An Economic Analysis of Intra-Household Bargaining Power

Jean-Paul Chavas*, and Matthew J. Klein*

April 17, 2020

Abstract: This paper investigates the economics of the intra-household bargaining process with a focus on the sharing rule and the determinants of power within the household. We apply Harsanyi's (1986) generic games framework to a limited-commitment collective model where a family's real consumption opportunities depend on each partners' specified outside option. We demonstrate uniqueness and existence of the solution by applying Rubinstein, Safra, and Thomson's (1992) appeals-immunity solution concept. We demonstrate that the expected value of bargaining power is semi-parametrically identified in an ordinal framework; the researcher need not assume specific utility or social welfare functional forms. In addition to empirical tractability and analytical precision, this update to the collective model provides a more nuanced view of intra-household bargaining power: partner's with a greater capacity to specify more damaging threats will have more control over the decision-making process, regardless of whether they act on those threats. (JEL: D1, D6, D7)
Keywords: Sharing Rule, Outside Options, Endogenous Threats, Generic Game, Bargaining Power

Acknowledgements: The authors would like to thank Bradford Barham, Paul Dower, and Jared Hutchins for their helpful comments. This paper was previously titled "Estimates of Household Bargaining Power when Outside Options are Endogenous."

*University of Wisconsin - Madison, Department of Agricultural and Applied Economics.

1 Introduction

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

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

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

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

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

Lewbel and Pendalur (2019) incorporate violence into the collective model developed by Browning, Chiappori, and Lewbel (2013) and use the identification strategy developed in Dunbar, Lewbel, and Pendakur (2013). Violence shifts the family from an efficient consumption technology to an inferior one (imposes higher shadow prices for the family). However, they are unable to derive a demand function for violence, which limits the predictive capacity of their model. Furthermore, they only demonstrate identification when threats take a binary value (either the husband abused his wife or he didn't). They are unable to study how more damaging threats might impact power dynamics, or explain why certain levels of violence obtain in equilibrium. They are unable to assess what effect the possibility of women filing for divorce might have on power dynamics.

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

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

frontier inward). Regardless of implementation, more damaging threats grant the partner who made them more control over the decision making process (moving the collective outcome along the Pareto frontier).

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

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

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

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

The partners engage in a "generic game" where they first jointly determine the rules of engagement and then play according to these rules (See Harsanyi, 1986). In order to demonstrate that a unique solution exists to this application of generic games to the household, we use the "appeals-immunity" solution concept introduced by Rubinstein, Safra, and Thomson (1992). As such it is useful to define a generic game in notation that is similar to theirs. Let a *bargaining problem* by a triple $< \theta, \succsim_f, \succsim_m > \in (\Sigma_f \times \Sigma_m) \times \rho^2$, where θ is a bargaining outcome defined over the set of possible outcomes $\Sigma_f \times \Sigma_m$; and $\succsim_f$ and $\succsim_m$ are partners' preferences over states of the world induced by their bargaining outcomes, which are elements of the set of permissible preferences, ρ. The permissible actions are threats of non-cooperation that each partner specifies but does not necessarily need to act on. They include physical or emotional violence, divorce, and milder punishments like neglecting agreed-upon duties or chores.

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

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

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

The nested bargaining game solved in the third stage, then, is a bijection mapping from the partners' choices in the first two stages to the optimal division of resources between the two partners, $\eta^* \in [0, 1]$; the n-vector of household demand functions, $z^*(p, y, \eta^*) \in \mathbb{R}^n$ where p and y are prices and family income; the set of partners' n-vectors of consumption amounts, $\{x_f^*(p, \eta^* y), x_m^*(p, (1 - \eta^*)y\} \in \mathbb{R}^{2n}$; and labor supply functions for each partner, $l_f^*, l_m^* \in [0,24]^2$. This bijection is $B(\theta)$: $\Sigma_f \times \Sigma_m \times \{0,1\}^2 \rightarrow [0,1] \times \mathbb{R}_+^{3n} \times [0,24]^2$. That is, for any given set of solutions to the first two stages, $\{\theta_f, \theta_m, I_f, I_m\}$, the unique corresponding nested game $B(\theta_f, \theta_m, I_f, I_m)$ is induced which, in turn, has the unique solution $\{\eta^*, z^*(p,y,\eta^*), \{x_f^*(p,\eta^* y), x_m^*(p,(1-\eta^*)y\}, l_f^*, l_m^* \Big| \theta\}$.

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

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

The Nested Two-Stage Collective Model

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

The central assumption in the nested model is that the partners reach a (conditionally) Pareto efficient outcome (as in Lewbel and Pendakur, 2019). The specified threats and whether they are carried out determines how far off of the best possible Pareto frontier a family's outcome will be. If neither partner acts on their threat, then the family is on a frontier that is strictly above the frontier that obtains when one or both partners carry out their threat. This is because acting on threats reduces the quality of the family's production and consumption technologies. Let $F(l_f, l_m|\theta)$

be the family's production function so that the total family income is $y = F(l_f, l_m|\theta)$, and let $A(x_f, x_m|\theta)$ be the family's consumption technology so that household demand functions are given by $z = A(x_f, x_m|\theta)$. This consumption technology relates the household purchases to the amount of each good that family members consume, accounting for externalities that accrue when the family consumes locally public goods, and for economies of scope within the household.[2] Implementing threats results in a lower Pareto frontier for the family by reducing the quality of F and A. Families are conditionally efficient in the sense that they consume along the frontier that their family faces, and the frontier is endogenous in that it results from the partners' choices of threats and implementation.

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

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

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

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

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

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

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

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

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

$$\frac{\partial \tilde{U}}{\partial V_f}\frac{\partial V_f}{\partial \eta} = -\frac{\partial \tilde{U}}{\partial V_m}\frac{\partial V_m}{\partial \eta} \tag{2}$$

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

Limited Commitment

If partners do not act on their threat, the family solves bargaining problem $\mathbb{B}(\theta_f, \theta_m, I_f = 0, I_m = 0)$, and each partner gets corresponding indirect utility from the efficient production and consumption technologies, and the resulting sharing rule, $V_f(p, \eta^* y|\theta_f, \theta_m, I_f = 0, I_m = 0)$ and $V_m(p, (1-\eta^*)y|\theta_f, \theta_m, I_f = 0, I_M = 0)$. As such, each partner prefers to implement their threat ($I_f = 1 \succ_f I_f = 0$ and $I_m = 1 \succ_m I_m = 0$) if the corresponding utility function representations of their preferences yield higher values when they implement their threats. They solve the following optimization problems:

$$I_f^* = \underset{I_f \in \{0,1\}}{argmax}\big[V_f(p, \eta^* y|\theta_f, \theta_m, I_f = 1, I_m), V_f(p, \eta^* y|\theta_f, \theta_m, I_f = 0, I_m)\big] \tag{3}$$

$$I_m^* = \underset{I_m \in \{0,1\}}{argmax}\big[V_m(p, (1-\eta^*)y|\theta_f, \theta_m, I_f, I_M = 1), V_m(p, (1-\eta^*)y|\theta_f, \theta_m, I_f, I_M = 0)\big]$$

The *collective allocation problem with limited commitment* is $\mathbb{B}(\theta_f, \theta_m, I_f^*, I_m^*)$. The solution to this problem is, conditional on the specified threats, the indicators for whether partners act on their threats, the optimal demand and supply functions, and the optimal division of resources, conditional on the quality of the production function and consumption technology. It would be optimal to reduce the total surplus that the family has access to by acting on the specified threat if the resulting change in the sharing rule increases an individuals' total consumption. That is, the increase in their "slice of the pie" more than offsets the decrease in the "total size of the pie." To be explicit, the limited-commitment collective problem that the family solves is

$$\max_{\eta\in[0,1]} \quad \tilde{U}(V_f(p,\eta y|\theta_f,\theta_m,I_f^*,I_m^*),V_m(p,(1-\eta)y|\theta_f,\theta_m,I_f^*,I_m^*)) \text{ subject to} \tag{4}$$

$$A(x_f,x_m|\theta_f,\theta_m,I_f^*,I_m^*)=z \text{ (Consumption Technology Constraint)}$$

$$y=F(l_f,l_m|\theta_f,\theta_m,I_f^*,I_m^*) \text{ (Production Constraint)}$$

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

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

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

$$I_f^*=\underset{I_f\in\{0,1\}}{argmax}\big[V_f(p,\eta^*y|\theta_f,\theta_m,I_f=1,I_m),V_f(p,\eta^*y|\theta_f,\theta_m,I_f=0,I_m)\big]$$

$$I_m^*=\underset{I_m\in\{0,1\}}{argmax}\big[V_m(p,(1-\eta^*)y|\theta_f,\theta_m,I_f,I_M=1),V_m(p,(1-\eta^*)y|\theta_f,\theta_m,I_f,I_M=0)\big]$$

When deciding what threats to choose, partners consider what corresponding nested game, $B(\theta_f,\theta_m,I_f^*,I_m^*)$, will result. They face trade offs: a more damaging threat grants more control over the decision making process but, if implemented, reduces total family surplus to a larger extent.

The First Stage of the Generic Game: Choosing Threats

While Harsanyi is relatively vague about how these threats damage players, we can specify how these damages occur more clearly in the family context.

Definition 1: Let θ_i, $\bar{\theta}_i \in (\Sigma_f\times\Sigma_m)\times\{0,1\}^2$ be two potential threats for partner $i\in\{f,m\}$. The threat θ_i is more damaging *than $\bar{\theta}_i$ if one or both of the following is true:*

1. *for $i\neq j\in\{f,m\}$, $F(l_f,l_m|\theta_i,\theta_j,I_i=1,I_j)<F(l_f,l_m|\bar{\theta}_i,\theta_j,I_i=1,I_j)$ (The* more damaging *threat, if carried out, lowers family income.)*
2. *for $i\neq j\in\{f,m\}$, $L(A(\theta_i,\theta_j,I_i=1,I_j))>L(A(\bar{\theta}_i,\theta_j,I_i=1,I_j))$ (The* more damaging *threat, if carried out, increases the shadow prices each person faces in their optimization problems.)*

Whenever someone specifies a more damaging threat, it must be because they gain additional control over the decision making process by doing so. So for any θ_f, $\bar{\theta}_f$ $\in\Sigma_f\times\{0,1\}$, and for any $\theta_m\in\Sigma_m\times[0,1]$, where θ_f is *more damaging* than $\bar{\theta}_f$,

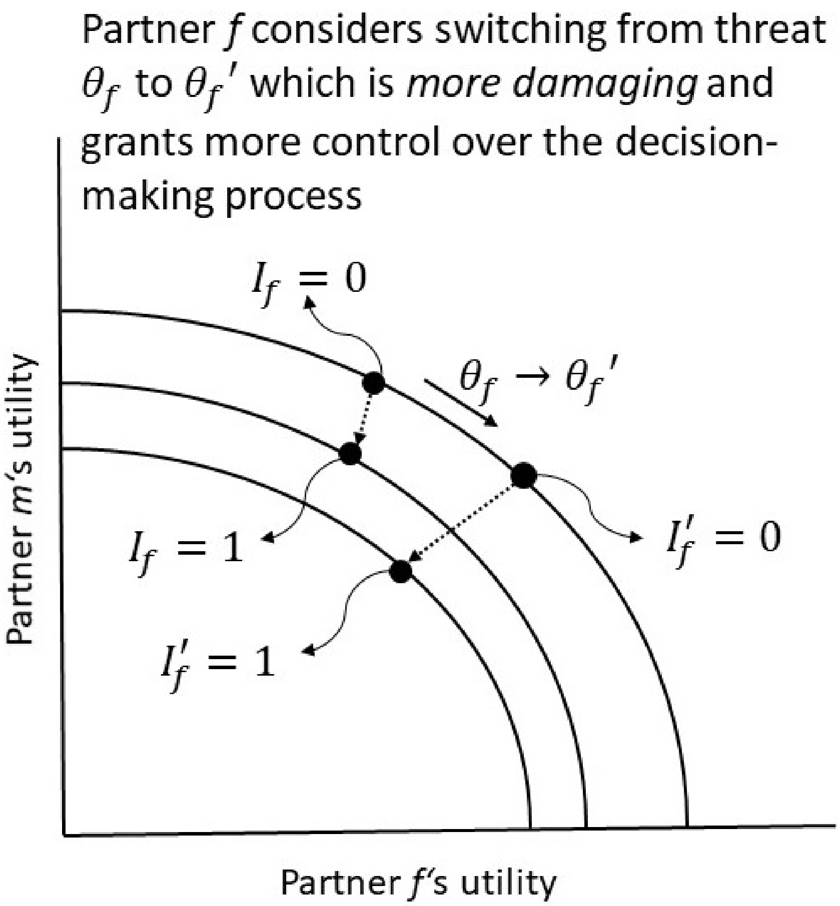

Figure 1: Partner f considers two different threats. The *more damaging* threat grants more bargaining power to partner f, shifting them along the fully efficient Pareto frontier if not implemented. If implemented, the *more damaging* threat moves the family to a lower Pareto frontier than the less damaging threat would, if it were implemented.

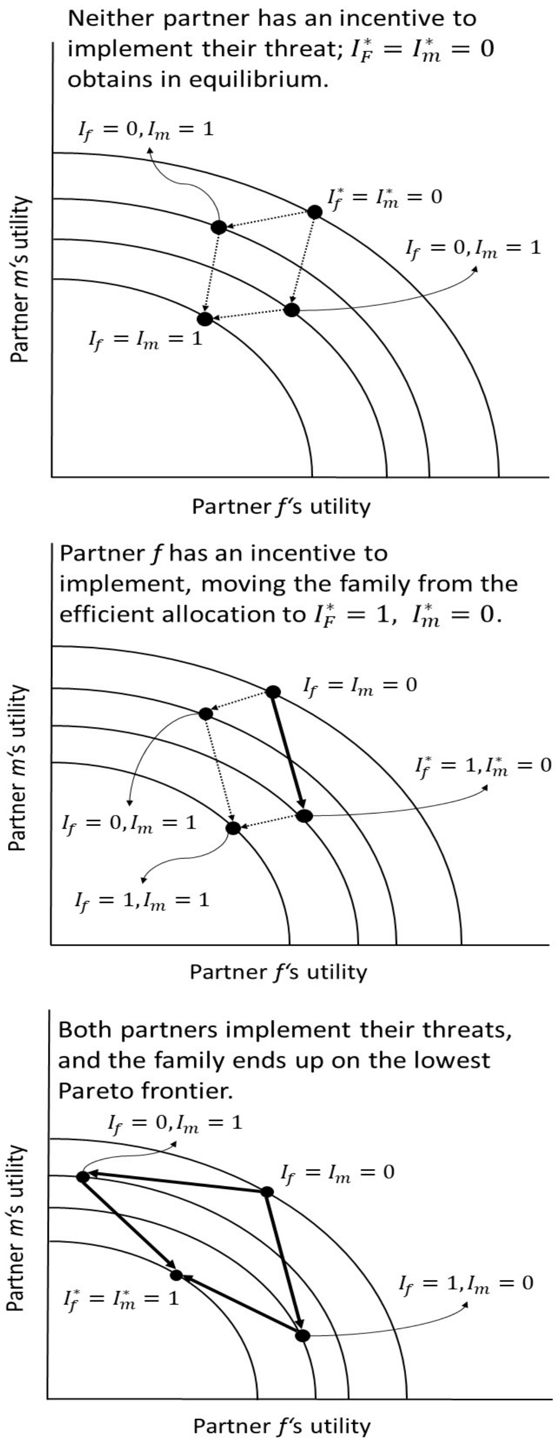

Figure 2: In the top panel, the threats that the partners have specified results in the fully-efficient equilibrium. In the middle panel, one of the partners has an incentive to implement their threat, and the family ends up on an intermediate Pareto frontier. In the bottom panel, both partners have an incentive to implement their threat and end up on the lowest Pareto frontier for their family.

$\eta(\theta_f, \theta_m, I_f, I_m) > \eta(\bar{\theta}_f, \theta_m, I_f, I_m)$. The symmetric statement is true for partner m. If more damaging threats did not confer more control over the decision making process, the optimal threat would be the least damaging one and it would never be implemented. See Figures 3.1 and 3.2 for graphical depictions of the choices in stages one (the threats) and two (implementation).

Partners choose their threats from the space of possible threats to maximize their payouts in the resulting *collective allocation problem with limited commitment.* We can set up this constrained maximization problem using the cardinal utility functions, and subsequently on ordinal preferences using bijections between threats and allocations. We first describe the cardinal problem, then formally define these bijections, then set up the corresponding ordinal problems, and finish with the definition of *individual-rationality.* The partners solve:

$$\theta_f^* = \underset{\theta_f \in \Sigma_f}{argmax} \quad V_f(L(A(\theta_f, \theta_m, I_f^*, I_m^*), \eta^*(\theta_f, \theta_m) F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)) \tag{5}$$

$$\theta_m^* = \underset{\theta_m \in \Sigma_m}{argmax} \quad V_m(L(A(\theta_f, \theta_m, I_f^*, I_m^*), (1 - \eta^*(\theta_f, \theta_m)) F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*))$$

The choice of a threat influences utility in three ways. A more punishing threat moves the resource share in a person's favor, and more punishing threats damage the consumption and production technologies more. These optimization problems have the following first order conditions (suppressing notation):

$$\begin{aligned} -\frac{\partial V_f}{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_f} &= \\ \frac{\partial V_f}{\partial \eta^*(\theta_f, \theta_m)} \frac{\partial \eta^*(\theta_f, \theta_m)}{\partial \theta_f} &\times \\ \frac{\partial V_f}{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_f} \end{aligned} \tag{6}$$

$$\begin{aligned} -\frac{\partial V_m}{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial L(A(\theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_m} &= \\ \frac{\partial V_m}{\partial \eta^*(\theta_f, \theta_m)} \frac{\partial \eta^*(\theta_f, \theta_m)}{\partial \theta_m} &\times \\ \frac{\partial V_m}{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)} \frac{\partial F(l_f^*, l_m^* | \theta_f, \theta_m, I_f^*, I_m^*)}{\partial \theta_m} \end{aligned} \tag{7}$$

These first order condition in (3.6) equate the marginal benefit that accrues to partner f from having a larger resource share to the marginal disutility that accrues because the Pareto frontier for the family shifts inwards when the threats are acted on. Equation (3.7) gives the corresponding relationship for partner m.

There may be multiple pairs (θ_f^*, θ_m^*) such that (3.6) and (3.7) hold simultaneously. The threats in any such pair are *individually* $-$ *rational* in $< \theta, \succsim_f^{sm}, \succsim_m^{sm} >$. Note that any pair (θ_f^*, θ_m^*) that satisfies (5) and (6) implies that the couple solves a distinct *collective allocation problem with limited commitment* $\mathbb{B}(\theta_f^*, \theta_m^*, I_f^*, I_m^*)$. Each nested game has a unique solution, including specific private consumption equivalents and labor allocations for each partner. As such, each pair is associated with a specific level of indirect utility for each partner. The differential operators relating the pair (θ_f^*, θ_m^*) to consumption and labor outcomes are:

$$\Delta_f(\theta_f^*, \theta_m^*, \succsim_f, \succsim_m, \tilde{U}) =$$
$$\left\{(x_f^*(L(A), \eta^* y), l_f^*((L(A), \eta^* y)) \in \mathbb{R}_+^n \times [0, 24] \middle| \text{conditions (3.2) and (3.6) hold}\right\}$$

$$\Delta_m(\theta_f^*, \theta_m^*, \succsim_f, \succsim_m, \tilde{U}) =$$
$$\left\{(x_m^*(L(A), (1-\eta^*) y), l_m^*(L(A), (1-\eta^*) y)) \in \mathbb{R}_+^n \times [0, 24] \middle| \text{conditions (3.2) and (3.7) hold}\right\}$$

As such, we can define preferences directly over threats for each player. By defining preferences over threats, we can define the concept of *individual-rationality* of threats without referencing the cardinal utility function representation of underlying preferences, and without needing to specify a functional form for the family's social welfare function.[3]

Definition 2: Let $<\theta, \succsim_f^{sm}, \succsim_m^{sm}> \in (\Sigma_f \times \Sigma_m) \times (\rho^{sm})^2$ *be a* bargaining problem. *Let* $\theta^* \in (\Sigma_f \times \Sigma_m) \times [0,1]^2$ *be a potential* solution. *The potential* solution θ *is* individually rational *if both of the following are true:*

1. $\forall \theta_m \in \Sigma_m$, and $\theta_f^* \in \Sigma_f$, $\forall \bar{\theta}_f \in \Sigma_f \setminus \{\theta_f^*\}$, $\theta_f^* \succsim_f \bar{\theta}_f \iff \Delta_f(\theta_f^*, \theta_m, \succsim_f, \succsim_m, \tilde{U}) \succsim_f \Delta_f(\bar{\theta}_f, \theta_m, \succsim_f, \succsim_m, \tilde{U})$ and

2. $\forall \theta_f \in \Sigma_f$, and $\theta_m^* \in \Sigma_m$, $\forall \bar{\theta}_m \in \Sigma_m \setminus \{\theta_m^*\}$, $\theta_m^* \succsim_m \bar{\theta}_m \iff \Delta_m(\theta_f, \theta_m^*, \succsim_f, \succsim_m, \tilde{U}) \succsim_m \Delta_m(\theta_f, \bar{\theta}_m, \succsim_f, \succsim_m, \tilde{U})$.

This definition states that people prefer a threat if they prefer the corresponding outcome from the specific resulting nested game. Let $(\Sigma_f^* \times \Sigma_m^*) \subset (\Sigma_f \times \Sigma_m)$ be the set of all potential *solutions* that are *individually-rational* for the *bargaining game* $< \theta, \succsim_f^{sm}, \succsim_m^{sm} > \in (\Sigma_f \times \Sigma_m) \times (\rho^{sm})^2$.

[3] In general, children's utility would enter into the social welfare function (and potentially also into parents' preferences directly), extending the model to do so does not change our core results or suggested estimation strategy. It may expand the set of threats that people have at their disposal to include threats over custody.

2.1 Formally Defining Appeals Immunity

Consider induced utility functions, which map from the *solution* space to the real numbers, and describe a person's willingness to pay to transition from a reference *solution* to an alternative. Let $u(\succsim_i, \theta^r, \theta)$ be such a function for $i \in \{f, m\}$, where θ^r is a reference solution and θ is a specified alternative, $u : \rho \times \left((\Sigma_f^* \times \Sigma_m^*)^* \times \{0,1\}^2\right)^2 \rightarrow \mathbb{R}$. Because both outcomes, θ^r and θ, are associated with specific private equivalent consumption bundles for each person, and because people have preferences over these consumption bundles, partners indirectly have preference relations over θ^r and θ. People always prefer the outcome where their partner does not carry out the threat to the outcome where their partner chooses to carry out their threat. Note that $u(\succsim_i, \theta^r, \theta)$ is differentiable.[4]

The value $u(\succsim_i, \theta^r, \theta)$ is a compensating variation value. It makes that partner indifferent between two pairs of threats by changing the value of the consumption bundles the right amount. For partner f, this value solves: $x_f(p, \eta^{*r}y + u(\succsim_f, \theta^r, \theta)) \sim_f x_f(p, \eta^* y)$. Similarly, for partner m, this compensating variation amount is the one that makes them indifferent between the consumption bundle associated with the reference bargaining solution, and the alternative: $x_m(p, (1-\eta^{*r})y + u(\succsim_m, \theta^r, \theta)) \sim_m x_m(p, (1-\eta^*)y)$. When $\theta^r \succ_i \theta \; \forall i \in \{f, m\}$, $u(\succsim_i, \theta^r, \theta) < 0$, when $\theta \succ_i \theta^r$, $u(\succsim_i, \theta^r, \theta) > 0$, and when $\theta^r \sim_i \theta$, $u(\succsim_i, \theta^r, \theta) = 0$. People are willing to pay a positive amount to switch to a prefer ed solution, and have to be paid to be willing to switch from a preferred solution to one they consider to be worse.

Definition 3: Let $<\theta, \succsim_f, \succsim_m> \in (\Sigma_f \times \Sigma_m) \times \rho^2$, *be a bargaining problem. Let* θ, $\theta^r \in (\Sigma_f \times \Sigma_m) \times \{0,1\}^2$ *be two potential* solutions *to* $<\theta, \succsim_f, \succsim_m>$. *The reference outcome* θ^r *is* immune to the appeal of the alternative θ *in* $<\theta, \succsim_f, \succsim_m>$ *if one of the following is true:*

1. $\theta^{*r} \succ_f \theta^*$ and $x_f(p, \eta^{*r}y) \succ_f x_f(p, \eta^* y + \tau_m)$ where $\tau_m \in [0, u(\succsim_m, \theta^r, \theta)]$, or
2. $\theta^{*r} \succ_m \theta^*$ and $x_m(p, (1-\eta^{*r})y) \succ_m x_m(p, (1-\eta^*)y + \tau_f)$ where $\tau_f \in [0, u(\succsim_f, \theta^r, \theta)]$

This definition simply says that there is no transfer that one partner could give to the other that makes both happy to switch from the reference threats to an alternative potential solution. In the first condition, partner m wants to switch and can offer some transfer amount to partner f equal to or less than their compensating variation value, but this is not enough to induce partner f to switch. In the second

[4]Formally, for some $\alpha \in [0,1]$, $u(\succsim_i, \alpha\theta^r + (1-\alpha)\theta, \theta) = \alpha u(\succsim_i, \theta^r, \theta)$. Convex mixtures of the reference and alternatives result in averages over willingness to pay amounts in a smooth manner. When the reference threat changes by a small amount, the willingness to pay changes by a small amount as well.

condition, the opposite scenario plays out. Partner f wants to switch but partner m does not, and there is not a feasible transfer from f to m that makes m prefer to switch. In these cases, the reference outcome will remain the family's outcome despite the appeal to switch to an alternative.

Consider an example based on Ramos' (2016) work on violence and power in the household. In her study context, northern Ecuador in 2011, 35% of married men abuse their partners and only a small fraction of those relationships end in divorce. In this context, men's threat would be some form of physical, sexual, or emotional abuse (or some combination), and women's threat could be divorce or some mild form of non-cooperation. Ramos shows that this violence increases men's resource shares but reduces women's productive capacity, reducing the amount of surplus available to the family and allocating a larger amount of this surplus to men.[5] Take a reference bargaining solution where men choose to physically abuse their partners (15% of households in this context) and women choose divorce as their threat but do not act on it ($I_f = 0$, $I_m = 1$). Consider an alternative with the same specified threats, but where neither partner acts on them ($I_f = 0$, $I_m = 0$). This realistic reference *bargaining solution* is immune to the appeal of this alternative if women are unable to provide a transfer (monetary or in kind) to their partners that compensates men for their loss of power, conditional on the total surplus being greater in the alternative. In the setting where we observe violence, men are getting so much more of the family's surplus that it is not possible for women to appeal successfully.

2.2 Uniqueness of the Appeals-Immune, Individually-Rational Solution

We now turn to a discussion of existence and uniqueness of a *solution* to the bargaining problem $< \theta, \succsim_f, \succsim_m >$. We follow Hanany and Safra (2000) to show that, for preferences satisfying certain assumptions, there is only one possible reference solution, θ^*, such that, for any possible alternative θ, θ^* is immune to the appeal of the alternative θ in $< \theta, \succsim_f, \succsim_m >$.[6] The following assumptions on the threat spaces, and the preferences, must hold for existence and uniqueness.

Assumption 3.3.1 — Compactness: Consider $\theta^a, \theta^b \in (\Sigma_f \times \Sigma_m)$. *Consider the infi-*

[5] Our model is mildly more general than hers because it makes divorce endogenous, let's women have threats as well, and is defined over the ordinal set of preferences, as opposed to the cardinal space of utility function representations. An application of this model to her study context might be fairly similar to the model she builds and estimates. Allowing women to threaten divorce would likely result in a prediction of the sharing rule that favors men slightly less than what she finds.

[6] The following assumptions are correspond to assumptions DOM, Q, CCE, and H in Rubinstein, Safra, and Thomson (1992).

nite sequence of solutions $\{\theta^a, \alpha_1\theta^a+(1-\alpha_1)\theta^b, \alpha_2\theta^a+(1-\alpha_2)\theta^b, ..., \theta^b\}$ $\forall \alpha_1, \alpha_2, \ldots, \alpha_\infty \in [0,1]$ *such that* $\alpha_1 < \alpha_2 < \cdots < \alpha_\infty$. *Then* $\theta_a \in (\Sigma_f^* \times \Sigma_m^*) \implies \theta_b \in (\Sigma_f^* \times \Sigma_m^*)$. *That is,* $(\Sigma_f^* \times \Sigma_m^*)$ *contains all of its limit points.*

Assumption 3.3.2 — Convexity: $\forall i \in \{f, m\}$, $\forall \alpha \in [0,1]$ $\forall \theta^{*a}, \theta^{*b} \in (\Sigma_f^* \times \Sigma_m^*)$, $\theta^{*a} \succ_i \theta^{*b} \implies \alpha\theta^{*a} + (1-\alpha)\theta^{*b} \succ_i \theta^{*b}$.

Assumption 3.3.3 — Concavity of Induced Utility Functions: Consider $\theta^a, \theta^b \in (\Sigma_f \times \Sigma_m)$, $\theta^a \succ_i \theta^b$ *for* $i \in \{f, m\}$, $\alpha \in (0,1)$. *The induced utility function* $u(\succsim_i, \theta^a, \theta^b)$ *is* concave *if* $u(\succsim_i, \alpha\theta^a + (1-\alpha)\theta^b, \theta^b) > \alpha u(\succsim_i, \theta^a, \theta^b) + (1-\alpha)u(\succsim_i, \theta^b, \theta^b) = \alpha u(\succsim_i, \theta^a, \theta^b) + 0 = \alpha u(\succsim_i, \theta^a, \theta^b)$.

Assumption 3.3.4 — Loss Aversion: Preferences that have induced utilities that satisfy the following property are said to be loss-averse *preferences: for* $\theta^r \succ_i \theta$, $u(\succsim_i, \theta, \theta^r) < |u(\succsim_i, \theta^r, \theta)|$. *That is, person i needs more compensation to be willing to move from a good reference to a bad alternative than they would be willing to pay to move from that bad alternative to the good reference.*

Assumption 3.3.5 — Weakly Symmetric Loss Aversion (WSLA): Consider two possible solutions, $\theta, \theta^r \in (\Sigma_f \times \Sigma_m))$. *When partners in a marriage have preferences that jointly exhibit the following property, they are said to be* weakly symmetric in loss aversion:

1. $\succsim_i$ *is* loss averse *for* $i \in \{f, m\}$, *and*
2. *for* $\theta^r \succsim_f \theta$ *and* $\theta \succsim_m \theta^r$, $u(\succsim_f, \theta^r, \theta) \leq |u(\succsim_m, \theta^r, \theta)|$, *and*
3. *for* $\theta^r \succsim_m \theta$ *and* $\theta \succsim_f \theta^r$, $u(\succsim_m, \theta^r, \theta) \leq |u(\succsim_f, \theta^r, \theta)|$.

Let the set of preferences that satisfies assumptions 3.3.2 - 3.3.5 be denoted ρ^{WSLA}. Let the compact subset of *individually-rational solutions* be denoted $(\Sigma_f^* \times \Sigma_m^*)^C$. Formally, assumptions 3.3.4 and 3.3.5 play the same role as the bounds assumption in (E2) of Theorem 3.2 of Hanany and Safra (2000), and their risk aversion assumption. In the setting with lotteries, the risk aversion assumption is almost the exact analog of the loss aversion in the setting without uncertainty. The WSLA assumption bounds the induced utilities that result from appeals to move away from the appeals-immune *solution*, as we see in the proof of the following proposition:

Proposition - Uniqueness: Let $< \theta, \succsim_f, \succsim_m > \in (\Sigma_f^* \times \Sigma_m^*)^C \times (\rho^{sm} \cap \rho^{WSLA})^2$ *be a* bargaining problem. *Let* $\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*} \times \{0,1\}^2$ *be a potential* solution *to*

$< \theta, \succsim_f, \succsim_m >$. *For all possible alternatives,* $\theta \in ((\Sigma_f \times \Sigma_m) \times [0,1]^2) \setminus \{\theta^*\}$, θ^* *is immune to the appeal of* θ *in* $< \theta, \succsim_f, \succsim_m >$ *if:*

$$\theta^* = \underset{\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*}}{argmax} \prod_{i \in \{f,m\}} u(\succsim_i, \theta^*, \theta).$$

Proof. $\theta^* = \underset{\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*}}{argmax} \prod_{i \in \{f,m\}} u(\succsim_i, \theta^*, \theta)$ *is unique because* $u(\succsim_i, \theta^*, \theta) \in \mathbb{R}_+$ *and* $u(\succsim_i, \theta^*, \theta)$ *is concave. At this point,*

$$\theta^* = \underset{\theta^* \in (\Sigma_f^* \times \Sigma_m^*)^{C*}}{argmax} \prod_{i \in \{f,m\}} u(\succsim_i, \theta^*, \theta) \iff$$

$$\frac{\partial u(\succsim_f, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_f^*} = -\frac{\partial u(\succsim_m, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_f^*} \text{ and}$$

$$-\frac{\partial u(\succsim_f, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_m^*} = \frac{\partial u(\succsim_m, (\theta_f^*, \theta_m^*), (\theta_f, \theta_m))}{\partial \theta_m^*}$$

These first order conditions say that at the solution, $\theta^* = (\theta_f^*, \theta_m^*)$, specifying an alternative reference results in a marginal benefit (higher willingness to pay to switch from the reference to the alternative) to one partner equal to the marginal cost (higher willingness to pay to avoid switching from the reference to the alternative) to the other player. We can move away from the equilibrium in one of four directions of interest, corresponding to the two conditions in Definition 1. When we move away from the equilibrium, the equalities in the first order conditions hold as inequalities. Because preferences satisfy WSLA, we know which direction the inequality will go in for each case. To formalize the idea of "moving away from the equilibrium," consider also some convex mixture of θ^* and θ. These four conditions are, for some small $\alpha \in (0,1)$:

1. Partner f prefers the change in partner f's threat, and partner m does not: $(\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*) \succ_f (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_m ((\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$. If this is true, and because preferences satisfy the WSLA assumption, the FOCs imply $u(\succsim_f, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*)) < |u(\succsim_m, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*))|$ and so no transfer, τ_f, exists such that $(\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$ successfully appeals θ^*.

2. Partner f prefers the change in partner m's threat, and partner m does not: $(\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m) \succ_f (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_m (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m))$. If this is true, and because preferences satisfy the WSLA assumption, then the FOCs imply $u(\succsim_f, (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m), (\theta_f^*, \theta_m^*)) < |u(\succsim_m, (\theta_f^*, (\alpha\theta_m^* + (1 -$

$\alpha)\theta_m), \theta_m^*), (\theta_f^*, \theta_m^*))|$ and so no transfer, τ_f, exists such that $(\theta_f^*, (\alpha\theta_m^* + (1 - \alpha)\theta_m)$ successfully appeals θ^*.

3. Partner m prefers the change in partner f's threat, and partner f does not: $((\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*) \succ_m (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_f ((\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$. If this is true, and because preferences satisfy the WSLA assumption, then the FOCs imply $|u(\succsim_f, (\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*))| > u(\succsim_m, (\alpha\theta_f^* + (1 - \alpha)\theta_f), \theta_m^*), (\theta_f^*, \theta_m^*))$ and so no transfer, τ_m, exists such that $(\alpha\theta_f^* + (1-\alpha)\theta_f), \theta_m^*)$ successfully appeals θ^*.

4. Partner m prefers the change in partner m's threat, and partner f does not: $(\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m) \succ_m (\theta_f^*, \theta_m^*)$ and $(\theta_f^*, \theta_m^*) \succ_f (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m))$. If this is true, and because preferences satisfy the WSLA assumption, then the FOCs imply $|u(\succsim_f, (\theta_f^*, (\alpha\theta_m^* + (1-\alpha)\theta_m), (\theta_f^*, \theta_m^*))| > u(\succsim_m, (\theta_f^*, (\alpha\theta_m^* + (1 - \alpha)\theta_m), \theta_m^*), (\theta_f^*, \theta_m^*))$ and so no transfer, τ_m, exists such that $(\theta_f^*, (\alpha\theta_m^* + (1 - \alpha)\theta_m)$ successfully appeals θ^*.

Because $(\Sigma_f^* \times \Sigma_m^*)^C$ is compact, and because the above argument holds for any possible reference, the θ^* that maximizes the Nash product of induced utilities over WSLA preferences is the unique *appeals-immune solution* to the *bargaining problem* $< \theta, \succsim_f, \succsim_m > \in (\Sigma_f^* \times \Sigma_m^*)^C \times (\rho^{WSLA})^2$. $\blacksquare$

This model with endogenous threats and resource shares allows us to understand power dynamics in the family more fully. Who ever is capable of specifying a more damaging threat will have more control over the household decision making process. Compared to the standard collective model, this generic game with a nested limited commitment problem allows us to understand why we observe certain sharing rules for each family. It allows us to measure power dynamics without assuming that only one partner can specify or act on a threat that reduces household efficiency.

3 Identification in a Panel Data Setting

In this section, we discuss how to bring the model in Section 2 to data. We demonstrate the the expected value of bargaining power, conditional on partners threats and whether they act on them, is semi-parametrically identified in a panel data setting. The key to this strategy is to treat preferences as random variables in order to avoid strong functional form assumptions for utility and social welfare. In this case, the sharing rule is a function of random variables, and so is a random variable itself. Making specific assumptions on the distribution of preferences gives a functional form for the expected value of power, conditioned on endogenously determined outside options. We note which parameters are to be recovered from the data in the derivation, then give specific details for how researchers can do so in Section 3.3.2.

3.1 Deriving the Functional Form of the Expected Value of Women's Intra-Household Bargaining Power

First, note that the optimization problems in (3.3) can be rewritten without reference to the indirect utility functions by applying a compensating variation approach within the nested game with limited commitment, $B(\theta_f, \theta_m, I_f^*, I_m^*)$. For any threat, there is some transfer value that makes partners indifferent between the prices they face in the collective allocation ($I_i = 0$) and the inefficient allocation ($I_i = 1$) for $i \in \{f, m\}$. These transfer amounts make the partners indifferent between the price settings they face in both settings, but the real income they would have at their disposal may still differ across the two collective and less efficient equilibria. Denote these transfer amounts γ_f and γ_m, such that

$$V_f(L(A(\theta_f, \theta_m, I_f = 1, I_m)), \eta y) = V_f(L(A(\theta_f, \theta_m, I_f = 0, I_m)), \eta y + \gamma_f) \text{ and}$$

$$V_m(L(A(\theta_f, \theta_m, I_f, I_m = 1)), (1-\eta)y) = V_m(L(A(\theta_f, \theta_m, I_f, I_m = 0)), (1-\eta)y + \gamma_m).$$

By holding prices constant across the indirect utility functions, only the second argument in the function varies with different threats. The impact on shadow incomes is captured by the changes in the sharing rule and the production function, and the impact on the price setting is captured by the compensating variation amount. The indirect utility functions are strictly increasing in their second argument. As such, the problems in (3.3) can be re-written in a form that does not explicitly reference the indirect utility function. In this formulation, the individuals consider the real (shadow) income they would have at their disposal in the collective and inefficient allocations and chooses to implement their threat if doing so gives a greater real income. That is, the partners solve the following, equivalent problems, where both options are measured in dollar amounts:

$$\max_{I_f \in \{0,1\}} \Big(\eta(\theta_f, \theta_m, I_f = 1, I_m) F(\theta_f, \theta_m, I_f = 1, I_m), \\ \eta(\theta_f, \theta_m, I_f = 0, I_m) F(\theta_f, \theta_m, I_f = 0, I_m) + \gamma_f \Big)$$

$$\max_{I_m \in \{0,1\}} \Bigg((1-\eta(\theta_f, \theta_m, I_f, I_m = 1))F(\theta_f, \theta_m, I_f, I_m = 1),$$

$$(1-\eta(\theta_f, \theta_m, I_f, I_m = 0)F(\theta_f, \theta_m, I_f, I_m = 0) + \gamma_m \Bigg)$$

We observe I_f^* and I_m^*, so we know which of the two options has a higher value. There are four possibilities: the fully efficient outcome $(I_f = 0, I_m = 0)$, he implements his threat but she does not $(I_f = 0, I_m = 1)$, she implements her threat but he does not $(I_f = 1, I_m = 0)$, and both implement their threats $(I_f = 1, I_m = 1)$. In each case, we get family-specific bounds on the possible values that the sharing rule can take. These bounds, when paired with a joint distributional assumption on the partners' preferences, result in a functional form for the expected value of the sharing rule that the econometrician can bring to data.

We derive the bounds next, then describe the remaining estimation steps: (1) assuming a distribution on preferences, (2) using a Heckman (1979) selection approach to recover estimates of threats for those who do not implement threats, (3) using a propensity score matching approach to recover the real (shadow) incomes in the unobserved cases, and (4) recovering the remaining parameters that determine the family-specific bounds on the sharing rule by estimating a fixed effects model using the simulated method of moments. To make equations more legible, let $\theta^e = (\theta_f, \theta_m, I_f = 0, I_m = 0)$, $\theta^{her} = (\theta_f, \theta_m, I_f = 1, I_m = 0)$, $\theta^{him} = (\theta_f, \theta_m, I_f = 0, I_m = 1)$, and $\theta^{both} = (\theta_f, \theta_m, I_f = 1, I_m = 1)$.

Case 1: The Fully Efficient Outcome: Because we observe $(I_f = 0, I_m = 0)$, we know that

$$\eta(\theta^e)F(\theta^e) + \gamma_f > \eta(\theta^{her})F(\theta^{her}) \text{ and}$$

$$(1-\eta(\theta^e))F(\theta^e) + \gamma_m > (1-\eta(\theta^{him}))F(\theta^{him}).$$

Solving these two inequalities for $\eta(\theta^e)$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^e) \in \left[\frac{\eta(\theta^{her})F(\theta^{her}) - \gamma_f}{F(\theta^e)}, 1 - \frac{(1-\eta(\theta^{him}))F(\theta^{him}) - \gamma_m}{F(\theta^e)} \right].$$

Here, the econometrician observes $F(\theta^e)$ and will estimate values for θ^{her} and θ^{him} using a Heckman (1979) selection approach, and for $F(\theta^{her})$ and $F(\theta^{him})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^{her})$, $\eta(\theta^{him})$, γ_f, and γ_m from a simulated method of moments approach. These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Case 2: Only Partner f Implements Their Threat: Because we observe $(I_f = 1, I_m = 0)$, we know that

$$\eta(\theta^e)F(\theta^e) + \gamma_f < \eta(\theta^{her})F(\theta^{her}) \text{ and}$$
$$(1 - \eta(\theta^{her}))F(\theta^{her}) + \gamma_m > (1 - \eta(\theta^{both}))F(\theta^{both}).$$

Solving these two inequalities for $\eta(\theta^{her})$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^{her}) \in \left[\frac{\eta(\theta^e)F(\theta^e) + \gamma_f}{F(\theta^{her})}, 1 - \frac{(1 - \eta(\theta^{both}))F(\theta^{both}) - \gamma_m}{F(\theta^{her})}\right].$$

Here, the econometrician observes $F(\theta^{her})$ and will estimate values for θ^e and θ^{both} using a Heckman (1979) selection approach, and for $F(\theta^e)$ and $F(\theta^{both})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^e)$, $\eta(\theta^{both})$, and $\gamma \equiv \gamma_m - \gamma_f$ from a simulated method of moments approach (which we will explain next). These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Case 3: Only Partner m Implements Their Threat: Because we observe $(I_f = 0, I_m = 1)$, we know that

$$\eta(\theta^{him})F(\theta^{him}) + \gamma_f > \eta(\theta^{both})F(\theta^{both}) \text{ and}$$

$$(1 - \eta(\theta^e))F(\theta^e) + \gamma_m < (1 - \eta(\theta^{him}))F(\theta^{him}).$$

Solving these two inequalities for $\eta(\theta^{him})$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^{him}) \in \left[\frac{\eta(\theta^{both})F(\theta^{both}) - \gamma_f}{F(\theta^{him})}, 1 - \frac{(1 - \eta(\theta^e))F(\theta^e) + \gamma_m}{F(\theta^{him})}\right].$$

Here, the econometrician observes $F(\theta^{him})$ and will estimate values for θ^e and θ^{both} using a Heckman (1979) selection approach, and for $F(\theta^e)$ and $F(\theta^{both})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^e)$, $\eta(\theta^{both})$, and $\gamma \equiv \gamma_m - \gamma_f$ from a simulated method of moments approach. These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Case 4: Both Partners Implement Their Threats: Because we observe $(I_f = 1, I_m =$

1), we know that

$$\eta(\theta^{him})F(\theta^{him}) + \gamma_f < \eta(\theta^{both})F(\theta^{both}) \text{ and}$$

$$(1-\eta(\theta^{her}))F(\theta^{her}) + \gamma_m < (1-\eta(\theta^{both}))F(\theta^{both}).$$

Solving these two inequalities for $\eta(\theta^{both})$, the sharing rule that results in equilibrium, gives the following bounds:

$$\eta(\theta^{both}) \in \left[\frac{\eta(\theta^{him})F(\theta^{him}) + \gamma_f}{F(\theta^{both})}, 1 - \frac{(1-\eta(\theta^{her}))F(\theta^{her}) + \gamma_m}{F(\theta^{both})}\right].$$

Here, the econometrician observes $F(\theta^{him})$ and will estimate values for θ^e and θ^{both} using a Heckman (1979) selection approach, and for $F(\theta^e)$ and $F(\theta^{both})$ using a propensity score matching approach. They can recover estimates of $\eta(\theta^e)$, $\eta(\theta^{both})$, and $\gamma \equiv \gamma_m - \gamma_f$ from a simulated method of moments approach. These estimates and observed values pin down the family specific bounds on the possible values of the sharing rule.

Distributional Assumptions on Preferences: At this point, we must make some assumption on preferences. Assuming that a specific functional form for utility functions adequately captures the preferences for all individuals in the sample is a strong, parametric assumption. We can make a weaker, semi-parametric assumption: that preferences and the social welfare function are random variables drawn according to some joint probability distribution function from $(\rho^{sm})^2 \times S$, where S is the set of permissible social welfare functions. Doing so implies a particular distribution of the sharing rule on the set determined by the observed values (I_f, I_m). This is because a (linear) bijection relates the specified threats, the pair (I_f, I_m), and the drawn random variables $(U_f, U_m, \tilde{U})$.[7]

A particularly tractable and weak distributional assumption is that preferences are conditionally uniformly distributed on some subset of $(\rho^{sm})^2 \times S$. Consider an arbitrarily drawn partition of $(\rho^{sm})^2 \times S$, with $c \in \mathbb{N}$ subsets $C \subset (\rho^{sm})^2 \times S$, where $\underset{i\in[1,c]}{\cup} C_i = (\rho^{sm})^2 \times S$. For each family in the population, let the triple $(U_f, U_m, \tilde{U})$ be uniformly drawn from the same subset, C_i. This is a very weak "similarity between partners" assumption since the partition can be drawn in any way (any number of subsets, with subsets of any size). This distributional assumption on preferences results in a uniform distribution of the sharing rule on the family-specific bounds derived above. As such, the estimator of the sharing rule can be written as follows in each of the four cases:

[7] For proofs that this bijective function exists and is linear, see Klein and Barham (2018).

1. $\mathbb{E}[\eta(\theta^{e})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{her})F(\theta^{her})-(1-\eta(\theta^{him}))F(\theta^{him})-\gamma_f+\gamma_m}{F(\theta^{e})}\right)$
2. $\mathbb{E}[\eta(\theta^{her})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{e})F(\theta^{e})-(1-\eta(\theta^{both}))F(\theta^{both})+\gamma_f+\gamma_m}{F(\theta^{her})}\right)$
3. $\mathbb{E}[\eta(\theta^{him})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{both})F(\theta^{both})-(1-\eta(\theta^{e}))F(\theta^{e})-\gamma_f-\gamma_m}{F(\theta^{him})}\right)$
4. $\mathbb{E}[\eta(\theta^{both})\big|(U_f, U_m, \tilde{U}) \overset{\text{Unif}}{\sim} C_i] = \frac{1}{2} + \frac{1}{2}\left(\frac{\eta(\theta^{him})F(\theta^{him})-(1-\eta(\theta^{her}))F(\theta^{her})+\gamma_f-\gamma_m}{F(\theta^{both})}\right)$

3.2 Identification

Summary of requirements: The researcher must have access to panel data with at least two waves for each household. Each wave must have demographic information about each partner (for instance their age, earnings, and education), and information about the threats that partners might specify (for instance, domestic abuse or divorce). From this information, the researcher can learn the threats partners specify, and whether they implement them. The values that the econometrician can estimate are the threats for partners who do not act on them, the real incomes for each household in the context we do not observe (for instance, in the case where $\{I_m = 0, I_f = 0\}$ when we observe $\{I_m = 1, I_f = 0\}$), and the expected value of the sharing rule. The assumptions required are distributional assumptions on preferences, distributional assumptions on the sharing rule in the cases we do not observe, and $H + 2$ distributional assumptions on regression coefficients to be estimated using the simulated method of moments (McFadden, 1989), where H is the number of households in the sample. These assumptions on the coefficient distributions are analogous to a single assumption on the distribution of the error term in a fixed effects model with H constraints on the intercepts, estimated using constrained least squares.

An optional, but helpful, data feature is an instrument that is correlated with how damaging each partners' threat is, and uncorrelated with whether or not they implement their threat. Lewbel and Pendakur (2019) use the thickness of each person's walls (a proxy for how likely it is that neighbors overhear domestic abuse) as an instrument in a similar setting. For instance, if you're worried about being overheard, you might specify a threat that does not entail loud volumes. If so, implementing the threat is uncorrelated with the likelihood of being overheard. Access to such an instrument makes estimation arguably less biased.

Summary of the strategy: There are four steps in estimation. First, the economist must quantify how damaging each threat is — i.e. specify a function that maps from the space of threats to the real numbers. The second step is to predict the threat for those who the researcher does not observe acting on a threat. The third is to fit a household fixed effects model with family-specific restrictions on each intercept. The fourth is to use the recovered regression coefficients and threat values to recover the expected value of women's intra-household bargaining power.

Step 1 of the estimation strategy: The researcher can quantify how damaging each threat is by using violence scales, like the one Ramos (2016) uses. This approach generates a scale from 0 to 1 that describes the severity of the threat. Each possibility is given some (normalized) value, and the score is the sum of the quantified individual actions comprising a threat. For instance, a women who is punched and verbally threatened with a weapon would have a higher score than a woman who is only threatened with a weapon. A women who is only punched would have a higher score than a woman who is only threatened.

In order to incorporate all types of threats into our model, the econometrician must also assign values to actions like divorce and the inefficient allocation of productive assets. This is a necessarily subjective exercise, but is based on the objective change in real incomes that each partner would have if the threat is exercised. Ideally, the threat scale would assign higher values to threats that reduce the real income available to an individual's partner by a greater amount. Since divorce (and any threat that ends cohabitation) results in both partners consuming at market prices, and consuming according to their own private incomes instead of a shared income, divorce can be given a higher scale value than violence. We suggest that any threat which only reduces the quality of the production function be given the lowest values in the scale, any threat that only reduces the quality of the consumption technology be given the next highest values, and any threat that affects both be given higher values. Various scales could be used to test the robustness of the results to various scale specifications.

Step 2 of the estimation process: The econometrician can recover the expected values of the threats that are not implemented using a Heckman (1979) selection approach. This allows us to recover estimates of these threats based on observable demographic characteristics, and an unobservable selection parameter, which describes the differences in specified threats between those who choose to implement them, and those who do not implement them. Disaggregating the sample by gender allows the econometrician to recover regression coefficients and selection parameters that differ across men and women. As such, estimate the following two models using full information maximum likelihood:

$$I_{f,h,t} = X_{f,h,t}\beta_{f,h,t} + z_{f,h,t}\beta_{instrument,f,h,t} + \epsilon_{f,h,t} \tag{8}$$

$$\theta_{f,h,t} = X_{f,h,t}\beta_{f,h,t} + \epsilon_{f,h,t}$$

$$I_{m,h,t} = X_{m,h,t}\beta_{m,h,t} + z_{m,h,t}\beta_{instrument,m,h,t} + \epsilon_{m,h,t} \tag{9}$$

$$\theta_{m,h,t} = X_{m,h,t}\beta_{m,h,t} + \epsilon_{m,h,t}$$

where t indexes time, h indexes the household that partners f and m belong to at the start of the panel, $z_{f,h,t}$ is an instrument for I_f (that is $cov(I_{f,h,t}, z_{f,h,t}) \neq 0$ and $cov(\theta_{f,h,t}, z_{f,h,t}) = 0$), $z_{m,h,t}$ is an instrument for I_m (that is $cov(I_{m,h,t}, z_{m,h,t}) \neq 0$ and $cov(\theta_{m,h,t}, z_{m,h,t}) = 0$), $X_{f,h,t}$ is a matrix of demographic characteristics for the women in the sample, and $X_{m,h,t}$ is a matrix of demographic characteristics for the men in the sample.[8] The results are the (nuisance) regression coefficients that relate observable characteristics to the severity of the specified threat, and the likelihood that the threat is carried out, plus the (nuisance) selection coefficient estimates that describe how much threats differ on average between those who implement them and those who do not. The predictions for threats of those who do not implement them are the fitted values from these Heckman (1979) selection models: $\hat{\theta}_{f,h,t}$ and $\hat{\theta}_{m,h,t}$. The researcher can use the fitted values for all individuals for consistency.

Step 3 of the estimation process: the econometrician must recover the shadow incomes for individuals in each case. They can use the predicted threats as propensity scores in a propensity score matching approach (Rosenbaum and Ruben, 1983). In each case, the econometrician needs to recover unobserved income values. We suggest using a nearest neighbor matching approach to recovering these parameters. For instance, if the researcher observes $(I_f = 0, I_m = 0)$ for some household, then they need to recover values for $F(\theta^{him})$ and $F(\theta^{her})$. We suggest using the income vale for the nearest household (nearest in the value of $\hat{\theta}_{f,h,t}$) such that that neighboring household has observed value $I_f = 1$ for the estimate of $F(\theta^{her})$. Symmetrically, we suggest using the income value for the closest neighbor in the value of $\hat{\theta}_{f,h,t}$ (where that neighbor has observed value $I_m = 1$) to recover the estimate of $F(\theta^{him})$. Denote these matching estimates as $\hat{F}(\hat{\theta}^{him})$ and $\hat{F}(\hat{\theta}^{her})$.

Step 4 in the estimation process: at this point, the remaining unknown parameter values in the bounds are the sharing rule values in the unobserved cases, and the compensating variation values, γ_f and γ_m. In order to recover these values, we suggest a fixed effects strategy. Since the approach is the same across all four cases, we only describe it for the fully efficient case. The first step in this panel estimation is to notice that the functional form for the expected value of the equilibrium sharing

[8] The model does not require instruments, but in that case it does require an additional joint normality assumption.

rule can be rearranged so that income is a linear function of the estimated shadow incomes in the unobserved cases, and unknown parameters. In the fully efficient case, this gives the following expression, suppressing notation:

$$F(\theta^e) = \frac{\eta(\theta^{her})F(\theta^{her}) - (1-\eta(\theta^{him}))F(\theta^{him}) - \gamma_f + \gamma_m}{\mathbb{E}[\eta(\theta^e)] - 1}$$

the researcher can estimate this equation using a fixed effects approach, with the following definitions for regression coefficients:

$$F(\theta^e)_{h,t} = \beta_{0,h} + \beta_1 \hat{F}(\hat{\theta}^{her})_{h,t} + \beta_2 \hat{F}(\hat{\theta}^{him})_{h,t} + \epsilon_{h,t} \text{ such that} \tag{10}$$

$$\beta_{0,h} \equiv \frac{\gamma_m - \gamma_f}{\mathbb{E}[\eta(\theta^e)] - 1}$$

$$\beta_1 \equiv \frac{\eta(\theta^{her})}{\mathbb{E}[\eta(\theta^e)] - 1}$$

$$\beta_2 + \epsilon_{h,t,2} \equiv \frac{1-\eta(\theta^{him})}{\mathbb{E}[\eta(\theta^e)] - 1}$$

$$\epsilon_{h,t,2} \equiv \epsilon_{h,t}\hat{F}(\hat{\theta}^{him})$$

$$\eta^e \in [0,1], \quad \eta^{her} \in [0,1], \quad \text{and } \eta^{him} \in [0,1].$$

Estimating (3.10) using the simulated method of moments requires the following assumptions: distributional assumptions for the H intercept parameters, $\beta_{0,h}$; distributional assumptions for the two slope parameters, β_1 and β_2; and distributional assumptions for the off-equilibrium sharing rule values, $\eta(\theta^{her})$ and $\eta(\theta^{him})$. Under these assumptions, there are two remaining unknowns per household in (3.10), and at least two equations per household in the panel. After the simulated method of moments estimation, the estimates of the regression coefficients can be plugged into the definitions in (9) to recover the expected value of the sharing rule in the equilibrium, and the difference between the two partners' compensating variation values, $\gamma_m - \gamma_f$. The constraints that the sharing rules only take values between 0 and 1 limit the regression coefficients' supports.

We suggest the following procedure for selection assumptions for the regression coefficients. First fit the regression in (3.10) without taking into account the adding up constraints ($\eta^e \in [0,1], \quad \eta^{her} \in [0,1], \quad \text{and } \eta^{him} \in [0,1]$). This gives naive regression coefficients. Assume that the simulated method of moments coefficients for (3.10) are truncated normal distributions with means given by the naive estimators, and variance given by the corresponding naive standard errors. For the off-equilibrium sharing rule values, we recommend assuming uniform distributions on the unit interval. With a large number of simulations, this process will give consistent estimators of the equilibrium sharing rule and the compensating variation value difference.

4 Conclusion

It is well known that policy makers can influence intra-household bargaining power by increasing the value of one partners' outside option (e.g. Mazzocco, 2007; Voena, 2015). In the standard limited commitment model, when policies change the value of the outside option so much that one partner newly prefers to leave the marriage, bargaining power in the family is renegotiated until that partner is indifferent between leaving and staying. Only drastic changes to the outside option result in a change in bargaining power, since most changes will be insufficient to make either partner prefer divorce to the current marriage contract. As such, power is expected to be fairly static, and changes in power dynamics are often expected to be small (see, e.g., Lise and Yamada, 2019).

When partners can specify outside options besides divorce, changes in power can be fairly frequent. Partners might subtly shift their specified threat, resulting in a different sharing rule, in response to smaller changes to relative outside options. Changes that would not cause the participation constraint to bind if the outside option is divorce might cause it to bind if the outside option is changing the amount of time allocated to labor (as in Walther, 2018) or the savings technology used (as in Schaner, 2015).

A different picture of power in the family emerges. Policy makers can hope to influence power dynamics even with "small" changes to the relative value of outside options. Events like minor salary increases, or even holiday bonuses, can change the balance of power in the family. Power, once collective models make outside options endogenous, appears to be highly variable and contested, not static.

Can policy makers influence partners' choices to specify certain threats, or pass policies that cause individuals to chose to not implement their threats? Any policy that expands individuals' space of possible threats to include more attractive possible outside options will accomplish the former. Any policy that changes the costs and benefits of implementing may achieve the latter. For instance, higher ratios of female-to-male police officers in the United States increases the likelihood that women report domestic abuse, increasing the cost of implementing that threat for men, and reducing the likelihood that they implement their threat (Miller and Segal, 2019). Our model predicts that such a policy would increase women's bargaining power.

In general, we would expect families to be more efficient in settings where men and women have more equal rights and opportunities. Any policy that increases equality in institutions outside of the family — like labor markets or divorce courts — gives the less powerful partner additional recourse against more damaging threats.

This changes the threats that obtain in equilibrium, and shifts the bargaining power dynamic towards equality within the family.

Reference List

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

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

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

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

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

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

De Mel, S., McKenzie, D., & Woodruff, C. (2009). Are Women More Credit Constrained? Experimental Evidence on Gender and Microenterprise Returns. American Economic Journal: Applied Economics, 1(3), 1-32.

Donni, O., & Molina, J. A. (2018). Household Collective Models: Three Decades of Theoretical Contributions and Empirical Evidence. IZA Discussion Paper No. 11915. Available at SSRN: https://ssrn.com/abstract=3286175

Hanany, E., & Safra, Z. (2000). Existence and Uniqueness of Ordinal Nash Outcomes. Journal of Economic Theory, 90(2), 254-276.

Klein, M. J., & Barham, B. L. (2019). Point Estimates of Household Bargaining Power Using Outside Options. In University of Wisconsin-Madison, Staff Paper No. 590 March 2018.

Harsanyi, J. C. (1986). Rational Behaviour and Bargaining Equilibrium in Games and Social Situations. CUP Archive.

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

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

Ligon, E. A. (2011). Dynamic Bargaining in Households (with an Application to Bangladesh). Available at SSRN 1776810.

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

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

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

Miller, A. R., & Segal, C. (2019). Do Female Officers Improve Law Enforcement Quality? Effects on Crime Reporting and Domestic Violence. The Review of Economic Studies, 86(5), 2220-2247.

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

Rosenbaum, P. R., & Rubin, D. B. (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika, 70(1), 41-55.

Rubinstein, A., Safra, Z., & Thomson, W. (1992). On the Interpretation of the Nash Bargaining Solution and its Extension to Non-Expected Utility Preferences. Econometrica, 1171-1186.

Schaner, S. (2015). Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings. American Economic Journal: Applied Economics, 7(2), 135-74.

Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household. Journal of Political Economy, 104(5), 1010-1046.

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

Walther, S. (2018). Noncooperative Decision Making in the Household: Evidence From Malawi. Journal of Development Economics, 134, 428-442.