University of Wisconsin-Madison
Department of Agricultural & Applied Economics

Staff Paper No. 590 Revised, July 2019

Point Estimates of Household Bargaining Power Using Outside Options

By

Matthew J. Klein and Bradford L. Barham

AGRICULTURAL &

APPLIED ECONOMICS

STAFF PAPER SERIES

Point Estimates of Household Bargaining Power Using Outside Options

Matthew J. Klein and Bradford L. Barham*

July 2, 2019

Abstract: We demonstrate a new way to recover point estimates of bargaining power from a collective model of the family with limited commitment. Estimating bargaining power for each family requires that researchers first model each partners' earnings opportunities and predict their incomes in their relevant outside options. This strategy has several distinct advantages over other contemporary measurement options, such as not needing assumptions on utility functional forms, distribution factors, or private assignable goods. We use our method to generate the first point estimates of women's power for a majority of the conjugal couples in the Progresa cash transfer experiment. We find that the mother in the median family had 25% the decision making power that her partner had before the cash transfer program, and 74% after treatment. We find that this empowerment caused a substantial improvement in family diet at both the intensive and extensive margins. (*JEL* D13, I15, I38)
Keywords: Collective Model, Bargaining Power, Outside Option, Human Capital

Acknowledgements: The authors thank Jean-Paul Chavas, Seth Gitter, Brian Gould, Gustavo Gordillo, Ana Paula Melo, James Manley, Travis McArthur, Craig McIntosh, Daniel Phaneuf, Esteban Quinones, Laura Schechter, Emilia Tjernström, Denni Tommasi, and anonymous referees for their suggestions, assistance, and encouragement. Josh Alfonso helped with the construction of the panel datasets. All errors or omissions are the sole responsibility of the authors.

*The authors are affiliated with the University of Wisconsin, Madison, Department of Agricultural and Applied Economics. Contact them at mjklein2@wisc.edu and bradford.barham@wisc.edu. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. Declarations of interest: none.

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).

However, research on this subject is hampered by a fundamental measurement challenge: power levels are not observable.[1] 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 consumption 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' joint decision-making process is subject to limited commitment; and (4) parents' outside options can either be divorce or an inefficient equilibrium in the family. 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 bargaining power for each family.

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 rely on predicting the earnings each partner would have in their outside option.

Our primary contribution is to demonstrate semi-parametric identification of these new estimators and suggest estimation strategies for them, increasing the tools available to re-

[1] See Doss (2013) for a discussion.

searchers aiming to study power and gender in the family. Our estimators have several advantages over other contemporary measurement options, which we outline in a literature review section. These benefits stem from the fact that these outside-option earnings predictions are built on models of individuals' characteristics, not on models of household characteristics (like the demand for certain goods). The strength of the collective model is that partners are allowed to have distinct preferences; our semi-parametric identification and estimation strategy preserves this characteristic.

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/Oportunidaes/Prospera) empower women? We provide the first 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 Rubelcava, 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, Tommasi (2019) and Sokullu and Valente (2017) have contributed to this literature by estimating resource shares for mothers, fathers, and children in this experimental sample. Valente and Sokullu (2017) 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 just 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 similar cash transfers learn from the unambiguously positive results we document.

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 here on)

[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.

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]

These models, and recent extensions by Wolf (2016), Dunbar, Lewbel, and Pendakur (2017), Chiappori and Kim (2017), and Sokullu and Valente (2017), 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 for a small portion of their sample.

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 divorce legislation can influence bargaining power. Lise and Yamada (2014) 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. We avoid this assumption by following a different prediction strategy, which we elaborate in the next section.

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

[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.

[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.

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 transfer, 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.[8]

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 or allocation data. Our primary agenda is to provide 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

[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."

[8]However, while there is a bijection between this experimental estimator and the sharing rule in the collective model, this method does not allow for the explicit recovery of the sharing rule.

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.

3 The Model

We develop the model in four steps. First, we introduce 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 options more explicit. Last, we children to the model as purely public goods (as in Blundell, Chiappori, and Meghir, 2005) in order to generate hypotheses on children's welfare and their mother's bargaining power.

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:

$$\underset{x_f, x_m}{max} \quad \tilde{U}[U_f(x_f), U_m(x_m)] \text{ subject to } p_f' x_f + p_m' x_m = y \tag{1}$$

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, $\eta \in [0, 1]$:

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

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

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:

$$\underset{\eta \in [0,1]}{max} \quad \tilde{U}[V_f(\eta y, p), V_m((1-\eta)y, p)]. \tag{2}$$

The solution to (2) results in the demand functions that also solve (1). We will proceed from the formulation of the collective model in (2).

Note that individuals are only consuming private goods in (1) and (2). Introducing the consumption technology in section 3.3, and the purely public good in section 3.4, relaxes this strong assumption completely. This allows us to include the wellbeing 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, and so it is challenging to otherwise incorporate their utility functions into a bargaining model like this one.

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.

3.2 Limited Commitment

The decision makers live in a family with limited commitment to cooperation. 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). 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 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 limited commitment relaxes the assumption of *ex ante* Pareto efficiency

used in the collective model with full commitment above.

Partners choose "In" or "Out" simultaneously. If both players choose "In" they get payoffs from the collective model according to the sharing rule. If one or both choose "Out" they both get the outside option payoffs. These problems can be written as $max\{V_f(\eta y, p), V_f(y_f^o, p)\}$ 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 $max\{\eta y, y_f^o\}$ and $max\{(1-\eta)y, y_m^o\}$.

We can write this as a single program, which contains the model in (2) as a special case (as in, e.g., Mazzocco, 2007).[9] The household's problem is given in (3), which has the corresponding Kuhn-Tucker formulation given in (4), and which results in the following first order conditions (with surpressed notation):

$$\max_{\eta \in [0,1]} \quad \tilde{U}[V_f(\eta y, p), V_m((1-\eta)y, p)] \text{ s.t. } \eta y \geq y_f^o \text{ and } (1-\eta)y \geq y_m^o, \tag{3}$$

$$\max_{\eta} \quad L = \tilde{U}[V_f(\eta y, p), V_m((1-\eta)y, p)] + \lambda_1(\eta y - y_f^o) + \lambda_2((1-\eta)y - y_m^o) + \lambda_3\eta + \lambda_4(1-\eta), \tag{4}$$

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

$$[\lambda_1] : \lambda_1 \geq 0; \quad \eta y - y_f^o \geq 0; \quad \lambda_1(\eta y - y_f^o) = 0$$

$$[\lambda_2] : \lambda_2 \geq 0; \quad (1-\eta)y - y_m^o \geq 0; \quad \lambda_2((1-\eta)y - y_m^o) = 0$$

$$[\lambda_3] : \lambda_3 \geq 0; \quad \eta \geq 0; \quad \lambda_3\eta = 0$$

$$[\lambda_4] : \lambda_4 \geq 0; \quad 1-\eta \geq 0; \quad \lambda_4(1-\eta) = 0$$

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:

[9] If there is full commitment (players cannot chose the outside option), players always choose "In", the participation constraints never bind, and the solution to (3) also solves (2).

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

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 $[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}]$ 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 $[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}]$ includes the cases where the participation constraints bind as special cases.[10]

Power could be recovered from $[\eta]$ alone. By assuming functional forms for each partner's utility functions, and a functional form for the social welfare function, and by assuming an interior solution, 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 QUAIDS demand functions, and a logit functional form for the sharing rule; DLP assume PIGLOG indirect utility functions, and that there are similarity 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 DLP (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 utility function describes individuals' 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 outside options. 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 $[\lambda_1]$, and $[\lambda_2]$, and the prediction models $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$, pin down the set $[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}]$ for each family. The estimator of this set is $[\frac{F(X_f,\psi_f)}{y}, 1 - \frac{F(X_m,\psi_m)}{y}]$, and will have family-specific upper and lower bounds. To move from an estimate on the set of possible bargaining power dynamics a family may have, to an estimator of the exact power dynamic a family has, we

[10] 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.

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. We make no assumption on what function describes any persons preferences.

Our semiparametric approach is to restrict the possible distributions that describe the population's set of utility functions. One especially useful, 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 useful because, if partners' can have any utility function, then any value of $\eta \in [\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}]$ can solve $[\eta]$, and so there is a uniform distribution on this set of possible power dynamics. As such, a natural estimator for bargaining power would be the expected value of this distribution, which is its central point. This uniformity assumption on preferences, then, translates to a family-specific uniform distribution on possible power dynamics, and an estimator for bargaining power:

$$\hat{\eta} = \mathbb{E}\left[\eta \quad \middle| \eta \sim \text{Uniform}\left[\frac{y_f^o}{y}, 1 - \frac{y_m^o}{y}\right]\right] = \frac{1}{2} + \frac{1}{2}\frac{F(X_f, \psi_f) - F(X_m, \psi_m)}{y} \tag{5}$$

However, it may be that some preference 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 partition, c_l, with 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 partition g_l. We can refine our second main assumption so that each persons' 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 (5) obtains from the same line of reasoning as before.

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 (5). It doesn't matter whether some types of utility functions are more common than others - that amounts to more families drawing form certain subsets of C.

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. See Appendix A for technical details and a proof.

This is a useful estimator because it does not require that researchers assume a particular functional form for utility, make any assumptions on how C is partitioned, or observe which subset of C a family's utility functions are drawn from. Instead, the researcher must estimate earnings for each partner in their outside option. Assuming a functional form for earnings (or data driven parameter selection technique) is more justifiable than assuming a utility functional form, since the researcher can find the method that minimizes mean squared error, potentially subject to using an unbiased estimator. In addition, economists can draw from a large literature on modeling earnings (e.g. Mincer 1974 and 1975; Heckman, 1977).

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 (5), note that it will depend on outside options even if neither partners' participation constraint

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 "In" versus "Out" 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." As such, these shadow prices are invariant to the household's allocation.

Then, couples choices of "In" and "Out" 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: $max\{V_f(\eta y, A(p)), V_f(y_o^f, p)\}$ 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 individual's gains from trade when the outside option is divorce.

Substituting in the value functions with the compensating variation parameters allows us to write the partners' problems as $max\{V_f(\eta y + \gamma_f, p), V_f(y_o^f, p)\}$ 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 write these problems as: $max\{\eta y + \gamma_f, y_f^o\}$ and $max\{(1-\eta)y + \gamma_m, y_m^o\}$. If the outside option is divorce, the problem in (3) becomes:

$$\underset{\eta\in[0,1]}{max} \quad \tilde{U}[V_f(\eta y, A(p)), V_m((1-\eta)y, A(p))] \text{ s.t. } \eta y + \gamma_f \geq y_f^o \text{ and } (1-\eta)y + \gamma_m \geq y_m^o \quad (6)$$

and if the outside option is an inefficient equilibrium in the family, the household's optimization problem is:

$$\underset{\eta\in[0,1]}{max} \quad \tilde{U}[V_f(\eta y, A(p)), V_m((1-\eta)y, A(p))] \text{ s.t. } \eta y \geq y_f^o \text{ and } (1-\eta)y \geq y_m^o \quad (7)$$

In the case that the outside option is an inefficient equilibrium in the family, then partners face the same Lindahl prices regardless of their choice of "In" or "Out." In this case, the semi-parametric identification results from (3) hold. In the case that the outside option is divorce, identification is slightly different. The first order conditions for the Lagrangian for (6) are given by $[\lambda_3]$ and $[\lambda_4]$ above, and:

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

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

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

From these first order conditions for (6), we know that an interior solution ($\lambda_3 = \lambda_4 = 0$) will solve $[\eta']$ and be bounded on $[\frac{y_f^o - \gamma_f}{y}, 1 - \frac{y_m^o - \gamma_m}{y}]$.

Parametric estimation of this model is far more challenging than it is for the model in (3). Assuming functional forms for utility and the social welfare function, and a form for A, $[\eta]$ 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, we point out that an alternative approach is available. Power is semi-parametrically point identified on a subset of the first order conditions and the prediction models: $[\eta']$, $[\lambda_1']$, $[\lambda_2']$, $F(X_f, \psi_f)$ and $F(X_m, \psi_m)$. However, a draw back of the semiparametric approach in this case is that the parameters γ_f and γ_m are not separately identified. Only their differ-

ence is identified. Applying the same logic about the distribution of the η that solves the equilibrium condition in $[\eta']$, we can write the functional form for the estimator for η as the conditional expectation:

$$\hat{\eta} = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_f, \psi_f) - F(X_m, \psi_m) + (\gamma_m - \gamma_f)}{y}\right). \tag{8}$$

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.

We suggest a panel-regression strategy to identifying γ and η for each family in equation (8). The additional identifying assumption is that observations exist for more than one period. To unpack this suggestion, note that rearranging (8) gives:

$$y = \frac{F(X_f, \psi_f) - F(X_m, \psi_m) + (\gamma_m - \gamma_f)}{2\hat{\eta} - 1}. \tag{9}$$

That is, income is linear in the difference between outside options 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 (9) can be re-written as in (10) for each period:

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

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

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

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

$$y_{h,t} = \beta_{0,h} + \beta_1(F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})) + \epsilon_{h,t} \tag{10}$$

[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.

In estimating (10), the econometrician must also specify the constraint that $\beta_1 + \epsilon_{1,h,t} \in [0, 1]$, otherwise the first order conditions $[\lambda_3]$ and $[\lambda_4]$ could be violated.[12] Estimating equation (10) 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 into our definitions above gives our estimators as a function of the fixed effect regression estimates:

$$\hat{\eta}_{h,t} = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_{f,h,t}, \psi_{f,h,t}) - F(X_{m,h,t}, \psi_{m,h,t})}{y_{h,t} - \tilde{\beta}_{0,h}}\right) \tag{11}$$

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

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. In the semi-parametric estimation approach, 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 (8) 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 (9) reduces to the form in (5) if $\gamma = 0$. The non-cooperative inefficient equilibrium in the family is the same as a divorce in which there

[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.

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.

3.4 Equality and Investment in Children

In this section, we derive the testable hypotheses for our empirical application. The primary question we want to answer with our model and data is whether increased gender equality in the household results in more investments in children. To gain insight into this topic, consider a purely public good in the family Q, which represents children's well-being. Partners now get utility from consumption and increased Q, and their utility functions are $U_f(x_f, Q)$ and $U_m(x_m, Q)$. The market price of increasing Q is p_Q, and the household transforms this to a Lindahl price $A(p_Q)$ via the consumption technology M.

Denote the demand function for Q as $h_Q(A(p), y, \eta)$. Consider two separate sharing rule values, $\eta' = \frac{1}{2}$ and $\eta'' < \eta'$. Under η' partners have equal decision making power. Under η'' women have less decision making power than their partners. We can learn about investmetns and equality by studying the difference $h_Q(A(p), y, \eta') - h_Q(A(p), y, \eta'')$. We want to know the conditions under which this difference is positive. To answer this question, consider the families problem:

$$\underset{\eta \in [0,1]}{max} \quad \tilde{U}[V_f(\eta y, A(p)), V_m((1-\eta)y, A(p))] \text{ s.t.} \tag{12}$$

$$V_f(\eta y, A(p)) = \underset{x_f, Q}{max} \quad U_f(x_f, Q) \text{ subject to } A(p_f)' x_f + A(p_Q)Q = \eta y,$$

$$V_m((1-\eta)y, A(p)) = \underset{x_m, Q}{max} \quad U_m(x_m, Q) \text{ subject to } A(p_m)' x_m + A(p_Q)Q = (1-\eta)y,$$

Denote the individuals' demand for Q as $h_{Q,f}(\eta y, A(P))$ and $h_{Q,m}((1-\eta)y, A(P))$. Then $h_Q(A(p), y, \eta) = h_{Q,f}(\eta y, A(P)) + h_{Q,m}((1-\eta)y, A(P))$. We can relate the individuals' demand for Q to their indirect utility functions by Roy's identity:

$$h_{Q,f}(\eta y, A(p)) = -\frac{\frac{\partial V_f(\eta y, A(p))}{\partial A(p_Q))}}{\frac{\partial V_f(\eta y, A(p))}{\partial \eta y}}, \text{ and}$$

$$h_{Q,m}((1-\eta)y, A(p)) = -\frac{\frac{\partial V_m((1-\eta)y, A(p)}{\partial A(p_Q))}}{\frac{\partial V_m((1-\eta)y, A(p))}{\partial (1-\eta)y}}.$$

By plugging in the individuals' demand, we know that the difference of interest, $h_Q(A(p), y, \eta') - h_Q(A(p), y, \eta'')$, will be positive when:

$$-\left(\frac{\frac{\partial V_f(\eta' y, A(p))}{\partial A(p_Q))}}{\frac{\partial V_f(\eta' y, A(p))}{\partial \eta' y}} + \frac{\frac{\partial V_m((1-\eta')y, A(p)}{\partial A(p_Q))}}{\frac{\partial V_m((1-\eta')y, A(p))}{\partial (1-\eta')y}}\right) > \left(\frac{\frac{\partial V_f(\eta'' y, A(p))}{\partial A(p_Q))}}{\frac{\partial V_f(\eta'' y, A(p))}{\partial \eta'' y}} + \frac{\frac{\partial V_m((1-\eta'')y, A(p)}{\partial A(p_Q))}}{\frac{\partial V_m((1-\eta'')y, A(p))}{\partial (1-\eta'')y}}\right) \tag{13}$$

Equation (13) tells us the condition under which increasing equality will increase investments in children. An increase in equality will induce an increase in partner f's individual demand for children's well-being, and a corresponding decrease in partner m's demand.[14] Household investments in children will increase when the increase in partner f's demand is larger than the decrease in partner m's demand. Equation (13) tells us that whether this occurs is a function of the relative marginal costs of investment, as well as the relative marginal benefits.

There are two approaches to testing whether equality maximizes investment in children, corresponding to the parametric and semi-parametric strategies introduced above. If researchers assume a functional form for V_f, V_m, and $\tilde{U}$ then they can estimate the values of the left hand side and right hand side of the above (perhaps using the estimation approach in BCL, which requires the researcher to also assume a functional form for η) to see if the inequality in (13) is supported by the sample data.

If the semi-parametric form strategy is used, then researchers can use the observed household consumption choices in a regression framework. Including the estimate of the resource share, and it's square, as explanatory variables in a regression can allow researchers to recover estimates of $\frac{\partial Q}{\partial \eta}$ and $\frac{\partial^2 Q}{\partial \eta^2}$.[15] In a situation where some exogenous shock increases the median family's power dynamic towards equality, studying only the first derivative is sufficient for understanding equality and investment in children in that sample. These estimates can be used to determine the power level at which investments in children are maximized in the sample data.

A natural question is whether preferences for children's well-being affect parents' divorce decisions. We leave for future research an extension of this model that makes the value of parents' outside options conditional on their preferences for Q. Under divorce, children are split between parents according to some legal ruling and familial agreement. Children might get some disutility from their parents decision, they might face different Lindahl prices after divorce, and, further, investments might decrease given that parents optimize over y_o^f and

[14] This is simply because demand for a good is increasing in the income available to the individual. It does not stem from assumptions about differences between men's and women's preferences for investing in their children. If we let f denote men's utility and m denote women's, the same argument holds. That is, the operative idea is equality, not some gender-specific element of preferences.

[15] If researchers are concerned about making parametric assumptions on the demand functions in these regression specifications, they can estimate a kernel density approach as in, for example, Cherchye et al (2015).

y_m^o, not ηy and $(1-\eta)y$. In essence, a parent's threat of divorce is less credible if they are concerned that divorce will have negative repercussions for their children. As such, parents' relative preferences for Q will partially determine household power dynamics, since they determine the relative concern over these negative repercussions.[16]

4 Application to Mexico's Welfare and Empowerment Program

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; Rubelcava, Teruel, and Thomas, 2009) and because nutrition is an important in children's human capital accumulation. Further, there are documented links between power and diet in other contexts (e.g. Thomas, 1990), and we'd like to see if the positive relationships found elsewhere are present in this context too. We employ our semi-parametric estimation strategy, using Heckman selection models to predict y_f^o and y_m^o. We present the empirical analogs our our hypotheses at the end of this section, and the results in Section 5.

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

[16]This also allows for an analysis of how differences in perceptions between partners influences how credible the threat of divorce is, and thus bargaining power. If both partners know that one believes divorce will have a large negative effect on their kids, and the other believes that the effect will be small, the household bargaining power will reflect this difference in beliefs. This may explain bargaining power and divorce rate differences for children who's partners got divorced.

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: 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%
Guerrero	7.4%	8.0%	7.5%
Hidalgo	15.8%	16.1%	16.4%
Michoacán	11.7%	12.0%	12.0%
Puebla	16.7%	16.0%	14.9%
Queretaro	4.9%	5.0%	4.7%
San Luis Potosí	16.4%	16.8%	16.7%
Veracruz	27.0%	26.5%	27.8%
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.[17] The median household

[17]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

in our sample has a combined income of 265 pesos/week in 1997;[18] 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 for additional descriptive statistics.[19]

In Table 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.[20] 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).[21] 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. Hoddinott

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. 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.

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

[19] 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.

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

[21] 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.

Table 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)

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 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%).

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 report 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 owing 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 "out" in our 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) choose "out." Casique (2001) reports similar modes of bargaining in poor, rural families.[22]

4.3 Estimating Bargaining Power Levels for Each Household

We first fit a probit model on the decision to work, then, using a Heckman selection

[22]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."

Table 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)

approach, we incorporate the differences between those who work and those who do not in our earnings regression. The probit is given in equation (14) and the earnings regression is given in equation (15). We estimate these using a Full Information Maximum Likelihood approach.

$$L_{i,t} = X^p \beta^p + \epsilon_{i,t} \tag{14}$$

$$ln(E_{i,t}) = X^e \beta^e + \epsilon_{i,t} \tag{15}$$

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. The units for the left hand side of equation (15) 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 (14) 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 statistics, the household poverty index calculated by the Government of Mexico, interaction terms and village-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.[23] Except for the fixed effects, we display all variables and estimates in Tables 4 and 5. 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 5.

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

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

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 their 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 of bargaining dynamics per household per period. We present the distributions of bargaining power over time in Figure 1. We plot the control and treatment groups separately to emphasize the effect of treatment on the treated. Both groups start with medians close to 0.20. The interpretation of this is 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 the say that their male partner has. 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 8 in Appendix C, we present summary statistics on decision making patterns in the household, private assignable good expenditures, and views on women's rights. These additional summary statistics provide a more complete backdrop for understanding women's bargaining power in this setting. Particularly telling are the facts that one third of women need permission from their partners to leave the home in 1997, and that number increases to 45% later. Further, couples tend to make decisions together, but views on women's rights typically favor men. For instance, 77.5% of respondents in 1997 agreed with the statement "Women should be obedient to men."[24] These summary stats suggest that men may have substantially more bargaining power than women, giving us more confidence in the distributions in Figure 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) by testing

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

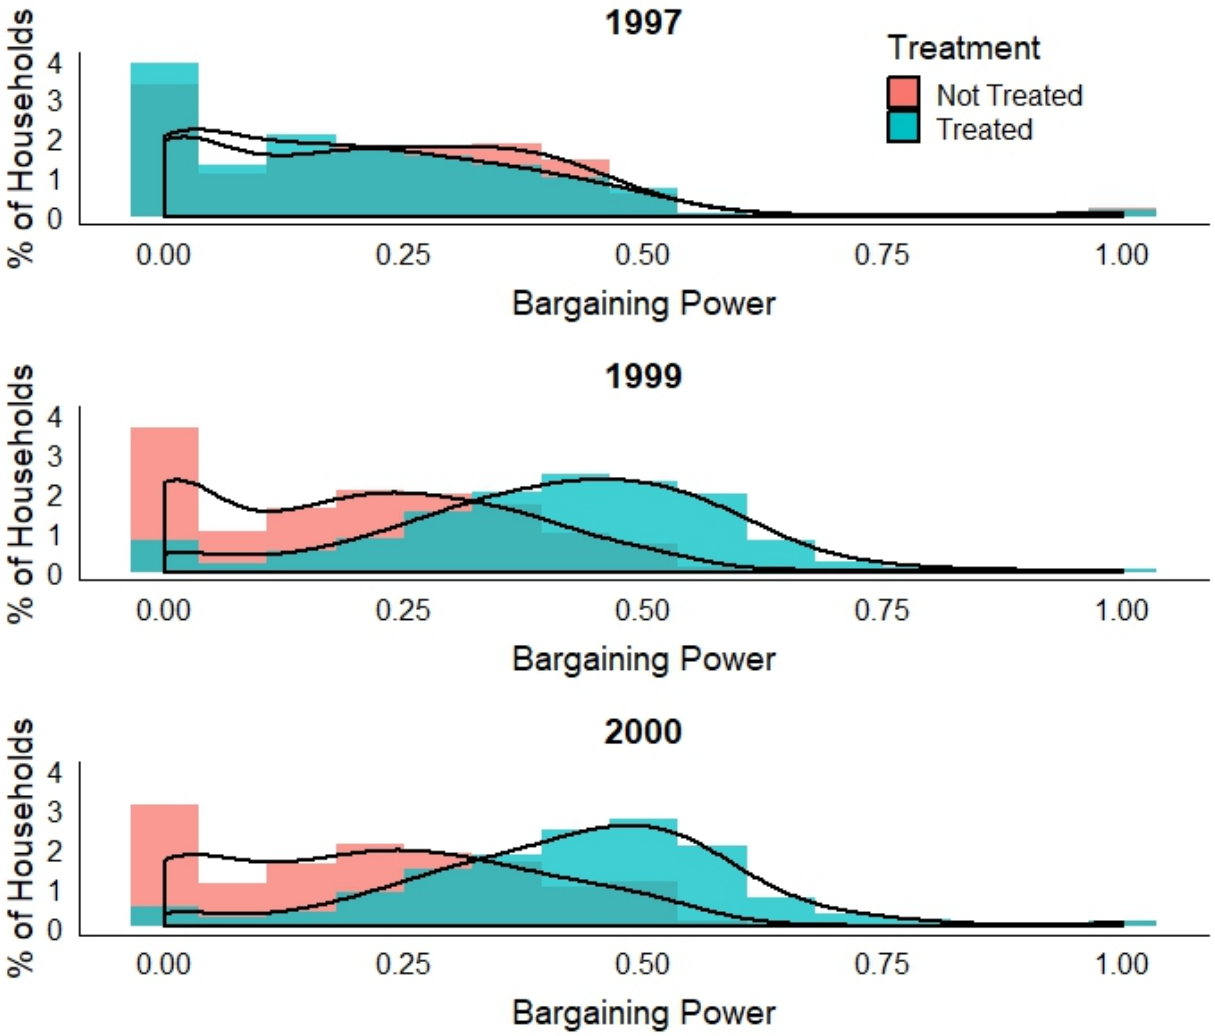

Figure 1: Distribution of bargaining dynamics, $\hat{\eta}$, over time across all households.

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 sure that the parallel trends assumption required for identifying the causal marginal effect holds. 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 by block bootstrapping and constructing a percentile-t confidence interval. If the confidence interval does not cover zero, we can conclude that the program had a significant impact on bargaining power.

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.[25] 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 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 (16) and (17):

$$Y^b_{h,t} = \beta^b_1 \hat{\eta}_{h,t} + \beta^b_2 X^b_{h,t} + \beta^b_3 \xi_h + \beta^b_4 \zeta_t + \epsilon_{h,t} \tag{16}$$

$$Y^c_{h,t} = \beta_1 \hat{\eta}_{h,t} + \beta_2 X_{h,t} + \beta_3 \xi_h + \beta_4 \zeta_t + \epsilon_{h,t} \tag{17}$$

The dependent variable in equation (16) is a dummy variable for whether or not a family

[25] 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.

consumed a food item in the past week. The dependent variable in equation (17) is the number of days per week that a family consumed the item.[26] 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, number of household kids ages 5-16, and number of children under the age of 5, own price, substitute prices, complement prices,[27] and staple prices.[28] 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]The superscript b denotes that equation (16) uses a binary dependent variable and disambiguates between the extensive and intensive margin investigations. We use a superscript c in equation (17) to disambiguate between total household income Y and the number of days per week that a food is consumed.

[27]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.

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

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 (16) 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 5 and 6. 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. 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.

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 4 and 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,

[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.

Table 4: 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)

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

Table 5: 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
σ	0.786*** (0.052)	0.555*** (0.003)
Note:		*p<0.1; **p<0.05; ***p<0.01

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.

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

We present demand regression results for chicken, eggs, leafy greens, and bananas in Table 6. The income coefficient estimates are generally positive and significant. The same price coefficients (not displayed) estimates are negative. They are not always significant, like 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. However, the standard errors for bargaining

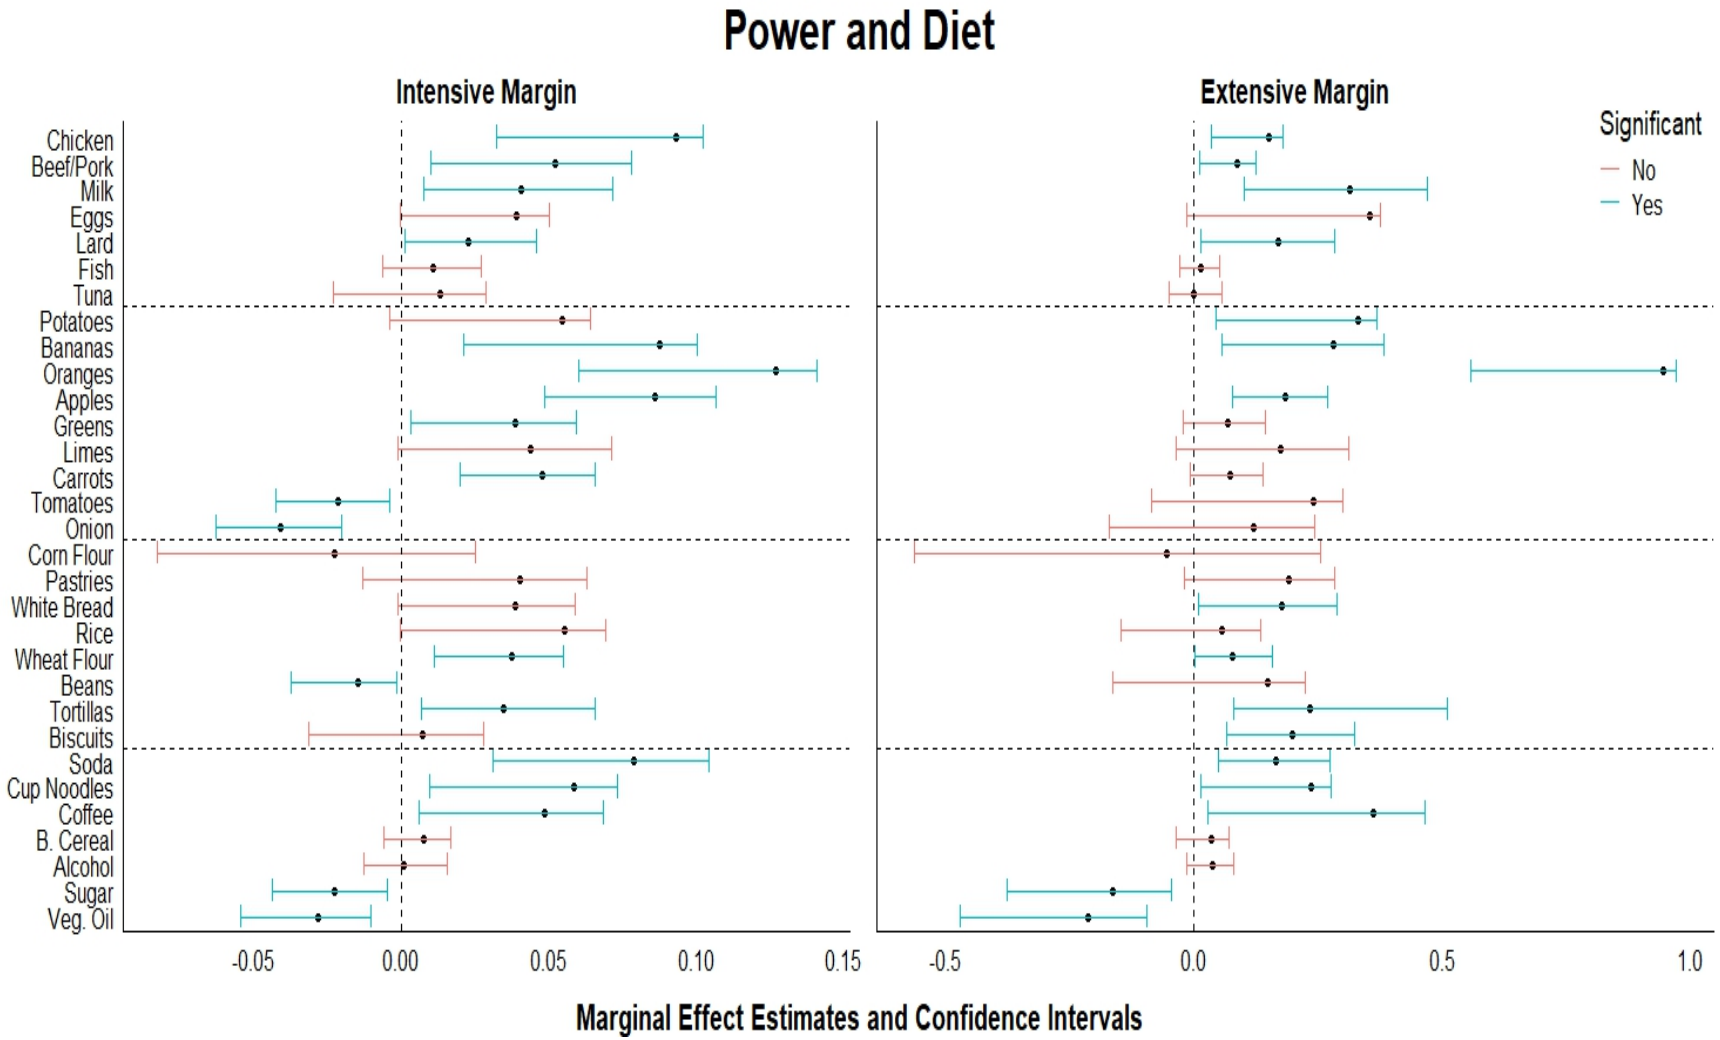

Figure 2: A Marginal Change in Bargaining Power Increases Net Family Consumption of Healthy Foods

power in Table 6 are incorrect - the percentile intervals provide correct measures of accuracy.

We generate marginal effect estimates from these thirty-one demand regressions and report them in Figure 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, 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 7. 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 $\hat{\delta} \times M$ 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.

6 Discussion

Two more comments are worth making here. First, it is possible that the model outlined in this section explains the unusual results documented in Hoehn-Velasco and Penglase

Table 7: 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.

(2019). They find that Mexico's introduction of no-fault divorce (from previously held stricter divorce legislation) reduced women's resource shares. This is puzzling since, as Voena (2015) documents in the USA, this type of legislation is typically thought to improve women's bargaining positions. 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 out 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 3 should obtain from the DLP family-specific systems of Engel curves for private assignable goods. This is an interesting avenue for future research, since estimating Engel curves at the same time as (5) (or (11) 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 - a very regrettable empirical reality (Bargain, Lacroix, and Tiberti, 2018). It could be that adding in limited commitment to the DLP model solves this problem. As a final suggestion, it might be useful to predict each family's private assignable good purchases to overcome the practice challenge of censored data. This would more closely align the DLP model with the one developed here.

7 Conclussion

We derive two new methods to estimate household bargaining power. Identification and estimation rely on predicting the earnings individuals would have in their relevant outside options - in our application our prediction approach is to use fitted values from Heckman selection models. We build on a rich literature connecting the outside option to bargaining power to derive the functional form for power (Nash, 1953; Becker, 1981; Mazzocco, 2007).

The primary benefit is that researchers can achieve point identification using assignable income data while allowing utility functions to take any form. In addition, the estimation method we employ explicitly accounts for censoring, which is a practical challenge to recovering resource shares from collective models of demand.

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, define the value of the outside option as the value of resources at each partners' disposal in that outside option. Third, predict each partners' earnings from the labor market, rent, transfers from friends and family, government programs, and other income generating opportunities. Researchers can use a Heckman selection approach to formally model unobservable characteristics that influence individuals' earnings ability. Fourth, depending on which estimator is appropriate, construct the measure of bargaining power using the functional forms we provide in equations (5) and (11).

We apply this measurement strategy to shed light on a twenty year old policy question of great importance. For the treatment group, we show that the median female decision maker's bargaining power before Progresa was 0.203, and that it rose to 0.406 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 substantially at the intensive and extensive margins.

In light of the Government of Mexico cancelling Prospera (the most recent iteration of Progresa) recently, 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 suggest that other gender targeted programs likely had similar empowerment effects, which may have subsequently increased investments in children.

We see this bargaining power measure as a complement to existing measurement methodologies 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 disadvantages of using our method are

that researchers must take a stand on what the outside option is, and how it affects power. None of the methods detailed in the essays listed above require the assertion that household decisions are subject to limited commitment.

We see our contribution as fitting into a larger discussion about how to measure women's empowerment in practice, a topic of perennial debate. Decision making within the household is one area where power matters, and so measuring power in this context sheds some light on women's empowerment in general. In 2017, the Abdul Latif Jameel Poverty Action Lab (J-PAL) featured a series of interviews on the topic, from prominent researchers and practitioners. Rachel Glennerster and Claire Walsh discussed practical challenges with using typical decision-making modules in impact evaluations, like the type used by Adato et al (2000). Sarah Baird and Danielle Moore discussed mixing qualitative and quantitative methods to generate more detailed, context-specific understandings of power relations. In 2018, J-PAL researchers Rachel Glennerster, Claire Walsh, and Lucia Diaz-Martin published "A Practical Guide to Measuring Women's and Girls' Empowerment in Impact Evaluations." This guide focuses on survey and non-survey instruments, but does not cover structural model of demand or outside options, or the bounded methods developed by Cherchye et al (2015, 2017). The nine non-survey methods they discuss include games, observing participation in community meetings, using vignette's, and learning from partners' biomarkers. Our method and these other structural methods provide additional measurement options, each with their own advantages and disadvantages.

Bibliography

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.

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.

Becker, G. S. (1981). A Treatise on the Family. National Bureau of Economic Research, Inc.

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.

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

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., González-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). Condiftional 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: Journal of the Econometric Society, 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.

De Barbieri, T. (1984) Mujeres y Vida Coitdiana, 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. (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.

Glennerster, R., Walsh, C., & Diaz-Martin, L. (2018) A Practical Guide to Measuring Women's and Girl's Empowerment in Impact Evaluations. A Abdul Latif Jameel Poverty Action Lab Guide. Retrieved from https://www.povertyactionlab.org/practical-guide-measuring-

women-and-girls-empowerment-impact-evaluations on 07.24.2018.

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, http://www.sciencedirect.com/science/article/pii/S0305750X08003057

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. (1977). Sample Selection Bias as a Specification Error (With an Application to the Estimation of Labor Supply Functions).

Hlavac, Marek (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 Côte 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

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

Lise, J., & Yamada, K. (2014). Household Sharing and Commitment: Evidence From Panel Data on Individual Expenditures and Time Use (No. W14/05). IFS Working Papers.

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). http://ebrary.ifpri.org/cdm/ref/collection/p15738coll2/id/128190

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.

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). Intra-Household Experiments: A Survey. Journal of Economic Surveys, 32(1), 134-175.

Nash, J. (1953). Two-person cooperative games. Econometrica: Journal of the Econometric Society, 128-140.

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.(2017) Individual Consumption in Collective Households: Identification Using Panel Data with an Application to PROGRESA, Unpublished Mimeo.

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.

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.

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 3.2 and 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 model in section 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 utility functions, C.

Another way to think of Assumption 1 is to index all of the elements in C by the unaccountably infinite real numbers on the unit interval. Then, each person has a utility function that corresponds to their preferences, and an index number that corresponds to that function. Assumption 1 is analogous to assuming that this index is uniformly distributed on the unit interval.

Under Assumption 1, 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 $[\eta]$. 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). 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 to predict outside option values for.

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 both Assumption 1 and the weaker Assumption 1A, the estimator of η is given in equation (5).

Assumption 2: 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 estimation and identification strategy, and under Assumptions 1 and 2, we can write the system of equations to be estimated as

$$\hat{\eta}_h = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_{f,h}\psi_{f,h}) - F(X_{m,h}, \psi_{m,h)}}{y_h}\right) \quad \forall h \in 1, \ldots, H \tag{18}$$

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

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

where Γ is a control function and $\nu_{i,j}$ 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 1 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 1*: Let Assumptions 1A and 2 hold, and consider interior solutions to (3). Then the sharing rule is point identified on (18).*

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 1A, $H > K$. As such, the number of parameters to be estimated (18) is $2K + H$ and the number of equations is $3H$. Therefore, the model is identified.*

Now consider the more challenging identification problem presented by the model in equation (7). Now the model is identified if there are at least two waves of data to estimate

a fixed effects model.

Assumption 3: 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:

$$\hat{\eta}_{h,t} = \frac{1}{2} + \frac{1}{2}\left(\frac{F(X_{f,h,t}\psi_{f,h,t}) - F(X_{m,h,t},\psi_{m,h,t})}{\tilde{y}_{h,t} - \tilde{\beta}_{0,h}}\right) \quad \forall h \in \{1,\ldots,H\} \text{ and } \forall t \in \{1,\ldots,T\} \tag{19}$$

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

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

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

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

Proposition 2*: Let Assumptions 1A, 2, and 3 hold and consider interior solutions to (7). Then the sharing rule is point identified on (19).*

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 (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 $\tilde{\epsilon}_h$. As such, the number of parameters to be estimated in (19) is $T(2K+H)+1$ and the number of equations is $3HT$. By Assumption 3, $T \geq 2$, and so the system in (19) identifies a unique value of eta for each household and each period.*

Abstractly, 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 sharing rule η_h.

Appendix B: Counterfactual Distributions

We can use our structural model to study the ramifications of alternative gender-targeting transfers. We can use our predicted outside option values and different Progresa transfer schemes to understand how the bargainig 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 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 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 3.

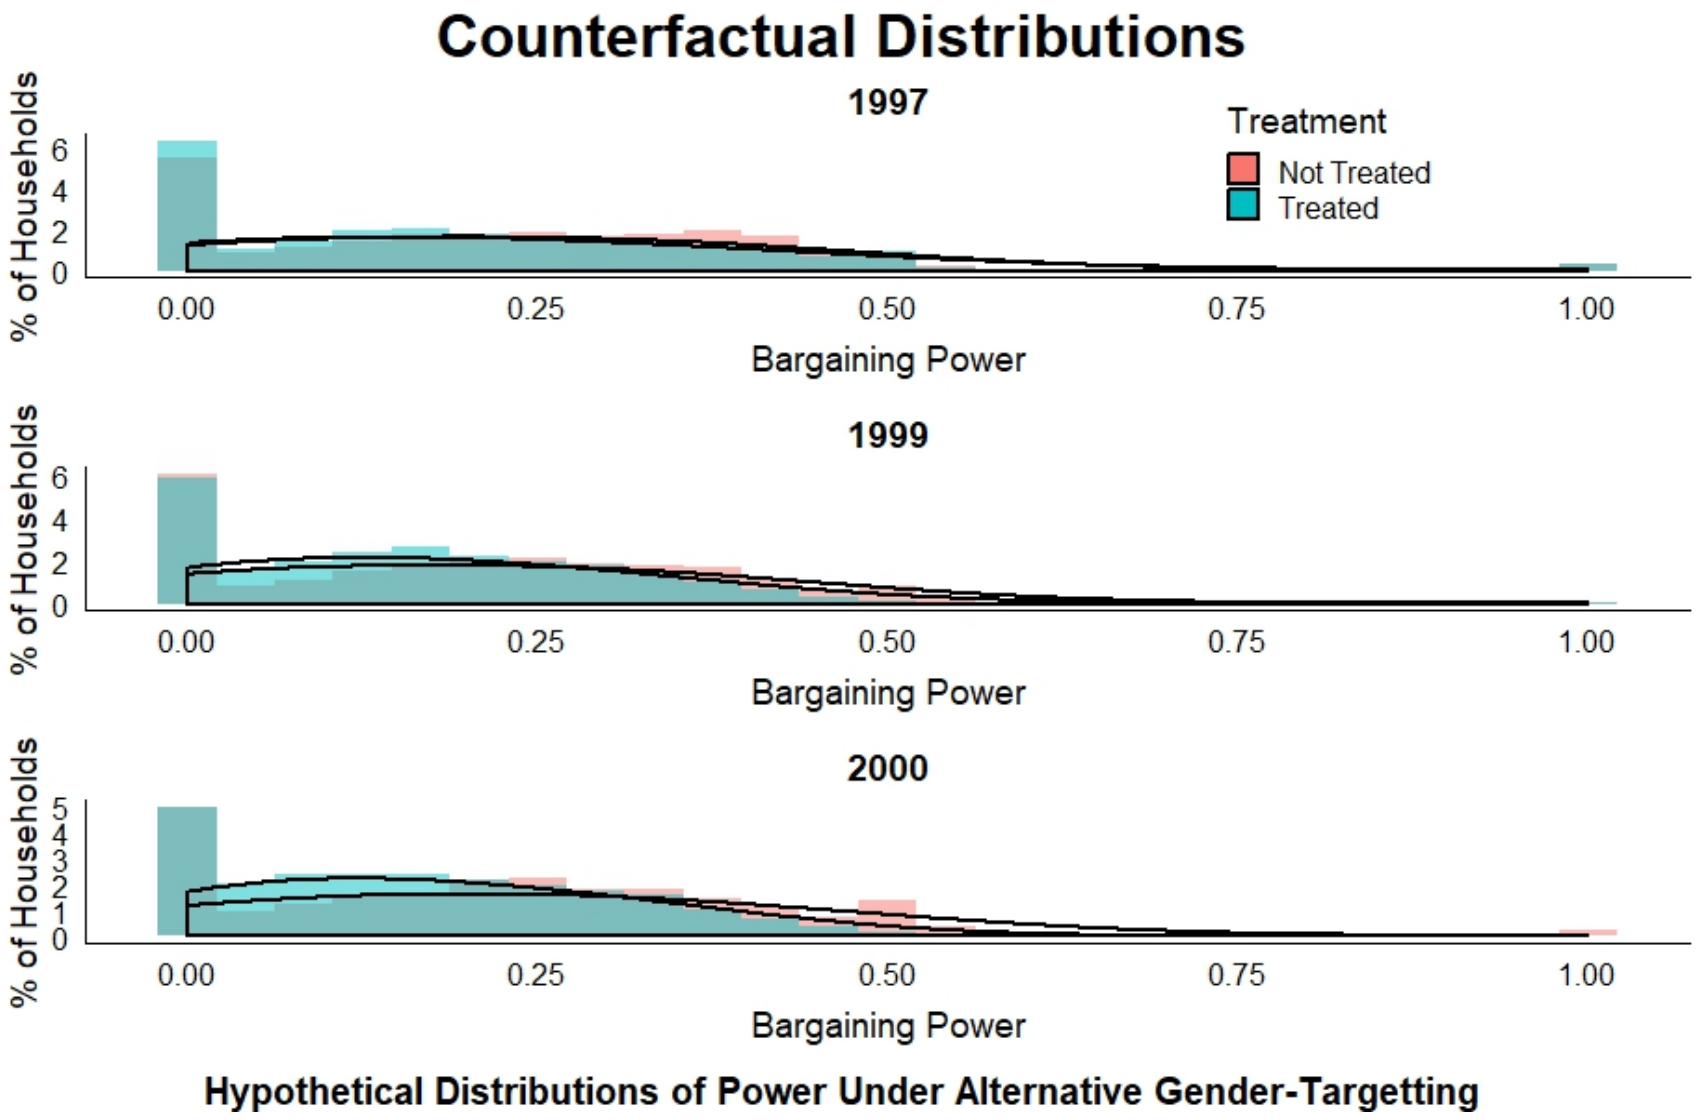

Figure 3: Unobserved Counterfactual Distributions Where Men Get Transfer Instead of Women

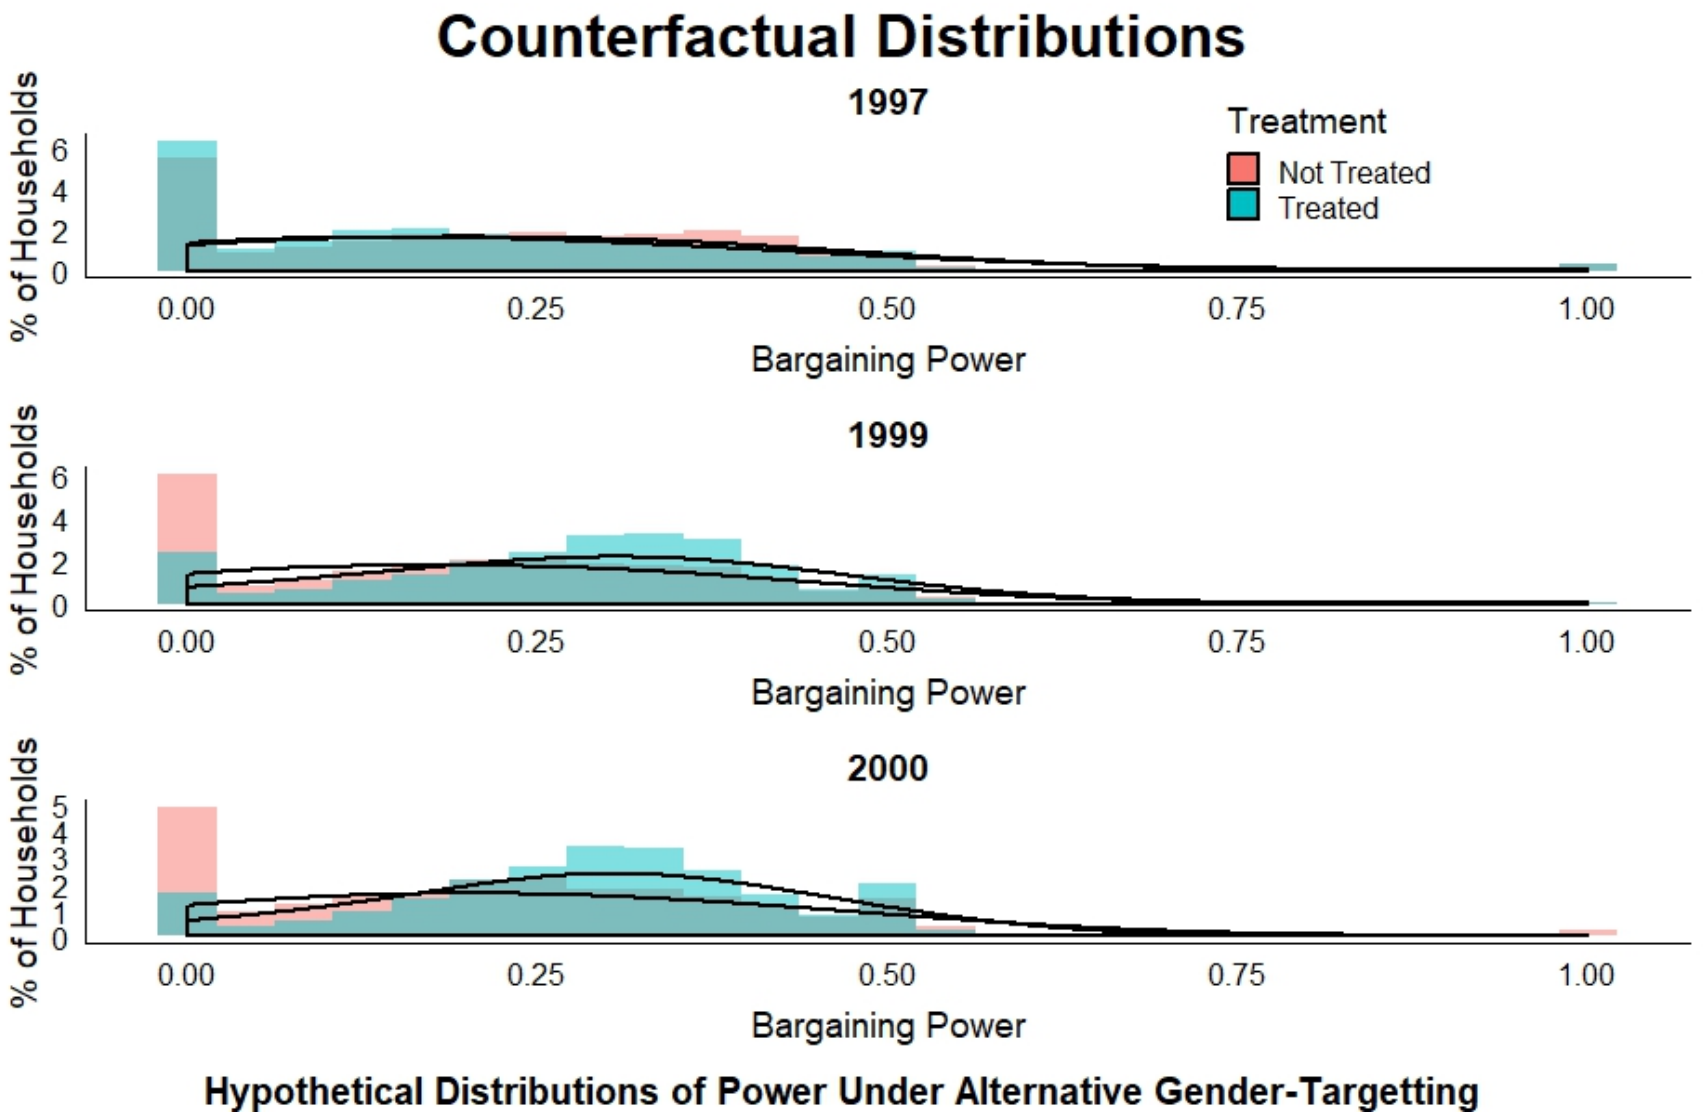

Figure 4: Unobserved Counterfactual Distributions Where Men and Women Split Transfer Equally

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 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 3 are reasonable.

In Table 8, 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 8 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 8 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 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.

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

The last portion of Table 8 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. 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.

Table 8: 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?"