DISCUSSION PAPER SERIES

DP16482

Feed the children

Laurens Cherchye, Pierre-André Chiappori, Bram De Rock, Charlotte Ringdal and Frederic Vermeulen

DEVELOPMENT ECONOMICS

LABOUR ECONOMICS

PUBLIC ECONOMICS

ISSN 0265-8003

Feed the children

Laurens Cherchye, Pierre-André Chiappori, Bram De Rock, Charlotte Ringdal and Frederic Vermeulen

Discussion Paper DP16482
Published 25 August 2021
Submitted 24 August 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Feed the children

Abstract

To understand the household decision-making process regarding food expenditures for children in poor households in Nairobi, we conduct an experiment with 424 married couples. In the experiment, the spouses (individually and jointly) allocated money between themselves and nutritious meals for one of their children. First, we find strong empirical support for individual rationality and cooperative behavior. Second, our results suggest that women do not have stronger preferences for children's meals than men. Third, the spouses' respective bargaining positions derived from consumption patterns strongly correlate with more traditional indicators. Finally, we document significant heterogeneity both between individuals and intra-household decision processes.

JEL Classification: C14, D13, C90

Keywords: collective model, Intra-household allocation, Experiment, Kenya, children

Laurens Cherchye - laurens.cherchye@kuleuven.be
KU Leuven and CEPR

Pierre-André Chiappori - pc2167@columbia.edu
Columbia University

Bram De Rock - bderock@ulb.ac.be
Université libre de Bruxelles, ECARES, KU Leuven, Department of Economics and CEPR

Charlotte Ringdal - charlotte.ringdal@cmi.no
Chr. Michelsen Institute (CMI), Norway

Frederic Vermeulen - frederic.vermeulen@kuleuven.be
University of Leuven and CEPR

Feed the children*

Laurens Cherchye† Pierre-André Chiappori‡ Bram De Rock§

Charlotte Ringdal¶ Frederic Vermeulen‖

August 18, 2021

Abstract

To understand the household decision-making process regarding food expenditures for children in poor households in Nairobi, we conduct an experiment with 424 married couples. In the experiment, the spouses (individually and jointly) allocated money between themselves and nutritious meals for one of their children. First, we find strong empirical support for individual rationality and cooperative behavior. Second, our results suggest that women do not have stronger preferences for children's meals than men. Third, the spouses' respective bargaining positions derived from consumption patterns strongly correlate with more traditional indicators. Finally, we document significant heterogeneity both between individuals and intra-household decision processes.

JEL Classification: C14, D13, C91, C92.

Keywords: collective model, intra-household allocation, experiment, Kenya, children.

*A pre-analysis plan was registered at the Evidence in Governance and Politics registry before we collected the data (https://osf.io/fj6yr). The project received IRB approval from Kenya Medical Research Institute (KEMRI) on November 6, 2017. *Acknowledgements:* We are grateful to Pauline Wanjeri and the team of experimenters at Busara, Kenya for assistance during the data collection.

†Department of Economics, University of Leuven (KU Leuven). E-mail: laurens.cherchye@kuleuven.be. Laurens Cherchye gratefully acknowledges the Fund for Scientific Research - Flanders (FWO) and the Research Fund KU Leuven for financial support.

‡Columbia University. E-mail: pc2167@columbia.edu.

§ECARES, Université Libre de Bruxelles and Department of Economics, University of Leuven (KU Leuven). E-mail: bram.de.rock@ulb.be. Bram De Rock gratefully acknowledges FWO and FNRS for their support.

¶Chr. Michelsen Institute (CMI), Norway. E-mail: charlotte.ringdal@cmi.no. Charlotte Ringdal gratefully acknowledges financial support from Sara van Dam Project Grant of the Royal Netherlands Academy of Arts and Sciences and from a European Research Council (grant ERC-StG-637915) and the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675.

‖Department of Economics, University of Leuven (KU Leuven). E-mail: frederic.vermeulen@kuleuven.be. Frederic Vermeulen gratefully acknowledges the Fund for Scientific Research - Flanders (FWO) and the Research Fund KU Leuven for financial support.

1 Introduction

Globally, almost 200 million children under 5 are undernourished and more than 340 million children suffer from hidden hunger (i.e., lack of vitamins and other essential micro-nutrients) (UNICEF, 2019). This problem is especially severe in Eastern and Southern Africa where more than 40% of children are undernourished and nearly 70% suffer from hidden hunger. For the poor living in urban slums, eating well has turned out to be even more challenging than for the poor in rural areas as access to nutritious and healthy food is limited. For instance, about 25% of a households' food expenditures are on street food, which is deemed less healthy (UNICEF, 2019). The example of Kenya is quite clear. More than 50% of the three-year-old children in urban slums are undernourished (UNICEF, 2012) - despite the fact that households in the slums of Nairobi spend almost 50% of their income on food (World Bank, 2016).

To combat the undernutrition of children, it is crucial to understand the mechanisms that generate this outcome. Quite obviously, some gaps remain in our knowledge on the spending patterns of households in slums, in particular with respect to food for children. Recent advances have stressed that household behavior depends on a variety of factors. While standard economic variables like income and market prices play an important role, adults in households may also have different preferences on how to spend the household's income - and particularly how much should be spent on children. In such cases, the intra-household bargaining process that determines who has the biggest say in the household decisions becomes an important channel that influences the spending pattern. For instance, it is often claimed that wives tend to care more about their children's health than their husbands. As a result, policy interventions that tend to shift decision power within the household in women's favor can be expected to improve children's welfare, including their nutrition.

Our objective in this paper is to disentangle the different determinants of poor households' spending patterns, with a special emphasis on the spending on food for children. This exercise faces clear difficulties. From a theoretical perspective, one needs a model allowing one to analyze the intra-household decision-making process in a context where adults may have diverging preferences on how much income should be spent on food for children - and where, as a result, shifts in the balance of power between them may have important consequences. This requirement rules out the traditional, unitary model, where a household is assumed to maximize well-behaved and constant preferences subject to some constraints. The two main candidates are therefore the collective approach (Chiappori, 1988, 1992), which assumes that household consumption is the Pareto efficient outcome of a bargaining process inside households (see for instance Blundell et al., 2005; Bourguignon et al., 2009; Browning and Chiappori, 1998; Cherchye et al., 2007, 2011b, 2012; Chiappori, 1988; Chiappori and Ekeland, 2009), and the non-cooperative approach, which studies the Nash equilibrium of a non-cooperative game involving private contributions to the public goods. Numerous empirical studies, focusing on both developed and developing countries, take the collective approach as a starting point (see

Baland and Ziparo, 2018; Donni and Molina, 2018, for a recent survey and critical discussion). Yet, several articles concentrate on non-cooperative models in many different settings (see for instance Bergstrom et al., 1986; Browning et al., 2009; Doepke and Tertilt, 2018; Gobbi, 2018; Konrad and Lommerud, 1995; Lewbel and Pendakur, 2019).

From an empirical perspective, the modeling of households' spending decisions, in particular the spending on food for children, faces a number of challenges. Households' consumption decisions are characterized by a lot of heterogeneity, both observed and unobserved. Moreover, individual shares in (food) consumption are almost never observed; typically, household surveys on spending behavior only record (household-level) aggregate expenditures on goods and services, in particular food.[1] An important advantage of the collective model is that it nevertheless allows one to recover individual shares of household resource under specific assumptions. In practice, different identification strategies have been developed to identify resource shares in this context; see Chiappori and Ekeland (2009), Cherchye et al. (2011b), and Dunbar et al. (2013), for three different general approaches (respectively a differentiable approach, a revealed preference approach, and an approach relying on specific parametric assumptions). In addition, identifying individual preferences, a prerequisite to obtain more insights into the different drivers that determine the spending pattern of poor households, requires sufficient price variation. Unfortunately, surveys with detailed household consumption information have most often a cross-sectional or short panel nature with no or little variations in relative prices.

To address these theoretical and empirical challenges, we set up a lab experiment that allows us to identify, household per household, the decision-makers' individual preferences, their individual resource shares and the share going to food for children, as determined by the household's intra-household bargaining process. The experiment was conducted with 424 monogamous married couples in Nairobi. It consisted of two main parts. In the first part of the experiment, each spouse made individual decisions where they chose between money for themselves and nutritious meals for one of their children. In the second part of the experiment, the spouses sat together and made the same decisions but now between money for the two of them, which is split up in a follow-up question, and nutritious meals for one of their children. In each part, individuals or couples were presented with 11 different choice sets. Each choice set consisted of between three and seven consumption bundles containing money for individual use and meals for one of their children. Across the 11 choice sets, the price of meals for children and the total budget available was varied in such a way that a substantial relative price variation is observed in the data. The experiment was explicitly designed to investigate (i) whether men and women have different preferences over the choice sets, (ii) how preferences are aggregated when decisions are made jointly, and (iii) whether individuals and couples satisfy the theoretical restrictions implied by rationality and/or Pareto efficiency.

This design provides an optimal framework to analyze the different channels that influ-

[1] As recent exceptions, see for instance Browning and Goertz (2012), Cherchye et al. (2012) and Bargain et al. (2021).

ence poor households' spending patterns with a main focus on food for children. Importantly, we will apply different and complementary methods to obtain new insights in this spending pattern. First, we derive a theoretical property that allows a direct horse race between the cooperative and the non-cooperative framework. Specifically, in a collective model the amounts parents *jointly* decide to spend on any public good (in our case, on children) sits *between* what each of them would spend, would (s)he be the unique decision maker. On the contrary, a non-cooperative decision process typically generates a joint amount that is *inferior to both* individually chosen amounts. This last prediction is clearly falsified in our data, suggesting that parents do not behave non-cooperatively.

Next, we apply both (nonparametric) revealed preference techniques to check the (household-specific) validity of the collective model and to identify its structural components, and a more standard (parametric) approach to obtain more detailed insights into the spouses' diverging preferences on how to spend the couple's resources, and the intra-household bargaining process that leads to the observed choices.

Our study offers some surprising insights. Firstly, when deciding alone, husbands allocate on average a higher budget share to meals for one of their children than their wives do, and thus a lower share to own money. When deciding jointly, spouses on average choose a budget share of meals for one of their children that is between the budget shares when the choice is made individually. This implies that the budget shares of own money drop substantially when spouses jointly decide compared to when they determine their own share individually. Secondly, our revealed preference analysis shows that the unitary model provides an adequate explanation of the individual choices for most of the individuals, while the collective model does a good job in explaining the joint choices. Finally, the parametric analysis shows that the observed intra-household allocation of expenditures to own money and meals for one of the children in the experiment is driven by a complex interplay between individual preferences that differ across spouses (with males' propensities to spend on children that are somewhat higher than females') and a bargaining process in which females have a somewhat better position. As will be discussed next, these results deviate from earlier results obtained in other empirical settings.

Our paper relates to several strands of the literature. Firstly, it relates to the literature on the role of women's bargaining power in household decision-making on the spending on children. These studies mostly focus on whether money in the hands of women have a different effect on household expenditure patterns than money in the hands of men, and specifically how it relates to spending on children. Earlier cross-sectional studies typically show that increasing women's bargaining power tends to boost spending on children (see for instance Doss, 2006; Hoddinott and Haddad, 1995; Lundberg et al., 1997; Thomas, 1990, 1993). Since the Oportunidades (former Progresa) program was launched in Mexico in 1997, conditional cash transfer programmes (CCT) have been studied extensively; studies generally find a positive effect of CCTs on the budget share spent on food (see for instance Attanasio and Lechene, 2014; Bobonis, 2009). A possible problem with these studies is that transfers are usually given to women only; it is there-

fore difficult to distinguish bargaining power influences from standard income effects (see for instance Angelucci and Attanasio, 2013; Attanasio and Lechene, 2002; Rubalcava et al., 2009). Recently, however, conditional and unconditional cash transfers that randomize the gender of the recipient have received more attention. While the conclusions are somewhat mixed (Akresh et al., 2016; Armand et al., 2020; Benhassine et al., 2015; Haushofer and Shapiro, 2016), Armand et al. (2020) find that in households with a female recipient, a higher budget share is spent on food and the quality of the diet is improved. We add to this literature by studying (i) whether men and women have different preferences when it comes to spending on food for children, and (ii) who, and under which conditions, has most say when decisions are made jointly. When it comes to the budget share spent on food or quality of diet, the limitation with most studies is that they only observe this at the household level and not at the individual level. In other words, one cannot directly observe whether the higher budget share (or the improved diet) is given to the children in the household or the adults. Our study makes it possible to look at spending on food for children directly, and relate this both to our new test of non-cooperative behavior and to the collective model.

Second, our study relates to the emerging literature studying household decision-making in the lab. For a review of the literature, see Munro (2018). Closest to our paper are Carlsson et al. (2012, 2013) and Ringdal and Sjursen (2021). Carlsson et al. (2012, 2013) study time and risk preferences when decisions are made individually and jointly. They find that men have more weight in the joint decision than women, and that the joint degree of risk aversion and discount factor are in-between the individual degrees of risk aversion and discount factors. Ringdal and Sjursen (2021) analyze the effect of increasing women's control over an endowment in the lab on investments in children's education. We differ from these studies in several ways. Firstly, we focus on the allocation to children and not time or risk preferences. Secondly, we have a within-subject design that allows us to look at both individual preferences and the aggregation of preferences in each couple. Moreover, we impose no restriction (beyond Pareto efficiency) on how the joint decision is made.

Finally, our study relates to the literature on testing whether household decisions are Pareto efficient. In developing countries, most of the evidence stems from agricultural surveys of farmer households in West Africa (Akresh, 2008; Duflo and Udry, 2004; Udry, 1996). These studies reject efficiency based on households not allocating resources efficiently across plots and spouses not insuring each other fully, although these conclusions have been challenged (see for instance Rangel and Thomas, 2019). On the consumption side, most of the data are from surveys in developed countries that are aggregated at the household level. These studies mostly reject the unitary model in favor of the collective model (see for instance Browning and Chiappori, 1998; Cherchye et al., 2009). Rangel and Thomas (2019) obtain a similar result is obtained for the developing country Burkina Faso. In the lab, Kaldor-Hicks efficiency has been tested by using standard games such as public good games (see Munro, 2018, for a review). Most of these studies reject Kaldor-Hicks efficiency, but as Munro (2018) notes, studies in

developed countries tend to show a higher degree of efficiency than studies in developing countries. To our knowledge, only one study has tested Pareto efficiency in the lab. Bateman and Munro (2005) consider individual and joint decisions between lotteries in the UK and reject Pareto efficiency of joint decisions. Building on the design of Bateman and Munro (2005), we are, to our knowledge, the first to test for Pareto efficiency in the lab in a developing country. Moreover, by using the task based on Harbaugh et al. (2001), we can use a revealed preference approach to test whether the collective model provides a good fit of the data, which is followed by the complementary estimation of a parametric model (allowing us to identify individual preferences and bargaining power measures) to understand household decision-making better.

The rest of the article is organized as follows. In Section 2, we describe the theoretical framework. Section 3 introduces the methods that we will rely on to analyze the experiment's results. Section 4 describes the experimental design and the sample. In Section 5, we give some descriptive statistics on the experimental results obtained. In Section 6, we conduct a structural analysis of the choice behavior in our experiment by means of both nonparametric revealed preference methods and more standard parametric methods. We conclude in Section 7. The online Appendix presents some additional empirical results.

2 Theoretical framework

In this section, we will discuss the theoretical framework that we will rely on in the experimental and empirical analysis. After introducing some notation, we will first present a simple test for non-cooperative decision making. Subsequently, we introduce the collective approach to joint decision making, which basically assumes Pareto efficiency of the decision outcomes.

2.1 The basic framework

2.1.1 Individual preferences

In the experiment, individuals (where the superscript f refers to the female, while the superscript m refers to the male) have to decide, independently or jointly, on the spending of exogenously determined amounts of money to two different commodities in different choice situations: the own consumption of both spouses (denoted by respectively q^f and q^m), and a number of meals for one of their children (Q). Each spouse is assumed to have own preferences, represented by a utility function of the form $u^i\left(q^i, Q\right)$ for $i = f, m$.[2] Throughout, we will

[2]In particular, individuals are assumed egoistic, in the sense that they do not care for their spouse's welfare. This assumption is mostly innocuous. Allowing for altruism (i.e., i's utility entering j's welfare) would not affect the qualitative conclusions, since we mostly assume efficiency; as it is well known, decisions that are efficient for altruistic agents are efficient for egoistic agents as well. We, however, assume away direct externalities, whereby a person's private consumption would directly enter the spouse's utility; this issue will be discussed (and tested) later on.

assume that spouses' utility functions are locally non-satiated, concave and non-decreasing in their arguments.

We define i's marginal willingness to pay for the public good as:

$$MWP^i\left(q^i,Q\right)=\frac{\partial u^i/\partial Q}{\partial u^i/\partial q^i},$$

and we assume that preferences are regular, in the following sense:

Assumption R (regularity): *$MWP^i\left(q^i,Q\right)$ is strictly increasing in q^i and strictly decreasing in Q.*

Assumption R implies in particular that both commodities are normal in the usual sense.

2.1.2 Individual decisions

We first consider the first stage of the experiment, during which individuals independently decide over the allocation between their own private consumption and expenditures on children. For simplicity, we consider a standard budget constraint:

$$q^i+PQ=x, \tag{1}$$

where the price of own private consumption is normalized to one, P is the price of the public good and x denotes the individual's total expenditures. Assuming an interior solution, first order conditions give:

$$MWP^i\left(q^i,Q\right)=P,\ i=f,m. \tag{2}$$

This condition, together with the budget constraint (1), determines i's optimal choice, which we denote $\left(q_D^i\left(P,x\right),Q_D^i\left(P,x\right)\right)$, $i=f,m$ (where the subscript D stands for 'dictator').[3]

2.2 The non-cooperative model: a simple test

A first task is to distinguish between the collective approach, based on individual rationality and Pareto efficiency, and the non-cooperative approach.

2.2.1 Non-cooperative models

In non-cooperative models, individuals, when deciding jointly, play a private contribution game by independently choosing their contribution to public good expenditures. The natural equilib-

[3]Corner solutions, where an individual allocates *all* available income to either own or the child's consumption, may also obtain; they are respectively characterized by inequality of the type $MWP^i\left(q^i,Q\right)\leq P$ or $MWP^i\left(q^i,Q\right)\geq P$.

rium concept is then (non-cooperative) Nash; that is, spouses each maximize their own utility, taking their partner's contribution as given.

In the most simple and static version (see for instance Bergstrom et al., 1986; Browning et al., 2009; Cherchye et al., 2011a), individuals are each endowed with their own income, out of which they fund their private contribution. Since this description does not exactly correspond to our experimental setting, and many real-life settings, where incomes are not individually provided in the joint decision phase, we use an extension proposed by Chiappori and Naido (2020), in which individuals behave as if they had previously shared their total income between them according to some (virtual) income distribution rule that is not observed by the econometrician.[4] In practice, thus, there exist $\tilde{x}^f$ and $\tilde{x}^m$, with $\tilde{x}^f + \tilde{x}^m = x$, such that agent i's program is:

$$\max_{q^i, Q^i} u^i \left(q^i, Q^f + Q^m \right) \text{ under } q^i + PQ^i = \tilde{x}^i.$$

Let $\left(q_N^f, q_N^m, Q_N \right)$ denote the Nash equilibrium allocation as a function of (P,x). Well-known results state that an equilibrium must be of either of two types. Let us first focus on an equilibrium of the first type. When the solution is interior, in the sense that both agents contribute to the public good, first order conditions are:

$$MWP^i \left(q_N^i, Q_N \right) = P, \; i = f, m. \tag{3}$$

These conditions, together with the *aggregate* budget constraint:

$$q_N^f + q_N^m + PQ_N = x,$$

fully determine the equilibrium. In particular, the latter only depends on total income x, not on the distribution rule $\left(\tilde{x}^f, \tilde{x}^m \right)$: this is the well-known income pooling result.

In the second type of Nash equilibrium only one agent contributes to the public good. Assume, for instance, that she is the only contributor (the alternative case is obviously symmetric); then $\tilde{x}^m = q_N^m$ and $\tilde{x}^f = x - \tilde{x}^m = q_N^f + PQ_N$. In particular, the regularity condition **R** implies that

$$Q_N(P,x) = Q_D^f \left(P, \tilde{x}^f \right) < Q_D^f(P,x).$$

In words, the non-cooperative expenditure on children, $Q_N(P,x)$, must be smaller than what she would choose if she were the only decision maker.

2.2.2 A simple test

When there is income pooling we obtain the following intuitive result.

Proposition 1. *In the non-cooperative framework and under the regularity assumption R, if*

[4] Note that the resulting model is, if anything, more general (hence harder to reject) than the standard one.

both spouses contribute to the public good, then:

$$Q_N(P,x) \leq \min\left(Q_D^f(P,x), Q_D^m(P,x)\right).$$

Proof. The various budget constraints give:

$$q_N^f + q_N^m + PQ_N = x = q_D^f + PQ_D^f = q_D^m + PQ_D^m.$$

Assume, now, that $Q_N(P,x) > Q_D^i(P,x)$ for at least one i. Then $q_N^i(P,x) < q_D^i(P,x)$ follows from the budget constraint. The regularity assumption then implies that:

$$MWP^i\left(q_N^i(P,x), Q_N(P,x)\right) < MWP^i\left(q_D^i(P,x), Q_D^i(P,x)\right).$$

Since both sides equal P by the first order conditions (2) and (3), we have a contradiction. □

In words, under non-cooperative provision of the public good, the couple chooses, at an interior solution, a level of children expenditures that is less than what any of the spouses would choose if (s)he were the unique decision maker, endowed with the same total income. This result translates, in a spectacular way, a standard conclusion: in the absence of cooperation, public goods are underprovided, because individuals fail to consider the benefits their contribution generates for their partner.

For the second type of Nash equilibrium, we obtained a similar result for the wife. It is in principle possible, though, that $Q_N(P,x)$ is larger than the *husband's* dictatorial choice $Q_D^m(P,x)$. That would be the case, for instance, if the husband cared much less about children than the wife, in the sense that, at the Nash equilibrium allocation (and despite the fact that she is the only contributor to current children expenditures), his marginal willingness to pay for the children was still inferior to the price P, so that each marginal dollar he receives would be exclusively devoted to his own private consumption. While such highly asymmetric situations are in principle possible, we consider them as very unlikely in our context, especially since spouses, when deciding alone, tend to devote similar amounts to the children (if anything, his contribution, when deciding alone, is on average larger than hers, as we shall see below).

We therefore expect that, should non-cooperation be the rule, the jointly provided level of children consumption would in most cases be less than what both he and she would choose as unique decision makers. This prediction can readily be tested on our data. The verdict is clear. On average, it is satisfied in less than 5% of the observations; and whatever the decision under consideration, the percentage of situations where it is *not* violated never exceeds 7.1%. We therefore conclude that, in the vast majority of cases, the non-cooperative model does not provide an adequate description of household behavior.[5]

[5]The previous results can readily be extended to altruistic individuals. The proof, as well as an illustration based on Cobb-Douglas preferences, is available from the authors on request.

2.3 The collective approach

From now on, we shall thus concentrate on the collective approach. That is, we assume that agents make their common decision based on a process that always generates Pareto efficient outcomes. Following the collective tradition, we are agnostic about the true nature of the process. It may be that individuals bargain over the final outcome, and that the process can be summarized by an axiomatic concept such as Nash bargaining, Kalai-Smorodinsky, or any other. Or it could be that agents play a non-cooperative bargaining game à la Rubinstein, the outcome of which is efficient under symmetric information; and other interpretations are possible as well. In any case, we simply make the (reduced-form) assumption that there exists some $\mu^f \in [0,1]$ such that the household's final decision solves the program:

$$\max \mu^f u^f\left(q^f, Q\right) + \left(1-\mu^f\right) u^m\left(q^m, Q\right)$$

under the budget constraint

$$q^f + q^m + PQ = x,$$

where x denotes the household's total expenditures. Importantly, the spouses' Pareto weights μ^f and $(1-\mu^f)$ may depend on the prices and total expenditure that apply to the specific choice situation at hand. This collective model with a public good is essentially the one characterized by Blundell et al. (2005). Importantly, this model implies that external effects of one spouse's private consumption on the other spouse's utility are ruled out. We will explicitly test whether this assumption is valid in the revealed preference analysis below.

3 Methodology

As indicated above, individual rationality and Pareto efficiency are the core assumptions underlying the collective model. To check the validity of these assumptions for our experimental context, we will apply revealed preference methods. These methods are nonparametric, in the sense that they do not rely on a specific functional form for preferences or demand. As such, they can be seen as robust. Moreover, they can be applied to a context with only a limited number of observations. The testable implications of the standard unitary model (where households are assumed to behave as a single decision-maker) have been derived in Afriat (1967) and Varian (1982). Cherchye et al. (2011b) derived testability and identification results for the collective model with both public and private consumption. This model will be the starting point of the current study, where the spending on children's food will act as a public good in the parents' preferences. In the following subsection, we will describe the collective model that applies to our experimental setting, and the corresponding revealed preference conditions.

Beyond the mere testing whether the implications of the theoretical model are satisfied, we also would like to identify the role of the parent's individual preferences in the determination of

the spending on food for children. In principle, one can rely on revealed preference methods to (set) identify these preferences (see for instance Blundell et al., 2015; Cherchye et al., 2011b; Varian, 1982). These methods are often cumbersome, though, and require a lot of price variation to obtain empirically meaningful bounds on individual indifference curves. Therefore, we will complement our revealed preference analysis with a collective consumption model that can be applied by standard econometric tools. This analysis will allow us to point identify preferences and the intra-household allocation process as summarized by the sharing rule, and the individuals' bargaining power. It will also make it easier to relate observable household characteristics to individuals' bargaining power in order to analyze the latter's determinants. Obviously, this analysis will require stronger (parametric) assumptions than those of the revealed preference analysis. The more standard empirical model, and the identification strategy that we will rely on, will be discussed in a second subsection.

3.1 Revealed preference analysis

In one part of the experiment, each spouse, alone and independently, is faced with a number of choice situations. In each choice situation, the spouse has to decide on the allocation of an exogenously determined amount of money to own consumption and meals for one of the couple's children, where commodities have exogenously determined prices. In another part of the experiment, couples have to make joint decisions on how to allocate exogenously determined amounts of money to both spouses' own consumption and meals for one of their children. The individual choice situations will be analyzed by the standard unitary model discussed in Section 2.1, while the couple's choices will be analyzed by the collective model introduced in Section 2.3. The expenditure on food for children are a public good, given that both spouses are supposed to care, possibly to various degrees, for their child. Analyzing the observed choices by means of this collective model will allow us to check the validity of the Pareto efficiency assumption (along with individual rationality of the spouses).

3.1.1 Individual decisions

In the experiment, each spouse is faced individually with different choice situations denoted by $t \in \{1,\ldots,T\}$. The price of the own consumption of the spouse is normalized to one, while the price of one meal, denoted by P_t, varies across each choice situation. The total expenditure in situation t is denoted by x_t^i, where $x_t^i = q_t^i + P_t Q_t$ and $i = f,m$. For each spouse i, we thus have the data set:

$$S^i = \{(P_t; q_t^i, Q_t), t = 1,\ldots,T\}.$$

Varian (1982), based on Afriat (1967), showed that there exists a utility function that rationalizes the data set if and only if S^i satisfies the following conditions, suitably adapted to our

specific setting.

Definition 1 (**Afriat conditions for the unitary model**). *Let* $S^i = \{(P_t; q_t^i, Q_t), t = 1, \ldots, T\}$ *be a data set. There exists a utility function* u^i *that provides a* unitary rationalization *of* S^i*, if there exist* u_t^i, $\lambda_t^i \in \mathbb{R}_{++}$ *that satisfy, for all choice situations r and t, the* Afriat conditions*:*

$$u_r^i - u_t^i \leq \lambda_t^i[(q_r^i + P_t Q_r) - (q_t^i + P_t Q_t)].$$

A revealed preference test of the unitary model checks whether the given data set S^i satisfies the Afriat conditions. Essentially, this requires checking feasibility of a set of constraints that are linear in continuous unknowns, which can be done through standard linear programming techniques. If the data set satisfies the Afriat conditions, then there exists a rationalizing utility function (and vice versa). Varian (1982) has shown that the Afriat conditions are equivalent to the Generalized Axiom of Revealed Preference (GARP), in the following sense: a data set S^i satisfies the Afriat conditions if and only if it satisfies GARP. Given this equivalence, we will interchangeably refer to the Afriat conditions and GARP when discussing our empirical results in Section 6.

3.1.2 Joint decisions

In the joint choice part of the experiment, the spouses in a couple have to jointly decide on the spending of exogenously determined amounts of money. Using the notation introduced above, this results in a total expenditure $x_t = q_t^f + q_t^m + P_t Q_t$ and the data set of a couple is:

$$S = \{(P_t; q_t^f, q_t^m, Q_t), t = 1, \ldots, T\}.$$

Under the Pareto efficiency assumption, the revealed preference conditions of the collective model in Section 2.3 boil down to the existence of two spouse-specific datasets $S^f = \{(P_t^f; q_t^f, Q_t), t = 1, \ldots, T\}$ and $S^m = \{(P_t^m; q_t^m, Q_t), t = 1, \ldots, T\}$, where $P_t^f + P_t^m = P_t$ (for all t), that simultaneously satisfy the Afriat conditions for the unitary model. The spouse-specific prices for the public good P_t^f and P_t^m can be interpreted as Lindahl prices and are not observed. This implies that the Afriat conditions are nonlinear in the unknowns $u_t^f, \lambda_t^f, P_t^f, u_t^m, \lambda_t^m$ and P_t^m. The unobserved prices P_t^f and P_t^m also imply that the sharing rule, which describes how the couple's total expenditure x_t is allocated to both spouses' shares x_t^f and x_t^m, remains unobserved. More specifically, both spouses' shares, defined by respectively $x_t^f = q_t^f + P_t^f Q$ and $x_t^m = q_t^m + P_t^m Q$, where $x_t^f + x_t^m = x_t$ cannot be directly observed by the choices made by the couple.

Cherchye et al. (2011b) have shown that one can check the validity of the above collective model and set identify the sharing rule and the spouse-specific prices of the public good by means of an integer programming (IP) characterization.[6] Given our purpose, we restrict atten-

[6]This integer programming formulation builds on the equivalence between the Afriat conditions and GARP

tion to the testing of the collective model under study. The characterization makes use of binary variables $b^i_{tr} \in \{0,1\}$ $(i = f,m)$, where $b^i_{tr} = 1$ implies that the bundle (q^i_t, Q_t) is "revealed preferred" to the bundle (q^i_r, Q_r) by spouse i (meaning that $u^i(q^i_t, Q_t) \geq u^i(q^i_r, Q_r)$). They showed that spouse-specific utility functions that rationalize the couples' data set exist if and only if S satisfies the following conditions.

Definition 2 (IP conditions for the collective model). *Let $S = \{(P_t; q^f_t, q^m_t, Q_t), t = 1, \ldots, T\}$ be a data set. There exist utility functions u^f and u^m that provide a* collective rationalization *of S if and only if there exist $P^f_t, P^m_t, x^f_t, x^m_t \in \mathbb{R}_+$ and $b^f_{rt}, b^m_{rt} \in \{0,1\}$ that satisfy, for all choice situations r, t and u, the* IP conditions*:*

$$
\begin{aligned}
&\textit{(i) } P^f_t + P^m_t = P_t; \\
&\textit{(ii) } x^f_t = q^f_t + P^f_t Q_t \textit{ and } x^m_t = q^m_t + P^m_t Q_t; \\
&\textit{(iii) } x^f_r - (q^f_t + P^f_r Q_t) < x_r b^f_{rt} \textit{ and } x^m_r - (q^m_t + P^m_r Q_t) < x_r b^m_{rt}; \\
&\textit{(iv) } b^f_{ru} + b^f_{ut} \leq 1 + b^f_{rt} \textit{ and } b^m_{ru} + b^m_{ut} \leq 1 + b^m_{rt}; \\
&\textit{(v) } x^f_t - (q^f_r + P^f_t Q_r) \leq x_t \left(1 - b^f_{rt}\right) \textit{ and } x^m_t - (q^m_r + P^m_t Q_r) \leq x_t \left(1 - b^m_{rt}\right).
\end{aligned}
$$

A revealed preference test of the collective model verifies whether the data set S satisfies the IP conditions. In this case, this requires checking feasibility of set of linear constraints that are characterized by binary integer variables, which can be done through integer programming methods.

We assumed above that the spouses' utility functions are the same in the individual and joint choice parts of the experiment. This implies that a spouse's utility function does not depend on the other spouse's own consumption. We can check the validity of this assumption by simultaneously applying the IP conditions in Definition 2 to the spouses' data sets gathered in the individual and joint choice experiments.[7] This is what we will do in our following empirical application.

Finally, the testable conditions for unitary and collective rationalization in Definitions 1 and 2 are "sharp" in nature: either a data set is exactly consistent with the conditions or it is not. When these sharp revealed preference tests result in a rejection, it is possible to quantify how close the observed choice behavior is to (exact) rationalizability. In this study, we make use of the Afriat index (after Afriat, 1973), which measures the fraction by which observed expenditures must be decreased to make the data set consistent with the model under study. The index is a number between 0 and 1. The smaller the fraction, the further the observed choices are from those that result from (exact) rationality. This Afriat index is often interpreted as a nonparametric goodness-of-fit measure for the consumption model that is evaluated (see

that we discussed above. We refer to Cherchye et al. (2011b) for more details.

[7] In doing so, we use that a decision of spouse i in the individual choice experiment can be represented as a joint decision with private consumption and the price of the public good fully assigned to i (i.e., $P^i = P$ and $q^j = P^j = 0$ for $j \neq i$).

also Choi et al., 2014, for more details).[8]

3.2 Standard demand analysis

The revealed preference analysis of each couple's individual and joint data sets will result in (i) knowing whether the couple's choice behavior is consistent with the collective model that assumes individual rationality and Pareto efficiency, and (ii) knowing whether the independence assumption we made on the individual utility functions is satisfied.

This revealed preference analysis is complemented by a more standard collective demand analysis, which will be parametric in nature. It can be interpreted as an attempt to point identify the spouses' individual preferences and the sharing rule, under the comforting assumption that the revealed preference analysis demonstrated the validity of the collective model. The more standard demand analysis will allow us to analyze in more detail how spouses decide on the allocation of the household's resources to, among others, meals for one of their children.

3.2.1 Individual decisions

In what follows, we will assume that each spouse's choices in the different experimental choice situations can be described by the Almost Ideal Demand System of Deaton and Muellbauer (1980). This model is relatively flexible in price and income effects, and yet sufficiently parsimonious to be applicable to our experimental setting that contains only 11 choice situations per spouse.

Let w_q^i and w_Q^i be the budget shares of own consumption and meals for one of the children of spouse i, respectively. The demand model, with the theoretical restrictions of adding-up, homogeneity and symmetry imposed, equals (for all t):

$$\begin{aligned} w_{q;t}^i &= \alpha^i + \beta^i \ln\left(\frac{x_t^i}{a^i(P_t)}\right) + \gamma^i \ln P_t, \\ w_{Q;t}^i &= 1 - \alpha^i - \beta^i \ln\left(\frac{x_t^i}{a^i(P_t)}\right) - \gamma^i \ln P_t, \\ \text{where } \ln a^i(P_t) &= \alpha_0^i + (1-\alpha^i)\ln P_t - \frac{1}{2}\gamma^i \ln^2 P_t. \end{aligned} \tag{4}$$

To analyze the spouses' preferences with respect to the allocation of resources to own consumption and meals for one of their children in a detailed way, we will make use of elasticities. Income elasticities are equal to:

$$\varepsilon_{q;t}^i = 1 + \frac{\beta^i}{w_{q;t}^i} \text{ and } \varepsilon_{Q;t}^i = 1 - \frac{\beta^i}{w_{Q;t}^i}.$$

[8] We refer to Choi et al. (2014) for a formal definition of the Afriat index of the unitary model. For the collective model, we use a common Afriat index for the IP conditions of the two spouses; see Bruyneel et al. (2012) for a similar application.

The uncompensated own price elasticities of spouse i, while accounting of the normalized price of own consumption, equal:

$$\varepsilon^i_{qq;t} = \frac{-\gamma^i - \beta^i(\alpha^i + \gamma^i \ln P_t)}{w^i_{q;t}} - 1 \text{ and } \varepsilon^i_{QQ;t} = \frac{-\gamma^i + \beta^i(1 - \alpha^i - \gamma^i \ln P_t)}{w^i_{Q;t}} - 1.$$

Finally, the compensated own price elasticities are equal to:

$$\overline{\varepsilon}^i_{qq;t} = \varepsilon^i_{qq;t} + \varepsilon^i_{q;t} w^i_{q;t} \text{ and } \overline{\varepsilon}^i_{QQ;t} = \varepsilon^i_{QQ;t} + \varepsilon^i_{Q;t} w^i_{Q;t} \ ,$$

where the latter must be negative given the negative-semidefiniteness of the Slutsky matrix.

The indirect utility function of spouse i that gives rise to the above demand equations equals:

$$\ln u^i_t = \frac{\ln x^i_t - \ln a^i(P_t)}{P_t^{-\beta^i}},$$

whereas spouse i's cost or expenditure function equals:

$$x^i_t = \left(u^i_t\right)^{\left(P_t^{-\beta^i}\right)} a^i(P_t).$$

Given adding-up, the parameters of the above demand system can be estimated by means of only one of the two demand functions. To allow for optimization error, we add an additive error term to the budget share of own consumption and estimate its unknown parameters by means of Blundell and Robin's (1999) iterated linear least squares estimator. Importantly, the application of the model is individual specific, which implies that we do not make any homogeneity assumptions across spouses in couples or across couples, which guarantees maximal heterogeneity in choice behavior. Of course, the exogenously set price and income variation, which is quite substantial, helps to obtain sufficiently precise estimates.

3.2.2 Joint decisions

In our following empirical analysis, we will assume that spouses' individual preferences remain constant in the individual and joint choices. This implies that the demand system that describes a spouse's choices in the joint decision equals the one of equation (4), where the market price P_t of meals for children, which act as a public good, is replaced by the spouse-specific Lindahl price P^i_t:

$$
\begin{aligned}
w^i_{q;t} &= \alpha^i + \beta^i \ln\left(\frac{x^i_t}{a^i(P^i_t)}\right) + \gamma^i \ln P^i_t \qquad (5)\\
w^i_{Q;t} &= 1 - \alpha^i - \beta^i \ln\left(\frac{x^i_t}{a^i(P^i_t)}\right) - \gamma^i \ln P^i_t \\
\text{where } \ln a^i(P^i_t) &= \alpha^i_0 + (1-\alpha^i)\ln P^i_t - \frac{1}{2}\gamma^i \ln^2 P^i_t.
\end{aligned}
$$

Recall that the share of spouse i in the couple's total expenditure equals $x^i_t = q^i_t + P^i_t Q$. Let us, without any loss of generality, parameterize the spouses' Lindahl price as $P^f_t = \theta^f_t P_t$ and $P^m_t = \left(1-\theta^f_t\right) P_t$, where $\theta^f_t \in [0,1]$. Note that the spouse's budget shares on the public good are equal to:

$$
w^f_{Q;t} = \frac{\theta^f_t P_t Q_t}{q^f + \theta^f_t P_t Q_t} \text{ and } w^m_{Q;t} = \frac{\left(1-\theta^f_t\right) P_t Q_t}{q^m + \left(1-\theta^f_t\right) P_t Q_t}. \qquad (6)
$$

Under the given assumption of equality of preferences in the individual and joint choices, the only unknown parameter that remains in each choice situation t is θ^f_t. In practice, we estimate the unknown parameters as follows. We first substitute the female's and male's budget shares of the public good in the left-hand side of equation (5) by the corresponding functions in (6) to write the couple's chosen number of meals in a choice situation t as functions g^i $(i = f,m)$ of the already identified spouse-specific preference parameters, the share in the Lindahl price, the own consumption and the price of the public good:

$$
Q_t = g^f(\theta^f_t; q^f, P_t, \alpha^f_0, \alpha^f, \beta^f, \gamma^f) \text{ and } Q_t = g^m(\theta^f_t; q^m, P_t, \alpha^m_0, \alpha^m, \beta^m, \gamma^m). \qquad (7)
$$

Given Pareto efficiency, both functions in (7) should obtain the same function value in the optimum. Of course, optimization error may result in some deviation. Therefore, we estimate the unknown parameters θ^f_t by minimizing the distance between the functions $g^f(.)$ and $g^m(.)$ in each choice situation. We impose that θ^f_t must be a scalar between 0 and 1 through the function $\theta^f_t = \frac{\exp(\widetilde{\theta}^f_t)}{1+\exp(\widetilde{\theta}^f_t)}$, where $\widetilde{\theta}^f_t$ is the unknown parameter to be estimated in practice.

Once the unknown parameters θ^f_t are estimated, we have all the information necessary to conduct a thorough demand and welfare analysis: individual preferences (which are represented by indirect utility and cost functions), the sharing rule, and the Lindahl prices. This will allow us to analyze how moving from individual to joint choices affects the number of meals for children, and which factors (female versus male preferences and bargaining power) determine the observed changes.

4 Experimental design

The experiment was conducted at Busara Center for Behavioral Economics in Nairobi, Kenya, from December 2017 to May 2018. It was constructed by means of z-tree (Fischbacher, 2007). In this section we first describe the main experimental task and the survey conducted after the experiment. We then provide some descriptive statistics about the sample. The next section will discuss the main experimental results on the intra-household allocation of resources.

4.1 Public good task

Our design is similar to Harbaugh et al. (2001). The participants were faced with 11 different choice sets consisting of between three and seven consumption bundles. Each bundle consisted of an amount of money and a number of weeks with meals for one of their children. The choice sets are described in Table 1. As can be seen, the (implicitly used) total budget varies from Ksh 200 (approx US$ 1.97) to Ksh 900 (approx US$ 8.87), and the (implicitly used) prices of one week of meals varies from Ksh 33.33 to Ksh 300.

In the first part of the task, the participants made individual choices between money for themselves and meals for their child. In the second part, the participants made the choices together with their spouse. We randomized whether the couples used the wife's or the husband's computer for the joint choices. Furthermore, in the joint choice situations we asked the spouses to decide on the division of the money between the two of them. We randomized whether the distribution of money was made before or after the public good decisions.

Before the participants started with the individual part of the public good task, they were given the name of the child chosen to receive the meals. Among the couple's children between 6 years and 14 years we randomly choose one of them. For each week of meals, the child got a nutritious meal three times. Each meal consisted of meat (chicken, goat, beef or fish), vegetables and starch (ugali, potatoes, rice, or cooking banana). In the individual part, each participant made choices between money for themselves and weeks of meals for their child. In the joint decision, the husband and the wife made the same 11 decisions together. After the public good task is completed, the spouses return to their individual computer. To ensure that the participants made a careful decision, they had 30 seconds to think about the decision before making their choice. Before both the individual and joint decisions, the participants were asked a set of control questions to ensure that they understood the task.

In total, each couple made 33 decisions of which one decision was randomly chosen for payment. To ensure that both spouses could choose their preferred consumption bundle in the individual decision, we provided plausible deniability through a luck of the draw mechanism. There was a 10% probability of getting an outcome where the participants' decisions were irrelevant. In this case the computer would choose one of the possible options the participants

could have chosen for payment.[9] We also ensured that spouses sat sufficiently apart so that they could not communicate or see their partners' decisions.

Table 1: Decisions in the public good task

	Option 1	**Option 2**	**Option 3**	**Option 4**	**Option 5**	**Option 6**	**Option 7**
Budget 1	(200,0)	(100, 3)	(0,6)				
Budget 2	(300,0)	(200, 2)	(100, 4)	(0,6)			
Budget 3	(300,0)	(200, 3)	(100, 6)	(0,9)			
Budget 4	(400,0)	(300, 1)	(200, 2)	(100,3)	(0,4)		
Budget 5	(400,0)	(300, 2)	(200, 4)	(100,6)	(0,8)		
Budget 6	(500,0)	(400, 1)	(300, 2)	(200,3)	(100,4)	(0,5)	
Budget 7	(600,0)	(300, 1)	(0, 2)				
Budget 8	(600,0)	(400, 1)	(200,2)	(0,3)			
Budget 9	(600,0)	(500, 1)	(400,2)	(300,3)	(200,4)	(100,5)	(0,6)
Budget 10	(800,0)	(600, 1)	(400,2)	(200,3)	(0,4)		
Budget 11	(900,0)	(600, 1)	(300,2)	(0,3)			

The table displays the different bundles the participants could choose between in each decision. The first number refers to payments in cash and the second number refers to the number of weeks with meals the child would receive.

4.2 Survey

In the survey, we elicited time preferences, asked questions about income, consumption, time-use, and family network, and did a short version of the Big-5 personality test.

Time preferences were elicited (hypothetically) by asking participants to choose between Ksh 1000 today and Ksh 1250 after 1 month (see Ashraf et al., 2006). If the participants choose the sooner option, they were faced with a second question where they had to choose between Ksh 1000 today and Ksh 1500 after 1 month. If the participants still choose the sooner option, they were asked how much they would have to be given in order to wait. We also asked these questions with 6 months versus 7 months delay.

We implemented a detailed consumption module where participants answered how much the household had spent on 13 food objects in the past week, how much the participant individually had spent on items such as airtime, transport, clothing, electricity, and water in the past month, and how much the participant individually had spent on long-term goods such as religious expenses, school fees, medical expenses and household durables in the past year. The time-use module included time spent on 10 different activities, including work, transport, household chores, helping family or non-family members, activities with children and church attendance. For family network, we asked the participants about how large their family was (living parents, siblings, aunts, uncles, cousins, nieces and nephews), how many family members lived in Nairobi and how many lived in the same area as the participants. Finally, we used a 10-item list for the Big-5 personality test (see Rammstedt and John, 2007).

[9] See Ashraf (2009) for a similar procedure.

When recruited, the spouse answered a set of questions about decision-making in the household.[10] These questions asked who in the household had the final say when it comes to the wife's health care, large household expenditures, household purchases for daily needs, visiting family, daily meals and taking a child to the doctor.

4.3 Sample

The sample consists of 424 couples from Kangemi and Kibera in Nairobi.[11] Both Kangemi and Kibera are informal settlements in Nairobi, Kenya, situated 12 and 6 kilometers, respectively from Nairobi city center and 8 and 1 kilometers, respectively, from the experimental site. Estimates of the population ranges from 20,000 to 100,000 in Kangemi (Oyunga-Ogubi et al., 2009; Wikipedia, 2015) and is around 250,000 in Kibera (Kibera.org.uk, 2019). The population consists of different tribes, ethnicities and cultures. A large fraction of the population in both settlements live in extreme poverty with a majority not having access to neither electricity nor piped water (Kibera.org.uk, 2019; Overseers' Education Centre, 2012).

A requirement to participate in the experiment was that the couples were married or living as husband and wife.[12] Furthermore, we required that neither of the spouses had participated in an experiment previously. In addition, we required that they had a child between the age of 6 years and 14 year who were in Kangemi/Kibera during the day (does not need to be in school) and had the possibility to attend the meals. We set the lower limit at 6 years old to make sure that a guardian did not need to follow the child to the venue where the meals were served. Participants from Kangemi were paid a Ksh 300 show-up fee plus a bonus of Ksh 50 if they arrived on time, while participants from Kibera were paid a Ksh 200 show-up fee plus a bonus.[13] On top of the show-up fee, each participant was paid Ksh 176 and each child received 1.4 weeks of meals. Of the randomly chosen children, 53% are girls and they are, on average, 9 years old.

5 Descriptive analysis

Let us first focus on the chosen allocations by the spouses in the individual and the joint decision situations of the experiment. Table 3 provides the average shares allocated to the child (Panel A), the wife (Panel B) and the husband (Panel C). First, when making an individual decision,

[10] We do not know whether the question was asked to the husband or the wife.

[11] Originally, the study was set to Kangemi, but because of problems with recruitment, the last 21% (90 couples) were recruited from Kibera instead. We did a sensitivity analysis of our main results in this respect and concluded this had limited impact on our conclusions. The results can be obtained upon request.

[12] These couples are referred to as "come-we-stay" marriages. They are very common, and the couple (and the community) refers to them as husband and wife.

[13] The show-up fee varies because of the distance to the experimental site. At the end of the experiment, the participants were paid Ksh 200 to cover transport, and given a certificate for the meals they have earned for their child. The remaining money payments to be made were sent through M-Pesa. M-Pesa is an SMS-based money-transferring system that can be used for most financial transactions in Kenya.

Table 2: Descriptive statistics

	Women		Men		
	Mean	Observations	Mean	Observations	p-value
Demographics					
Age	33.86	422	37.27	423	0.000
From Kangemi	0.79	424	0.79	424	
Children	2.86	413	2.74	413	0.423
Years married	11.68	413	11.88	413	0.696
Age married	20.92	413	25.19	413	0.000
Education					
No education	0.01	398	0.01	409	0.973
Primary school	0.55	398	0.44	409	0.003
Secondary school	0.40	398	0.48	409	0.019
University	0.05	398	0.07	409	0.155
Employment and income					
Self employed	0.29	413	0.26	413	0.392
Work for family member	0.04	413	0.07	413	0.130
Work for someone outside family	0.11	413	0.27	413	0.000
Chose not to work	0.01	413	0.03	413	0.057
Could not find work	0.53	413	0.35	413	0.000
Did not try to find work	0.02	413	0.02	413	0.634
Daily income (Ksh)	381	413	782	413	0.000
Weekly income (Ksh)	1539	413	2938	413	0.000
Time preferences					
Discount factor	18.24	399	19.14	401	0.310
Decision-making power					
# decisions wife is involved in	4.46	424	4.46	424	

Note: p-value is from a two-sample two-sided t-test between men and women.

husbands allocate a significantly higher share to the child's meals than their wives: on average, husbands allocate 37% of the experimental budget to meals for one of their children, whereas their spouses opt for a share of 32% on average (the p-value of a two-sided test equals 0.00). The same observation applies to the median of the share of food for one of the children, which is substantially higher for husbands than for wives. Second, the average share allocated to the child in the joint decision (35%) is between the above reported shares allocated to the child in the spouses' individual decision situations. This is a very interesting observation, since it implies that both spouses' shares of own money for the joint decisions are substantially lower than those associated with the individual decisions: for wives, the average share drops from 68% to 36%, while for their husbands, it drops from 63% to 30%. Finally, these figures also imply that females are allocated a significantly higher share of own money than males in the joint decision situations on average (36% versus 30%, with a p-value of 0.00 of a two-sided test).

Table 3: Allocation decision

	Panel A: Allocation to child				
	Mean	Standard deviation	1st Quartile	Median	3rd Quartile
Individual decision, wife	0.32	0.34	0.00	0.25	0.50
Individual decision, husband	0.37	0.34	0.00	0.33	0.67
Joint decision	0.35	0.33	0.00	0.33	0.50
	Panel B: Allocation to wife				
	Mean	Standard deviation	1st Quartile	Median	3rd Quartile
Individual decision, wife	0.68	0.34	0.50	0.75	1.00
Joint decision	0.36	0.25	0.17	0.33	0.50
	Panel C: Allocation to husband				
	Mean	Standard deviation	1st Quartile	Median	3rd Quartile
Individual decision, husband	0.63	0.34	0.33	0.67	1.00
Joint decision	0.30	0.22	0.13	0.29	0.48

Note: The table displays descriptive statistics on the allocation to child (Panel A), the wife (Panel B), and the husband (Panel C) including the mean, standard deviation, first quartile, median and third quartile. Shares allocated to the husband and the wife in the joint decision are calculated in two steps. First, for each decision we calculate the share allocated to either spouse. Then, for the amount allocated to either spouse (between Ksh 100 and Ksh 900), we use the decision on how to allocate the money between spouses to calculate the shares.

These results are in stark contrast to what we would expect based on other studies focusing on the intra-household allocation of resources. They are reinforced when we have a more detailed look at the average shares in the different choice situations. In all choice situations described in Table 4, the husband's allocation to the child's meals is higher than the wife's on average. Moreover, in all joint decision situations, husbands have on average a lower share of own money than their wives. These striking results are all but evidence in favor of the hypothesis that women have a stronger preference for meals for their child than men do, or that wives have a worse bargaining position than their husbands.

The above tables show a lot of heterogeneity across all decisions (see Figure A.1 in the

Table 4: Allocation decision

	Individual, wife		Individual, husband		Joint		
	Child	Wife	Child	Husband	Child	Wife	Husband
Decision 1	0.35	0.65	0.37	0.63	0.39	0.33	0.28
Decision 2	0.34	0.66	0.35	0.65	0.34	0.35	0.31
Decision 3	0.32	0.68	0.37	0.63	0.34	0.35	0.31
Decision 4	0.31	0.69	0.38	0.62	0.34	0.36	0.30
Decision 5	0.30	0.70	0.36	0.64	0.32	0.37	0.31
Decision 6	0.30	0.70	0.38	0.62	0.32	0.45	0.36
Decision 7	0.32	0.68	0.37	0.63	0.35	0.36	0.30
Decision 8	0.33	0.67	0.39	0.61	0.38	0.34	0.28
Decision 9	0.30	0.70	0.36	0.64	0.32	0.38	0.30
Decision 10	0.31	0.69	0.39	0.61	0.34	0.36	0.30
Decision 11	0.33	0.67	0.40	0.60	0.35	0.36	0.29

Note: The table displays the mean allocation to the child, the wife, and the husband when the wife makes an individual decision (Columns (1)-(2)), when the husband makes an individual decision (Columns (3)-(4)) and when the decision is made jointly (Columns (5)-(8). Shares allocated to the husband and the wife in the joint decision are calculated in two steps. First, for each decision we calculate the share allocated to either spouse. Then, for the amount allocated to either spouse (between Ksh 100 and Ksh 900), we use the decision on how to allocate the money between spouses to calculate the shares.

online Appendix for a further illustration). The above averages, for example, hide the fact that the experimental budget is in some cases almost entirely allocated to one of the spouses in a couple, while in other households almost the whole budget goes to meals for one of the children. For instance, in 2.8% (1.7%) of the couples, the share of own money for the wife (husband) is, on average, at least 90% over the different choice situations. In 6.9% of the households, the share of meals for the child exceeds 90% on average over the different choice situations. All this shows that it is indeed very important to account of the heterogeneity across individuals' preferences in analyses such as ours, as we indicated in the Introduction.

To further analyze this, Table 5 shows that the difference between the husband's and the wife's individual decisions remains robust when adding control variables. Specifically, for our most general specification, changing the decision-maker from the husband to the wife decreases the share allocated to the child by 5.8 percentage points ($p < 0.001$). We also find that the share allocated to the child negatively correlates with the number of children ($p = 0.057$) and the husband's discount factor ($p = 0.069$), and positively correlates with whether the couple lives in Kangemi ($p < 0.001$) and the wife's age ($p = 0.029$).

Overall, men allocate more to children than women in 31.6% of the couples while women allocate more to the child in 20.6% of the couples. Table A.1 in the online Appendix shows that the fraction of couples where the husband allocates more to the child varies between 22.2% in decision 1 to 37.1% in Decision 6, and the fraction of couples where the wife allocates more to the child varies from 17.7% in Decision 7 to 24.8% in Decision 9. Despite this heterogeneity,

Table 5: Individual allocation decisions

	(1)	(2)	(3)	(4)
Wife's decision	-0.055	-0.055	-0.053	-0.058
	(0.014)	(0.014)	(0.014)	(0.015)
Price of food, 100s		0.008	0.007	0.006
		(0.005)	(0.005)	(0.005)
Real income in decision, 100s		-0.002	-0.001	-0.001
		(0.002)	(0.002)	(0.002)
Years married			0.001	-0.000
			(0.002)	(0.002)
Number of children			-0.006	-0.007
			(0.003)	(0.003)
Living in Kangemi			0.204	0.197
			(0.015)	(0.026)
Age chosen child			0.001	-0.001
			(0.005)	(0.006)
Chosen child is female			0.013	0.005
			(0.028)	(0.030)
Weekly income, W				-0.004
				(0.005)
Weekly income, H				-0.003
				(0.004)
Age, H				-0.001
				(0.002)
Age, W				0.004
				(0.002)
Secondary school or more, H				-0.009
				(0.026)
Secondary school or more, W				-0.030
				(0.021)
Unemployed, H				0.032
				(0.025)
Unemployed, W				0.003
				(0.021)
Discount factor, H				-0.002
				(0.001)
Discount factor, W				-0.001
				(0.001)
Observations	9306	9306	9064	7986
R^2	0.263	0.263	0.267	0.289

Note: The table displays OLS regressions with fixed effects at the session level. Dependent variable: share allocated to child. Standard errors clustered at the session level in parenthesis.

for all decisions, it is more likely that the husband allocates more to the child than that the wife allocates more.

Next, Table 6 displays the predictors of the wife allocating more to the child (column (1)), and the husband allocating more to the child (column (2)). First, we see that the likelihood that the wife allocates more to the child than the husband increases with the income in the experiment ($p = 0.003$), years married ($p = 0.026$), age ratio (defined as age wife divided by age husband) ($p < 0.001$), and whether the husband is unemployed while the wife is employed ($p = 0.007$). Further, the likelihood decreases with whether they live in Kangemi ($p = 0.018$), the price of food in the experiment ($p = 0.001$), whether the chosen child is a girl ($p = 0.007$) and the wife's discount factor ($p < 0.001$). Second, the likelihood that the husband allocates more to the child than the wife increases with the real income in the experiment ($p < 0.001$), the age of the chosen child ($p < 0.001$), whether the chosen child is female ($p = 0.013$) and decreases with the price of food ($p < 0.001$) and the income ratio ($p < 0.001$). The correlations with the age and income ratios might indicate that couples where the effect of changing the decision-maker from the husband to the wife is driven by bargaining power and not preferences. In the next section, we will investigate this further through the lens of our structural model, hereby fully respecting the heterogeneity in choice behavior across households.

6 Structural analysis of the experimental results

In this section, we will check whether the experimental results can be rationalized by the structural models discussed in Sections 2 and 3. We start by assessing the empirical performance of the unitary model for the individual and the joint decisions and, subsequently, of the collective model for the joint decisions. By using revealed preference methods, we can test the theory at the level of individual households and thus account of full heterogeneity across households. In a following step, we will conduct a parametric analysis of our experimental results. In particular, we will compare spouses' preferences over money for own consumption and meals for one of their children, and we will investigate the variation in spouses' bargaining positions across households.

6.1 Revealed preference analysis

Table 7 shows the test results for the unitary model. We find that 86% and 84% of women and men, respectively, pass the Afriat conditions for the unitary model, which were given in Definition 1. This can be considered as a very good fit, especially when taking into account that the experiment involves a substantial variation in the relative price of a child's meal. The average Afriat indices are equal to 97% for both females and males and, thus, high in general. For those individuals not being be perfectly rational (i.e. having an Afriat index that is strictly smaller than 1), the Afriat indices are equal to on average 79% for women and on average 81%

Table 6: Predictors of likelihood wife allocates more or husband allocates more to child, marginal effects

	W allocates more	H allocates more
Price of food, 100s	-0.018	-0.056
	(0.012)	(0.014)
Real income in decision, 100s	0.005	0.032
	(0.005)	(0.006)
Years married	0.004	-0.005
	(0.001)	(0.001)
Number of children	0.000	0.004
	(0.003)	(0.003)
Living in Kangemi	-0.151	0.094
	(0.054)	(0.057)
Age chosen child	-0.006	0.028
	(0.003)	(0.003)
Chosen child is female	-0.052	0.057
	(0.013)	(0.015)
Income ratio (W/H)	0.001	-0.004
	(0.000)	(0.001)
Age ratio (W/H)	0.138	-0.096
	(0.027)	(0.034)
Secondary school or more, H	0.043	-0.064
	(0.014)	(0.016)
Secondary school or more, W	-0.013	-0.057
	(0.015)	(0.016)
H unemployed, W employed	0.038	0.031
	(0.019)	(0.022)
W unemployed, H employed	0.018	-0.002
	(0.016)	(0.018)
Discount factor, H	-0.001	0.001
	(0.001)	(0.001)
Discount factor, W	-0.003	-0.000
	(0.001)	(0.001)

Note: The table displays the marginal effects from a multinominal regression with fixed effects at the session level. Dependent variable: likelihood that spouses allocate the same to the child (base category), likelihood that the wife allocates more to the child than the husband (column (1)), and likelihood that the husband allocates more to the child than the wife (column (2)). Standard errors in parenthesis.

for men. They implicate that the experimental incomes need to be reduced by on average 21% for women and 19% for men who are not perfectly rational according to the GARP test.

The good fit of the unitary model for the individual decisions is not due to a lack of power. This can be concluded on the basis of a power analysis. The latter evaluates the probability of detecting an alternative hypothesis to the unitary model under study. We calculated two power measures. The first is based on Bronars (1987), who used Becker (1962)'s notion of irrational behavior as the alternative hypothesis to the unitary model. The alternative hypothesis states that individuals randomly choose consumption bundles that exhaust the available budget. The first power measure gives the proportion of uniform random draws that result in an Afriat index strictly less than one (and thus that the alternative of random behavior is detected by a rejection of GARP). The random behavior underlying the first power measure allows consumption bundles that can be quite unrealistic. Therefore, we also constructed a second power measure, which gives the proportion of random draws from the observed distribution of choices that result in an Afriat index that is strictly less than one.[14] The results in the table show that both power measures are very high for the female and male individual distributions: random behavior remains undetected in less than 5% of the cases. The good fit of the unitary model for the individual decisions cannot be dedicated to the fact that the GARP test does not have enough power for our data. Interestingly, also 88% of the couples pass the GARP test. Since also this result is not due to a lack of power, we might conclude that the unitary model also performs well to rationalize the joint decisions of couples.

Table 7: Revealed preference test of the unitary model

	Afriat conditions		Power	
	Pass-rate	Afriat index	Random, uniform	Random, bootstrap
Individual decision wife	0.86	0.97	0.96	0.98
Individual decision husband	0.84	0.97	0.95	0.99
Joint decision	0.88	0.98	0.96	0.99

Note: Column (1) displays the pass-rate of individual decisions and joint decisions for the Afriat conditions (or, equivalently, GARP), with the average Afriat index in Column (2). Columns (3) and (4) show the power measures discussed in the main text.

Table 8 gives the results associated with the revealed preference test of the collective model as described in Sections 2 and 3. It turns out that 93.8% of the couples satisfy the IP conditions associated with the collective model (see Definition 2). The average Afriat index over all households is equal to 99.7%, while it is equal to 95.2% for the couples who cannot be perfectly rationalized by means of the collective model focused on. This implies that the fit of the collective model is better than the fit of the unitary model applied to couples. This is the case for both the proportion of households that satisfy the respective revealed preference restrictions and the Afriat indices for couples that cannot be rationalized. It cannot be ruled out, though, that the better fit of the collective model is due to the fact that the power of the associated test

[14]This second power measure is similar in nature to the so-called bootstrap method for quantifying the power of revealed preference tests that was suggested by Harbaugh et al. (2001) and Andreoni and Miller (2002).

is lower than that of the unitary test. The power associated with the IP test, shown in Table 8, turns out to be lower than that of the GARP test applied to couples as shown in Table 7.

As described in Sections 2 and 3, we assume that the spouses' utility functions are the same in the individual and the joint decisions. This implies that individual utility functions only depend on own consumption and meals for one of the children. Or, in other words, that there are no externalities associated with the other spouse's consumption. To evaluate the validity of this assumption, we simultaneously checked the IP conditions of the collective model in Definition 2 for the spouses' data sets gathered in the individual and joint choice experiments. Note that this test is based on 33 choices: 11 individual choices of the wife, 11 individual choices of the husband and 11 joint choices. The results associated with this test are also reported in Table 8. Obviously, the percentage of households that can be fully rationalized is substantially lower than when the individual and joint decisions are separately analyzed: the 33 choices of 49.9% of the couples can be rationalized by our collective model. The average Afriat index equals 91%, while it equals 82% for those households that cannot be completely rationalized. This is lower than the respective numbers associated with the individual and joint decisions when analyzed independently from each other. Still, given the fact that three times as many observations need to be rationalized, we consider this a good fit of our collective model without externalities associated with the other spouse's own consumption. Note also that the power of the latter test is very high: random behavior is quasi always detected when both individual and joint choices are considered.

Table 8: Revealed preference test of the collective model

	IP conditions		Power	
	Pass-rate	Afriat index	Random, uniform	Random, bootstrap
Joint decision	0.94	1.00	0.43	0.50
Individual and joint decisions	0.50	0.91	0.99	1.00

Column (1) displays the pass-rate for the collective IP restrictions, with the average Afriat index in Column (2). Columns (3) and (4) show the power measures discussed in the main text.

Overall, the revealed preference results suggest that the combination of a unitary model for the individual decisions and the collective model for the joint decisions (with the same individual utility functions across decision situations) does a good job in the rationalization of the observed experimental household choices. Consequently, we will use this as a basis for our following parametric analysis of the experimental results.

Before proceeding to this parametric analysis, we use our revealed preference results to measure the validity of the Pareto efficiency assumption (i.e., the main distinguishing assumption of the collective model). More specifically, we investigate whether certain characteristics of the couple can predict whether the joint decision is efficient or not. We use two measures of efficiency: the Afriat index reported in Table 8 and, following our argument in Section 2,

whether the jointly chosen quantity of the public good is situated between the spouses' individually chosen quantities (which is the case in 89% of decisions; see Table A.2). Table 9 displays the results. Whether the couple lives in Kangemi reduces both the Afriat index and the likelihood that the joint decision is between the two individual decisions. Furthermore, if the wife received secondary schooling or more, the Afriat index is higher. The likelihood that the joint decision is between the two individual decisions increases with the price of food and is reduced with the real income in the decisions. A possible interpretation is that when the issue is "more serious" (in the sense that the fraction of total income at stake is larger), individual decisions appear to deviate less from efficiency.

Table 9: Predictors of efficiency

	Afriat index (1)	Between (2)
Years married	-0.001	0.000
	(0.002)	(0.001)
Number of children	0.001	0.003
	(0.002)	(0.002)
Living in Kangemi	-0.102	-0.140
	(0.019)	(0.023)
Age chosen child	-0.002	-0.002
	(0.003)	(0.004)
Chosen child is female	0.009	0.019
	(0.014)	(0.020)
Income ratio (W/H)	0.000	-0.000
	(0.000)	(0.000)
Age ratio (W/H)	0.019	0.025
	(0.045)	(0.031)
Secondary school or more, H	0.014	-0.012
	(0.019)	(0.021)
Secondary school or more, W	0.040	0.030
	(0.019)	(0.021)
H unemployed, W employed	-0.043	0.013
	(0.029)	(0.038)
W unemployed, H employed	0.011	-0.031
	(0.019)	(0.023)
Discount factor, H	0.001	0.001
	(0.001)	(0.001)
Discount factor, W	0.001	0.000
	(0.001)	(0.001)
Price of food, 100s		0.034
		(0.008)
Real income in decision, 100s		-0.010
		(0.005)
Observations	359	3949
R^2	0.247	0.067

Note: The table displays OLS regressions with fixed effects at the session level. Dependent variable: Afriat index of joint decision (column (1)), likelihood that joint decision is between the two individual decisions (column (2)). Standard errors clustered at the session level in parenthesis.

6.2 Parametric analysis

Based on the above revealed preference analysis, we can conclude that the unitary model does a good job in explaining the individual decisions, while the collective model can rationalize the joint decisions of husbands and wives. Given these results, we will now conduct a more standard parametric demand analysis to further investigate the experimental results obtained.

6.2.1 Individual preferences

Elasticities

Let us first discuss the individual preferences of husbands and wives over money for own consumption and meals for one of their children. To obtain these individual preferences, we estimate the demand system given in Equation (4) on the 11 individual decision situations of each respondent. Note that this allows full heterogeneity across individuals; not only between men and women, but also between all individuals of a given gender.[15]

Table 10 displays wives' and husbands' income, compensated own price and uncompensated own price elasticities with respect to the private good (money for themselves, q) and the public good (meals for one of their children, Q). All elasticities are evaluated for the mean budget and mean prices in the experiment. The results indicate that, on average, the private good is a luxury good for women, while it is a necessity for men. For the public good, we see the reverse: a 1% increase in income increases men's demand for meals for one of their children with on average 1.13%, while it increases women's demand by on average 0.94%. The difference between the two income elasticities is not significant, though ($p = 0.10$). Note also that there is quite some variation across individuals. Moreover, the median income elasticities are equal to 1 for both wives and husbands.

For the uncompensated own price elasticity, we see that if the price of the public good increases by 1%, men's demand for the public good decreases by on average 2.86% and women's demand decreases by on average 3.34%, ceteris paribus ($p = 0.10$). The average compensated own price elastiticy for the public good of wives and husbands equals on average respectively -2.98 and -2.48 ($p = 0.09$). These results indicate that the demand for the public good is elastic for both wives and husbands. The same applies to the demand of the private good, albeit to a somewhat lesser extent on average. Once again, it is clear that there is quite some variation across individuals in terms of the price sensitivity of both goods.

Taken together, the results on the income and price elasticities suggest that women and men do not have substantial differences in their preferences over money for themselves and meals for one of their children, on average. To analyze this further, we made cross tabulations of the income elasticities and uncompensated own price elasticities within couples. The results can be found in Tables A.4 - A.7 in the online Appendix. Although there is some bunching in the

[15] See Table A.3 in the online Appendix for the average parameter estimates across couples.

Table 10: Elasticities

	Panel A: Wife's decision				
	Mean	Std. Dev	1st Quartile	Median	3rd Quartile
Income q	1.09	1.10	0.89	1.00	1.37
Income Q	0.94	1.92	0.39	1.00	1.21
Compensated q	-1.88	3.76	-2.00	-0.72	-0.00
Compensated Q	-2.98	4.49	-3.88	-1.00	-0.93
Uncompensated q	-2.62	3.71	-2.74	-1.16	-1.00
Uncompensated Q	-3.34	4.37	-4.32	-1.46	-1.00
	Panel B: Husband's decision				
	Mean	Std. Dev	1st Quartile	Median	3rd Quartile
Income q	0.97	1.22	0.73	1.00	1.21
Income Q	1.13	1.61	0.71	1.00	1.35
Compensated q	-1.86	3.94	-2.07	-0.72	-0.09
Compensated Q	-2.48	4.25	-2.47	-1.00	-0.61
Uncompensated q	-2.48	3.80	-2.74	-1.10	-1.00
Uncompensated Q	-2.86	4.22	-3.20	-1.25	-1.00

Note: The table displays income elasticities, compensated elasticities and uncompensated elasticities for the private and public good, respectively.

third quartiles of the income elasticities, and in the second quartiles of the uncompensated price elasticities, it does not seem that the elasticities of wives and husbands are strongly correlated with each other. Overall, only about 40% of the spouses have elasticities that are in the same quartile. This may suggest that preferences regarding children may not play an important role in the corresponding matching process.

Lindahl prices

Another way to look at spouses' relative preferences for meals for one of their children is by means of their spouse-specific Lindahl prices. Given Pareto efficiency at the couple's level, the Lindahl price of an individual can be interpreted as that individual's willingness to pay for a meal for one of the children.

Table 11 reports on the relative Lindahl prices for each choice situation, and on average over all the choice situations. On average, there turns out to be no significant difference between men and women's relative Lindahl prices over all the choice situations ($p = 0.36$). Some variation in this respect can be observed, though: in six choice situations, women's average Lindahl prices are slightly higher than men's; in the remaining five choice situations, men's average Lindahl prices are somewhat higher than women's.

The quantile information in Table 11 reveals a substantial variation over households. This is also demonstrated by the results in Table 12, which show the proportion of households in which the wife has a higher Lindahl price than her husband. In only 2 out of 11 decisions, it is

Table 11: Women's relative Lindahl prices

	Mean	Std. Dev	1st Quartile	Median	3rd Quartile
Decision 1	0.48	0.29	0.28	0.50	0.67
Decision 2	0.49	0.25	0.37	0.50	0.64
Decision 3	0.51	0.29	0.31	0.50	0.71
Decision 4	0.49	0.23	0.38	0.50	0.62
Decision 5	0.51	0.28	0.30	0.50	0.71
Decision 6	0.52	0.25	0.38	0.51	0.68
Decision 7	0.49	0.26	0.29	0.50	0.69
Decision 8	0.50	0.25	0.35	0.50	0.63
Decision 9	0.54	0.27	0.39	0.53	0.71
Decision 10	0.52	0.25	0.37	0.50	0.66
Decision 11	0.50	0.25	0.35	0.50	0.65
Average	0.50	0.13	0.47	0.50	0.54

most likely that the wife has a higher Lindahl price than her husband.

Taken together, both the experimental results and the parametric analysis of these results suggest that women *do not* have a stronger preference for children's meals than men do; rather, they either have similar preferences or men have stronger preferences for children's meals.

Table 12: Proportion of couples where wife's Lindahl price is higher

	W higher Lindahl price
Decision 1	44.0
Decision 2	45.9
Decision 3	49.6
Decision 4	45.2
Decision 5	49.4
Decision 6	51.5
Decision 7	45.6
Decision 8	45.4
Decision 9	56.0
Decision 10	48.2
Decision 11	48.5
Average	48.1

6.2.2 Bargaining power

The collective model described in Sections 2 and 3 has the spouses' individual preferences (represented by their utility functions) and their bargaining process as main building blocks.

After the discussion of the individual preferences in the former subsections, we now focus attention on the bargaining process.

Several bargaining power measures can be derived on the basis of our collective model. The sharing rule is defined by the spouses' shares in a couple's total expenditure. In Section 3, we showed that the share of spouse i in the couple's total expenditure equals $x_t^i = q_t^i + P_t^i Q$. This share accounts of both the private consumption of a spouse and the expenditure on the public good (meals for one of the couple's children, which are valued at that spouse's Lindahl price). In percentage terms, the relative share of spouse i is equal to $xr_t^i = \frac{x_t^i}{x_t^f + x_t^m}$.

In what follows, we will also analyze the conditional sharing rule, which captures the share of a spouse in the couple's total private expenditure. The conditional share of spouse i in the couple's total private expenditure is thus equal to q_t^i given the normalized price of the spouse's own money. In percentage terms, the relative conditional share of spouse i is equal to $xcr_t^i = \frac{q_t^i}{q_t^f + q_t^m}$.

Finally, we have the spouses' Pareto weights, μ_t^f and $(1-\mu_t^f)$, which are associated with the spouses' utility functions in the maximization problem described in Sections 2 and 3. It can be shown that a spouse's Pareto weight is equal to that spouse's marginal utility of income divided by the sum of both spouses' marginal utilities of income. Clearly, this implies that the Pareto weight is a cardinal measure that depends on the representation of the spouses' individual preferences. This is in contrast with the sharing rule and the conditional sharing rule.

Table 13: Relative sharing rule, conditional sharing rule and bargaining weights

	Mean	Std. Dev	Min	1st Quartile	Median	3rd Quartile	Max
Relative sharing rule	0.53	0.22	0.00	0.45	0.50	0.64	1.00
Conditional sharing rule	0.54	0.20	0.00	0.48	0.50	0.61	1.00
Bargaining weight	0.47	0.31	0.00	0.17	0.50	0.74	1.00

In Table 13, we give summary statistics on the relative sharing rule, the relative conditional sharing rule and the wife's Pareto weight across the eleven choice situations.[16] The relative share of women is equal to on average 0.53. This implies that women receive on average 53% of the couple's total expenditure, which is to be compared to an average share of 47% for men. The average difference between both spouses' shares thus equals 6 percentage points, which is relatively big. Note also the important variation across households, which is demonstrated by the quantile information in the table. The higher bargaining power of females in our experiment is also confirmed by the conditional sharing rule. Also here, Table 13 shows that the relative share of women in the couple's private expenditures is higher than that of men on average. Again, the percentage point difference between the women's relative share and the

[16] As can bee seen in Tables A.8 - A.11 in the online Appendix, there is little variation across the decisions.

men's equals 8, which is a relatively big difference. For completeness, we also show the average Pareto weight of females in Table 13. Once again, it should be stressed that this bargaining power measure has a cardinal nature and that no conclusions should be drawn from the level of the measure. The table shows, though, that also here a big variation can be observed across households. In some households, husbands almost entirely decide on how the household's total expenditure is allocated in some choice situations, whereas in other households, choices are almost entirely driven by wives. This is corroborated by the results based on the (conditional) sharing rule as mentioned before.

Before we continue with an overall discussion of the results obtained, let us have a look at the external validity of our bargaining power measures. To do this, we will analyze the association of our structural bargaining power measures with variables that are often thought to have an impact on individuals' bargaining positions. Examples are relative incomes and education levels, and the size of spouses' social networks (see for instance Attanasio and Lechene, 2014; Bobonis, 2009; Browning et al., 1994). Further, we will check whether there is a connection between our bargaining power measures and an indicator on who has most say in household decisions. Next, we will investigate whether there is a relation between spouses' time use (time spent on personal care and leisure) and our bargaining power measures. This because spouses' leisures are often explained variables in collective models, and as such may be connected to our experimental results (see for instance Blundell et al., 2005; Cherchye et al., 2015, 2011b, 2012; Chiappori, 1988, 1992). In a similar vein, we will analyze whether there is an association between the bargaining power measures obtained through the experiment and the private consumption of both husbands and wives obtained through the questionnaire attached to the experiment.

Table A.12 shows that there is an important and significant positive association between the wife's relative sharing rule and her say in household decisions: the more decisions in which she has a say, the higher her relative share in the resources provided by the experiment. On the contrary, there is a small negative association between her relative sharing rule and her time spent on personal activities. Here one might expect a positive sign in the sense that a better bargaining position is reflected in both a higher private consumption share and more leisure if both commodities are normal goods. These results are corroborated by the results in Table A.13, which focus on the relative conditional sharing rule. As might be expected given the cardinal nature of the Pareto weight, there seems to be no strong association between this bargaining power measure and other indicators according to the results in Table A.14.

6.3 Discussion

The above analysis based on our collective model obtained some interesting findings. Firstly, we did not find any evidence based on the experiment that women in urban informal settlements near Nairobi have stronger preferences in favor of spending money on meals for one of their

children than their husbands. This is somewhat surprising in light of existing studies that focus on relative preferences of mothers and fathers with respect to children's expenditures (see the references in the Introduction). Further, it seems that the bargaining position of women is not worse than that of men on average. On the contrary: we found that women have on average sharing rules that are more favorable to them (although of course, some variation in this respect could be found across households). Again, this contrasts with some other studies in African countries (see for instance Dunbar et al., 2013).

This may beg the question how our results can be explained. Given the particular set-up of our experiment, one possible explanation of men having equally strong or stronger preferences for meals for one of their children than women is that women may have better information about their children's nutritional status. To test for this, we use differences in reported food expenditures on children between men and women as a proxy for knowledge about their children's nutritional status. If men report that less is spent on food (overall or on children) than women, then they might underestimate the nutritional status of their children and therefore might allocate more to the child.

We analyze this hypothesis through the results shown in Table 14. In Column (1) we use an indicator variable that is equal to 1 if the wife reports higher food expenditures than her husband, while in Column (2) we use an indicator variable taking the value 1 if the wife reports higher food expenditures on *children* than her husband. It is clear from the results in the table that the fact that women allocate less to meals for one of their children on average (see the economically and statistically significant coefficient associated with "Wife decision") is not significantly affected by whether men report higher expenditures than their wives. Consequently, the above hypothesis seems to be rejected.

Table 14: Heterogeneity: Average share allocated to child

	Var =	
	W higher food expenditures	W higher food expenditures on children
Wife's decision	-0.081	-0.079
	(0.026)	(0.019)
Wife's decision × Var	0.036	0.056
	(0.035)	(0.034)
Var	-0.005	-0.015
	(0.033)	(0.034)
Observations	7986	7986
R^2	0.293	0.293

Note: OLS regressions with fixed effects at the session level. Dependent variable: share allocated to the child. Controls: wife's age, husband's age, wife's education, husband's education, whether wife is unemployed, whether husband is unemployed, whether couple lives in Kangemi, years married, household size, number of children, age of child chosen to receive meals, gender of child chosen to receive meals, weekly food expenses estimated by the husband (in 1000' of Ksh), weekly food expenses on children by the husband (in 1000' of Ksh), wife's and husband's weekly income (in 1000' of Ksh). Standard errors clustered at the session level in parenthesis.

7 Conclusion

In this paper, we describe an experimental approach aimed at understanding the decision process regarding children expenditures among poor households living in a slum in Nairobi. The design allows us to isolate individual preferences and to concentrate on the decision process. Interestingly, our approach imposes no restriction on the nature or the magnitude of preference heterogeneity, whether between genders or even between individuals of the same gender: individual preferences are estimated *individual by individual*, on the basis of the various choices they are facing.

The main conclusions of our work can be summarized as follows. First, standard theory appears to provide a good description of household behavior, as observed in our experiment. Individual decisions mostly satisfy standard axioms of individual rationality. Regarding the household's joint decision process, a non-cooperative approach seems largely counterfactual. On the contrary, a collective model, which postulates efficient behavior, fits the data quite well. Interestingly, estimates of the spouses' respective bargaining position derived from consumption patterns exhibit a strong, positive correlation with more traditional indicators, such as the respective importance of men and women in decision making (as reported by the individuals themselves).

Secondly, a structural comparison between individual and joint decisions finds a surprising degree of symmetry between genders. On average, there is no significant difference between genders in terms of preferences for children. When choosing independently, spouses tend, in a majority of cases, to allocate the same amount to children; when their choices differ, husbands, if anything, are willing to allocate slightly more than wives. Similarly, the various indicators of bargaining power that we derive do not suggest a significant imbalance in favor of one gender. The sharing rule, the conditional sharing rule and the Pareto weights are, on average, almost equally distributed between husband and wife. These findings strongly contrast with several existing results in the literature, which tend to find either much stronger children preference from women or very unequal intra-household allocations. Whether this discrepancy is due to the specific nature of our experiment or reflects some deeper cultural characteristics remains to be seen, and deserves further investigation. At the very least, our findings suggest that standard prejudices (e.g., 'more money to the wife is always good for children') should be handled with care. Social norms can widely vary with the cultural context.

Finally, while the previous conclusions apply on average, we also observe a very significant level of heterogeneity, both between individuals (including of the same gender) and between intra-household decision processes. To give but one example, while the average sharing rule sits around 1/2 (indicating equal sharing) for all decisions, the standard deviation is about 1/4 in all cases, suggesting that, in a number of couples, the allocation is grossly imbalanced. All in all, the picture provided by this analysis appears much more complex, and ultimately much richer, than a superficial analysis might have suggested.

References

Afriat, S. (1973). "On a system of inequalities in demand analysis: An extension of the classical method," *International Economic Review*, 14: 460–472.

Afriat, S. N. (1967). "The construction of a utility function from expenditure data," *International Economic Review*, 8(1): 67–77.

Akresh, Richard (2008). "(In)efficiency in intrahousehold allocations," Manuscript, department of Economics, University of Illinois.

Akresh, Richard, Damien de Walque, and Harounan Kazianga (2016). "Evidence from a randomized evaluation of the household welfare impacts of conditional and unconditional cash transfers given to mothers or fathers," World Bank Policy Research Working Paper. Washington, DC: World Bank.

Andreoni, James and John Miller (2002). "Giving according to GARP: An experimental test of the consistency of preferences for altruism," *Econometrica*, 70(2): 737–753.

Angelucci, Manuela and Orazio Attanasio (2013). "The demand for food of poor urban mexican households: Understanding policy impacts using structural models," *American Economic Journal: Economic Policy*, 5(1): 146–178.

Armand, Alex, Orazio Attanasio, Pedro Carneiro, and Valerie Lechene (2020). "The effect of gender-targeted conditional cash transfers on household expenditures: Evidence from a randomized experiment," *The Economic Journal*, 130(631): 1875–1897.

Ashraf, Nava (2009). "Spousal control and intra-household decision making: An experimental study in the Philippines," *American Economic Review*, 99(4): 1245–77.

Ashraf, Nava, Dean Karlan, and Wesley Yin (2006). "Tying odysseus to the mast: Evidence from a commitment savings product in the Philippines," *Quarterly Journal of Economics*, 121(2): 635–672.

Attanasio, O. P. and V. Lechene (2014). "Efficient responses to targeted cash transfers," *Journal of Political Economy*, 122(1): 178–222.

Attanasio, Orazio and Valérie Lechene (2002). "Tests of income pooling in household decisions," *Review of Economic Dynamics*, 5(4): 720–48.

Baland, Jean-Marie and Roberta Ziparo (2018). *Intra-Household Bargaining in Poor Countries*, Oxford University Press.

Bargain, Olivier, Guy Lacroix, and Luca Tiberti (2021). "Intrahousehold resource allocation and individual poverty: Assessing collective model predictions against direct evidence on sharing," Technical report, IZA DP No. 14406.

Bateman, Ian and Alistair Munro (2005). "An experiment on risky choices amongst households," *The Economic Journal*, 115(502): 176–189.

Becker, Gary S. (1962). "Irrational behavior and economic theory," *Journal of Political Economy*, 70(1): 1–13.

Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen (2015). "Turning a shove into a nudge? A labeled cash transfer for education," *American Economic Journal: Economic Policy*, 7(3): 86–125.

Bergstrom, T., L. Blume, and H.Varian (1986). "On the private provision of public goods," *Journal of Public Economics*, 29: 25–49.

Blundell, Richard, Martin Browning, Laurens Cherchye, Ian Crawford, Bram De Rock, and Frederic Vermeulen (2015). "Sharp for sarp: Nonparametric bounds on counterfactual demands," *American Economic Journal: Microeconomics*, 7(1): 43–60.

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir (2005). "Collective labor supply with children," *Journal of Political Economy*, 113(6).

Blundell, Richard and Jean Marc Robin (1999). "Estimation in large and disaggregated demand systems: an estimator for conditionally linear systems," *Journal of Applied Econometrics*, 14(3): 209–232.

Bobonis, Gustavo J. (2009). "Is the allocation of resources within the household efficient? New edivdence from a randomized experiment," *Journal of Political Economy*, 117(3): 453–503.

Bourguignon, Francois, Pierre-Andre Chiappori, and Costas Meghir (2009). "Efficient intrahousehold allocations and distribution factors: Implications and identification," *Review of Economic Studies*, 76: 503–528.

Bronars, Stephen (1987). "The power of nonparametric tests of preference maximization," *Econometrica*, 55(3): 693–98.

Browning, Martin, François J. Bourguignon, Pierre-André Chiappori, and Valérie Lechene (1994). "Income and outcomes: A structural model of intrahousehold allocation," *Journal of Political Economy*, 102(6): 1067–96.

Browning, Martin and Pierre-André Chiappori (1998). "Efficient intra-household allocations: A general characterization and empirical tests," *Econometrica*, 66(6): 1241–1278.

Browning, Martin, Pierre-André Chiappori, and Valérie Lechene (2009). "Distributional effects in household models: Separate spheres and income pooling," *Economic Journal*, 120(545): 786–799.

Browning, Martin and Mette Goertz (2012). "Spending time and money within the household," *Scandinavian Journal of Economics*, 114(3): 681–704.

Bruyneel, Sabrina, Laurens Cherchye, and Bram De Rock (2012). "Collective consumption models with restricted bargaining weights: an empirical assessment based on experimental data," *Review of Economics of the Household*, 10: 395 – 421.

Carlsson, Fredrik, Haoran He, Peter Martinsson, Ping Qin, and Matthias Sutter (2012). "Household decision making in rural China: Using experiments to estimate the influences of spouses," *Journal of Economic Behavior & Organization*, 84: 525–36.

Carlsson, Fredrik, Peter Martinsson, Ping Qin, and Matthias Sutter (2013). "The influence of spouses on household decision making under risk: an experiment in rural china," *Experimental Economics*, 16: 383–401.

Cherchye, Laurens, Bram De Rock, Arthur Lewbel, and Frederic Vermeulen (2015). "Sharing rule identification for general collective consumption models," *Econometrica*, 832: 2001–2041.

Cherchye, Laurens, Thomas Demuynck, and Bram De Rock (2011a). "Revealed preference analysis of noncooperative household consumption," *Economic Journal*, 121: 1073–1096.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen (2007). "The collective model of household consumption: A nonparametric characterization," *Econometrica*, 75: 553–574.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen (2009). "Opening the black box of intra-household decision-making: Theory and non-parametric empirical tests of general collective models," *Journal of Political Economy*, 117: 1074–1104.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen (2011b). "The revealed preference approach to collective consumption behavior: testing and sharing rule recovery," *Review of Economic Studies*, 78: 176–198.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen (2012). "Married with children. a collective labor supply model with detailed time use and intrahousehold expenditure information," *American Economic Review*, 102: 3377–3405.

Chiappori, Pierre-Andre (1988). "Rational household labor supply," *Econometrica*, 56(1): 63–90.

Chiappori, Pierre-André (1992). "Collective labor supply and welfare," *Journal of Political Economy*, 100(3): 437–67.

Chiappori, Pierre-Andre and I. Ekeland (2009). "The micro economics of efficient group behavior identification," *Econometrica*, 77: 763–799.

Chiappori, Pierre-André and J. Naido (2020). "The engel curves of non-cooperative households," *Economic Journal*, 130: 653–674.

Choi, Syngjoo, Shachar Kariv, Wieland Müller, and Dan Silverman (2014). "Who is (more) rational?" *American Economic Review*, 104(6): 1518–1550.

Deaton, Angus and John Muellbauer (1980). "An almost ideal demand system," *American Economic Review*, 70(3): 312–326.

Doepke, M. and M. Tertilt (2018). "Women's empowerment, the gender gap in desired fertility, and fertility outcomes in developing countries," *AEA Papers and Proceedings*.

Donni, Olivier and José Alberto Molina (2018). "Household Collective Models: Three Decades of Theoretical Contributions and Empirical Evidence," IZA Discussion Paper No. 11915.

Doss, Cheryl (2006). "The effects of intrahousehold property ownership on expenditure patterns in Ghana," *Journal of African Economies*, 15(1): 149–80.

Duflo, Esther and Christopher Udry (2004). "Intrahousehold resource allocation in Côte d'Ivoire: Social norms, separate accounts and consumption choices," Working Paper 10498. Cambridge, MA: National Bureau of Economic Research.

Dunbar, Geoffrey R., Arthur Lwebel, and Krishna Pendakur (2013). "Children's resources in collective households: Identification, estimation and an application to child poverty in Malawi," *American Economic Review*, 103(1): 438–471.

Fischbacher, Urs (2007). "z-tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10(2): 171–78.

Gobbi, Paula (2018). "Childcare and commitment within households," *Journal of Economic Theory*, 176.

Harbaugh, William T., Kate Krause, and Timothy R. Berry (2001). "GARP for kids: On the development of rational choice behavior," *American Economic Review*, 91(5): 1539–45.

Haushofer, Johannes and Jeremy Shapiro (2016). "The short-term impact of unconditional cash transfers to the poor: Experimental evidence from Kenya," *Quarterly Journal of Economics*, 131(4): 1973–20.

Hoddinott, John and Lawrence Haddad (1995). "Does female income share influence household expenditures? evidence from cote d'ivoire," *Oxford Bulletin of Economics and Statistics*, 57(1): 77–96.

Kibera.org.uk (2019). "Kibera facts & information," `https://www.kibera.org.uk/facts-info/`.

Konrad, K.A. and K.E. Lommerud (1995). "Family policy with non-cooperative families," *Journal of Scandinavian Economics*, 97: 581–601.

Lewbel, Arthur and Krishna Pendakur (2019). "Inefficient collective households: Cooperation and consumption," Technical report.

Lundberg, Shelly, Robert A. Pollak, and Terence J. Wales (1997). "Do husbands and wives pool their resources? Evidence from the United Kindom child benefit," *Journal of Human Resources*, 32(3): 463–80.

Munro, Alistair (2018). "Intra-household experiments: A survey," *Journal of Economic Surveys*, 32(1): 134–175.

Overseers' Education Centre (2012). "Kangemi area," `https://overseerseducationcentre.wordpress.com/2012/11/19/kangemi-area/`.

Oyunga-Ogubi, M.A., N.J. Waudo, A. Afullo, and S.O. Oiye (2009). "Street food in Nairobi, Kenya: Their role as a source of micronutrmicro in low income groups." *African Journal of Food, Agriculture, Nutrition and Development*, 9(4).

Rammstedt, Beatrice and Oliver P. John (2007). "Measuring personality in one minute or less: A 10-item short version of the big five inventory in english and german," *Journal of Research in Personality*, 41: 203–212.

Rangel, Marcos and Duncan Thomas (2019). "Decision-making in complex households," .

Ringdal, Charlotte and Ingrid Hoem Sjursen (2021). "Household bargaining and spending on children: Experimental evidence from tanzania," *Economica*, forthcoming.

Rubalcava, Luis, Graciela Teruel, and Duncan Thomas (2009). "Investments, time preferences, and public transfers paid to women," *Economic Development and Cultural Change*, 57(3): 507–38.

Thomas, Duncan (1990). "Intra-household resource allocation: An inferential approach," *Journal of Human Resources*, 25(4): 635–64.

Thomas, Duncan (1993). "The distribution of income and expenditure within the household," *Annals of Economics and Statistics*, 29: 109–35.

Udry, Christopher (1996). “Gender, agricultural production, and the theory of the household,” *Journal of Political Economy*, 104(5): pp. 1010–1046.

UNICEF (2012). “Children in an urban world,” The state of the world 2012.

UNICEF (2019). “Children, food and nutrition. growing well in a changing world,” The State of the World’s Children 2019.

Varian, Hal R. (1982). “The nonparametric approach to demand analysis,” *Econometrica*, 107(4): 945–972.

Wikipedia (2015). “Kangemi,” Accesses 15.11.2017.

World Bank (2016). “Kenya urbanization review,” World Bank.