Essays on Gender and Power in the Family

By

Matthew J. Klein

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Agricultural and Applied Economics)

at the

UNIVERSITY OF WISCONSIN-MADISON

2020

Date of final oral examination: March 25^{th}, 2020

The dissertation is approved by the following members of the Final Oral Committee:

Bradford Barham, Professor, Agricultural and Applied Economics

Jean-Paul Chavas, Anderson-Bascom Professor, Agricultural and Applied Economics

Paul Dower, Assistant Professor, Agricultural and Applied Economics

Lauren Papp, Vaughan Bascom Professor, Women, Family and Community

Laura Schechter, Professor, Agricultural and Applied Economics

To my parents.

Abstract

> "Power, properly understood, is the ability to achieve purpose. It is the strength required to bring about social, political, or economic changes."
>
> - Martin Luther King Jr. (1967)

When people make decisions as a group, power dynamics determine the extent to which each person's preferences are represented in the final outcome. In the family, when partners make joint decisions like where to live, how many children to have, what food to eat, how much education their children should get and what schools they should attend, or whether to take in an elderly parent, power plays an important role. In theory, and as is often borne out in empirical tests, families with more equal decision makers have better health and investment outcomes.

In the three essays in this dissertation, I expand economists' ability to quantitatively measure power dynamics in families and document some of the causes and consequences of inequality. I develop this measurement methodology in Chapter 1, and extend it in Chapter 3. My key insight is that we can recover estimates of women's bargaining power in the family from panel data on individuals' earnings opportunities under weak assumptions on utility functions. This increases the number of contexts where economists can study power in the family. In Chapter 3, I extend the model to cover cases where family members engage in socially inefficient behaviors, like domestic abuse and hoarding productive assets.

I find that inequality in institutions outside of the household — including labor markets, legal frameworks, customs, and court rulings — is reproduced in power dynamics within the family. When men have systematic advantages in these areas, they also have superior bargaining positions within the family. This gives policy makers a useful tool: to increase equality in the family, increase equality across institutions in society more broadly.

I find that inequality has serious consequences. Using experimental data from Mexico, in Chapter 1, I document that more unequal families have worse diets. Using observa-

tional data from Malawi, in Chapter 2, I find that more unequal families have members who are more likely to contract malaria.

Reference List

King Jr, M. L. (2010). Where Do We Go From Here: Chaos or Community? Beacon Press. Boston

Acknowledgments

Brad Barham, my advisor, has believed in me and my research agenda since the first day we met. I pitched the idea that became Chapter 1 in our first meeting, and he cleared my plate as a research assistant so that I could focus on the project. Without Brad's belief, support, patience, and advice, this document would not be possible, and I would not be the economist I am today. So, first and foremost, I'd like to thank Brad.

I am also deeply indebted to the other faculty in our department who helped me grow, and held my hand as I stumbled through the course work. I am especially thankful for the guidance and help offered by Laura Schechter and Jean-Paul Chavas. In addition, I would like to extend a heartfelt thank you to Barbara Forrest and Mary Treleven.

This is a volume about how married couples make choices for themselves and their kids. I would like to thank my parents for their unwavering support and good humor. This dissertation is dedicated to them. I would also like to thank my brother and my partner, Kendall, for their love and for being goobers.

I would like to thank everyone who helped me with the research in this document, both academically and as friends, not already listed, including (in alphabetical order of last name) Joshua Alfonso, Gabriela Arteaga, Courtney and Dan Bolinson, Colin Bos, Oliver Browne, Mitchell Campbell, Michael Carter, Eduardo Cenci, Joshua Deutschmann, Gaurav Doshi, Paul Dower, Robin Flick, Jeremy Foltz, Seth Gitter, Jakob Gnirke, Gustavo Gordillo, Brian Gould, Arielle von Hippel, Jared Hutchins, Irene Jaqz, Alexey Kalinin, Erik Katovich, Do Soon Kim, Emma Knatterud-Johnson, Dhanesh Krishnarao, Greta Landis, James Manley, Travis McArthur, Craig McIntosh, Ana Paula Melo, Grant Nord, Daniel Phaneuf, Thomas Prendergast, Esteban Quiñones, Katharine Scott, Kaitlyn Sims, Andrew Stevens, Adam Theising, Emilia Tjernström, Denni Tommasi, Justin Winikoff, Catherine Witherspoon, and Yuexuan Wu. You all have made my life in Madison very enjoyable. You've all helped me learn and grow as a person and as an economist.

An extra-special, extra shout out to my 204 N Livingston crew and my job market buddies.

Contents

1 Point Estimates of Bargaining Power Using Outside Options 1

1.1 Introduction 2

1.2 Literature Review 6

1.3 The Model 9

1.3.1 The Collective Model of the Family 10

1.3.2 Choosing Between the Outside Option and the Collective Allocation 12

1.3.3 Consumption Technology 18

1.4 Application to Progresa 24

1.4.1 Progresa and Descriptive Statistics 24

1.4.2 The Relevant Outside Option Does Not Include Divorce 29

1.4.3 Estimating Bargaining Power Levels for Each Household 30

1.4.4 Hypotheses and Testing 34

1.5 Results: The Magnitude of Progresa's Empowerment and the Nutritional Ramifications 37

1.5.1 Caveats 40

1.6 Discussion 42

1.7 Conclusion 44

2 Gender Equality Can Reduce the Malaria Burden in Malawi 68

2.1 Introduction 69

2.2 A Model of the Causes of Inequality in the Family 74

2.2.1 Empirical Strategy 82

2.3 Intra-Household Inequality and Its Causes in Malawi 84
2.4 Differences in Women's Bargaining Power in Matrilineal and Patrilineal Communities 94
2.4.1 Traditions in Malawi 94
2.4.2 Power Dynamics Across Matrilineal and Patrilineal Communities 97
2.5 The Consequence of Intra-Household Inequality in Malawi: Increased Malaria Incidence 102
2.5.1 Malaria in Malawi 103
2.5.2 Women's Empowerment Can Reduce The Malaria Burden in Malawi 105
2.6 Robustness Checks 109
2.6.1 Two Alternative Bargaining Power Measures 111
2.7 Policy Recommendation 112
2.8 Conclusion 115
2.9 Appendix A: Robustness Checks for our Measure of Bargaining Power . . 124
2.9.1 Constructing a Proxy for Bargaining Power Using Principal Component Analysis 124
2.9.2 Applying the PCA Index Method to the Malawi Panel Data . . . 126
2.9.3 Dunbar, Lewbel, and Pendakur's (2013) Collective Model of the Family 126
2.9.4 Applying the DLP Method to the Malawi Panel Data 130

3 Estimates of Household Bargaining Power when Outside Options are Endogenous 143

3.1 Introduction 144
3.2 A Model of Household Decision Making with Limited Commitment and Endogenous Threat Points 147
3.2.1 Formally Defining Appeals Immunity 158
3.2.2 Uniqueness of the Appeals-Immune, Individually-Rational Solution 160
3.3 Identification in a Panel Data Setting 164

3.3.1 Deriving the Functional Form of the Expected Value of Women's Intra-Household Bargaining Power 164
3.3.2 Identification 169
3.4 Conclusion 174

Chapter 1

Point Estimates of Bargaining Power Using Outside Options

Coauthored with Bradford Barham

Abstract: We develop a new way to recover point estimates of women's intra-household bargaining power. Our measure is semi-parametrically identified in a collective model from differences in the shadow price and income settings that partners face in the family versus those they face in their outside options. We use a compensating variation approach to estimate shadow price differences, and we predict earnings to estimate shadow income differences. We generate the first estimates of bargaining power for couples in Mexico's Progresa cash transfer program. Before treatment, the typical wife had 25% of her husband's power; after, 74%. Women's empowerment subsequently improved family diet. (*JEL* D13, I15, I38)

Keywords: Bargaining Power, Outside Option, Human Capital

1.1 Introduction

Women systematically have less agency and power than men, across institutions and across countries (Malapit et al, 2014; Jayachandran, 2015; Hanmer and Klugmen, 2016). Eliminating this power imbalance has the capacity to increase aggregate welfare at many levels. Within the family, empowering women may lead to increased investments in children, generating lasting improvements in their livelihoods (as in Thomas, 1990; or Duflo, 2003). At a national level, empowering women may accelerate growth (Duflo, 2012).

Research on this subject is hampered by a fundamental measurement challenge: power levels are not observable. If researchers could obtain point estimates of men's and women's power, they could directly document the extent of this societal imbalance. Paired with typical program evaluation tools, social scientists could then study how various policy interventions influence power dynamics, and measure the marginal effects of a change in bargaining power on key development outcomes, like investments in children's human capital. Developing new methods to identify and estimate power levels is an important contribution to designing and evaluating policy initiatives aimed at empowering women.

In this paper, we derive two new estimators of women's bargaining power in the family. To do so, we construct a collective model of family decision making with four key characteristics: (1) parents have distinct preferences and make Pareto efficient allocation decisions on their family's behalf; (2) a household-specific consumption technology relates the purchase of goods in the market to individuals' consumption, accounting for externalities; (3) parents' also choose whether to accept this jointly-made decision or to act on their outside option; and (4) parents' outside options can either be divorce, or behavior that induces an inefficient equilibrium in the family. The (shadow) earnings that partners enjoy differ across the collective and private allocations. If the outside option entails divorce, then (shadow) prices differ across settings as well since the partners abandon their shared consumption technology. Combined with predictions for the earnings that each partner would enjoy in their outside option, a subset of the equilibrium conditions from this structural model semi-parametrically identify a point estimate of the female decision-maker's bargaining power for each family.

In our derivation, we first define parents' utility functions as random variables, in a similar fashion to random utility models (McFadden, 1974). Instead of assuming that one functional form correctly describes every person's preferences, we make the weaker assumption that there is a particular distribution of preferences across the population. This modeling choice entails a trade off. Instead of identifying power itself, we identify the expected value of power for each family. We argue that the benefits of this approach outweigh the costs. We show that each of the infinitely many possible utility functions that researchers can impose corresponds to a unique functional form for bargaining power, and thus a different parameter value in estimation. It is not possible to know if the choice of a utility function introduces misspecification error, and assumptions that this does not occur are very strong. Treating utility functions as random variables allows us to circumvent this common issue in other approaches to the study of intrahousehold bargaining power.

The second key step in our derivation is to treat the family-specific consumption technology from a price theoretic perspective (Weyl, 2019). A central challenge in studying familial bargaining power is that individuals' well-being within the family is a function of positive externalities, and how much they benefit from economies of scope in the collective allocation. As Browning, Chiappori, and Lewbel (2013) point out, these benefits of communal living can be modeled with a family-specific consumption technology. This implies that individuals face family-specific Lindahl prices for consuming "private good equivalents" of household goods. We use a compensating variation approach to understand how well off people would be in the case of divorce, where they would face market prices instead of these distinct (and hard to identify) Lindahl prices. These transfer values measure a broad range of prices that cover market-valued and non-market-valued goods.[1] We demonstrate how to recover the difference between partners' compensating variation values. The bargaining power measure when the outside option is divorce is a function of

[1]For instance, they capture how families split assets at the time of divorce. They also capture costs from any potential stigma associated with divorce. They capture the opportunity cost imposed by marriages being between two people. That is, marrying someone comes with the "price" that you cannot marry someone else, which might bring some utility. The compensating variation value amalgamates all of these price concerns into a single parameter.

this difference.

Which of the two estimators a researcher should use depends on whether the outside option in their study context is divorce, or an inefficient equilibrium within the family. The estimator when the outside option is divorce is paired with an estimator for the partners' relative compensating variation values which make them indifferent between the price regime they face inside (Lindahl) and outside (market) the family.

Regardless of which outside option researchers believe best fits their study setting, identification and estimation also rely on predicting the earnings each partner would have in their outside option. Our identification results hold regardless of how the researcher chooses to predict these earnings. The prediction method could be dynamic, could account for censoring in labor markets, could explicitly model earnings differences between single and married people, and may employ a data-driven model-selection method like the lasso.

Our primary contribution is to demonstrate semi-parametric identification of these new bargaining power estimators and offer estimation strategies for them, increasing the tools available to researchers aiming to study gender and power in the family. Our estimators have several advantages over other contemporary measurement options, which we outline in a literature review section. The theoretical advancement stems from the main strength of the collective model: partners are allowed to have distinct utility functional forms. Our semi-parametric identification strategy preserves this characteristic. Improvements in empirical tractability stem from the fact that we do not need to model household demand functions to estimate power dynamics, which can lead to weak-identification issues when demand for privately assignable goods is censored (Tommasi and Wolf, 2018).

We apply our model to answer an outstanding question of great, and recently renewed, policy relevance: to what extent did the Government of Mexico's gender-targeted cash transfer program (Progresa/Oportunidades/Prospera) empower women? We provide point estimates of women's decision-making power within the family for a vast majority of conjugal partners in the well-studied cash transfer experiment, prior to treatment and during the experiment. Then, we exploit the exogenous variation in power levels to identify the causal relationship between women's decision-making power and household demand for healthy foods.

We find that the median woman's bargaining power directly before the cash transfer is 0.20, meaning that men had four times more decision making power than women. We estimate that the large cash transfer program - which increased the median woman's income by thirteen fold and her contribution to the household income by 1000% (Angelucci, 2008) - increased the median female recipients' bargaining power to 0.42. The program more than doubled women's bargaining power. We find that this empowerment had positive and significant ramifications for household diet, increasing the likelihood that a family consumes animal products by 12.6%, and fruits and vegetables by 6%. The program's total effect on demand for healthy food on the extensive margin is about 14% attributable to an empowerment effect, and 86% attributable to an income effect.

Our second contribution is to measure the distributional effects of this famous gender-targeted cash transfer program. Many authors have analyzed whether Progresa empowered women. However, the size of Progresa's empowering effect is still unknown. The nature of the program's effects on the division of surplus within the family is an open question. Adato et al (2000) provided the first evidence on this topic, summarizing interview accounts that suggested women gained more than men from the program, but also reporting that women were less likely to autonomously determine how their incomes were spent. Attanasio and Lechene (2002), Bobonis (2009), and Rubalcava, Teruel, and Thomas (2009) document that Progresa income is spent differently from other income - suggesting that Progresa empowered women - but Handa et al (2009) find the opposite. In some contexts, women were more likely to be subjected to certain forms of intimate partner violence (Angelucci, 2008; Bobonis et al 2013; Bobonis et al 2015).

Most recently, Sokullu and Valente (2018) and Tommasi (2019) have contributed to this literature by estimating resource shares for mothers, fathers, and children in this experimental sample. Sokullu and Valente (2018) adapt the Dunbar, Lewbel, and Pendakur (2013) approach to a panel data setting, and find that Progresa decreased women's resource shares and increased children's. Tommasi (2019) uses a more typical application of the Dunbar, Lewbel, and Pendakur approach and finds that mothers' and children's resource shares increased while fathers' decreased.[2]

This lack of a clear set of program effects on women's bargaining power has serious ramifications. This popular cash transfer program was recently cancelled after more than 20 years of distributing benefits. Prior to its cancellation, this cash transfer served one fourth of Mexican families. More than 60 governments and NGOs worldwide have emulated Progresa. Clearly measuring how much Progresa empowered women can give policy makers a straightforward understanding of how gender-targeting benefits can promote equality. It is possible that this program would not have been cancelled if the literature provided a clearer understanding of the program's empowerment effects. We hope that policy makers designing and pursuing similar cash transfers learn from the unambiguously positive results we document.

1.2 Literature Review

Economists employ four basic strategies to recover point estimates of bargaining power. The first is to estimate a structural model of household demand, as in Browning, Chiappori, and Lewbel (2013). The second is to estimate a structural model of partners' outside options, as in Voena (2015). The third is to use proxy variables, or indices of proxy variables, as in the papers reviewed by Doss (2013). The fourth is to elicit bargaining power from an experiment or game, as in the papers reviewed by Munro (2018). In this literature review, we briefly discuss each approach and its advantages and disadvantages, though a thorough treatment of the subject is beyond the scope of this essay (Donni and Molina, 2018).[3]

Researchers can recover family-specific bargaining power estimates from models of household demand. For instance, Browning, Chiappori, and Lewbel (2013, BCL from

[2]It is possible that the difference is owing to an identification challenge posed by high degrees of censoring in the private assignable goods data. Tommasi and Wolf (2018) point out that censoring can lead to flat Engle curves when there is a large degree of censoring, and in this data, only about 10%-15% of individuals live in households that meet the data requirements. It is possible that a model that explicitly accounts for this censoring would give a different result than these authors find. In the future, updating the Dunbar, Lewbel, and Pendakur (2013) method to explicitly model selection into the group of families that meet their data requirements may allow for broader external validity and improved estimator behavior.

[3]See Chiappori and Mazzocco (2017) for a much more comprehensive review of collective models in general, and of recovering power from structural models.

here on) assume that couples and singles have the same preferences, estimate demand systems for single men, single women, and couples, and identify power as a function of the parameters from this system. Dunbar, Lewbel, and Pendakur (2013, DLP from here on) update the BCL approach by analyzing Engel curves for private assignable goods. They show that family-specific estimates of bargaining power can be backed out from the system of Engel curves for goods like clothing and shoes. They do not have to make assumptions about the preferences of singles and couples, but do need to make weaker assumptions on how preferences are similar across people in the family.[4] Their contribution rests on the insight that private assignable goods enter into the family-specific consumption technology in known ways, and so focusing the analysis on these goods reduces the complexity of the identification problem.

These models, and recent extensions by Wolf (2016), Dunbar, Lewbel, and Pendakur (2017), Chiappori and Kim (2017), and Sokullu and Valente (2018), are promising ways to recover bargaining power in the family. However, recovering resource share estimates from demand models as in DLP suffers from a crucial drawback: few households actually meet the data requirements to estimate the model. In the Progresa data, only about 10-15% of families meet the data requirements for the DLP method. At worst, this threatens identification (Tommasi and Wolf, 2018). At best, it means that researchers can only recover bargaining-power estimates for a small portion of their sample, limiting the statistical power and representativeness of estimates.

The sharing rule can also be recovered from models of partners' outside options, as in the static Nash models of Manser and Brown (1980), and McElroy and Horney (1981); and the dynamic models in Mazzocco (2007) and Ligon (2011).[5] In the version of the collective model with limited commitment, the bargaining power level in the family is a function of the initial dynamic, and the history of updates to partners' outside options. Economists can use this additional information about the sharing rule to achieve identification. For instance, Voena (2015) estimates partners' outside option values to demonstrate how

[4]They also show how identification holds under the assumption that families of the same size have similar preferences.

[5]The static models are nested within the collective model as special cases. As such, using the collective model is more appealing since it does not require researchers to assume the families bargaining structure.

divorce legislation can influence bargaining power. Lise and Yamada (2018) estimate partners' outside options (with an assumption on equality to center bargaining power estimates) in Japan, and corroborate Mazzocco's (2007) theory that unexpected shocks to partners' outside options change family bargaining dynamics.[6] These papers typically assume a functional form for utility in order to predict the value of partners' outside options.

Possibly the most common approach is to use proxy variables, or indices of proxy variables, in place of a derived estimator of power. These methods are easy to implement and offer straightforward interpretation, but suffer from endogeneity problems and are not comprehensive. These estimates can be useful, especially in conjunction with structural models.[7] For instance, researchers have used relative education (e.g. Lundberg and Ward-Batts, 2000), relative assets at the time of marriage (e.g. Doss, 1996), female income share (e.g. Hoddinott and Haddad, 1995), identity of reported primary decision maker (e.g. Ashraf, Karlan, and Yin, 2010), expenditures on gender-specific goods (Lundberg, Pollak, and Wales, 1997), and indices that combine these and other pieces of information (e.g. Ewerling et al 2018), sometimes in arbitrary ways (e.g. Schaner, 2017).

An increasingly popular option to measure bargaining power is to use an experiment (Munro, 2018). For example, Dosman and Adamowicz (2006) elicit preferences for each partner about where to go camping, and then observe the couples actual decision. They find that couples in their sample typically pick the location that women preferred. Carlsson et al (2013) analyze many individual and joint risk and timing decisions, and find that men in their sample typically have more power. Almås et al (2018) elicit women's willingness to pay for two separate cash transfers, one to them and one to their partners. They deduce the couple's sharing rule from the difference between these amounts. These approaches are attractive since they do not require researchers to make assumptions on utility functional forms.

[6] See the summary of limited commitment models in Chiappori and Mazzocco (2017) for additional examples of papers that estimate bargaining power via recovery of the outside options.

[7] Doss (2013) reviews this literature, and writes that "It is not possible to measure women's bargaining power; bargaining power is fundamentally unobservable. At best, we can find good proxies for women's bargaining power."

Our new method for measuring power falls closest to the group of papers that estimate structural models of partners' outside options. The main contributions in our model are to recover estimates without assuming utility functional forms, and to incorporate family-specific consumption technologies into the household's problem. In doing so, we account for challenging issues that arise when the outside option is divorce - for instance, the problem of individuals having different expected values of remarrying. We do not need to use a distribution factor as in Dunbar, Lewbel, and Pendakur (2017), and we utilize information about each spouse separately, as opposed to using household-level demand data. Our primary agenda is to provide an additional empirically tractable option for researchers to use in their studies, and thereby facilitate research on gender and power.

As a final note, it is also worth mentioning that several other studies estimate outside-option earnings. Examples include Blundell et al (2007) who estimate a selection model (as we do in our application) to study the effects of relative wages on the sharing rule; Ramos (2016) who studies the relationship between predicted earnings, bargaining power, domestic violence, and cash transfers; and Cherchye et al (2018) who predict outside-option earnings values to predict which households get divorced in Malawi. Our paper fits in nicely with this group of essays that predict earnings for each partner to learn about bargaining dynamics.

1.3 The Model

We develop the model in three steps. First, we summarize the basic collective model of the family, introduced by Apps and Rees (1988) and Chiappori (1988 and 1992). Second, we add in participation constraints (as in Mazzocco, 2007) to relate outside options to bargaining power. Third, we add in a consumption technology in the family (as in Browning, Chiappori, and Lewbel, 2013) in order to make our concept of the outside option more explicit. We provide proofs and additional details in Appendix A.

In section 1.3.2 we derive the estimator that is appropriate to use when studying contexts where the outside option does not entail divorce, i.e. where divorce rates are low because of legal or cultural barriers. In this case, we assume that the only differences between the socially-efficient collective allocation and the socially-inefficient outside option is captured by differences in shadow earnings across settings, not differences in the shadow prices. This is a weak assumption in some contexts, like in our analysis of poor families living in the Mexican countryside in the 1990s. Benería and Roldán (1987) describe different, common models of household finance in Mexico in 1981 and 1982. About half of the husbands in their sample hide income from their wives. This income-hiding behavior likely leads to lower purchases of public goods, and is an example of the type of outside option we consider in our application. There is no reason to suspect that the shadow prices that partners face — which are directly related to the consumption technology they collectively employ — differ depending on whether men hide income from their partners. In any study, we recommend ethnographic research to determine whether such an assumption is likely to hold.

In section 1.3.3, we derive the estimator that is appropriate to study contexts where divorce is common. This requires us to deal with challenging theoretical problems. For instance, people have idiosyncratic expected values of remarrying. This affects how credible their threat of divorce is, and thus influences power. To solve issues of this nature, we recover the relative compensating variation values that make partners indifferent between the consumption technology used in marriage, and the one they would use after the household dissolves.

In both derivations, we make a weak, and particularly tractable, conditional uniformity assumption on the pair of parents' utility functions. This assumption implies that the expected value of a family's power is the central point of their family-specific subset of the unit interval. This expected value is our estimator of power. We see this approach as a theoretical improvement over assuming that a particular utility function describes peoples' preferences since the former requires a strong assumption that there is no model misspecification.

1.3.1 The Collective Model of the Family

Consider a representative household with two decision makers, indexed by subscripts f and m. They have distinct preferences over their consumption of n-vectors of goods,

x_f and x_m, which have market prices $p = (p_f, p_m)$. The individuals have distinct, monotonically increasing, continuous, strictly quasi-concave, and twice-differentiable utility functions $U_f(x_f)$ and $U_m(x_m)$. Denote household income as y, and let $\tilde{U}[U_f(x_f), U_m(x_m)]$ be some twice-differentiable social welfare function that is strictly increasing in both of its arguments. Household allocations are assumed to be Pareto efficient, and so households solve the program:

$$\max_{x_f, x_m} \quad \tilde{U}[U_f(x_f), U_m(x_m)] \text{ subject to } p_f' x_f + p_m' x_m = y \tag{1.1}$$

A set of Marshallian demand functions solves each households' problem. Because this program results in a Pareto efficient allocation, by the second welfare theorem, it can be implemented by a decentralized economy as well (Chiappori, 1992). This decentralized approach is typically written in two stages, and introduced with the second stage first (as in, e.g. Chiappori and Mazzocco, 2017). In the second stage, both partners solve individual optimization problems subject to prices, household income, and a sharing rule, $\eta \in [0, 1]$:

$$\max_{x_f} \quad U_f(x_f) \text{ subject to } p_f' x_f = \eta y$$

$$\max_{x_m} \quad U_m(x_m) \text{ subject to } p_m' x_m = (1 - \eta) y$$

Each partner solves their problem and gets indirect utility $V_f(\eta y, p)$ and $V_m((1 - \eta)y, p)$. The first stage of this problem is to pick the sharing rule η that maximizes the weighted sum of indirect utility functions:

$$\max_{\eta \in [0,1]} \quad \tilde{U}[V_f(\eta y, p), V_m((1 - \eta) y, p)]. \tag{1.2}$$

The solution to (1.2) results in the demand functions that also solve (1.1). We will proceed from the formulation of the collective model in (1.2).

Note that individuals are only consuming private goods in (1.1) and (1.2). Introducing the consumption technology in section 1.3.3 relaxes this strong assumption completely.

This allows us to include the well-being of other family members in the households problem (inside the decision-makers' utility functions) regardless of whether these individuals have credible outside options. Infants, for example, do not have outside options but their well-being, by assumption, can enter into parents' utility functions.

As in Browning, Chiappori, and Lewbel (2013), we interpret the sharing rule η as bargaining power, since it determines the division of surplus in the family, and because η is not subject to arbitrary cardinalizations of the utility functions.

1.3.2 Choosing Between the Outside Option and the Collective Allocation

Each partner can act on an outside option, which can either be to divorce their partner, or to move the household to an inefficient equilibrium *without dissolving the family* (Lundberg and Pollak, 1993).[8] That is, each partner either chooses between the collective allocation and divorce, or between the collective allocation and some non-divorce outside option. In general, these three contexts (the status quo, divorce, and an inefficient equilibrium in the family) can vary in terms of the income and pricing setting that individuals face. For now, we consider only a choice between the status quo of marriage, and some generic outside option where the price setting remains constant, but the income setting changes. We discuss changes in both the price setting and the income setting after introducing the family's consumption technology in section 1.3.3.

If partners cooperate, they split the common resource according to the family sharing rule, and attain indirect utilities of $V_f(\eta y, p)$ and $V_m((1-\eta)y, p)$. If they act on the exit option, they receive some alternative indirect utilities $V_f(y_f^o, p)$ and $V_m(y_m^o, p)$, where y_f^o and y_m^o are the incomes each partner would have in the outside option. These outside option allocations are generally inefficient in that they lead to the under-provision of public goods. As such, introducing these outside options as participation constraints relaxes the assumption of *ex ante* Pareto efficiency used in the collective model with full

[8] An example of this type of inefficienct equilibrium within the family is one where men are violent, or threaten violence, to ensure a specific allocation is reached (Ramos, 2016; Lewbel and Pendakur, 2019). According to estimates from the World Health Organization, one in three women experience physical or sexual violence from an intimate partner in their lifetime (UN World Health Organization, 2014).

commitment above. These choices can be written as

$$max\{V_f(\eta y, p), V_f(y_f^o, p)\} \text{ and}$$

$$max\{V_m((1-\eta)y, p), V_m(y_m^o, p)\}$$

.

Since these indirect utility functions are strictly increasing in their first arguments, holding prices constant (for the time being), this couple of problems simplifies to

$$max\{\eta y, y_f^o\} \text{ and}$$

$$max\{(1-\eta)y, y_m^o\}.$$

We can write these two choices, and the choice of how to allocate the family's resources in (1.2), as a single program. This program contains the model in (1.2) as a special case (as in, e.g., Mazzocco, 2007). The household's problem is given in (1.3), which has the corresponding Kuhn-Tucker formulation given in (1.4), and which results in the following first order conditions (with suppressed notation):

$$\underset{\eta\in[0,1]}{max} \quad \tilde{U}[V_f(\eta y, p), V_m((1-\eta)y, p)] \text{ s.t. } \eta y \geq y_f^o \text{ and } (1-\eta)y \geq y_m^o, \tag{1.3}$$

$$\underset{\eta}{max} \quad \tilde{U}[V_f(\eta y, p), V_m((1-\eta)y, p)] + \lambda_1(\eta y - y_f^o) + \lambda_2((1-\eta)y - y_m^o) + \lambda_3\eta + \lambda_4(1-\eta), \tag{1.4}$$

$$\frac{\partial\tilde{U}}{\partial V_f}\frac{\partial V_f}{\partial\eta} + \frac{\partial\tilde{U}}{\partial V_m}\frac{\partial V_m}{\partial\eta} + y(\lambda_1 - \lambda_2) + (\lambda_3 - \lambda_4) = 0 \tag{η}$$

$$\lambda_1 \geq 0; \quad \eta y - y_f^o \geq 0; \quad \lambda_1(\eta y - y_f^o) = 0 \tag{λ_1}$$

$$\lambda_2 \geq 0; \quad (1-\eta)y - y_m^o \geq 0; \quad \lambda_2((1-\eta)y - y_m^o) = 0 \tag{λ_2}$$

$$\lambda_3 \geq 0; \quad \eta \geq 0; \quad \lambda_3\eta = 0 \tag{λ_3}$$

$$\lambda_4 \geq 0; \quad 1-\eta \geq 0; \quad \lambda_4(1-\eta) = 0 \tag{λ_4}$$

The first equilibrium condition tells us that the household sharing rule is the one that sets the marginal gain to partner f from an increase in η equal to the marginal loss to partner m. Solving for η requires us to consider multiple cases, corresponding to the complimentary slackness conditions. The solution can be written as a piecewise function in three of four possible cases:

$$\eta^* = \begin{cases} \frac{y_f^o}{y} & \lambda_1 > 0, \lambda_2 = 0 \\ 1 - \frac{y_m^o}{y} & \lambda_1 = 0, \lambda_2 > 0 \\ \frac{1}{2} & \lambda_1 > 0, \lambda_2 > 0, y_f^o = y_m^o = \frac{y}{2} \end{cases}$$

In the fourth possible case neither participation constraint binds and $\lambda_1 = \lambda_2 = 0$. Consider only interior solutions ($\lambda_3 = \lambda_4 = 0$), then η^* is bounded on the interval $\left[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}\right]$ and solves $\frac{\partial \tilde{U}}{\partial V_f}\frac{\partial V_f}{\partial \eta^*} = -\frac{\partial \tilde{U}}{\partial V_m}\frac{\partial V_m}{\partial \eta^*}$. The primary solution of interest is the fourth case, where neither participation constraint binds, since the set of values $\left[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}\right]$ includes the cases where the participation constraints bind as special cases.[9]

Power *could* be recovered from equation (η), the first order condition for program (1.4) with respect to η, alone. By assuming functional forms for each partner's utility functions, and a functional form for the social welfare function, η can be recovered for each household. There is parametric identification in this collective model, as usual. Researchers can estimate this model given assumptions on these functional forms. This is the approach that BCL and DLP use. BCL assume utility functions that generate quadratic

[9] It also contains the case where both participation constraints bind as a special case where the outside options are equal to each other, and to half of family income.

almost ideal (QUAIDS) demand functions, and a logit functional form for the sharing rule; DLP assume price independent generalized log (PIGLOG) indirect utility functions, and that there are similarities in household demand for different family members' private assignable goods (or similarities across families of different sizes). These authors derive estimators for bargaining power that are conditional on these assumptions. However, it is unclear whether resource share estimates are robust to different assumptions about utility functions. Our theoretical advancement, as those in Chiappori and Kim (2017) and Dunbar, Lewbel, and Pendakur (2017), is to recover estimates of bargaining power while relaxing these assumptions on utility functional forms.

We suggest pursuing an alternative, semi-parametric identification approach. Instead of assuming a single functional form to describe an individual's preferences (which is not verifiable), we make two key, alternative assumptions. These allow us to make small theoretical advancements, and to achieve large gains in empirical tractability.

The first key identifying assumption in this approach is that there are identified models of earnings in the outside option. Let $y_f^o = F(X_f, \psi_f)$ and $y_m^o = F(X_m, \psi_m)$, where X is a set of person specific observable characteristics and ψ is latent ability. The first order conditions (λ_1), and (λ_2), and the prediction models $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$, pin down the set $\left[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}\right]$ for each family. The estimator of this set is $\left[\frac{F(X_f,\psi_f)}{y}, 1 - \frac{F(X_m,\psi_m)}{y}\right]$, and will have family-specific upper and lower bounds. We leave the choice of how to predict y_f^0 and y_m^0 to the econometrician. Certain modeling approaches will be more appropriate for each context. Regardless of how the econometrician predicts these earnings, though, the estimator that gives the family-specific bounds is $\left[\frac{F(X_f,\psi_f)}{y}, 1 - \frac{F(X_m,\psi_m)}{y}\right]$. To move from partial identification to point identification, we need a second key assumption.

Our second key assumption is that there exists some (unobservable) distribution on the infinite set of utility functions that individuals might have. Denote this set of functions as C, and the distribution on C as G. We treat each person's preferences as a random variable, instead of assuming a single function describes preferences for the population. That is, when sampling a family from the population, the researcher observes demand and supply data, as well as characteristics of the family, like the age of its members. In addition, each family has an unobservable set of preferences and power dynamics, drawn from some distributions. Each person in the population can have distinct preferences, summarized by distinct functional forms.

Our semiparametric approach is to restrict the possible probability distributions that describe the likelihood of observing any pair of utility functions for decision makers in a family. One especially tractable, though likely not correct, assumption is that G is a uniform distribution so that any utility function is equally likely to describe an individual's preferences. This uniformity assumption is tractable because it implies that η is uniformly distributed on $\left[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}\right]$.[10] As such, a natural estimator for bargaining power would be the expected value of this distribution, which is its central point:

$$\hat{\eta} = \mathbb{E}\left[\eta \quad \middle| \eta \sim \text{Uniform}\left[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}\right]\right] = \frac{1}{2} + \frac{1}{2}\frac{F(X_f, \psi_f) - F(X_m, \psi_m)}{y} \tag{1.5}$$

However, it may be that some preferences are more common than others. G might be unimodal. To allow for some preferences to be more common than others, we must refine the assumption we make on G.

Let there be some subsets of C that form a partition, and denote each of these subsets as $c_l \subset C$ for $l \in \{1, \ldots, L\}$, where L is an arbitrarily large number. For each subset, let there be a corresponding distribution, g_l, that has support equal to c_l. For each family, let the partners' utility functions be random variables drawn from a single probability distribution g_l. As such, partners will have similar preferences, in the sense that they have random draws from the same distribution, g_l, on the same subset of C, c_l. More than one family may have partners with utility functions drawn from a particular subset, c_l.

In this way, instead of making a single, strong uniformity assumption on G, we can make assumptions on how partners' preferences are drawn from their family-specific par-

[10] This statement requires some unpacking, which we do in Appendix A. In short, there is a functional $\Gamma : C \times C \to [\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}]$ such that $\eta^* = \Gamma(\tilde{U}(U_f, U_m))$. This functional is a differential operator on the social welfare function, $\tilde{U}$, since it returns the value of η that solves (3) for any pair of utility functions, set of prices, income, and outside option values. Therefor it is a linear operator. As such, η is a linear transformation of two uniformly distributed variables, and so η is uniformly distributed.

tition g_l. We can refine our second main assumption so that each person's utility function is a random variable that is conditionally uniformly distributed. For each decision maker within a family, assume that their preferences are drawn uniformly from their family-specific partition, c_l. The estimator in (1.5) obtains from the same line of reasoning as before (A linear differential operator relates the randomly distributed utility functions to a family's power dynamic, and so power itself is a uniformly distributed variable).

Since we need not make any assumptions on how this group of L subsets partitions C, this assumption is very weak. Regardless of how the partition is drawn, the assumption that partners' preferences are random variables uniformally distributed on the same subset of permissible utility functions is sufficient to derive the estimator in (1.5). It doesn't matter whether some types of utility functions are more common than others in the population - this (unverifiable assertion) amounts to more families drawing form certain subsets of C. This refinement allows for this possibility, but does not enforce it. When estimating the model under this joint conditional uniformity assumption, the econometrician does not need to recover any information about the partition on C, or about partners' preferences.

This estimator is semi-parametrically identified (as opposed to non-parametrically identified) because it requires us to assume functional forms for $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$. It is semi-parametric identified (instead of parametrically identified) because it does not require us to assume functional forms for utility, or the social welfare function.

The functional form for our estimator has several nice features. When the two decision makers have equal outside options, the sharing rule is 0.5, and decision makers are equal. The partner with the higher outside has more decision making power, and the difference from equality is the difference between the outside option values scaled by household resources. Interestingly, when household income is higher relative to the outside option values, the differences between partners' outside options matters less in determining the sharing rule.

As a final comment on the estimator in (1.5), note that it will update in response to a change in the value of outside options even if neither partners' participation constraint binds because of that change. The bounds on the values that power can take in each family are informative. This is in contrast to limited commitment models (which are typically used in an intertemporal setting), where power only updates if the change to the value of outside options causes one partner's participation constraint to bind. In this way, the formulation of power in this model is more similar to what would obtain from a repeated game framework, where the relative credibility of each partner's threat point determines the equilibrium allocation.

1.3.3 Consumption Technology

A key feature of the family, and one of the reasons why households form, is that individuals can jointly consume many goods. They may share car rides, so that gasoline and vehicle maintenance costs are shared. It is cheaper to heat one home and let all family members consume the heat, than to heat the many separate homes of individuals living apart. Food waste may be reduced if a couple jointly consume meals, compared to the individuals cooking for themselves.

This fact is particularly important to model when the outside option is divorce. In this context, as we will formalize in this section, individuals face a different set of prices in the family and outside of it. In the family, they face Lindahl prices. If they choose to dissolve the partnership, they face market prices. As such, the choice of whether to get a divorce also depends on the differences between these Lindahl prices and market prices.

BCL model these Lindahl prices by introducing a consumption technology to the family's problem. They assume some function exists that relates household purchases z to individual equivalents, x_f and x_m, $z = M(x_f + x_m)$. Since identification using our strategy does not require the use of consumption data, we do not focus on details of this function, but simply note that some function exists that relates market purchases to household consumption. As BCL point out, this is an application of Becker's (1965) home goods model.

For each family, these consumption technologies generate a set of shadow prices that are weakly less than market prices. Denote the vector-valued function that relates the market prices to shadow prices for each family as $A(p)$, where each element of $A(p)$ is weakly smaller than the corresponding value in p. As DLP point out on page 442, "each member faces the same shadow prices because the degree to which a good can be shared is an attribute of the good, rather than an attribute of the consumer."

Then, couples' choices whether to divorce or not will also depend on the jointness of goods while they are living in the family. The comparisons they make when divorce is the outside option are now:

$$max\{V_f(\eta y, A(p)), V_f(y_o^f, p)\}\text{and}$$

$$max\{V_m((1-\eta)y, A(p)), V_m(y_o^m, p)\}$$

.

We cannot simplify these problems as we did before, but need to use a compensating variation argument to proceed. There is some set of transfers (γ_f, γ_m) that would make individuals indifferent between consuming at the Lindahl prices and the market prices. We know that, since families face the same market prices and the same consumption technology M, the difference between γ_f and γ_m is owing to the differences in their utility functions. The compensating variation for partner f is the value γ_f that solves $V_f(\eta y, A(p)) = V_f(\eta y + \gamma_f, p)$. For partner m, γ_m solves $V_m((1-\eta)y, A(p)) = V_m((1-\eta)y + \gamma_m, p)$. The values γ_f and γ_m summarize the degree to which people benefit from using the household consumption technology. Paired with the sharing rule and household income, they characterize individuals' gains from trade when the outside option is divorce. Substituting in the indirect utility functions with the compensating variation parameters allows us to write the partners' problems as

$$max\{V_f(\eta y + \gamma_f, p), V_f(y_o^f, p)\} \text{ and}$$

$$max\{V_m((1-\eta)y + \gamma_m, p), V_m(y_o^m, p)\}.$$

Since these functions are strictly increasing in their first argument, we can simplify these problems:

$$max\{\eta y + \gamma_f, y_f^o\} \text{ and}$$

$$max\{(1-\eta)y + \gamma_m, y_m^o\}.$$

In the case that the outside option is an inefficient equilibrium in the family, partners face the same Lindahl prices regardless of their choice over the collective allocation or the outside option. If the outside option is divorce, the problem in (1.3) becomes:

$$\max_{\eta\in[0,1]} \quad \tilde{U}[V_f(\eta y, A(p)), V_m((1-\eta)y, A(p))] \text{ s.t. } \eta y+\gamma_f \geq y_f^o \text{ and } (1-\eta)y+\gamma_m \geq y_m^o \quad (1.6)$$

The first order conditions for the Lagrangian for (1.6) are given by (λ_3) and (λ_4) above, and:

$$0 = \frac{\partial\tilde{U}(V_f,V_m)}{\partial V_f(\eta y, A(p))}\frac{\partial V_f(\eta y, A(p))}{\partial \eta} + \frac{\partial\tilde{U}(V_f,V_m)}{\partial V_m((1-\eta)y, A(p))}\frac{\partial V_m((1-\eta)y, A(p))}{\partial \eta} + y(\lambda_1-\lambda_2)+(\lambda_3-\lambda_4) \quad (\eta')$$

$$\lambda_1 \geq 0; \quad \eta y + \gamma_f - y_f^o \geq 0; \quad \lambda_1(\eta y + \gamma_f - y_f^o) = 0 \quad (\lambda_1')$$

$$\lambda_2 \geq 0; \quad (1-\eta)y + \gamma_m - y_m^o \geq 0; \quad \lambda_2((1-\eta)y + \gamma_m - y_m^o) = 0 \quad (\lambda_2')$$

From these first order conditions for (1.6), we know that the equilibrium η^* will solve the equation (η') and be bounded on $\left[\frac{y_f^o-\gamma_f}{y}, 1-\frac{y_m^o-\gamma_m}{y}\right]$.

Parametric estimation of this model is far more challenging than it is for the model in (1.3). Assuming functional forms for utility and the social welfare function, and a form for A, (η) gives one equation per household, but far more parameters to be recovered. For each household, there is a bargaining power measure and the elements of A to estimate. This is a very hard problem to solve, and the key contributions of BCL and DLP are to present new ways to do so using household demand data.

As before, an alternative approach is available. The expected value of power is identified on a subset of the first order conditions and the prediction models: (η'), (λ_1'), (λ_2'), $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$. However, a drawback of the semiparametric approach in this case is that the parameters γ_f and γ_m are not separately identified. Only their difference is identified. Applying the same conditional-uniformity assumption on the joint distribution of partners' preferences as before, we can write the functional form for the expectation of η as:

$$\hat{\eta} = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_f, \psi_f) - F(X_m, \psi_m) + (\gamma_m - \gamma_f)}{y}\right). \tag{1.7}$$

One interesting feature of this model is that, with divorce as the outside option, the relative gains from joint consumption also determine the sharing rule. If $\gamma_f > \gamma_m$, then partner f gains more from the common consumption technology than partner m does, and divorce is a less credible threat for partner f than it is for partner m. This reduces partner f's bargaining power. This is an intuitive result, and this difference is identified in our semi-parametric approach.

Estimating (1.7) requires an additional step, compared to estimating (1.5), and has an additional data requirement. The additional data requirement is that observations exist for more than one period. The additional step is to estimate a linear function of family income. Rearranging (1.7) gives:

$$y = \frac{F(X_f, \psi_f) - F(X_m, \psi_m) + (\gamma_m - \gamma_f)}{2\hat{\eta} - 1}. \tag{1.8}$$

Income is linear in the difference between outside option earnings predictions and the parameters of interest. Let subscript t denote a time period, and let there be data for each household in $T \geq 2$ time periods.[11] Define the following parameters for individuals f and m in household h and period t, so that the linear equation in (1.8) can be re-written as in (1.9) for each period:

[11]BCL require $T \geq 13$ for their main identification proof, since they must recover far more information to learn about the consumption technology than we need to learn about $\gamma_m - \gamma_f$. Requiring more than one time period is a weak data requirement by comparison.

$$\gamma_{h,t} \equiv \gamma_{m,h,t} - \gamma_{f,h,t}$$

$$\beta_{0,h} \equiv \frac{\gamma_{h,t}}{2\hat{\eta}_{h,t} - 1}$$

$$\beta_1 + \epsilon_{1,h,t} \equiv \frac{1}{2\hat{\eta}_{h,t} - 1}$$

$$\epsilon_{h,t} \equiv \epsilon_{1,h,t}(F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t}))$$

$$y_{h,t} = \beta_{0,h} + \beta_1(F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})) + \epsilon_{h,t} \tag{1.9}$$

In estimating (1.9), the econometrician must also specify the constraint that $\beta_1 + \epsilon_{1,h,t} \in [-1, 1]$, otherwise the first order conditions (λ_3) and (λ_4) could be violated.[12] This adds an additional H constraints to the fixed effects model, where H is the number of households in the sample. Estimating equation (1.9) with this constrained household fixed-effects strategy gives us one intercept estimate per household, $\tilde{\beta}_{0,h}$; one slope estimate, $\tilde{\beta}_1$; and one error term for each household in each period, $\tilde{\epsilon}_{h,t}$.[13] Plugging these sample estimates into our definitions of β_0 and β_1 gives our estimators as a function of the fixed effect regression estimates:

$$\hat{\eta}_{h,t} = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})}{y_{h,t} - \tilde{\beta}_{0,h}}\right) \tag{1.10}$$

$$\gamma_{h,t} = \tilde{\beta}_{0,h}(2\hat{\eta}_{h,t} - 1)$$

The parameters $\gamma_{h,t}$ and $\eta_{h,t}$ vary over time, while $\beta_{0,h}$ does not. As such, one technical

[12] Estimating this fixed effects model with constraints may proceed, for example, using the simulated method of moments to select the estimates that minimize the sum of squared errors from a family-specific permissible support.

[13] We denote these estimates with tildes instead of hats because they obtain from a constrained least squares problem, not an ordinary least squares problem.

assumption built into this model is that $\gamma_{h,t}$ and $\eta_{h,t}$ must vary in a constant proportion over time. That is, by assumption, the left hand side of the following equation does not change over time, even while the parameters on the right hand side do: $\hat{\beta}_{0,h} = \frac{\hat{\gamma}_{h,t}}{2\hat{\eta}_{h,t}-1}$. This technical condition does not provide additional insight into the how the model works, but does need to hold for the fixed effects estimation process to correctly recover a point estimate of women's bargaining power in the family.

An important feature of this model is that we can incorporate the expected value of remarrying in partners' choices to stay married or get divorced (as in Cherchye et al 2018). We can also incorporate some loss of utility from stigma associated with divorce. These parameters cannot be estimated separately, but power can be recovered conditional on the differences between partners' parameter values. Note that parameters that add or subtract some value from the outside option utility would enter linearly into the participation constraints. As such, they enter the constrained least squares model via $\gamma = (\gamma_m - \gamma_f) + (S_m - S_f) + (R_m - R_f)$ where S_f and S_m are the partner-specific stigma values, and R_f and R_m are partner-specific expected gains from remarrying. Then, we could estimate (1.9) to recover the sharing rule, and γ, which now has a more convoluted interpretation. So, estimates of power can be recovered even with this additional set of unknown parameters added to the model, since the number of parameters to be estimated (η and γ) does not change.

When divorce is the relevant threat point, this model can account for considerations like stigma and the conditions in the marriage market. In fact, we can incorporate any number of linear parameters into the participation constraints. This model is appropriate for studying power when these considerations are thought to be important, but is inappropriate for studying the specific elements comprising γ.

This exercise allows us to formalize the difference between the two outside options. Notice that the functional form of the estimator in (1.8) reduces to the form in (1.5) if $\gamma = 0$. The non-cooperative inefficient equilibrium in the family is the same as a divorce in which there is no change in the partners' use of the family consumption technology, no possibility of remarrying someone else in the future, and no stigma to reduce partners' utility levels.

1.4 Application to Progresa

We analyze Mexico's welfare program, Progresa, because the government sought to increase women's relative decision-making power, and yet no point estimates of bargaining power have been recovered for a majority of the families in the experiment's sample. We briefly discuss the welfare program and setting since detailed accounts are readily available (e.g. Parker and Todd, 2017). We focus on food expenditures since they constitute a large share of household expenditures (Attanasio and Lechene, 2002; Rubalcava, Teruel, and Thomas, 2009) and because nutrition is essential to children's human capital accumulation. Further, there are positive, documented links between power and diet in other contexts (e.g. Thomas, 1990), and we'd like to see if they are present in this context too. We estimate the expected value of power for each family in (1.5), using Heckman selection models to predict y_f^o and y_m^o. We present the empirical analogs of our hypotheses at the end of this section and the results in Section 5. We conclude the application with a discussion of its shortcomings.

1.4.1 Progresa and Descriptive Statistics

In 1997, the Mexican government surveyed rural populations with the goal of identifying which households were eligible to receive Progresa benefits. Communities were randomly assigned treatment or control status, successfully creating ex-ante comparable treatment and control groups (Behrman and Todd, 1999). In May of 1998, eligible households in treatment communities began receiving benefits. Over time, control communities were phased into the program and, by November of 2000, all eligible households were receiving benefits. Comparing eligible households in early-treatment and late-treatment communities from 1998 to 2000 allows us to conduct causal inference.

Progresa distributed three sizable transfers (a scholarship, a grant to buy school supplies, and a nutritional grant) to families who met the health and education conditions. These three grants summed to be as large as 20% of control families' expenditures (Hoddinott and Skoufias, 2004), and were paid directly to each households' female head. Additionally, female household heads were required to attend health courses that covered (among many other, non-diet topics) the importance of healthy diets. Parker, Rubalcava, and Teruel (2008) provide a detailed description of the payment schedules and features, and synthesize some findings about the program's effectiveness.

Table 1.1: Summary Statistics

	1997	1999	2000
Mean 16-70 Y.O. Working Female Earnings (pesos/week)	122	158	199
Mean 16-70 Y.O. Working Male Earnings (pesos/week)	146	176	238
Female Labor Force Participation	13%	8%	7%
Male Labor Force Participation	84%	85%	72%
Mean HH Progresa Transfer (pesos/week)	0	94	120
Median Female Household Head Age	38	41	41
Median Male Household Head Age	43	46	46
Median Female Household Head Education Years	2	2	2
Median Male Household Head Education Years	3	3	3
Median Number of Children Per Household	2	2	3
Mean Total Household Income (pesos/week)	265	332	421
Mean Total Household Income in 1997 Pesos/Week	265	237	273
Percent Speaking an Indigenous Language	27.9%	27.8 %	28.5%
Number of Households	15,968	16,002	14,252

In addition to earnings and labor market decisions, we utilize information about household structure, diet, and migration collected in the Progresa survey.[14] The median house-

[14] A note on the sample we analyze and dataset construction follows. We analyze a specific subset of the data - households with a clear male and female head. This subset allows us to explore bargaining dynamics between two spouses, as opposed to accidentally including households in our analysis that prescribe to an alternate bargaining arrangement, such as between three people, or between two individuals not married to each other. An example of the latter might be a household whose decisions are jointly made by a male household head and his father, who lives in the home with the couple. To identify this subset of conjugal couples, we utilize information from the survey about who the head is, who their spouse is, and the order in which the head listed household members when asked to divulge who lives with them. We keep households whose ordering is household head, spouse, then others. We drop households with only one household head or whose individual ordering has someone listed before the head's spouse. We also drop households with individuals that did not respond to the income and labor questions, or responded that they did not know. In all, we drop about 40% of the households in the Progresa sample across these three waves (the total number of households observations is 78,072). We drop households containing earners in the top and bottom 1%, enabling better predictions. The total number of households in the each period is displayed in Table 1. We only use three waves of data, the baseline, November 1999, and November 2000, because the income modules are different in the remaining waves. Our analysis relies heavily on directly comparing earnings information over time, so we opt to only use waves with the exact same earnings questions.

hold in our sample has a combined income of 265 pesos/week in 1997;[15] has a male worker in an agricultural position; has 2 children, one of whom is under the age of five, and has two literate adults. See Table 1.1 for additional descriptive statistics.[16]

In Table 1.2, we provide a breakdown of how all men and women ages 16-65 (not just household heads) earn income by year and gender. We report the average transfer amount in inflation-adjusted pesos per week in parentheses. We construct our Progresa payment variables from survey responses on whether a household member received any of the three Progresa transfers and, if so, which member and how much they received.

In the market, the primary sources of income are employment as a laborer, and entrepreneurship. Government transfers are the most important sources of income for women, with Progresa providing more for women than any other source of income.[17] Procampo has almost the same reach, but gives cash transfers almost exclusively to men. In general, men have more opportunities to generate income, men garnered more income when they participated in the market or received a transfer, and opportunities decreased for everyone from 1997 to 2000. We use all eighteen sources of earning incomes in our analysis.

The Progresa survey includes 7-day recall information on whether specific foods were consumed in the household and the frequency of consumption. We examine thirty-one of these foods, individually and at times grouped into food types (e.g., animal products, and fruits and vegetables).[18] The binary and count nature of these data are reflected in our empirical models where we estimate linear probability models and Poisson regressions of dietary outcomes.

Progesa's impact on diet is well documented: while nutrition worsened for the average household from March 1998 to November 1999, Progresa had a mitigating effect.

[15]In 1997, one peso was worth $0.11.

[16]Values in this table are not adjusted for inflation, which was substantial in our study period. This does not affect our analysis since we use year fixed effects in all regressions, and (in the construction of the power measures) individuals face the same inflation settings. See Skoufias (2001) for comprehensive summary statistics concerning household structure and income. See Hoddinott, Skoufias, and Washburn (2000) for a comprehensive summary of household diet in the sample.

[17]Roughly 30% of female household heads in this sample report receiving the Progresa transfer.

[18]There are 36 food types total, but information is not collected in all waves for some. For 4 of the foods that we do not examine, fewer than 1% of households reported consuming them at all.

Hoddinott and Skoufias (2004) document that treatment households consume 6.4% more calories than control households in November 1999. Attanasio and Lechene (2014) provide corroborating evidence - treatment households' income share spent on food does not decrease even given the substantive increase in total household income associated with the transfer. Typically, they point out, income shares spent on food are decreasing in income. They infer that women's empowerment offsets this typically-observed effect. Behrman and Hoddinott (2005) also provide supporting evidence, recording a significant reduction in stunting for treatment household children, thereby inferring that Progresa improved diets. We present summary statistics for diet in Table 1.3. Treatment households were more likely to consume a wide range of healthy food items in November of 1999 than control households. Some examples of the larger differences are chicken (57.4% of treatment versus 50.5% of control), eggs (88.1% versus 82.9%), oranges (50.2% versus 42.2%), and bananas (50% versus 43.9%).

Table 1.2: Percent of Men and Women Ages 16-70 Receiving Market and Transfer Income by Year and Source (Inflation Adjusted Mean Peso/Week Amounts in Parentheses)

	1997		1999		2000	
	Men	Women	Men	Women	Men	Women
N	22634	22807	22404	22094	21126	21584
Paid Market Activity						
Ag. Laborer	51.6 (130.5)	2.8 (105.40)	60.5 (123.61)	2.6 (92.73)	52.2 (138.19)	1.8 (118.16)
Non-Ag. Laborer	10.2 (275.26)	4.9 (177.97)	8.9 (233.22)	3.4 (159.5)	7.6 (259.35)	2.7 (176.01)
Entrepreneur	10.6 (139.87)	3.7 (87.83)	8.2 (58.3)	1.5 (61.29)	7.9 (150.84)	2.6 (84.43)
Ejiditario	8.5 (98.97)	0.1 (43.34)	6.1 (70.65)	0.1 (36.08)	3.3 (127.85)	0.1 (83.47)
Manager	0.2 (214.67)	0.2 (168.85)	0.2 (177.74)	0.1 (134.55)	0.2 (224.43)	0.1 (181.19)
Other	0.6 (115.24)	0.2 (94.1)	0.2 (114.87)	0.2 (60.64)	0.3 (268.09)	0.1 (108.58)
Additional Job	5.9 (47.50)	0.9 (40.17)	2.7 (32.96)	0.5 (35.48)	0.9 (112.94)	0.4 (46.93)
Sold Products in Market	2.6 (107.90)	0.5 (101.19)	4.0 (110.33)	0.8 (75.68)	1.2 (77.27)	0.3 (79.17)
Unpaid Market Activity						
Family Work (No Pay)	7.9	3.9	3.4	1.4	0.4	0.1
Work (No Pay)	0.2	0.2	0.1	0	0.1	0
No Participation	10	84	12.3	90.3	27.9	92.4
Transfers						
Procampo	22.9 (34.65)	0.9 (27.47)	18.0 (32.44)	0.9 (28.86)	16.3 (25.26)	1.1 (34.14)
Progresa	0	0	0	19.7 (63.50)	0	18.5 (74.51)
Money from Neighbors	2.1 (119.27)	1.7 (106.28)	0.9 (103.37)	2.0 (131.26)	0.9 (99.47)	2.4 (119.76)
Pension	0.3 (222.36)	0	0.3 (189.13)	0	0.2 (212.96)	0
Disability Payment	0.1 (225.60)	0	0.1 (146.33)	0	0.1 (112.23)	0.2 (133.04)
Property Rents	0.4 (62.69)	0	0.1 (125.15)	0	0	0
Scholarship	0.2 (33.98)	0.1 (30.71)	0	4.3 (26.88)	0	0
Bank Interest	0.1 (63.86)	0	0	0	0	0
Other Transfer	0.7 (109.75)	0.1 (77.73)	0.2 (128.13)	0.4 (27.31)	0.5 (96.03)	0.2 (110.75)
Gov Credit Program	0	0	1.5 (27.02)	0.8 (27.77)	0.5 (22.45)	0.1 (16.24)

Table 1.3: Percent of Households Consuming Each Food (with Mean Days/Week Frequency)

Names	1997	1999 Control	1999 Treatment	2000
Tomatoes	89.5 (4.27)	95.4 (5.67)	97.5 (5.70)	97.2 (5.70)
Onions	90.1 (4.73)	95.6 (5.89)	97.3 (6.02)	96.8 (5.88)
Potatoes	67.5 (1.54)	55.4 (1.82)	61.7 (1.98)	62.1 (1.87)
Carrots	19.7 (0.33)	7.20 (0.19)	8.00 (0.21)	8.60 (0.20)
Leafy Greens	25.7 (0.42)	5.90 (0.15)	6.30 (0.14)	15.6 (0.39)
Oranges	74.0 (2.76)	42.2 (1.67)	50.2 (2.05)	54.5 (2.35)
Bananas	61.1 (1.52)	43.9 (1.28)	50.0 (1.44)	51.7 (1.45)
Apples	23.8 (0.44)	12.7 (0.31)	15.9 (0.40)	15 (0.35)
Limes	49.9 (1.62)	34.0 (1.57)	37.3 (1.66)	32.1 (1.33)
Chicken	53.5 (0.69)	50.5 (0.74)	57.4 (0.86)	61.9 (0.90)
Beef & Pork	33.7 (0.42)	23.9 (0.34)	28.7 (0.42)	22.9 (0.31)
Eggs	89.0 (3.22)	82.9 (3.54)	88.1 (3.72)	89.7 (3.68)
Milk	41.7 (1.71)	23.8 (1.17)	23.3 (1.12)	27.7 (1.26)
Fish	11.6 (0.15)	1.40 (0.02)	1.40 (0.02)	1.70 (0.02)
Canned Tuna	26.7 (0.34)	10.0 (0.15)	12.3 (0.19)	7.40 (0.1)
Lard	18.4 (0.87)	13.2 (0.86)	14.8 (0.93)	11.1 (0.69)
Tortialls	89.2 (6.17)	88.8 (6.15)	87.1 (6.03)	88.9 (6.13)
Corn Flour	64.6 (4.25)	28.7 (1.87)	33.6 (2.23)	30.6 (1.98)
White Bread	40.8 (1.07)	14.1 (0.48)	15.6 (0.51)	12.1 (0.38)
Mexican Patries	61.0 (1.52)	34.3 (1.05)	41.3 (1.30)	32 (0.95)
Wheat Flour	14.5 (0.36)	5.00 (0.21)	6.00 (0.26)	3.70 (0.15)
Cup Noodles	80.9 (1.9)	70.8 (1.85)	76.3 (2.07)	68.5 (1.71)
Rice	77.1 (1.36)	64.3 (1.72)	67.9 (1.85)	72.0 (1.8)
Biscuits	43.3 (1.02)	10.1 (0.31)	13.2 (0.42)	12.6 (0.39)
Beans	92.7 (5.07)	96.9 (5.92)	96.8 (5.80)	96.9 (5.68)
Breakfast Cereals	3.20 (0.09)	1.40 (0.08)	1.60 (0.09)	1.60 (0.08)
Pastries	4.90 (0.09)	0.40 (0.01)	0.5 (0.01)	0.4 (0.01)
Soda	47.1 (1.03)	18.1 (0.41)	20.2 (0.44)	26.3 (0.54)
Alcohol	6.70 (0.15)	2.00 (0.04)	2.00 (0.04)	1.60 (0.05)
Coffee	87.6 (5.44)	70.6 (4.7)	73.7 (4.88)	67.3 (4.35)
Sugar	91.3 (5.93)	97 (6.64)	97.4 (6.67)	94.2 (6.29)
Vegetable Oil	84.9 (5.42)	88.7 (6.02)	89.7 (6.07)	88.8 (5.89)

1.4.2 The Relevant Outside Option Does Not Include Divorce

The relevant outside option in this context is not divorce, but the inefficient equilibrium within the family. In the Progresa sample, fewer than 1% of individuals reported having had a divorce. Mexico is one of the most Catholic countries in the world, and the Catholic tradition places stigma on divorce. The low levels of divorce in the Progresa sample data could be attributed to the large stigma attached to the practice.

In addition, we know about cultural practices shaping household finance from detailed mixed methods studies (Benería and Roldán, 1987; and Casique, 2001). These studies document that control of household income, and the social ramifications of this control, is gendered. Benería and Roldán document men's ability to hide income, which roughly corresponds to choosing to a family being in the outside option allocation in this framework. Benería and Roldán write on page 119 that "husbands' different and better-paid class position outside the household is translated into a commanding position within the family/household context." They report evidence that the men in more than half of the families in their sample (from Mexico City) hide income from their partner. Casique (2001) reports similar modes of bargaining in poor, rural families.[19]

1.4.3 Estimating Bargaining Power Levels for Each Household

Because most women have zero earnings, we use a Heckman selection model to predict earnings in the case of the outside option. This will not always be the best way to predict outside options, but in this case it is critical to model the unobservable differences between individuals who earn some positive amount, and those who do not. The identification results from section 1.3 do not depend on the prediction method used.

We fit a probit model on the decision to earn some amount from the market activities in Table 1.2. Then, we estimate the differences between those who work and those who do not in our earnings regression using the inverse Mills ratio. The probit is given in equation (1.11) and the earnings regression is given in equation (1.12). We estimate these using a Full Information Maximum Likelihood approach.

[19]She writes on page 30, "Particularly among low-income women, the idea of the man as the main breadwinner and family authority is quite clear (García and Oliveria, 1994; De Barbieri, 1984) and this socially extended conception heavily determines their interpretations about their own rights and obligations. The implications of this fact are considerable if we consider that the possibility of translating women's empowerment [outside the home] to relatively more power in the household may depend on whether or not women see female power as legitimate."

$$L_{i,t} = X^p \beta^p + \epsilon_{i,t} \tag{1.11}$$

$$ln(E_{i,t}) = X^e \beta^e + \epsilon_{i,t} \tag{1.12}$$

Variables L and E are a market participation dummy and continuous earnings, respectively. L is equal to one if people responded that they worked last week, or sold items at the market last week, or worked for the family business last week, or made products to be sold, or did chores for money for another family, or that they have a job but didn't work last week, or worked on the family farm with compensation. L is equal to zero if respondents indicated that they did not work and do not have a job, or worked for the family without pay.[20] The units for the left hand side of equation (1.12) are the natural log of pesos/week. We convert back to pesos/week for our predicted values. We predict individuals earnings, not wages, so that we thereby avoid having to estimate how much a person would work if they decided to earn income from any of the above sources.

Subscript i indexes individuals and t indexes time. The covariates in (1.11) include age, age squared, other unearned income like pensions and bank interest, number of children ages 5-16, number of children under the age of five, number of years of education, a literacy dummy, number of male and female adults in the household, village-specific migration statistics, the household poverty index calculated by the Government of Mexico, interactions between family composition and migration variables, and state-year fixed effects. The covariates included in the earnings regression are summarized by X^e and include many of the same variables. To satisfy the exclusion restriction, we drop some of the village level variables from X^e. These include the percent of women in a village who are able to leave the home without permission from their partner, percent of women in a village that require accompaniment if they leave the home, and the percent of individuals in a village who think women should be able to work outside of the home.[21] Except for

[20]See Parker and Skoufias (2000) for a detailed description of individuals' time use in the Progresa experimental data. In general, unpaid family work is farm labor, and is often seasonal.

[21]We only have data for the first and last waves for the last variable dropped to meet exclusion restriction.

the fixed effects, we display all variables and estimates in Tables 1.7 and 1.8 in Appendix D. We estimate the selection and outcome equations using two subsets of the data, one including only women ages 16-65 and one including only men ages 16-65. We discuss these regression results in Section 1.5, though the estimates from (1.11) and (1.12) are nuisance parameters.

Since we observe the requirements for various transfers (e.g. the Progresa transfer requires that parents have children of certain ages) we set the predicted values of the transfers equal to the observed values of the transfers, and denote them as $\hat{\tau}_f$ and $\hat{\tau}_m$. These are, for each partner, the sum of their unearned income from sources like Progresa, Procampo (an agricultural subsidy given almost entirely to men), bank interest, pensions, land rent, and other sources.

We let the predicted outside option earnings be the sum of individuals' predicted labor market incomes and predicted transfers. We construct our household bargaining power point estimate using the two household heads' outside option values,

$$\hat{\eta} = \frac{1}{2} + \frac{1}{2}\frac{(X_f\hat{\beta}_f + \hat{\tau}_f) - (X_m\hat{\beta}_m + \hat{\tau}_m)}{y}.$$

The result is a single measure per household per period. We present the distributions of bargaining power over time in Figure 1.1. We plot the control and treatment groups separately to emphasize the effect of treatment. Both groups start with medians close to 0.20. This means that the median female decision maker has one fourth the amount of say that their male partner does. The control group's median stays stable over time while the treatment group's median increases to 0.42. This increases the median female decision maker's relative decision making power to roughly seventy percent of that of their male partner. Eligible households in control regions were phased into the program in 2000, shifting the treatment distribution again. This randomized, transfer-induced variation over time within households allows us to estimate the causal marginal effects of a change in power on family diet.

In Table 1.5 in Appendix C, we present summary statistics on household decision making patterns, private assignable good expenditures, and views on women's rights. These

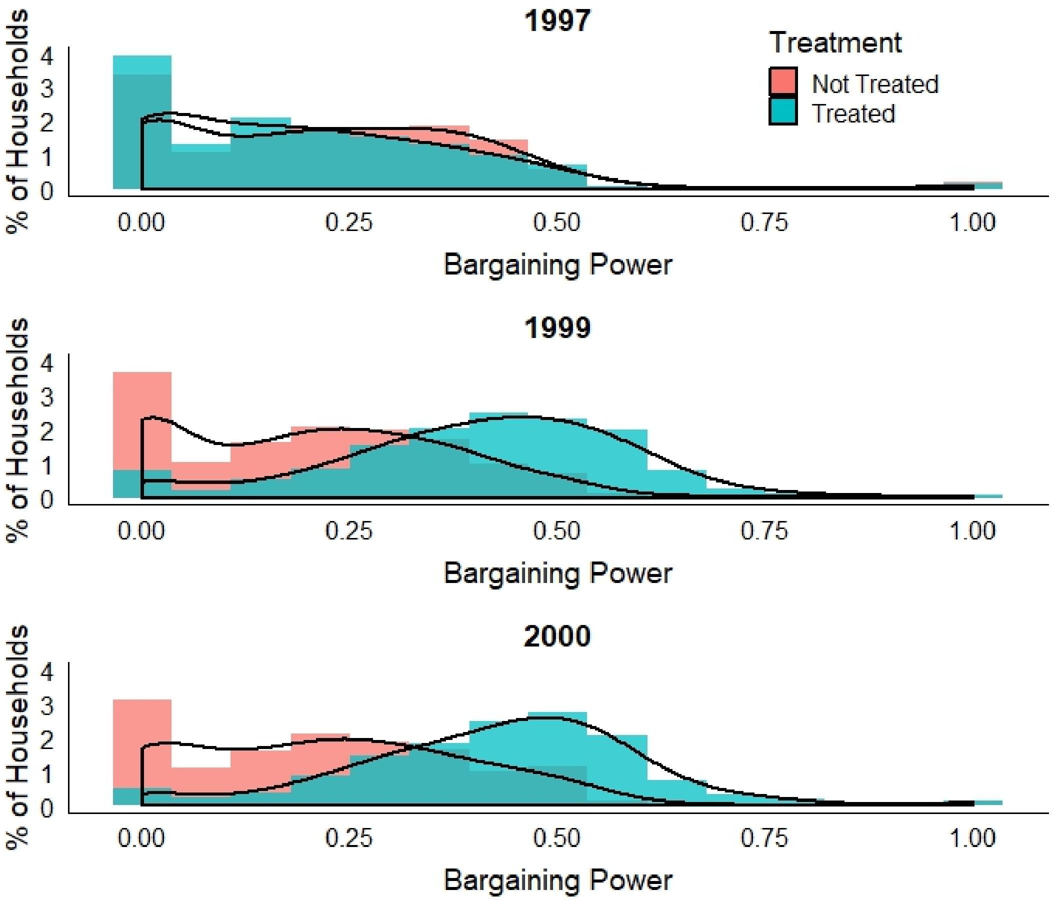

Figure 1.1: Distribution of bargaining dynamics, $\hat{\eta}$, over time across all households. We use a Gaussian kernel density with bin-width equal to four times the standard deviation of the distribution (Wickham, 2016).

summary statistics provide a rich backdrop for understanding women's bargaining power in this setting. Particularly telling are the facts that one third of women needed permission from their partners to leave the home in 1997, and that number increased to 45% later. Though couples tended to make decisions together, views on women's rights favored men. For instance, 77% of respondents in 1997 agreed with the statement "Women should be obedient to men."[22] These data suggest that men had substantially more bargaining power than women, corroborating the results in Figure 1.1.

1.4.4 Hypotheses and Testing

With the level estimates of bargaining power for each household and in each period, we can begin testing hypotheses. First, we formalize the visual evidence in Figure (1.1) by testing the null hypothesis that Progresa did not change power in the treatment group against the alternative hypothesis that it did. We calculate a simple difference-in-difference estimator and use a block bootstrap (clustered at the household level) algorithm to generate a confidence interval on this average treatment effect estimate. Our second and third hypotheses are on whether Progresa's empowerment effect subsequently improved family diet.

We exploit Progresa's randomization to construct our difference-in-differences estimator. Because of the randomization, we can be confident that the parallel trends assumption required for identifying the causal marginal effect holds.[23] We examine the treatment and control groups in the baseline and the second wave, since control group eligible families were phased into the program in the third wave. The means of the control distributions in 1997 and 1999 were $\bar{\hat{\eta}}^c_{1997} = 0.225$ and $\bar{\hat{\eta}}^c_{1999} = 0.203$. The means of the treatment distributions were $\bar{\hat{\eta}}^t_{1997} = 0.197$ and $\bar{\hat{\eta}}^t_{1999} = 0.406$. We can use these to construct a classic difference-in-difference estimator, $\hat{\delta} = (\bar{\hat{\eta}}^t_{1999} - \bar{\hat{\eta}}^t_{1997}) - (\bar{\hat{\eta}}^c_{1999} - \bar{\hat{\eta}}^c_{1997}) = 0.231$. We test whether this estimator is statistically different from zero at the 95% level using a block bootstrap algorithm to construct a percentile-t confidence interval. If the confidence interval does not cover zero, we reject the null hypothesis that the program had no effect

[22] The survey reads: "Las mujeres deben obedecer a los hombres: Acuerdo, Desacuerdo, o No Sabe."

[23] See Behrman and Todd (1999) for balance tables. They fail to reject the null that control and treatment localities are the statistically indistinguishable along many observable dimensions.

on women's bargaining power in the family.

The remainder of our hypotheses focus on how this change in bargaining power impacted household allocations. First, we test whether the marginal effect of a change in power on diet at the extensive margin is statistically different from zero. That is, we examine whether families with more equal power dynamics are more likely to eat healthy foods at all. Second, we examine the intensive margin: do families with more equal decision makers eat healthy foods more times per week? We run two demand regressions for each of the thirty-one food types that we have panel data on. We examine the extensive and intensive margins using an OLS estimation and a Poisson regression.[24] We choose to examine demand for each food item, instead of grouping the foods, because this approach is more informative, and because it enables analysis at the extensive margin. See Table 1.3 for the set of food items listed according to the categories “fruits and vegetables," “animal products," “pulses and grains," and “miscellaneous." These demand regressions are given by regressions (1.13) and (1.14):

$$Y^b_{h,t} = \beta^b_1 \hat{\eta}_{h,t} + \beta^b_2 X^b_{h,t} + \beta^b_3 \xi_h + \beta^b_4 \zeta_t + \epsilon_{h,t} \tag{1.13}$$

$$Y^c_{h,t} = \beta_1 \hat{\eta}_{h,t} + \beta_2 X_{h,t} + \beta_3 \xi_h + \beta_4 \zeta_t + \epsilon_{h,t} \tag{1.14}$$

The dependent variable in equation (1.13) is a dummy variable for whether or not a family consumed a food item in the past week. The dependent variable in equation (1.14) is the number of days per week that a family consumed the item.[25] The subscript h denotes the household and t denotes period, as before. In both equations, $\hat{\eta}$ is our estimate of bargaining power. The matrix $X_{h,t}$ is a set of control variables including log household earnings (which includes the Progresa transfer and other government transfers), number of household kids ages 5-16, and number of children under the age of 5, own price,

[24]With a total of 62 demand regressions and 62 primary hypotheses tested, the expected number of false positives is 3.4 using a 95% confidence interval.

[25]The superscript b denotes that equation (1.13) uses a binary dependent variable and disambiguates between the extensive and intensive margin investigations. We use a superscript c in equation (1.14) to disambiguate between total household income Y and the number of days per week that a food is consumed.

substitute prices, complement prices,[26] and staple prices.[27] We use locality specific prices when available. When they are not available, we construct a municipality level average. In the rare cases that no municipality average is available, we use state level averaged prices. The inclusion of the household earnings soaks up the income effect associated with Progresa.

The primary feature of these regressions are the household, ξ_h, and year, ζ_t, fixed effects which force all identifying variation to come from changes within the household over time. Our results are based on households that experienced a change in bargaining dynamics and a change in consumption patterns. These fixed effects also control for time-invariant household idiosyncrasies, perhaps including preferences. The time fixed effect controls for period specific factors like inflation.

With this approach, the marginal effect of power on diet at the extensive margin is $\hat{\beta}_1^b$. At the intensive margin, the marginal effect is $M \equiv \frac{1}{H}\sum_1^H(\hat{\beta}_1 e^{X_H\hat{\beta}})$, the derivative of the Poisson regression with respect to $\hat{\eta}$ evaluated at the means of the control variables. We test our second and third hypotheses with block bootstrapping since $\hat{\eta}$ is an estimate, not an observed value.

We can unpack this effect further. Denote treatment as z. Then, average treatment effect of Progresa on diet, via it's effect on power, is $\frac{\partial Q^*}{\partial z} = \frac{\partial Q^*}{\partial \eta}\frac{\partial \eta}{\partial z}$. Our difference-in-difference estimator, $\hat{\delta}$, gives an estimate of $\frac{\partial \eta}{\partial z}$. Our variables M^b and $\hat{\beta}_1^b$ estimate $\frac{\partial Q^*}{\partial \hat{\eta}}$. As such, we can generate estimates of $\frac{\partial Q^*}{\partial z}$ by multiplying $\hat{\delta} \times \hat{\beta}_1^b$. By comparing these values to simple t-tests of the Program effects on diet at the extensive margin, we can discern the magnitude of the empowerment effect relative to the overall effect. There are two primary design-driven ways that Progresa impacted diet: via empowerment and via income. As such, we can demonstrating the importance of empowerment in explaining demand relative to the income effect.

[26]We include all other food items that come from the same food group. For fruits and vegetables, we include prices for onions, tomatoes, oranges, potatoes, bananas, limes, leafy greens, and apples. For pulses and grains, we include tortillas, cup noodles, rice, digestive biscuits, white bread, wheat flour, and breakfast cereals. For animal products, we include prices for eggs, chicken, milk, beef/pork, tuna/sardines, and lard. For the miscellaneous food items we include prices for sugar, coffee, soda, vegetable oil, and cup cakes. We do not have price data for alcohol, lamb/goat, carrots, corn flour, and fish.

[27]For all regressions we include prices for milk, beans, rice, and eggs.

1.5 Results: The Magnitude of Progresa's Empowerment and the Nutritional Ramifications

We confidently reject all three of our null hypotheses. Progresa's income transfers increased women's bargaining power by roughly 100% from 1998 to 1999. These increases positively and significantly affect the probability and frequency of consuming many food items, with the largest effects occurring for healthy foods, such as fruits, vegetables, and proteins. We first discuss the earnings estimations for males and females and then the constructed bargaining power estimates. Then, we examine the estimates from equations (1.13) for four example foods, chicken, eggs, leafy greens, and bananas. We bootstrap the $\hat{\beta}_1$ coefficients from these regressions to build the confidence intervals used in hypothesis testing.

Female and male earnings estimation results are reported in Tables 1.7 and 1.8. While the estimates from (1.11) and (1.12) are nuisance parameters, assessing whether the parameter estimates are sensible is useful in analyzing the quality of the model predictions, and thus the quality of the bargaining power estimates. A perusal of the coefficient estimates in both stages reveals outcomes consistent with similar Mincerian participation-earnings models. For example, earnings are positively associated with age (for women), education, and older (but not younger) children.[28] Being a household head decreases the amount earned in the market, likely by reducing the number of hours worked. The coefficient estimate on the inverse Mills ratio is positive and significant for women, which is consistent with other two stage models that find the unobserved characteristics that positively influence market participation choices also positively affect earnings estimates. This indicates that employment is an important determinant of female bargaining power; we include the inverse Mills ratios in predicted earnings. The distribution of predicted earnings for all women who work in all years (corrected for inflation) has median 111.75 pesos/weeks and standard deviation of 75.91 pesos/week. The values for men are 153.1 and 72.95 respectively.

[28]See Mincer (1974 and 1975).

Our first estimate of bargaining power derived from these estimates precedes treatment. At that time, the median bargaining power estimate for sample households was 0.216. We generate a percentile-t confidence interval for the difference-in-differences estimator, $\hat{\delta} = 0.231$, of [0.203, 0.242]. We strongly reject the null hypothesis of no average treatment effect. Progresa empowered women by a substantial amount.[29]

This increase is due to the government's gender targeting. Figures 1.4 and 1.5 in Appendix B shows what would have occurred under two counterfactual transfer patterns, one where the transfer was split evenly between spouses by the government and one where the transfer went completely to men. In a hypothetical context where the transfers had been given to men, women's bargaining power would have decreased. In the hypothetical with the split transfer, women's power increased but to a lesser extent. These different hypothetical situations give some insight into contexts where women's bargaining power decreases. They are relevant, despite being hypothetical, since some of Mexico's transfers go almost entirely to men, like Procede and Procampo. The counterfactuals in Appendix B provide intuition for the impact of those programs on women's bargaining power.

We present demand regression results for chicken, eggs, leafy greens, and bananas in Table 1.6 in Appendix D. The income coefficient estimates are generally positive and significant. The same price coefficients (not displayed) estimates are negative. They are not always significant, because the household and time fixed effects absorb the prices effects to a large degree. The coefficient estimate for the effect of bargaining power on consumption is positive. These results are consistent with rejecting hypothesis 2 at the intensive margin of no significant effect of bargaining power on dietary outcomes. All controls and standard errors in 1.6 are nuisance parameters.

We generate marginal effect estimates from these thirty-one demand regressions and report them in Figure 1.2 along with their 95% confidence intervals. A change in bargaining power has a positive, large, and statistically significant effect on the consumption of many food items, especially healthy ones like chicken, beef and pork, milk, carrots,

[29] Adato et al (2000) present summary statistics from Progresa's decision making module, meant to shed light on household bargaining dynamics. Our estimates are intuitive given those statistics. It is plausible that women have one fourth the say that their husbands given the decision making patterns in the Progresa data. We comment further in Appendix B.

green leaf vegetables, oranges, bananas, and apples. For example, the marginal effect of a bargaining power increase on chicken is 0.09 for the probability of consumption, and 0.15 for the number of days per week that the food is consumed. This translates to an increase at the extensive margin that families consumed chicken because of Progresa's empowerment effect of 2%, and at the intensive margin of 1.6 days per year. Overall, the marginal effect of power on diet at the extensive margin is significantly different from zero for eighteen of thirty-one foods, and at the intensive margin for seventeen foods. Individually, the effects of Progresa on diet via increased gender equality may seem small, but jointly across healthy foods, the effects are large.

We break down the total program effects at the intensive margin into the income and empowerment components in Table 1.4. Progresa changed the likelihood that families consume 22 foods, which we list in column 1. In columns 2 and 3, we present the likelihood that the treatment and control groups consumed each of these foods in 1999. In columns 4 and 5, we give the differences in the means and the t-statistic on the difference. In column 6 we report $\hat{\delta} \times \hat{\beta}_1^b$ and in column 7 we report the empowerment effect divided by the total effect. If the confidence intervals do not cover zero for the marginal effect of power on extensive demand for a food, we calculate the values in columns 6 and 7. If the confidence interval for a food covers zero, we say that the total program effect is entirely owning to the income effect.

For 13 of these foods, we find that the change in power mediated the total change. Overall, we see that 14% of Progresa's influence on the extensive-margin demand for animal products, and 6% of the influence on extensive-margin demand for fruits and vegetables, was caused by Progresa's empowerment effect. We also see that empowerment caused families to consume less tomatoes and onions, while the income effects went the other directions. These foods are very commonly consumed so this suggests that empowerment effects trend towards more balanced diets. Families with relative gender equality are more likely to substitute staples like onions for less common healthy foods like apples and oranges.

1.5.1 Caveats

Even after controlling for household and year specific factors in our analysis, there may be factors that we are unable to model that vary across time within the family, and which influence demand. In particular, Progresa had several conditions that we do not model. Some of the program effect that we've currently attributed to changes in power and in income could be, in reality, owing to these conditions. These conditions include sending your child to school, routinely visiting the doctor for health check ups, and attending health classes. In addition, infants and pregnant mothers received health supplements to be consumed as porridge or smoothies. These conditions and additional treatment elements might influence demand for various foods in ways that we do not capture with our fixed effects models.

It is possible that we have attributed some of the program effects on food demand that are caused by the conditions to the changes in power and income. Both of the marginal effects that we estimate may be biased upward. Which estimate is more biased is unclear. We suggest, as an avenue for future research, that economists model the relationships between income, power, and program conditions more completely, as in see Baird, McIntosh, and Özler (2011).

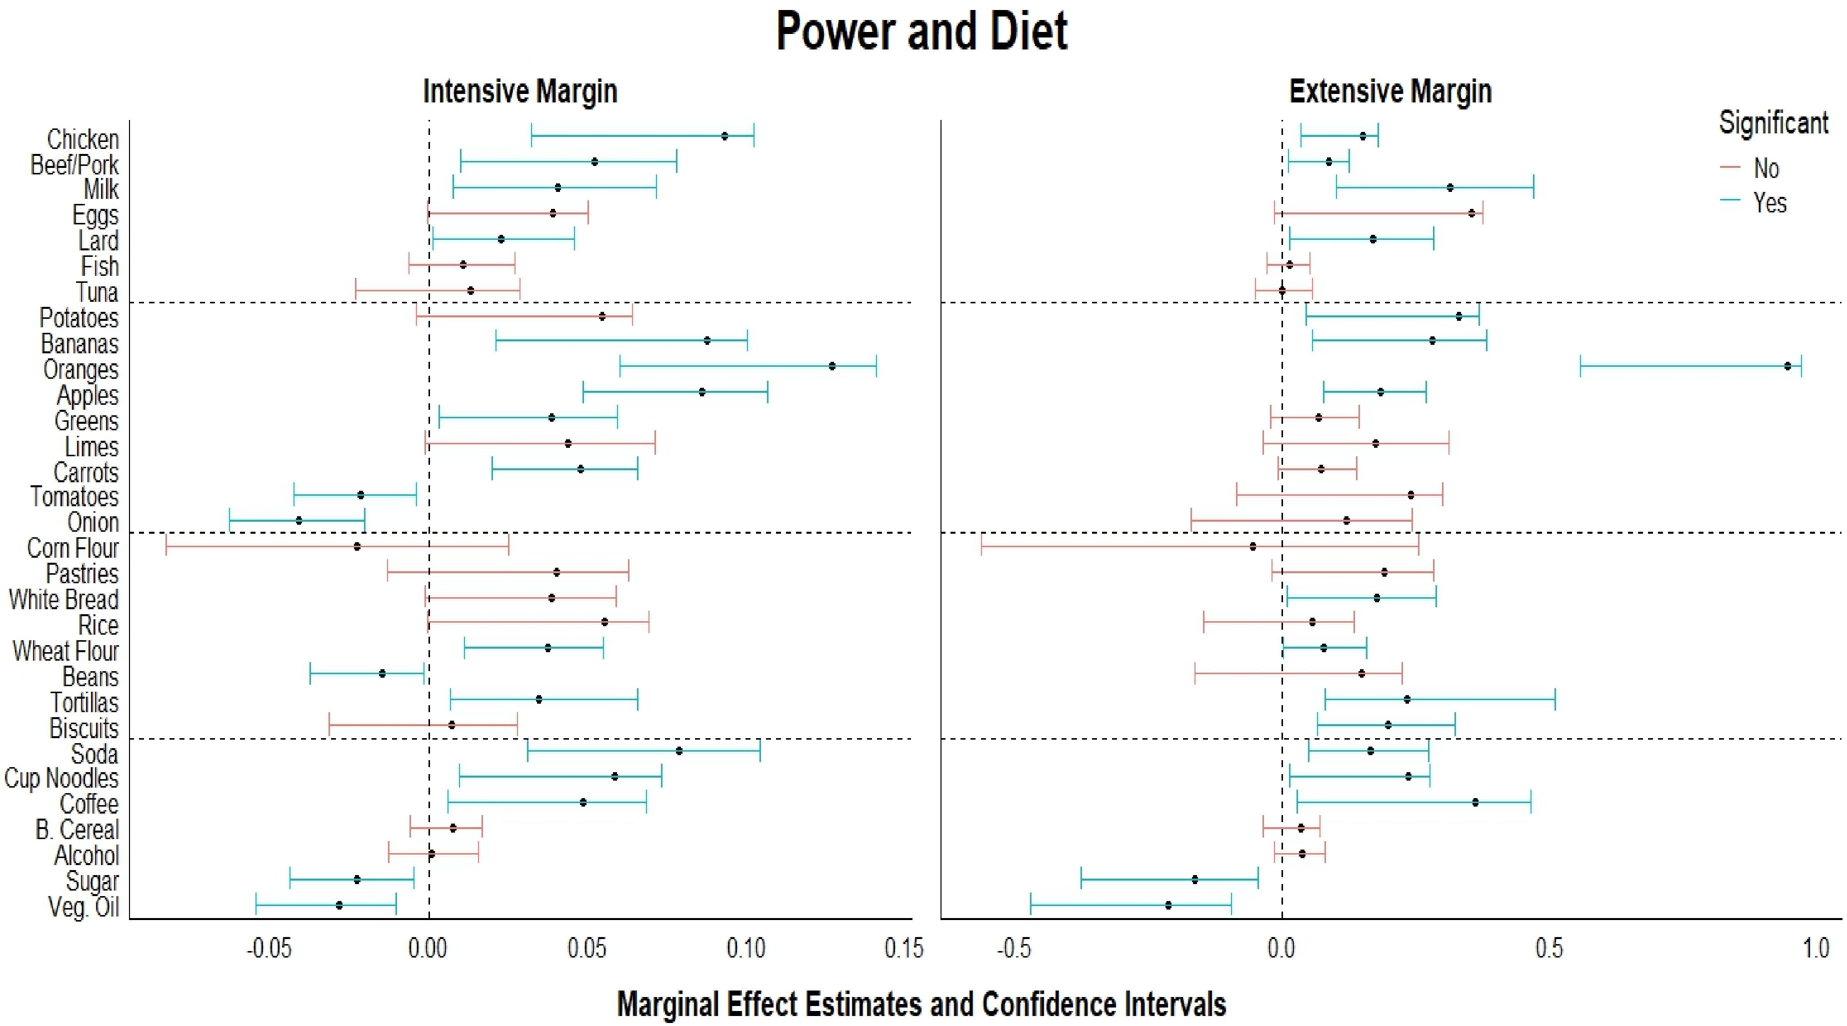

Figure 1.2: A Marginal Increase in Bargaining Power Increases Net Family Consumption of Healthy Foods

Table 1.4: Percent of Progresa's (Experimental) Effect Explained by the Change in Power

	1999 Treatment	1999 Control	Diff	t-Stat	$\hat{\delta} \times \hat{\beta}_1^b$	% Power
Animal Products						
Chicken	57.4	50.5	6.9	6.966	2.124	30.787
Beef/Pork	28.7	23.9	4.8	5.418	1.193	24.855
Lard	14.8	13.2	1.6	2.323	0.523	32.716
Fruits & Vegetables						
Tomatoes	97.5	95.4	2.1	6.098	-0.488	-23.238
Onions	97.3	95.6	1.7	4.793	-0.929	-54.649
Oranges	50.2	42.2	8	8.038	2.898	36.23
Bananas	50	43.9	6.1	6.209	1.998	32.757
Apples	15.9	12.7	3.2	4.585	1.962	61.313
Other						
Tortillas	87.1	88.8	-1.7	-2.704	0.789	-46.434
Wheat Flour	6.1	5.0	1	2.116	0.86	85.969
Cup Noodles	76.3	70.8	5.5	6.34	1.339	24.352
Soda	20.2	18.1	2.1	2.747	1.797	85.562
Coffee	73.7	70.6	3.1	3.461	1.111	35.826
Pure Income Effect						
Potatoes	61.7	55.4	6.3	6.466	0	0
Tuna	12.3	10	2.3	3.633	0	0
Eggs	88.1	82.9	5.2	7.746	0	0
Limes	37.3	34	3.3	3.446	0	0
Corn Flour	33.6	28.7	4.9	5.273	0	0
White Bread	15.6	14.1	1.5	2.152	0	0
Pan de Dulce	41.3	34.3	7	7.16	0	0
Rice	67.9	64.3	3.6	3.852	0	0
Biscuits	13.2	10.1	3.1	4.836	0	0

Note: The values in columns 2, 3, 4, 6, and 7 are all percents. The zeroes in the Pure Income Effect section, columns 6 and 7, reflect that the confidence intervals for the marginal effect of an increase in women's bargaining power in the household for those foods covers zero.

1.6 Discussion

Two more comments are worth making here. First, it is possible that the model outlined in this paper explains the puzzling results documented in Hoehn-Velasco and Penglase (2019). They find that Mexico's introduction of no-fault divorce (from previouslyheld stricter divorce legislation) had no effect on women's resource shares. This is puzzling since this type of legislation is typically thought to improve women's bargaining positions. Voena (2015) documents a positive effect of the introduction of no-fault divorce in the USA. The differences in outcomes from similar policies in Mexico and the United States presents a puzzle.

We suggest a possible solution to this puzzle. In the USA, both before and after the introduction of no-fault divorce, the relevant outside option was divorce (not an inefficient allocation within the family). As such, this policy improved the value of women's outside option values and increased their bargaining power. In Mexico, the introduction of no-fault divorce, in contrast, changed the relevant outside option from an inefficient equilibrium in the family to divorce. As such, new factors became relevant in determining power, captured by the parameter γ. Even if the value of $y_f^0 - y_m^0$ increased because of the policy, women's bargaining power could decrease overall if γ had a larger absolute value than this increase.

Testing this hypothesis in future research can help policy makers predict the effects of divorce legislation. If they find that, as in Mexico, this legislation may reduce women's bargaining power, they may consider pairing it with counter-acting policies. Such policies would necessarily help women in the case of divorce.

Second, this model is compatible with the DLP estimation method. The same resource shares that solve the constrained optimization problems in section 1.3 should obtain from the DLP family-specific systems of Engel curves for private assignable goods, if the additional assumptions (PIGLOG form, additional similarities across types of households) on utility functions do not introduce model-misspecification error. This is an interesting avenue for future research, since estimating Engel curves at the same time as (1.5) (or (1.10) depending on the context-dependant definition of the outside option) could provide more stable estimates. The DLP method has been shown to give different estimates depending on which private assignable good the researcher uses (Bargain, Lacroix, and Tiberti, 2018). It could be that adding in outside options and participation constraints to the DLP model solves this problem.

1.7 Conclusion

We derive two new methods to estimate household bargaining power. Identification and estimation rely on recovering differences between the shadow prices and shadow incomes that individuals face in the collective allocation and in their outside option. We build on a rich literature connecting the outside option to bargaining power to derive our estimators (Nash, 1950 and 1953; Basu, 2006; Mazzocco, 2007; Ligon, 2011; etc.). The primary benefit of our approach is that researchers can achieve point identification using assignable income data while avoiding strong utility functional form assumptions.

The steps in estimation are as follows: first, define the outside option for each family member as a non-cooperative household equilibrium or divorce, depending on which is a more credible threat in the study context. Second, predict each partners' earnings in the outside option, possibly following our example and using data from labor markets, rent, transfers from friends and family, government programs, and other income generating opportunities. Third, depending on which estimator is appropriate, construct the measure of bargaining power using the functional forms we provide in equations (1.5) and (1.10). If estimating (1.10), use a constrained fixed effects approach to estimate (1.9), and use the resulting regression coefficients to construct the bargaining power estimate.

We apply this measurement strategy to shed light on a twenty year old policy question of great importance. We show that the median female decision maker's bargaining power before Progresa was 0.20, and that it rose to 0.42 because of Progresa's cash transfer. This transfer increased women's income share in the family by 1000%, and doubled her decision making power. Because of this shift in household power, household demand for healthy food items increased at the intensive and extensive margins.

In light of the recent decision by the Government of Mexico to cancel Prospera (the latest name for Progresa), we hope that other researchers apply our method to the many other gender-targeted conditional cash transfers that exist across the world. Parker and Todd (2017) report that conditional cash transfer programs "have now been implemented in over sixty countries on five continents, ranging from among the poorest countries in the world, such as Malawi, to recent initiatives in developed countries including England and the United States." Our model can be applied to understand the distributional effects of these programs. We hope that such research is conducted prior to these governments deciding whether to cancel their programs. We would hypothesize based on evidence from Progresa that other gender targeted programs might also have large female empowerment effects, which could in turn increase investments in children through a variety of mechanisms including nutrition.

We see our approach to measuring bargaining power as a complement to existing approaches like those developed by Dunbar, Lewbel, and Pendakur (2013 and 2017), Chiappori and Kim (2017), Cherchye et al (2015 and 2017), and Almås et al (2018). By approaching the identification problem from a different angle, we hope we have broadened the set of tools available to researchers studying the household. The primary difference between our method and those that recover resource shares from demand data is that we do not assume a functional form for utility in estimation.

The primary disadvantage of our method is that researchers must take a stand on what the outside option is. In our application, we assume that the consumption technology a family uses is invariant to whether the participation constraints hold or not. This assumption would fail if, for example, women move in with their parents for some time as a response to intimate partner violence. Future research could relax the dichotomous nature of the estimators developed in this paper. In theory, a spectrum of outside options could be accommodated by the type of estimators we build here, with compensating variation values capturing reductions in the quality of the consumption technology in the inefficient equilibrium. Lewbel and Pendakur (2019) make progress in this direction, though they do not incorporate outside options into their analysis and rely on an instrumental variables approach for estimation. We suggest that researchers advance our framework in such a fashion in future analyses.

Reference List

Adato, M., De la Briere, B., Mindek, D., & Quisumbing, A. (2000). The Impact of PROGRESA on Women's Status and Intrahousehold Relations. Washington, DC: International Food Policy Research Institute.

Almås, I., Armand, A., Attanasio, O., & Carneiro, P. (2018). Measuring and Changing Control: Women's Empowerment and Targeted Transfers. The Economic Journal, 128(612), F609-F639.

Angelucci, M. (2008). Love on the Rocks: Domestic Violence and Alcohol Abuse in Rural Mexico. The B.E. Journal of Economic Analysis Policy, 8(1), pp. -. Retrieved 23 Jun. 2019, from doi:10.2202/1935-1682.1766

Apps, P. F., Rees, R. (1988). Taxation and the Household. Journal of Public Economics, 35(3), 355-369.

Ashraf, N., Karlan, D., & Yin, W. (2010). Female Empowerment: Impact of a Commitment Savings Product in the Philippines. World Development, 38(3), 333-344.

Attanasio, O., & Lechene, V. (2002). Tests of Income Pooling in Household Decisions. Review of Economic Dynamics, 5(4), 720-748.

Attanasio, O. P., & Lechene, V. (2014). Efficient Responses to Targeted Cash Transfers. Journal of Political Economy, 122(1), 178-222.

Baird, S., McIntosh, C., & Ãzler, B. (2011). Cash or condition? Evidence from a cash transfer experiment. The Quarterly journal of economics, 126(4), 1709-1753.

Bargain, O., Lacroix, G., & Tiberti, L. (2018). Validating the Collective Model of Household Consumption Using Direct Evidence on Sharing. Partnership for Economic Policy Working Paper, (2018-06).

Basu, K. (2006). Gender and Say: A Model of Household Behaviour with Endogenously Determined Balance of Power. The Economic Journal, 116(511), 558-580.

Becker, G. S. (1965), A Theory of the Allocation of Time. Economic Journal, 75, 493-517.

Behrman, J. R., & Hoddinott, J. (2005). Programme Evaluation with Unobserved Heterogeneity and Selective Implementation: The Mexican PROGRESA Impact on Child Nutrition. Oxford Bulletin of Economics and Statistics, 67(4), 547-569.

Behrman, J. R., & Todd, P. E. (1999). Randomness in the Experimental Samples of PROGRESA (Education, Health, and Nutrition Program). International Food Policy Research Institute, Washington, DC.

Benería, L., & Roldán, M. (1987). The Crossroads of Class and Gender: Industrial Homework, Subcontracting, and Household Dynamics in Mexico City. University of Chicago Press.

Blundell, R., Chiappori, P. A., Magnac, T., & Meghir, C. (2007). Collective Labour Supply: Heterogeneity and Non-Participation. The Review of Economic Studies, 74(2), 417-445.

Bobonis, G. (2009). Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment. Journal of Political Economy, 117(3), 453-503. doi:10.1086/600076

Bobonis, G. J., GonzÃ¡lez-Brenes, M., and Roberto, C. (2013) Public Transfers and Domestic Violence: The Roles of Private Information and Spousal Control. American Economic Journal: Economic Policy, 5 (1): 179-205.

Bobonis, G. J., Castro, R., & Morales, J. S. (2015). Conditional Cash Transfers for Women and Spousal Violence: Evidence of the Long-Term Relationship From the Oportunidades Program in Rural Mexico. Inter-American Development Bank.

Browning, M., Chiappori, P. A., & Lewbel, A. (2013). Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power. Review of Economic Studies, 80(4), 1267-1303.

Carlsson, F., Martinsson, P., Qin, P., & Sutter, M. (2013). The influence of spouses on household decision making under risk: an experiment in rural China. Experimental Economics, 16(3), 383-401.

Cherchye, L., Cosaert, S., De Rock, B., Kerstens, P. J., & Vermeulen, F. (2017). Individual Welfare Analysis for Collective Households.

Cherchye, L., De Rock, B., Lewbel, A., & Vermeulen, F. (2015). Sharing Rule Identification for General Collective Consumption Models. Econometrica, 83(5), 2001-2041.

Cherchye, L., de Rock, B., Walther, S & Vermeulen, F. (2018) Where Did it Go Wrong? Marriage and Divorce in Malawi. Available at SSRN:
https://ssrn.com/abstract=2759255 or http://dx.doi.org/10.2139/ssrn.2759255

Chiappori, P. A. (1988). Rational Household Labor Supply. Econometrica, 63-90.

Chiappori, P. A. (1992). Collective Labor Supply and Welfare. Journal of Political Economy, 100(3), 437-467.

Chiappori, P. A., & Kim, J. H. (2017). A Note on Identifying Heterogeneous Sharing Rules. Quantitative Economics, 8(1), 201-218.

Chiappori, P. A., & Mazzocco, M. (2017). Static and Intertemporal Household Decisions. Journal of Economic Literature, 55(3), 985-1045.

d'Aspremont, C., & Ferreira, R. D. S. (2014). Household Behavior and Individual Autonomy: An Extended Lindahl Mechanism. Economic Theory, 55(3), 643-664.

d'Aspremont, C., & Ferreira, R. D. S. (2019). Enlarging the Collective Model of Household Behavior: A Revealed Preference Analysis. Economic Theory, 68(1), 1-19.

De Barbieri, T. (1984) Mujeres y Vida Cotidiana, México, Fondo de Cultura Economíca (FCE) and Instituto de Investigaciones Sociales-UNAM.

Dosman, D., & Adamowicz, W. (2006). Combining Stated and Revealed Preference Data to Construct an Empirical Examination of Intrahousehold Bargaining. Review of Economics of the Household, 4(1), 15-34.

Doss, C. (1996). Women's Bargaining Power in Household Economic Decisions: Evidence from Ghana (No. 13517). University of Minnesota, Department of Applied Economics.

Doss, C. (2013). Intrahousehold Bargaining and Resource Allocation in Developing Countries. The World Bank Research Observer, 28(1), 52-78.

Donni, O., & Molina, J. A. (2018). Household Collective Models: Three Decades of Theoretical Contributions and Empirical Evidence.

Duflo, E. (2003). Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. The World Bank Economic Review, 17(1), 1-25.

Duflo, E. (2012). Women Empowerment and Economic Development. Journal of Economic Literature, 50(4), 1051-79.

Dunbar, G. R., Lewbel, A., & Pendakur, K. (2013). Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi. American Economic Review, 103(1), 438-71.

Dunbar, G. R., Lewbel, A., & Pendakur, K. (2017). Identification of Random Resource Shares in Collective Households Without Preference Similarity Restrictions. Bank of Canada.

Ewerling, F., Lynch, J. W., Victora, C. G., van Eerdewijk, A., Tyszler, M., & Barros, A. J. (2017). The SWPER Index for Women's Empowerment in Africa: Development and Validation of an Index Based on Survey Data. The Lancet Global Health, 5(9), e916-e923.

García, B. & Oliveria, O. (1994) Trabajo Femenino y Vida Familiar en México. México D.F., El Colegio de México.

Handa, S. Peterman, A., Davis, B., and Stampini, M. (2009) Opening Up Pandora's Box: The Effect of Gender Targeting and Conditionality on Household Spending Behavior in Mexico's Progresa Program. World Development, ISSN 0305-750X

Hanmer, L., & Klugman, J. (2016). Exploring Women's Agency and Empowerment in Developing Countries: Where Do We Stand?. Feminist Economics, 22(1), 237-263.

Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. Econometrica, 47: 153-161.

Hlavac, M. (2015). Stargazer: Well-Formatted Regression and Summary Statistics Tables. R package version 5.2. http://CRAN.R-project.org/package=stargazer

Hoddinott, J., & Haddad, L. (1995). Does Female Income Share Influence Household Expenditures? Evidence from CÃ´te d'Ivoire. Oxford Bulletin of Economics and Statistics, 57(1), 77-96.

Hoddinott, J., & Skoufias, E. (2004). The Impact of PROGRESA on Food Consumption. Economic Development and Cultural Change, 53(1), 37-61.

Hoddinott, J., Skoufias, E., & Washburn, R. (2000). The Impact of PROGRESA on Consumption: A Final Report. International Food Policy Research Institute, Washington, DC.

Jayachandran, S. (2015). The Roots of Gender Inequality in Developing Countries. Annual Review of Economics 2015 7:1, 63-88

Klein, M. J., McArthur, T. (2018) Gender and Land in Mexico. Working Paper, AgEcon Search Link: https://ageconsearch.umn.edu/record/274259

Lewbel, A., & Pendakur, K. (2019). Inefficient Collective Households: Abuse and Consumption.

Ligon, E. A. (2011). Dynamic bargaining in households (with an application to Bangladesh).

Lise, J., & Yamada, K. (2018). Household Sharing and Commitment: Evidence From Panel Data on Individual Expenditures and Time Use. The Review of Economic Studies, 86(5), 2184-2219.

Lundberg, S., & Pollak, R. A. (1993). Separate Spheres Bargaining and the Marriage Market. Journal of Political Economy, 101(6), 988-1010.

Lundberg, S. J., Pollak, R. A., & Wales, T. J. (1997). Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. Journal of Human Resources, 463-480.

Lundberg, S. J. & Ward-Batts, J. (2000) Saving for Retirement: Household Bargaining and Household Net Worth, Working Papers wp004, University of Michigan, Michigan Retirement Research Center.

Malapit, H. J., Sproule, K., Kovarik, C., Meinzen-Dick, R. S., Quisumbing, A. R., Ramzan, F., Hogue, E. & Alkire, S. (2014). Measuring Progress Toward Empowerment: Women's Empowerment in Agriculture Index: Baseline Report. International Food Policy Research Institute (IFPRI).

Manser, M., & Brown, M. (1980). Marriage and Household Decision-Making: A Bargaining Analysis. International Economic Review, 31-44.

Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and A Test of Commitment. The Review of Economic Studies, 74(3), 857-895.

McElroy, M. B., & Horney, M. J. (1981). Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand. International Economic Review, 333-349.

McFadden, D.(1974). Conditional Logit Analysis of Qualitative Choice Behavior. *Frontiers in Econometrics*, ed. by P. Zarembka, New York: Academic Press, 1974, 105-142.

Mincer, J. (1974). Schooling, Experience, and Earnings. Human Behavior & Social Institutions No. 2.

Mincer, J. (1975). Education, Experience, and the Distribution of Earnings and Employment: An Overview. In Education, Income, and Human Behavior (pp. 71-94). NBER.

Munro, A. (2018). IntraâHousehold Experiments: A Survey. Journal of Economic Surveys, 32(1), 134-175.

Nash, J. (1950). The bargaining problem. Econometrica: Journal of the Econometric Society, 155-162.

Nash, J. (1953). Two-person cooperative games. Econometrica, 128-140.

Parker, S., & Skoufias, E. (2000). The Impact of PROGRESA on Work, Leisure and Time Allocation. Washington, DC: International Food Policy Research Institute.

Parker, S. W., & Todd, P. E. (2017). Conditional Cash Transfers: The Case of Progresa/Oportunidades. Journal of Economic Literature, 55(3), 866-915.

Parker, S., Rubalcava, L., & Teruel, G. (2008). Evaluating Conditional Schooling and Health Programs, ch. 62, p. 3963-4035 in Schultz, T. Paul and Strauss, John A. eds., Elsevier, http://EconPapers.repec.org/RePEc:eee:devchp:5-62.

Ramos, A. (2016). Household Decision Making with Violence: Implications for Conditional Cash Transfer Programs. Universitat Autonoma de Barcelona.

Rubalcava, L., Teruel, G., and Thomas, D. (2009) Investments, Time Preferences, and Public Transfers Paid to Women. Economic Development and Cultural Change 2009 57:3, 507-538

Schaner, S. (2017). The Cost of Convenience? Transaction Costs, Bargaining Power, and Savings Account Use in Kenya. Journal of Human Resources, 52(4), 919-945.

Skoufias, E. (2001). PROGRESA and its Impacts on the Human Capital and Welfare of Households in Rural Mexico: A Synthesis of the Results of an Evaluation by IFPRI. International Food Policy Research Institute, Washington, DC.

Sokullu, S., & Valente, C. (2018). Individual Consumption in Collective Households: Identification Using Panel Data with an Application to PROGRESA. Unpublished Manuscript, 3(5), 12.

Thomas, D. (1990). Intra-household Resource Allocation: An Inferential Approach. Journal of Human Resources, 635-664.

Tommasi, D. (2019). Control of Resources, Bargaining Power and the Demand of Food: Evidence from PROGRESA. Journal of Economic Behavior Organization, 161, 265-286.

Tommasi, D., & Wolf, A. (2018). Estimating Household Resource Shares: A Shrinkage Approach. Economics Letters, 163, 75-78.

UN World Health Organization (2014) Global Status Report on Violence Prevention 2014. Wold Health Organization, United Nations Development Programme, and United Nations Office on Drugs and Crime, Geneva, Switzerland.

Voena, A. (2015). Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?. The American Economic Review, 105(8), 2295-2332.

Weyl, G. (2019). Price Theory. Journal of Economic Literature, 57 (2): 329-84.

Wickham, H. (2016) ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York.

Wolf, A. (2016). Resource Allocation in Couples: A Collective Model with Prior Information. ECARES Working Papers.

Appendix A: Identification Details

In this Appendix, we give more details on the semiparametric identification strategies discussed in sections 1.3.2 and 1.3.3. The key identification results is that, regardless of which strategy the researcher employs, the subset set of first order conditions analyzed and the prediction models provide more equations than there are parameters to estimate. First consider the pair of parents' utility functions in the model in section 1.3.2.

Assumption 1: People's utility functions are uniformly drawn from the uncountably infinite set of all monotonically increasing, twice differentiable, and strictly quasi-concave functions, C.

Define a functional $\Gamma : C \times C \rightarrow [0,1]$, which maps from the set of possible utility functions for a couple to the unit interval. This function relates the realized utility functions for each partner to the power dynamic in their family, $\Gamma(\tilde{U}(U_f, U_m)) = \eta$. We know that the η that obtains from Γ will be the value that solves $[\eta]$ (or $[\eta']$ if the relevant outside option is divorce).

To show that the η that solves $[\eta]$ (or $[\eta']$ depending on the outside option) exists and is unique, we must clarify that η is actually a function of prices and income. It is a Marshallian demand function. As such it is more appropriate to write $\mathbb{H} : R^{n+1} \rightarrow [0,1]$ and $\Gamma : C \times C \rightarrow \mathbb{H}$, where $\mathbb{H}$ is the space of Marshallian demand functions for power, which map from the space of incomes and prices to the unit interval. Then, $\eta^* = \eta^*(y/p) \in \mathbb{H}$. As such, this functional is a differential operator on the social welfare function:

$$\Gamma(\tilde{U}(U_f, U_m)) = \left\{ \eta^*(y/p) \;\middle|\; 0 = \frac{\partial \tilde{U}}{\partial U_f}\frac{\partial U_f}{\partial \eta^*} + \frac{\partial \tilde{U}}{\partial U_m}\frac{\partial U_m}{\partial \eta^*} \right\} \tag{1.15}$$

Proposition 1*: Let Assumption 1 hold, then Γ is bijective.*

Proof: Let Assumption 1 hold. Each element of C is a quasi-concave and twice differentiable function, and the constraints in (3) and (6) are linear, so Kuhn-Tucker's theorem applies to the operator in (15). Then the solution to a family's constrained optimization problem, η^*, exists and is unique by Kuhn-Tucker's theorem. □

Assumption 2: A family's social welfare function, $\tilde{U}$, is linear.

Proposition 2: *Let Assumptions 1 and 2 hold. Γ is a linear operator.*

Proof: Linearity implies and is implied by $\Gamma(\tilde{U}(U_{f,1},U_{m,1})+\tilde{U}(U_{f,2},U_{m,2})) = \Gamma(\tilde{U}(U_{f,1},U_{m,1}))+\Gamma(\tilde{U}(U_{f,2},U_{m,2}))$, and $\Gamma(\alpha\tilde{U}(U_f,U_m)) = \alpha\Gamma(\tilde{U}(U_f,U_m))$ for some $\alpha \in C$. Consider the first implication:

$$\Gamma(\tilde{U}(U_{f,1},U_{m,1}) + \tilde{U}(U_{f,2},U_{m,2})) =$$

$$\left\{\eta_1^* + \eta_2^* \;\middle|\; 0 = \left(\frac{\partial\tilde{U}}{\partial U_{f,1}}\frac{\partial U_{f,1}}{\partial\eta_1^*} + \frac{\partial\tilde{U}}{\partial U_{m,1}}\frac{\partial U_{m,1}}{\partial\eta_1^*}\right), \quad 0 = \left(\frac{\partial\tilde{U}}{\partial U_{f,2}}\frac{\partial U_{f,2}}{\partial\eta_2^*} + \frac{\partial\tilde{U}}{\partial U_{m,2}}\frac{\partial U_{m,2}}{\partial\eta_2^*}\right)\right\}$$

$$= \left\{\eta_1^* \;\middle|\; 0 = \frac{\partial\tilde{U}}{\partial U_{f,1}}\frac{\partial U_{f,1}}{\partial\eta_1^*} + \frac{\partial\tilde{U}}{\partial U_{m,1}}\frac{\partial U_{m,1}}{\partial\eta_1^*}\right\} + \left\{\eta_2^* \;\middle|\; 0 = \frac{\partial\tilde{U}}{\partial U_{f,2}}\frac{\partial U_{f,2}}{\partial\eta_2^*} + \frac{\partial\tilde{U}}{\partial U_{m,2}}\frac{\partial U_{m,2}}{\partial\eta_2^*}\right\}$$

$$= \Gamma\tilde{U}_1 + \Gamma\tilde{U}_2.$$

So the first assertion holds. The second assertion holds as well:

$$\Gamma(\alpha\tilde{U}) = \left\{\alpha\eta^* \;\middle|\; 0 = \alpha\left(\frac{\partial\tilde{U}}{\partial U_f}\frac{\partial U_f}{\partial\eta^*} + \frac{\partial\tilde{U}}{\partial U_m}\frac{\partial U_m}{\partial(\eta^*)}\right)\right\}$$

$$= \alpha\left\{\eta^* \;\middle|\; 0 = \frac{\partial\tilde{U}}{\partial U_f}\frac{\partial U_f}{\partial\eta^*} + \frac{\partial\tilde{U}}{\partial U_m}\frac{\partial U_m}{\partial(\eta^*)}\right\}$$

$$= \alpha\Gamma\tilde{U}.$$

□

Proposition 3: *Let Assumptions 1 and 2 hold. If the outside option is not divorce, then η^* is uniformly distributed on $[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}]$. If the outside option is divorce, then η^* is uniformly distributed on $[\frac{y_f^o+\gamma_f}{y}, 1 - \frac{y_m^o+\gamma_m}{y}]$.*

Proof: From the first order conditions to (3) and (6), we know that η^* is bounded on either $[\frac{y_f^o}{y}, 1-\frac{y_m^o}{y}]$ (if the outside option is not divorce) or on $[\frac{y_f^o+\gamma_f}{y}, 1-\frac{y_m^o+\gamma_m}{y}]$ (if the outside option is divorce). Let these sets be non-empty. values of power outside these intervals will lead to the outside option allocation instead of the collective allocation. Because Γ is a linear transformation of a pair of bivariate uniformly distributed random variables, $\eta^* = \Gamma(\tilde{U}(U_f, U_m))$ is also a uniformly distributed random variable. The probability distribution of η^* can be written $P(\eta^* \in [\frac{y_f^0}{y}, 1-\frac{y_m^0}{y}]) = P(\{U_f, U_m\} \in \Gamma^{-1}(\tilde{U}(U_f, U_m)))$ with probability density function

$$f(\eta) = \begin{cases} \frac{1}{(1-\frac{y_m^0}{y})-(\frac{y_f^0}{y})} & \text{if } \eta \in [\frac{y_f^0}{y}, 1-\frac{y_m^0}{y}] \\ 0 & \text{if else.} \end{cases}$$

□

We can slightly relax Assumption 1 to reflect the likely possibility that some preferences are more likely to obtain in a population than others. For instance, for cultural reasons, people may prefer some foods over others. As technologies evolve, people may prefer some modes of entertainment or travel over others. If there are broad trends of this kind in the population, then asserting that any utility function is equally likely is too strong. Instead of a uniform distribution assumption, we can make a conditionally uniform distribution assumption:

Assumption 1A: Within a single family, both partners' utility functions are uniformly drawn from a subset of C, c_l.

Under Assumptions 1A and 2, there is a uniform probability distribution on the family-specific set $[\frac{y_f^o}{y}, 1-\frac{y_m^o}{y}]$ of values of η that solve $[\eta]$ (when the outside option is not divorce). The expected value of the empirical analog of $[\frac{y_f^o}{y}, 1-\frac{y_m^o}{y}]$ is the estimator of η and is given in equation (5). When the outside option is divorce, the estimator is given by equation (10).

Define the number of families in a sample to be H, and denote a specific family with subscript h. Each household has two decision makers, f and m, so there are a total of $2H$ individuals for which to predict earnings in the outside outside option.

Assumption 3: The functions that relate observable characteristics and latent ability to the individuals' income in the outside options, $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$, are identified.

Using the semi-parametric identification strategy, and under Assumptions 1A, 2, and 3, we can write the system of equations to be estimated as

$$\hat{\eta}_h = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_{f,h}\psi_{f,h}) - F(X_{m,h}, \psi_{m,h})}{y_h}\right) \quad \forall h \in 1, \ldots, H \tag{1.16}$$

$$F(X_{h,j}, \psi_{h,j}) = \beta' X_{h,j} + \beta_\psi \hat{\psi}_{h,j} + \epsilon_{h,j} \quad \forall h \in \{1, \ldots, H\} \text{ and } \forall j \in \{m, f\}$$

$$\hat{\psi}_{h,j} = \Delta(X_{h,j}) + \nu_{h,j} \quad \forall h \in \{1, \ldots, H\} \text{ and } \forall j \in \{m, f\}$$

where Δ is a control function and $\nu_{h,j}$ is an error term. Denote the number of population parameters to be estimated in the prediction models as K so that, by estimating the models separately for men and women, there are a total of $2K$ parameters to estimate in $F(X_{i,j}, \psi_{i,j})$ $\forall i \in \{1, \ldots, H\}$ and $\forall j \in \{m, f\}$. Some subset of the population parameters to recover in predicting each partners' outside option earnings may only appear in the control function. For instance, you may include an instrumental variable in the control function.

The key identification requirement in this model is that the researcher can predict how much income a person would have in the outside option. Assumption 2 simply states that the models $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$ have more equations than parameters: $H > K$. The degrees of freedom in each model are $D = H - K > 0$. Abstractly speaking, this assumption that a model of earnings data is identified is analogous to the assumptions about demand functions in Browning, Chiappori, and Lewbel (2013) and Dunbar, Lewbel, and

Pendakur (2013) who assume that demand functions and Engels' curves, respectively, are identified.

Proposition 4*: Let Assumptions 1A, 2, and 3 hold.Let the outside option be an inefficient equilibrium in the family that does not end cohabitation. Then the estimator of bargaining power is point identified on (1.16).*

Proof: In the model to predict y_f^o, there are H equations and K parameters to be recovered. In the model to predict y_m^o, there are H equations and K parameters to be recovered. Under Assumption 3, $H > K$. As such, the number of parameters to be estimated (1.16) is $2K + H$ and the number of equations is $3H$. Therefore, the model is identified. $\square$

Now consider the more challenging identification problem presented by the model in equation (6). Now the model is identified if there are at least two waves of data to estimate a fixed effects model.

Assumption 4: For each household, there are at least two waves of data: $T \geq 2$.

Under Assumptions 1A-3, the system of equations to be estimated is:

$$\hat{\eta}_{h,t} = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_{f,h,t}\psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t)}}{\tilde{y}_{h,t} - \tilde{\beta}_{0,h}}\right) \quad \forall h \in \{1, \ldots, H\} \text{ and } \forall t \in \{1, \ldots, T\} \tag{1.17}$$

$$y_{h,t} = \beta_{0,h} + \beta_1(F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})) + \epsilon_{h,t} \quad \forall h \in \{1, \ldots, H\} \text{ and } \forall t \in \{1, \ldots, T\}, \text{where:}$$

$$\beta_1 + \frac{\epsilon_{h,t}}{F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})} \in [-1, 1]$$

$$F(X_{h,j,t}, \psi_{h,j,t}) = \beta' X_{h,j,t} + \beta_{\psi,t}\hat{\psi}_{h,j,t} + \epsilon_{h,j,t} \quad \forall h \in \{1, \ldots, H\} \text{ and } \forall j \in \{m, f\} \text{ and } \forall t \in \{1, \ldots, T\}$$

$$\hat{\psi}_{h,j} = \Delta(X_{h,j}) + \nu_{h,j} \quad \forall h \in \{1, \ldots, H\} \text{ and } \forall j \in \{m, f\} \text{ and } \forall t \in \{1, \ldots, T\}$$

Proposition 5: *Let Assumptions 1A, 2, 3, and 4 hold. Let the outside option include divorce. Then the estimator for bargaining power is point identified on (1.17).*

Proof: In the model to predict y_f^o, there are HT equations and KT parameters to be recovered. In the model to predict y_m^o, there are HT equations and KT parameters to be recovered. Under Assumption 1, $H > K$. In the linear equation given by (1.10), there are H intercept parameters and one population-level slope parameter to estimate. There are HT equations in this system. These parameter estimates identify the household specific error term $\tilde{\epsilon}_{h,t}$. As such, the number of parameters to be estimated in (1.17) is $T(2K + H) + 1$ and the number of equations is $3HT$. By Assumption 4, $T \geq 2$, and so the system in (1.17) identifies a unique value of η for each household and each period. □

As a final comment, note that the key to identification in this model is the same as the key to identification in DLP. In their model to recover resource shares from household consumption of private assignable goods, they set up a system of Engel curves for each family and the constraint that resource shares sum to one. They estimate population parameters to identify the Engel curves, and back out household-level parameters from the household specific system of equations with the sample estimates plugged in. Likewise, we set up a system of family level equations that includes population parameters. After estimating the population parameters and plugging in the sample estimates, we can solve the system of equations for each household to find their household specific power dynamic, η_h.

Appendix B: Counterfactual Cash Transfer Targeting Schemes

We can use our structural model to study the ramifications of alternative gender-targeting transfers. We can use our predicted earnings in the outside option, and different Progresa transfer schemes, to understand how the bargaining power impacts depend explicitly on gender targeting. Two interesting counterfactuals are the case where the transfers went completely to fathers, and the case where the transfers were split equally between partners. We hold all else equal in these counterfactuals, like employment decisions and the value of the transfers for each family in each period.

Under the first counterfactual, men's outside options are increasing since they get an additional monthly income source. The distributions of bargaining power across families and time under this hypothetical transfer are given by Figure 1.3. The median female bargaining power decreases over three years, and the magnitude of the change is reduced. Giving the same transfers to fathers empowers them less since their outside options are already relatively more valuable than their partners'.

If the government had split the transfer equally between partners, then both women's and men's outside options would improve. However, women's would improve more, since they start from a position of relative disadvantage. Thus, under this split regime, the median power increases from 1997 to 2000, but by less than when the transfers go to women completely. The counterfactual distributions are given in Figure 1.4.

This exercise sheds light on *why* bargaining power increases from 1997-2000 in the observed sample. It is because of the government's targeting transfers to mothers. If the government had prescribed to an alternative targeting method, then the transfers could easily have benefited men on average, or made relatively smaller but economically significant difference for bargaining relations. So our measure confirms previous studies' findings of an increase because of the gender targeting (and transfer magnitudes).

These hypothetical distributions also provide some intuition on how the Government of Mexico's other transfer programs might influence power dynamics. Some of them are implicitly gendered, like Procampo and Procede, and so we would expect them to generate changes in power dynamics for the median household. These two transfers were based on land management patterns, and in 1997 approximately 15% of land in Mexico was managed by women (Klein and McArthur, 2018) and so approximately 85% of these two government transfers went to men. If we were evaluating these two implicitly gendered programs that favor men, instead of Progresa, we would expect a decrease in women's bargaining power, as in Figure 1.3.

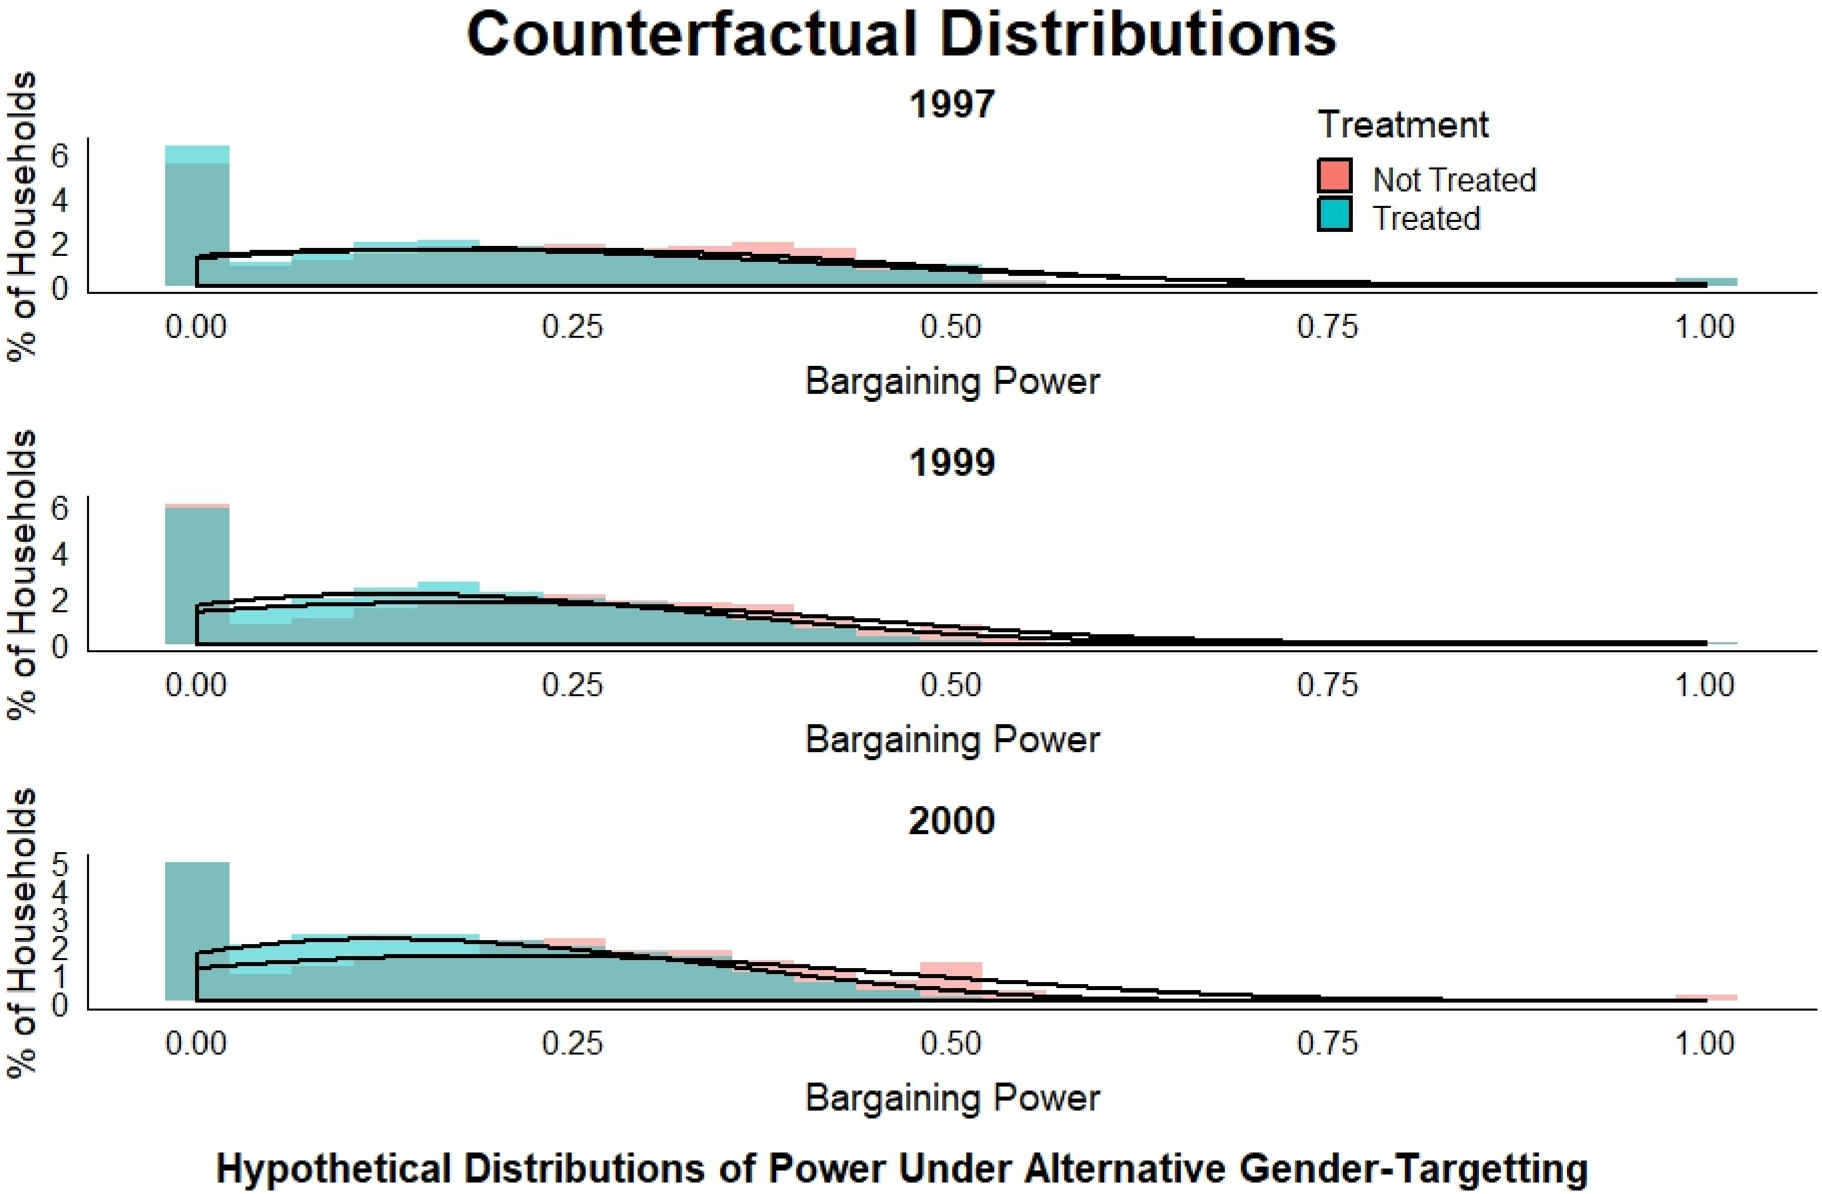

Figure 1.3: Unobserved Counterfactual Distributions Where Men Get Transfer Instead of Women

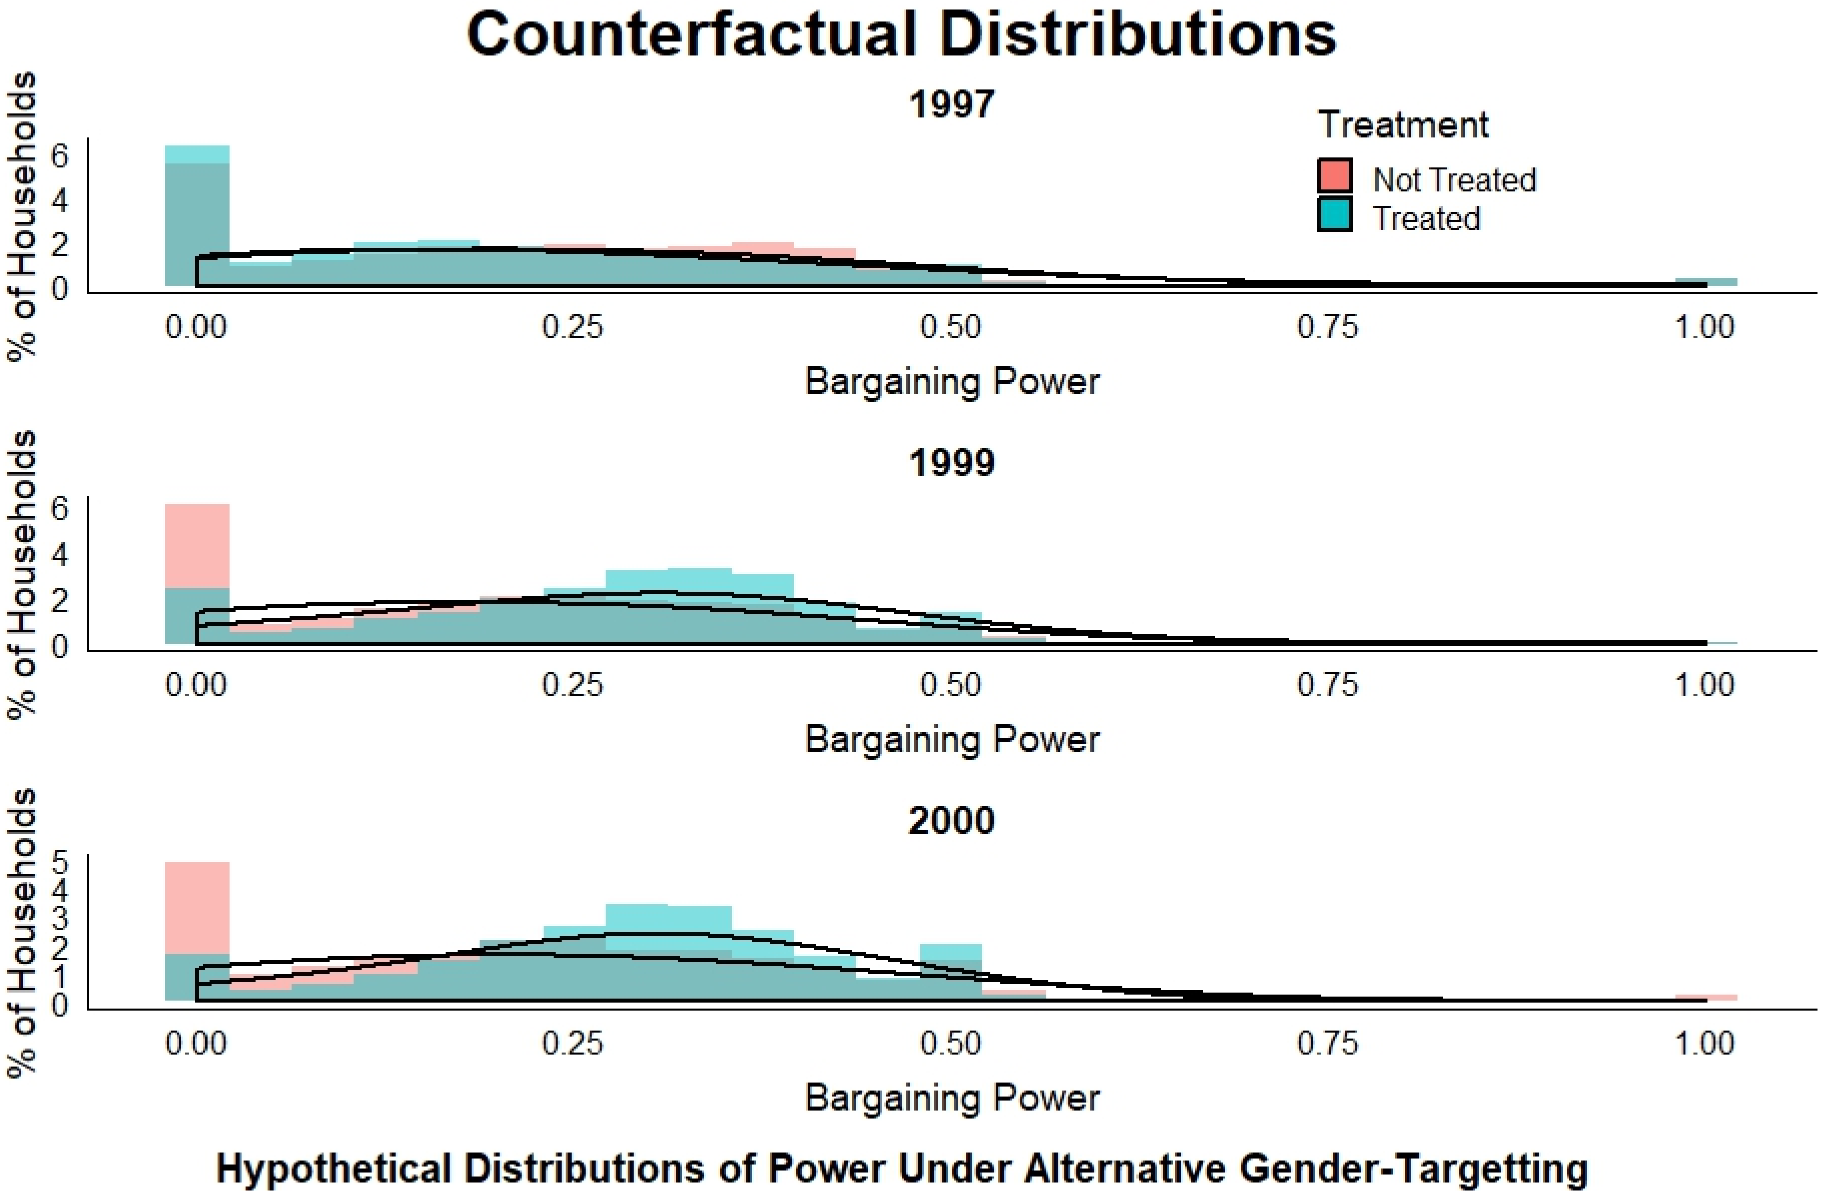

Figure 1.4: Unobserved Counterfactual Distributions Where Men and Women Split Transfer Equally

Appendix C: Additional Context on Power Dynamics in Rural Mexico 1997-2000

The longitudinal survey administered in order to evaluate Progresa included information on household heads' decision making practices, their views on women's rights in society, and on consumption choices regarding select private assignable goods. This information allows us to provide additional context for our study. Along with the employment and transfer information in Table 1.2, the summary statistics presented in this appendix provide a more complete backdrop for understanding power dynamics in the context we study. We can assess whether the bargaining power estimates we generate by estimating the model in Section 1.3 are reasonable.

In Table 1.5, we report the additional decision-making information, categorized by module, in the baseline and in March of 1999, one year after the program began and during the Progresa experiment. We present the summary stats in the second wave according to treatment status, with the third column ("1999 T") relating to the treatment group and the fourth relating to the control group ("1999 C").

The first component of Table 1.5 reports expenditures on shoes and clothing for girls, boys, women, and men. Private assignable goods are a key puzzle piece in understanding household power dynamics for two reasons: they are observable and they form the basis of the resource share estimation strategy in DLP. In 1999, the treatment group spent twice as much on children's shoes and clothing than the control group. Women and men spent the same amount on these goods. We analyze the relationship between women's bargaining power and these expenditures below.

The next three components of Table 1.5 are devoted to questions about who makes decisions in the family. The response options are either that the couple decides together, that the woman decides by herself, or that the man decides by himself. We present three groups of dummy variables to show these dynamics.[30]

Couples typically make decisions together. Some unusual patterns in this context are

[30]There is a forth category of "no response" which we do not include.

that women are eight times more likely than men to control small livestock in 1997 (as noted by Rubalcava, Teruel, and Thomas, 2009); men are much more likely than women to decide by themselves how women's income is spent after Progresa is implemented (as noted in Table 5 of Adato et al, 2000); and, in all categories, men seem to have more autonomy in the treatment group in 1999 than their peers in the control group. The relationship between government transfers and decision making, then, seems to be a complicated one. Adato et al (2000) solve this problem by augmenting their work with a rigorous qualitative analysis.

The last portion of Table 1.5 presents summary statistics for views on women's rights. The responses for these questions are either to agree with the statement, disagree with the statement, or to neither agree nor disagree with the statement. The original Spanish language statements are presented in the table notes. We present dummy variables equal to one when the respondent agreed with the statement. For the most part, the responses stay constant over time, with the main exceptions being increases in the likelihood that women need permission to leave the home and that women should give an opinion on community affairs.

All in all, these summary stats point to low bargaining power for women who do not receive Progresa, and a complex relationship between power and agency in the family.[31] Future research can explore the relationship between power and agency further, as in the essays on control and power by Basu (2006) and Ramos (2016), among others. See d'Aspremont and Ferreira (2014 and 2019) for theoretical treatments of agency and power in the family.

[31] Initial explorations into the empirical relatinship between agency and $\hat{\eta}$ did not yield clear patterns.

Table 1.5: Bargaining Power and Decision Making Household Means by Year

	Baseline	1999 T	1999 C
Bargaining Power	0.217	0.406	0.203
Private Assignable Goods			
Spending on Girls' Shoes and Clothing (pesos/6 months)	196.19	142.603	72.489
Spending on Boys' Shoes and Clothing (pesos/6 months)	208.384	155.204	80.286
Spending on Women' Shoes and Clothing (pesos/6 months)	199.858	87.779	102.437
Spending on Mens' Shoes and Clothing (pesos/6 months)	327.738	111.031	110.196
Female Autonomy			
Telling Children to go to School	0.091	0.091	0.08
Spending Women's Marginal Income	0.053	0.018	0.02
Household Repairs	0.02	0.033	0.034
Buy Children's Shoes	0.052	0.05	0.046
Telling Sick Children to go to the Doctor	0.096	0.096	0.086
Control Small Livestock	0.161	NA	NA
Control Garden Products	0.057	NA	NA
Male Autonomy			
Telling Children to go to School	0.101	0.11	0.096
Spending Women's Marginal Income	0.175	0.359	0.334
Household Repairs	0.3	0.308	0.267
Buy Children's Shoes	0.235	0.196	0.181
Telling Sick Children to go to the Doctor	0.096	0.098	0.088
Control Small Livestock	0.019	NA	NA
Control Garden Products	0.096	NA	NA
Joint Decision Making			
Telling Children to go to School	0.762	0.747	0.719
Spending Women's Marginal Income	0.734	0.576	0.558
Household Repairs	0.637	0.608	0.604
Buy Children's Shoes	0.669	0.705	0.67
Telling Sick Children to go to the Doctor	0.764	0.757	0.723
Control Small Livestock	0.299	NA	NA
Control Garden Products	0.28	NA	NA
Dummy Variables for Views on Women's Rights			
Do Women Need Permission to Leave the Home?	0.33	0.46	0.452
Is a Woman Place in the Home?	0.539	0.54	0.52
Should Women be Obedient?	0.775	0.724	0.705
Should Women Voice Opinion to Community?	0.788	0.831	0.844
Should Women Have Jobs Outside the Home?	0.73	0.735	0.753
Should Men and Women have Equal Rights?	0.896	0.894	0.907
Should Women Have Their Own Opinions?	0.875	0.898	0.906

Notes: The original text for the "Dummy Variables for Views on Women's Rights" variables are as follows: (1) "El lugar de la mujer está en la casa," (2) "Las mujeres deben obedecer a los hombres," (3) "Las mujeres deben opinar en asuntos de la comunidad," (4) "Las mujeres deberían tener un trabajo fuera de casa", (5) "Las mujeres y los hombres deben tener los mismos derechos," (6) "Las mujeres deben tener su propia opinión", and (7) "Tiene usted que pedirle permiso a su esposo para visitar a sus parientes o amigas?"

Appendix D: Additional Tables

Table 1.6: Increased Equality Improves Family Diet

	Dependent Variable: Weekly Consumption Dummy			
	Chicken	Eggs	Leafy Greens	Bananas
$\hat{\eta}_{h,t}$	0.093	0.039	0.039	0.087
	(0.026)	(0.019)	(0.019)	(0.027)
Log HH Income	0.032	0.017	0.008	0.025
	(0.007)	(0.005)	(0.005)	(0.007)
# Kids	−0.006	0.005	−0.001	−0.003
	(0.006)	(0.004)	(0.005)	(0.007)
# Young Kids	−0.020	0.002	−0.004	−0.014
	(0.006)	(0.004)	(0.005)	(0.006)
Animal Product Prices	Yes	Yes	No	No
Fruit/Veg Prices	No	No	Yes	Yes
Staple Prices	Yes	Yes	Yes	Yes
Observations	32,764	32,764	32,764	32,764
R^2	0.454	0.390	0.443	0.426
Adjusted R^2	0.156	0.057	0.140	0.113

Table 1.7: Selection Regression Results

	Dependent variable:	
	Earnings Dummy Variable	
	Women	Men
Constant	−0.849 (0.197)	−2.318 (0.161)
Age	0.040 (0.006)	0.112 (0.004)
Age Squared	−0.001 (0.0001)	−0.001 (0.0001)
Other Income Dummy	0.375 (0.164)	0.596 (0.074)
Asihn of Other Income	−0.030 (0.036)	−0.165 (0.017)
Number of Kids	0.023 (0.016)	−0.019 (0.012)
Number of Kids Ages	0.049 (0.136)	0.200 (0.114)
Education	0.030 (0.004)	−0.027 (0.003)
Literate Dummy	0.045 (0.032)	0.159 (0.027)
Other Gov Transfer Dummy	−0.003 (0.046)	0.035 (0.095)
Indigenous Language Dummy	−0.133 (0.047)	−0.021 (0.053)
Spanish and Ind. Lang. Dummy	0.079 (0.047)	0.031 (0.053)
Household Head Dummy	−0.565 (0.037)	0.987 (0.034)
Number Female Adults	0.051 (0.017)	−0.026 (0.014)
Number Male Adults	−0.091 (0.019)	−0.041 (0.013)
Gov. Poverty Index	−0.0002 (0.015)	0.016 (0.012)
Gov. Poverty Dummy	−0.002 (0.015)	−0.021 (0.012)
Number Female Kids	0.019 (0.036)	−0.024 (0.028)
Number Male Kids	−0.061 (0.027)	−0.079 (0.023)
Prop. Village Migrates MEX	−0.054 (0.218)	0.829 (0.148)
Prop. Village Migrates USA	−1.701 (0.739)	−1.952 (0.633)
Num Male Adults * Prop. MEX Mig	0.302 (0.452)	−0.806 (0.401)
Num Male Adults * Prop. USA Mig	0.580 (0.325)	−0.316 (0.223)
Num Female Adults * Prop. MEX Mig	0.313 (0.183)	−0.269 (0.167)
Num Male Adults * Prop. USA Mig	−0.104 (0.303)	0.429 (0.271)
Progresa Control Group Dummy	−0.082 (0.037)	0.052 (0.021)
Female HH Head's Progresa Income	0.018 (0.008)	
Women's Job View Proportion	−0.035 (0.080)	0.131 (0.060)
Need Permission Proportion	−0.388 (0.154)	0.281 (0.127)
Need Accompaniment Proportion	0.121 (0.118)	0.115 (0.094)
ER1 Interaction	0.085 (0.062)	0.044 (0.047)
ER2 Interaction	0.027 (0.135)	−0.136 (0.115)
ER3 Interaction	−0.263 (0.098)	−0.053 (0.084)
State-by-year Fixed Effects	Yes	Yes
Observations	43,871	43,179
Log Likelihood	−14,777.360	−41,838.460
ρ	0.562 (0.104)	−0.720 (0.013)

Table 1.8: Earnings Regression Results

	Dependent variable:	
	Log Earnings	
	Women	Men
Constant	3.255 (0.235)	5.230 (0.042)
Age	0.026 (0.008)	−0.003 (0.002)
Age Squared	−0.0003 (0.0001)	0.00003 (0.00002)
Other Income Dummy	−0.086 (0.209)	−0.204 (0.024)
Asihn of Other Income	0.031 (0.046)	0.025 (0.006)
Number of Kids	0.021 (0.021)	0.011 (0.005)
Number of Kids Ages 0-5	−0.049 (0.021)	−0.014 (0.004)
Education	0.057 (0.006)	0.025 (0.001)
Literate Dummy	−0.047 (0.043)	−0.036 (0.010)
Other Gov Transfer Dummy	−0.088 (0.061)	−0.001 (0.029)
Indigenous Language Dummy	−0.071 (0.066)	−0.224 (0.019)
Spanish and Ind. Lang. Dummy	−0.080 (0.066)	0.053 (0.019)
Household Head Dummy	−0.421 (0.073)	−0.079 (0.014)
Number Female Adults	−0.016 (0.022)	0.021 (0.006)
Number Male Adults	−0.033 (0.026)	−0.002 (0.005)
Gov. Poverty Index	0.001 (0.020)	−0.005 (0.005)
Gov. Poverty Dummy	−0.013 (0.020)	0.002 (0.005)
Number Female Kids	−0.003 (0.051)	−0.025 (0.011)
Number Male Kids	0.121 (0.036)	0.123 (0.009)
Prop. Village Migrates MEX	0.117 (0.281)	−0.092 (0.061)
Prop. Village Migrates USA	0.135 (0.993)	0.793 (0.237)
Num Male Adults * Prop. MEX Mig	−2.165 (0.613)	−0.469 (0.144)
Num Male Adults * Prop. USA Mig	−0.007 (0.469)	−0.113 (0.099)
Num Female Adults * Prop. MEX Mig	0.987 (0.230)	0.092 (0.070)
Num Male Adults * Prop. USA Mig	0.109 (0.411)	−0.104 (0.120)
Progresa Control Group Dummy	−0.128 (0.048)	−0.031 (0.008)
Female HH Head's Progresa Income	0.002 (0.012)	
State-by-year Fixed Effects	Yes	Yes
Observations	43,871	43,179
Log Likelihood	-14777.36	-41838.46
ρ	0.786 (0.052)	0.555 (0.003)

Chapter 2

Gender Equality Can Reduce the Malaria Burden in Malawi

Coauthored with Bradford Barham and Yuexuan Wu

Abstract: We provide the first quantitative evidence that increasing intra-household equality reduces malaria transmission in Malawi. We estimate women's bargaining power using a structural model that relates gender discrimination in institutions outside of the family to power dynamics within the family. Women in Malawi, we find, are significantly disadvantaged in labor markets and in the case of divorce. These disadvantages translate into weak intra-household bargaining positions. We exploit geographic variation in cultural norms in Malawi that confer distinct privileges to women to identify the effect of an increase in women's bargaining power on malaria contraction rates. An increase in women's bargaining power significantly decreases the likelihood that a family member contracts malaria. These results are robust to the use of bargaining power measures derived from other collective models of the family, and to different specifications of our two stage least squares model. (JEL D1, I14, I15)

Keywords: Bargaining Power, Women's Empowerment, Malaria, Sharing Rule

2.1 Introduction

Women's ability to influence decisions within their family - their intra-household bargaining power - is a reflection of gender inequality in institutions outside of the household.[1] For example, discrimination against women in labor markets can reduce their decision-making power in the family (Lise and Yamada, 2019). This inequality in household decision-making processes has consequences. When one parent has a majority of the bargaining power, the family tends to invest less in public goods, leading to reductions in childrens' health and human capital accumulation.[2] Explicitly linking the causes and consequences of inequality in the family can lead to insightful new policy recommendations (e.g. Calvi, 2019).

In the summer of 2019, the Director General of the World Health Organization called for new policies to combat malaria worldwide (WHO a, 2019). Malaria control efforts have reduced the global malaria burden in the last two decades, but standard prevention measures, like distributing bed nets, face diminishing returns.[3] Devising new strategies to further combat malaria is paramount for continuing to reduce the disease's prevalence (WHO b, 2019).[4]

In this paper, we explore the extent to which increasing gender equality across society can reduce the malaria burden by changing *how* married partners jointly make decisions. We study the Malawi LSMS panel data from 2010 and 2013 in order to gain empirical insights. We choose this setting for two reasons. First, the malaria burden is a leading public health concern in Malawi. Roughly 30% of Malawi's population is in-

[1]For theoretical analyses of this topic, see Mazzocco (2007) and Ligon (2011). For empirical documentation of this stylized fact, see Benería and Roldán (1987), Voena (2015), and the evidence reviewed in Doss (2013) and Chiappori and Mazzocco (2017).

[2]For a theoretical analysis on the topic, see Blundell Chiappori, and Meghir (2005). For empirical documentation of this stylized fact, see Thomas (1990), Beegle, Frankenberg, and Thomas (2001), Duflo (2003), Attanasio and Lechene (2010), Li and Wu (2011), Lépine and Strobl (2013), Calvi (2017), Calvi, Lewbel, and Tommasi (2017), Lowes (2017), Sraboni and Quisumbing (2018), Holland and Rammohan (2019), Klein and Barham (2019), and Tommassi (2019). See Annan et al. (2019) for a cross country study within Subsaharan African. See Richards et al (2013) and Doss (2013) for broad literature reviews that discuss family members' health and power dynamics between household heads.

[3]See Ranson and Lissenden (2016) for evidence that mosquitoes are becoming resistant to standard insecticides used to treat bed nets.

[4]In 2018, $663 million were invested in research and development of such methods (WHO b, 2019), reflecting the need for, and potential of, new technologies and approaches.

fected annually, which poses serious morbidity and mortality costs (Fernando, Rodrigo, and Rajapakse, 2010; World Health Organization, 2018). Second, Malawi is intersected by Africa's "matrilineal belt;" both matrilineal and patrilineal communities can be found in Malawi. Across these groups, substantial benefits in inheritance, marriage, and divorce accrue either to men or to women systematically.[5] We exploit this variation in advantages outside of the family to identify the relationship between power dynamics in the family and the likelihood that members of that family contract malaria.[6]

We build and estimate two closely related models in order to fully link inequality in institutions outside of the home to the malaria contraction rates that families face. The first is a collective model where families have idiosyncratic consumption technologies relating household purchases to individuals' consumption (as in Browning, Chiappori, and Lewbel, 2013). In this model, parents also choose whether to remain married or divorce and, if they remain married, they decide together how to allocate their shared budget. We follow the semi-parametric identification strategy developed in Klein and Barham (2019) to recover the expected value of bargaining power from differences in (shadow) prices and (shadow) incomes that partners face in the case of divorce, relative to those they face within the marriage. These differences are functions of inequality in institutions outside of the family. For instance, in patrilineal communities in Malawi, divorced women typically reimburse their partner's family for the bride price they received at marriage. This custom is not observed in matrilineal communities. This added price of divorce changes the relative value of partners' outside options, thereby influencing power dynamics within the family.

Estimating this model on the full sample of Malawian households, we find that the median female decision maker has about one third of the bargaining power that her partner has. We reject the null hypothesis that men and women have equal decision-

[5]For further discussion, see Peters (1997), Reniers (2003), Takane (2008), Ewing et al (2016), Lowes (2017), and Walther (2018).

[6]See Lowes (2017) for a recent study showing how health outcomes differ across these cultural groups. She finds that children have better health outcomes in general in matrilineal communities. See Walther (2018) for a recent study showing that men and women supply labor and invest in land differently across these cultural groups, specifically in order to try to gain influence within the family. Many historical and ethnographic studies predate these recent economic studies; for example, Phiri (1983).

making power (t-statistic = 69.75). This gender gap in bargaining power is driven by large differences in labor market opportunities and by the functioning of other institutions as captured by partners' relative compensating variation values. The median man's earnings are more than three times larger than the median woman's earnings. In the median household, the institutional benefits that men enjoy are valued at roughly nine months of earnings for the female household head.

Next, we study differences across matrilineal and patrilineal communities in Malawi. We would expect women to have higher bargaining power dynamics in matrilineal communities, since women are relatively less disadvantaged there. In fact, the gap between men's and women's observed earnings is lower in matrilineal communities. Also, the difference between men's and women's (estimated) compensating variation values is lower in matrilineal communities (men are relatively better off in the case of divorce in patrilineal communities). Combined, these differences result in a 30% increase in the average women's bargaining power in completely matrilineal communities compared to completely patrilineal communities.

To formalize this difference, we regress the percent of matrilineal families in a community on women's bargaining power in the family. We find that as the percent increases, women's bargaining power increases as well. The significant regression result forms the basis of our second model - a two-stage least squares (2SLS) approach relating women's bargaining power in the family to malaria contraction rates. Using this 2SLS model, we find that a standard deviation increase in women's bargaining power in the family would result in a 58% reduction in malaria contraction rates. Under the assumption that the gender of the parent who benefits from these inheritance, marital, and divorce customs only influences malaria contraction rates through bargaining power dynamics and observed controls, we can interpret this as a causal result. Women's empowerment, we find, is a viable way to reduce the malaria burden in Malawi.

We conduct two counterfactual analyses to understand how changes in gender equality across society could empower Malawian women, and subsequently cause a reduction in malaria incidence. First, we study how a cash transfer targeted to women — similar to the *Progresa/Oportunidades/Prospera* transfer in Mexico — could increase their expected earnings in the case of divorce, improving the relative value of women's outside options and increasing equality in the family. We find that a $25 monthly transfer to women in 2010 (when Malawi's GDPPC was $480) would have empowered women enough to result in equality in decision-making power between men and women in the family and a 80% reduction in malaria contraction rates. Second, we study how an additional year of schooling for women would influence earnings outcomes, and thus power dynamics in the family. We find that additional school does not have a large return (one year would increase female bargaining power less than 4%), since earnings opportunities are limited and educational attainment is very low to begin with. We conclude that a cash transfer in this setting would empower women to a greater degree, and reduce malaria more effectively, than an intervention attempting to increase women's educational attainment.

To study the robustness of our results to various structural modeling assumptions, we estimate two additional measures of bargaining power in the family. These have very different modeling approaches. Finding corroborating results in these robustness checks may alleviate concerns that our findings are an artifact of the semi-parametric identification strategy in Klein and Barham (2019). A major difference in our main specification and the robustness-check specifications is that the decision to divorce is endogenous in our model.

Our first robustness check is to estimate the Dunbar, Lewbel, and Pendakur (2013, "DLP" from hereon out) model of the family, which allows the econometrician to recover the family's sharing rule from private assignable goods information. We follow DLP in estimating Engle curves for clothing and shoes for mothers, fathers, sons, and daughters in the LSMS data. The second robustness check is to construct a proxy for power from partners' demographic information. This kind of reduced-form approach is commonly used, perhaps more frequently than structural modeling approaches.[7] In both of these checks,

[7]Recent examples of studies using indices include those by Saha and Sangwan (2019), who study women's power and access to microfinance; Schaner (2017), whose index and qualitative verification strategy is the only survey-based method to measure power that is referenced in the Abdul Latif Jameel Poverty Action Lab's "A Practical Guide to Measuring Women's and Girls' Empowerment in Impact Evaluations" (Glennerster, Walsh, and Diaz-Martin, 2018); and Mishra and Sam (2016), who use an index created from reported decision making patterns. Mishra and Sam (2016) provide additional examples

women in the median family have significantly higher bargaining power in matrilineal communities than they do in patrilineal communities. We use these alternative measures in our 2SLS model, and find broadly corroborating results.

To probe the robustness of our 2SLS model results, we test two key controls, a geographic fixed effect and a local malaria incidence measure. Malaria contraction rates and our instrument vary spatially in Malawi; thus, controlling for the geographic region of households captures this coincidental covariation.[8] By including the percent of neighboring families with a member who has had malaria in the last two weeks, we incorporate a strong proxy for the malaria risk factors that families face. Our results are robust to both specification changes separately and combined.

In addition, we study two mechanisms underlying our results, though data limitations preclude some potentially insightful analyses. We know that women's bargaining power influences family allocations in ways that reduce the likelihood of malaria contraction for family members. We hypothesize that increased equality in the family leads to increased household demand for bed nets. However, we are unable to reject the null hypothesis that this is not the case, likely because bed nets are already common in this setting. We hypothesize that families with more equal partners are more likely to have private toilet facilities and thus be potentially less exposed to infected mosquitoes. We reject the null hypothesis that this is not the case. In addition, we hypothesize that more equal partners might invest more in children's nutrition, and may allocate more resources to caring for sick household members (as documented in focus group interviews by Ewing et al, 2016). We leave empirical investigation of these potential mechanisms for future research.

We conclude that policies that increase gender equality across society can significantly reduce the malaria burden in Malawi. We suggest that current anti-malaria programs incorporate a gender component, and that empowerment programs designed explicitly to reduce malaria transmission rates be considered. While we find that reallocating some

---

and ground their index in Agarwal's (1997) conception of women's power and agency. See Ewerling et al (2017) for an example of a study that uses principal components analysis to generate weights for a bargaining power index.

[8] We control for geography at the "regional" level, of which there are three in Malawi. For reference, there are counties in the United States of America that are larger geographically than the regions in Malawi.

level of anti-malarial funding to women's empowerment programs would further reduce the malaria burden, we leave an analysis of the optimal anti-malarial program portfolio to future research.

We invite policy makers to skip from here to Section 2.7. In Section 2.2, we build our model of family decision making and describe the estimation process. In Section 2.3, we present the results of estimating this model for the full sample of Malawian households in 2010 and 2013. In Section 2.4, we describe the matrilineal and patrilineal cultural differences more fully, and present results from estimating our model using data from these communities separately. In Section 2.5, we explore these differences by estimating a 2SLS model that links women's bargaining power to malaria incidence. In Section 2.6, we implement the robustness checks described above. In Section 2.7, we develop our two counterfactual analyses. We conclude in Section 2.8, and provide the robustness checks and mechanisms analysis in appendices.

2.2 A Model of the Causes of Inequality in the Family

In this section, we formalize the link between household allocations, power dynamics in the family, and inequality in institutions outside of the home. We build the model and empirical approach in five steps. First, we briefly provide a summary of the basic collective model of the family, which is at the core of our analytical framework. This standard model relates parents' preferences to household demand functions. Second, we decentralize the family's problem following Chiappori (1988) to set up a comparison between price and income settings for each partner in the collective allocation and in the outside option. Third, we incorporate these comparisons into the family's optimization problem as participation constraints. Fourth, we introduce the type of assumptions required for our semi-parametric identification strategy. These are assumptions on the distribution of preferences across the population. Fifth, we make specific distributional assumptions. Under these assumptions, we can recover the expected value of women's bargaining power in each family from information on their earnings opportunities outside of the home, and gender discrimination outside of the home in other contexts. A full proof of semi-parametric identification in this setting can be found in Klein and Barham (2019).

Step 1: The Collective Model of Family-Decision Making Processes

Families make many decisions jointly, like where to live, what to eat, which school to send a child to, and when to send a child to the doctor. All of these choices affect family members' health, and the likelihood of someone in the family contracting malaria. Modeling how parents' preferences, and their decision-making power, relate to observed choices is important in understanding the take-up of preventative measures, investments in children, and ultimately, the likelihood that any family member becomes ill. To model power in the family, economists can use the collective model of household behavior (Apps and Rees, 1988; Chiappori, 1988, 1992).[9]

Consider a family with two decisions makers, indexed f and m, and a child, indexed c. Families have total income y. They jointly determine how to spend this budget, and purchase a $k-$vector of goods, denoted $z = (z^1, \ldots, z^k)$ with prices $p = (p^1, \ldots, p^k)$. Each family has a distinct consumption technology that converts these household-level purchases into consumption goods for each family member. Denote the amount that each person consumes of each good $x_\tau = (x_\tau^1, \ldots, x_\tau^k)$ for individuals of type $\tau \in \{c, f, m\}$. Each family member has an increasing, quasi-concave, twice differentiable utility function that represents their preferences over their private consumption amounts. In addition, assume that parents choose the allocation without explicit input from their child. They consider their dependent's utility over consumption when making choices, and the child takes their parents' decisions as given. Then, parents' utility functions are over consumption and their child's indirect utility. As such, the child's utility function is $U_c(x_c)$, and the parents utility functions are $U_f(x_f, U_c(x_c))$ and $U_m(x_m, U_c(x_c))$.

Denote the family-specific consumption technology as $A(z)$ such that $x_f + x_m + x_c = A(z)$. This consumption function captures economies of scale that result from cooperation,

[9]This model is elaborated on by Browning et al. (1994), Browning and Chiappori (1998), Blundell, Chiappori, and Meghir (2005), and Chiappori and Ekeland (2006).

and the positive externalities that obtain when families consume pubic goods. There are public goods within the family. It is cheaper to heat one home than it is to heat many, for example. There may be less spillage when one person cooks for the family, compared to each member cooking separately.[10] Note that expenditures that make one person less likely to contract malaria, or which help a sick person clear the contagious disease faster, have a positive externality: they also reduce the chances that other family members become ill. As such, expenditures on a bed net that only cover one person are actually a public good. This consumption technology captures these positive externalities.

Assume that the family chooses z so that no alternative allocation can be made that makes someone better off, and everyone else at least as well off; in other words, the family reaches a Pareto efficient allocation. Assume that families consider real prices instead of nominal prices (no money illusion). Then there exists some family-specific social welfare function, $\tilde{U}(U_f(x_f, U_c(x_c)), U_m(x_m, U_c(x_c)))$, that is twice differentiable and strictly increasing in each of its arguments, so that the family's problem is:

$$\max_{z,x_f,x_m,x_c} \tilde{U}(U_f(x_f, U_c(x_c)), U_m(x_m, U_c(x_c))) \text{ such that } p'z = y \text{ and } x_f + x_m + x_c = A(z) \tag{2.1}$$

and the solution to this program gives demand functions for each individual, $x^*_\tau(p/y)$.

Step 2: A Two-Stage Formulation of the Collective Model

By assumption, the solution to (2.1) is a Pareto efficient allocation.[11] As such, by the

[10]To gain intuition for how this consumption technology works, consider two examples. First, if the family purchases one unit of heating for their dwelling, and everyone spends an equal amount of time enjoying the heated dwelling, then each person consumes a unit of heating. This is a public, non-rival good. Consider a second example. If the family purchases a bed net and the parents sleep under it, leaving the child unprotected, the parents each consume one unit of bed netting. The child has a positive externality of living in a space with a lower human reservoir of the disease, however, and so gains utility from their parents consumption of the bed net. The consumption technology would capture the positive externality associated with the parents consumption of this partially public, rival good.

[11]This assumption can be relaxed. Lewbel and Pendakur (2019) use this same model set up and a conditional-Pareto-efficiency assumption to decentralize the model. That approach can be applied here, but does not need to be. See Chavas and Klein (2020) for such an extension.

second welfare theorem, the same allocation will be reached by a decentralized economy subject to the optimal division of resources. Expressing the problem using the decentralized formulation is useful because it introduces the notion of bargaining power we use in this paper: the sharing rule.

This alternative formulation of the problem comprises two steps and is solved by backwards induction. In the first step, the household divides the common resources between the two decision makers. In the second step, the decision makers privately allocate their (shadow) income to maximize their utility. The resulting set of demand functions is the same that results from (2.1), as shown by Chiappori (1988 and 1992).

In a decentralized economy with public goods, "Lindahl" prices clear the market.[12] Because each family uses an idiosyncratic consumption technology, each family faces a corresponding family-specific set of Lindahl prices in this decentralized optimization program (Browning, Chiappori, and Lewbel, 2013). Within a household, partners face the same (shadow) prices since they use the same consumption technology. Let $L(p) \leq p$ denote the Lindahl prices in the family, such that there is an invertible function relating A and L (that is, for each consumption technology there is a unique set of corresponding Lindahl prices, and vice versa). Lindahl prices are lower than market prices when the consumption goods have positive externalities.

Partners have different (shadow) incomes, which are determined by the division of resources. Let $\eta \in [0,1]$ be the share of the family income which partner f has control over. Then, each partner solves:

$$\underset{x_f,x_c}{max} \quad U_f(x_f, U_c(x_c)) \text{ s.t. } \eta y = L(p)'(x_f + x_c) \tag{2.2}$$

$$\underset{x_f,x_c}{max} \quad U_m(x_m, U_c(x_c)) \text{ s.t. } (1-\eta) y = L(p)'(x_m + x_c) \tag{2.3}$$

The solutions to these private optimization problems are the demand functions $x_f^*(L(p), \eta y)$

[12]Lindahl prices clear individual-level markets in a public goods setting. See Lindahl (1919) for the original treatment of this topic, and Mas-Colell, Whinston, and Green (1995) for a modern summary of the concept.

and $x_m^*(L(p),(1-\eta)y)$, and $x_c^*(L(p),y,\eta)$ so that parents attain indirect utility from a given η equal to $V_f(L(p),\eta y)$ and $V_m(L(p),(1-\eta)y)$. Here, η is the degree to which partners' preferences are expressed in the aggregate allocation, a natural measure of bargaining power. The second step is to solve for the division of surplus, η, that maximizes the families social welfare function:

$$\underset{\eta\in[0,1]}{max} \quad \tilde{U}\big(V_f(p,\eta y),V_m(p,(1-\eta)y)\big). \tag{2.4}$$

The solution is the family's optimal division of resources between decision makers, η^*, and is a function of real (Lindahl) prices: $\eta^* = \eta^*(y/L(p))$. The solutions obtained by solving (2.2)-(2.4) using backwards induction are the same solutions to the problem in (2.1). That is, both ways to formulate the program give the same optimal household demand, $z^*(p,y)$, and individuals' consumption (Browning, Chiappori, and Lewbel, 2013).

Step 3: Introducing the Possibility of Divorce

Let decision makers also consider divorce. Since half of marriages in Malawi end in divorce within 20 years of forming (Reniers, 2003), this is an important feature of the data generating process. The relative attractiveness of seeking a divorce will determine the individuals' bargaining positions — the value of η^* that solves (2.4). The value of each partners' outside options will depend on the differences in the price and income settings they face in the household setting and in the case of divorce. Outside of the household, individuals only have their private income to allocate instead of a shared budget. Outside of the household, they face market prices instead of Lindahl prices. These differences result in different indirect utilities in the collective allocation and the outside option.

Denote the amount of income an individual would have if he or she divorced as y_m^0 and y_f^0 respectively. These may be different than the earnings they have in the marriage, since their labor supply decisions will be a function of their preferences alone, instead of the outcome of a bargaining process. The indirect utility values that individuals would get in the case of divorce are $V_f(p,y_f^0)$ and $V_m(p,y_m^0)$, whereas the indirect utilities that they would get within the family are $V_f(L(p), \eta y)$ and $V_m(L(p), (1-\eta)y)$. As such, a partner chooses to stay married if the indirect utility that partner receives from marriage is larger than the indirect utility from divorce. The partners solve problems $max\{V_f(L(p), \eta y), V_f(p, y_f^0)\}$ and $max\{V_m(L(p), (1-\eta)y), V_m(p, y_0^m)\}$. These decisions depend on bargaining power in the family, the value of consuming at Lindahl prices instead of market prices, and incomes in the outside options.

To simplify these problems, note that there exists a pair of compensating variation values (γ_f, γ_m) that makes the decision makers indifferent between consuming at the Lindahl and market prices: $V_f(L(p), \eta y) = V_f(p, \eta y + \gamma_f)$ and $V_m(L(p), (1-\eta)y) = V_m(p, (1-\eta)y + \gamma_m)$. By substituting these into the individuals' decisions, their problems will be over functions that differ only in their second arguments. Since their indirect utility functions are increasing in the second arguments, the individuals' decisions to stay married can be written as $max\{\eta y + \gamma_f, y_0^f\}$ and $max\{\eta y + \gamma_m, y_0^m\}$. These problems are not only simpler, they are also measured in dollar amounts, not in utils. These problems can be incorporated into the family's constrained optimization problem as participation constraints:

$$\max_{\eta\in[0,1]} \quad \tilde{U}\big(V_f(L(p), \eta y), V_m(L(p), (1-\eta)y)\big) \tag{2.5}$$

$$\text{s.t. } \eta y + \gamma_f \geq y_0^f \text{ and } (1-\eta)y + \gamma_m \geq y_0^m.$$

By the Kuhn-Tucker theorem, we know that this problem will have a unique solution that satisfies the following properties, suppressing notation:

$$\frac{\partial \tilde{U}}{\partial V_f}\frac{\partial V_f}{\partial \eta^*} = -\frac{\partial \tilde{U}}{\partial V_m}\frac{\partial V_m}{\partial \eta^*} \tag{2.6}$$

$$\eta^*(y/L(p)) \in \left[\frac{y_f^0 - \gamma_f}{y}, 1 - \frac{y_m^0 - \gamma_m}{y}\right].$$

Identifying the sharing rule from (2.6) is possible using a parametric approach. Researchers can assume functional forms for utility functions, the social welfare function, and the consumption technology, and solve for a numerical solution to (2.6). This is the approach in Browning, Chiappori, and Lewbel (2013). A semi-parametric approach is also available. We can avoid these functional form assumptions by making a weaker distributional assumption, and recovering the functional form for the expected value of the solution to (2.6). This is the approach in Klein and Barham (2019), and the approach we take here.

Step 4: Functional Analysis of the Family's Constrained Optimization Problem When Parents' Utility Functions Are Random Variables

Let partners' utility functions be random variables. The researcher does not know what they are before collecting data from a family (and even then cannot observe them) but might know something about the distribution of preferences in a population. For cultural reasons, for example, people in certain families may be more likely to have particular political views or religious beliefs. They may be more likely to consume certain types of food, or be more likely to prefer certain forms of entertainment. Let partners' utility functions be elements from the infinite, compact set of all quasi-concave, increasing, twice-differentiable functions, C. Picking one (of the infinite) functions to describe preferences for each partner would be a parameteric approach to identification. We use a semi-parametric approach by assuming that (U_f, U_m) has a particular joint distribution on $C \times C$.

If utility functions are random variables, then η is a random variable. To see this, note that η is a function of real total income, so $\eta : \mathbb{R}_+ \to \left[\frac{y_f^0 - \gamma_f}{y}, 1 - \frac{y_m^0 - \gamma_m}{y}\right]$. Denote the set of all possible functions $\eta(y/L(p))$ as $\mathbb{H}$. There exists a functional, $\Gamma : C \times C \to \mathbb{H}$, that relates the realized utility functions to the parents' bargaining positions:

$$\Gamma(\tilde{U}(V_f, V_m)) = \left\{ \eta^*(y/L(p)) \;\middle|\; 0 = \frac{\partial \tilde{U}}{\partial V_f}\frac{\partial V_f}{\partial \eta^*} + \frac{\partial \tilde{U}}{\partial V_m}\frac{\partial V_m}{\partial \eta^*} \right\} \tag{2.7}$$

This function is a linear differential operator on the social welfare function. That is, the sharing rule is a linear transformation of random variables. As such, making assumptions about the joint distributions of utility functions on C directly implies a distribution of $\eta^*(y/L(p))$ on $\left[\frac{y_f^0-\gamma_f}{y}, 1-\frac{y_m^0-\gamma_m}{y}\right]$. This in turn implies a specific functional form for the expected value of η^* for any particular family, on the family-specific subset of the unit interval.

Step 5: Making an Assumption on the Joint Distributions of Utility Functions, and Deriving an Estimator for η

We follow Klein and Barham (2019) in making a weak "conditional uniformity" assumption. If we assume that all utility functions (and all preference relations) are equally likely in a population, we would be asserting that no type of preference is more common than others. However, there are patterns that are often predictable in people's preferences, like enjoying certain types of food. Therefore, it seems too strong to assume that all preferences are equally likely. In addition, we might expect that some quasi-concave, increasing, differentiable functions do not describe anyone's preferences. To capture these likely states of the world, we can relax the assumption that all preferences are equally likely by making a conditional uniformity assumption.

Let there be a partition of C comprising an arbitrarily large number of sets. Each of these sets is a subset of C, each has an empty intersection with any other subset, and the union of all subsets in the partition is equal to C. This partition can be drawn in any way. Assume that each family's pair of utility functions is jointly, uniformly drawn from the same subset. Then, saying that some types of preferences are more common than others in a population is analogous to saying that more families draw from particular subsets of C than others.

This imposes a weak "similarity between partners" condition, since their utility functions will be drawn from the same subset of C. This is a weak assumption since this subset can be arbitrarily large (it could be unaccountably infinite) or small (a singleton, in which case partners have identical preferences). In the case of an infinitely large subset, partners' preferences could be arbitrarily similar or dissimilar. Any subset could comprise any "type" of utility function, and exclude any other "type." We need not recover any information about this partition in practice.

This is empirically tractable. Under this conditional uniformity assumption, the expected value of η is the central point of the family-specific bounds on possible power dynamics. It is given by

$$\mathbb{E}\left[\eta_f \middle| \eta_f \sim \text{Uniform}\left[\frac{y_f^o - \gamma_f}{y}, 1 - \frac{y_m^o - \gamma_m}{y}\right]\right] = \frac{1}{2} + \frac{1}{2}\frac{y_0^f - y_0^m + (\gamma_m - \gamma_f)}{y}. \tag{2.8}$$

This functional form has several intuitive properties. When partners have equally valuable outside options, $\eta_f = \frac{1}{2}$ and partners have equal decision-making power. When one partner has a more valuable outside option, they have more power in the family. This partner has additional decision-making power equal to the magnitude of their advantage in the outside option, scaled by household income.

2.2.1 Empirical Strategy

The econometrician can bring the parametric expression of the expected value of bargaining power in equation (2.8) to data in the following way. First, they can predict y_0^f and y_0^m using one of several methods. We use a Heckman selection function. So

$$y_0^f = F(X_f, \psi_f) \text{ and}$$

$$y_0^m = F(X_m, \psi_m)$$

where X_τ is a set of observable characteristics of people and places, and ψ_τ is their latent ability. Second, note that the parameter of interest, η_f, is linear in income, predicted earnings in the outside options, and the difference $\gamma_m - \gamma_f$. We can rewrite equation (2.8) as:

$$y = \frac{F(X_f, \psi_f) - F(X_m, \psi_m) + \gamma}{2\mathbb{E}[\eta_f] - 1} \quad \text{where} \tag{2.9}$$

$$\gamma \equiv \gamma_m - \gamma_f.$$

In order to recover the two unknown parameters from this single equation, we need to use a panel data approach. That is, we need to observe the known parameters in at least two periods of time, denoted with subscript t. Let h denote a specific household. Under a regularity condition on the proportion of η_f to γ over time, we can estimate the following fixed effects model using constrained least squares to recover estimates of γ and η_f:

$$y = \beta_{0,h} + \beta_1(F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})) + \epsilon_{h,t} \text{ s.t.} \tag{2.10}$$

$$\beta_{0,h} = \frac{\gamma_{h,t}}{2\mathbb{E}[\eta_{f,h,t}] - 1}, \text{ and}$$

$$\beta_1 + \epsilon_{1,h,t} = \frac{1}{2\mathbb{E}[\eta_{f,h,t}] - 1}, \text{ and}$$

$$\epsilon_{h,t} = \epsilon_{1,h,t}(F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})), \text{ and}$$

$$\mathbb{E}[\eta_{f,h,t}] \in [0, 1],$$

where the proportionality condition is imposed by the first constraint. This fixed effects model is simply a way of defining terms so that we can fit the model to data and impose the constraint that power be a number between zero and one.

Estimating this model gives parameter estimates of the family-specific fixed effects, $\tilde{\beta}_{0,h}$, and the population-specific slope parameter, $\tilde{\beta}_1$. Plugging these estimates back into the constraints in (2.10) gives functional forms of $\hat{\eta}_{f,h,t}$ and $\hat{\gamma}_{h,t}$ for each family in each period:

$$\hat{\eta}_{f,h,t} = \frac{1}{2} + \frac{1}{2}\frac{F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})}{y_{h,t} - \tilde{\beta}_{0,h}} \tag{2.11}$$

$$\hat{\gamma}_{h,t} = \tilde{\beta}_{0,h}(2\hat{\eta}_{f,h,t} - 1).$$

The intercepts and slope estimates calculated by fitting the model in equation (2.10) are nuisance parameters, and are simply discarded after constructing the estimate of bargaining power. See Klein and Barham (2019) for identification proofs and a more detailed derivation.

The econometrician can estimate equation (2.10) using one of (at least) two equally viable strategies. They can estimate a fixed effects model where they separately specify the constraints on each family's intercept. This is conceptually straightforward but may be computationally challenging. Alternatively, the econometrician can recover consistent estimates of the intercept and slope parameters in equation (2.10) using the simulated method of moments (SMM, see McFadden [1989] for asymptotic properties of SMM estimators). Here, the econometrician must specify a distribution to draw each intercept parameter from, and a distribution for the slope parameter. This is a set of $H + 1$ distributional assumptions, where H is the total number of households in the sample. This large set of assumptions is functionally equivalent to the single assumption on the distribution of the error term in a more straightforward fixed effects regression. This approach has advantages in practice since the econometrician can explicitly define the support for each intercept estimate such that the constraints in equation (2.10) must be met. This is simply an alternative way of specifying the set of feasible solutions to the constrained least squares problem in equation (2.10).

2.3 Intra-Household Inequality and Its Causes in Malawi

We estimate equation (2.10) using the LSMS Integrated Household Panel data from 2010 and 2013.[13] The first step is to predict each partners' earnings in the case of divorce. To do this, and because there is moderate censoring in women's earnings, we estimate two separate Heckman selection models. The first model is fit on women ages 16-65, and the second model is fit on men ages 16-65. This allows us to recover gender-specific selection parameters, and to model different levels of selection into earnings-generating activities by gender. As such, the first regression in each Heckman selection model is a regression of whether or not an individual earns some amount in the labor market on the observable

[13] Download data from the National Statistical Office of Malawi and the World Bank here: https://microdata.worldbank.org/index.php/catalog/2248.

characteristics X_f and X_m. The second reduced form equation is of individuals' earnings on observable characteristics and the selection parameters.

To construct our measure of individuals' total earnings, we sum up each persons' reported earnings from the labor market, from growing crops, from entrepreneurship, from raising livestock, and from other sources. These represent the sum of ways to earn incomes surveyed in the LSMS data. We report the average earnings by category for household heads and their spouses in Table 2.1.[14] We report the percent of household heads and spouses who report earning a positive amount in parentheses. For instance, 77% of adult women earn some positive amount, with 29% of adult women earning incomes from wage labor.

We also present the t-statistics for tests of the null hypotheses that men and women earn equal amounts from each source, and the t-statistics for whether men and women are equally likely to engage in each earnings opportunity. In general, men earn more than women from the same economic activity, and are more likely to engage in lucrative activities like wage labor and entrepreneurship.

We plot the distributions of men's and women's earnings, along with the medians and averages of each distribution in Figure 2.1. The average earnings are higher than the medians for each category, since high-earning outliers drive up the average. We winsorize the earnings variable in our Heckman regressions and in the SMM estimation procedure.

A majority of household heads (77% of women and 95% of men) earn some amount of income from at least one of the eight sources listed in Table 2.1. The most common sources of earnings are rainy season agricultural sales. About 38% of women and 67% of men earn income by selling crops like corn, rice, and tobacco. The next most common is earning a wage in the labor market, and after that, "other sources" is the most common.

The most lucrative occupation for women is as a wage earner. The typical woman who works in the labor market earns far more than the typical woman who earns income from any other source. The next most lucrative profession for women is being an entrepreneur. For men, entrepreneurship has the highest return, followed by earning wages. In gen-

[14]Note that one US dollar was roughly equal to 150 Malawian Kwacha in 2010, and that nation wide GDPPC was roughly $480 in 2010 USD.

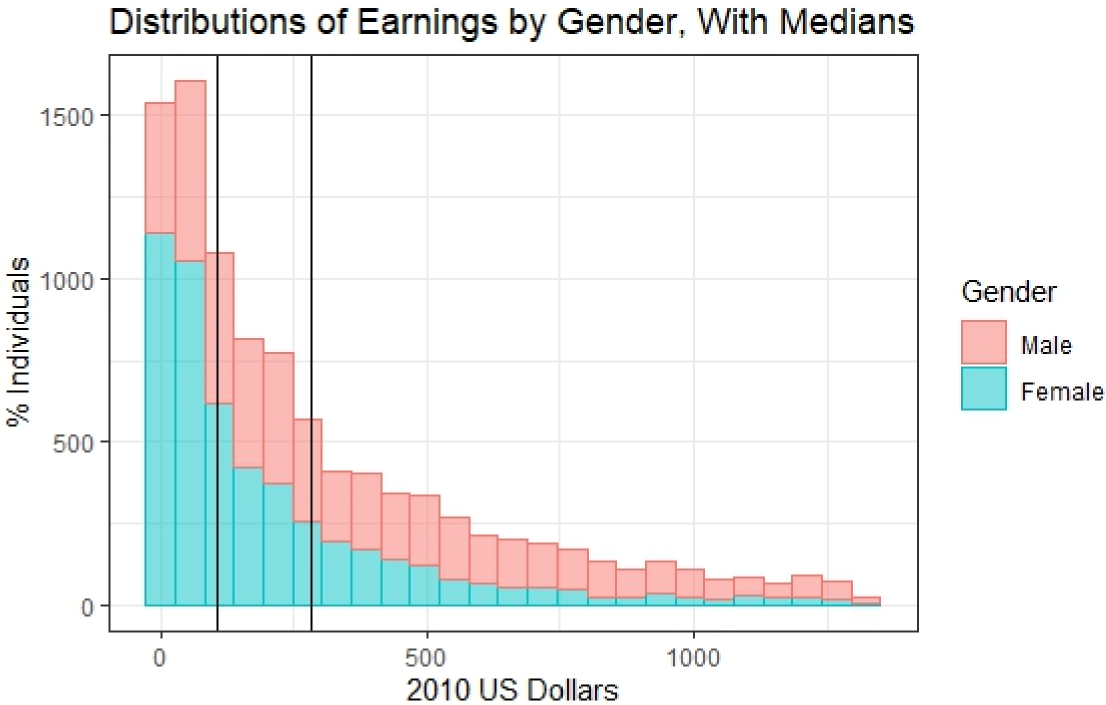

Figure 2.1: Pooled distributions of earnings in 2010 and 2013, by gender, for individuals ages 16-65 in Malawi. The median woman's earnings amount is about 19,500 Kwacha, about \$130 in 2010 USD. The median man's earnings amount is about 64,400, \$425 in 2010 USD. To improve the quality of the visualization, we truncate the data in this figure so that it does not include an income earner who earns more than the top 1% female earner.

Table 2.1: Household Head Earnings Summaries (in 2010 Kwacha) by Gender (Percentage of Individuals with Positive Earnings in Parentheses) with Test Statistics for Gendered Differences

	Female Head		Male Head	t-Statistic
Total Individual Income	97,284.65	$<$	237,657.5	-8.85
	(77.16%)		(94.64%)	(-29.72)
Income from Labor Market	125,613.88	$<$	210,320.7	-5.26
	(28.74%)		(53.55%)	(-28.96)
Income from Rainy Season Crop Sales	20,782.35	$<$	88,240.49	-6.95
	(38.62%)		(66.88%)	(-33.00)
Income from Entrepreneurship	92,561.59	$<$	239,016.16	-2.09
	(12.49%)		(17.89%)	(-8.36)
Income from Other Sources	27,734.09	$<$	29,339.1	-0.36
	(34.97%)		(34.29%)	(0.80)
Income from Livestock Sales	3,135.95	$<$	4,013.76	-0.75
	(15.71%)		(16.91%)	(-1.82)
Income from Livestock Products	1,010.28	$<$	1,219.04	-0.20
	(3.36%)		(3.32%)	(0.12)
Income from Dry Season Crop Sales	198.28	$<$	392.05	-2.81
	(2.48%)		(3.07%)	(-2.01)
Income from Tree Crop Sales	265.68	$<$	527.96	-1.81
	(6.36%)		(5.89%)	(1.11)

eral, interacting with the labor market generates more earnings than selling agricultural products, like livestock or tree crops.

These models are given by the pair of selection and earnings regressions, where i indexes the individual, and $g \in \{f, m\}$ indexes gender and t indexes the year:

$$P_{i,g,t} = \beta^{P}_{i,g,t} X_{i,g,t} + \epsilon_{i,g,t} \tag{2.12}$$

$$ln(E_{i,g,t}) = \beta^{E}_{i,g,t} X_{i,g,t} + \epsilon_{i,g,t}$$

These regressions include time and district fixed effects; age and age squared; the size of plots an individual manages;[15] a dummy for whether the person is the head of the household; family size; the number of men, women, boys, and girls in the family; literacy dummies for English and Chichewa; a dummy variable for whether the individual is currently living where he or she was born; and a dummy variable equal to one for

[15]We drop the top 5% of landholders from these regressions since they drastically reduce the quality of the predictions if included. These outliers manage far larger tracks of land than the median farmer.

Table 2.2: Household Variable Means (with Standard Deviations in Parentheses)

	2010	2013
% Families with Malaria in Last 2 Weeks	29.73%	25.72%
% Live in Rural Area	74.71%	74.84%
Total Expenditures (in 10,000s MWK, 2010)	81.44 (98.58)	81.28 (88.75)
Food Expenditure Share	58.02% (14.57)	59.91% (13.92)
Plot Size in Acres	1.68 (1.58)	1.60 (1.60)
Family Characteristics		
Family Size	4.97 (2.25)	5.18 (2.26)
# of Boys	1.07 (1.08)	0.94 (1.05)
# of Girls	1.10 (1.06)	0.96 (1.05)
# of Children Ages 0-5	0.74 (0.81)	0.62 (0.74)
Fathers		
Age	40.08 (14.30)	40.68 (14.24)
Education (in Years)	6.94 (4.83)	7.35 (5.01)
Chichewa Literacy	76.71%	81.75%
English Literacy	48.03%	51.77%
Mothers		
Age	37.25 (15.58)	37.78 (15.37)
Education (in Years)	4.71 (4.29)	5.24 (4.51)
Chichewa Literacy	59.42%	66.71%
English Literacy	29.34%	33.65%
Water Access		
% of Families w/ Water in House	16.20%	16.85%
% of Families w/ Water Public	77.33%	78.53%
% of Families w/ Water Natural	5.85%	4.03%
% of Families w/ a Private Latrine	65.53%	71.5%
Dwellings		
% of Families w/ Grass Roof	59.09%	53.87%
% of Families w/ Iron Roof	40.63%	45.73%
% of Families w/ Mud Floor	65.36%	62.29%
% of Families w/ Cement Floor	31.80%	34.81%
% of Families w/ Earth Wall	15.16%	10.40%
% of Families w/ Brick Wall	29.02%	27.98%
% of Families w/ Fired Brick Wall	45.22%	54.6%

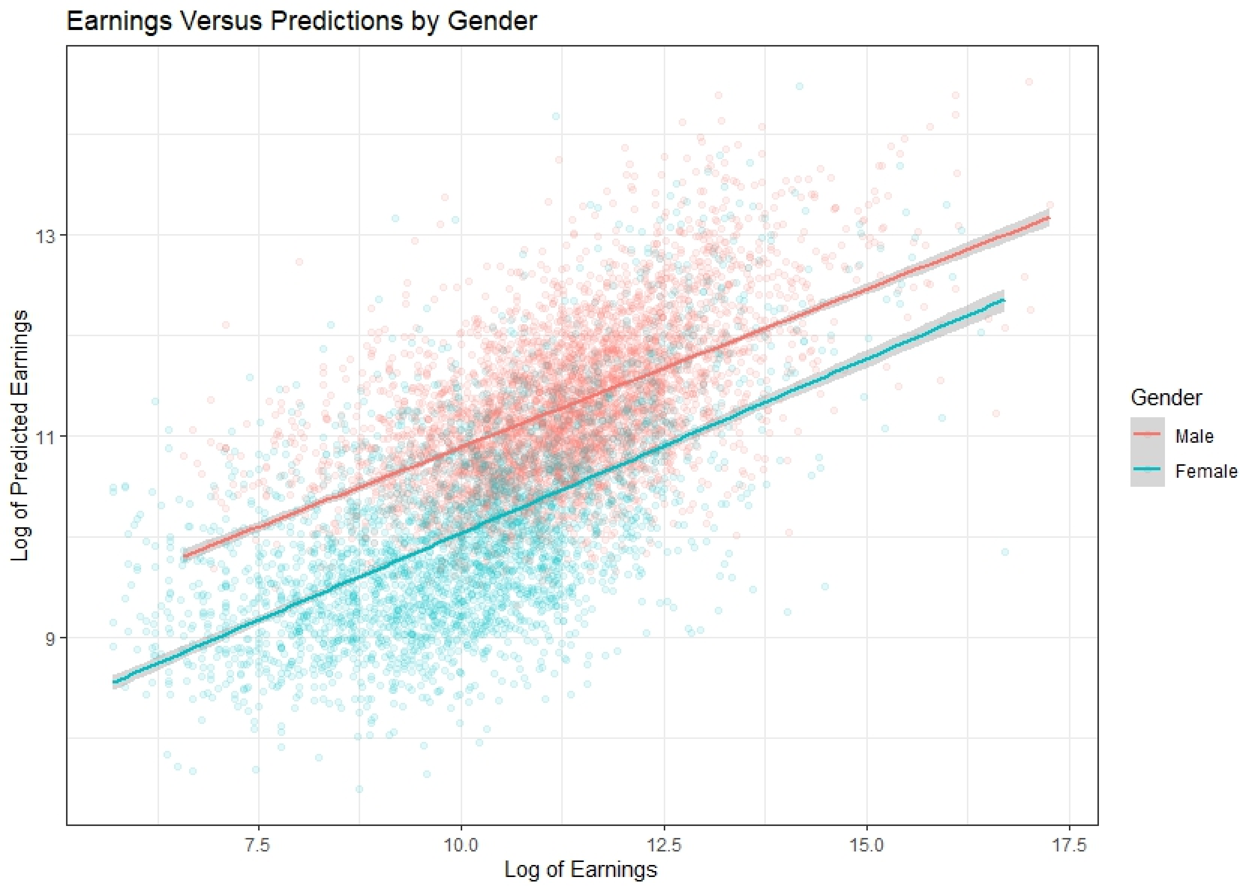

Figure 2.2: Log predictions versus log earnings for people who earn some positive amount. Women have lower observed and predicted earnings than men on average. Linear fits with 95% confidence interval generated using ggplot2 in R (Wickham, 2016).

individuals living in rural areas. We present the summary statistics for these variables, and several others that become salient in our analysis of malaria contraction rates, in Table 2.2. They provide a good sketch of families' living situations in Malawi in 2010 and 2013.

To meet the exclusion restriction for Heckman models, we also include three additional variables in the selection regressions that we do not include in the earnings regressions: distance to daily and weekly markets for each community, and the number of kids under the age of five in the household. People may be less likely to bring crops and other goods to sell in markets that are far away, but once they made the sunk cost investment of traveling to the market, they will work some amount unrelated to the distance from their home. We fit the Heckman regressions using full information maximum likelihood estimation.

We present the predictions versus observed earnings in Figure 2.2, and the selection and earnings regression results in in Appendix B. The results for women and men are in Tables 2.10 and 2.11, respectively. From the selection regressions, we see that being male, being a household head, being older, and having more land are positively correlated with the likelihood that an individual earns some amount. From the earnings regressions we see that household heads earn more, those with more education earn more, families in rural areas have lower earnings, and entrepreneurs earn more. The inverse Mills ratio is positive and significant for women. This suggests that women who choose not to earn would earn less than their peers who do choose to work, if they entered the labor market. They may be incentivized to do so if they choose to divorce. The exclusion restrictions for women are significant, and so these estimates are identified off of the instruments. This is not true for men, but censoring is not a major characteristic of the data generating process for men. These estimates are identified based on the same normality assumptions in typical OLS regressions.

The median predictions for women's earnings, conditional on women working, is 18,270.0 Kwacha, compared to the observed median of 21,688.4 Kwacha. The median prediction for women who do not earn is 8,388.9. The median prediction for men who earn is 60,711.0 Kwacha while the observed median is 67,334.0. The fact that men's predicted earnings are much higher suggest that power dynamics will favor men on average, since the first difference in the numerator of the fraction in (2.11) will be negative. We set the predicted earnings for each household head in each wave equal to the fitted values from the earnings regressions in (2.12):

$$F(X_{f,h,t}, \psi_{f,h,t}) = e^{\hat{\beta}^{E}_{f,t} X_{f,t}}$$

$$F(X_{m,h,t}, \psi_{m,h,t}) = e^{\hat{\beta}^{E}_{m,t} X_{m,t}}$$

Then, we plug in these decision-maker-specific predicted earnings into (2.10), and estimate the fixed effects model using the simulated method of moments.

This estimation process is as follows. In order to produce the $H + 1$ distributional assumptions for the intercept and slope parameters, we fit an unconstrained OLS model to the regression in (2.10). The constraint in (2.10) will bind for any family whose OLS intercept estimate falls into the region of parameter estimates that imply $\eta \notin [0, 1]$. We assume that the distribution of each intercept is a truncated normal, centered on the OLS estimate but with a support that includes only permissible intercept estimates. The standard deviation of these truncated normal variables is equal to the standard error from the OLS estimation. We assume the slope estimate is normally distributed around the OLS estimate, with standard deviation equal to the standard error from the OLS estimation. That is, we specify one truncated normal distribution per family, and one sample-specific normal distribution for the slope estimate. We run 1,000 simulations. The SMM-estimates are those that minimize the sum of squared errors in (2.10) subject to the constraints binding. Denote the household-specific SMM intercept estimates as $\tilde{\beta}_{0,h}$. Then the bargaining power estimate for each family is given by (2.11).

We estimate $\hat{\eta}_{f,h,t}$ for 2045 families in our pooled sample (45% of families headed by a married couple).[16] The median bargaining power measure in this model is 0.26 and the standard deviation is 0.15 (Figure 2.3). This means that the typical wife has 33% of the power that her partner has. A woman with a standard deviation more power than the median woman has about 75% of the power that her husband has. We reject the null hypothesis that men and women have equal bargaining power in favor of the alternative hypothesis that women have less power. The t-statistic for this test is -69.75.

An interesting feature of this model is that we can recover the intrahousehold differences in partner's compensating variation amounts — i.e. the degree to which each partner prefers to use the family's consumption technology instead of a private technology, measured in 2010 Kwacha. If we can think of predicted earnings as the flow of wealth in the outside option, we can think of these compensating variation amounts as the stocks of wealth given divorce. For instance, if the government gave one partner a large cash

[16]The main reason we drop families is if they only appear in one wave of the panel (1479 households, or 34% of the data). We only observe 2,886 families in both time periods. The next biggest reason for dropping families is missing or partial data. In addition, we also drop the 36 families with more than ten members, and the 32 families with more than five rooms in their dwelling. Finally, we drop the two families in the sample who have more than three kids under the age of five. As noted, we Winsorize the income and land variables to improve the earnings predictions.

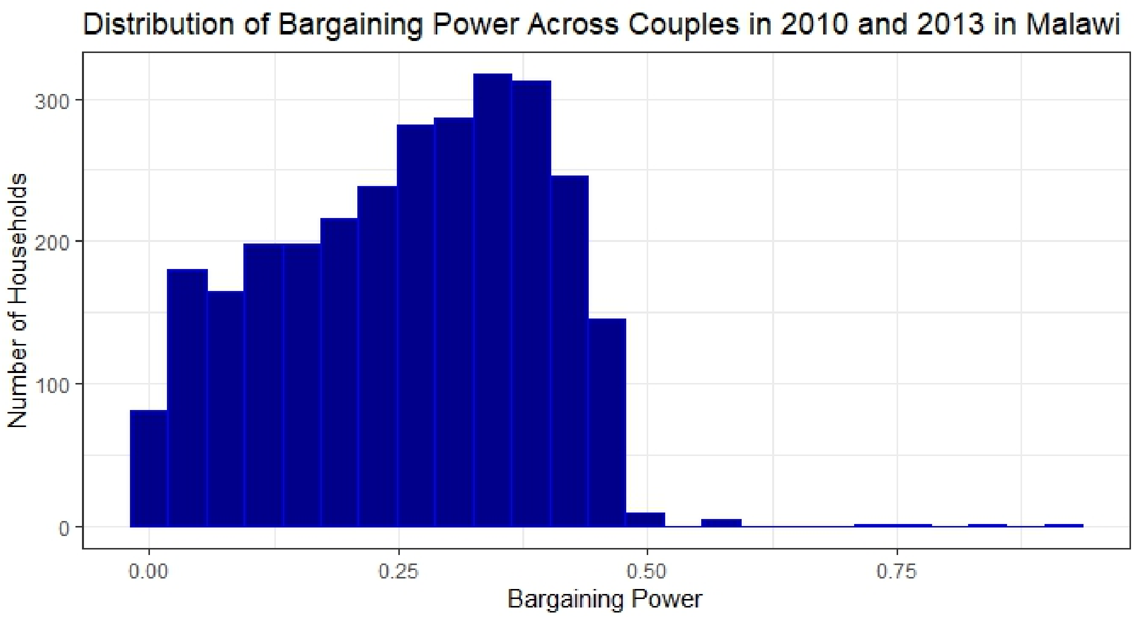

Figure 2.3: Distribution of $\hat{\eta}$ across couples. A measure of $\frac{1}{2}$ indicates that men and women have equal bargaining power in the family. Values lower than $\frac{1}{2}$ indicate that men have more power. When the value is 0, men have all of the decision-making power in the partnership.

transfer at the time of divorce, then that partners' "stock" benefits from divorce would increase, his or her threat point would be more credible, and his or her bargaining power would increase.

We plot the pooled distribution of $\hat{\gamma}$ across couples in Figure 2.4. If $\hat{\gamma} = 0$, the male and female decision makers in the family have the same compensating variation amounts, and their power dynamics will be completely determined by their relative earnings potentials. If the sign of $\hat{\gamma}$ is positive, it indicates one of two things for each family. Either men are better off than women using their private technologies, or both partners would be harmed by divorce but men would be harmed less. The value of $\hat{\gamma}$ indicates that the magnitude of the difference.

The median value of this parameter in the population is 17,541 Malawian Kwacha, or roughly 110 dollars (at 2010 currency values). This indicates that either men are better off than their partners in divorce, or at least they are harmed less by negative phenomena like mental duress and the loss of use of the communal consumption technology. The magnitude of $\hat{\gamma}$ in the median family is economically meaningful, roughly 80% of the median annual income for women. This benefit for men, coupled with their higher predicted earnings in the case of divorce, translate into higher bargaining power in the family for

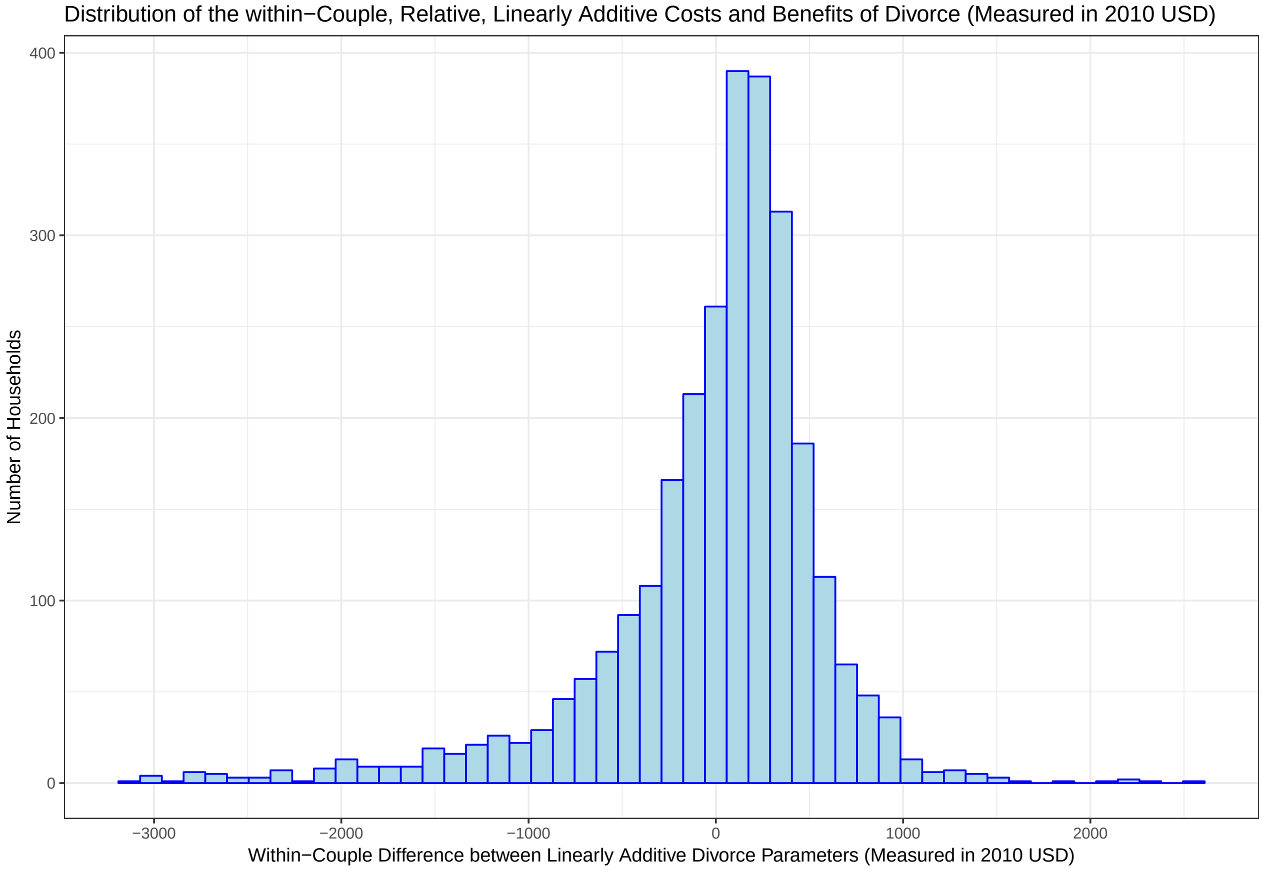

Figure 2.4: Distribution of $\hat{\gamma} \equiv \gamma_m - \gamma_f$ estimates. The units are 2010 USD. Positive values indicate that men are better off than women in the case of divorce, not accounting for earnings flows. The husband in the median family would have a more valuable (or less damaging) one-time gain (or one-time loss) than his wife if they chose to get a divorce.

them.

For many families, $\hat{\gamma}$ takes a negative value. The tails in the distribution are large. This could be because the family consumption technology is particularly valuable for some individuals, compared to consuming at market prices. For instance, if one partner got to keep the family house in the case of divorce, the family $\hat{\gamma}$ value would be skewed in that person's favor, and divorce would be a relatively more credible threat for them. In the tails, these differences will determine power dynamics to a greater extent, since they will make divorce more or less credible threats for one partner relative to the other.

2.4 Differences in Women's Bargaining Power in Matrilineal and Patrilineal Communities

In this section, we first discuss the qualitative differences across matrilineal and patrilineal communities in Malawi. These differences, and qualitative analyses conducted by ethnographers, lead to the hypothesis that men have relatively more bargaining power in the family in patrilineal communities. We then test this hypothesis by estimating our model on the two populations separately, and conducting a simple means comparison.

2.4.1 Traditions in Malawi

Across the world, far more societies observe patrilineal descent traditions than matrilineal descent traditions. However, in Malawi and the greater Zambezi River region, matrilineal descent is relatively common (Davison, 1997; Giuliano and Nunn, 2018). Both cultural practices are observed within the country, with a greater adherence to matrilineal and matrilocal traditions in the south. In our sample, 42% of families practiced the matrilineal and matrilocal tradition in the average community, with a standard deviation of 29 percentage points.

The particular set of customs we study is called *chikamwini*, and is a bundle of matrilineal inheritance, bride price payment, and matrilocal relocation practices. For families that practice the *chikamwini* tradition, lineage is traced through mothers. Individuals are part of their mother's kinship group, and inheritances are passed from mothers to their children. Additionally, at the time of marriage, husbands move into their wives dwellings, often moving away from the community they grew up in, and moving to the community their wife is a part of already. At the time of divorce, should divorce occur, there are no particular expectations concerning the division of assets.

In families that adhere to a patrilocal and patrilineal tradition, these three practices favor men instead of women. Lineage is traced through fathers and inheritances are passed from fathers to to their children. Wives move to husbands communities at the time of marriage, often leaving their social network and joining their husband's established community. At the time of divorce, should it occur, women have to pay back a portion of the bride price they received at marriage.

Together, these three differences significantly alter the position of women in the family. They have better social networks in matrilineal communities, and so may have an easier time earning incomes there than in patrilineal communities. Men might experience the opposite. Having an additional price built into divorce shifts the prices people face in the marriage versus the outside option, and makes them differ across matrilineal and patrilineal communities. We would expect to see this reflected in the $\hat{\gamma}$ parameter we estimate, and we capture such differences. And if women are in control of organizing the inheritance, it may also confer some additional measure of ownership to them, so that they are more likely to keep assets in the case of divorce (also shifting $\hat{\gamma}$ in matrilineal communities). In sum, relative to women in patrilineal communities, we would expect women in matrilineal communities to have more bargaining power. Our model will capture the differences between the two communities if they are reflected in earnings opportunities (afforded by different social network strength) and/or in prices associated with divorce (captured by the compensating variation parameters). It is possible that women have less bargaining power than men in both types of communities. In this case, we would expect there to be greater equality in matrilineal communities.

There is a substantial literature documenting the effects of these kinship practices on family life. For instance, Lowes (2017) shows that matrilineal partners are more likely to hide income from each other in games-in-the-field experiments. She uses a geographical regression discontinuity design to show that families just within Africa's matrilineal belt have healthier family members than families just outside of the belt. Walther (2018) shows that men and women make different investment and labor choices across matrilineal and patrilineal communities in Malawi. These differences confirm her model's predictions that individuals will invest in the value of their outside options in order to gain more control over the family decision-making process.

Ewing et al. (2016) study the qualitative relationship between bargaining power, matrilineal practices, and malaria contraction. They conduct focus group interviews in the southern Chiwawa district and find that couples bargain over when to send a child for medical attention. Women with more bargaining power are able to quickly mobilize family resources to seek medical help for sick children. This reduces the likelihood of superinfection (reinfection while already infected), and speeds the child's transition from "sick with symptoms" to "sick and receiving treatment" to "temporarily immune." Ewing et al. (2016) report that families with less equal partners often wait longer, to see whether a child will recover, before taking them to a doctor, sometimes waiting so long that the child collapses. The timing of medical treatment is just one of the many ways in which bargaining dynamics are salient in understanding who contracts malaria, who suffers from superinfection, and who in the population might be temporarily immune or susceptible. This mechanism seems likely to explain our results since the variable we study is whether any family member was sick in the last two weeks. If more equal parents take their kids to the doctor more quickly, than the length of time they are sick decreases, and the surveyors are less likely to encounter that family during the duration of the illness, mechanically.

Gottlieb and Robinson (2016) study the relationships between gender, political engagement, and kinship traditions. They construct an index of political engagement and show that women in matrilineal communities have higher index values. They also summarize the link between household bargaining power and kinship traditions succinctly:

> "Finally, women in matrilineal societies are likely to have greater intra-household bargaining power vis-a-vis their spouses. This is certainly true when a couple resides matrilocally and a women is surrounded by her family. But it is also likely to be the case no matter the residence location, since matrilineal women have greater exit options than patrilineal women, for whom bridewealth would have been paid, effectively limiting a woman's ability to return home after a failed marriage or spousal death (Schatz, 2002, 2005). As a result, marriage bonds tend to be weaker and divorce rates higher in matrilineal groups (Schatz, 2002), presumably allowing women more power within the marriage (Phiri, 1983)."

Kerr (2005) studies power in Malawi as well. She conducted interviews and focuses on a case study involving a single family. Kerr draws a connection between gendered institutions and traditions, and bargaining power in the household, writing "Wives' unequal position is thus due to a lack of entitlements, such as land, access to employment, support from kin and the state. Some differences between this area of northern Malawi and other studies from central and southern Malawi are due to the different entitlements, particularly control over land and income, which speaks to the enduring implications of different lineage systems in the region." Our approach to identifying the relationship between malaria and matrilineality follows this line of reasoning.

2.4.2 Power Dynamics Across Matrilineal and Patrilineal Communities

Our data permits a study of matrilineality at the community level: the Malawi LSMS panel from 2010 to 2013 has information about the percent of families that practice *chikamwini* and other traditions, but not household-level data on adherence to this practice. We use the percent of families that adhere to the *chikamwini* practices as our instrument for bargaining power. A majority of communities in our sample (62%) are a mixture of matrilineal and matrilocal households or patrilineal and patrilocal households. Some are purely patrilocal and patrilineal (26%) and the remainder are fully matrilineal and matrilocal (12%). We define a matrilocal community as one where at least half of its residents practice the *chikamwini* tradition. About 45% of families in our sample live in matrilocal communities.

In this section, we test the null hypothesis that the average women's intra-household bargaining power in matrilineal communities is the same as the average women's in patrilineal communities against the alternative that it is not. We estimate the model in Section 2 on these two groups separately, and compare means. We reject this null hypothesis ($t =$ 11.29). Women in matrilineal communities have more bargaining power on average (0.28 versus 0.23), but still less power than their husbands. The family power dynamic is closer to equality in matrilineal communities.

We decompose these differences in bargaining power across communities, and find

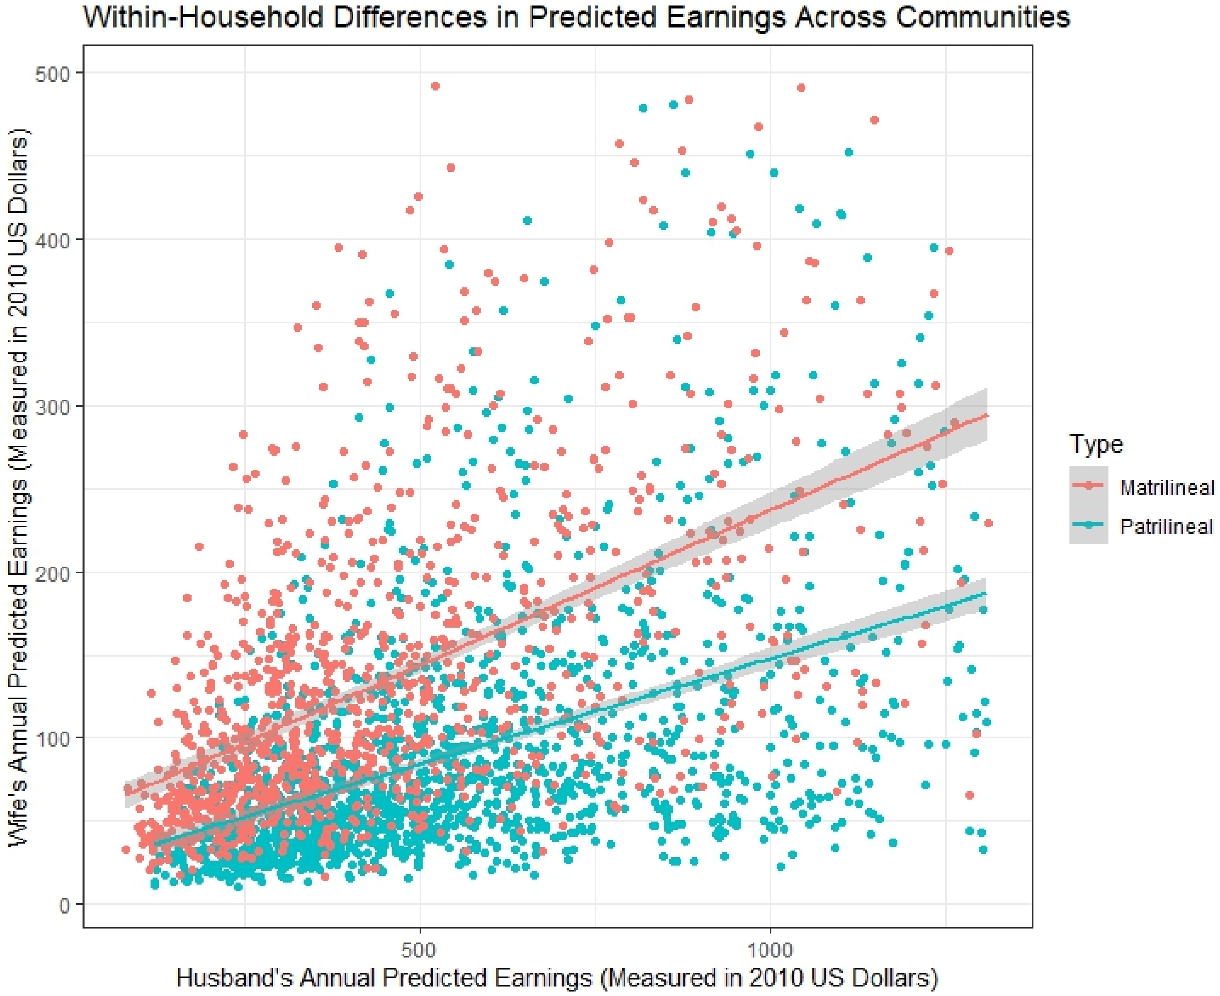

Figure 2.5: Within-household differences in predicted earnings by gender and community type. Women have higher relative predicted earnings in matrilineal communities on average. Linear fits with 95% confidence intervals generated using R's ggplot package (Wickham, 2016). We only display the bottom portion of the distributions to improve the quality of the visualizations.

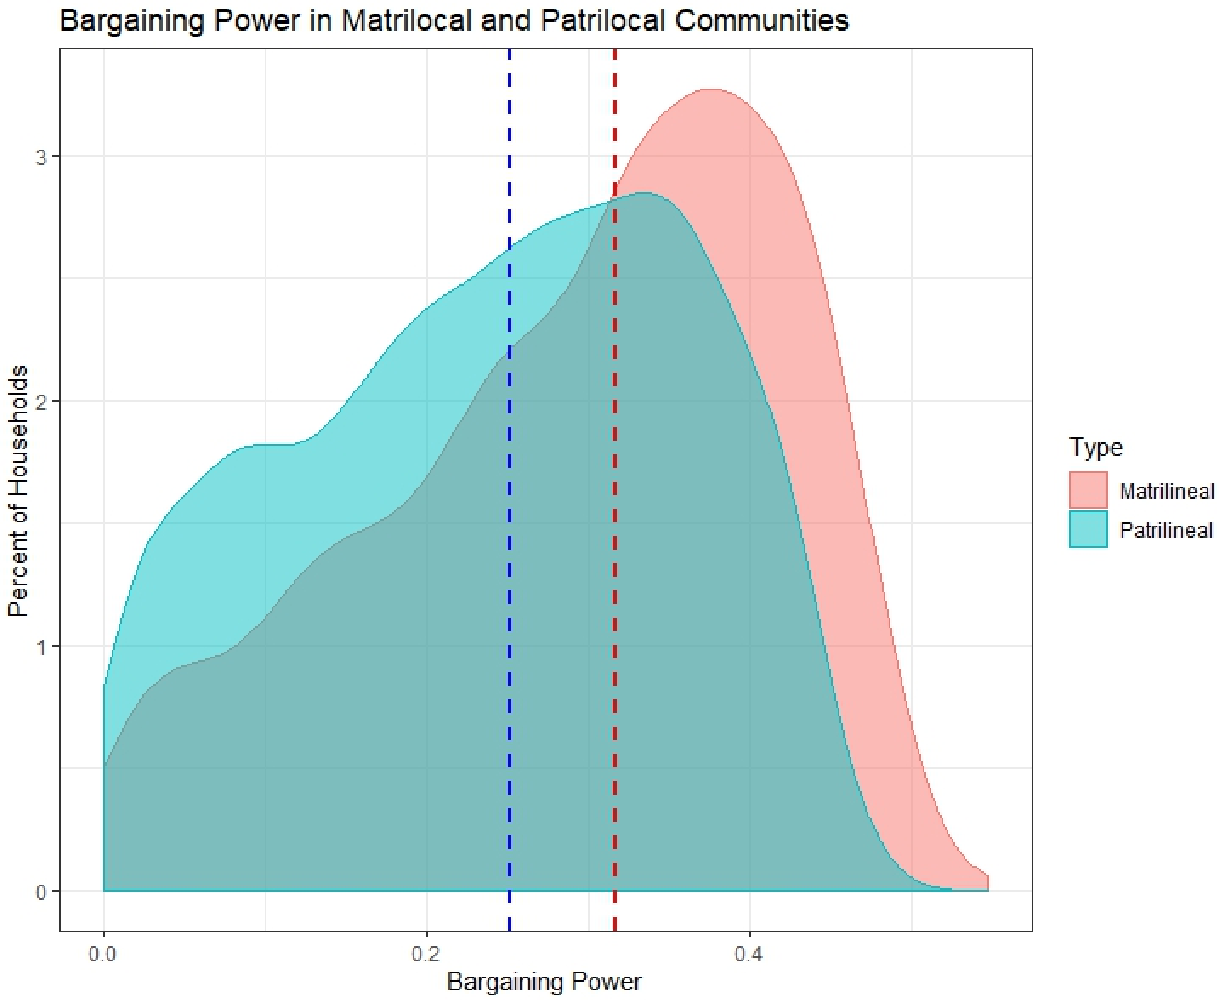

Figure 2.6: Distribution of $\hat{\eta}$ across couples in matrilineal and patrilineal communities with medians given by the dashed vertical lines.

that they are due to both differences in earnings possibilities and in differences in the γ parameter. All three differences between patrilineal and matrilineal communities — relocation, kinship-identity, and bridal price reimbursement — seem to be important in determining bargaining power in the family.

In Figure 2.5, we plot relative within-household predictions for men's and women's earnings in matrilineal (red) and patrilineal (blue) communities. In general, the ratio of women's predicted earnings to their husbands' predicted earnings is higher in matrilineal communities. These differences are potentially attributable to patterns of relocation. In matrilineal communities, men relocate and women have stronger relative social networks. These systematic differences in the degree of connectivity result in a reduction in the gender wage gap in matrilineal communities that is not present in patrilineal communities.

In Figure 2.6, we plot the histograms of intra-household bargaining power in matrilineal (red) and patrilineal (blue) communities. Women tend to have more bargaining power in matrilineal communities, but still less than their partners. We reject the null hypothesis that women have the same average bargaining power across community types, and the null hypotheses that men and women have equal bargaining power in either community type.

In addition, we can test how matrilocal relocation practices influence the relative costs of divorce. The $\hat{\gamma}$ estimates should be lower in communities where men, on average, move to their wives' communities. At the time of divorce a man would, all else equal, have a larger distance to travel to his new post-divorce home. This should translate to a (relatively) worse negative shock than what his ex-wife experiences, at least along this relocation dimension. In addition, women are expected to reimburse their husband's family for a portion of the bride price in the case of divorce in patrilineal communities, adding an explicit difference in the cost of divorce across communities.

We test the null hypothesis that communities wherein more than 50% of the families practice matrilocal relocation have the same mean γ as communities where fewer than 50% of the families practice matrilocal relocation. We reject this null hypothesis ($p = 0.003$) in favor of the alternative hypothesis that the γ values are different. The mean value of $\hat{\gamma}$ is almost 10,000 Kwacha lower in primarily matrilocal communities, and the mean is negative, favoring women. This suggests that costs associated with relocating after marriage are a significant determinant of the relative values of outside options, and thus of women's intra-household bargaining power.

Table 2.3: Regressing the Percent of Matrilineal Families in a Community on Women's Intra-Household Bargaining Power

	Bargaining Power		
	(1)	(2)	(3)
Instrument	0.334^{***}	0.039^{*}	0.030^{**}
	[0.45, 0.49]	<0.016, 0.040>	[0.009, 0.039]
Normalized Family Expenditures		-0.016^{*}	-0.011
		(0.009)	(0.009)
Family Size		-0.004^{***}	-0.004^{**}
		(0.001)	(0.001)
Number of Kids 0-5 yo		0.004	0.004
		(0.003)	(0.003)
Number of Rooms in Dwelling		0.003	0.002
		(0.003)	(0.003)
Year Fixed Effects	NO	YES	YES
Region Fixed Effects	NO	NO	YES
Dwelling Material Fixed Effect	NO	YES	YES
Observations	2,705	2,705	2,705
F Statistic	49.19	46.51	13.52
R^2	0.451	0.819	0.822

Note: 95% bootstrapped confidence intervals in square brackets. Controls include family size, per capita monthly total household expenditures, number of children under the age of 5 in the family, number of rooms in the dwelling, and dwelling material fixed effects. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$. Note that the percentile confidence interval for the first column's parameter estimate is skewed upward, reflecting skew in the empirical distribution of the true parameter value. The interval covers the true population parameter value with 95% likelihood, by definition. The reported F statistics are calculated for regressions with and without the instrument only, not with and without all controls.

In Table 2.3, we regress the percent of families in a community that adhere to the *chikamwini* tradition on bargaining power.[17] In column (1) we do not include any controls; in column (2) we include family size, the number of children in the household, the number of rooms in the dwelling, dummy variables for what the dwelling is constructed from, and year fixed effects as controls. In column (3) we include these basic controls and add in regional fixed effects. We bootstrap the coefficient for the instrument, clustering at the household level and using 500 iterations. We report 95% confidence intervals in square brackets, and 90% confidence intervals in pointed brackets.

Across models, we see a strong correlation between community-level matrilineality and women's bargaining power in the family. The point estimates are consistently positive and significant, and the F-statistics for the instrument are large. With the controls, the point estimate of the relationship between the instrument and bargaining power is consistent with the differences observed in Figure 2.6. Women's median bargaining power in a completely patrilocal and patrilineal community (0.23) is about 30% lower the median women's bargaining power in a completely matrilocal and matrilineal community (0.33).

The percent of families in a community adhering to the *chikamwini* tradition is a strong instrument for women's intra-household bargaining power. Under the assumption that the gender-identity and kinship-identity of the parent who relocates are only related to malaria contraction rates via their effect on household bargaining processes — after controlling for geography, family structure, and income — the percent of families in a community adhering to the *chikamwini* tradition is also a valid instrument for women's intra-household bargaining power.

2.5 The Consequence of Intra-Household Inequality in Malawi: Increased Malaria Incidence

In this section, we present some background information on malaria in Malawi, give a few reasons why equality might change household decision making patterns in ways that

[17]We Winzorize the power variable, and drop families in the top 1% of the following variables: number of kids under the age of five, number of family members, and the number of rooms in the dwelling.

reduce malaria transmission rates, present naive (biased) correlations between predicted malaria contraction rates and intra-household equality, and give the results for our 2SLS model.

2.5.1 Malaria in Malawi

Malaria is one of the most common and lethal diseases in Malawi. From 1980 to 2010, the estimated annual number of deaths from malaria increased from 6,089 to 13,953 (Murray et al, 2012). The population tripled in that time frame, the life expectancy increased by 12 years, and the number of deaths per 1000 attributable to malaria increased by about one-third, from 48 to 67 per 1,000. In 2010, the prevalence (number of cases divided by the population size) of malaria in Malawi was 43% (Malawi Ministry of Health, 2014). Over the next four years, the government and NGOs scaled up antimalaria campaigns, and prevalence fell to 33% in 2014 (Malawi Ministry of Health, 2014). In our sample, 27% of families reported that at least one household member contracted malaria in the last two weeks. Prevalence rates are similar today.

Multiple mosquito populations carry the *P. falciparum* infection in Malawi. Mzilahowa et al (2012) study the number of infected bites per year — the entomological inoculation rate — that people in the Chikwawa district of Malawi suffer. They find that *Anopheles funestus* and three species of *A. gambiae* can pass the disease to humans, and that people typically get 15 infected bites per month. These different types of mosquitoes exhibit different behaviors, so that one type of intervention (e.g., interventions targeted to reduce biting within the home) cannot completely reduce the entomological inoculation rate to zero. Multiple different types of interventions are required.

The human relationship to the *P. falciparum* parasite is complex. In general, people can be in any one of five possible states of infection: susceptible; partially immune; infected with symptoms and not receiving treatment; infected without symptoms and not receiving treatment; or infected and currently being treated, typically with symptoms (Griffin et al, 2010). The movement from one status to another depends on a large number of environmental factors, as well as a huge variety of choices that individuals and families make. These environmental factors and human decisions influence the likelihood and rate at which people pass from one state to another via three mechanisms: the entymological inoculation rate, the transmission efficiency (the likelihood of contracting malaria when bitten by an infected mosquito), and the time required to clear the disease (Smith et al, 2005).

Contemporary control and elimination strategies in Malawi influence malaria prevalence via all three channels. The primary methods of control over the last decade have been the rollout of insecticide-treated bed nets, indoor residual spraying, and artemisinin-based combination therapy, often abbreviated ACT (World Health Organization, 2018). In 2015 and 2016 alone, international organizations and the government distributed more than 10 million nets to the people of Malawi, more than one net for every two people (World Health Organization, 2018). Now, the government is working in tandem with the World Health Organization to rollout the RTS,S malaria vaccine (van den Berg et al, 2019). Bednets and spraying influence the entymological inoculation rate; the vaccine influences transmission efficiency; and ACT influences the rate of recovery. The take-up of all of these technologies might also be a function of decision-making power, since they promote public health within the family.

In addition to understanding transmission as a function of malaria prevention campaigns, we can understand the disease prevalence as a function of family decision-making patterns. Decisions that parents make jointly can also influence all three key mechanisms that mediate malaria transmission. Decisions like where to live, the degree of investment in the dwelling structure, and whether to share a latrine with a neighbor or use a private latrine,[18] can all influence the entymological inoculation rate. Decisions concerning nutrition[19] can influence transmission efficiency. Infants enjoy a brief period of immu-

[18]Sharing a latrine with a neighbor, as opposed to using a private latrine (flush toilet, ventilated improved pit latrine, or outhouse with a roof), exposes family members to an increased risk of being bitten by an infected mosquito, since individuals must travel to more heavily-traversed areas and stay there for a period of time. Dawaki et al (2016) report that families with a flush toilet are 70% less likely to contract malaria than families who use pit or ground dug latrines in Nigeria. Ayele, Zewotir, and Mwambi (2012) find that toilet type matters for malaria contraction rates in Ethiopia and Bamaga et al (2014) find the same using data from Yemen.

[19]There are strong and positive links between gender equality in the family, and quality of diet (e.g., Sraboni and Quisumbing, 2018; Klein and Barham, 2019). There are strong and negative links between nutritional attainment and malaria contraction. There are two pathways for this connection: better health

nity from chemicals inherited from their mothers (Sehgal, Siddiqui, and Alpers, 1989). The extent of this brief immunity may be a function of family decisions that influence maternal health. Decisions about when to seek medical help and how much to invest in treatment can influence the rate of clearing the disease.

2.5.2 Women's Empowerment Can Reduce The Malaria Burden in Malawi

We present two types of findings in this section. First, we predict the likelihood that a household will have at least one member with malaria and plot it against women's intra-household bargaining power in order to see the correlation between equality and malaria contraction rates. We predict this likelihood by fitting values to the regression in column (3) of Table 2.3, excluding bargaining power but including the percent of matrilocal and matrilineal families in a community. We plot this relationship in Figure 2.7. Then, we fit three 2SLS models, where the three regressions in Table 2.3 serve as the first stage regressions. We present these results in Table 2.4. We block bootstrap the coefficient for instrumented bargaining power, clustering at the household level and repeating 500 times.[20] We report the 95% confidence intervals in square brackets.

In Figure 2.7, we see that there is a negative correlation between women's empowerment in Malawi and malaria contraction rates. The purple line is fitted to the entire sample, and has a significant and negative slope. We also see heterogeneity in this relationship across matrilineal and patrilineal communities. We show this relationship for households in matrilineal (red) and patrilineal (blue) communities. In matrilineal communities, an increase in women's bargaining power is more correlated with malaria than in patrilineal communities. This figure demonstrates that increased equality is correlated with a reduction in malaria transmission rates in Malawi.

These correlations are likely biased in unpredictable ways. A simple regression of

reduces the time required to clear the disease and the likelihood of contracting malaria from an infected bite (Lakkam and Wein, 2015).

[20] All block bootstrap algorithms include the Heckman and SMM estimation steps to recover outside option values and bargaining power parameters.

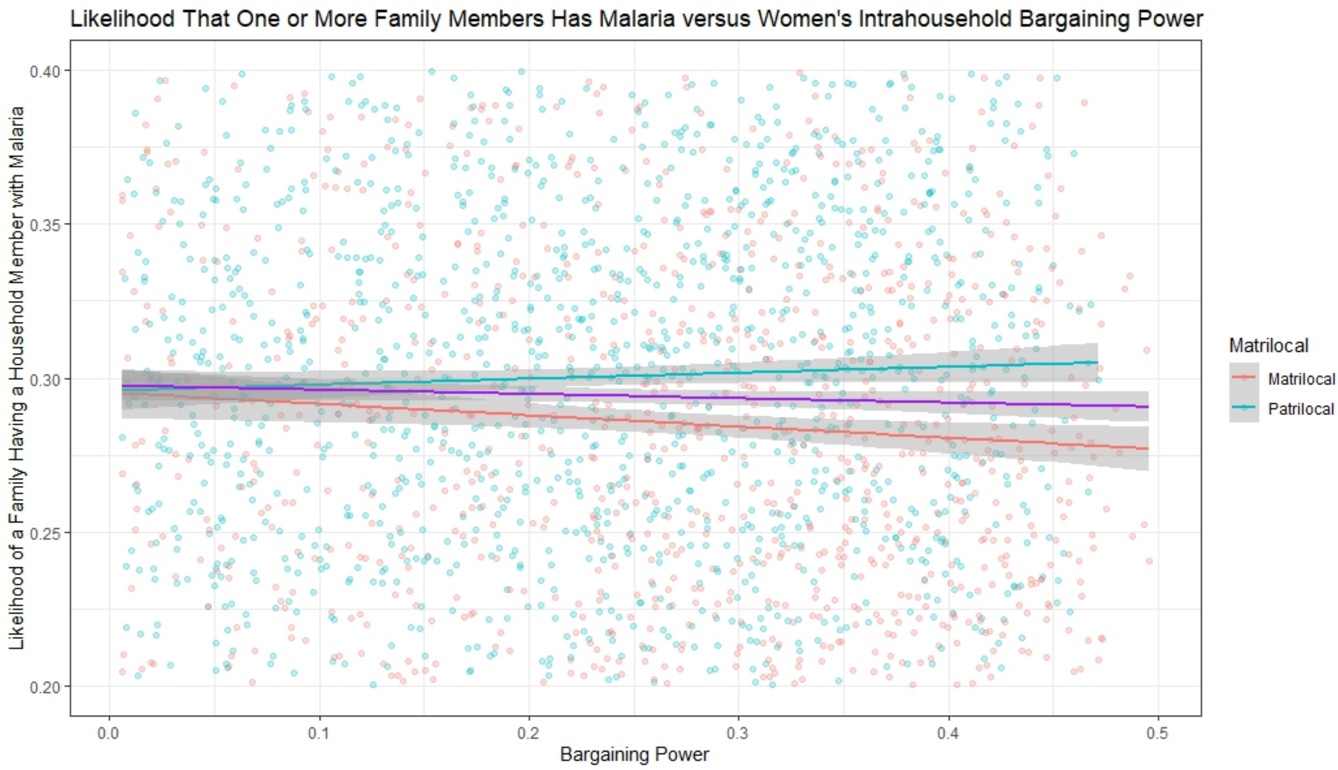

Figure 2.7: Correlation between women's intra-household bargaining power and malaria contraction rates in matrilineal and patrilineal communities in Malawi. The middle, purple fitted line is for the entire sample. The higher, red fitted line is for matrilineal communities only. The lower, blue line is for patrilineal communities only.

Table 2.4: An Increase in Women's Intra-household Bargaining Power Would Reduce Malaria Contraction Rates in Malawi

	Bargaining Power		
	(1)	(2)	(3)
Instrumented Women's Bargaining Power	-3.429^{***}	-2.946^{**}	-8.527^{**}
	[-2.72, -2.13]	[-12.55, -0.22]	[-41.04, -5.63]
Normalized Family Expenditures		-0.041	-0.159
		(0.175)	(0.177)
Family Size		0.082^{***}	0.069^{**}
		(0.027)	(0.028)
Number of Kids 0-5 yo		0.143^{**}	0.152^{***}
		(0.058)	(0.058)
Number of Rooms in Dwelling		-0.002	0.015
		(0.047)	(0.048)
Year Fixed Effects	NO	YES	YES
Region Fixed Effects	NO	NO	YES
Dwelling Material Fixed Effect	NO	YES	YES
Observations	2,716	2,716	2,716
Log Likelihood	$-1,749.014$	$-1,638.334$	$-1,630.665$
Akaike Inf. Crit.	3,500.029	3,294.669	3,283.330

Note: 95% bootstrapped confidence intervals in square brackets. Controls include family size, per capita monthly total household expenditures, number of children under the age of 5 in the family, number of rooms in the dwelling, and dwelling material fixed effects. *p<0.1; **p<0.05; ***p<0.01. Note that the percentile confidence interval for the first column's parameter estimate is skewed downward, reflecting skew in the empirical distribution of the true parameter value. The interval covers the true population parameter value with 95% likelihood.

women's bargaining power on a dummy variable equal to one when someone in the family has malaria will suffer from omitted variable bias. For instance, HIV is relatively common in Malawi. Men and women with more equal decision-making power may be more likely to engage in safe sex practices. This would make more equal parents less likely to have HIV/AIDS, which weakens people's immune systems and increases their risk of infection. Roughly one in ten people in Malawi in 2004 were HIV positive (estimates from the Malawi Demographic and Helath Survey of 2004, reported in Greenwood et al, 2019). Women that are more empowered might be more likely to engage in safe sexual practices (equality might promote sexual health), and whether individuals are HIV/AIDS positive influences the likelihood that they contract malaria by increasing the transmission effectiveness of *P. falciparum* from infected mosquitoes to humans. Without this control, a simple regression will not give us the causal effect of an increase in women's bargaining power on malaria transmission rates.

In fact, without observing each person's susceptibility to infection, the number of infected bites they suffer per week, and their ability to clear the disease, any simple regression may produce biased marginal effect estimates. In these simple regressions, we would identify the marginal effect of interest if we had correctly controlled for all variables that are correlated with power and malaria contraction rates. This is a strong assumption, and almost certainly does not hold (e.g. we do not have reliable data on individuals' HIV/AIDS status). Instead, we make a weaker assumption to identify the causal effect of an increase in equality on malaria contraction rates. We assume that the gender and kinship-identity of the parent receiving cultural benefits in relocation at the time of marriage, kinship-identity, and bridal price reimbursement practices are uncorrelated with malaria contraction rates, after controlling for geography, power, and other observable household characteristics.

In Table 2.4, under this exclusion restriction assumption, we show that there is a strong causal relationship between women's empowerment in Malawi and malaria contraction rates. A standard deviation increase in women's bargaining power would increase the average from 0.25 to 0.38, a point where women would have about 60% of the decision-making power that their partners have. This standard deviation increase in women's bargaining power would result in a 58% decrease in the likelihood that any family member contracts malaria, according to the results in column (3) of Table 2.4. The results in columns (1) and (2) give decreases of 29% and 25% respectively. We conclude that an increase in women's bargaining power would cause a reduction in the malaria burden in Malawi. Our 2SLS results would be biased if there exists some variable that we have not controlled for that is correlated with the percent of families in a community that adhere to a matralineal and matrilocal tradition, and also with the number of family members that had malaria at some time in the last two weeks. Our identifying assumption is that no such variable exists.

2.6 Robustness Checks

Having an instrument that varies at the community level prevents us from including community fixed effects in our 2SLS model.[21] These fixed effects would better control for variation across communities in the conditions that lead to malaria contraction, like the amount of standing water, and the prevalence of infected mosquitoes. Controlling for these differences would make our exclusion restriction assumption more likely to hold.

In order to control for these differences across communities, we can include the percent of a family's neighbors that have at least one member with malaria. This is a very close proxy for the likelihood that a family member contracts malaria. However, it is a "bad" control in the sense that the percent of a family's neighbors with malaria depends on whether the family has members with malaria, which itself depends on the family's power dynamic. So, introducing this variable will introduce some bias to our 2SLS estimates. Introducing this variable has a tradeoff: it controls for local geographic features and so strengthens our model, but it is a "bad" control and so weakens our model. Whether this variable improves our model on net is challenging to know. As such, we include this variable as a robustness check.

[21] Our geographical fixed effects are at a higher level. We control for regions in Malawi, of which there are three. For reference, the United States is 83 times larger than Malawi, and there are counties in the United States that are larger than Malawi's regions.

Table 2.5: Controlling for the Percent of Neighbors With At Least One Member That Had Malaria in the Last Two Weeks

	Dependent variable:	
	First Stage	Second Stage
Instrument	0.046**	
	[0.005, 0.045]	
Instrumented Women's Bargaining Power		−4.851*
		<-26.356, -1.104>
Normalized Income	−0.003	−0.081
	(0.009)	(0.179)
Family Size	−0.003**	0.091***
	(0.001)	(0.028)
Number of Kids 0-5 yo	0.005	0.156**
	(0.003)	(0.061)
Number of Rooms in Dwelling	−0.003	0.009
	(0.003)	(0.050)
% Nieghbors with Malaria	0.036**	2.545***
	(0.017)	(0.322)
Year Fixed Effects	YES	YES
Region Fixed Effects	YES	YES
Dwelling Material Fixed Effects	YES	YES
Hospital in Community Fixed Effect	YES	YES
Observations	2,716	2,716
R^2	0.816	
Adjusted R^2	0.815	
Log Likelihood		−1,596.486
Akaike Inf. Crit.		3,218.971
F Statistic	32.49	

Note: 95% bootstrapped confidence intervals in square brackets and 90% bootstrapped confidence intervals in pointed brackets. Controls include family size, per capita monthly total household expenditures, percent of neighbors with malaria in the last two weeks, number of children under the age of 5 in the family, number of rooms in the dwelling, dwelling material fixed effects. *p<0.1; **p<0.05; ***p<0.01. Note that the percentile confidence interval for the first column's parameter estimate is skewed downward, reflecting skew in the empirical distribution of the true parameter value. The interval covers the true population parameter value with 95% likelihood. The reported F statistic is calculated for regressions with and without the instrument only, not with and without all controls.

We present the results for this robustness check in Table 2.5. The first stage results are positive and significant at the 95% confidence level, and the F-statistic is large (32.49). The second stage coefficient is negative, large, and significant at the 90% confidence level. According to this result, a one standard deviation increase in women's bargaining power would result in a decrease in the likelihood that at least one family member had malaria in the last two weeks by 37.85%. This robustness check suggest that our standard 2SLS model results are not driven by coincidental correlation between malaria prevalence and traditions *within* or *across* Malawi's regions.

2.6.1 Two Alternative Bargaining Power Measures

We also produce two alternative measures of women's bargaining power in the family and use them in our 2SLS model. The first alternative is a proxy variable constructed from information on individuals' education level, age, income, landholdings, and literacy, and using principal component analysis to reduce this data to a single parameter. The second alternative is a ratio of women's resource shares to the couples' summed resource shares, where we recover resource shares from the standard DLP model. We give detailed explanations of how we derive both alternatives in Appendix A. We prefer our main specification to these two alternatives because we can use it to measure the linkages between the instrument, outside options, and bargaining power. This makes our main specification relatively attractive because the threat of divorce is an important determinant of bargaining power in this context.[22]

The average of the resource share ratio measure is 0.45, meaning that women have roughly 80% of the decision-making power that their partners have. The average of the PCA measure is 0.41, meaning that women have roughly 70% of the decision-making power that their partners have. We reject the null hypothesis that men and women have equal bargaining power in the median family using both alternatives. We hypothesize that there are substantive differences in the degree of inequality across the three measures of

[22]See Cherchye et al (2018) for an analysis of how the benefits of divorce (captured in our model via the differences in shadow prices and incomes in the two settings) predict divorce rates in the LSMS 2010 and 2013 panel data.

bargaining power that we estimate because divorce is an important mediating factor in power dynamics in the family in Malawi, and only our preferred measure actually models this feature of partners' decision-making processes.

We use these alternative measures in our 2SLS model, and find similar first-stage coefficient estimates for the relationship between the instrument and the power measure. However, for the resource share measure, the coefficient estimate is not significantly different from zero — the percent of families in a community that adhere to a matrilineal and matrilocal tradition is a weak instrument for the ratio of a woman's resources share to the couples summed resource shares.

In the second stage regressions for the PCA model, we find corroborating results: a standard deviation increase in women's bargaining power would result in an 80% reduction in the likelihood that the family has one or more members with malaria. We also explore naive correlations between the DLP measure and malaria contraction rates and find results similar to those in Figure 2.7: women's empowerment is negatively correlated with the likelihood that family members contract malaria.

2.7 Policy Recommendation

We can conduct a counterfactual analysis to gain insights about the potential effects of an unconditional cash transfer targeted to female decision makers. In particular, we can add some amount of income to women's outside option values (and the family's total income), then recalculate bargaining power using equations and (2.10) and (2.11). We can use the marginal effect estimates from column (3) of Table 2.3 to analyze how this empowerment would translate into a reduction of malaria prevalence.

We can analyze the various effects of the same sized cash transfer in matrilineal and patrilineal communities separately by estimating equations (2.10) and (2.11) for the two groups separately, as in Section 4.2. We know from the first stage results of our 2SLS model that power dynamics will favor men more in partilineal communities, and so would expect the same sized cash transfer to engender a larger increase in women's bargaining power in these communities.

We analyze hypothetical transfer amounts that range from four USD a month to 54 USD a month, and report the results in Figure 2.8. The horizontal axis in Figure 2.8 shows the amount of additional bargaining power women would have at certain transfer levels. The vertical axis reports the predicted malaria prevalence in that hypothetical context. Four dollars a month, given directly to the female household head, would generate small increases in women's bargaining power, with larger increases in patrilineal communities. Doubling the transfer size more than doubles the empowerment effects, and begins to generate large, meaningful decreases in malaria transmission. A transfer of roughly 25 USD per month to mothers would equate women's predicted earnings and men's, and produce gender equality in the family. This large increase in gender equality in the family would translate to a 80% reduction in malaria transmission. This magnitude transfer in 2010 would have tripled the median women's income. This is an expensive but powerful anti-malarial measure.

We propose that such a cash transfer be rolled out, using a RCT to test effects in a trial group and scaling up accordingly. Such a RCT could enable the study of spillovers within communities by randomizing the presence and values of transfer at both the community and family level. This would also allow researchers to test whether the curve in Figure 2.8 is empirically supported, in both steepness and curvature. The percent of families given cash transfers in each community could potentially be varied in order to measure how empowerment spillovers influence the reduction of malaria transmission. That is, researchers could answer how untreated families benefit from the neighboring female decision-makers receiving cash transfers. We would expect that there would be positive spillover effects since malaria is a contagious disease.

This transfer could be a conditional cash transfer (CCT) or an unconditional cash transfer (UCT). A hypothetical condition for mothers receiving the transfer could be having children who routinely get checked by doctors and found not to have malaria. It is possible that such conditionality would increase the malaria-reduction effect of an empowerment-oriented cash transfer. Communities could be randomly allocated to a CCT or a UCT to study the effects of such conditions, as in Baird, McIntosh, and Özler (2011).[23] Additionally, conditions could vary randomly across communities to determine which are most effective.

A cash transfer would be an expensive way to reduce the malaria burden relative to other approaches, like distributing bed nets. This is because the effect of any transfer-based empowerment program will be proportional to the costs of living in the case of divorce. However, cash transfers targeted to female household heads have multiple benefits. First, they increase gender equality, which is a social good in itself. Second, they can have broader health effects. For instance, it is possible that mothers who receive transfers will take children to doctors more frequently, as documented in the focus group interviews conducted by Ewing et al (2016). This may lead to diagnoses and treatment of other conditions. Third, cash transfers also have income effects. Families' budget sets shift outward in addition to power dynamics changing, and families can have higher consumption and investment levels.

In addition, we can artificially vary the number of years of schooling that women have and assess the effects on predicted earnings. Increasing a year of schooling makes women more likely to work (see Table 2.10) and earn more when they work (see Table 2.11). These changes translate into about a 4% increase in women's bargaining power, equivalent to the empowerment realized by a \$2/month transfer to women. Mothers have very low levels of education to start with, about 5 years on average (compared to 7 years for fathers). Women's primary occupation, when they earn some amount, is agriculture. An additional year of schooling (relative to an additional year of experience) does not increase women's earnings prospects substantially in Malawian in 2010 and 2013. For this reason, considering intra-household bargaining power dynamics, increasing women's

[23]These authors ran a cash transfer experiment in Malawi. They gave cash transfer to girls and parents at the same time, and larger transfers to parents (4-10 USD per month) than to girls (1-5 USD per month). If we used our model to study how this transfer would lead to empowerment for girls, we would find that everyone's outside options are improving in this set of transfers. Since the transfer for parents are larger, it is challenging to predict, in theory, how these transfers would empower girls. Further, defining girls' outside options is challenging. If parents' outside options are divorce, and girls' outside options do not include divorce, then the expected value of bargaining power is a complex function of parents' compensating variation values. In practice, a transfer targeted to one person in the family is more likely to result in an increase in their bargaining power. Modeling this is an interesting avenue of future research, however, since children may participate in bargaining over schooling outcomes (Huntington-Klein, 2018).

educational attainment is a limited policy tool.

There is potentially a trade off between investing in empowerment programs and investing in other anti-malaria efforts, like providing free indoor residual spraying. As such, it is natural to ask what the optimal allocation of funding is over the set of all possible anti-malaria projects. Our analysis suggests that the optimal portfolio of anti-malaria spending has some positive expenditures allocated to women's empowerment programs. Current funding for anti-malaria programs could be efficiently reorganized such that some positive amount is spent on empowerment programs. We leave for future research an analysis of the optimal split of resources across all possible anti-malaria projects in Malawi.

2.8 Conclusion

The Malawian government, with support from international organizations like the World Health Organization and the USAID-funded President's Malaria Initiative, is working to reduce the malaria burden in many ways. Over the last five years, donors spent more than $100 million to control and eliminate malaria in Malawi (World Health Organization, 2018). Funded projects include the dissemination of treated bed nets, indoor residual spraying, improved treatment for those diagnosed with Malaria (World Health Organization, 2018), text message campaigns to improve medical professionals' performance (Steinhardt et al, 2019), and, recently, a vaccination campaign (van den Berg et al, 2019). All of these tools are useful in reducing the malaria burden, and in tandem may be more effective than individual efforts. Nonetheless, the malaria burden remains high. Additional tools to reduce malaria transmission may be necessary to continue lowering malaria prevalence, and one day, eradicate the disease in Malawi.

We show that increasing gender equality reduces malaria transmission. We estimate that a standard deviation increase in women's decision-making power would reduce the likelihood that a family member contracts malaria by 58%. Some examples of programs that have increased women's bargaining power elsewhere are gender-targeted cash transfers (Duflo, 2003; and Tommasi, 2019), increased equality of asset division and rights in the case of divorce (Voena, 2015), and vocational training (Bandiera et al, 2018). We suggest that the government of Malawi and other donors working to eradicate malaria consider these policies as additional methods to combat the disease. Like the text message campaigns documented by Steinhardt et al (2019), these types of interventions are not based on public health or epidemiological research, but on social science. Combining methods with foundations in many disciplines may be key to reducing the malaria burden further.

We suggest pairing conventional public health policies with women's empowerment policies, piloting their coupling with regional randomized controlled trials, and scaling up in accordance with the experimental results. This work could be conducted in tandem with the Department of Gender Affairs in Malawi. In addition, we hope that future studies analyze the external validity of our results to other countries in the region. It is possible that increased gender equality could reduce the malaria burden elsewhere in sub-Saharan Africa, and around the world.

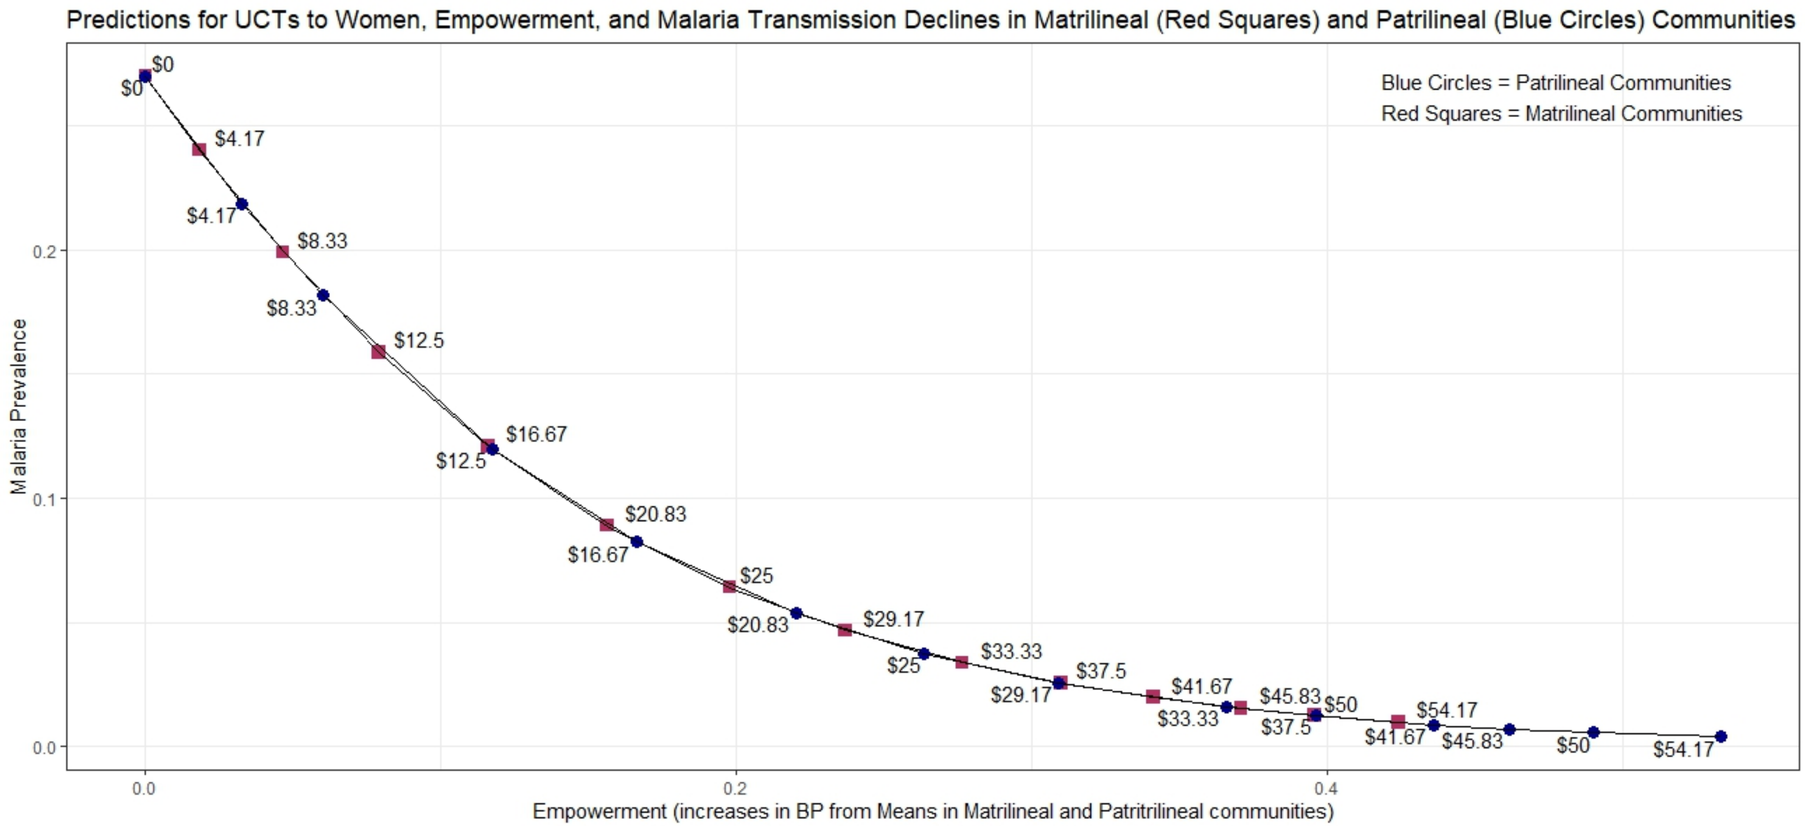

Figure 2.8: Counterfactual analyses of cash transfers to female decision makers. Larger transfer generate larger increases in women's bargaining power, which translate into larger reductions in malaria prevalence. The red squares are counterfactual analysis results for matrilineal communities, defined as communities where more than 50% of families adhere to a matrilineal and matrilocal tradition. The blue circles report the results for patrilineal communities. We use the marginal effect reported from the model in column (3) of Table 2.3 to consider how empowerment translates into malaria reductions.

Reference List

Agarwal, B. (1997). "Bargaining" and Gender Relations: Within and beyond the Household. Feminist Economics, 3(1), 1-51.

Alam, Md Moshi Ul & Monica. (2018). Women Empowerment and Household Sanitation: Bounds in ATE when Treatment Group is Partially Observed. Working Paper.

Annan, J. Donald, A., Goldstein, M., Gonzalez M., Paula L., Koolwal, G. (2019). Taking Power: Women's Empowerment and Household Well-Being in Sub-Saharan Africa (English). Policy Research Working Paper; no. WPS 9034. Washington, D.C.: World Bank Group.

Apps, P. F., & Rees, R. (1988). Taxation and the Household. Journal of Public Economics, 35(3), 355-369.

Attanasio, O., & Lechene, V. (2010). Conditional Cash Transfers, Women and the Demand for Food (No. 10, 17). IFS working papers.

Attanasio, O., & Lechene, V. (2014). Efficient Responses to Targeted Cash Transfers. Journal of Political Economy, 122(1), 178-222.

Ayele, D. G., Zewotir, T. T., & Mwambi, H. G. (2012). Prevalence and Risk Factors of Malaria in Ethiopia. Malaria Journal, 11(1), 195.

Baird, S., McIntosh, C., & Özler, B. (2011). Cash or Condition? Evidence From a Cash Transfer Experiment. The Quarterly Journal of Economics, 126(4), 1709-1753.

Bamaga, O. A., Mahdy, M. A., Mahmud, R., & Lim, Y. A. (2014). Malaria in Hadhramout, A Southeast Province of Yemen: Prevalence, Risk Factors, Knowledge, Attitude and Practices (KAPs). Parasites & Vectors, 7(1), 351.

Bandiera, O., Buehren, N., Burgess, R., Goldstein, M., Gulesci, S., Rasul, I., & Sulaiman, M. (2018). Women's Empowerment in Action: Evidence from a Randomized Control Trial in Africa. World Bank.

Bargain, O., Lacroix, G., & Tiberti, L. (2018). Validating the Collective Model of Household Consumption Using Direct Evidence on Sharing. Partnership for Economic Policy Working Paper, (2018-06).

Beegle, K., Frankenberg, E., & Thomas, D. (2001). Bargaining Power Within Couples and Use of Prenatal and Delivery Care in Indonesia. Studies in Family Planning, 32(2), 130-146.

Benería, L., & Roldán, M. (1987). The Crossroads of Class and Gender: Industrial Homework, Subcontracting, and Household Dynamics in Mexico City. University of Chicago Press.

Blundell, R., Chiappori, P. A., & Meghir, C. (2005). Collective Labor Supply with Children. Journal of Political Economy, 113(6), 1277-1306.

Browning, M., Bourguignon, F., Chiappori, P. A., & Lechene, V. (1994). Income and Outcomes: A Structural Model of Intrahousehold Allocation. Journal of Political Economy, 102(6), 1067-1096.

Browning, M., & Chiappori, P. A. (1998). Efficient Intra-Household Allocations: A General Characterization and Empirical Tests. Econometrica, 1241-1278.

Browning, M., Chiappori, P. A., & Lewbel, A. (2013). Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power. Review of Economic Studies, 80(4), 1267-1303.

Calvi, R. (2017). Why Are Older Women Missing in India? The Age Profile of Bargaining Power and Poverty. Working Paper.

Calvi, R., Lewbel, A., & Tommasi, D. (2017). LATE With Mismeasured or Misspecified Treatment: An Application to Women's Empowerment in India. Working Paper.[1341].

Cherchye, L., De Rock, B., Walther, S., & Vermeulen, F. (2018). Where Did It Go Wrong? Marriage and Divorce in Malawi. in *Marriage and Divorce in Malawi* (July 1, 2018) Available at SSRN: https://ssrn.com/abstract=2759255 or http://dx.doi.org/10.2139/ssrn.27

Chavas, J. P., & Klein, M. J. (2020). Estimating Intra-Household Bargaining Power When Outside Options Are Endogenous (No. 595). University of Wisconsin, Agricultural and Applied Economics.

Chiappori, P. A. (1988). Rational Household Labor Supply. Econometrica, 63-90.

Chiappori, P. A. (1992). Collective Labor Supply and Welfare. Journal of Political Economy, 100(3), 437-467.

Chiappori, P. A., & Ekeland, I. (2006). The Microeconomics of Group Behavior: General Characterization. Journal of Economic Theory, 130(1), 1-26.

Chiappori, P. A., & Mazzocco, M. (2017). Static and Intertemporal Household Decisions. Journal of Economic Literature, 55(3), 985-1045.

Davison, J. (1997). Gender, Lineage, and Ethnicity in Southern Africa. Boulder, CO: Westview Press.

Dawaki, S., Al-Mekhlafi, H. M., Ithoi, I., Ibrahim, J., Atroosh, W. M., Abdulsalam, A. M., ... & Ahmed, A. (2016). Is Nigeria Winning the Battle Against Malaria? Prevalence, Risk Factors and KAP Assessment Among Hausa Communities in Kano State. Malaria Journal, 15(1), 351.

Doss, C. (2013). Intrahousehold Bargaining and Resource Allocation in Developing Countries. The World Bank.

Duflo, E. (2003). Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. The World Bank Economic Review, 17(1), 1-25.

Dunbar, G. R., Lewbel, A., & Pendakur, K. (2013). Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi. American Economic Review, 103(1), 438-71.

Ewerling, F., Lynch, J. W., Victora, C. G., van Eerdewijk, A., Tyszler, M., & Barros, A. J. (2017). The SWPER index for women's empowerment in Africa: development and validation of an index based on survey data. The Lancet Global Health, 5(9), e916-e923.

Ewing, V. L., Tolhurst, R., Kapinda, A., Richards, E., Terlouw, D. J., & Lalloo, D. G. (2016). Increasing Understanding of the Relationship Between Geographic Access and Gendered Decision-Making Power for Treatment-Seeking for Febrile Children in the Chikwawa District of Malawi. Malaria Journal, 15(1), 521.

Fernando, S. D., Rodrigo, C., & Rajapakse, S. (2010). The 'Hidden' Burden of Malaria: Cognitive Impairment Following Infection. Malaria Journal, 9(1), 366.

Giuliano, P., & Nunn, N. (2018) Ancestral Characteristics of Modern Populations. Economic History of Developing Regions, 33 (1), 1-17.

Glennerster, R., Walsh, C., & Diaz-Martin, L. (2018). A Practical Guide to Measuring Women's and Girls' Empowerment in Impact Evaluations. Gender Sector, Abdul Latif Jameel Poverty Action Lab.

Gottlieb, J., & Robinson, A. L. (2016). The Effects of Matrilineality on Gender Differences in Political Behavior Across Africa. University of California, Berkeley, mimeo.

Greenwood, J., Kircher, P., Santos, C., & Tertilt, M. (2019). An Equilibrium Model of the African HIV/AIDS Epidemic. Econometrica, 87(4), 1081-1113.

Griffin, J. T., Hollingsworth, T. D., Okell, L. C., Churcher, T. S., White, M., Hinsley, W., ... & Ghani, A. C. (2010). Reducing Plasmodium Falciparum Malaria Transmission in Africa: a Model-Based Evaluation of Intervention Strategies. PLoS Medicine, 7(8), e1000324.

Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. Econometrica, 153-161.

Heijnen, M., Cumming, O., Peletz, R., Chan, G. K. S., Brown, J., Baker, K., & Clasen, T. (2014). Shared Sanitation Versus Individual Household Latrines: A Systematic Review of Health Outcomes. PloS one, 9(4), e93300.

Henningsen, A., & Hamann, J. D. (2007). Systemfit: A Package for Estimating Systems of Simultaneous Equations in R. Journal of Statistical Software, 23(4), 1-40.

Holland, C., & Rammohan, A. (2019). Rural Women's Empowerment and Children's Food and Nutrition Security in Bangladesh. World Development, 124, 104648.

Huntington-Klein, N. (2018). College Choice as a Collective Decision. *Economic Inquiry*, 56(2), 1202-1219.

Kerr, R. B. (2005). Food Security in Northern Malawi: Gender, Kinship Relations and Entitlements in Historical Context. Journal of Southern African Studies, 31(1), 53-74.

Klein, M. J., & Barham, B. L. (2019). Point Estimates of Household Bargaining Power Using Outside Options (No. 590). University of Wisconsin, Agricultural and Applied Economics.

Kublin, J. G., Patnaik, P., Jere, C. S., Miller, W. C., Hoffman, I. F., Chimbiya, N., ... & Molyneux, M. E. (2005). Effect of Plasmodium Falciparum Malaria on Concentration of HIV-1-RNA in the Blood of Adults in Rural Malawi: a Prospective Cohort Study. The Lancet, 365(9455), 233-240.

Lakkam, M., & Wein, L. M. (2015). Analyzing The Nutrition-Disease Nexus: The Case of Malaria. Malaria Journal, 14(1), 479.

Lépine, A., & Strobl, E. (2013). The Effect of Women's Bargaining Power on Child Nutrition in Rural Senegal. World Development, 45, 17-30.

Li, L., & Wu, X. (2011). Gender of Children, Bargaining Power, and Intra-household Resource Allocation in China. Journal of Human Resources, 46(2), 295-316.

Lindahl, E. (1919). Die Greechtigkeit der Besteuring. Lund: Gleerup [English Translation: "Just Taxation: — A Positive Solution. In "Classics in the Theory of Public Finance" Edited by Musgrave, R. A. and Peacock, T. A.. London: Macmillan, 1958.]

Lindblade, K. A., Mwandama, D., Mzilahowa, T., Steinhardt, L., Gimnig, J., Shah, M., ... & Zoya, J. (2015). A Cohort Study of the Effectiveness of Insecticide-Treated Bed Nets to Prevent Malaria in an Area of Moderate Pyrethroid Resistance, Malawi. Malaria Journal, 14(1), 31.

Lowes, S. (2017). Matrilineal Kinship and Spousal Cooperation: Evidence From The Matrilineal Belt. Unpublished Manuscript.
URL: https://scholar.harvard. edu/files/slowes/files/lowes_matrilineal.pdf.

Lise, J., & Yamada, K. (2019). Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use. The Review of Economic Studies, 86(5), 2184-2219.

Macdonald, G. (1957). The Epidemiology and Control of Malaria. Bull World Health Organ. 15:613â26.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). Microeconomic Theory (Vol. 1). New York: Oxford university press.

Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. The Review of Economic Studies, 74(3), 857-895.

McFadden, D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. Econometrica, 995-1026.

Ministry of Health. (2014). Malawi National Malaria Indicator Survey. Lilongwe, Malawi.

Mishra, K., & Sam, A. G. (2016). Does Women's Land Ownership Promote Their Empowerment? Empirical Evidence From Nepal. World Development, 78, 360-371.

Murray, C. J., Rosenfeld, L. C., Lim, S. S., Andrews, K. G., Foreman, K. J., Haring, D., . . . & Lopez, A. D. (2012). Global Malaria Mortality Between 1980 and 2010: A Systematic Analysis. The Lancet, 379(9814), 413-431.

Mzilahowa, T., Hastings, I. M., Molyneux, M. E., & McCall, P. J. (2012). Entomological Indices of Malaria Transmission in Chikhwawa District, Southern Malawi. Malaria Journal, 11(1), 380.

Peters, P. E. (1997). Against the Odds: Matriliny, Land and Gender in the Shire Highlands of Malawi. Critique of Anthropology, 17(2), 189-210.

Phiri, K. M. (1983). Some Changes in the Matrilineal Family System among the Chewa of Malawi. The Journal of African History 24(2):257â274.

Ranson, H., & Lissenden, N. (2016). Insecticide Resistance in African Anopheles Mosquitoes: A Worsening Situation That Needs Urgent Action to Maintain Malaria Control. Trends in Parasitology, 32(3), 187-196.

Reniers, G. (2003). Divorce and Remarriage in Rural Malawi. Demographic Research Special Collection 1.6.

Richards, E., Theobald, S., George, A., Kim, J. C., Rudert, C., Jehan, K., & Tolhurst, R. (2013). Going Beyond The Surface: Gendered Intra-Household Bargaining as a Social Determinant of Child Health and Nutrition in Low and Middle Income Countries. Social Science & Medicine, 95, 24-33.

Saha, B., & Sangwan, N. (2019). Credit Where Credit's Due: The Enabling Effects of Empowerment in Indian Microfinance. World Development, 122, 537-551.

Schaner, S. (2017). The Cost of Convenience? Transaction Costs, Bargaining Power, and Savings Account Use in Kenya. Journal of Human Resources, 52(4), 919-945.

Schatz, E. J. (2002). Numbers and Narratives: Making Sense of Gender and Context in Rural Malawi. Dissertations available from ProQuest. AAI3043950. https://repository.upenn.edu/dissertations/AAI3043950

Schatz, E. J. (2005). âTake Your Mat and Go!': Rural Malawian Women's Strategies in the HIV/AIDS Era. Culture, Health, and Sexuality 7(5):479â492.

Sehgal, V. M., Siddiqui, W. A., & Alpers, M. P. (1989). A Seroepidemiological Study to Evaluate the Role of Passive Maternal Immunity to Malaria in Infants. Transactions of the Royal Society of Tropical Medicine and Hygiene, 83, 105-106.

Smith, D. L., Dushoff, J., Snow, R. W., & Hay, S. I. (2005). The Entomological Inoculation Rate and Plasmodium Falciparum Infection in African Children. Nature, 438(7067), 492.

Sokullu, S., & Valente, C. (2018). Individual Consumption in Collective Households: Identification Using Panel Data with an Application to PROGRESA. Unpublished Manuscript, 3(5), 12.

Sraboni, E., & Quisumbing, A. (2018). Women's Empowerment in Agriculture and Dietary Quality Across the Life Course: Evidence from Bangladesh. Food Policy, 81, 21-36.

Steinhardt, L. C., Mathanga, D. P., Mwandama, D., Nsona, H., Moyo, D., Gumbo, A., ... & Troell, P. (2019). The Effect of Text Message Reminders to Health Workers on Quality of Care for Malaria, Pneumonia, and Diarrhea in Malawi: A Cluster-Randomized, Controlled Trial. The American Journal of Tropical Medicine and Hygiene, 100(2), 460-469.

Stopnitzky, Y. (2017). No Toilet, No Bride? Intrahousehold Bargaining in Male-Skewed Marriage Markets in India. Journal of Development Economics, 127, 269-282.

Takane, T. (2008). Customary Land Tenure, Inheritance Rules, and Smallholder Farmers in Malawi. Journal of Southern African Studies, 34(2), 269-291.

Thomas, D. (1990). Intra-Household Resource Allocation: An Inferential Approach. Journal of Human Resources, 635-664.

Tommasi, D., & Wolf, A. (2018). Estimating Household Resource Shares: A Shrinkage Approach. Economics Letters, 163, 75-78.

Tommasi, D. (2019). Control of Resources, Bargaining Power and the Demand of Food: Evidence from PROGRESA. Journal of Economic Behavior & Organization, 161, 265-286.

van den Berg, M., Ogutu, B., Sewankambo, N. K., Biller-Andorno, N., & Tanner, M. (2019). RTS,S Malaria Vaccine Pilot Studies: Addressing the Human Realities in Large-Scale Clinical Trials. Trials, 20(1), 316.

Voena, A. (2015). Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?. American Economic Review, 105(8), 2295-2332.

Walther, S. (2018). Noncooperative Decision Making in the Household: Evidence from Malawi. Journal of Development Economics.

Wickham, H. (2016). ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York.
https://openknowledge.worldbank.org/handle/10986/30902

World Health Organization. (2018) World Malaria Report 2018. Geneva. Licence: CC BY-NC-SA 3.0 IGO.

World Health Organization. (2019a) Rising to the Challenge of Malaria Eradication, https://www.who.int/news-room/detail/23-08-2019-rising-to-the-challenge-of-malaria-eradication. Accessed 01/13/2020.

World Health Organization Strategic Advisory Group on Malaria Eradication. (2019b) Malaria Eradication: Benefits, Future Scenarios and Feasibility. Executive Summary. Geneva: World Health Organization. WHO/CDS/GMP/2019.10 Licence: CC BY-NC-SA 3.0 IGO.

2.9 Appendix A: Robustness Checks for our Measure of Bargaining Power

In addition to studying the robustness of our results to various 2SLS modeling specifications, we assess whether our results are robust to the use of alternative bargaining power measurement methodologies. The two methods that we report here — using a proxy measure and estimating DLP's collective model of the family — are relegated to robustness checks simply because divorce is not endogenous in these models. The instrument makes divorce more or less valuable, which we can measure directly using the model presented in Section 2, but only indirectly using the models presented here. In this appendix we provide an overview of the derivations for these two measures, summary statistics for the measures recovered from the Malawi LSMS 2010 and 2013 panel data, and the robustness results for studying power and malaria contraction rates using these alternatives.

2.9.1 Constructing a Proxy for Bargaining Power Using Principal Component Analysis

Our first alternative approach to measuring power in the family is to create a proxy variable from information about partners. We can select variables to include in the index by following studies on which characteristics influence power dynamics (e.g. Agarwal, 1997; and Doss, 2013). In our measure, we combine five commonly used proxy variables: individuals' educations, earnings, land holdings, literacy, and age.

We construct the bargaining power proxy variable in four steps: (1) calculate (two) within-couple, normalized values of relative land holdings, income, education, (English) literacy, and age; (2.14) use PCA to calculate the weights for female household heads and calculate index values for them in each year in the sample; (3) repeat the second step for male household heads; (4) take the ratio of the female decision maker's values to the sum of couple's values.

We carry out the normalization in step (1) in the following way. Let $x_{f,h}$ and $x_{m,h}$ be the vector of descriptive variables for male and female household heads in a family h.

We calculate two family-specific values, one for the father and one for the mother. The values of each element of this vector of five variables, for $t \in \{f, m\}$, are:

$$\chi_{t,h} = \begin{cases} 1 & \text{if } x_{t,h} >= x_{-t,h} \\ \frac{x_{t,h}}{x_{-t,h}} & \text{if } else \end{cases}$$

Partners get values 1 if they have a higher variable value, and some value between zero and one otherwise. We set the value of the element of χ relating to literacy to one-half when partners have the same literacy. We then fit a PCA model for men and women separately, generating loadings for each gender.

Let $\hat{\beta}_f^{PCA}$ and $\hat{\beta}_m^{PCA}$ be the vector of loadings from the first PCA. We calculate a value for each female and male decision maker that most efficiently aggregates information about them $\hat{\beta}_f^{PCA}\chi_{f,h}$ and $\hat{\beta}_m^{PCA}\chi_{m,h}$. Then, we let our measure of bargaining power from this exercise be:

$$\hat{\eta}_h^{PCA} = \frac{\hat{\beta}_f^{PCA}\chi_{f,h}}{\hat{\beta}_m^{PCA}\chi_{m,h} + \hat{\beta}_f^{PCA}\chi_{f,h}} \tag{2.13}$$

A key advantage of this approach is that we do not need any assumptions on utility functions, or the functional forms of earnings models or demand functions. A key disadvantage is that we have made a very strong assumption about the functional form of bargaining power. To use this method, we must assume that the functional form in equation (2.13) is exactly correct. We must assume that the estimate is not missing any relevant information, and the ratio is the correct functional form. This reduced form approach to measuring power is relatively easy to construct and can inform us broadly about gender inequality in the family, but should be used with caution. In tandem with our two structural methods for recovering power, it is a useful tool for exploring bargaining power inequality and malaria incidence in Malawi.

2.9.2 Applying the PCA Index Method to the Malawi Panel Data

We displayed the summary statistics for the variables that we use to construct the PCA index for bargaining power in Tables 2.1 and 2.2. Men in this sample are older and have higher incomes, more educational attainment, and a higher likelihood of being literate. In addition, men manage twice as much land as women.[24]

The loadings we estimate for women are 0.37 for land, 0.40 for income, 0.59 for education, 0.16 for age, and 0.58 for literacy. For men, these values are 0.26, 0.22, 0.67, 0.01, and 0.65. For both men and women, the variance across families in the values of χ are largely driven by income and literacy. For men, these two variables explain more of the variation, relative to the values for women.

The average index value for men is 0.78 and for women is 0.56. The mean of the bargaining power variable is 0.41. We calculate this index for 98% of the families in our sample with a married male and female household head (4,487 families pooled across both 2010 and 2013). Using this method, we reject the null hypothesis that men and women have equal values of bargaining power, in favor of the alternative hypothesis that women have less bargaining power than men. The t-statistic on this test has value -63.78. The corresponding 95% confidence interval on the difference between men's and women's PCA index values is [-0.16, -0.17].

2.9.3 Dunbar, Lewbel, and Pendakur's (2013) Collective Model of the Family

Our third approach to measuring women's bargaining power in the family is to apply DLP's model to recover the sharing rule for each family. In DLP's (2013) extension, family members are assumed to have "price-independent generalizable log" (PIGLOG) utility functions and are assumed to fully commit to household decisions (divorce and noncooperative behaviors are not possible). The family production function is assumed to be linear so that some matrix A relates household-level purchases to individual consumption equivalents, $z = A[x_f + x_m + x_c]$. The matrix A is such that the Lindahl prices

[24] The test statistic on the difference between men's and women's land management is 4.93.

that family members face are weakly less than market prices (to account for the positive consumption externalities within households).

Families are assumed to consume some private assignable goods.[25] If a family consumes some positive amount of a private assignable good for each member, then that family meets the data requirements for the DLP method. Typical examples of private assignable goods are clothing and shoes. For private assignable goods, the need to address the problem of externalities when calculating resource shares is obviated: the matrix A simply has a 1 on the diagonal for private goods, and 0's elsewhere. As such, demand functions for private assignable goods are invariant to changes in the family-specific matrix, A. One of the primary contributions of DLP is to extend the BCL model so that the researcher does not need to estimate A.

DLP consider the family's Engel curves (demand functions that do not depend on prices) for private assignable goods. They are able to recover resource shares from a family's system of Engel curves given one of two additional assumptions. The first is that Engel curves have the same shape across people within in a family, at least at low levels of y. This is their "similarity across people" (SAP) assumption. The second is that individuals in families of the same size have similar preferences. This is their "similarity across types [of families]" assumption. We assume the former in our analysis and so do not provide the functional forms for Engel curves under the latter. Under these assumptions on utility functional forms, the Engel curves for private assignable goods in the household are

$$W_m(y) = \eta_m(\delta_m + \beta ln(\eta_m)) + \eta_m \beta ln(y), \tag{2.14}$$

$$W_f(y) = \eta_f(\delta_f + \beta ln(\eta_f)) + \eta_f \beta ln(y),$$

$$W_c(y) = \eta_c(\delta_c + \beta ln(\eta_c)) + \eta_c \beta ln(y),$$

where $W_t(y)$ is the family's budget share devoted to the private assignable good for a

[25]Recall that a private good is one that only one member can consume. An assignable good is one that the econometrician can ascribe with certainty to a certain individual.

person of type t. The SAP assumption lets a single β describe the slopes of the Engel curves. As such, fitting (2.14) using the system of seemingly unrelated regressions gives slope parameters that vary only by an individuals' resource share. The first component of the right-hand side of the Engel curve is set equal to the intercept, a nuisance parameter. The system of equations to fit to demand data, then, is:

$$W_m(y) = \beta_{0,m} + \beta_{1,m} ln(y) + \epsilon_m, \tag{2.15}$$

$$W_f(y) = \beta_{0,f} + \beta_{1,f} ln(y) + \epsilon_f$$

$$W_c(y) = \beta_{0,c} + \beta_{1,c} ln(y) + \epsilon_c$$

Fitting this model gives sample-level estimates of the population parameters for each Engel curve. The econometrician is left with a system of sample-level and family-level observed parameters, plus the resource share values, and one nuisance parameter. This system is:

$$W_m(y) = \hat{\beta}_{0,m} + \hat{\beta}_{1,m} ln(y) \text{ s.t. } \hat{\beta}_{1,m} = \hat{\eta}_m \beta, \tag{2.16}$$

$$W_f(y) = \hat{\beta}_{0,f} + \hat{\beta}_{1,f} ln(y) \text{ s.t. } \hat{\beta}_{1,f} = \hat{\eta}_f \beta,$$

$$W_c(y) = \hat{\beta}_{0,c} + \hat{\beta}_{1,c} ln(y), \text{ s.t. } \hat{\beta}_{1,c} = \hat{\eta}_c \beta,$$

$$\hat{\eta}_m + \hat{\eta}_f + \hat{\eta}_c = 1.$$

The sharing rule and some additional nuisance parameter, β, are identified in (2.16). For each family, there is an Engel curve per "type" of person, and the additional equation dictating that resource shares sum to one. For instance, in a family with no children, there are three total equations, and partners' resource shares can be recovered. For a family with five children, there will be four total equations, and the sharing rule can be recovered between the mother, the father, and the group of children collectively. Additional "types" (e.g., "daughters" and "sons") can be added to the model as long as the econometrician observes positive family spending on private assignable goods for them. That is, the survey must ask families about expenditures on "daughter-specific" and "son-specific" goods, in addition to "mother-specific" and "father-specific" goods.

We construct a measure of power from these resource shares: the ratio of women's resource shares to men's and women's summed resource shares. This measure captures the division of the gains from trade in a marriage:

$$\hat{\eta}^{RS} = \frac{\hat{\eta}_f}{\hat{\eta}_f + \hat{\eta}_m}.$$

There are reasons to be cautious when using this method. The first is that censoring in reported private assignable goods consumption threatens identification (Tommasi and Wolf, 2018). In general, most families do not meet the data requirements for the resource share method. In the 12-month recall data that DLP study, 75% of families do not meet the requirements to estimate the sharing rule. In Tommasi's (2019) application to Progresa data, almost 90% of families are censored. Likewise, in our Malawi data, about 90% of families are censored. This degree of censoring can lead to flat Engel curves, and weak identification of the sharing rule (Tommasi and Wolf, 2018).

In addition, families that meet the data requirements are likely to be systematically more wealthy than families who do not. Therefore, sharing rule estimates may not be externally valid for families that do not purchase a positive amount of PAGs for each type of member. This has several downsides. One is that studying program effects on the sharing rule is limited to studying the local effects to these systematically wealthier families. More directly, the sharing rule cannot be recovered or studied for the majority of the population. Third, as shown by Bargain, Lacroix, and Tiberti (2018), resource share estimates may not be robust to different choices of private assignable goods. While the DLP method can tell us a great deal about distribution of resources and power within the family, it is perhaps best interpreted as one measure of power among several. Estimating additional models of decision-making power, as we do, is thus valuable in assessing the robustness of our analysis of women's empowerment and malaria transmission.

2.9.4 Applying the DLP Method to the Malawi Panel Data

We pool observations across sample years to estimate resource shares, and use the IHS-prepared expenditure value as our right-hand side variable in (2.15). We assume that clothing and shoes are private assignable goods, and use the budget share spent on the sum of these goods as left-hand side variables for the system in (2.15).

We estimate this system of nonlinear equations using three-stage least squares, instrumenting for household consumption with a polynomial of the community's averaged labor market wages (following Attanasio and Lechene, 2010, 2014; Sokullu and Valente, 2018; and Tommasi, 2019).[26] Instrumenting for consumption allows us to plausibly identify the causal relationship between budget shares going to private assignable goods and household consumption, washing out measurement error. Following Calvi (2017), we include the food budget share in our system of equations. This increases stability in the regressions, since the errors in the regressions of private assignable goods are likely to be similar to each other but (relatively) dissimilar to the errors in the regression on the food budget share.

We scale the intercept and slope parameters by the PCA index power measure, so that our regressions include underlying preference parameters. As such, the intercept and slope of each Engel curve varies by the household's index value, which depends on ten underlying preference parameters.[27] That is, households with the same BP index value have the same slope and coefficient estimates in the system of equations we solve. Letting $\hat{\eta}_h^{PCA}$ represent a household's index value, the system we estimate is given by

$$W_{food}(y) = \beta_{0,food}\hat{\eta}_h^{PCA} + \beta_{1,food}\hat{\eta}_h^{PCA}ln(y) + \epsilon_{food} \tag{2.17}$$

$$W_f(y) = \beta_{0,f}\hat{\eta}_h^{PCA} + \beta_{1,f}\hat{\eta}_h^{PCA}ln(y) + \epsilon_f$$

[26] We use R's systemfit to estimate the 3SLS model (Henningsen and Hamann, 2007). Three-stage least squares is simply a way to fit a system of seemingly unrelated equations where the same endogenous variable is present in each equation, and an instrumental variables approach is used to identify the corresponding regression coefficients.

[27] Different authors use different preference parameters. For instance, Calvi (2017) uses a linear index of 25 preference parameters in her model of families with children, and an index with 22 parameters for childless couples. Tommasi (2019) uses an index with ten demographic controls plus time and location fixed effects. DLP estimate their model with 18 preference parameters incorporated into each regression in their system (under their assumption of similarity across types of people within the family).

$$W_m(y) = \beta_{0,m}\hat{\eta}_h^{PCA} + \beta_{1,m}\hat{\eta}_h^{PCA} ln(y) + \epsilon_m$$

$$W_g(y) = \beta_{0,g}\hat{\eta}_h^{PCA} + \beta_{1,g}\hat{\eta}_h^{PCA} ln(y) + \epsilon_g$$

$$W_b(y) = \beta_{0,b}\hat{\eta}_h^{PCA} + \beta_{1,b}\hat{\eta}_h^{PCA} ln(y) + \epsilon_b.$$

We report the regression results for the most common family type in Table 2.9, Appendix B. The slope coefficients are significantly different from zero for all of the private assignable goods, and so we can be sure that resource shares are identified in this data. In addition, we learn that the budget share spent on food is decreasing in income, as is generally observed (Attanasio and Lechene, 2014). We also observe that the budget shares spent on private assignable goods are decreasing in family income. By assumption, they bring utility to only one person. It may be a more efficient aggregate welfare maximization strategy to spend more money on public goods as family income increases.

Then, we plug in the sample-specific 3SLS estimates, and the family-specific expenditures, structure, and $\hat{\eta}_h^{PCA}$ into the system of equations given in (2.17). Denote the number of women in a family as F, of men as M, of boys as B, and of girls as G. We solve this system for each household using the generalized method of moments to minimize the sum of squared differences of the left-hand sides from zero:

$$W_f(y) - \hat{\beta}_{0,f}\hat{\eta}_h^{PCA} - \hat{\beta}_{1,f}\hat{\eta}_h^{PCA} ln(\frac{\eta_f y}{F}) = 0 \tag{2.18}$$

$$W_m(y) - \hat{\beta}_{0,m}\hat{\eta}_h^{PCA} - \hat{\beta}_{1,m}\hat{\eta}_h^{PCA} ln(\frac{\eta_m y}{M}) = 0$$

$$W_g(y) - \hat{\beta}_{0,g}\hat{\eta}_h^{PCA} - \hat{\beta}_{1,g}\hat{\eta}_h^{PCA} ln(\frac{\eta_g y}{G}) = 0$$

$$W_b(y) - \hat{\beta}_{0,b}\hat{\eta}_h^{PCA} - \hat{\beta}_{1,b}\hat{\eta}_h^{PCA} ln(\frac{\eta_b y}{B}) = 0$$

$$1 - \eta_f - \eta_m - \eta_g - \eta_b = 0$$

The set of parameter estimates $\{\hat{\eta}_f, \hat{\eta}_m, \hat{\eta}_g, \hat{\eta}_b\}$ that minimizes the sum of squared distances from zero is the set of resource share estimates for each household. System (2.18) is specifically for families with some positive number of men, women, boys, and girls. When a family has a different structure, we study the corresponding subset of the equations in system (2.18). For instance, in a household without boys, we would drop the fourth of five equations in (2.18), as well as the boys' resource share term in the fifth equation.

We can recover the sharing rule for a total of 821 families in our pooled sample (11% of all families in our sample, and 18% of families headed by a married couple). The median resource shares are 24.63% for women, 31.02% for men, 24.44% for girls, and 11.94% for boys. In the household that has the median resource share for women, the female household head has a resource share of 24.63%, the male household head has a resource share of 33.9%, the single boy has a resource share of 17.69%, and the two girls split a resource share of 23.7%. In this family, women consume less surplus than men, and girls consume less than boys.

Analyzing resource shares by family size reveals that women sacrifice more consumption as their family grows than men do. The ratio of the female decision maker's resource share to the male decision maker's resource share in the median family with no children is 1.08. In these 67 families, women have slightly higher point estimates of resource shares, but we fail to reject the null that these 67 men and women have equal resource shares (with p-value 0.51). As the number of children grows, this value drops substantially. With one and two children, these ratio values are 0.81 and 0.82, respectively. With three children, this value is 0.76. With four, it is 0.61. This same pattern is documented by Dunbar, Lewbel, and Pendakur (2013) using Malawi IHS data from 2004.

The vast majority of families in our sample have children. For these families, the average female resource share is 0.24. The average male resource share is 0.28. The t-statistic for the test that the distributions of men's and women's resource shares are the same is -4.78. We reject the null hypothesis that men and women in families with children have equal resource shares, in favor of the alternative hypothesis that women have small resource shares. The bargaining power measure we construct, $\hat{\eta}^{RS}$, has a median of 0.45. We reject the null hypothesis that men and women have an equal proxy value on average. The t-statistic for this test is -4.68.

We display the first and second stage results for these two models in Table 2.6, where the controls are family size, normalized family expenditures, number of rooms in the dwelling, dwelling materials fixed effects, a dummy for whether there is a health clinic in the community, and region-year fixed effects. We also include the percent of neighbors with malaria in the last two weeks in these robustness checks.[28] We find that the coefficient estimates for our instrument are similar across these two models, and similar to the results in Table 2.3. However, with the low number of observations for the DLP method, we are unable to reject the null hypothesis that this measure is different from zero. As such, we report correlations for the DLP measure, and causal relationships for the PCA proxy variable.

In Table 2.7, we report four regression results. The first two columns show naive OLS regression results with and without controls. The correlations reported in columns (1) and (2) are similar to those in Figure 2.7, where we report correlations using our preferred method. The third column reports the same regression result as in Table 2.3, column (3), but using the PCA measure instead of our preferred model's measure. The regression results there roughly confirm those reported in the main body of this paper. A standard deviation increase in this bargaining power measure would result in an 80% reduction in the likelihood that a family has at least one member with malaria. We also report a 2SLS model where the second stage is a Poisson model, and the left hand side variable is the number of family members with malaria in the last two weeks. These results (here and for our main model) are nearly identical to the second-stage-Logit results.

In Table 2.8, we report the results for regressing women's relative resource shares on a dummy variable equal to one when one or more family members had malaria in the last two weeks. The negative correlations in this table are similar to those in Figure 2.7, where we present the negative correlations between the preferred measure and malaria contraction likelihood.

[28]We do not test whether these robustness checks for the first identification strategy are themselves robust to dropping this "bad" control Including this variable, in any case, gives more conservative estimates of the marginal effect of an increase in equality in the family on malaria contraction rates.

Table 2.6: First-Stage Regression Results

Dependent Variable: Women's Bargaining Power		
	DLP	PCA
Instrument	0.014	0.022***
	[-0.035, 0.050]	[0.013, 0.033]
Controls	Yes	Yes
Region-Year Fixed Effects	Yes	Yes
Observations	660	2,824
R^2	0.765	0.936
Adjusted R^2	0.760	0.936
F Statistic	28.66***	863.86***
	(df = 15; 645)	(df = 15; 2809)

Note: 95% bootstrapped confidence intervals in brackets. Controls include family size, per capita monthly total household expenditures, percent of neighbors with malaria in the last two weeks, number of children under the age of 5 in the family, number of rooms in the dwelling, dwelling material fixed effects, and a dummy variable for whether there is a health clinic or hospital in the community. The first column reports the results using the power estimates from the Dunbar, Lewbel, and Pendakur (2013) model. The second column reports the results from the Klein and Barham (2019) model, and the third reports the results using the index constructed from principal components analysis. The reported F statistics are calculated for regressions with and without the instrument only, not with and without all controls. *p<0.1; **p<0.05; ***p<0.01.

Table 2.7: An Increase in Women's Bargaining Power Decreases the Malaria Burden for the Family (Results Using the Index Constructed Using PCA, $\eta^{PCA}_{f,h,t}$)

	Dependent Variable: Malaria Contraction			
	Not IV	Not IV	IV	IV
	logistic	*logistic*	*logistic*	*Poisson*
	(1)	(2)	(3)	(4)
Women's Bargaining Power	-1.962^{***}	-0.305		
	(0.097)	(0.385)		
Instrumented Power			-9.335^{**}	-9.297^{**}
			[-19.230, -2.760]	[-16.393, -3.796]
Controls	NO	YES	YES	YES
Region-Year Fixed Effects	NO	YES	YES	YES
Observations	2,793	2,793	2,790	2,790
Log Likelihood	−1,710.054	−1,635.279	−1,632.109	−2,142.635
Akaike Information Criterion	3,422.107	3,300.558	3,294.218	4,315.270

Note: 95% Bootstrapped Confidence Intervals in Brackets. Controls include family size, per capita monthly total household expenditures, percent of neighbors with malaria in the last two weeks, number of children under the age of 5 in the family, number of rooms in the dwelling, dwelling material fixed effects, and a dummy variable for whether there is a health clinic or hospital in the community. *p<0.1; **p<0.05; ***p<0.01.

Table 2.8: An Increase in Women's Bargaining Power is Negatively Correlated With the Malaria Burden for the Family (Dunbar, Lewbel, and Pendakur [2013] Model Results, $\eta^{RS}_{f,h,t}$)

	Dependent Variable: Malaria Contraction	
	Not IV	Not IV
	logistic	*logistic*
	(1)	(2)
Women's Bargaining Power	-1.549^{***}	-0.591^{*}
	(0.169)	(0.353)
Controls	NO	YES
Region-Year Fixed Effects	NO	YES
Observations	652	652
Log Likelihood	−403.526	−380.344
Akaike Information Criterion	809.052	790.687

Note: controls include family size, per capita monthly total household expenditures, percent of neighbors with malaria in the last two weeks, number of children under the age of 5 in the family, number of rooms in the dwelling, dwelling material fixed effects, and a dummy variable for whether there is a health clinic or hospital in the community. *p<0.1; **p<0.05; ***p<0.01.

Appendix B: Tables Reporting Nuisance Parameters

Table 2.9: Regression Results for Eligible Families with Mothers, Fathers, Sons, and Daughters, Used to Back Out Resource Shares Using the DLP Method

Dependent Variable	$\hat{\beta}_0$	$\hat{\beta}_1$
$W_{food}(y)$	6.068***	-0.371***
	(1.038)	(0.075)
$W_f(y)$	0.046***	-0.003***
	(0.015)	(0.001)
$W_m(y)$	0.032**	-0.002*
	(0.014)	(0.001)
$W_b(y)$	0.023***	-0.001**
	(0.009)	(0.001)
$W_g(y)$	0.043***	-0.003***
	(0.012)	(0.001)
N	397	
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 2.10: Selection Regression Results

	Earnings Dummy Variable	
	Women	Men
Intercept	−2.790*** (0.205)	−1.507*** (0.244)
Distance to Daily Market	0.005* (0.003)	0.003 (0.004)
Distance to Weekly Market	0.007 (0.004)	0.006 (0.005)
Number Kids Under 5	0.023 (0.021)	0.051* (0.028)
Head Dummy	0.735*** (0.058)	1.481*** (0.070)
Plot Size	0.387*** (0.031)	0.445*** (0.057)
Age	0.116*** (0.010)	0.119*** (0.011)
Age Squared	−0.001*** (0.0001)	−0.001*** (0.0001)
Rural Dummy	0.109 (0.074)	−0.126 (0.090)
Same-Community-Rural Interaction	0.271*** (0.042)	−0.059 (0.048)
Same Community	0.115 (0.108)	−0.280** (0.118)
Chichewa Literacy	−0.121** (0.053)	0.084 (0.077)
English Literacy	−0.087* (0.047)	−0.131** (0.059)
Highest Grade Completed	0.023*** (0.007)	−0.016** (0.007)
Number Men	−0.201*** (0.020)	−0.092*** (0.020)
Number Women	0.001 (0.023)	−0.089*** (0.023)
Number Girls	−0.061*** (0.018)	−0.065*** (0.021)
Number Boys	−0.014 (0.018)	−0.051** (0.021)
2013 Dummy	0.144*** (0.044)	0.154*** (0.053)
District Fixed Effects	Yes	Yes
Inverse Mills Ratio	0.924*** (0.294)	−0.307 (0.211)
Observations	7,029	7,458
Log Likelihood	−3,691.607	−2,522.757
χ^2 (df = 47)	2,348.833***	4,176.181***
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 2.11: Earnings Regression Results

	Earnings Continuous Variable	
	Log Earnings	
	Women	Men
Intercept	3.062*** (0.725)	6.879*** (0.369)
Head Dummy	0.957*** (0.171)	0.883*** (0.175)
Plot Size	0.361*** (0.052)	0.141*** (0.022)
Age	0.193*** (0.023)	0.112*** (0.013)
Age Squared	−0.002*** (0.0003)	−0.001*** (0.0002)
Rural Dummy	−0.419*** (0.106)	−0.267*** (0.071)
Same-Community-Rural Interaction	0.215*** (0.066)	−0.218*** (0.039)
Same Community	−0.031 (0.166)	−0.071 (0.133)
Chichewa Literacy	−0.289*** (0.068)	−0.060 (0.060)
English Literacy	−0.052 (0.067)	−0.080* (0.046)
Highest Grade Completed	0.155*** (0.009)	0.100*** (0.006)
Number Men	−0.203*** (0.042)	0.093*** (0.022)
Number Women	0.101*** (0.034)	0.063*** (0.024)
Number Girls	−0.061** (0.024)	0.012 (0.017)
Number Boys	−0.069*** (0.023)	−0.025 (0.017)
2013 Dummy	0.333*** (0.048)	0.221*** (0.034)
District Fixed Effects	Yes	Yes
Observations	3,661	5,154
R^2	0.982	0.989
Adjusted R^2	0.981	0.989
Residual Std. Error	1.371 (df = 3615)	1.176 (df = 5108)

Note: *p<0.1; **p<0.05; ***p<0.01

Appendix C: Mechanisms Analysis

We take an initial look at two possible explanations for the negative relationship between women's decision-making power and malaria transmission rates: bed net use and private latrine use. It is possible that more equal parents are better able to enforce bed net use for children. They may also invest more in this preventative technology, purchasing more comfortable nets or more nets per capita. It is also possible that women with more power are able to allocate collective resources toward improving household hygiene and sanitation facilities. Total health benefits may accrue disproportionately to women but may also result in reduced exposure to carrier mosquitoes for all family members. We run 2SLS models to study these mechanisms, with the same first stages as reported in Table 2.5, and with logistic regressions for the second stage. We report the regression results in Table 2.12. The assumptions required for identification are that the prevalence of matrilocal and matrilineal marital traditions are uncorrelated with household-level latrine sharing or bed net use, respectively.

We find that an increase in bargaining power increases the likelihood that a family has a private bathroom. Women with less power are far more likely to share a bathroom with members of another family. Sharing a bathroom may increase exposure to mosquitoes that have bitten neighbors. Increased privacy may be a good in itself and may reduce malaria transmission. Ayele, Zewotir, and Mwambi (2012), Bamaga et al (2014), and Dawaki et al (2016) report that the type of toilet that a family uses makes a large difference for malaria contraction rates. Having a toilet inside the house versus a latrine outside the household is a particularly large explanatory factor for malaria prevalence. In rural Malawi, almost all families have latrine outside of the dwelling (roughly 3% have flush toilets). For families with latrines, less than ten percent have ventilated and improved latrines, while the rest have more basic facilities. About 30% of families share a facility with other community members. Individuals who use private facilities likely come into contact with fewer infected mosquitoes than individuals who use communal facilities. As such, latrine type is an important factor explaining etymological innoculation rates.

Further, having a private restroom may be linked to other positive health outcomes.

For instance, people with private latrines are less likely to suffer from diarrhea (Heijnen et al, 2014). Stopnitzky (2017), and Alam and Monica (2018), among others, also document the link between increased privacy and gender equality in the family.

However, we find that gender equality is not linked to the likelihood that families use bed nets. Bed nets are very common in Malawi, largely because of the Malawian government's, and international organizations', efforts to disseminate preventative technologies in the last few decades. More than half of Malawi's families have at least one bed net for every two members (World Health Organization, 2018). Between 2009 and 2015, more than 20 million nets were disseminated in Malawi (Lindblade et al 2015). This dissemination effort likely did not target families based on their power dynamic. This would explain why this technology use is not a function of gender equality in the family in this context.

We suggest three more mechanisms that could help to explain our results, but leave their analysis for future research. The most obvious possible mechanism is that more equal parents invest more in other preventative measures like indoor residual spraying, or health behaviors like regularly treatments during pregnancy. The second is that men and women with more equal decision-making power may be more likely to engage in safe sex practices. This would make more equal parents less likely to have HIV/AIDS, which weakens people's immune systems and increases their risk of infection. Roughly one in ten people in Malawi in 2004 were HIV positive, and so this may be a mechanism explaining our findings (estimates from the Malawi Demographic and Helath Survey of 2004, reported in Greenwood et al, 2019). The third is that mothers who contract malaria during pregnancy may be more likely to have children who contract malaria, and that mothers' contraction is negatively correlated with gender equality.

Table 2.12: Mechanisms

	Dependent variable:	
	Toilets	Nets
Instrumented Women's Bargaining Power	13.620**	3.176
	[12.11, 82.97]	[-7.17, 22.51]
Normalized Family Expenditures	0.185	1.056***
	(0.178)	(0.216)
Family Size	0.175***	−0.084***
	(0.027)	(0.031)
Percent of Neighbors with Temp Malaria	−0.644**	1.116***
	(0.316)	(0.387)
Number of Kids 0-5 yo	−0.123**	0.376***
	(0.060)	(0.074)
Hospital Dummy	0.080	0.188
	(0.093)	(0.118)
Number of Rooms in Dwelling	0.464***	0.058
	(0.053)	(0.062)
Dwelling Material Fixed Effects	Yes	Yes
Region-Year Fixed Effects	Yes	Yes
Observations	2,609	2,789
Log Likelihood	−1,585.940	−1,163.899
Akaike Information Criterion	3,201.880	2,357.799

Note: 95% Bootstrapped Confidence Intervals in Brackets. Controls include family size, per capita monthly total household expenditures, percent of neighbors with malaria in the last two weeks, number of children under the age of 5 in the family, number of rooms in the dwelling, dwelling material fixed effects, and a dummy variable for whether there is a health clinic or hospital in the community. *p<0.1; **p<0.05; ***p<0.01.

Chapter 3

Estimates of Household Bargaining Power when Outside Options are Endogenous

Coauthored with Jean-Paul Chavas

Abstract: We generalize the collective model of the family to let partners choose their outside options, whether to act on them, and how to allocate shared resources. To do so, we apply Harsanyi's (1986) generic games framework to a limited-commitment collective model where a family's real consumption opportunities depend on each partners' specified outside option. We demonstrate uniqueness and existence of the solution by applying Rubinstein, Safra, and Thomson's (1992) appeals-immunity solution concept. We demonstrate that the expected value of bargaining power is semi-parametrically identified in an ordinal framework; the research need not assume specific utility or social welfare functional forms. In addition to empirical tractability and analytical precision, this update to the collective model provides a more nuanced view of intra-household bargaining power: partners with a greater capacity to specify more damaging threats will have more control over the decision-making process, regardless of whether they act on those threats. (JEL: D1, D6, D7)

Keywords: Sharing Rule, Outside Options, Endogenous Threats, Generic Game, Bargaining Power

3.1 Introduction

Economists have used the collective model of the family (Apps and Rees, 1988; Chiappori 1988, 1992) to empirically study power dynamics in the family for the last three decades (Donni and Molina, 2018). A key explanation for why certain power dynamics obtain in equilibrium is that partners' next best alternative to the collective allocation — their "outside options" — are more or less valuable (Mazzocco, 2007; Ligon, 2011). The partner with a more valuable outside option has more bargaining power in the family.

However, the collective model has two critical shortcomings that limit economists' ability to measure power dynamics. First, researchers must make strong assumptions about the nature of each partners' outside options. The standard approach is to assume the outside option is divorce, or some nebulous inefficient allocation within the family (Chiappori and Mazzocco, 2017). Assuming that each person's outside option takes the same form introduces misspecification error, and threatens identification. Letting each person endogenously choose their outside option in a generic game framework (Harsanyi, 1986) fixes this problem. Second, researchers must make strong functional form assumptions for social welfare, utility, production, and consumption functions, further introducing the possibility of misspecification error. Proceeding from an ordinal framework instead of a cardinal framework solves this problem.

In this paper, we address these two shortcomings of the collective model. We generalize the collective model with limited-commitment to a generic games framework, and we demonstrate semi-parametric identification of the sharing rule in an ordinal setting. To show that the solution is unique, we use the appeals-immune solution concept developed by Rubinstein, Safra, and Thomson (1992) and advanced by Hanany and Safra (2000). Then we suggest an estimation strategies that allows the econometrician to recover the expected value of intra-household bargaining power, conditional on each partners' threats and whether they act on them, that does not require economists to specify functional forms for utility, social welfare, production, and consumption functions.

To consider how these advancements contribute to economists' ability to measure bargaining power in the family, consider two recent, innovative papers on domestic abuse and power in the family — Ramos (2016), who studies data from northern Ecuador in 2011, and Lewbel and Pendakur (2019) who study data from Bangladesh in 2015. In both papers, men are allowed to choose some level of violence in the family (though only Ramos [2016] recovers a demand function for violence). The benefit of violence is that men gain control over the decision-making process — their "resource share" increases. The cost is that women are harmed, which is a negative outcome in itself, of course, and can also reduce productive possibilities for the household. By choosing to be violent with some positive frequency, men consume a larger portion of a smaller amount of surplus, and more overall than if they had not been violent. These papers advance our ability to measure power dynamics in inefficient families.[1]

However, both papers are subject to the limitations of the collective model already discussed. Ramos (2016) assumes that only men can specify and act on threats. In fact, women may be able to take some recourse to violence, like moving in with their parents temporarily, allocating labor or assets in a socially-inefficient way, or filing for a divorce. By assuming that women have no outside option, Ramos (2016) may overestimate the benefits that accrue to men from violence. Further, her work may be subject to misspecification error in the functional forms for utility and production functions, and the elasticity of women's labor allocations to violence.

Lewbel and Pendalur (2019) incorporate violence into the collective model developed by Browning, Chiappori, and Lewbel (2013) and use the identification strategy developed in Dunbar, Lewbel, and Pendakur (2013). Violence shifts the family from an efficient consumption technology to an inferior one (imposes higher shadow prices for the family). However, they are unable to derive a demand function for violence, which limits the predictive capacity of their model. Furthermore, they only demonstrate identification

[1]Similar arguments about the individual rationality of socially-inefficient behavior could be applied to explain a broad range of observed behaviors. Udry (1996) documents intra-household inefficiencies in the allocation of fertilizer across men's and women's agricultural plots using data from 1981-1985 in Burkina Faso. Reallocating some fertilizer from men's plots to women's would increase household income by 6% on average. de Mel, McKenzie, and Woodruff (2009) find evidence of inefficient demand for productive assets more broadly in experimental data from Sri Lanka, and Schaner (2015) documents that families in Kenya engage in inefficient savings behavior when partners discount future consumption at different rates. Walther (2018) documents partners choosing inefficient labor allocations in Malawi in order to gain more control over the decision-making process. See Basu (2006) for a theoretical treatment of the topic.

when threats take a binary value (either the husband abused his wife or he didn't). They are unable to study how more damaging threats might impact power dynamics, or explain why certain levels of violence obtain in equilibrium. They are unable to assess what effect the possibility of women filing for divorce might have on power dynamics.

The generalization of the collective model we develop in this paper solves these problems. In our generic game, we let both individuals in a couple choose their outside option (or "threat" in Harsanyi's language) from a large set that includes divorce, physical and emotional violence, damage to productive assets or the consumption technology, or some combination of these. We relate these threats to each other by quantifying their impact on the real shadow income that each person has at their disposal in the collective setting (or individually if they divorce), which is similar to the approach in Lewbel and Pendakir (2019). If implemented, more damaging threats reduce the total surplus available to the family more (move the Pareto frontier inward). Regardless of implementation, more damaging threats grant the partner who made them more control over the decision making process (moving the collective outcome along the Pareto frontier).

Our model provides analytical clarity on the role of threats in determining the quantity and division of surplus within the family. Power, it becomes clear, is an individual's capacity to specify a more damaging threat in order to gain additional control over the family decision-making process. The family's sharing rule is determined by power dynamics within the family. Also, unlike the models in the two example papers, power in our framework is influenced by the threats that partners specify, regardless of whether they act on them. Families where men might be violent, but are not, will have different power dynamics from families where men do not have an inclination towards violence. The two models we've discussed, and other collective models, are unable to make this distinction. So Harsanyi's generic games framework adds nuance to the collective model's capacity to analyze power dynamics within families, as well as adding empirical tractability.

We also introduce a semi-parametric identification strategy, which enables researchers to bring this model to data. This strategy is ordinal, it does not require cardinal assumptions on utility functions or the social welfare function. Instead, researchers assume that preferences are jointly distributed across families in the population in some way. This distribution assumption implies a functional form for the expected value of the sharing rule in the family, which can be recovered from panel data with information on threats and violence, demographic characteristics, partners' earnings profiles, and at least two waves of information. Researchers can instrument for violence in order to correct for measurement error.

3.2 A Model of Household Decision Making with Limited Commitment and Endogenous Threat Points

Consider two married partners, indexed f and m, who bargain over consumption and production decisions, who may choose to be non-cooperative to some degree, and who can end the partnership via divorce. The choices that the individuals make will depend on their preferences, the bargaining power each partner has, prices, and income. The choices they make can shape the lives of their dependants. We build a model with these features and derive the households' demand functions for consumption goods, derive individuals' demand functions for violence and other forms of non-cooperation, predict whether divorce will occur, and show how the distribution of surplus within the family is related to the threat and use of violence, and how well-off partners are in the case of divorce.

The partners engage in a "generic game" where they first jointly determine the rules of engagement and then play according to these rules (See Harsanyi, 1986). In order to demonstrate that a unique solution exists to this application of generic games to the household, we use the "appeals-immunity" solution concept introduced by Rubinstein, Safra, and Thomson (1992). As such it is useful to define a generic game in notation that is similar to theirs. Let a *bargaining problem* by a triple $< \theta, \succsim_f, \succsim_m > \in (\Sigma_f \times \Sigma_m) \times \rho^2$, where θ is a bargaining outcome defined over the set of possible outcomes $\Sigma_f \times \Sigma_m$; and $\succsim_f$ and $\succsim_m$ are partners' preferences over states of the world induced by their bargaining outcomes, which are elements of the set of permissible preferences, ρ.

The permissible actions are threats of non-cooperation that each partner specifies but does not necessarily need to act on. They include physical or emotional violence, divorce, and milder punishments like neglecting agreed-upon duties or chores.

In this application of Harsanyi's generic games, the rules that partners decide on are their outside options, and the nested game is a limited-commitment collective model of the family. As such, the model has three sequential stages, and can be solved by backwards induction. In the first stage, both partners chose their threats, θ_f and θ_m, from their space of permissible threats, Σ_f and Σ_m. Specifying a more damaging threat moves the eventual collective allocation in a players favor.

In the second stage, partners consider the pair of threats and chose whether or not to carry their threat out. Denote this binary choice with an indicator variable equal to one when a partner chooses to implement their threat: $I_f, I_m \in \{0, 1\}$. Acting on a more damaging threat reduces total family surplus more by reducing the quality of the production function, the consumption technology, or both. This establishes the nested game that the partners play, by determining what technologies are available to them. The bargaining outcome can be written $\theta = \{\theta_f, \theta_m, I_f, I_m\}$.

In the third stage, the partners solve a collective model of the family subject to the consumption and production technologies available to them. This collective model can be written using a two stage formulation that optimally splits the family resources according to a sharing rule, and then lets players choose allocations in a disaggregated way. This third stage has two sub-stages. In total, there are three choices per player, and one household-level choice: the choice of (1) a threat, (2) whether to implement the threat, (3) how to allocate a scare budget and time, and (4) how to divide household resources between the two decision-makers.

The nested bargaining game solved in the third stage, then, is a bijection mapping from the partners' choices in the first two stages to the optimal division of resources between the two partners, $\eta^* \in [0, 1]$; the n-vector of household demand functions, $z^*(p, y, \eta^*) \in \mathbb{R}^n$ where p and y are prices and family income; the set of partners' n-vectors of consumption amounts, $\{x_f^*(p, \eta^* y), x_m^*(p, (1-\eta^*)y\} \in \mathbb{R}^{2n}$; and labor supply functions for each partner,

$l_f^*, l_m^* \in [0,24]^2$. This bijection is $B(\theta)$: $\Sigma_f \times \Sigma_m \times \{0,1\}^2 \rightarrow [0,1] \times \mathbb{R}_+^{3n} \times [0,24]^2$. That is, for any given set of solutions to the first two stages, $\{\theta_f, \theta_m, I_f, I_m\}$, the unique corresponding nested game $B(\theta_f, \theta_m, I_f, I_m)$ is induced which, in turn, has the unique solution $\{\eta^*, z^*(p, y, \eta^*), \{x_f^*(p, \eta^* y), x_m^*(p, (1-\eta^*)y\}, l_f^*, l_m^* \Big| \theta\}$.

In this section, we introduce each step of the problem in reverse order, starting with the third stage where partners take the threats and implementation choices as given. A unique solution to the problem in this stage is guaranteed to exist under restrictions on preferences and the solution space by the Kuhn-Tucker theorem. Then we work backwards to the choice of whether to implement, which induces a distinct bargaining game $B(\theta)$. Working backwards again, we come to the first stage where partners chose their outside options, which will be the unique appeals-immune and individually-rational threats.

While each partners' choice to implement their threat is individually rational, implementation is not socially optimal. We discuss how such a decision can be viewed as a "bargaining failure" by re-examining the solution to the nested game when players cannot act on their threats. This enforces a socially optimal solution (an allocation that lays on the outer most Pareto frontier) and results in a different sharing rule. This solution might obtain in the presence of a mediator, for example.

The Nested Two-Stage Collective Model

Consider the set of preferences with quasi-convex, twice differentiable, and increasing utility function representations $\succsim_i^{sm} \in \rho^{sm} \subset \rho$. Restrict attention to the class of bargaining problems of this nature: $< \theta, \succsim_f^{sm}, \succsim_m^{sm} >$. The two partners have some cardinal utility function representations of their ordinal preferences, defined over private-equivalent consumption bundles and time allocated to labor, $U_f(x_f, l_f)$ and $U_m(x_m, l_m)$. We discuss the solution to the nested game in this cardinal framework, then in a more general, ordinal framework. We achieve this by formalizing the link between preferences over consumption and labor allocations, and the chosen threats. We use this link to define ordinal preferences over threats themselves, not the resulting consumption bundles.

The central assumption in the nested model is that the partners reach a (conditionally) Pareto efficient outcome (as in Lewbel and Pendakur, 2019). The specified threats and whether they are carried out determines how far off of the best possible Pareto frontier a family's outcome will be. If neither partner acts on their threat, then the family is on a frontier that is strictly above the frontier that obtains when one or both partners carry out their threat. This is because acting on threats reduces the quality of the family's production and consumption technologies. Let $F(l_f, l_m|\theta)$ be the family's production function so that the total family income is $y = F(l_f, l_m|\theta)$, and let $A(x_f, x_m|\theta)$ be the family's consumption technology so that household demand functions are given by $z = A(x_f, x_m|\theta)$. This consumption technology relates the household purchases to the amount of each good that family members consume, accounting for externalities that accrue when the family consumes locally public goods, and for economies of scope within the household.[2] Implementing threats results in a lower Pareto frontier for the family by reducing the quality of F and A. Families are conditionally efficient in the sense that they consume along the frontier that their family faces, and the frontier is endogenous in that it results from the partners' choices of threats and implementation.

Because partners reach a Pareto efficient allocation (decisions over labor and consumption), the household's problem can be written using Chiappori's (1988) decentralized format. The family behaves as though it optimized some social welfare function, $\tilde{U}(U_f(x_f, l_f), U_m(x_m, l_m))$, which is increasing, differentiable, and concave in both of its arguments. By the second welfare theorem, the family's problem can be written in two stages: first, the family splits total resources between the two partners, then the two partners chose their private consumption demand and labor supply. Since partners make decisions over the consumption of public goods (i.e. goods that have externalities for other household members), they face Lindahl prices instead of market prices (Browning, Chiappori, and Lewbel, 2013). Let the Lindahl prices induced by the family-specific and threat-dependant consumption technology be $L(A(\theta)) \in \mathbb{R}^n_+$, such that $L(A(\theta)) \leq p$ when there are positive externalities. Note that the prices do not depend on the allocation

[2]See Browning, Chiappori, and Lewbel (2013) for an extended description of the characteristics of such a technology.

choices, but are simply functions of the quality of the consumption technology and the goods that consumers may purchase. Then given θ, the family solves $B(\theta)$:

$$\max_{\eta\in[0,1]} \quad \tilde{U}(V_f(p,\eta y|\theta), V_m(p,(1-\eta)y|\theta)) \text{ subject to} \tag{3.1}$$

$$A(x_f, x_m|\theta) = z \text{ (Consumption Technology Constraint)}$$

$$y = F(l_f, l_m|\theta) \text{ (Production Constraint)}$$

$$p'z = y \text{ (Household Budget Constraint)}$$

$$V_f(p,\eta y|\theta) = \max_{\substack{x_f\in\mathbb{R}^n_+,\\ l_f\in[0,24]}} U_f(x_f, l_f) \text{ s.t. } L(A(\theta))'x_f = \eta y \text{ (Her Choices)}$$

$$V_m(p,(1-\eta)y|\theta) = \max_{\substack{x_m\in\mathbb{R}^n_+,\\ l_m\in[0,24]}} U_m(x_m, l_m) \text{ s.t. } L(A(\theta))'x_m = (1-\eta)y \text{ (His Choices)}$$

The solution to $\mathbb{B}(\theta)$ is the set of household demand functions $z^*(p,y)$, the individuals' private consumption equivalents $x_f^*(L(A),\eta y)$ and $x_m^*(L(A(\theta)),(1-\eta)y)$, labor supply decisions l_f^* and l_m^*, and the sharing rule $\eta^*(\theta)$. A unique solution is guaranteed to exist by the Lagrange multiplier theorem. The first order condition of (1) with respect to η gives the optimality condition for the division of resources between partners, suppressing notation for legibility:

$$\frac{\partial\tilde{U}}{\partial V_f}\frac{\partial V_f}{\partial\eta} = -\frac{\partial\tilde{U}}{\partial V_m}\frac{\partial V_m}{\partial\eta} \tag{3.2}$$

This condition equates the marginal gain for the family from allocating the marginal unit of the family's resources to partner f's control, to the marginal loss to the family that occurs because partner m is no longer in control of that portion of the family's resource.

Limited Commitment

If partners do not act on their threat, the family solves bargaining problem $\mathbb{B}(\theta_f, \theta_m, I_f =$

$0, I_m = 0$), and each partner gets corresponding indirect utility from the efficient production and consumption technologies, and the resulting sharing rule, $V_f(p, \eta^* y|\theta_f, \theta_m, I_f = 0, I_m = 0)$ and $V_m(p, (1-\eta^*)y|\theta_f, \theta_m, I_f = 0, I_M = 0)$. As such, each partner prefers to implement their threat ($I_f = 1 \succ_f I_f = 0$ and $I_m = 1 \succ_m I_m = 0$) if the corresponding utility function representations of their preferences yield higher values when they implement their threats. They solve the following optimization problems:

$$I_f^* = \underset{I_f \in \{0,1\}}{argmax} \big[V_f(p, \eta^* y|\theta_f, \theta_m, I_f = 1, I_m), V_f(p, \eta^* y|\theta_f, \theta_m, I_f = 0, I_m)\big] \tag{3.3}$$

$$I_m^* = \underset{I_m \in \{0,1\}}{argmax} \big[V_m(p, (1-\eta^*)y|\theta_f, \theta_m, I_f, I_M = 1), V_m(p, (1-\eta^*)y|\theta_f, \theta_m, I_f, I_M = 0)\big]$$

The *collective allocation problem with limited commitment* is $\mathbb{B}(\theta_f, \theta_m, I_f^*, I_m^*)$. The solution to this problem is, conditional on the specified threats, the indicators for whether partners act on their threats, the optimal demand and supply functions, and the optimal division of resources, conditional on the quality of the production function and consumption technology. It would be optimal to reduce the total surplus that the family has access to by acting on the specified threat if the resulting change in the sharing rule increases an individuals' total consumption. That is, the increase in their "slice of the pie" more than offsets the decrease in the "total size of the pie." To be explicit, the limited-commitment collective problem that the family solves is

$$\underset{\eta \in [0,1]}{max} \quad \tilde{U}(V_f(p, \eta y|\theta_f, \theta_m, I_f^*, I_m^*), V_m(p, (1-\eta)y|\theta_f, \theta_m, I_f^*, I_m^*)) \text{ subject to} \tag{3.4}$$

$$A(x_f, x_m|\theta_f, \theta_m, I_f^*, I_m^*) = z \text{ (Consumption Technology Constraint)}$$

$$y = F(l_f, l_m|\theta_f, \theta_m, I_f^*, I_m^*) \text{ (Production Constraint)}$$

$$p'z = y \text{ (Household Budget Constraint)}$$

$$V_f(p, \eta y|\theta_f, \theta_m, I_f^*, I_m^*) = \underset{\substack{x_f \in \mathbb{R}^n_+, \\ l_f \in [0,24]}}{max} U_f(x_f, l_f) \text{ s.t. } L(A(\theta_f, \theta_m, I_f^*, I_m^*))' x_f = \eta y$$

$$V_m(p,(1-\eta)y|\theta_f,\theta_m,I_f^*,I_m^*) = \max_{\substack{x_m\in\mathbb{R}^n_+,\\ l_m\in[0,24]}} U_m(x_m,l_m) \text{ s.t. } L(A(\theta_f,\theta_m,I_f^*,I_m^*))'x_m = (1-\eta)y$$

$$I_f^* = \underset{I_f\in\{0,1\}}{argmax}\big[V_f(p,\eta^*y|\theta_f,\theta_m,I_f=1,I_m),V_f(p,\eta^*y|\theta_f,\theta_m,I_f=0,I_m)\big]$$

$$I_m^* = \underset{I_m\in\{0,1\}}{argmax}\big[V_m(p,(1-\eta^*)y|\theta_f,\theta_m,I_f,I_M=1),V_m(p,(1-\eta^*)y|\theta_f,\theta_m,I_f,I_M=0)\big]$$

When deciding what threats to choose, partners consider what corresponding nested game, $B(\theta_f,\theta_m,I_f^*,I_m^*)$, will result. They face trade offs: a more damaging threat grants more control over the decision making process but, if implemented, reduces total family surplus to a larger extent.

The First Stage of the Generic Game: Choosing Threats

While Harsanyi is relatively vague about how these threats damage players, we can specify how these damages occur more clearly in the family context.

Definition 1: Let θ_i, $\bar{\theta}_i \in (\Sigma_f\times\Sigma_m)\times\{0,1\}^2$ be two potential threats for partner $i\in\{f,m\}$. The threat θ_i is more damaging *than $\bar{\theta}_i$ if one or both of the following is true:*

1. *for $i\neq j\in\{f,m\}$, $F(l_f,l_m|\theta_i,\theta_j,I_i=1,I_j) < F(l_f,l_m|\bar{\theta}_i,\theta_j,I_i=1,I_j)$ (The* more damaging *threat, if carried out, lowers family income.)*

2. *for $i\neq j\in\{f,m\}$, $L(A(\theta_i,\theta_j,I_i=1,I_j)) > L(A(\bar{\theta}_i,\theta_j,I_i=1,I_j))$ (The* more damaging *threat, if carried out, increases the shadow prices each person faces in their optimization problems.)*

Whenever someone specifies a more damaging threat, it must be because they gain additional control over the decision making process by doing so. So for any θ_f, $\bar{\theta}_f \in \Sigma_f\times\{0,1\}$, and for any $\theta_m\in\Sigma_m\times[0,1]$, where θ_f is *more damaging* than $\bar{\theta}_f$, $\eta(\theta_f,\theta_m,I_f,I_m) > \eta(\bar{\theta}_f,\theta_m,I_f,I_m)$. The symmetric statement is true for partner m. If more damaging threats did not confer more control over the decision making process, the optimal threat would be the least damaging one and it would never be implemented. See

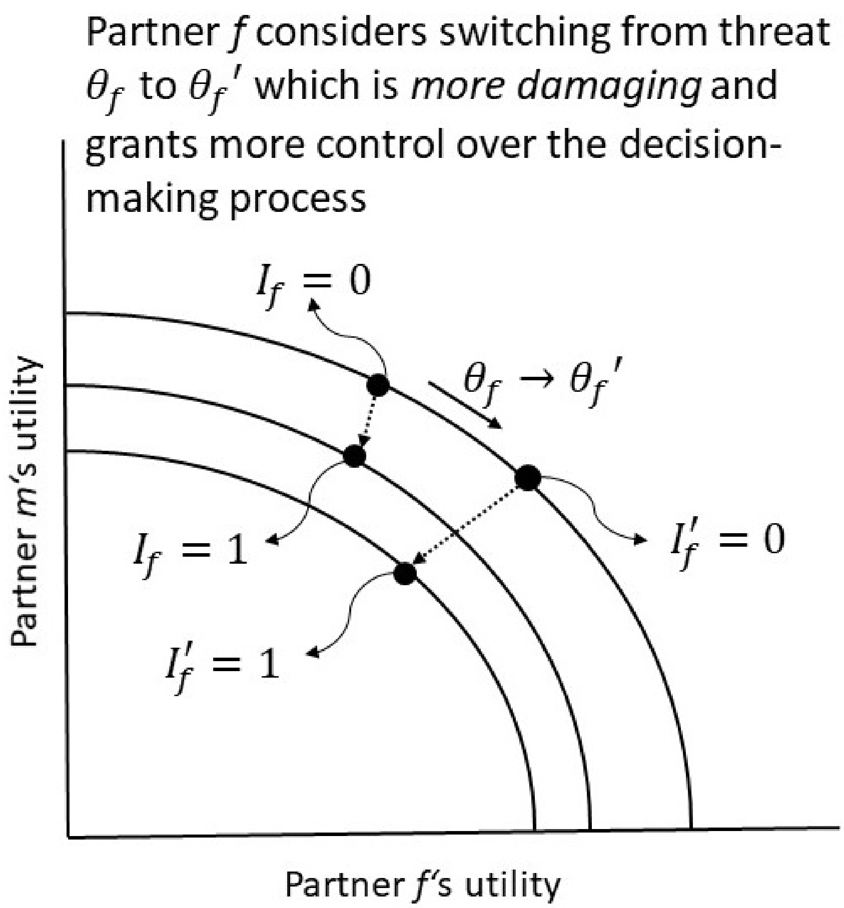

Figure 3.1: Partner f considers two different threats. The *more damaging* threat grants more bargaining power to partner f, shifting them along the fully efficient Pareto frontier if not implemented. If implemented, the *more damaging* threat moves the family to a lower Pareto frontier than the less damaging threat would, if it were implemented.

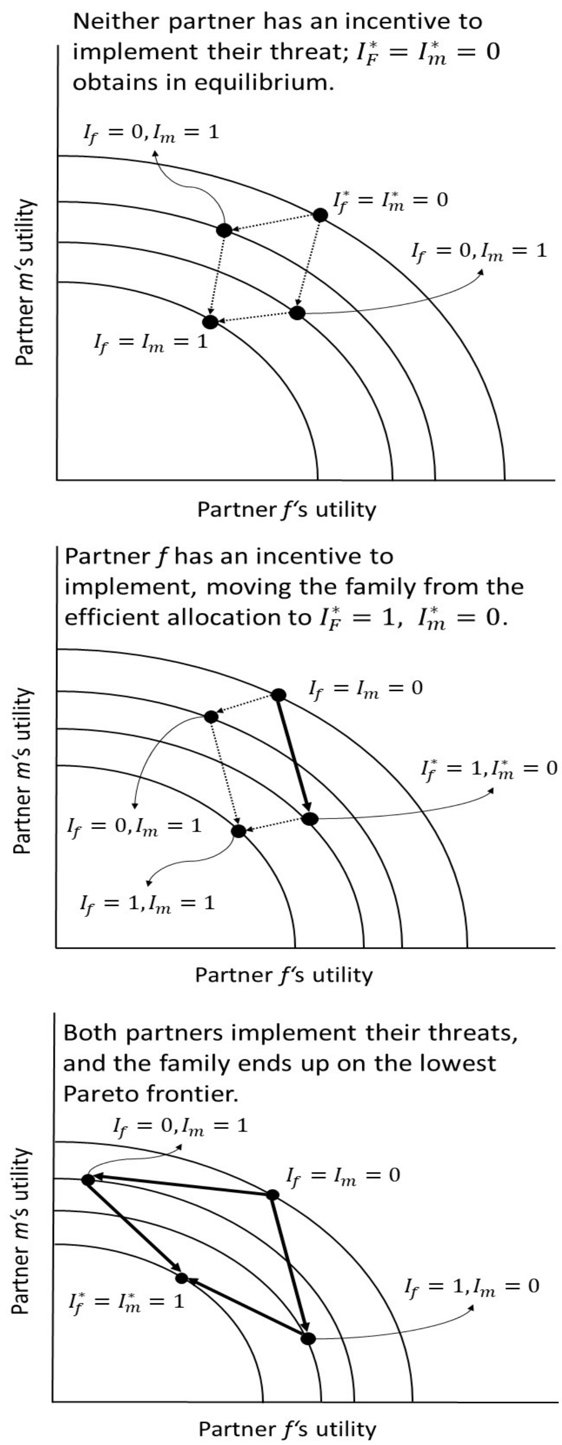

Figure 3.2: In the top panel, the threats that the partners have specified results in the fully-efficient equilibrium. In the middle panel, one of the partners has an incentive to implement their threat, and the family ends up on an intermediate Pareto frontier. In the bottom panel, both partners have an incentive to implement their threat and end up on the lowest Pareto frontier for their family.

Figures 3.1 and 3.2 for graphical depictions of the choices in stages one (the threats) and two (implementation).

Partners choose their threats from the space of possible threats to maximize their payouts in the resulting *collective allocation problem with limited commitment*. We can set up this constrained maximization problem using the cardinal utility functions, and subsequently on ordinal preferences using bijections between threats and allocations. We first describe the cardinal problem, then formally define these bijections, then set up the corresponding ordinal problems, and finish with the definition of *individual-rationality*. The partners solve:

$$\theta_f^* = \underset{\theta_f \in \Sigma_f}{argmax} \quad V_f(L(A(\theta_f, \theta_m, I_f^*, I_m^*), \eta^*(\theta_f, \theta_m) F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)) \tag{3.5}$$

$$\theta_m^* = \underset{\theta_m \in \Sigma_m}{argmax} \quad V_m(L(A(\theta_f, \theta_m, I_f^*, I_m^*), (1 - \eta^*(\theta_f, \theta_m)) F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*))$$

The choice of a threat influences utility in three ways. A more punishing threat moves the resource share in a person's favor, and more punishing threats damage the consumption and production technologies more. These optimization problems have the following first order conditions (suppressing notation):

$$\begin{aligned} &- \frac{\partial V_f}{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_f} = \\ &\qquad \frac{\partial V_f}{\partial \eta^*(\theta_f, \theta_m)} \frac{\partial \eta^*(\theta_f, \theta_m)}{\partial \theta_f} \times \\ &\qquad \frac{\partial V_f}{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_f} \end{aligned} \tag{3.6}$$

$$- \frac{\partial V_m}{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_m} = \frac{\partial V_m}{\partial \eta^*(\theta_f, \theta_m)} \frac{\partial \eta^*(\theta_f, \theta_m)}{\partial \theta_m} \times \frac{\partial V_m}{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_m} \quad (3.7)$$

These first order condition in (3.6) equate the marginal benefit that accrues to partner f from having a larger resource share to the marginal disutility that accrues because the Pareto frontier for the family shifts inwards when the threats are acted on. Equation (3.7) gives the corresponding relationship for partner m.

There may be multiple pairs (θ_f^*, θ_m^*) such that (3.6) and (3.7) hold simultaneously. The threats in any such pair are *individually−rational* in $< \theta, \succsim_f^{sm}, \succsim_m^{sm} >$. Note that any pair (θ_f^*, θ_m^*) that satisfies (3.6) and (3.7) implies that the couple solves a distinct *collective allocation problem with limited commitment* $\mathbb{B}(\theta_f^*, \theta_m^*, I_f^*, I_m^*)$. Each nested game has a unique solution, including specific private consumption equivalents and labor allocations for each partner. As such, each pair is associated with a specific level of indirect utility for each partner. The differential operators relating the pair (θ_f^*, θ_m^*) to consumption and labor outcomes are:

$$\Delta_f(\theta_f^*, \theta_m^*, \succsim_f, \succsim_m, \tilde{U}) =$$

$$\left\{ (x_f^*(L(A), \eta^* y), l_f^*((L(A), \eta^* y)) \in \mathbb{R}_+^n \times [0, 24] \middle| \text{conditions (3.2) and (3.6) hold} \right\}$$

$$\Delta_m(\theta_f^*, \theta_m^*, \succsim_f, \succsim_m, \tilde{U}) =$$

$$\left\{ (x_m^*(L(A), (1-\eta^*) y), l_m^*(L(A), (1-\eta^*) y)) \in \mathbb{R}_+^n \times [0, 24] \middle| \text{conditions (3.2) and (3.7) hold} \right\}$$

As such, we can define preferences directly over threats for each player. By defining preferences over threats, we can define the concept of *individual-rationality* of threats without referencing the cardinal utility function representation of underlying preferences, and without needing to specify a functional form for the family's social welfare function.[3]

Definition 2: Let $<\theta, \succsim_f^{sm}, \succsim_m^{sm}> \in (\Sigma_f \times \Sigma_m) \times (\rho^{sm})^2$ *be a* bargaining problem. *Let* $\theta^* \in (\Sigma_f \times \Sigma_m) \times [0,1]^2$ *be a potential* solution. *The potential* solution θ *is* individually rational *if both of the following are true:*

1. $\forall \theta_m \in \Sigma_m$, and $\theta_f^* \in \Sigma_f$, $\forall \bar{\theta}_f \in \Sigma_f \setminus \{\theta_f^*\}$, $\theta_f^* \succsim_f \bar{\theta}_f \iff \Delta_f(\theta_f^*, \theta_m, \succsim_f, \succsim_m, \tilde{U}) \succsim_f \Delta_f(\bar{\theta}_f, \theta_m, \succsim_f, \succsim_m, \tilde{U})$ and

2. $\forall \theta_f \in \Sigma_f$, and $\theta_m^* \in \Sigma_m$, $\forall \bar{\theta}_m \in \Sigma_m \setminus \{\theta_m^*\}$, $\theta_m^* \succsim_m \bar{\theta}_m \iff \Delta_m(\theta_f, \theta_m^*, \succsim_f, \succsim_m, \tilde{U}) \succsim_m \Delta_m(\theta_f, \bar{\theta}_m, \succsim_f, \succsim_m, \tilde{U})$.

This definition states that people prefer a threat if they prefer the corresponding outcome from the specific resulting nested game. Let $(\Sigma_f^* \times \Sigma_m^*) \subset (\Sigma_f \times \Sigma_m)$ be the set of all potential *solutions* that are *individually-rational* for the *bargaining game* $< \theta, \succsim_f^{sm}, \succsim_m^{sm}> \in (\Sigma_f \times \Sigma_m) \times (\rho^{sm})^2$.

3.2.1 Formally Defining Appeals Immunity

Consider induced utility functions, which map from the *solution* space to the real numbers, and describe a person's willingness to pay to transition from a reference *solution* to an alternative. Let $u(\succsim_i, \theta^r, \theta)$ be such a function for $i \in \{f, m\}$, where θ^r is a reference solution and θ is a specified alternative, $u : \rho \times \left((\Sigma_f^* \times \Sigma_m^*)^* \times \{0,1\}^2\right)^2 \to \mathbb{R}$. Because both outcomes, θ^r and θ, are associated with specific private equivalent consumption bundles for each person, and because people have preferences over these consumption bundles, partners indirectly have preference relations over θ^r and θ. People always prefer the outcome where their partner does not carry out the threat to the outcome where their partner chooses to carry out their threat. Note that $u(\succsim_i, \theta^r, \theta)$ is differentiable.[4]

[3]In general, children's utility would enter into the social welfare function (and potentially also into parents' preferences directly), extending the model to do so does not change our core results or suggested estimation strategy. It may expand the set of threats that people have at their disposal to include threats over custody.

[4]Formally, for some $\alpha \in [0,1]$, $u(\succsim_i, \alpha\theta^r + (1-\alpha)\theta, \theta) = \alpha u(\succsim_i, \theta^r, \theta)$. Convex mixtures of the reference and alternatives result in averages over willingness to pay amounts in a smooth manner. When the reference threat changes by a small amount, the willingness to pay changes by a small amount as well.

The value $u(\succsim_i, \theta^r, \theta)$ is a compensating variation value. It makes that partner indifferent between two pairs of threats by changing the value of the consumption bundles the right amount. For partner f, this value solves: $x_f(p, \eta^{*r}y + u(\succsim_f, \theta^r, \theta)) \sim_f x_f(p, \eta^* y)$. Similarly, for partner m, this compensating variation amount is the one that makes them indifferent between the consumption bundle associated with the reference bargaining solution, and the alternative: $x_m(p, (1-\eta^{*r})y + u(\succsim_m, \theta^r, \theta)) \sim_m x_m(p, (1-\eta^*)y)$. When $\theta^r \succ_i \theta \; \forall i \in \{f, m\}$, $u(\succsim_i, \theta^r, \theta) < 0$, when $\theta \succ_i \theta^r$, $u(\succsim_i, \theta^r, \theta) > 0$, and when $\theta^r \sim_i \theta$, $u(\succsim_i, \theta^r, \theta) = 0$. People are willing to pay a positive amount to switch to a prefer ed solution, and have to be paid to be willing to switch from a preferred solution to one they consider to be worse.

Definition 3: Let $< \theta, \succsim_f, \succsim_m > \in (\Sigma_f \times \Sigma_m) \times \rho^2$*, be a bargaining problem. Let* θ*,* $\theta^r \in (\Sigma_f \times \Sigma_m) \times \{0,1\}^2$ *be two potential* solutions *to* $< \theta, \succsim_f, \succsim_m >$*. The reference outcome* θ^r *is* immune to the appeal of the alternative θ *in* $< \theta, \succsim_f, \succsim_m >$ *if one of the following is true:*

1. $\theta^{*r} \succ_f \theta^*$ and $x_f(p, \eta^{*r}y) \succ_f x_f(p, \eta^* y + \tau_m)$ where $\tau_m \in [0, u(\succsim_m, \theta^r, \theta)]$, or
2. $\theta^{*r} \succ_m \theta^*$ and $x_m(p, (1-\eta^{*r})y) \succ_m x_m(p, (1-\eta^*)y + \tau_f)$ where $\tau_f \in [0, u(\succsim_f, \theta^r, \theta)]$

This definition simply says that there is no transfer that one partner could give to the other that makes both happy to switch from the reference threats to an alternative potential solution. In the first condition, partner m wants to switch and can offer some transfer amount to partner f equal to or less than their compensating variation value, but this is not enough to induce partner f to switch. In the second condition, the opposite scenario plays out. Partner f wants to switch but partner m does not, and there is not a feasible transfer from f to m that makes m prefer to switch. In these cases, the reference outcome will remain the family's outcome despite the appeal to switch to an alternative.

Consider an example based on Ramos' (2016) work on violence and power in the household. In her study context, northern Ecuador in 2011, 35% of married men abuse their partners and only a small fraction of those relationships end in divorce. In this context, men's threat would be some form of physical, sexual, or emotional abuse (or some combination), and women's threat could be divorce or some mild form of non-cooperation. Ramos shows that this violence increases men's resource shares but reduces women's productive capacity, reducing the amount of surplus available to the family and allocating a larger amount of this surplus to men.[5] Take a reference bargaining solution where men choose to physically abuse their partners (15% of households in this context) and women choose divorce as their threat but do not act on it ($I_f = 0$, $I_m = 1$). Consider an alternative with the same specified threats, but where neither partner acts on them ($I_f = 0$, $I_m = 0$). This realistic reference *bargaining solution* is immune to the appeal of this alternative if women are unable to provide a transfer (monetary or in kind) to their partners that compensates men for their loss of power, conditional on the total surplus being greater in the alternative. In the setting where we observe violence, men are getting so much more of the family's surplus that it is not possible for women to appeal successfully.

3.2.2 Uniqueness of the Appeals-Immune, Individually-Rational Solution

We now turn to a discussion of existence and uniqueness of a *solution* to the bargaining problem $< \theta, \succsim_f, \succsim_m >$. We follow Hanany and Safra (2000) to show that, for preferences satisfying certain assumptions, there is only one possible reference solution, θ^*, such that, for any possible alternative θ, θ^* is immune to the appeal of the alternative θ in $< \theta, \succsim_f, \succsim_m >$.[6] The following assumptions on the threat spaces, and the preferences, must hold for existence and uniqueness.

Assumption 3.3.1 — Compactness: Consider $\theta^a, \theta^b \in (\Sigma_f \times \Sigma_m)$. *Consider the infinite se-*

[5] Our model is mildly more general than hers because it makes divorce endogenous, let's women have threats as well, and is defined over the ordinal set of preferences, as opposed to the cardinal space of utility function representations. An application of this model to her study context might be fairly similar to the model she builds and estimates. Allowing women to threaten divorce would likely result in a prediction of the sharing rule that favors men slightly less than what she finds.

[6] The following assumptions are correspond to assumptions DOM, Q, CCE, and H in Rubinstein, Safra, and Thomson (1992).

quence of solutions $\{\theta^a, \alpha_1\theta^a + (1-\alpha_1)\theta^b, \alpha_2\theta^a + (1-\alpha_2)\theta^b, ..., \theta^b\}$ $\forall \alpha_1, \alpha_2, \ldots, \alpha_\infty \in [0,1]$ *such that* $\alpha_1 < \alpha_2 < \ldots < \alpha_\infty$. *Then* $\theta_a \in (\Sigma_f^* \times \Sigma_m^*) \implies \theta_b \in (\Sigma_f^* \times \Sigma_m^*)$. *That is,* $(\Sigma_f^* \times \Sigma_m^*)$ *contains all of its limit points.*

Assumption 3.3.2 — Convexity: $\forall i \in \{f, m\}$, $\forall \alpha \in [0,1]$ $\forall \theta^{*a}, \theta^{*b} \in (\Sigma_f^* \times \Sigma_m^*)$, $\theta^{*a} \succ_i \theta^{*b} \implies \alpha\theta^{*a} + (1-\alpha)\theta^{*b} \succ_i \theta^{*b}$.

Assumption 3.3.3 — Concavity of Induced Utility Functions: Consider $\theta^a, \theta^b \in (\Sigma_f \times \Sigma_m)$, $\theta^a \succ_i \theta^b$ *for* $i \in \{f, m\}$, $\alpha \in (0,1)$. *The induced utility function* $u(\succsim_i, \theta^a, \theta^b)$ *is* concave *if* $u(\succsim_i, \alpha\theta^a + (1-\alpha)\theta^b, \theta^b) > \alpha u(\succsim_i, \theta^a, \theta^b) + (1-\alpha)u(\succsim_i, \theta^b, \theta^b) = \alpha u(\succsim_i, \theta^a, \theta^b) + 0 = \alpha u(\succsim_i, \theta^a, \theta^b)$.

Assumption 3.3.4 — Loss Aversion: Preferences that have induced utilities that satisfy the following property are said to be loss-averse *preferences: for* $\theta^r \succ_i \theta$, $u(\succsim_i, \theta, \theta^r) < |u(\succsim_i, \theta^r, \theta)|$. *That is, person i needs more compensation to be willing to move from a good reference to a bad alternative than they would be willing to pay to move from that bad alternative to the good reference.*

Assumption 3.3.5 — Weakly Symmetric Loss Aversion (WSLA): Consider two possible solutions, $\theta, \theta^r \in (\Sigma_f \times \Sigma_m))$. *When partners in a marriage have preferences that jointly exhibit the following property, they are said to be* weakly symmetric in loss aversion:

1. $\succsim_i$ *is* loss averse *for* $i \in \{f, m\}$, *and*

2. *for* $\theta^r \succsim_f \theta$ *and* $\theta \succsim_m \theta^r$, $u(\succsim_f, \theta^r, \theta) \leq |u(\succsim_m, \theta^r, \theta)|$, *and*

3. *for* $\theta^r \succsim_m \theta$ *and* $\theta \succsim_f \theta^r$, $u(\succsim_m, \theta^r, \theta) \leq |u(\succsim_f, \theta^r, \theta)|$.

Let the set of preferences that satisfies assumptions 3.3.2 - 3.3.5 be denoted ρ^{WSLA}. Let the compact subset of *individually-rational solutions* be denoted $(\Sigma_f^* \times \Sigma_m^*)^C$. Formally, assumptions 3.3.4 and 3.3.5 play the same role as the bounds assumption in (E2) of Theorem 3.2 of Hanany and Safra (2000), and their risk aversion assumption. In the setting with lotteries, the risk aversion assumption is almost the exact analog of the loss aversion in the setting without uncertainty. The WSLA assumption bounds the induced utilities that result from appeals to move away from the appeals-immune *solution*, as we see in the proof of the following proposition:

Proposition - Uniqueness: Let $<\theta, \succsim_f, \succsim_m> \in (\Sigma_f^* \times \Sigma_m^*)^C \times (\rho^{sm} \cap \rho^{WSLA})^2$ *be a* bargaining problem. *Let* $\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*} \times \{0,1\}^2$ *be a potential* solution *to* $<\theta, \succsim_f, \succsim_m>$. *For all possible alternatives,* $\theta \in ((\Sigma_f \times \Sigma_m) \times [0,1]^2) \setminus \{\theta^*\}$, θ^* *is immune to the appeal of* θ *in* $<\theta, \succsim_f, \succsim_m>$ *if:*

$$\theta^* = \underset{\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*}}{argmax} \prod_{i \in \{f,m\}} u(\succsim_i, \theta^*, \theta).$$

Proof. $\theta^* = \underset{\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*}}{argmax} \prod_{i \in \{f,m\}} u(\succsim_i, \theta^*, \theta)$ *is unique because* $u(\succsim_i, \theta^*, \theta) \in \mathbb{R}_+$ *and* $u(\succsim_i, \theta^*, \theta)$ *is concave. At this point,*

$$\theta^* = \underset{\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*}}{argmax} \prod_{i \in \{f,m\}} u(\succsim_i, \theta^*, \theta) \iff$$

$$\frac{\partial u(\succsim_f, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_f^*} = -\frac{\partial u(\succsim_m, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_f^*} \text{ and}$$

$$-\frac{\partial u(\succsim_f, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_m^*} = \frac{\partial u(\succsim_m, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_m^*}$$

These first order conditions say that at the solution, $\theta^* = (\theta_f^*, \theta_m^*)$, specifying an alternative reference results in a marginal benefit (higher willingness to pay to switch from the reference to the alternative) to one partner equal to the marginal cost (higher willingness to pay to avoid switching from the reference to the alternative) to the other player. We can move away from the equilibrium in one of four directions of interest, corresponding to the two conditions in Definition 1. When we move away from the equilibrium, the equalities in the first order conditions hold as inequalities. Because preferences satisfy

WSLA, we know which direction the inequality will go in for each case. To formalize the idea of "moving away from the equilibrium," consider also some convex mixture of θ^* and θ. These four conditions are, for some small $\alpha \in (0,1)$:

1. Partner f prefers the change in partner f's threat, and partner m does not: $(\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*) \succ_f (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_m ((\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$. If this is true, and because preferences satisfy the WSLA assumption, the FOCs imply $u(\succsim_f, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*)) < |u(\succsim_m, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*))|$ and so no transfer, τ_f, exists such that $(\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$ successfully appeals θ^*.

2. Partner f prefers the change in partner m's threat, and partner m does not: $(\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m) \succ_f (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_m (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m))$. If this is true, and because preferences satisfy the WSLA assumption, then the FOCs imply $u(\succsim_f, (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m), (\theta_f^*, \theta_m^*)) < |u(\succsim_m, (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m), \theta_m^*), (\theta_f^*, \theta_m^*))|$ and so no transfer, τ_f, exists such that $(\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m)$ successfully appeals θ^*.

3. Partner m prefers the change in partner f's threat, and partner f does not: $((\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*) \succ_m (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_f ((\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$. If this is true, and because preferences satisfy the WSLA assumption, then the FOCs imply $|u(\succsim_f, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*))| > u(\succsim_m, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*))$ and so no transfer, τ_m, exists such that $(\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$ successfully appeals θ^*.

4. Partner m prefers the change in partner m's threat, and partner f does not: $(\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m) \succ_m (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_f (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m))$. If this is true, and because preferences satisfy the WSLA assumption, then the FOCs imply $|u(\succsim_f, (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m), (\theta_f^*, \theta_m^*))| > u(\succsim_m, (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m), \theta_m^*), (\theta_f^*, \theta_m^*))$ and so no transfer, τ_m, exists such that $(\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m)$ successfully appeals θ^*.

Because $(\Sigma_f^* \times \Sigma_m^*)^C$ is compact, and because the above argument holds for any possible reference, the θ^* that maximizes the Nash product of induced utilities over WSLA preferences is the unique *appeals-immune solution* to the *bargaining problem* $< \theta, \succsim_f, \succsim_m >$

$\in (\Sigma_f^* \times \Sigma_m^*)^C \times (\rho^{WSLA})^2$. ■

This model with endogenous threats and resource shares allows us to understand power dynamics in the family more fully. Who ever is capable of specifying a more damaging threat will have more control over the household decision making process. Compared to the standard collective model, this generic game with a nested limited commitment problem allows us to understand why we observe certain sharing rules for each family. It allows us to measure power dynamics without assuming that only one partner can specify or act on a threat that reduces household efficiency.

3.3 Identification in a Panel Data Setting

In this section, we discuss how to bring the model in Section 3.2 to data. We demonstrate the the expected value of bargaining power, conditional on partners threats and whether they act on them, is semi-parametrically identified in a panel data setting. The key to this strategy is to treat preferences as random variables in order to avoid strong functional form assumptions for utility and social welfare. In this case, the sharing rule is a function of random variables, and so is a random variable itself. Making specific assumptions on the distribution of preferences gives a functional form for the expected value of power, conditioned on endogenously determined outside options. We note which parameters are to be recovered from the data in the derivation, then give specific details for how researchers can do so in Section 3.3.2.

3.3.1 Deriving the Functional Form of the Expected Value of Women's Intra-Household Bargaining Power

First, note that the optimization problems in (3.3) can be rewritten without reference to the indirect utility functions by applying a compensating variation approach within the nested game with limited commitment, $B(\theta_f, \theta_m, I_f^*, I_m^*)$. For any threat, there is some transfer value that makes partners indifferent between the prices they face in the collective allocation ($I_i = 0$) and the inefficient allocation ($I_i = 1$) for $i \in \{f, m\}$. These transfer amounts make the partners indifferent between the price settings they face in both settings, but the real income they would have at their disposal may still differ across the two collective and less efficient equilibria. Denote these transfer amounts γ_f and γ_m, such that

$$V_f(L(A(\theta_f, \theta_m, I_f = 1, I_m)), \eta y) = V_f(L(A(\theta_f, \theta_m, I_f = 0, I_m)), \eta y + \gamma_f) \text{ and}$$

$$V_m(L(A(\theta_f, \theta_m, I_f, I_m = 1)), (1-\eta)y) = V_m(L(A(\theta_f, \theta_m, I_f, I_m = 0)), (1-\eta)y + \gamma_m).$$

By holding prices constant across the indirect utility functions, only the second argument in the function varies with different threats. The impact on shadow incomes is captured by the changes in the sharing rule and the production function, and the impact on the price setting is captured by the compensating variation amount. The indirect utility functions are strictly increasing in their second argument. As such, the problems in (3.3) can be re-written in a form that does not explicitly reference the indirect utility function. In this formulation, the individuals consider the real (shadow) income they would have at their disposal in the collective and inefficient allocations and chooses to implement their threat if doing so gives a greater real income. That is, the partners solve the following, equivalent problems, where both options are measured in dollar amounts:

$$\underset{I_f \in \{0,1\}}{max} \Bigg(\eta(\theta_f, \theta_m, I_f = 1, I_m) F(\theta_f, \theta_m, I_f = 1, I_m), \\ \eta(\theta_f, \theta_m, I_f = 0, I_m) F(\theta_f, \theta_m, I_f = 0, I_m) + \gamma_f \Bigg)$$

$$\underset{I_m \in \{0,1\}}{max} \Bigg((1 - \eta(\theta_f, \theta_m, I_f, I_m = 1)) F(\theta_f, \theta_m, I_f, I_m = 1), \\ (1 - \eta(\theta_f, \theta_m, I_f, I_m = 0) F(\theta_f, \theta_m, I_f, I_m = 0) + \gamma_m \Bigg)$$

We observe I_f^* and I_m^*, so we know which of the two options has a higher value. There are four possibilities: the fully efficient outcome $(I_f = 0, I_m = 0)$, he implements his threat but she does not $(I_f = 0, I_m = 1)$, she implements her threat but he does not $(I_f = 1, I_m = 0)$, and both implement their threats $(I_f = 1, I_m = 1)$. In each case, we get family-specific bounds on the possible values that the sharing rule can take. These bounds, when paired with a joint distributional assumption on the partners' preferences, result in a functional form for the expected value of the sharing rule that the econometrician can bring to data.

We derive the bounds next, then describe the remaining estimation steps: (1) assuming a distribution on preferences, (2) using a Heckman (1979) selection approach to recover estimates of threats for those who do not implement threats, (3) using a propensity score matching approach to recover the real (shadow) incomes in the unobserved cases, and (4) recovering the remaining parameters that determine the family-specific bounds on the sharing rule by estimating a fixed effects model using the simulated method of moments. To make equations more legible, let $\theta^e = (\theta_f, \theta_m, I_f = 0, I_m = 0)$, $\theta^{her} = (\theta_f, \theta_m, I_f = 1, I_m = 0)$, $\theta^{him} = (\theta_f, \theta_m, I_f = 0, I_m = 1)$, and $\theta^{both} = (\theta_f, \theta_m, I_f = 1, I_m = 1)$.

Case 1: The Fully Efficient Outcome: Because we observe $(I_f = 0, I_m = 0)$, we know that

$$\eta(\theta^e)F(\theta^e) + \gamma_f > \eta(\theta^{her})F(\theta^{her}) \text{ and}$$

$$(1 - \eta(\theta^e))F(\theta^e) + \gamma_m > (1 - \eta(\theta^{him}))F(\theta^{him}).$$

Solving these two inequalities for $\eta(\theta^e)$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^e) \in \left[\frac{\eta(\theta^{her})F(\theta^{her}) - \gamma_f}{F(\theta^e)}, 1 - \frac{(1 - \eta(\theta^{him}))F(\theta^{him}) - \gamma_m}{F(\theta^e)}\right].$$

Here, the econometrician observes $F(\theta^e)$ and will estimate values for θ^{her} and θ^{him} using a Heckman (1979) selection approach, and for $F(\theta^{her})$ and $F(\theta^{him})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^{her})$, $\eta(\theta^{him})$, γ_f, and γ_m from a simulated method of moments approach. These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Case 2: Only Partner f Implements Their Threat: Because we observe $(I_f = 1, I_m = 0)$, we know that

$$\eta(\theta^e)F(\theta^e) + \gamma_f < \eta(\theta^{her})F(\theta^{her}) \text{ and}$$

$$(1 - \eta(\theta^{her}))F(\theta^{her}) + \gamma_m > (1 - \eta(\theta^{both}))F(\theta^{both}).$$

Solving these two inequalities for $\eta(\theta^{her})$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^{her}) \in \left[\frac{\eta(\theta^e)F(\theta^e) + \gamma_f}{F(\theta^{her})}, 1 - \frac{(1 - \eta(\theta^{both}))F(\theta^{both}) - \gamma_m}{F(\theta^{her})}\right].$$

Here, the econometrician observes $F(\theta^{her})$ and will estimate values for θ^e and θ^{both} using a Heckman (1979) selection approach, and for $F(\theta^e)$ and $F(\theta^{both})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^e)$, $\eta(\theta^{both})$, and $\gamma \equiv \gamma_m - \gamma_f$ from a simulated method of moments approach (which we will explain next). These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Case 3: Only Partner m Implements Their Threat: Because we observe $(I_f = 0, I_m = 1)$, we know that

$$\eta(\theta^{him})F(\theta^{him}) + \gamma_f > \eta(\theta^{both})F(\theta^{both}) \text{ and}$$

$$(1 - \eta(\theta^e))F(\theta^e) + \gamma_m < (1 - \eta(\theta^{him}))F(\theta^{him}).$$

Solving these two inequalities for $\eta(\theta^{him})$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^{him}) \in \left[\frac{\eta(\theta^{both})F(\theta^{both}) - \gamma_f}{F(\theta^{him})}, 1 - \frac{(1-\eta(\theta^{e}))F(\theta^{e}) + \gamma_m}{F(\theta^{him})}\right].$$

Here, the econometrician observes $F(\theta^{him})$ and will estimate values for θ^e and θ^{both} using a Heckman (1979) selection approach, and for $F(\theta^e)$ and $F(\theta^{both})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^e)$, $\eta(\theta^{both})$, and $\gamma \equiv \gamma_m - \gamma_f$ from a simulated method of moments approach. These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Case 4: Both Partners Implement Their Threats: Because we observe $(I_f = 1, I_m = 1)$, we know that

$$\eta(\theta^{him})F(\theta^{him}) + \gamma_f < \eta(\theta^{both})F(\theta^{both}) \text{ and}$$

$$(1-\eta(\theta^{her}))F(\theta^{her}) + \gamma_m < (1-\eta(\theta^{both}))F(\theta^{both}).$$

Solving these two inequalities for $\eta(\theta^{both})$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^{both}) \in \left[\frac{\eta(\theta^{him})F(\theta^{him}) + \gamma_f}{F(\theta^{both})}, 1 - \frac{(1-\eta(\theta^{her}))F(\theta^{her}) + \gamma_m}{F(\theta^{both})}\right].$$

Here, the econometrician observes $F(\theta^{him})$ and will estimate values for θ^e and θ^{both} using a Heckman (1979) selection approach, and for $F(\theta^e)$ and $F(\theta^{both})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^e)$, $\eta(\theta^{both})$, and $\gamma \equiv \gamma_m - \gamma_f$ from a simulated method of moments approach. These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Distributional Assumptions on Preferences: At this point, we must make some assumption on preferences. Assuming that a specific functional form for utility functions adequately captures the preferences for all individuals in the sample is a strong, parametric assumption. We can make a weaker, semi-parametric assumption: that preferences and the social welfare function are random variables drawn according to some joint probability distribution function from $(\rho^{sm})^2 \times S$, where S is the set of permissible social welfare functions. Doing so implies a particular distribution of the sharing rule on the set determined by the observed values (I_f, I_m). This is because a (linear) bijection relates the specified threats, the pair (I_f, I_m), and the drawn random variables $(U_f, U_m, \tilde{U})$.[7]

A particularly tractable and weak distributional assumption is that preferences are conditionally uniformly distributed on some subset of $(\rho^{sm})^2 \times S$. Consider an arbitrarily drawn partition of $(\rho^{sm})^2 \times S$, with $c \in \mathbb{N}$ subsets $C \subset (\rho^{sm})^2 \times S$, where $\underset{i\in[1,c]}{\cup} C_i = (\rho^{sm})^2 \times S$. For each family in the population, let the triple $(U_f, U_m, \tilde{U})$ be uniformly drawn from the same subset, C_i. This is a very weak "similarity between partners" assumption since the partition can be drawn in any way (any number of subsets, with subsets of any size). This distributional assumption on preferences results in a uniform distribution of the sharing rule on the family-specific bounds derived above. As such, the estimator of the sharing rule can be written as follows in each of the four cases:

1. $\mathbb{E}[\eta(\theta^{e})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{her})F(\theta^{her}) - (1-\eta(\theta^{him}))F(\theta^{him}) - \gamma_f + \gamma_m}{F(\theta^{e})}\right)$
2. $\mathbb{E}[\eta(\theta^{her})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{e})F(\theta^{e}) - (1-\eta(\theta^{both}))F(\theta^{both}) + \gamma_f + \gamma_m}{F(\theta^{her})}\right)$
3. $\mathbb{E}[\eta(\theta^{him})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{both})F(\theta^{both}) - (1-\eta(\theta^{e}))F(\theta^{e}) - \gamma_f - \gamma_m}{F(\theta^{him})}\right)$
4. $\mathbb{E}[\eta(\theta^{both})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{him})F(\theta^{him}) - (1-\eta(\theta^{her}))F(\theta^{her}) + \gamma_f - \gamma_m}{F(\theta^{both})}\right)$

3.3.2 Identification

Summary of requirements: The researcher must have access to panel data with at least two waves for each household. Each wave must have demographic information about each partner (for instance their age, earnings, and education), and information about the threats that partners might specify (for instance, domestic abuse or divorce). From this information, the researcher can learn the threats partners specify, and whether they implement them. The values that the econometrician can estimate are the threats for partners who do not act on them, the real incomes for each household in the context

[7]For proofs that this bijective function exists and is linear, see Klein and Barham (2018).

we do not observe (for instance, in the case where $\{I_m = 0, I_f = 0\}$ when we observe $\{I_m = 1, I_f = 0\}$), and the expected value of the sharing rule. The assumptions required are distributional assumptions on preferences, distributional assumptions on the sharing rule in the cases we do not observe, and $H + 2$ distributional assumptions on regression coefficients to be estimated using the simulated method of moments (McFadden, 1989), where H is the number of households in the sample. These assumptions on the coefficient distributions are analogous to a single assumption on the distribution of the error term in a fixed effects model with H constraints on the intercepts, estimated using constrained least squares.

An optional, but helpful, data feature is an instrument that is correlated with how damaging each partners' threat is, and uncorrelated with whether or not they implement their threat. Lewbel and Pendakur (2019) use the thickness of each person's walls (a proxy for how likely it is that neighbors overhear domestic abuse) as an instrument in a similar setting. For instance, if you're worried about being overheard, you might specify a threat that does not entail loud volumes. If so, implementing the threat is uncorrelated with the likelihood of being overheard. Access to such an instrument makes estimation arguably less biased.

Summary of the strategy: There are four steps in estimation. First, the economist must quantify how damaging each threat is — i.e. specify a function that maps from the space of threats to the real numbers. The second step is to predict the threat for those who the researcher does not observe acting on a threat. The third is to fit a household fixed effects model with family-specific restrictions on each intercept. The fourth is to use the recovered regression coefficients and threat values to recover the expected value of women's intra-household bargaining power.

Step 1 of the estimation strategy: The researcher can quantify how damaging each threat is by using violence scales, like the one Ramos (2016) uses. This approach generates a scale from 0 to 1 that describes the severity of the threat. Each possibility is given some (normalized) value, and the score is the sum of the quantified individual actions comprising a threat. For instance, a women who is punched and verbally threatened with a weapon would have a higher score than a woman who is only threatened with a weapon. A women who is only punched would have a higher score than a woman who is only threatened.

In order to incorporate all types of threats into our model, the econometrician must also assign values to actions like divorce and the inefficient allocation of productive assets. This is a necessarily subjective exercise, but is based on the objective change in real incomes that each partner would have if the threat is exercised. Ideally, the threat scale would assign higher values to threats that reduce the real income available to an individual's partner by a greater amount. Since divorce (and any threat that ends cohabitation) results in both partners consuming at market prices, and consuming according to their own private incomes instead of a shared income, divorce can be given a higher scale value than violence. We suggest that any threat which only reduces the quality of the production function be given the lowest values in the scale, any threat that only reduces the quality of the consumption technology be given the next highest values, and any threat that affects both be given higher values. Various scales could be used to test the robustness of the results to various scale specifications.

Step 2 of the estimation process: The econometrician can recover the expected values of the threats that are not implemented using a Heckman (1979) selection approach. This allows us to recover estimates of these threats based on observable demographic characteristics, and an unobservable selection parameter, which describes the differences in specified threats between those who choose to implement them, and those who do not implement them. Disaggregating the sample by gender allows the econometrician to recover regression coefficients and selection parameters that differ across men and women. As such, estimate the following two models using full information maximum likelihood:

$$I_{f,h,t} = X_{f,h,t}\beta_{f,h,t} + z_{f,h,t}\beta_{instrument,f,h,t} + \epsilon_{f,h,t} \tag{3.8}$$

$$\theta_{f,h,t} = X_{f,h,t}\beta_{f,h,t} + \epsilon_{f,h,t}$$

$$I_{m,h,t} = X_{m,h,t}\beta_{m,h,t} + z_{m,h,t}\beta_{instrument,m,h,t} + \epsilon_{m,h,t} \tag{3.9}$$

$$\theta_{m,h,t} = X_{m,h,t}\beta_{m,h,t} + \epsilon_{m,h,t}$$

where t indexes time, h indexes the household that partners f and m belong to at the start of the panel, $z_{f,h,t}$ is an instrument for I_f (that is $cov(I_{f,h,t}, z_{f,h,t}) \neq 0$ and $cov(\theta_{f,h,t}, z_{f,h,t}) = 0$), $z_{m,h,t}$ is an instrument for I_m (that is $cov(I_{m,h,t}, z_{m,h,t}) \neq 0$ and $cov(\theta_{m,h,t}, z_{m,h,t}) = 0$), $X_{f,h,t}$ is a matrix of demographic characteristics for the women in the sample, and $X_{m,h,t}$ is a matrix of demographic characteristics for the men in the sample.[8] The results are the (nuisance) regression coefficients that relate observable characteristics to the severity of the specified threat, and the likelihood that the threat is carried out, plus the (nuisance) selection coefficient estimates that describe how much threats differ on average between those who implement them and those who do not. The predictions for threats of those who do not implement them are the fitted values from these Heckman (1979) selection models: $\hat{\theta}_{f,h,t}$ and $\hat{\theta}_{m,h,t}$. The researcher can use the fitted values for all individuals for consistency.

Step 3 of the estimation process: the econometrician must recover the shadow incomes for individuals in each case. They can use the predicted threats as propensity scores in a propensity score matching approach (Rosenbaum and Ruben, 1983). In each case, the econometrician needs to recover unobserved income values. We suggest using a nearest neighbor matching approach to recovering these parameters. For instance, if the researcher observes $(I_f = 0, I_m = 0)$ for some household, then they need to recover values for $F(\theta^{him})$ and $F(\theta^{her})$. We suggest using the income vale for the nearest household (nearest in the value of $\hat{\theta}_{f,h,t}$) such that that neighboring household has observed value $I_f = 1$ for the estimate of $F(\theta^{her})$. Symmetrically, we suggest using the income value for the closest

[8]The model does not require instruments, but in that case it does require an additional joint normality assumption.

neighbor in the value of $\hat{\theta}_{f,h,t}$ (where that neighbor has observed value $I_m = 1$) to recover the estimate of $F(\theta^{him})$. Denote these matching estimates as $\hat{F}(\hat{\theta}^{him})$ and $\hat{F}(\hat{\theta}^{her})$.

Step 4 in the estimation process: at this point, the remaining unknown parameter values in the bounds are the sharing rule values in the unobserved cases, and the compensating variation values, γ_f and γ_m. In order to recover these values, we suggest a fixed effects strategy. Since the approach is the same across all four cases, we only describe it for the fully efficient case. The first step in this panel estimation is to notice that the functional form for the expected value of the equilibrium sharing rule can be rearranged so that income is a linear function of the estimated shadow incomes in the unobserved cases, and unknown parameters. In the fully efficient case, this gives the following expression, suppressing notation:

$$F(\theta^e) = \frac{\eta(\theta^{her})F(\theta^{her}) - (1 - \eta(\theta^{him}))F(\theta^{him}) - \gamma_f + \gamma_m}{\mathbb{E}[\eta(\theta^e)] - 1}$$

the researcher can estimate this equation using a fixed effects approach, with the following definitions for regression coefficients:

$$F(\theta^e)_{h,t} = \beta_{0,h} + \beta_1 \hat{F}(\hat{\theta}^{her})_{h,t} + \beta_2 \hat{F}(\hat{\theta}^{him})_{h,t} + \epsilon_{h,t} \text{ such that} \tag{3.10}$$

$$\beta_{0,h} \equiv \frac{\gamma_m - \gamma_f}{\mathbb{E}[\eta(\theta^e)] - 1}$$

$$\beta_1 \equiv \frac{\eta(\theta^{her})}{\mathbb{E}[\eta(\theta^e)] - 1}$$

$$\beta_2 + \epsilon_{h,t,2} \equiv \frac{1 - \eta(\theta^{him})}{\mathbb{E}[\eta(\theta^e)] - 1}$$

$$\epsilon_{h,t,2} \equiv \epsilon_{h,t} \hat{F}(\hat{\theta}^{him})$$

$$\eta^e \in [0,1], \quad \eta^{her} \in [0,1], \quad \text{and } \eta^{him} \in [0,1].$$

Estimating (3.10) using the simulated method of moments requires the following assumptions: distributional assumptions for the H intercept parameters, $\beta_{0,h}$; distributional assumptions for the two slope parameters, β_1 and β_2; and distributional assumptions for the off-equilibrium sharing rule values, $\eta(\theta^{her})$ and $\eta(\theta^{him})$. Under these assumptions, there are two remaining unknowns per household in (3.10), and at least two equations per household in the panel. After the simulated method of moments estimation, the estimates of the regression coefficients can be plugged into the definitions in (9) to recover the expected value of the sharing rule in the equilibrium, and the difference between the two partners' compensating variation values, $\gamma_m - \gamma_f$. The constraints that the sharing rules only take values between 0 and 1 limit the regression coefficients' supports.

We suggest the following procedure for selection assumptions for the regression coefficients. First fit the regression in (3.10) without taking into account the adding up constraints ($\eta^e \in [0,1], \quad \eta^{her} \in [0,1], \quad$ and $\eta^{him} \in [0,1]$). This gives naive regression coefficients. Assume that the simulated method of moments coefficients for (3.10) are truncated normal distributions with means given by the naive estimators, and variance given by the corresponding naive standard errors. For the off-equilibrium sharing rule values, we recommend assuming uniform distributions on the unit interval. With a large number of simulations, this process will give consistent estimators of the equilibrium sharing rule and the compensating variation value difference.

3.4 Conclusion

It is well known that policy makers can influence intra-household bargaining power by increasing the value of one partners' outside option (e.g. Mazzocco, 2007; Voena, 2015). In the standard limited commitment model, when policies change the value of the outside option so much that one partner newly prefers to leave the marriage, bargaining power in the family is renegotiated until that partner is indifferent between leaving and staying. Only drastic changes to the outside option result in a change in bargaining power, since most changes will be insufficient to make either partner prefer divorce to the current marriage contract. As such, power is expected to be fairly static, and changes in power dynamics are often expected to be small (see, e.g., Lise and Yamada, 2019).

When partners can specify outside options besides divorce, changes in power can be fairly frequent. Partners might subtly shift their specified threat, resulting in a different sharing rule, in response to smaller changes to relative outside options. Changes that would not cause the participation constraint to bind if the outside option is divorce might cause it to bind if the outside option is changing the amount of time allocated to labor (as in Walther, 2018) or the savings technology used (as in Schaner, 2015).

A different picture of power in the family emerges. Policy makers can hope to influence power dynamics even with "small" changes to the relative value of outside options. Events like minor salary increases, or even holiday bonuses, can change the balance of power in the family. Power, once collective models make outside options endogenous, appears to be highly variable and contested, not static.

Can policy makers influence partners' choices to specify certain threats, or pass policies that cause individuals to chose to not implement their threats? Any policy that expands individuals' space of possible threats to include more attractive possible outside options will accomplish the former. Any policy that changes the costs and benefits of implementing may achieve the latter. For instance, higher ratios of female-to-male police officers in the United States increases the likelihood that women report domestic abuse, increasing the cost of implementing that threat for men, and reducing the likelihood that they implement their threat (Miller and Segal, 2019). Our model predicts that such a policy would increase women's bargaining power.

In general, we would expect families to be more efficient in settings where men and women have more equal rights and opportunities. Any policy that increases equality in institutions outside of the family — like labor markets or divorce courts — gives the less powerful partner additional recourse against more damaging threats. This changes the threats that obtain in equilibrium, and shifts the bargaining power dynamic towards equality within the family.

Reference List

Apps, P. F., & Rees, R. (1988). Taxation and the Household. Journal of Public Economics, 35(3), 355-369.

Basu, K. (2006). Gender and Say: A Model of Household Behaviour with Endogenously Determined Balance of Power. The Economic Journal, 116(511), 558-580.

Browning, M., Chiappori, P. A., & Lewbel, A. (2013). Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power. Review of Economic Studies, 80(4), 1267-1303.

Chiappori, P. A. (1988). Rational Household Labor Supply. Econometrica, 63-90.

Chiappori, P. A. (1992). Collective Labor Supply and Welfare. Journal of Political Economy, 100(3), 437-467.

Chiappori, P. A., & Mazzocco, M. (2017). Static and Intertemporal Household Decisions. Journal of Economic Literature, 55(3), 985-1045.

De Mel, S., McKenzie, D., & Woodruff, C. (2009). Are Women More Credit Constrained? Experimental Evidence on Gender and Microenterprise Returns. American Economic Journal: Applied Economics, 1(3), 1-32.

Donni, O., & Molina, J. A. (2018). Household Collective Models: Three Decades of Theoretical Contributions and Empirical Evidence. IZA Discussion Paper No. 11915. Available at SSRN: https://ssrn.com/abstract=3286175

Hanany, E., & Safra, Z. (2000). Existence and Uniqueness of Ordinal Nash Outcomes. Journal of Economic Theory, 90(2), 254-276.

Klein, M. J., & Barham, B. L. (2019). Point Estimates of Household Bargaining Power Using Outside Options. In University of Wisconsin-Madison, Staff Paper No. 590 March 2018.

Harsanyi, J. C. (1986). Rational Behaviour and Bargaining Equilibrium in Games and Social Situations. CUP Archive.

Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. Econometrica, 153-161.

Lewbel, A., & Pendakur, K. (2019). Inefficient Collective Households: Abuse and Consumption. Boston College Unpublished Manuscript

Ligon, E. A. (2011). Dynamic Bargaining in Households (with an Application to Bangladesh). Available at SSRN 1776810.

Lise, J., & Yamada, K. (2019). Household Sharing and Commitment: Evidence From Panel Data on Individual Expenditures and Time Use. The Review of Economic Studies, 86(5), 2184-2219.

Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. The Review of Economic Studies, 74(3), 857-895.

McFadden, D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. Econometrica 995-1026.

Miller, A. R., & Segal, C. (2019). Do Female Officers Improve Law Enforcement Quality? Effects on Crime Reporting and Domestic Violence. The Review of Economic Studies, 86(5), 2220-2247.

Ramos, A. (2016). Household Decision Making with Violence: Implications for Transfer Programs. mimeo. Barcelona: Universitat Autonoma de Barcelona.

Rosenbaum, P. R., & Rubin, D. B. (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika, 70(1), 41-55.

Rubinstein, A., Safra, Z., & Thomson, W. (1992). On the Interpretation of the Nash Bargaining Solution and its Extension to Non-Expected Utility Preferences. Econometrica, 1171-1186.

Schaner, S. (2015). Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings. American Economic Journal: Applied Economics, 7(2), 135-74.

Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household. Journal of Political Economy, 104(5), 1010-1046.

Voena, A. (2015). Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples? American Economic Review, 105(8), 2295-2332.

Walther, S. (2018). Noncooperative Decision Making in the Household: Evidence From Malawi. Journal of Development Economics, 134, 428-442.