman is willing to pay in order to gain control of an amount otherwise offered to her husband. While it is intuitive that such a measure should be related to the bargaining power of a woman within the household, it is useful to consider models of intra-household allocation of resources to interpret it more precisely. In this Section, we first provide a general discussion of the intuition behind the interpretation of our measure within a wide range of intra-household allocation models. We then turn to a more specific discussion within the collective model framework of intra-household allocation of resources. This is to make the intuition precise in a given context. In particular, within the context of this theoretical framework, we ask two questions: how does our measure change when we vary exogenously the bargaining power of the woman within the household, and how does the measure change with a targeted conditional cash transfer of the type we have in our sample.

2.1. Models of Intra-household Allocations

The process of resource allocation within the household can be, and has been, modelled in many different ways. Some models assume cooperative behaviour between household members, while others allow for non-cooperative behaviour and the waste of resources in the process of intra-household allocation.⁹ A growing literature applies these models to study household decision-making in a variety of specific contexts, such as the allocation of resources to children, labour supply, and fertility decisions (Hoddinott and Haddad, 1995; Chiappori *et al.*, 2002; Blundell *et al.*, 2005, 2007; Ashraf *et al.*, 2014; Voena, 2015).

⁹ See e.g. Browning *et al.* (2014) and Doepke and Tertilt (2014) for discussion of different model approaches.

Recently, several empirical papers have proposed rigorous testing of decision-making in different contexts, work that is informative about the suitability of alternative models of household decision-making (Ashraf, 2009; Ashraf *et al.*, 2014; Attanasio and Lechene, 2014). Depending on the context, one may prefer to focus on one particular group of models, either cooperative or non-cooperative ones. However, for the specific measure of power we are analysing, the intuitions from the non-cooperative and cooperative models are similar: with decreasing marginal utility of income, and with either caring preferences (i.e. one cares about the spouse's outcomes) or with the presence of public goods, one would expect that the more resources a spouse has control over, the less she is willing to pay in order to gain control over additional resources.¹⁰

In a unitary setting, however, it is assumed that the household acts as if there is only one decision-maker in the household and, therefore, a woman cannot gain more control through a targeted transfer. In the following we are more exact; we start by discussing explicitly the so-called unitary model, which is quite restrictive in assuming that there is a single decision-maker in the household (Becker, 1991). We move on to discuss the much used collective model of household decision-making (Alderman *et al.*, 1995). We show formally that whereas the unitary model would predict zero willingness to pay, the collective model predicts that the willingness to pay to control additional resources decreases with the control over the existing resources. The collective model can be derived from a cooperative or a non-cooperative setting, but the implemented allocations are always Pareto optimal. In our setting, we are agnostic about whether the allocation comes about through Nash or other types of bargaining or through some other decision-making process.

We assume that household decisions are carried out by two decision-makers, a woman (*A*) and her spouse (*B*), who decide how to allocate total household expenditure to different goods, either publicly or privately consumed. Let Q be the quantity consumed of a public good, say spending on children, and let q_A and q_B be quantities consumed of private goods for the woman and her spouse. The household budget constraint is given by the following:

$$PQ + p_A q_A + p_B q_B = x = x_A + x_B, \quad (1)$$

where P , p_A , and p_B are the prices for public and private consumption goods. Total household expenditure, x , equals income in this setting, since we do not consider savings. Household income is given by the sum of person *A* and person *B*'s incomes, x_A and x_B , respectively. Individual preferences are defined over private goods and public goods, and we assume that there is no direct caring for the spouse: $u^A(Q, q_A)$ is the utility function for person *A*, and $u^B(Q, q_B)$ is the utility function for person *B*.

A unitary model assumes that choices are made according to a 'unitary' household utility function $\tilde{U}(Q, q_A, q_B)$. A natural assumption is to impose that the household utility function respects individual preferences. One way to rationalise a unitary model

¹⁰ Note that in a non-cooperative setting with neither caring nor public goods, she should on the margin be willing to pay the whole amount as she gains no utility from her spouse receiving the money.

based on individual preferences is to assume that households maximise a weighted sum of individual preferences, in which the weights are fixed preference parameters:¹¹

$$\tilde{U}(Q, q_A, q_B) = \mu u^A(Q, q_A) + (1 - \mu) u^B(Q, q_B). \quad (2)$$

A central characteristic of this representation of a unitary model is that the weight does not vary with either prices or income. The demand for each individual commodity depends on prices and total household income only, and is independent of the distribution of income within the household. Such demand functions satisfy general conditions (the Slutsky conditions).

The collective model, on the other hand, assumes that resources are allocated efficiently but it allows the weights μ to depend on prices, income and distribution factors. In the literature, distribution factors are defined as variables that have an impact on the decision process but affect neither preferences nor the budget constraint. Distribution factors play a fundamental role in distinguishing the collective model from the unitary model. In a unitary model, such variables should not influence any household demands. In the collective setting, the household utility function can therefore be expressed as follows:

$$\tilde{U}(Q, q_A, q_B) = \mu(P, p_A, p_B, x, \mathbf{z}) u^A(Q, q_A) + [1 - \mu(P, p_A, p_B, x, \mathbf{z})] u^B(Q, q_B), \quad (3)$$

where \mathbf{z} is a vector of distribution factors. An often used example of a distribution factor is the share of income controlled or generated by one of the spouses.

The collective model has been used to understand the effects of targeting cash transfers to women (Attanasio and Lechene, 2002, 2014; Schady and Araujo, 2006). Targeted cash transfers affect not only total household income but also how it is distributed among household members. Within the framework of the unitary model, these transfers would affect household decisions only through the effect that they have on total income and the budget constraint. If we consider instead a collective model, then a targeted cash transfer could affect not only the budget constraint but also each household member's bargaining power, both through x and through each spouse's share of income. If the latter is a distribution factor affecting weights μ , then transfers will have an effect on commodity demands over and above any effects operating through total household income. As we discuss below, the context we study is that of a cash transfer which in some municipalities was targeted to women while in others it was targeted to the person who is registered as the head of household, who is a man in the large majority of cases.

The experiment we have executed and described above induces participants to consider explicitly the trade-off between the total amount of resources available to the household and those controlled by the participant. In what follows we make this link explicit. We note that while the collective model assumes efficiency, it does so conditional on given weights μ . Women with a willingness to pay for controlling certain resources of $s > 0$ are effectively sacrificing some resources. However, the behaviour of these women is not inconsistent with the collective model as, within the theoretical model we are considering, they are willing to pay s to change the weights μ . A choice of $s > 0$ is inconsistent with the unitary model.

¹¹ See e.g. Attanasio and Lechene (2014) for a similar representation and Samuelson (1956); Browning *et al.* (2014) for a general discussion of the unitary and collective model.

2.2. *Interpreting s as a Measure of Power Within the Collective Model*

Our measure of power within the marriage can be expressed, in the context of the collective model, in terms of an indifference condition, which we now derive. In what follows, to simplify the notation and the derivation of our results, we consider a version of the collective model without public goods. Public goods can, however, be re-introduced without affecting our results.¹²

Let us define the woman's share of total household income as $f = x_A / (x_A + x_B)$, where x^A and x^B are wife and husband's income, respectively. We assume that f , which is obviously affected by any transfer the wife receives, is a distribution factor and that it summarises the effect that husband and wife's income has on the Pareto weights μ .

Our experiment offers a payment either to the participant (a woman) or to her husband, and identifies the participant's willingness to pay to be the recipient of that amount. Let s denote her willingness to pay as a share of the total payment offered, E . We define f' to be the value of f that we would observe if the participant receives $(1-s)E$: $f' = [x_A + (1-s)E] / [x_A + x_B + (1-s)E]$. Similarly, define f'' to be the participant's share of resources when the husband receives E : $f'' = x_A / (x_A + x_B + E)$. Since this payment also affects the household's total income, we define x' to be the resources available when the wife receives $(1-s)E$, and x'' to be the resources available when the husband receives E : $x' = x_A + x_B + (1-s)E$ and $x'' = x_A + x_B + E$.

The indirect utility function for household member k , $v^k[x, \mu(x, f, \mathbf{z})]$ is a function of total expenditure x and the weight μ , which we assume to depend on x, f and some other distribution factors \mathbf{z} . For notational simplicity, we omit from the specification of the indirect utility (or the weights) the effect of prices, as they are not central for our argument. The indirect utility function can then be expressed as follows:

$$v^k[x, \mu(x, f, \mathbf{z})] = u^k(q_k^*), \quad k = A, B,$$

where q_k^* are the quantities maximising (3) subject to the household budget constraint. Relative to the indirect utility function that one would derive from a unitary model, we note that this expression depends on μ and, through it, on the distribution factors f and \mathbf{z} .

The indifference condition that defines our measure is then:

$$v^A[x', \mu(x', f', \mathbf{z})] = v^A[x'', \mu(x'', f'', \mathbf{z})]. \quad (4)$$

Under relatively mild conditions a sharing rule $\rho[x, \mu(x, f, \mathbf{z})]$ exists such that q_k^* solves the following problem (Bourguignon *et al.*, 2009):

$$\max u^k(q_k) \quad \text{s.t. } p_k q_k = y_k,$$

where $y_A = \rho[x, \mu(x, f, \mathbf{z})]$ and $y_B = x - \rho[x, \mu(x, f, \mathbf{z})]$.¹³ In this case, we can write the indirect utility function depending only on the sharing rule (and prices, which are kept implicit):

¹² For the case in which we use a sharing rule, we would have to consider a conditional sharing rule. One way to conceptualise such a conditional sharing rule is by considering a two-stage budgeting procedure where the couple in the first stage agrees on how much to spend on the public good, and then subsequently decides on spending on private goods.

¹³ With public goods one can instead define conditional sharing rules.

$$v^A[x, \mu(x, f, \mathbf{z})] = v^A\{\rho[x, \mu(x, f, \mathbf{z})]\}.$$

Our indifference condition, (4) can therefore be expressed in terms of the sharing rule:

$$\rho[x', \mu(x', f', \mathbf{z})] = \rho[x'', \mu(x'', f'', \mathbf{z})]. \quad (5)$$

In what follows, rather than a vector of distribution factors \mathbf{z} , we consider, without loss of generality, a single distribution factor z in addition to f . In this framework, we can prove the following Proposition.

PROPOSITION 1. *Suppose the weight $\mu(\cdot)$ does not depend on x and that the distribution factor z is such that $\partial\mu/\partial z > 0$. If $\partial^2\rho[x, \mu(f, z)]/\partial\mu^2 < 0$, $\partial^2\rho[x, \mu(f, z)]/\partial\mu\partial x > 0$ and $\partial^2\mu(f, z)/\partial z\partial f \leq 0$, then $ds/dz < 0$.*

Proof. See online Appendix D.

What the Proposition says is that our measure s is sensitive to changes in the bargaining power of women. In particular, an exogenous change in the bargaining power of women (driven by a distribution factor other than f), is reflected in a lower willingness to pay for control s . For this to be true, however, we need to have that the sharing rule is concave in μ , and that the factor we are considering is not a complement for f in the weight μ . Both are intuitive and natural conditions.

2.3. Willingness to Pay and Targeted Transfers Within the Collective Model

It is useful to explicitly introduce targeted transfers into the collective model in (1) and (3), in order to understand how they may affect the measure of empowerment we have implemented. Let T be the amount of the transfer, α the fraction of the transfer targeted to the woman in the household, and $1 - \alpha$ the fraction targeted to her spouse. A simple characterisation of our data is one in which α is either equal to 1 or 0, depending on whether the household resides in a municipality where payments are targeted to the woman or to the household head, respectively.

With a targeted transfer, the distribution factor f is given by: $f = (x_A + \alpha T) / (x_A + x_B + T)$, a measure we can construct for each household using survey data.

Introducing our measure into the picture and following the same argument as used above, we define f' to be the value of f that we would observe if the participant receives $(1 - s)E$: $f' = [x_A + \alpha T + (1 - s)E] / [x_A + x_B + T + (1 - s)E]$. Similarly, we define f'' to be the participant's share of resources when her husband receives E : $f'' = (x_A + \alpha T) / (x_A + x_B + T + E)$.

Substituting the new expressions for f' and f'' into (5), which is the indifference condition that defines our measure s , we can derive the following Proposition.

PROPOSITION 2. *If $\mu(\cdot)$ does not depend on x , $\partial^2\rho[x, \mu(f, z)]/\partial\mu^2 < 0$ and $\partial^2\rho[x, \mu(f, z)]/\partial\mu\partial x > 0$, then $\partial s/\partial\alpha < 0$.*

Proof. See online Appendix D.

This Proposition says that, under the stated conditions for the sharing rule, shifting a cash transfer from the husband to the wife, as in the programme we are considering, should reduce the reported s . We interpret this result as saying that, within the collective model, targeted transfers to women increase their bargaining power within the couple.

3. Data

As discussed in the introduction, the experiment was conducted with the respondents of a survey carried out in the winter of 2013, for the evaluation of the impact of targeting women rather than household heads with a conditional cash transfer (the Macedonian Conditional Cash Transfer for Secondary School Education, or MCCT). We are able to match the data from the experiment with this survey. In this Section, we describe the main features of our data.

3.1. *Macedonian CCT for Secondary Education*

The laboratory experiment was conducted on a sample of women living in households eligible for the MCCT. The MCCT is a social protection programme which aims to alleviate poverty, and to increase secondary school enrolment and completion rate among children in the poorest households of the population. It was first implemented by the Macedonian Ministry of Labour and Social Policy in the autumn of 2010 and provided cash transfers to poor households conditional on having school-age children attending secondary school at least 85% of the time.

The programme was offered to the beneficiaries of the Social Financial Assistance (SFA) benefit, which is the most significant income support programme in the country, accounting for around 0.5% of GDP and 50% of total spending on social assistance. SFA is a means-tested monetary transfer granted to people who are fit for work but nevertheless are unable to support themselves. The amount a household is entitled to under the SFA depends on household size and time spent in the programme, varying from 1,825 MKD per month (39.24 US\$) for a one-member household, to 4,500 MKD (96.75 US\$) for households with five or more members. The actual SFA payment to the household is the difference between the amount the household is entitled to and the actual household income. In this context, the potential MCCT payments are substantial: they amount to 1,000 MKD per month per child (21.50 US\$), and, conditional on being eligible for SFA, they are not related to income.

During the first three years of the MCCT (school years 2010/11, 2011/12 and 2012/13), a randomised field experiment was designed to test whether gender-targeted transfers generate differential outcomes in terms of household decision-making and human capital investment. The 84 municipalities in the country were randomly assigned to one of two groups.¹⁴ In the first group, payment of the CCT was made to

¹⁴ Random assignment was done after stratifying the sample of municipalities by population size. The Republic of Macedonia's 84 municipalities were first divided into seven groups, depending on population size, and then randomised into two groups of 42 municipalities each. In municipalities in the first group the grant is paid to the mother of the child, while in municipalities in the second group the grant is paid to the household head, regardless of gender.

the mother of the child, while in the second group it was made to the household head, who is generally male.¹⁵

Three waves of a household survey were collected to study the impacts of this experiment on household outcomes: one baseline and two follow-up surveys. Each survey contains detailed information on a variety of household outcomes (demographic characteristics, expenditures on durable and non durable goods, housing) and individual level information on household members (education, health, labour supply, time use). We add several sources of income to construct total household income: labour income, income from financial assistance (including assistance from the CCT programme) and assistance from family and friends. When available, we use income information for a given household from up to two survey rounds. The woman's income share, f , is then defined as the share of total parental income that can be attributed to the mother, such as her labour income and income transfers from her relatives. Sometimes it is not clear how to attribute a particular source of income to a household member. In the case of the SFA subsidy, for example, we attribute it to the household head, since the household head is the legal recipient of this transfer. However, our results are robust to reasonable changes in the definition of income shares. We discuss these issues in detail in online Appendix C.3.

In order to obtain information on the CCT transfer the households receive, we match administrative data from the CCT programme with each child who is enrolled in the programme and part of the survey. We then compute the amount of money transferred to each household in the first three years of the programme. We assign the CCT income to the household member eligible for the transfer. Further details on the programme and the evaluation can be found in Armand and Carneiro (2013).

3.2. *Laboratory Experiment*

During the summer of 2014, we invited a subset of the urban women who were sampled for the evaluation of the CCT to participate in the laboratory experiment we describe above. These women resided in urban areas and had to cohabit with a partner (the aim is to study the control over resources between spouses and, therefore, we excluded households with only one parent).¹⁶ An area is defined as urban if it is within a 10-kilometre radius from an urban settlement, as defined by Macedonian law. In particular, an urban settlement is defined as a 'compactly built up residential area with a population exceeding 3,000, which has a developed structure of various economic activities, which has over 51% of the workforce working in the secondary and tertiary sector, which has an urban physiognomy of zones for residence, recreation and green

¹⁵ According to the rulebook for acquiring the right to financial assistance, the household head is determined by the following ordered rules: if there is an employed person in the household, the household head would be the employed person; if there is a pensioner, the household head would be the pensioner; if no employed person or pensioner exists in the household, the household head is the unemployed person representing the household; for all other households, the Social Welfare Centre selects the household head as the person representing the household.

¹⁶ Note that 98% of the women are legally married and therefore we throughout refer to their partner as their husband.

area (parks), town square, street infrastructure, communal services, and which acts as a functional centre for the surrounding populated places’.

The experiment was carried out in 43 settlements, 10 of which are independent municipalities and form the capital city Skopje. An office location for the experiment was arranged in each of these settlements. By using an office, one avoids having to carry out the experiment in the household dwelling, where answers could have been affected by the family environment and where it could be difficult to isolate the woman from the presence of her husband.

Our sample includes a set of households eligible for the first year of the CCT (the 2010/11 academic year), and another set of households eligible for the third year (the 2012/13 academic year). Out of 906 selected women, 768 participated in the laboratory experiment, giving a fairly high response rate of 84.8%. In online Appendix C.1, we show that the response rate in municipalities where the CCT transfer is paid to the mother is not statistically different from the response rate in municipalities where it is paid to the head of household. The main predictors of response rate in our sample are husband’s employment, ethnicity and, for those who have the information available, baseline expenditure levels.

Table 2 presents the summary statistics for participants in the experiment in terms of age and education, ethnicity, household size and other household attributes. It also presents balance tests by displaying the differences in the characteristics for households residing in municipalities with different payment modalities for the CCT. On average, participants are 44 years of age and have 7.5 years of education. Their partners are slightly older (47 years old) and have higher levels of education (8.5 years of education). Average household size is about 4.7, the average number of children is 2.5 and the vast majority of couples are legally married (98%). Since all households in the experiment (and in the CCT) are recipients of SFA, we expect very few adult members to report to be working in the month prior to the interview. Only 9% of women and 18% of men report any employment during that period. About 14% of households are involved in farming and breeding (even though this is an urban sample). In terms of basic living standards, 90% of households have access to public water and electricity.

Figure 2 shows the distribution of participants’ willingness to pay (WTP) to gain control over the transfer in the experiment, as a fraction of the total amount offered in the experiment (E in the model). As discussed in Section 2, we interpret a high WTP as an indication of low bargaining power within the marriage.

The average WTP in this sample is equal to 19% of the maximum total amount offered in the experiment (E). It is worth noting that a substantial fraction of women (7%) reveals a willingness to pay above 0.9, meaning that they are willing to pay almost the whole amount in order to avoid that their husband becomes the recipient. This could be justified if they benefit very little, or even experience negative effects, from their husband’s increase in income.¹⁷ Nevertheless, a very high value for WTP could also reflect a situation where mothers gain substantial control over resources by increasing only slightly her contribution to total household income.

¹⁷ Examples are households where the husbands are alcoholic or have other consumption habits which make their spending non-beneficial or even invoke negative effects for the spouse, such as spending on prostitution.

Table 2
Participant Characteristics and Randomisation Balance

	Payment to household head (1)	Difference with payment to mother (2)
Age (husband)	47.20 (0.49)	-0.19 (0.65)
Age difference (h - w)	3.29 (0.26)	-0.24 (0.37)
Schooling (husband)	8.45 (0.31)	0.19 (0.36)
Schooling difference (h - w)	0.96 (0.18)	-0.26 (0.30)
Albanian	0.25 (0.07)	0.06 (0.11)
Macedonian	0.47 (0.07)	-0.05 (0.10)
Roma	0.18 (0.05)	0.02 (0.07)
Turk	0.10 (0.04)	-0.02 (0.04)
Household members	4.70 (0.12)	-0.04 (0.17)
Number of children	2.52 (0.09)	0.07 (0.15)
Legally married	0.98 (0.01)	-0.01 (0.01)
Worked (wife)	0.09 (0.03)	0.00 (0.03)
Worked (husband)	0.18 (0.04)	0.02 (0.05)
Male household head	0.84 (0.03)	-0.01 (0.05)
Muslim	0.56 (0.07)	0.04 (0.10)
Farmer/breeder	0.14 (0.04)	0.01 (0.05)
Access to public water and electricity	0.90 (0.02)	0.01 (0.03)
Living in Skopje	0.29 (0.11)	0.00 (0.16)
Living in main settlement	0.77 (0.05)	0.01 (0.07)
Observations	768	768
Enrolled in CCT (2012/3)	0.64 (0.03)	-0.05 (0.04)

Notes. Standard errors clustered at municipality level are presented in parenthesis (*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). This Table shows estimates of a linear regression of the variables indicated in the first column on the treatment indicator $mother_j$ and a constant. Column (1) presents estimates for the constant. Column (2) presents instead estimates of the coefficient for the treatment indicator $mother_j$. The inclusion of seven dummies for the randomisation strata (each dummy indicating a quantile for the municipality in the distribution of the population) does not affect the results.

There is also a large fraction of women who report a zero and negative willingness to pay (35%). The interpretation of a zero or negative WTP is not straightforward. A woman living in a unitary household would report a zero-WTP. A woman with a very

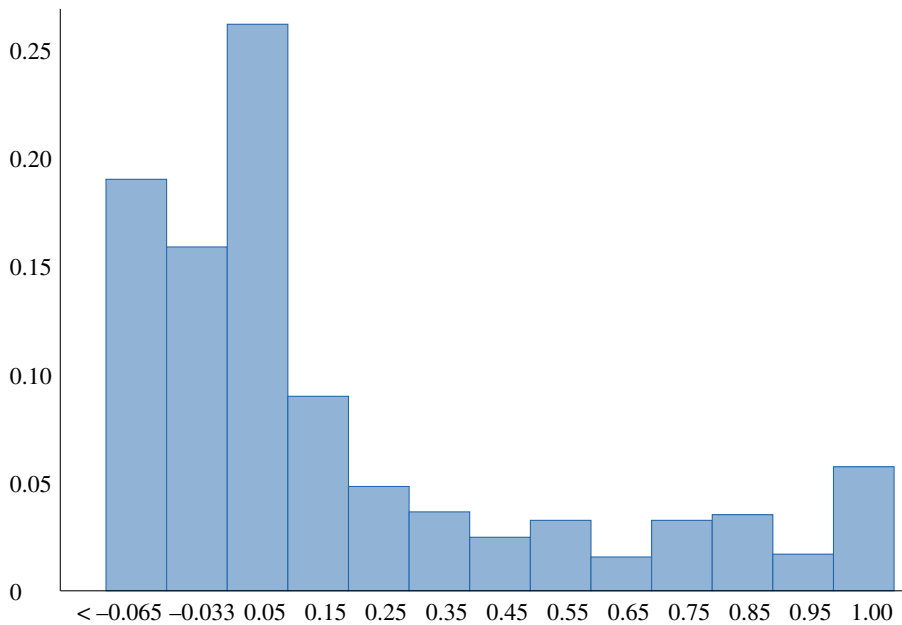


Fig. 2. *Willingness to Pay*

Notes. The Figure presents the distribution of participants by the share the participant is willing to pay in order to receive the payment. The left bar is representing all responses smaller or equal to -0.065 , which is the smallest censored share and is equal to 50 MKD divided by the maximum stake, 800 MKD. Colour Figure can be viewed at wileyonlinelibrary.com.

high level of bargaining power, so that she is effectively the sole decision-maker in the household, would also report a zero-WTP. This would be equivalent to a unitary model where household preferences coincide with the woman's preference. A zero-WTP could also be obtained in the case of a woman with no bargaining power, perhaps because of social norms. This would be equivalent to a unitary model where the preferences coincide with the husband's preference. It could also be the case that, for some women reporting zero-WTP, their bargaining power is very low and the amounts offered in the game are not sufficiently high to significantly change their relative bargaining weights.

This may be the case if women who have very little bargaining power initially, need to obtain a certain level of income before transfers of this kind matter, i.e. they need a 'big push' in order to reach the state where smaller contributions to household income matter for the power balance. Finally, an alternative explanation for a WTP equal to zero may be that some individuals are reluctant to make decisions and, hence, they shy away from decision-making if they are given the opportunity to do so. This explanation is consistent with evidence from recent experimental studies.¹⁸ For these reasons, the zero-WTP observations should be interpreted cautiously: throughout the empirical analysis we therefore present robustness analyses treating zero-WTP observation in different ways.

¹⁸ See Gartner and Sandberg (2015) and references therein.

3.3. Survey Questions: Traditional Measures of Empowerment

At the end of the laboratory experiment, we conducted an additional short survey. The survey included further questions about the within-household decision-making process, depression, domestic conflict and violence, prospects of future work opportunities, networks for financial assistance and private goods consumption by the participant and her spouse (cell phone bills, food for children, cigarette and alcohol consumption for both adults). The questions on decision-making are very similar to those traditionally used in many surveys (such as those used in the PROGRESA evaluation and analysed by Adato *et al.* (2000) in the study we cited above). In particular, we ask whether the main decision-maker for a variety of choices, such as schooling, expenditure on food or finance, is the husband, the wife, or whether decisions are made jointly by both spouses.¹⁹

Table 3 presents descriptive statistics for answers concerning household decision-making and domestic violence. In terms of decision-making, we observe that, as in many similar surveys, participants tend to report that most decisions are taken jointly with their husbands. In fact, 81% of participants report to be choosing together with their husband for decisions related to schooling, 58% for choices related to food expenditures and 87% for choices related to additional income. In terms of domestic violence, the share of women who find it justifiable, in some circumstances, to be beaten by their husbands, ranges from 10% (when the violence is linked to arguing against the husband) to 33% (when the violence is related to neglecting the children). Similar levels of acceptance of domestic violence are seen when the participant is asked whether domestic violence is widespread in her neighbourhood of residence. Only 30–35% of participants report that they argued with their husbands about the management of money or the children's discipline during the week before the interview. Our data show that a high fraction of participants presents symptoms of depression and a large fraction reports a very low perceived likelihood of divorce.

Using the information collected from survey questions, we build two indexes: a household decision-making (HDM) index and a domestic violence (DV) index.²⁰ The HDM index is constructed performing factor analysis (FA) using a polychoric correlation matrix on the variables concerning the identity of the decision-maker within the household when it comes to children's schooling, food expenses, household financial administration and extra income. These variables are coded as 0 if the woman is deciding, 1 if the two partners are deciding together and 2 if the man is deciding. Therefore, like our WTP variable, this index is decreasing in women's power.

The DV index is constructed using principal component analysis (PCA) on the selected dummy variables and we construct the index using the first component only. This index captures the participant's attitudes towards domestic violence and her perceptions of domestic violence in the neighbourhood of residence. We focus on whether the participant believes that domestic violence is acceptable if a woman argues

¹⁹ The full questionnaire can be found in online Appendix A.2. In asking these questions, we did not probe the participants on 'tie breakers' when they respond that they both decide on something, as done, for instance, by Ashraf (2009).

²⁰ Domestic violence has been studied in relation to targeted cash transfers as another measure of empowerment (or rather lack thereof; Angelucci, 2008; Bobonis *et al.*, 2013).

Table 3
Descriptive Statistics on Self-reported Empowerment Indicators

	(1) Mean	(2) Median	(3) SD	(4) Min	(5) Max
Wife decides about school	0.12	0.00	0.33	0	1
Wife decides about food	0.24	0.00	0.42	0	1
Wife decides about finance	0.27	0.00	0.44	0	1
Wife decides about extra amount	0.08	0.00	0.27	0	1
Husband decides about school	0.06	0.00	0.24	0	1
Husband decides about food	0.19	0.00	0.39	0	1
Husband decides about finance	0.30	0.00	0.46	0	1
Husband decides about extra amount	0.04	0.00	0.21	0	1
Both decide about school	0.81	1.00	0.39	0	1
Both decide about food	0.58	1.00	0.49	0	1
Both decide about finance	0.43	0.00	0.50	0	1
Both decide about extra amount	0.87	1.00	0.33	0	1
Violence justified for argument	0.10	0.00	0.29	0	1
Violence justified for going out	0.16	0.00	0.37	0	1
Violence justified for neglecting children	0.33	0.00	0.47	0	1
Violence justified for burning food	0.05	0.00	0.22	0	1
Violence common for argument	0.28	0.00	0.45	0	1
Violence common for going out	0.21	0.00	0.41	0	1
Violence common for neglecting children	0.20	0.00	0.40	0	1
Violence common for burning food	0.10	0.00	0.30	0	1
Argued about managing money	0.35	0.00	0.48	0	1
Argued about children's discipline	0.30	0.00	0.46	0	1
Depression index	13.01	13.00	6.34	0	30
Presence of depression symptoms	0.70	1.00	0.46	0	1
High likelihood of divorce in the neighbourhood	0.22	0.00	0.41	0	1

Notes. All answers are reported by the participant at the end of the WTP experiment. Depression index is based on CES-D10 test. Presence of depression symptoms is a dummy variable equal to one if the CES-D10 depression value is equal or larger than 10. For variables concerning likelihood, we refer to 'high likelihood' if answers were 'Very likely' and 'Somewhat likely', instead of 'Could happen', 'Unlikely' or 'Very unlikely'. For variables concerning decisions, we refer to 'wife (husband) decides' if the participant (partner) is the only decision-maker. The exact questions are reported in online Appendix A.2.

with her partner, goes out without informing her partner, neglects her children, or burns the food while cooking. In addition, we use information on whether the participant believes that these types of domestic violence are common in her neighbourhood.

Table 4 presents the factor loadings for the DV index and the HDM index. All items, which have roughly the same scale and variance, contribute in a similar way to each index. There is no item or set of items on which the index loads particularly. That said, for HDM, the loadings on decisions about food expenditures and financial administration in the household are a little larger than for the other items, which means that the index is more reflective of household decision-making about these issues than about the schooling of children, or about managing extra income. Similarly, one should also note that, for DV, the loadings on the questions about the prevalence of domestic violence are a little higher than those on questions about the acceptability of domestic violence, perhaps because there is less measurement error or random variability in these questions.

Figure 3 presents the distribution of the HDM and the DV indexes. The HDM index has a bell-shaped distribution where the DV index has a large mass point at 0,

Table 4
Factor Loadings for HDM and DV Indexes

	Factor loadings	
	HDM index (1)	DV index (2)
Participation in the decision about:		
School	0.510	
Food expenses	0.714	
Financial administration	0.701	
Managing an extra amount	0.516	
Violence justified for:		
Argument		0.280
Going out		0.289
Neglecting children		0.299
Burning food		0.234
Violence common for:		
Argument		0.380
Going out		0.447
Neglecting children		0.450
Burning food		0.386
Observations	768	768
Share of total variance explained	–	0.401

Notes. For the HDM index, factor loadings are computed performing factor analysis using a polychoric correlation matrix. Variables about household decision-making are coded as 0 if the woman is deciding, 1 if the two partners are deciding together and 2 if the man is deciding. For the DV index, the Table presents the principal component analysis factor loadings of the first component. ‘Violence justified for’ refers to the question ‘Sometimes a husband is annoyed or angered by things that his wife does. In your opinion, is a husband justified in hitting or beating his wife if ...’. ‘Violence common for’ refers to the question ‘In your neighbourhood, is it usual for husbands to beat the wives if ...’.

reflecting the pattern of answers to the underlying questions. For both indexes, there is large dispersion in the answers of different households.

3.4. *A Comparison of Empowerment Measures*

The empowerment measure we have collected, based on the willingness to pay for resource control, is novel in the literature. It is therefore interesting to compare it with more traditional measures. We do so by first computing the correlation between the WTP measure and the HDM and DV indexes. These correlations are reported in Table 5. In the left panel of the Table, we report the correlations using the entire sample. Remarkably, the correlation between the WTP and indexes based on traditional measures is negative and, in the case of the HDM index, significantly so. The correlation becomes smaller, but remains negative, when we omit from the sample the participants who revealed zero and negative WTP, as we can see in the right panel.

The negative correlation is, at least at first sight, puzzling, however, it is obvious that the two sets of measures are capturing very different factors. One possibility is that all the measures are affected by a variety of factors, one of which is ‘power’ within the household, which then load differently on to the observed variables.

This idea is captured by the following simple model. Let there be I candidate measures of empowerment, $(m_1, \dots, m_b, \dots, m_I)$. Each measure is assumed to reflect

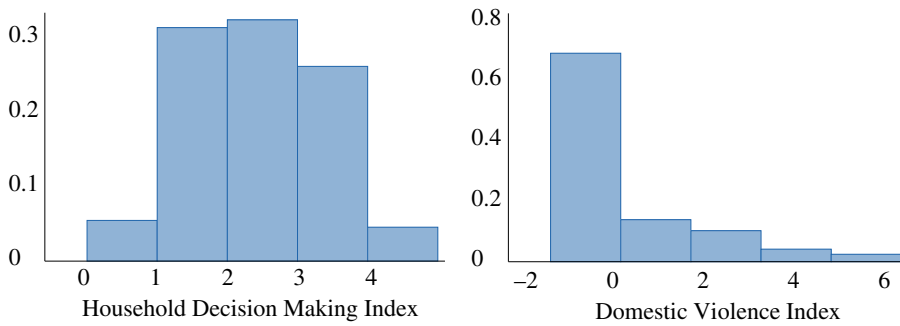


Fig. 3. *Distribution of Domestic Violence and Household Decision-making Indexes*

Notes. The graph presents the distribution of the HDM index (left panel) and of the DV index (right panel). Indexes are built using principal component analysis and using the first component only. See Table 4 for factor loadings. Colour Figure can be viewed at wileyonlinelibrary.com.

the true level of women empowerment, P , but may also reflect effects from one or more confounding factor(s), X , and pure random noise specific to each measure, ε_i . Based on this, we can express each measure as follows:

$$m_i = \alpha_i + \lambda_i P + \beta_i X + \varepsilon_i, \tag{6}$$

where λ_i gives the degree to which the measure reveals actual power in the household and β_i gives the effect of a potential confounding factor. A negative correlation between two measures can arise if different measures load some X s with opposite sign and these variables are sufficiently variable to overturn the positive correlation that would arise from different measures reflecting P .

We have attempted to estimate the model of (6) in our data. Since we only have three measures of empowerment, WTP, HDM and DV, we only allow for one unobservable factor in our model. We include several observable characteristics, as documented in Table 6, which shows the loadings on the unobservable factor and the coefficients on all the observable regressors.²¹

After controlling for several control variables the λ in the HDM is still negative and statistically different from zero, suggesting that, if the above story is reasonable, there may be additional determinants of the power measures we collect which are unobserved in our dataset and which may have impacts of opposite signs on different measures. In fact, looking at the observables, this may well be the case.

²¹ Let m_1 be the WTP measure, m_2 be the HDM index and m_3 the DV index. We estimate the following single-factor model:

$$\begin{aligned} m_1 &= \alpha_i + \lambda_1 P + \beta_1 X + \varepsilon_1, \\ m_2 &= \alpha_i + \lambda_2 P + \beta_2 X + \varepsilon_2, \\ m_3 &= \alpha_i + \lambda_3 P + \beta_3 X + \varepsilon_3, \end{aligned}$$

where λ is the loading on the unobservable factor and β is the coefficients on observable regressors. We normalise $\lambda = 1$ in the WTP equation. In addition, for estimation purposes, we impose a left censoring point at 0 and a right censoring point at 1 for WTP, a left censoring point at 0 for HDM index, and a left censoring point at -1 for DV index, in accordance with the data. P , ε_1 , ε_2 and ε_3 are i.i.d. normal. We estimate the model using maximum likelihood.

Table 5
Correlation Matrix for Empowerment Measures

	Include all observations		Exclude zero and negative-WTP	
	WTP (1)	HDM index (2)	WTP (3)	HDM index (4)
HDM index	-0.160***	-	-0.078	-
DV index	-0.009	0.096**	-0.035	0.067
Observations	768	768	500	500

Notes.***p < 0.01, **p < 0.05, *p < 0.1. The Table presents correlation coefficients between the different measures of empowerment. WTP is defined as the share of the transfer the participant is willing to pay to become the recipient instead of her husband. HDM index is built using information on who within the household is participating in decisions. DV index is built using information on participant's attitudes towards domestic violence. Details on the constructions of the indexes are presented in subsection 3.3. In columns (3) and (4) we restrict the sample by excluding the participants that in the laboratory experiment reported a zero or negative WTP.

Notice that some variables, such as the gender of the household head, religion, number of children, or ethnicity, correlate differently with the alternative measures. For instance, male-headed households has a significant negative association with WTP but a significant positive association with the HDM. Some of the other variables also show signs of opposite associations but many of these estimated effects are insignificant.

In the empirical analysis, we use the CCT transfer to measure exogenous changes in power. Note that, as we have exogenous variation from the field experiment, we can study how well our measure reflects power, even though we may not have good direct measures of the confounding factors, some of which may be unobserved. In Section 4, we evaluate our measure as well as the alternative measures derived from responses to survey questions about empowerment. To this end, we use both reduced form and instrumental variables approaches.

4. Empirical Analysis

Among participants, women residing in municipalities where the CCT payment was targeted to them have potentially been empowered by an increase in their share of household income, when compared to women living in municipalities where the CCT was paid to household heads. In this Section, we describe the empirical methodology used to test whether the payment modality of the CCT affects women's empowerment, and we present results using as outcomes the traditional measures of empowerment, and the WTP measure from the laboratory experiment.

Random assignment of the identity of the recipient of the transfer (mother *versus* head of household) across municipalities allows us to assess the impact of female empowerment on different variables, by comparing these variables between residents in the two types of municipalities. Let $mother_j$ be an indicator variable taking the value 1 if the cash transfer is offered to mothers in municipality j , and zero otherwise. The outcome of interest is a measure of participant i 's (living in municipality j) empowerment, y_{ij} .

Table 6
Single-factor Model Estimates for Empowerment Measures

	Empowerment measures		
	WTP	HDM index	DV index
	m_1 (1)	m_2 (2)	m_3 (3)
<i>P</i>	1	-12.701***	-2.338
	(constrained)	(4.421)	(1.682)
Age (husband)	-0.005 (0.003)	-0.015* (0.008)	0.023 (0.022)
Age difference (h - w)	0.005 (0.005)	0.000 (0.010)	-0.007 (0.029)
Schooling (husband)	0.008 (0.007)	0.005 (0.016)	-0.175*** (0.047)
Schooling difference (h - w)	-0.002 (0.007)	0.006 (0.016)	0.173*** (0.045)
Adult members	-0.003 (0.037)	-0.011 (0.067)	-0.241 (0.209)
Number of children	0.021 (0.016)	0.027 (0.037)	0.278*** (0.094)
Legally married	-0.199* (0.113)	-0.384* (0.219)	0.580 (0.748)
Male household head	-0.168*** (0.045)	0.445*** (0.099)	0.356 (0.312)
Muslim	0.159* (0.085)	0.186 (0.177)	-0.279 (0.508)
Living in Skopje	-0.109 (0.070)	-0.005 (0.142)	-0.915** (0.371)
Living in main settlement	0.025 (0.041)	-0.083 (0.091)	-0.044 (0.264)
Albanian	-0.317*** (0.090)	0.284 (0.198)	-0.024 (0.570)
Turkish	-0.168 (0.102)	-0.066 (0.228)	-0.279 (0.628)
Roma	-0.150* (0.089)	0.098 (0.186)	0.964* (0.549)
Regional controls	Yes	Yes	Yes
Observations	768	768	768

Notes. ***p < 0.01, **p < 0.05, *p < 0.1. The Table presents estimates for a single-factor model described by (6). The model is estimated with maximum likelihood and constraining λ_1 to be equal to 1. HDM index is built using information on who within the household is participating in decisions. DV index is built using information on participant's attitudes towards domestic violence. Details on the constructions of the indexes are presented in subsection 3.3. Controls include age and education of partner's, the gender of the household head, religion and ethnicity of the household, regional dummies, indicator variables for the type of settlement.

We estimate the following relationship:

$$y_{ij} = \beta_0 + \beta_1 mother_j + \mathbf{X}'_i \beta_2 + \mathbf{V}'_j \beta_3 + \varepsilon_{ij}, \tag{7}$$

where \mathbf{X}_i is a vector of participant, spouse, and household characteristics, \mathbf{V}_j is a vector of settlement and municipality characteristics, and ε_{ij} is the error term. We estimate this equation using least squares, allowing for clustered standard errors at the municipality level. The variables in X include household head's and partner's education, age and gender, ethnicity and religion of the household, and household size and composition.

Municipality controls include regional dummies, an indicator for whether the participant resides in the main settlement of the municipality and whether the municipality is part of the capital city (Skopje).

Not all women in the sample are actual recipients of the CCT, due to the imperfect take-up of the programme. In order to go beyond intent-to-treat estimates of (7) and estimate the impact of receiving a transfer paid to the mother as opposed to receiving a transfer paid to her husband, we need to address two potential sources of endogeneity in the take-up of the programme.

First, programme participation is voluntary. In the year just preceding the survey, only 62% of all eligible households enrolled in the programme. Such a low take-up of the programme is directly related to the decision of not enrolling the child in secondary school, which is a condition for receiving the transfer.²²

Second, there is a proportion of households where the person registered at the social welfare centre as the household head is the mother and not her husband. The choice of who to declare as household head could be driven by unobservable variables which also affect the outcome. This decision is likely to have been made prior to the introduction of CCT and is related to a prior application that the household made for SFA benefits.

Therefore, when estimating the impact on empowerment of cash transfers received by mothers, we also use an instrumental variables strategy (IV), where the instrument for cash transfers received by the mother is the modality of payment in each municipality, $mother_j$. In this Section, we focus on two different measures of take-up of the CCT programme: the mother's income share (corresponding directly to our theoretical discussion) and the total income from the CCT received by the mother in the first three years of the programme.

Let d_{ij} be the endogenous regressor we are interested in. Then we estimate the following model:

$$\begin{aligned} y_{ij} &= \beta_0 + \beta_1 d_{ij} + \mathbf{X}'_i \boldsymbol{\beta}_2 + \mathbf{V}'_j \boldsymbol{\beta}_3 + \varepsilon_{ij}, \\ d_{ij} &= \theta_0 + \theta_1 mother_j + \mathbf{X}'_i \boldsymbol{\theta}_2 + \mathbf{V}'_j \boldsymbol{\theta}_3 + \omega_{ij}, \end{aligned} \quad (8)$$

where \mathbf{X}_i is a vector of household characteristics, \mathbf{V}_j is a vector of municipality characteristics, and ε_{ij} and ω_{ij} are household-specific error terms. As above, we compute standard errors accounting for clustering at municipality level.

4.1. Traditional Measures of Empowerment

In this subsection, we analyse the survey measures of empowerment discussed above, and discuss how they are affected by gender-targeted CCT payments. We focus on the HDM and DV indexes, presented in subsection 3.3. Table 7 presents estimates for the

²² There is also a set of households that do not participate in the CCT because they lost the right to SFA, and therefore indirectly lost the right to apply to the CCT programme. We do not address this issue explicitly. However, when we match the eligible SFA population in 2010 (baseline) and the eligible SFA population in 2013 (second follow-up), we obtain fairly high match rates across different types of municipalities, suggesting that the severity of this problem may be uncorrelated with the identity of the CCT recipient.

Table 7
Effect of Targeting Payments to Women on HDM and DV Indexes

Dependent variable:	HDM index			DV index		
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	IV	OLS	IV	IV
Payment to mother	-0.112 (0.090)			-0.200 (0.172)		
Mother's income share		-0.596 (0.468)			-1.069 (0.902)	
Mother's CCT income			-0.007 (0.006)			-0.013 (0.010)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	768	768	768	768	768	768

Notes. Standard errors in parenthesis are clustered at municipality level (***p < 0.01, **p < 0.05, *p < 0.1). Columns (1) and (4) present estimates using (7) for the effect of targeting payments to mother on the HDM and the DV indexes. Columns (2)–(3) and (5)–(6) presents estimates for (8) on the effect of mother's income share and on total CCT transfer to the mother on willingness to pay. HDM index is built using information on who within the household is participating in decisions. DV index is built using information on participant's attitudes towards domestic violence. Details about the indexes are presented in subsection 3.3. Controls include age and education of partners, the gender of the household head, religion and ethnicity of the household, regional dummies, indicator variables for the type of settlement and indicator variables for the stake.

effect of gender-targeted payments where we consider these indexes as the outcome variables.

In columns (1) and (4), we focus on intent-to-treat estimates (equation (7)), corresponding to the impacts of residing in a municipality where CCT payments were transferred to mothers. In columns (2) and (5), we report IV estimates (equation (8)) of the effect of mother's income share (instrumented by the treatment variable) on the empowerment indexes, while in columns (3) and (6) we report corresponding IV estimates of the effect of mother's total CCT transfer. First-stage estimates are reported in the lower panel of Table 9. There are no statistically significant coefficients in this Table. However, the point estimates indicate that gender-targeted transfers increase empowerment among women who receive those transfers.²³

4.2. Laboratory-based Measure of Empowerment

We now turn our attention to the WTP measure from the laboratory experiment. In columns (1)–(3) of Table 8, we present ITT estimates of β_1 , using different combinations of controls. All estimates are negative and statistically significant, indicating that in municipalities where mothers are the recipients of the CCT (and

²³ Note that the raw correlation between our measure and the HDM is negative, i.e., the willingness to pay and this index are positively correlated (Table 5). This can be explained by the confounding factors in (6) affecting some of the measures. Denoting our new measure m_1 and one of the alternative measures m_2 , and considering the covariance between p and X to be positive (i.e. $\lambda_1 > 0$, $\beta_1 = 0$, $\lambda_2 \approx 0$ and $\beta_2 < 0$), the covariance between the two measures will be negative. We find no significant correlation between our measure and the DV index, which may again be a result of confounding factors, or may simply reflect that the DV index does not reveal (economic) power.

Table 8
ITT Estimates of the Effect of Targeting Payments on Willingness to Pay

	Dependent variable: willingness to pay					
	Include all observations			Exclude always husband and always herself		
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	OLS	OLS
Payment to mother	-0.057** (0.025)	-0.053** (0.024)	-0.053** (0.024)	-0.058*** (0.021)	-0.055*** (0.019)	-0.055*** (0.020)
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Ethnicity controls	No	Yes	Yes	No	Yes	Yes
Stake controls	No	No	Yes	No	No	Yes
R ²	0.055	0.074	0.074	0.060	0.082	0.083
Observations	768	768	768	576	576	576

Notes. Standard errors in parenthesis are clustered at municipality level (***p < 0.01, **p < 0.05, *p < 0.1). The Table presents estimates using (7) for the effect of targeting payments to mothers on the WTP. Dependent variable is WTP defined as the share of transfer the participant is willing to pay to become the recipient instead of her husband. Payment to mother is a dummy variable equal to one if the household resides in a municipality where the CCT transferred the money to mothers. In columns (4)–(6), the sample is restricted by excluding the participants who decided to always be the recipient or always choosing the husband to be the recipient. Controls include age and education of partner's, the gender of the household head, religion and ethnicity of the household, regional dummies, indicator variables for the type of settlement and indicator variables for the stake.

therefore are potentially more empowered), women are on average willing to pay a lower amount to keep the cash from the laboratory experiment than in municipalities where the CCT recipient is the head of household (and therefore, the level of empowerment of women is potentially smaller). Since a large proportion of women report extreme values for s_{ij} (either 0 or 1), columns (4)–(6) of Table 8 examine what happens when we exclude these extremes from our analysis. Although we are using a substantially smaller sample size, our results are essentially unchanged.

We then estimate (8) using the mother's income share as the main endogenous variable. The lower panel of Table 9 shows that the instrumental variable *mother_j* strongly predicts the mother's income share, which is 19 percentage points higher in municipalities where mothers are the recipients of the CCT, compared to municipalities where the household heads are recipients. Our IV estimates show that mother's income share has a significant effect on the willingness to pay to become the recipient in the laboratory experiment. Shifting all income from the male partner to the female partner in the household would decrease this willingness to pay by about 28 percentage points. A slightly larger effect is observed when we exclude from the sample those participants who decided to either always be the recipient themselves or always chose her husband to be the recipient but, overall, results are robust to the inclusion of controls in the model and to the exclusion of extreme values. If instead we use the amount of the total CCT transfer received by the mother as the explanatory variable, we estimate that an increase of 10,000 MKD in the transfer to the mother reduces her willingness to pay by around 3 percentage points.²⁴

²⁴ We obtain similar results when the endogenous variable is the number of years in which a mother received transfers from the CCT. See online Appendix C.4.

Table 9
Effect of Mother's Income Share and CCT Transfer on Willingness to Pay

	Dependent variable: willingness to pay			
	Include all observations		Exclude always herself and always husband	
	(1)	(2)	(3)	(4)
	IV	IV	IV	IV
Wife's income share	-0.283** (0.114)		-0.328*** (0.122)	
Wife's CCT income		-0.003** (0.002)		-0.003*** (0.001)
Controls	Yes	Yes	Yes	Yes
Observations	768	768	576	576
First stage results:				
Payment to mother	0.187*** 0.023	15.655*** 1.250	0.168*** 0.025	15.728*** 1.363
R ²	0.436	0.356	0.434	0.359
F-test of excluded instrument	68.750	156.805	44.869	133.201

Notes. Standard errors in parenthesis are clustered at municipality level (***p < 0.01, **p < 0.05, *p < 0.1). The Table presents estimates for (8) on the effect of mother's income share, defined as the share attributable to the mother of total household income, and on total CCT transfer to the mother on WTP. Dependent variable is WTP defined as the share of transfer the participant is willing to pay in order to receive the money instead of the partner receiving the money. Endogenous variables are instrumented using the dummy variable Payment to mother, equal to 1 if the household resides in a municipality where the CCT transferred the money to mothers. In columns (3)–(4), the sample is restricted by excluding the participants who decided to always be the recipient themselves or always choosing their husband to be the recipient. Controls include age and education of partners, the gender of the household head, religion and ethnicity of the household, regional dummies, indicator variables for the type of settlement and indicator variables for the stake.

4.2.1. *Censoring of willingness to pay*

The data from our laboratory experiment is left and right censored, since we could not elicit WTP over an infinitely large support. On one hand, a participant may have been willing to pay an even larger amount than the maximum proposed in order to secure that her husband becomes the recipient. On the other hand, we never allow the WTP to go above 1, which could be a possibility, if the participant is willing to pay in order to avoid that the partner becomes the recipient.

The exact censoring points in our data differ depending on the stakes.²⁵ The underlying willingness to pay (s_i^*) is therefore not observed beyond bounds which are determined by the initial stake (E_i) and by the rules of the experiment. In addition, the exact realisation of s_i is never observed as the software is designed to stop the sequence when the difference between two consecutive offered amounts with opposite decisions (to become the recipient or her husband to become the recipient) is smaller than 20 MKD.

²⁵ For example, when the initial stake is 600 MKD and the participant always chooses to be the recipient, the willingness to pay is right censored at 0.9375 (i.e. $(600 - 37.5)/600$). If instead the participant always chooses that the husband is the recipient, the answer is instead left censored at -0.0834 (i.e. $(600 - 650)/600$). Table B1 presents censoring points for each stake both in terms of last amount offered to the participant and in terms of the corresponding willingness to pay.

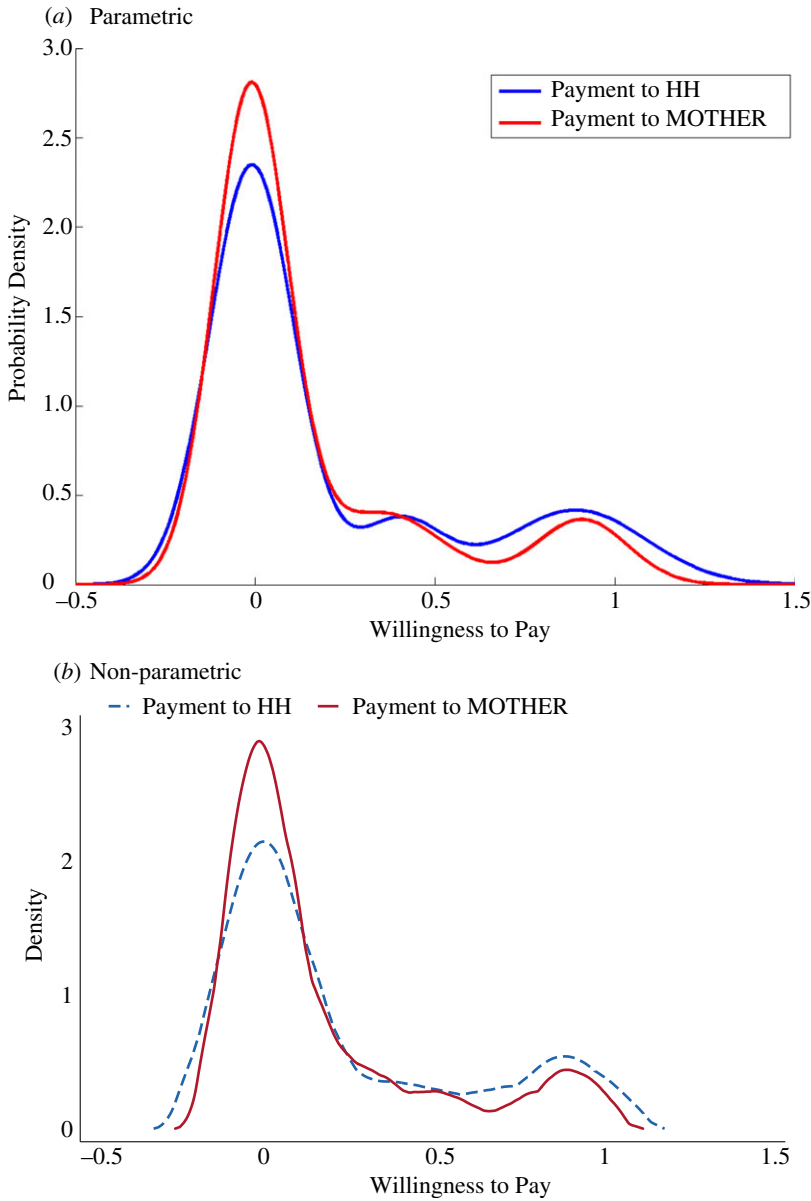


Fig. 4. *Distribution Fit For Willingness to Pay: Comparison of Treatment Groups. (a) Parametric; (b) Non-parametric*

Notes. The Figure shows a comparison of the distribution of WTP estimated for the two treatment groups (payment to household head and payment to mother). In panel (a), the distribution fit is computed assuming a mixture of three Normal distributions. The parameters are estimated using maximum likelihood and imposing multiple censoring points. Estimates are presented in Table 10. In panel (b), the distribution fit is estimated non-parametrically using Kernel density. Colour Figure can be viewed at wileyonlinelibrary.com.

Table 10
Parameters of Parametric Fit for Willingness to Pay: Distribution

	Payment to HH (1)	Payment to mother (2)
Mixture distribution		
Mean (μ)	0.218 (0.020)	0.155 (0.017)
SD	0.388	0.332
First component		
Weight (w_1)	0.690 (0.027)	0.720 (0.052)
Mean (μ_1)	-0.010 (0.008)	-0.012 (0.010)
SD (σ_1)	0.117 (0.009)	0.103 (0.009)
Second component		
Weight (w_2)	0.106 (0.037)	0.165 (0.055)
Mean (μ_2)	0.402 (0.045)	0.355 (0.064)
SD (σ_2)	0.118 (0.035)	0.165 (0.061)
Third component		
Weight (w_3)	0.204 (0.035)	0.115 (0.020)
Mean (μ_3)	0.891 (0.047)	0.909 (0.026)
SD (σ_3)	0.195 (0.056)	0.127 (0.029)
Observations	768	768
Test for equality of means [†] :		
Difference in means ($\mu_{HH} - \mu_M$)		-0.063**
t-test (p-value)		2.40 (0.017)

Notes. Standard errors in parenthesis. The distribution fit is computed assuming a mixture of three normal distributions. Parameters are estimated using maximum likelihood and imposing left/right censoring points and bounds for each observation. [†]Test for equality of the mean in the two groups is carried out assuming independence between the two groups and using a t-test = $(\mu_{HH} - \mu_M) / \sqrt{s.e.(\mu_{HH})^2 + s.e.(\mu_M)^2}$, where μ_{HH} is the mean for participants living in municipalities where the CCT payment was transferred to household heads and μ_M for participants living where the CCT payment was transferred to mothers. ***p < 0.01, **p < 0.05, *p < 0.1.

In this subsection, we estimate the full distribution of willingness to pay using the censored data and a maximum likelihood procedure which tries to fit a mixture of three normal distributions to the data, accounting for the censoring in the data. We estimate separate models for participants living in the two different types of municipalities, distinguished by the identity of the recipient of the CCT.²⁶

²⁶ Online Appendix B discusses in detail the methodology used to estimate the parameters of the distribution. We report the result for an unconditional version of the maximum likelihood estimation. We extend the estimation using a version conditional on observable characteristics and the results are unchanged.

The top panel of Figure 4 presents a comparison of the fitted distribution s_i for women residing in the two groups of municipalities. It is clear that those residing in municipalities where the CCT is paid to the mother have a lower s_i . The bottom panel of the Figure shows the non-parametric density fit to the raw data, which has more limited support because of the censoring.²⁷ The two pictures are very similar.

Table 10 presents the estimated parameters of the distributions. Since we fit a mixture of three normal distributions, we report the weight, the mean and the standard deviation of each of the three components. In addition, at the top of the Table, we also report the overall mean and standard deviation of the overall distribution.

In order to test for equality of means of s_i across the two groups of municipalities, we assume independence between these groups, and we use a standard two tailed t-test. There is a difference of 6 percentage points in s_i for participants in each type of municipality, which is statistically different from zero. The estimates from this parametric model are similar from the regression estimates from the previous Section, which did not account for censoring of the observations, providing further evidence of the robustness of our results.

5. Conclusion

Identifying the empowerment effect of targeted cash transfers is fundamental to understanding the effect of targeting women as an instrument for empowering women within households. In this article, we present a novel identification strategy to measure women's willingness to pay for receiving transfers. We report results from an economic experiment among female urban recipients of SFA who are also part of a household that was offered a CCT.

Targeted transfers alter household decision-making through (at least) two channels. First, the transfer has an effect on total household income which may affect bargaining positions for men and women directly. Second, the transfer has an effect on the share of resources attributable to each household member. Our laboratory experiment identified the values that make the women indifferent between being the recipient and her husband being the recipient, providing information about the trade-off the women makes between household income and empowerment. Our results show that women are, on average, willing to sacrifice some household income to receive the money and gain more power over resources. Note that this result means that the unitary model is generally rejected in our study. Our results further show that having already been empowered by the CCT (i.e. residing in a municipality where women were offered the CCT) leads, on average, to a lower willingness to sacrifice household income to gain power.

We provide evidence that our experimental set-up measures bargaining power, and do so in a more effective way than traditional survey-based measures.

²⁷ To test for equality of the distributions of WTP in municipalities where the CCT is transferred to mothers and municipalities where it is transferred to household heads (Figure 4), we perform a non-parametric two-sample Kolmogorov–Smirnov test. The p-value of the test is equal to 0.101 if we include all observations and 0.062 if we restrict the sample by excluding the participants who decided to always become the recipient or always letting their husband become the recipient. We need to highlight, though, that in case of censored data, results from a general version of the Kolmogorov–Smirnov test are conservative (Schumacher, 1984). This supports the rejection of equality of the two distributions.

*IIES – Stockholm University and Norwegian School of Economics
University of Navarra and Navarra Centre for International Development
University College London and IFS
University College London, IFS and NBER, and Cemmap*

Additional Supporting Information may be found in the online version of this article:

Appendix A. Design: Experiment and Survey.

Appendix B. Parametric Distribution Fit For Willingness to Pay.

Appendix C. Additional Data Analysis.

Appendix D. Proof of Propositions 1 and 2.

Data S1.

References

- Adato, M., de la Briere, B., Mindek, D. and Quisumbing, A.R. (2000). 'The impact of PROGRESA on women's status and intrahousehold relations', IFPRI.
- Alderman, H., Chiappori, P.A., Haddad, L., Hoddinott, J. and Kanbur, R. (1995). 'Unitary versus collective models of the household: is it time to shift the burden of proof?', *World Bank Research Observer*, vol. 10 (1217), pp. 1–19.
- Almås, I., Armand, A. and Attanasio, O. (2016). 'Private money or household money?', AEA RCT Registry. Available at: <https://www.socialscienceregistry.org/trials/1063/> (last accessed: 18 February 2017).
- Angelucci, M. (2008). 'Love on the rocks: domestic violence and alcohol abuse in rural Mexico', *BE Journal of Economic Analysis & Policy*, vol. 8(1), pp. 1017–1031.
- Armand, A. and Carneiro, P. (2013). 'Impact evaluation of conditional cash transfers for secondary school attendance in Macedonia: results from the first follow-up', Ministry of Labour and Social Policy, Republic of Macedonia.
- Ashraf, N. (2009). 'Spousal control and intra-household decision making: an experimental study in the Philippines', *American Economic Review*, vol. 99(4), pp. 1245–77.
- Ashraf, N., Field, E. and Lee, J. (2014). 'Household bargaining and excess fertility: an experimental study in Zambia', *American Economic Review*, vol. 104(7), pp. 2210–37.
- Attanasio, O. and Lechene, V. (2002). 'Tests of income pooling in household decisions', *Review of Economic Dynamics*, vol. 5(4), pp. 720–748.
- Attanasio, O.P. and Lechene, V. (2014). 'Efficient responses to targeted cash transfers', *Journal of Political Economy*, vol. 122(1), pp. 178–222.
- Becker, G.S. (1991). *A Treatise on the Family*, Cambridge, MA: Harvard University Press.
- Blundell, R., Chiappori, P.A. and Meghir, C. (2005). 'Collective labor supply with children', *Journal of Political Economy*, vol. 113(6), pp. 1277–1306.
- Blundell, R., Chiappori, P.A., Magnac, T. and Meghir, C. (2007). 'Collective labour supply: heterogeneity and non-participation', *Review of Economic Studies*, vol. 74(2), pp. 417–445.
- Bobonis, G.J., González-Brenes, M., Castro, R. (2013). 'Public transfers and domestic violence: the roles of private information and spousal control', *American Economic Journal: Economic Policy*, vol. 5(1), pp. 179–205.
- Bourguignon, F., Browning, M. and Chiappori, P.A. (2009). 'Efficient intra-household allocations and distribution factors: implications and identification', *Review of Economic Studies*, vol. 76(2), pp. 503–28.
- Browning, M. and Chiappori, P.A. (1998). 'Efficient intra-household allocations: a general characterization and empirical tests', *Econometrica*, vol. 66(6), pp. 1241–78.
- Browning, M., Chiappori, P.A. and Weiss, Y. (2014). *Economics of the Family*, New York, NY: Cambridge University Press.
- Chiappori, P.A. (1992). 'Collective labor supply and welfare', *Journal of Political Economy*, vol. 100(3), pp. 437–467.
- Chiappori, P.A., Fortin, B. and Lacroix, G. (2002). 'Marriage market, divorce legislation, and household labor supply', *Journal of Political Economy*, vol. 110(1), pp. 37–72.
- Doepke, M. and Tertilt, M. (2014). 'Does female empowerment promote economic development?', National Bureau of Economic Research.
- Doss, C. (2006). 'The effects of intrahousehold property ownership on expenditure patterns in Ghana', *Journal of African Economies*, vol. 15(1), pp. 149–80.

- Dufló, E. (2003). 'Grandmothers and granddaughters: old-age pensions and intrahousehold allocation in South Africa', *World Bank Economic Review*, vol. 17(1), pp. 1–25.
- Dufló, E. (2012). 'Women empowerment and economic development', *Journal of Economic Literature*, vol. 50(4), pp. 1051–79.
- Fiszbein, A., Schady, N.R. and Ferreira, F.H. (2009). *Conditional Cash Transfers: Reducing Present and Future Poverty*, Washington, DC: World Bank Publications.
- Gartner, M. and Sandberg, A. (2015). 'Is there an omission effect in prosocial behavior?', Stockholm School of Economics.
- Gudevold, T.J. and Kjørholt, I.E.H. (2016). 'The impact of information on female empowerment in low-income households in Kenya: a pilot study', Master's thesis, Norwegian School of Economics.
- Hoddinott, J. and Haddad, L. (1995). 'Does female income share influence household expenditures? Evidence from Côte d'Ivoire', *Oxford Bulletin of Economics and Statistics*, vol. 57(1), pp. 77–96.
- Lundberg, S.J., Pollak, R.A. and Wales, T.J. (1997). 'Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit', *Journal of Human Resources*, vol. 32(3), pp. 463–80.
- Samuelson, P.A. (1956). 'Social indifference curves', *Quarterly Journal of Economics*, vol. 70(1), pp. 1–22.
- Schady, N.R. and Araujo, M. (2006). *Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador*, Washington, DC: World Bank Publications.
- Schumacher, M. (1984). 'Two-sample tests of Cramér-von Mises- and Kolmogorov-Smirnov-type for randomly censored data', *International Statistical Review*, vol. 52(3), pp. 263–81.
- Thomas, D. (1990). 'Intra-household resource allocation: an inferential approach', *Journal of Human Resources*, vol. 25(4), pp. 635–64.
- Voena, A. (2015). 'Yours, mine, and ours: do divorce laws affect the intertemporal behavior of married couples?', *American Economic Review*, vol. 105(8), pp. 2295–332.
- Ward-Batts, J. (2008). 'Out of the wallet and into the purse: using micro data to test income pooling', *Journal of Human Resources*, vol. 43(2), pp. 325–51.