ECARES

The revealed preference approach to analyzing collective household consumption: a practical roadmap

Laurens Cherchye
Department of Economics, University of Leuven (KU Leuven)

Bram De Rock
ECARES, Université libre de Bruxelles and
Department of Economics, University of Leuven (KU Leuven)

Frederic Vermeulen
Department of Economics, University of Leuven (KU Leuven)

November 2026

ECARES working paper 2026-32

ECARES
ULB - CP 114/04
50, F.D. Roosevelt Ave., B-1050 Brussels BELGIUM
www.ecares.org

The revealed preference approach to analyzing collective household consumption: a practical roadmap

Laurens Cherchye* Bram De Rock† Frederic Vermeulen‡

November 30, 2026

Abstract

This chapter surveys the current state of the nonparametric revealed preference (RP) approach to analyzing collective consumption behavior, and synthesizes the methodological tools that are available. It provides a practical roadmap for testing, recovering, and quantifying collectively rational (i.e., Pareto efficient) household consumption behavior without relying on any parametric assumptions regarding individual utilities or demand. Building on the RP tradition for the classic unitary model, we outline a framework for analyzing collective household decision making with private and public consumption. The framework uses personalized prices and quantities to represent within-household allocation patterns, and employs the sharing rule to capture the distribution of resources across household members. We characterize collective rationality through nonparametric testable implications, and we show that this characterization provides a productive basis for (set) identification procedures using linear and integer programming. We also discuss methodological extensions that integrate RP methods with estimated demand systems, marriage market information, household production, and intertemporal decision environments. Finally, we briefly illustrate how the method can be used in empirical work to study sharing-rule bounds and poverty at the individual level.

*Department of Economics, University of Leuven (KU Leuven). E-mail: laurens.cherchye@kuleuven.be. Laurens Cherchye thanks the Fund for Scientific Research - Flanders (FWO) and the Research Fund KU Leuven for financial support. He gratefully acknowledges the Van Eeckhoudt-Ruttens family for supporting the 600 years of KU Leuven Chair on "Growing Up in Poverty".

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

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

1 Introduction

This chapter surveys the nonparametric revealed preference (RP) approach to identifying the collective model of household consumption behavior. The RP tradition offers a way to assess the empirical content of utility maximization without imposing parametric structure on utilities or demand, relying solely on preference information revealed by observed consumer choices. The concept originates from the pioneering contributions of Samuelson (1938, 1948), who introduced the Weak Axiom of Revealed Preference (WARP) to capture the restrictions on demand as the observable footprint of rational behavior, and Houthakker (1950), who formulated the Strong Axiom of Revealed Preference (SARP) to ensure the existence of a well-behaved utility representation consistent with observed demand. Building on these foundations, Afriat (1967) demonstrated how finite data sets can be tested for consistency with utility maximization, and how utility functions can be recovered without imposing functional form assumptions. Subsequently, Diewert (1973) offered theoretical refinements, and Varian (1982) unified and extended these insights into a transparent testing framework based on the Generalized Axiom of Revealed Preference (GARP) and linear inequalities commonly known as "Afriat inequalities". This seminal work provides a comprehensive toolkit for analyzing the behavior of single consumers without relying on parametric assumptions. Consequently, it enables a nonparametric analysis of household consumption within the framework of the unitary model, which treats the household as if it were a single decision maker.

The collective model refines this unitary perspective by acknowledging that households comprise multiple members who act as decision makers with potentially heterogeneous preferences. Within this framework, observed household behavior is modeled as the outcome of a bargaining process among these members, with the resulting allocation depending on their respective bargaining power. The foundations of this approach were laid by Chiappori (1988, 1992), who formalized the notion that observed household choices can be interpreted as Pareto efficient outcomes of an intrahousehold decision process, and established the associated testable implications for household demand. We refer to such Pareto efficient behavior as "collectively rational".[1] Later contributions, notably Browning and Chiappori (1998) and Chiappori and Ekeland (2006, 2009), provide a complete differential characterization of collectively rational behavior, generalizing the Slutsky symmetry condition of the unitary

[1] Collectively rational (or Pareto efficient) household consumption behavior is to be contrasted with noncooperative or semicooperative behavior. See Browning, Chiappori, and Lechene (2010), Lechene and Preston (2011) and d'Aspremont and Dos Santos Ferreira (2014) for differential characterizations of non-cooperative and semicooperative household consumption models. Nonparametric analogues are developed by Cherchye, Demuynck, and Rock (2011) and Cherchye, Cosaert, Demuynck, and De Rock (2020), who provide linear/integer programming formulations similar to those for the collective model (presented in Section 4.3).

model.[2] This approach is particularly well suited to parametric analyses of household consumption behavior, which rely on specific functional-form assumptions for individual utility functions and household demand.

In this chapter we survey the current state of the nonparametric RP approach to analyzing collectively rational household consumption, which complements this parametric work. This nonparametric alternative offers three key advantages. First, it provides testable implications for collective rationality that are based on observable price–quantity data rather than relying on strong parametric assumptions about individual utilities and household demand. Second, it makes it possible to identify in a constructive way the key primitives that shape observed collective behavior. An important example is the sharing rule that determines how resources are distributed within households. This can be done without relying on parametric assumptions that are usually impossible to verify. Third, the analysis remains computationally feasible by drawing on linear and integer programming methods. These methods can handle situations in which consumption quantities are not assignable, or only partially assignable, to individual household members. The website `revealedpreferences.org` provides a useful point of reference for this research agenda, including related papers, code, data, and RP tools.

At this point, it is useful to mention a few disadvantages of the RP approach *vis-à-vis* the differential approach. First, RP methods generally deliver set identification results rather than point-identification results. For example, the identification strategy discussed in this chapter gives upper and lower bounds on the sharing rule. This is not problematic when the bounds are close to each other, but the bounds can be wide when the data contain little relative price variation. This is related to a second limitation: the data sometimes lack the power to reject a behavioral hypothesis by means of RP restrictions. For example, if a consumer's budget lines do not cross in the data, then the GARP conditions discussed below can never be violated. Any behavior will then be consistent with the unitary model, even behavior where the consumer randomly chooses a commodity bundle that exhausts the available budget (see, e.g., Bronars (1987)). There are remedies for such cases. We will come back to these when we discuss extensions of the main model of this chapter.[3]

This chapter is organized as follows. Section 2 revisits the unitary model and the cele-

[2]More specifically, "differential characterization" refers to a set of conditions, expressed in terms of derivatives of the demand functions (such as substitution matrices or Slutsky-type derivatives), that fully describe when observed behavior can be generated by a particular model. In the collective setting, these conditions generalize the familiar Slutsky symmetry condition from the unitary model and provide derivative-based restrictions on how individual preferences and the within-household sharing rule jointly shape household demand responses.

[3]Another way to look at this is how much mileage can be obtained by assuming an explicit (non-verifiable) functional form in an empirical analysis, which goes beyond the theoretical model focused on.

brated Afriat Theorem, which provides the conceptual and computational template for the RP analysis. Section 3 introduces the notion of rational collective consumption for two-member households with private and public goods. It also addresses partial observability through feasible personalized quantities, and clarifies how private consumption with intra-household externalities can be treated as public quantities. Section 4 adapts the RP analysis to the collective setting (building on Cherchye, De Rock, and Vermeulen, 2007, 2011), introducing feasible personalized (Lindahl) prices, and specifying member-level RP conditions that ensure each member's choices are consistent with utility maximization at personalized budgets defined by the intrahousehold sharing rule. This yields a collective counterpart to unitary RP equivalences and Afriat-type inequalities. It also provides a productive basis for recovering intrahousehold allocations. In Section 5, we provide a selective survey of extensions that broaden the scope of nonparametric analyses of the collective approach. These include combining RP with estimated demand systems to obtain sharper bounds on sharing rules, embedding intrahousehold allocation patterns within the structure of individuals' marriage markets to link sharing rules to marital sorting and stability, incorporating time use and household production, and extending the analysis to dynamic environments characterized by intertemporal household consumption behavior with varying degrees of intrahousehold commitment. Together, these advances illustrate how RP tools can integrate structural insights and additional data to further enrich the empirical analysis of intrahousehold allocation patterns. Finally, Section 6 illustrates the empirical usefulness of the method. We focus on one specific application to sharing-rule bounds and individual poverty, which shows how the methodology can be used to move from RP restrictions to welfare-relevant conclusions.

In sum, this chapter provides a general introduction to nonparametric identification of the collective model. The goal is to equip applied researchers with a roadmap for testing collective rationality, recovering intrahousehold consumption allocations and sharing rules, and extending the analysis to richer environments, while preserving the robustness related to the nonparametric nature of the RP approach.

2 The unitary model

The unitary model assumes that a household acts as if it were a single decision maker. This implies that each observed household quantity bundle is assumed to maximize a single utility function subject to the corresponding household budget constraint. In order to define this more formally, we need to introduce some notation.

In the tradition of Afriat (1967), Diewert (1973) and Varian (1982) we start our analysis from a finite set of T price-quantity pairs. We denote the vectors of prices and quantities

associated with observation t by $\mathbf{p}_t$ and $\mathbf{q}_t$, where $\mathbf{p}_t \in \mathbb{R}^N_{++}$ and $\mathbf{q}_t \in \mathbb{R}^N_+$. This gives us the set of observations

$$S^{un} = \{(\mathbf{p}_t; \mathbf{q}_t)\,; t = 1, \ldots, T\}\,.$$

The following definition states the condition for rational consumption behavior in terms of the unitary model.

Definition 1 (unitary rationalization) *Let S^{un} be a set of observations. A utility function U provides a unitary rationalization of S^{un} if, for each observation t, we have*

$$U(\mathbf{q}_t) \geq U(\mathbf{z})$$

for all $\mathbf{z} \in \mathbb{R}^N_+$ satisfying the household budget constraint

$$\mathbf{p}_t'\mathbf{z} \leq \mathbf{p}_t'\mathbf{q}_t.$$

In words, every observed bundle $\mathbf{q}_t$ maximizes the utility function U given the household's budget constraint in choice situation t. Throughout, we assume utility functions that are locally non-satiated, continuous, monotonically increasing, and concave. Strictly speaking, the Afriat Theorem for the unitary model (see Theorem 1 below) only requires non-satiated utility functions; see Varian (1982) for a detailed discussion. However, our characterization of collectively rational behavior imposes continuity, monotonicity, and concavity, and therefore we maintain these properties throughout. For a discussion of nonparametric RP conditions for collectively rational consumption under weaker assumptions, we refer to Cherchye, De Rock, and Vermeulen (2010).

A central concept in the empirical RP analysis of the unitary model is the Generalized Axiom of Revealed Preference (GARP).

Definition 2 (GARP) *Let S^{un} be a set of observations. The set S^{un} satisfies GARP if there exist relations R_0 and R that meet:*

(i) if $\mathbf{p}_s'\mathbf{q}_s \geq \mathbf{p}_s'\mathbf{q}_t$, then $\mathbf{q}_s\ R_0\ \mathbf{q}_t$;

(ii) if $\mathbf{q}_s\ R_0\ \mathbf{q}_u$, $\mathbf{q}_u\ R_0\ \mathbf{q}_v$, …, $\mathbf{q}_z\ R_0\ \mathbf{q}_t$ for some (possibly empty) sequence $(u, v, \ldots, z)$, then $\mathbf{q}_s\ R\ \mathbf{q}_t$;

(iii) if $\mathbf{q}_s\ R\ \mathbf{q}_t$, then $\mathbf{p}_t'\mathbf{q}_t \leq \mathbf{p}_t'\mathbf{q}_s$.

In words, the quantities $\mathbf{q}_s$ are directly revealed preferred over the quantities $\mathbf{q}_t$ (i.e., $\mathbf{q}_s\ R_0\ \mathbf{q}_t$) if $\mathbf{q}_s$ were chosen when $\mathbf{q}_t$ were equally attainable (i.e., $\mathbf{p}_s'\mathbf{q}_s \geq \mathbf{p}_s'\mathbf{q}_t$); see condition

(i). Next, the relation R exploits transitivity of preferences; see condition (ii). Finally, condition (iii) imposes that the quantities $\mathbf{q}_t$ cannot be more expensive than revealed preferred quantities $\mathbf{q}_s$.

A key result in RP theory is that a set S^{un} of price–quantity observations is rationalizable by the unitary model if and only if it satisfies GARP. This fundamental insight is formalized and extended in the classic Afriat Theorem (Varian, 1982, building on Afriat, 1967):

Theorem 1 (Afriat Theorem) *Let S^{un} be a set of observations. The following statements are equivalent:*

(i) There exists a utility function U that provides a unitary rationalization of S^{un}.

(ii) The set S^{un} satisfies GARP.

(iii) For all $s, t \in \{1, \ldots, T\}$, there exist numbers U_t and $\lambda_t \in \mathbb{R}_{++}$ that meet the Afriat inequalities

$$U_s - U_t \leq \lambda_t \mathbf{p}_t' (\mathbf{q}_s - \mathbf{q}_t) .$$

In this result, condition (ii) implies that data consistency with GARP is necessary and sufficient for a unitary rationalization of the data, while condition (iii) provides an equivalent empirical characterization in terms of the "Afriat inequalities".

Theorem 1 gives us two distinct ways to verify whether a given data set S^{un} is unitary rationalizable. The first method of testing GARP was originally suggested by Varian (1982). This method consists of three steps, which correspond to the three conditions in Definition 2. The first step constructs the relation R_0 from the data set $S^{un} = \{(\mathbf{p}_t; \mathbf{q}_t) ; t = 1, \ldots, T\}$: $\mathbf{q}_s \; R_0 \; \mathbf{q}_t$ if and only if $\mathbf{p}_s' \mathbf{q}_s \geq \mathbf{p}_s' \mathbf{q}_t$. The second step computes the transitive closure R of R_0. Varian (1982) suggests using the algorithm described in Warshall (1962) for efficiently computing these transitive closures. The third step verifies whether $\mathbf{p}_t' \mathbf{q}_t \leq \mathbf{p}_t' \mathbf{q}_s$ whenever $\mathbf{q}_s \; R \; \mathbf{q}_t$. If this is the case, then the data set satisfies GARP, which means that the data set is unitary rationalizable. Conversely, if the data violate GARP, then we conclude that the data are not unitary rationalizable. Due to its efficiency, this GARP-based method is easy to use in applied work.

The second method for verifying unitary rationalizability directly tests the feasibility of the Afriat inequalities. These inequalities are linear in the unknowns U_t and λ_t for $t \in \{1, \ldots, T\}$, which means their feasibility can be checked using standard linear programming techniques. The data set S^{un} is unitary rationalizable if and only if a feasible solution $(U_t, \lambda_t)_{t=1}^T$ exists. Intuitively, the values U_t and λ_t represent the utility levels and marginal utilities of income associated with each observation t (that is, for the observed bundle $\mathbf{q}_t$)

that jointly rationalize the choices in S. See Afriat (1967) and Diewert (1973) for detailed discussions.

From an empirical perspective, the two approaches (based on GARP and on Afriat inequalities) complement each other: the GARP test is computationally efficient, while the Afriat inequalities provide additional structural information about preferences (in particular U_t and λ_t) rationalizing the observed consumption behavior.

3 The collective model

We now move beyond the unitary framework introduced in Section 2 and explicitly consider households with multiple decision makers. In particular, we focus on households composed of two members, indexed by $m \in \{1, 2\}$, who jointly determine their consumption choices. The environment contains N privately consumed goods and K publicly consumed goods. Private goods generate utility only for the individual who consumes them. Typical examples include personal clothing, individual meals, or a private smartphone. Public goods, by contrast, benefit all household members. Standard examples are housing services, heating, shared electricity use, or a family car that both members can use together.

Let $\mathbf{q} \in \mathbb{R}^N_+$ denote the vector of household-level quantities of private goods, and let $\mathbf{Q} \in \mathbb{R}^K_+$ represent the quantities of publicly consumed goods. For each member m, the vector $\mathbf{q}^m \in \mathbb{R}^N_+$ captures that member's privately consumed quantities, subject to the adding-up condition $\mathbf{q}^1 + \mathbf{q}^2 = \mathbf{q}$. We assume that each member's preferences are described by an individual-specific utility function $U^m(\mathbf{q}^m, \mathbf{Q})$, meaning that member m has preferences over her or his own private consumption and public consumption.

This specification rules out externalities in private consumption, meaning that one member does not directly derive utility from the other member's private consumption. Nevertheless, the framework can accommodate such externalities by reclassifying the relevant goods as public goods. For instance, if a private good n consumed by member m generates external benefits for the other member (e.g., clothing that contributes to shared household well-being), it can be included in $\mathbf{Q}$. In this way, $\mathbf{q}^m$ should be interpreted as capturing purely private consumption, while any privately consumed goods that create external effects are treated as part of the public consumption vector $\mathbf{Q}$.

As before, we assume an empirical analyst who observes T household choices characterized by prices and quantities. For each observation $t \in \{1, \ldots, T\}$, $\mathbf{p}_t \in \mathbb{R}^N_{++}$ and $\mathbf{q}_t \in \mathbb{R}^N_+$ denote prices and household quantities for private goods, while $\mathbf{P}_t \in \mathbb{R}^K_{++}$ and $\mathbf{Q}_t \in \mathbb{R}^K_+$ denote prices and quantities for public goods. In general, the analyst does not observe individual private consumption $\mathbf{q}^m_t$. If some private goods are observed at the individual

level, they are called assignable goods (see Bourguignon, Browning, and Chiappori (2009)). Formally, for each member m and observation t, define the observed assignable quantities $\mathbf{q}_t^{Am} \in \mathbb{R}_+^N$ as lower bounds for the unobserved quantities $\mathbf{q}_t^m$, i.e.,

$$\mathbf{q}_t^m \geq \mathbf{q}_t^{Am}.$$

Examples include exclusive goods that benefit only one member. More generally, assignable information may pertain to a subset of goods and/or household members. We summarize the observed data as

$$S = \{(\mathbf{p}_t, \mathbf{P}_t; \mathbf{q}_t, \mathbf{q}_t^{A1}, \mathbf{q}_t^{A2}, \mathbf{Q}_t); t = 1, \ldots, T\}.$$

In practice, we typically have $\mathbf{q}_t^{A1} + \mathbf{q}_t^{A2}$ is strictly smaller than $\mathbf{q}_t$, meaning that some private consumption cannot be assigned prior to analysis. It is even possible that $\mathbf{q}_t^{Am} = \mathbf{0}$, i.e., no assignable information is available. To address this, we introduce the concept of feasible personalized quantities.

Definition 3 (feasible personalized quantities) *Let S be a set of observations. For each observation t and member $m = 1, 2$, feasible personalized quantities $\mathfrak{q}_t^m \in \mathbb{R}_+^N$ satisfy*

$$\mathfrak{q}_t^m \geq \mathbf{q}_t^{Am} \quad \textit{and} \quad \mathfrak{q}_t^1 + \mathfrak{q}_t^2 = \mathbf{q}_t.$$

Using feasible personalized quantities, we define the condition for a collective rationalization of a data set. Intuitively, this requires that observed household consumption can be represented as a Pareto efficient outcome of some within-household bargaining process.

Definition 4 (collective rationalization) *Let S be a set of observations. A pair of utility functions (U^1, U^2) provides a collective rationalization of S if, for each observation t, there exist feasible personalized quantities $\mathfrak{q}_t^m$ and Pareto weights $\mu_t^m > 0$, $m = 1, 2$, such that*

$$\mu_t^1 U^1(\mathfrak{q}_t^1, \mathbf{Q}_t) + \mu_t^2 U^2(\mathfrak{q}_t^2, \mathbf{Q}_t) \geq \mu_t^1 U^1(\mathfrak{z}^1, \mathbf{Z}) + \mu_t^2 U^2(\mathfrak{z}^2, \mathbf{Z})$$

for all $\mathfrak{z}^m \in \mathbb{R}_+^N$ and $\mathbf{Z} \in \mathbb{R}_+^K$ satisfying the household budget constraint

$$\mathbf{p}_t'(\mathfrak{z}^1 + \mathfrak{z}^2) + \mathbf{P}_t'\mathbf{Z} \leq \mathbf{p}_t'\mathbf{q}_t + \mathbf{P}_t'\mathbf{Q}_t.$$

In words, a collective rationalization requires that observed choices maximize a weighted sum of members' utilities subject to the household budget. In this respect, each Pareto

weight reflects a member's bargaining power since it captures how important that member's preferences are for the household decision.

A final remark concerns the treatment of children within this framework. There are two conceptually distinct ways to incorporate them. One option is to model children as additional decision makers (see, for example, Dauphin et al., 2011). Cherchye, De Rock, and Vermeulen (2011) discuss collective RP analysis in the general case with $M \geq 2$ household members. This applies when children have their own decision power. A second option is to treat expenditures on children as part of the household's public consumption, so that their welfare is reflected indirectly through the vector $\mathbf{Q}$ of publicly consumed goods. This interpretation is particularly natural when children do not actively participate in decision making, yet their needs shape household-wide consumption of goods such as housing, heating, or shared services. The vector $\mathbf{Q}$ may also include public goods produced within the household, such as jointly generated services or time-intensive activities. Intrahousehold production can likewise contribute to children's welfare, in which case their utility depends not only on expenditures but also on time invested in childcare (see, for example, Blundell, Chiappori, and Meghir, 2005; Cherchye, Rock, and Vermeulen, 2012b). We will return to the topic of household production, including the role of time use, in Section 5.

4 Revealed preference analysis

We next introduce the RP characterization of rational household consumption behavior under the collective model introduced in Section 3. This characterization provides the basis for defining and recovering sharing rules. Because some primitives are typically unobserved, we also present an equivalent finite-dimensional integer programming (IP) formulation that is directly suited for empirical implementation. Specifically, this IP formulation enables us to test whether observed behavior is consistent with collective rationality and to identify the set of collective decision models, including sharing rules, that rationalize the observed choices. The logic mirrors the unitary case discussed in Section 2, where GARP and the Afriat inequalities provide equivalent tests and constructive representations of rational behavior; see Theorem 1.

4.1 Revealed preference characterization

To define the nonparametric RP condition for a collective rationalization of the data (as in Definition 4), we first introduce the concept of feasible personalized prices. These prices capture the share of the aggregate (market) price for publicly consumed goods that is borne

by each household member. For each public good, these personalized prices reflect the members' willingness to pay for the quantity consumed. Because these personalized prices sum to the observed market price of the public good, they can be interpreted as "Lindahl prices" characterizing the Pareto efficient provision of public goods.

Definition 5 (feasible personalized prices) *Let S be a set of observations. For each observation t and member $m \in \{1,2\}$, feasible personalized prices $\mathfrak{P}_t^m \in \mathbb{R}_+^K$ satisfy*

$$\mathfrak{P}_t^1 + \mathfrak{P}_t^2 = \mathbf{P}_t.$$

We can now adapt the GARP concept for the unitary model from Definition 2 to the collective setting with feasible personalized prices and quantities. As in the unitary case, GARP requires that observed choices are consistent with utility maximization under the relevant budget constraint. Here, the relevant budget constraint for each member aggregates private expenditure at market prices and public expenditure at personalized (Lindahl) prices.

Definition 6 (GARP with personalized prices and quantities) *Consider feasible personalized prices and quantities for a set of observations S. For member $m \in \{1,2\}$, the set $\{(\mathbf{p}_t,\, \mathfrak{P}_t^m;\, \mathfrak{q}_t^m, \mathbf{Q}_t)\,;\; t = 1,\ldots,T\}$ satisfies GARP if there exist relations R_0^m and R^m such that:*

(i) If $\mathbf{p}_s'\mathfrak{q}_s^m + (\mathfrak{P}_s^m)'\mathbf{Q}_s \geq \mathbf{p}_s'\mathfrak{q}_t^m + (\mathfrak{P}_s^m)'\mathbf{Q}_t$, then $(\mathfrak{q}_s^m, \mathbf{Q}_s)\; R_0^m\; (\mathfrak{q}_t^m, \mathbf{Q}_t)$.

(ii) If $(\mathfrak{q}_s^m, \mathbf{Q}_s)\; R_0^m\; (\mathfrak{q}_u^m, \mathbf{Q}_u)$, $(\mathfrak{q}_u^m, \mathbf{Q}_u)\; R_0^m\; (\mathfrak{q}_v^m, \mathbf{Q}_v), \ldots, (\mathfrak{q}_z^m, \mathbf{Q}_z)\; R_0^m\; (\mathfrak{q}_t^m, \mathbf{Q}_t)$ for some (possibly empty) sequence $(u, v, \ldots, z)$, then $(\mathfrak{q}_s^m, \mathbf{Q}_s)\; R^m\; (\mathfrak{q}_t^m, \mathbf{Q}_t)$.

(iii) If $(\mathfrak{q}_s^m, \mathbf{Q}_s)\; R^m\; (\mathfrak{q}_t^m, \mathbf{Q}_t)$, then

$$\mathbf{p}_t'\mathfrak{q}_t^m + (\mathfrak{P}_t^m)'\mathbf{Q}_t \leq \mathbf{p}_t'\mathfrak{q}_s^m + (\mathfrak{P}_t^m)'\mathbf{Q}_s.$$

The following proposition provides the RP characterization of collectively rational behavior; see Cherchye, De Rock, and Vermeulen (2011) for a detailed treatment. It mirrors the equivalences in Theorem 1 for the unitary model, with GARP and Afriat-type inequalities at the level of each household member under feasible personalized prices and quantities.

Proposition 1 *Let S be a set of observations. The following conditions are equivalent:*

(i) There exists a pair of utility functions U^1 and U^2 that provide a collective rationalization of S.

(ii) For each member $m \in \{1,2\}$, there exist feasible personalized prices and quantities such that the set $\{(\mathbf{p}_t, \mathfrak{P}_t^m; \mathfrak{q}_t^m, \mathbf{Q}_t)\,;\, t = 1,\ldots,T\}$ satisfies GARP.

(iii) For all $s,t \in \{1,\ldots,T\}$ and each member $m \in \{1,2\}$, there exist feasible personalized prices and quantities, and numbers $U_t^m > 0$ and $\lambda_t^m > 0$ that meet the Afriat inequalities

$$U_s^m - U_t^m \leq \lambda_t^m \Big[\big(\mathbf{p}_t' \mathfrak{q}_s^m + (\mathfrak{P}_t^m)' \mathbf{Q}_s\big) - \big(\mathbf{p}_t' \mathfrak{q}_t^m + (\mathfrak{P}_t^m)' \mathbf{Q}_t\big) \Big].$$

Condition (ii) states that collective rationality requires individual rationality (i.e., GARP consistency) of each member, evaluated at some feasible personalized prices and quantities. Condition (iii) gives the corresponding member-specific Afriat inequalities, paralleling condition (iii) in Theorem 1. In applications, the *true* personalized prices and personalized quantities are unobserved. Hence, the requirement is existential: there must exist *at least one* set of feasible personalized prices and feasible personalized quantities that satisfies the conditions.

In what follows, we will primarily work with condition (ii). The reason is that the Afriat inequalities in condition (iii) are nonlinear in the unknown utilities U_t^m, marginal utilities of income λ_t^m, and feasible personalized quantities $\mathfrak{q}_t^m$ and prices $\mathfrak{P}_t^m$, making them difficult to handle in empirical RP applications.

4.2 Sharing rule

Importantly for what follows, the result in Proposition 1 allows for a decentralized interpretation of collective rationality: rationality at the household level requires individual rationality at the member level. Given this, collectively rational consumption behavior can be represented as the outcome of a two-step allocation procedure. In the first step, a sharing rule distributes the aggregate household income across members. In the second step, each member maximizes her or his utility subject to the resulting income share and the member's personalized prices. This second step corresponds to the GARP requirement in condition (ii) in Proposition 1, which effectively implies that each household member behaves as a utility maximizer with a member-specific utility function. We emphasize that we do not assume that households *explicitly* implement a sharing rule. Rather, the two-step representation states that the outcome of the household allocation process can be *characterized* in this way.

Formally, we can restate the collective rationalization condition from Definition 4 as follows: a pair of utility functions (U^1, U^2) provides a collective rationalization of the data if, for each observation t, there exist feasible personalized quantities and prices such that,

for each member $m \in \{1, 2\}$,

$$U^m(\mathfrak{q}_t^m, \mathbf{Q}_t) \;\geq\; U^m(\mathfrak{z}^m, \mathbf{Z})$$

for all $\mathfrak{z}^m \in \mathbb{R}_+^N$ and $\mathbf{Z} \in \mathbb{R}_+^K$ that satisfy the member-specific budget constraint at personalized (Lindahl) prices,

$$\mathbf{p}_t'\mathfrak{z}^m + (\mathfrak{P}_t^m)'\mathbf{Z} \;\leq\; \mathbf{p}_t'\mathfrak{q}_t^m + (\mathfrak{P}_t^m)'\mathbf{Q}_t.$$

This member-specific utility maximization condition corresponds to the second step of the two-step representation outlined above. For the first step, the income share of member m in observation t is given by

$$\mathbf{p}_t'\mathfrak{q}_t^m + (\mathfrak{P}_t^m)'\mathbf{Q}_t,$$

which defines the sharing rule underlying the observed household consumption behavior. This motivates the following definition.

Definition 7 (feasible income shares) *Consider feasible personalized prices and quantities for a set of observations S such that, for each member $m \in \{1, 2\}$, the set $\{(\mathbf{p}_t, \mathfrak{P}_t^m; \mathfrak{q}_t^m, \mathbf{Q}_t); t = 1, \ldots, T\}$ satisfies GARP. Then the feasible income share for member $m \in \{1, 2\}$ at prices $(\mathbf{p}_t, \mathbf{P}_t)$ is*

$$\mathfrak{y}_t^m \;=\; \mathbf{p}_t'\mathfrak{q}_t^m + (\mathfrak{P}_t^m)'\mathbf{Q}_t.$$

This decentralized representation of collectively rational behavior follows from the Pareto efficiency assumption that underlies the within-household bargaining process and is formally related to the classic decentralization result for collective rationality without public goods (Chiappori, 1988, 1992). An important difference in our setting is that each member's preferences may depend not only on her or his private consumption but also on public consumption, which implies that personalized Lindahl prices may differ from observed market prices. Intuitively, the part of each member's income share associated with publicly consumed quantities reflects how well those public quantities align with the member's preferences for public goods (as captured by the Lindahl prices).

The sharing rule is central in this two-step representation and can be interpreted as an indicator of the members' bargaining power. Let $y_t = \mathbf{p}_t'\mathbf{q}_t + \mathbf{P}_t'\mathbf{Q}_t$ denote household income. A higher relative income share for member m, i.e., $\eta_t^m = \mathfrak{y}_t^m / y_t$, then signals greater bargaining power for that member; see, for instance, Browning, Chiappori, and Lewbel (2013). Conveniently, the sharing rule concept is independent of cardinal representations of preferences (in contrast to the Pareto weights in Definition 4), making it particularly useful in applications.

Still, one needs to be careful when drawing welfare conclusions from the sharing rule in the presence of public consumption. In the above definition, public consumption quantities are valued by individual Lindahl prices. These prices reflect individual marginal willingness to pay for the public goods, which themselves depend on the individuals' consumption bundles of private and public goods. Thus, an increase in a member's income share, combined with an increase in her Lindahl prices for public goods, does not necessarily imply that she is better off overall. In principle, the identification strategy described in the next section allows us to test specific hypotheses about the relationship between individuals' shares and their welfare levels. A complement to the above sharing rule that also explicitly accounts for public goods is the "conditional" sharing rule. It captures how the household shares income for private consumption conditional on the given level of public consumption. Unlike the sharing rule in Definition 7, it does not directly incorporate Lindahl prices for public goods. See, for example, Blundell, Chiappori, and Meghir (2005) and Chiappori and Meghir (2014) for discussion.

Whatever sharing-rule concept is used, a key question for applications is how to identify it from observed data. In what follows, we show that the Afriat-type characterization of collectively rational behavior in Proposition 1 provides a productive basis for addressing this identification question.

4.3 Integer programming characterization and identification

From the perspective of practical applications, conditions (ii) and (iii) in Proposition 1 are not directly useful because they are expressed in terms of unobserved personalized prices $\mathfrak{P}_t^m$ and personalized quantities $\mathfrak{q}_t^m$. First, regarding the GARP consistency requirement in (ii), note that given market prices $(\mathbf{p}_t, \mathbf{P}_t)$ and household quantities $(\mathbf{q}_t, \mathbf{Q}_t)$, there are infinitely many feasible specifications of the personalized quantities and prices. Any such feasible specification may induce distinct relations R_0^m and R^m in the GARP test, and it is practically impossible to verify all of these cases separately. Next, as argued above, the Afriat-type inequalities in condition (iii) are nonlinear in the unobserved objects U_t^m, λ_t^m, $\mathfrak{q}_t^m$ and $\mathfrak{P}_t^m$, which also complicates direct implementation.

Given this, Cherchye, De Rock, and Vermeulen (2011) provide an equivalent integer programming (IP) characterization of collectively rational consumption behavior by reformulating condition (ii) in IP terms. This formulation is amenable to empirical work and leverages modern solution methods tailored to integer programming; see, for example, Wolsey and Nemhauser (2014).[4] To obtain the IP representation, define binary variables $x_{st}^m \in \{0, 1\}$,

[4]Talla Nobibon, Cherchye, Crama, Demuynck, De Rock, and Spieksma (2016) provide an in-depth discus-

where $x^m_{st} = 1$ is interpreted as "$(\mathfrak{q}^m_s, \mathbf{Q}_s)$ R^m $(\mathfrak{q}^m_t, \mathbf{Q}_t)$" for a given feasible specification of personalized prices and quantities.

Proposition 2 *Let S be a set of observations. There exists a pair of utility functions U^1 and U^2 that provide a collective rationalization of S if and only if there exist $\mathfrak{P}^m_t \in \mathbb{R}^K_+$, $\mathfrak{q}^m_t \in \mathbb{R}^N_+$, $\mathfrak{y}^m_t \in \mathbb{R}_+$, and $x^m_{st} \in \{0,1\}$ such that for all $s,u,t \in \{1,\ldots,T\}$ and $m \in \{1,2\}$:*

(i) $\mathfrak{q}^1_t + \mathfrak{q}^2_t \;=\; \mathbf{q}_t \quad \textit{and} \quad \mathfrak{q}^m_t \;\geq\; \mathbf{q}^{Am}_t$,

(ii) $\mathfrak{P}^1_t + \mathfrak{P}^2_t \;=\; \mathbf{P}_t$,

(iii) $\mathfrak{y}^m_t \;=\; \mathbf{p}'_t\mathfrak{q}^m_t + (\mathfrak{P}^m_t)'\mathbf{Q}_t$,

(iv) $\mathfrak{y}^m_s - (\mathbf{p}'_s\mathfrak{q}^m_t + (\mathfrak{P}^m_s)'\mathbf{Q}_t) \;\leq\; y_s\, x^m_{st}$,

(v) $x^m_{su} + x^m_{ut} \;\leq\; 1 + x^m_{st}$,

(vi) $\mathfrak{y}^m_t - (\mathbf{p}'_t\mathfrak{q}^m_s + (\mathfrak{P}^m_t)'\mathbf{Q}_s) \;\leq\; y_t\,(1 - x^m_{st})$,

where $y_t \equiv \mathbf{p}'_t\mathbf{q}_t + \mathbf{P}'_t\mathbf{Q}_t$ denotes household income at observation t.

The conditions in this result can be interpreted as follows. Conditions (i)–(iii) specify feasible personalized prices, feasible personalized quantities, and feasible income shares, directly corresponding to Definitions 3, 5, and 7. Conditions (iv)–(vi) then impose, for each household member, the GARP requirements from Definition 6: condition (iv) defines the direct RP relations (i.e., if $(\mathfrak{q}^m_t, \mathbf{Q}_t)$ is affordable at $(\mathbf{p}_s, \mathfrak{P}^m_s)$ when evaluating $(\mathfrak{q}^m_s, \mathbf{Q}_s)$, then $x^m_{st} = 1$); condition (v) imposes transitivity of preferences; and condition (vi) implies GARP consistency (i.e., if $x^m_{st} = 1$, then $(\mathfrak{q}^m_t, \mathbf{Q}_t)$ cannot be strictly more expensive than $(\mathfrak{q}^m_s, \mathbf{Q}_s)$ at prices $(\mathbf{p}_t, \mathfrak{P}^m_t)$).[5]

Proposition 2 provides an operational necessary and sufficient test for collective rationality. If the IP conditions (i)–(vi) define an empty feasible region for a given data set, then no collective rationalization exists. Conversely, a non-empty feasible region implies that a collective rationalization does exist. Moreover, the same formulation yields a practical identification method: any feasible solution directly recovers member-specific personalized prices, personalized quantities, and income shares that rationalize the data. Upper and lower bounds on these unknowns can then be obtained by maximizing and minimizing them subject to the rationality conditions (i)–(vi). For instance, one can compute the maximal and

sion of computational aspects related to the IP characterization of collectively rational consumption behavior. These authors also present alternative procedures to enhance computational efficiency.

[5] In conditions (iv) and (vi), the constants y_s and y_t serve the standard "Big-M" function familiar from integer programming: they activate or deactivate the relevant GARP constraints depending on the value of the binary variable $x^m_{st} \in \{0,1\}$.

minimal value of the share $\mathfrak{y}_t^m$ that is compatible with the conditions. Conveniently, these bounds can be obtained using standard linear and integer programming tools. This procedure effectively set identifies the feasible values of the primitives that are consistent with a collective rationalization of the data. For a detailed discussion of this set identification approach, see Cherchye, De Rock, and Vermeulen (2011).

4.4 Illustration

We next illustrate the concepts introduced above and the set identification procedure. Specifically, we consider a stylized data set of prices and quantities that is rationalizable under the conditions of Proposition 2; for examples of non-rationalizable data sets, we refer the reader to Cherchye, De Rock, and Vermeulen (2007, 2012a). This example demonstrates that meaningful set identification is feasible even in a minimalistic setting with only two consumption observations. Naturally, richer data sets will generally yield tighter identification results.

Specifically, we consider a setting with two private goods ($N = 2$) and one public good ($K = 1$). The first private good is not assignable, while the second private good is fully assigned to the two members. We observe two household choices:

t	$\mathbf{p}_t$	$\mathbf{P}_t$	$\mathbf{q}_t$	$\mathbf{q}_t^{A1}$	$\mathbf{q}_t^{A2}$	$\mathbf{Q}_t$	y_t
1	$(4,1)'$	5	$(6,2)'$	$(0,1)'$	$(0,1)'$	3	41
2	$(3,7)'$	6	$(3,8)'$	$(0,5)'$	$(0,3)'$	2	77

.

In the first observation, household income equals $y_1 = \mathbf{p}_1'\mathbf{q}_1 + \mathbf{P}_1'\mathbf{Q}_1 = 4\times6+1\times2+5\times3 = 41$, while in the second observation it equals $y_2 = 3 \times 3 + 7 \times 8 + 6 \times 2 = 77$.

Feasible personalized prices, quantities and income shares. Because the second private good is fully assigned, condition (i) implies

$$\mathfrak{q}_{1,2}^1 = \mathfrak{q}_{1,2}^2 = 1, \qquad \mathfrak{q}_{2,2}^1 = 5, \qquad \mathfrak{q}_{2,2}^2 = 3,$$

for every feasible specification of personalized quantities. By contrast, the first private good can be freely split across the two members, subject to adding up. One feasible specification of the unobserved personalized quantities, personalized prices and income shares is therefore

t	$\mathfrak{q}_t^1$	$\mathfrak{q}_t^2$	$\mathfrak{P}_t^1$	$\mathfrak{P}_t^2$	$\mathfrak{y}_t^1$	$\mathfrak{y}_t^2$	η_t^1	η_t^2
1	$(2,1)'$	$(4,1)'$	5	0	24	17	24/41	17/41
2	$(3,5)'$	$(0,3)'$	0	6	44	33	44/77	33/77

.

This table directly illustrates conditions (i)–(iii). First, the personalized quantities add up to the observed household quantities and respect the assignable lower bounds. Second, the personalized prices for the public good add up to the observed market price. Third, the income shares are obtained by pricing each member's private consumption at market prices and public consumption at personalized prices. For example, for the first observation, we have

$$\mathfrak{y}_1^1 = \mathbf{p}_1'\mathfrak{q}_1^1 + (\mathfrak{P}_1^1)'\mathbf{Q}_1 = 4 \times 2 + 1 \times 1 + 5 \times 3 = 24,$$

and

$$\mathfrak{y}_1^2 = \mathbf{p}_1'\mathfrak{q}_1^2 + (\mathfrak{P}_1^2)'\mathbf{Q}_1 = 4 \times 4 + 1 \times 1 + 0 \times 3 = 17.$$

Together, these two income shares add up to household income in the first observation (i.e., $y_1 = 41$).

Rationalizability. The given specification of feasible personalized quantities and prices satisfies the rationalizability conditions in Proposition 2. To see this, we start by noting that it gives rise to the following cross-budget comparisons:

m	$\mathfrak{y}_s^m$	$\mathbf{p}_s'\mathfrak{q}_t^m + (\mathfrak{P}_s^m)'\mathbf{Q}_t$	x_{st}^m
1	24	27	$x_{12}^1 = 0$
1	44	13	$x_{21}^1 = 1$
2	17	3	$x_{12}^2 = 1$
2	33	37	$x_{21}^2 = 0$

.

m	$\mathfrak{y}_s^m$	$\mathbf{p}_s'\mathfrak{q}_t^m + (\mathfrak{P}_s^m)'\mathbf{Q}_t$	x_{st}^m
1	24	27	$x_{12}^1 = 0$
1	44	13	$x_{21}^1 = 1$
2	17	3	$x_{12}^2 = 1$
2	33	37	$x_{21}^2 = 0$

.

Let us explain this in detail for member 1. The personalized bundle of observation 1 is affordable at the personalized prices and income share of observation 2:

$$\mathbf{p}_2'\mathfrak{q}_1^1 + (\mathfrak{P}_2^1)'\mathbf{Q}_1 = 13 < 44 = \mathfrak{y}_2^1.$$

Hence, observation 2 is revealed preferred to observation 1 for member 1, and condition (iv) requires $x_{21}^1 = 1$. This implies that the corresponding condition (vi) becomes potentially binding, since its right-hand side is zero. Conversely, the personalized bundle of observation

2 is not affordable at the personalized prices and income share of observation 1:

$$\mathbf{p}_1'\mathfrak{q}_2^1 + (\mathfrak{P}_1^1)'\mathbf{Q}_2 = 27 > 24 = \mathfrak{y}_1^1.$$

Therefore, no revealed-preference relation from observation 1 to observation 2 is forced for member 1. Condition (iv) does not require $x_{12}^1 = 1$, and setting $x_{12}^1 = 0$ makes the corresponding condition (vi) non-binding. With only two observations, the transitivity condition (v) has no additional empirical content.

The same interpretation applies to member 2: observation 1 is revealed preferred to observation 2 because the observation 2 bundle is affordable at observation 1 prices and income share, while the reverse comparison is not forced.

Together with the trivial choices $x_{11}^1 = x_{22}^1 = x_{11}^2 = x_{22}^2 = 1$, the specified binary variables satisfy the revealed-preference restrictions in Proposition 2. For neither member do we obtain a cycle in which two different observations are revealed preferred to each other. The specified personalized quantities, personalized prices and income shares therefore define member-specific choices that are consistent with stable individual preferences.

Set identifying the sharing rule. We can also use the example to illustrate why Proposition 2 obtains bounds on the sharing rule. First, based on the assignable quantity information alone, member 1 must consume at least one unit of the second private good in the first observation. Member 2 must also consume at least one unit of that good. Assignable information alone therefore gives

$$\mathfrak{y}_1^1 \in \left[\mathbf{p}_1'\mathbf{q}_1^{A1},\ y_1 - \mathbf{p}_1'\mathbf{q}_1^{A2}\right] = [1, 40],$$

or equivalently, for the relative sharing rule, $\eta_1^1 \in [1/41, 40/41]$.

The RP restrictions tighten these bounds because they link the two observations through stable individual preferences. For instance, for member 1, any feasible specification satisfies

$$\begin{aligned}
&\mathfrak{y}_2^1 - \left(\mathbf{p}_2'\mathfrak{q}_1^1 + (\mathfrak{P}_2^1)'\mathbf{Q}_1\right) \\
&= 3(\mathfrak{q}_{2,1}^1 - \mathfrak{q}_{1,1}^1) + 7(\mathfrak{q}_{2,2}^1 - \mathfrak{q}_{1,2}^1) + \mathfrak{P}_2^1(Q_2 - Q_1) \\
&\geq 3(0-6) + 7(5-1) + 6(2-3) = 4 > 0.
\end{aligned}$$

Thus, condition (iv) implies that $x_{21}^1 = 1$ must hold in every feasible solution. Condition

(vi) then implies

$$\begin{aligned}\mathfrak{y}_1^1 &\leq \mathbf{p}_1'\mathfrak{q}_2^1 + (\mathfrak{P}_1^1)'\mathbf{Q}_2 \\ &= 4\mathfrak{q}_{2,1}^1 + \mathfrak{q}_{2,2}^1 + 2\mathfrak{P}_1^1 \\ &\leq 4 \times 3 + 5 + 5 \times 2 \\ &= 27.\end{aligned}$$

Hence, for this stylized data set, the lower bound remains the assignability lower bound, but the IP characterization yields a strictly smaller upper bound:

$$1 \leq \mathfrak{y}_1^1 \leq 27 \text{ or } \eta_1^1 \in \left[\frac{1}{41}, \frac{27}{41}\right].$$

Note that our illustration is based on only two observations. In general, additional observations, involving different consumption bundles under different prices and income levels, will yield tighter bounds by exploiting a larger set of revealed-preference restrictions. As such, the example should be viewed primarily as an illustration of the identification logic. In Section 6, we show that the same RP-based approach can also be fruitfully applied to real-world data, where richer datasets yield empirically informative bounds on intrahousehold resource allocation and individual well-being.

5 Extensions

We conclude our methodological discussion by outlining several extensions of the RP methodology described above. Our goal is to provide a concise and non-technical overview; for formal details and proofs, we refer to the cited papers. These extensions are modular and can be combined in practical applications, allowing researchers to design tools tailored to specific empirical questions. This flexibility highlights the broad applicability of RP methods for analyzing collective consumption behavior and enables rich, data-driven empirical analyses.

5.1 Impure goods

In the previous section, we explicitly distinguished between public goods (captured by $\mathbf{Q}_t$) and private goods (captured by $\mathbf{q}_t$), an approach commonly adopted in empirical applications. In many practical settings, however, some goods simultaneously serve private and public purposes, giving rise to "impure" goods. For example, "car use" may include a public component (e.g., a family outing) as well as a private component (e.g., one member's commute). If the analyst observes the private and public components of each good separately, then the same methodology as before can be applied. In that case, the privately consumed

component and the publicly consumed component of each impure good are formally treated as distinct goods.

When the private and public components of impure goods are not separately observed, Browning and Chiappori (1998) propose a collective consumption model that does not require an a priori classification of goods as public or private. Cherchye, De Rock, and Vermeulen (2007) develop the revealed-preference characterization of this model, deriving a testable necessary condition for data consistency as well as a complementary sufficient condition; these two conditions generally do not coincide. Building on this, Cherchye, De Rock, Sabbe, and Vermeulen (2008) formulate an integer-programming representation of the necessary condition in Cherchye, De Rock, and Vermeulen (2007), which is formally similar to the IP formulation in Proposition 2, and they demonstrate its practical implementation through an empirical application. Cherchye, De Rock, and Vermeulen (2011) show that this IP characterization provides a practical basis for (set) identifying sharing rules without requiring the analyst to specify ex ante the public or private nature of goods.

5.2 Demand systems and RP restrictions

It is also possible to integrate RP restrictions for rational consumption behavior with estimated demand systems to deliver sharp bounds on objects of interest that are jointly disciplined by sample information and RP inequalities. In this way, this extension provides a remedy for the lack of discriminatory power that is sometimes found in RP applications. Blundell, Browning, and Crawford (2003) and Blundell, Browning, and Crawford (2008) pioneered this approach for the unitary model, combining nonparametrically estimated Engel curves and RP constraints to yield informative counterfactual demand bounds. The key idea is that RP restrictions rule out demand reactions that violate rationality, while the detailed structure of estimated demand systems further tightens the set of feasible responses. The resulting bounds provide a sharply defined and maximally informative region for counterfactual demand behavior consistent with rational demand theory.

Building on this logic for the collective framework, Cherchye, De Rock, Lewbel, and Vermeulen (2015) develop a methodology to identify sharp sharing rule bounds that exploit both RP restrictions for collective rationality (in the spirit of Proposition 1) and richer information provided by considering household demand as a function of prices and household income. Subsequently, Cherchye, Cosaert, De Rock, Kerstens, and Vermeulen (2018) extend this methodology to derive sharp bounds on money-metric welfare measures for individual household members.[6] For empirical applications, the practical upshot is that researchers can

[6]Money-metric welfare measures refer to indicators of well-being that express a person's utility in monetary terms (i.e., the amount of money needed to reach a certain utility level). They allow comparisons

combine standard demand estimation with RP methods to recover tight, data-driven bounds on sharing rules and welfare metrics without relying on parametric utility assumptions.

5.3 Household consumption and marriage markets

Given the empirical observation that bargaining positions are linked to outside options, there is an active strand of the literature that explicitly links intrahousehold allocations to individuals' marriage markets and marital stability. Cherchye, Demuynck, De Rock, and Vermeulen (2017) and Browning, Cherchye, Demuynck, De Rock, and Vermeulen (2024) incorporate marriage market equilibrium conditions into collective models to connect sharing rules with observed sorting and match surplus, thereby tying within-household bargaining outcomes to marital matching dynamics. Marital stability refers to the persistence of a marriage when neither spouse has an incentive to dissolve the union.[7] Research linking intrahousehold allocations to marriage markets stresses that the distribution of resources within the household influences these incentives and thus shapes the stability of the match. This approach structurally relates marital stability to collective household decisions and provides conditions under which observed consumption bundles, prices, and matching patterns jointly restrict feasible intrahousehold sharing patterns.

Methodologically, these studies embed collective rationality within a broader equilibrium framework of matching and stability, thereby introducing additional observables that sharpen the identification of sharing rules and individual preferences. This framework also allows for a structural analysis of how features of individuals' marriage markets, which are typically treated as "distribution factors" in the literature on collective consumption models (Bourguignon, Browning, and Chiappori, 2009), influence within-household consumption allocations. Examples of such marriage-market characteristics include individuals' relative earning potential, education level, age or age differences, local sex ratios, and personal wealth. These factors can influence each person's outside options in their current marriage and therefore function as distribution factors in collective consumption models. Empirically, this approach enables the joint use of consumption and marriage market data to uncover richer aspects of intrahousehold allocation and its determinants.

of individual welfare using a common, money-based scale. See Chiappori and Meghir (2014) for a general discussion of money-metric measures of individual welfare in the context of collective consumption models.

[7]More formally, marital stability requires that the existing match satisfies two conditions: individual rationality and no blocking pairs. Individual rationality means that each spouse is at least as well off within the marriage as they would be on their own, so neither has an incentive to leave unilaterally. The no-blocking-pairs condition means that there is no alternative partner that either spouse would strictly prefer over their current match and who would also prefer them in return. When both conditions hold, no mutually preferable alternative match exists, and the marriage is considered stable. See Gale and Shapley (1962), Shapley and Shubik (1972) and Becker (1973) for seminal contributions.

5.4 Time use and household production

Collective models with time use and household production relax the purely consumptive interpretation of choices by allowing inputs (including both time use and market goods) to be transformed into produced commodities that enter members' utilities. Building on this idea, Cherchye, De Rock, Surana, and Vermeulen (2020) and Cherchye, De Rock, Vermeulen, and Walther (2021) incorporate intrahousehold production technologies alongside consumption, thereby extending identification to settings in which observed market purchases and time allocations jointly determine the production of both public and private commodities. A distinctive feature of this work is that it models how household investments of time and money generate public goods consumed jointly within the household, as well as private goods that may be partly consumed internally and potentially also sold outside the household.

This enriched framework broadens the scope of collective models by linking consumption, time allocation, and household production in a unified RP approach. Methodologically, this adapts RP logic to joint production–consumption settings while preserving testable implications analogous to Proposition 1, now incorporating both technological constraints and intrahousehold allocation rules. For empirical practice, this widens the applicability of RP-based collective analysis to data sets containing time-use measures or home-production proxies, enabling researchers to study how domestic production activities interact with market consumption and shape within-household resource allocations and individual welfare.

5.5 Intertemporal consumption and commitment

The static collective model studies Pareto efficient allocations period by period, whereas intertemporal collective models ask how households allocate consumption, labor supply, and savings when members face uncertainty over the life cycle. The central issue is the degree of commitment within the household. Under full commitment, spouses can commit ex ante to a state-contingent allocation rule. Unexpected wage, income, preference, or demographic shocks then affect choices through the budget and state variables, but they do not by themselves trigger a renegotiation of Pareto weights. Under limited commitment, by contrast, the allocation must remain acceptable to each spouse after shocks are realized. If a shock improves one spouse's outside option or weakens the other spouse's bargaining position, the relevant participation constraint may bind, leading to a renegotiation of Pareto weights and, therefore, of the sharing rule. Mazzocco (2007) provides an early collective characterization and empirical test of commitment in an intertemporal setting, and Chiappori and Mazzocco (2017) surveys the broader static and dynamic collective literature. More recently, Theloudis et al. (2024) develop a lifecycle collective labor supply framework that distinguishes full, lim-

ited, and no commitment by comparing the effects of contemporaneous and historical news on household behavior.

Adams, Cherchye, De Rock, and Verriest (2014) extend the above approach by formulating RP tests of intertemporal collective rationality under varying commitment regimes. Their analysis adapts RP-style restrictions to time-series consumption data and establishes conditions under which observed allocation sequences are consistent with dynamic collective behavior. Methodologically, this line of work generalizes static RP analysis to dynamic environments by incorporating intertemporal budget constraints and Euler-type optimality conditions into the collective framework. It thus allows rationality requirements and commitment assumptions to operate jointly over time. Empirically, it provides nonparametric tools for evaluating dynamic household consumption choices in panel or time-series data. This enables researchers to determine whether observed consumption paths can be reconciled with shifting intrahousehold bargaining positions and to examine how resource changes or economic shocks shape the intertemporal distribution of consumption and individual welfare.

6 Empirical illustration

We end this chapter with a specific illustration of what the RP methodology can deliver. The purpose is not to provide a full replication of the empirical analysis, but to show how the theoretical objects introduced above become useful empirical objects. The discussion is based on Cherchye, De Rock, Lewbel, and Vermeulen (2015), who combine estimated household demand functions with RP restrictions to obtain bounds on the intrahousehold sharing rule for a general collective model with impure goods (as discussed in Sections 5.1 and 5.2).

6.1 Data and empirical set-up

The application uses data from the 1999–2009 Panel Study of Income Dynamics (PSID). The sample consists of couples in which both spouses are between 25 and 65 and participate in the labor market. The analysis is conducted separately for childless couples and for couples with children. After removing households with missing information and a small number of extreme observations, the final samples contain 865 childless couples and 1,390 couples with children.

The empirical model is a joint consumption–labor supply model. Household full income is allocated to the leisure of both spouses and to three categories of nondurable consumption: food, housing, and other goods. Full income is defined as the sum of the spouses' maximum

possible labor income and nonlabor income, excluding savings and expenditures on durables. Wages are used as prices of leisure. This is important because wages are not treated merely as distribution factors: they enter the budget constraint and may also affect the Pareto weights. The analysis therefore explicitly allows for nonseparability between leisure and nondurable consumption.

This empirical set-up builds on the collective consumption framework developed in the previous sections. The household-level demand function is observed, while the individual quantities, personalized prices, and income shares are not. Male and female leisure are treated as private and non-assignable commodities that do not entail externalities. No restriction is imposed on the three nondurable goods: food, housing, and other goods may be public, private, affected by externalities, or mixtures of these possibilities. As such, the application is consistent with the general collective framework discussed above, in which the public/private nature of nondurable consumption need not be specified ex ante.

The empirical implementation compares different ways of estimating household demand: a nonparametric demand system, a QUAIDS specification, a QUAIDS specification that imposes the SR1 restriction of Browning and Chiappori (1998), and a QUAIDS specification with SR1 and observable taste shifters. In the analysis, the sharing rule bounds of each specification are compared to each other. In addition, these bounds are compared to 'naive' bounds. The latter are defined as follows: the naive lower bound of a spouse is defined as the bound associated with assigning all the household's nondurable consumption to the other spouse. This implies that the first spouse only consumes his or her leisure. The naive upper bound of a spouse contains not only the value of that spouse's leisure but also all the household's nondurable consumption. Such a comparison is informative because it shows where the empirical content comes from.

In the empirical analysis, the RP restrictions turn out to substantially tighten the naive bounds. Additional observable heterogeneity in the demand estimation step tightens the bounds further. For example, for childless couples, the average difference between the upper and lower bounds on the female income share falls from 17.52 percentage points for the naive bounds to 2.91 percentage points for the RP bounds based on QUAIDS with SR1 and taste shifters. More than 75 percent of childless couples have a difference between upper and lower bounds below about 4.5 percentage points under this last specification.

Given these findings, the discussion below focuses on the results obtained under the QUAIDS specification that incorporates observable taste shifters and imposes the SR1 restriction. Importantly, however, the main qualitative conclusions are similar across the alternative RP-restricted specifications. As we show, the resulting bounds are sufficiently informative to support meaningful conclusions regarding both intrahousehold resource al-

location and individual well-being despite the absence of point identification. Section 6.2 examines the implications for intrahousehold resource allocation, while Section 6.3 considers the corresponding implications for individual welfare, focusing in particular on individual poverty and associated gender differences in poverty.

6.2 Sharing-rule bounds: income and relative wages

These reported averages suggest that the absence of point identification need not preclude empirically relevant conclusions. As a first illustration, consider Figure 1, where the lower and upper bounds on the relative female income share are plotted against the logarithm of full household income, separately for childless couples and couples with children. In each figure, the blue dots give household-specific upper bounds, the red crosses give household-specific lower bounds, and the thick black lines give local averages.

While there is substantial heterogeneity across households, the bounds are visually close to each other for many households. Moreover, the trend lines are slightly downward sloping, but close to horizontal.

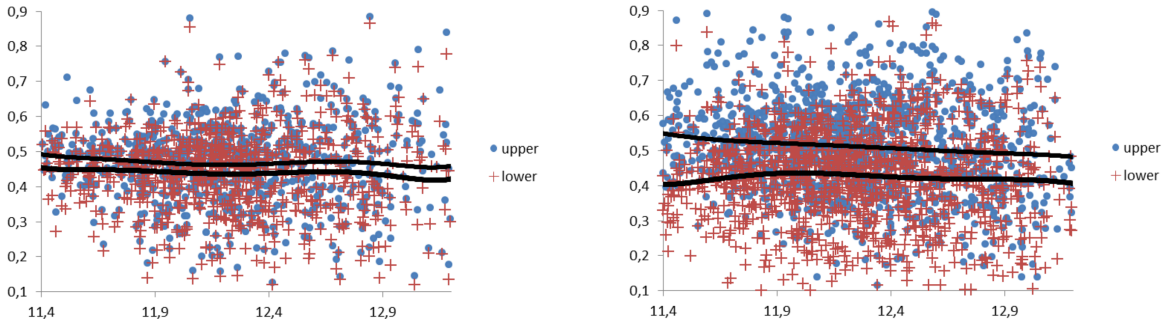

Figure 1: Relative sharing rule bounds and the logarithm of full income – childless couples (on the left) and couples with children (on the right). Source: Cherchye, De Rock, Lewbel and Vermeulen (2015, published Figures 3 and 4).

We next relate the same female income-share bounds to the relative wage, defined as the female wage divided by the male wage. This connects more directly to the bargaining interpretation of the sharing rule. Figure 2 shows a clear positive association: the relative female income share increases when the female wage rises relative to the male wage. This observation is consistent with the standard collective-model interpretation that a higher relative wage is associated with a stronger female bargaining position. Relative wages measure potential earnings outside marriage and can therefore proxy for outside options, which in turn may affect bargaining.

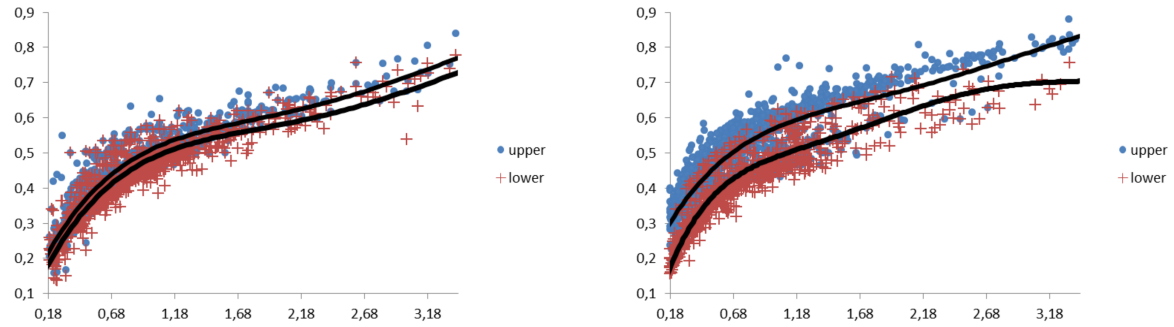

Figure 2: Relative sharing rule bounds and the relative wage – childless couples (on the left) and couples with children (on the right). Source: Cherchye, De Rock, Lewbel and Vermeulen (2015, published Figures 5 and 6).

Taken together, the two figures show that the recovered income-share bounds are sufficiently informative to uncover systematic relationships between intrahousehold resource allocation, household income, and relative wages. Beyond their interpretation as indicators of bargaining power, these set identified income shares can also be used to construct individual-level welfare measures, as we illustrate next.

6.3 Individual poverty

Because the sharing rule determines the resources available to each household member, the identified bounds on income shares immediately translate into bounds on individual welfare measures. Standard poverty measures are usually computed at the household level. This implicitly assumes that household resources are shared equally among members, which makes household income a sufficient indicator of each member's welfare. The collective approach relaxes this assumption by using the sharing-rule bounds to compute poverty rates at the individual level.

Cherchye, De Rock, Lewbel, and Vermeulen (2015) define the household poverty line as 60 percent of the median full income in the relevant sample. An individual is then classified as poor if her or his income share falls below half of this household poverty line. Since the sharing rule is set identified rather than point-identified, this procedure yields lower and upper bounds on individual poverty rates.

Table 1 reports the main results. For childless couples, the standard household poverty rate is 11.3 percent, while the individual poverty rate lies between 15.3 percent and 18.6 percent. For couples with children, the standard household poverty rate is 12.2 percent, while the individual poverty rate lies between 12.7 percent and 24.8 percent. In both samples, the individual-level poverty bounds are higher than the household poverty rate. This shows

that household-level poverty measures may understate poverty when resources are unequally shared within households.

Table 1: Poverty rates based on household income and sharing-rule bounds. Source: Cherchye, De Rock, Lewbel and Vermeulen (2015).

Sample	Households	All individuals	Females	Males
Childless couples	11.3%	15.3–18.6%	20.7–24.7%	12.3–12.4%
Couples with children	12.2%	12.7–24.8%	15.0–29.5%	10.3–20.2%

The gender-specific results are also informative. For childless couples, the female poverty rate is bounded between 20.7 percent and 24.7 percent, while the male poverty rate is bounded between 12.3 percent and 12.4 percent. For couples with children, the female poverty rate is bounded between 15.0 percent and 29.5 percent, while the male poverty rate is bounded between 10.3 percent and 20.2 percent. Overall, the results suggest that poverty within couples is more prevalent among women than among men in these data. This again shows the importance of an individual welfare analysis instead of a household-level analysis.

7 Conclusion

This chapter reviewed the nonparametric RP approach to analyzing collective household consumption behavior. The collective model generalizes the unitary model by assuming that each spouse in a couple is a utility-maximizing decision maker, with own preferences, who interacts with the other spouse to jointly choose the household's commodity bundle that is assumed to be Pareto efficient.

A first key ingredient of the RP framework is the notion of feasible personalized quantities for private goods and personalized (Lindahl) prices for public goods. These objects allow observed household choices to be represented as Pareto efficient allocations without imposing strong functional-form assumptions regarding individual utilities and household demand. A second essential component is the sharing rule, which summarizes the allocation of resources within the household and serves as an indicator of members' bargaining power and individual welfare. Unlike Pareto weights, this object is robust to cardinal rescaling of individual utilities.

Methodologically, the chapter presented an operational framework for nonparametric identification of personalized prices, quantities, and the sharing rule. At its core lies the RP characterization of the collective model in the spirit of the Afriat Theorem. This characterization provides testable restrictions that can be implemented using standard linear and integer programming techniques. Any feasible solution to these restrictions recovers person-

alized quantities, personalized prices, and intrahousehold sharing rules that rationalize the data. This enables nonparametric set identification of these unknowns from observed household consumption choices. The methodology also extends to settings with impure goods and to richer empirical environments that incorporate demand estimation, marriage market constraints, time use and household production, and intertemporal consumption considerations. The empirical illustration showed how these set identification results can be used in practice. In particular, sharing-rule bounds can be related to household income and relative wages, and they can be translated into individual-level poverty measures. This demonstrates that the RP approach is not only a testing device, but also a tool for welfare and policy analysis.

Where do we stand now? The RP approach to the collective model has evolved into a mature and versatile empirical framework. Researchers can test collective rationality using finite data sets, recover economically meaningful objects without parametric utility assumptions, and apply the framework to contexts that mix public and private goods or require additional structure, such as dynamic or market equilibrium elements. The tools are modular. They can be combined with econometric methods for demand estimation and scaled using off-the-shelf optimization solvers. As a result, the approach remains close to the data while delivering the interpretability needed for policy and welfare analysis based on the collective consumption model.

The next step is to establish this methodology as a standard approach for addressing empirical questions based on household data. The growing availability of granular microdata on household consumption and time use behavior makes this increasingly feasible. It should allow researchers to generate informative evidence on the decision-making processes that shape the distribution of welfare within households.

References

Adams, Abi, Laurens Cherchye, Bram De Rock, and Ewout Verriest. 2014. "Consume now or later? Time inconsistency, collective choice, and revealed preference." *American Economic Review* 104 (12): 4147–4183.

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

Becker, Gary S. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy* 81 (4) 813–846.

Blundell, Richard, Martin Browning, and Ian Crawford. 2003. "Nonparametric Engel curves and revealed preferences." *Econometrica* 71 (1): 205–240.

Blundell, Richard, Martin Browning, and Ian Crawford. 2008. "Best nonparametric bounds on demand responses." *Econometrica* 76 (6): 1227–1262.

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir. 2005. "Collective labor supply with children." *Journal of Political Economy* 113 (6): 1277–1306.

Bourguignon, François, Martin Browning, and Pierre-André Chiappori. 2009. "Efficient intra-household allocations and distribution factors: Implications and identification." *The Review of Economic Studies* 76 (2): 503–528.

Bronars, Stephen. 1987. "The power of nonparametric tests of preference maximization." *Econometrica* 55 (3): 693–698.

Browning, Martin, Laurens Cherchye, Thomas Demuynck, Bram De Rock, and Frederic Vermeulen. 2024. "Spouses with benefits: on match quality and consumption inside households." *KU Leuven Department of Economics Discussion Paper* DPS 24.09.

Browning, Martin, Pierre-André Chiappori, and Valérie Lechene. 2010. "Distributional effects in household models: Separate spheres and income pooling." *The Economic Journal* 120 (545): 786–799.

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

Browning, Martin, Pierre-André Chiappori, and Arthur Lewbel. 2013. "Estimating consumption economies of scale, adult equivalence scales, and household bargaining power." *Review of Economic Studies* 80 (4): 1267–1303.

Cherchye, Laurens, Sam Cosaert, Bram De Rock, Pieter Jan Kerstens, and Frederic Vermeulen. 2018. "Individual welfare analysis for collective households." *Journal of Public Economics* 166 98–114.

Cherchye, Laurens, Sam Cosaert, Thomas Demuynck, and Bram De Rock. 2020. "Group consumption with caring individuals." *The Economic Journal* 130 (627): 587–622.

Cherchye, Laurens, Bram De Rock, Arthur Lewbel, and Frederic Vermeulen. 2015. "Sharing rule identification for general collective consumption models." *Econometrica* 83 (5): 2001–2041.

Cherchye, Laurens, Bram De Rock, Jeroen Sabbe, and Frederic Vermeulen. 2008. "Nonparametric tests of collectively rational consumption behavior: an integer programming procedure." *Journal of Econometrics* 147 (2): 258–265.

Cherchye, Laurens, Bram De Rock, Khushboo Surana, and Frederic Vermeulen. 2020. "Marital matching, economies of scale and intrahousehold allocations." *Review of Economics and Statistics* 102 (4): 823–837.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2007. "The collective model of household consumption: a nonparametric characterization." *Econometrica* 75 (2): 553–574.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2010. "An Afriat theorem for the collective model of household consumption." *Journal of Economic Theory* 145 (3): 1142–1163.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2011. "The revealed preference approach to collective consumption behavior." *Review of Economic Studies* 78 (1): 176–198.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2012a. "Collective household consumption behavior: Revealed preference analysis." *Foundations and Trends in Econometrics* 4 (4): 225–312.

Cherchye, Laurens, Bram De Rock, Frederic Vermeulen, and Selma Walther. 2021. "Where did it go wrong? Marriage and divorce in Malawi." *Quantitative Economics* 12 (2): 505–545.

Cherchye, Laurens, Thomas Demuynck, Bram De Rock, and Frederic Vermeulen. 2017. "Household consumption when the marriage is stable." *American Economic Review* 107 (6): 1507–1534.

Cherchye, Laurens, Thomas Demuynck, and Bram De Rock. 2011. "Revealed preference analysis of non-cooperative household consumption." *The Economic Journal* 121 (555): 1073–1096.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2012b. "Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information." *American Economic Review* 102 (7): 3377–3405.

Chiappori, Pierre-André, and Ivar Ekeland. 2006. "The micro economics of group behavior: General characterization." *Journal of Economic Theory* 130 (1): 1–26.

Chiappori, Pierre-André, and Maurizio Mazzocco. 2017. "Static and Intertemporal Household Decisions." *Journal of Economic Literature* 55 (3): 985–1045. 10.1257/jel.20150715.

Chiappori, Pierre-André, and Costas Meghir. 2014. "Intra-household welfare."Technical report, Cowles Foundation Discussion Paper No. 1949.

Chiappori, Pierre-André. 1988. "Rational household labor supply." *Econometrica* 56 (1): 63–89.

Chiappori, Pierre-André. 1992. "Collective labor supply and welfare." *Journal of Political Economy* 100 (3): 437–467.

Chiappori, Pierre-André, and Ivar Ekeland. 2009. "The micro economics of efficient group behavior: Identification." *Econometrica* 77 (3): 763–799.

Dauphin, Anyck, Abdel-Rahmen El Lahga, Bernard Fortin, and Guy Lacroix. 2011. "Are children decision-makers within the household?" *The Economic Journal* 121 (553): 871–903.

Diewert, Erwin. 1973. "Afriat and revealed preference theory." *Review of Economic Studies* 40 (3): 419–425.

d'Aspremont, Claude, and Rodolphe Dos Santos Ferreira. 2014. "Household behavior and individual autonomy: an extended Lindahl mechanism." *Economic Theory* 55 (3): 643–664.

Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly* 69 (1) 9–15.

Houthakker, Hendrik. 1950. "Revealed preference and the utility function." *Economica* 17 (66): 159–174.

Lechene, Valérie, and Ian Preston. 2011. "Noncooperative household demand." *Journal of Economic Theory* 146 (2): 504–527.

Mazzocco, Maurizio. 2007. "Household intertemporal behaviour: A collective characterization and a test of commitment." *The Review of Economic Studies* 74 (3): 857–895.

Samuelson, Paul A. 1938. "A note on the pure theory of consumer's behavior." *Economica* 5 (17): 61–71.

Samuelson, Paul A. 1948. "Consumption theory in terms of revealed preference." *Economica* 15 (60): 243–253.

Shapley, Lloyd, and Martin Shubik. 1972. "The Assignment Game I: The Core." *International Journal of Game Theory* 1 111–130.

Talla Nobibon, Fabrice, Laurens Cherchye, Yves Crama, Thomas Demuynck, Bram De Rock, and Frits CR Spieksma. 2016. "Revealed preference tests of collectively rational consumption behavior: formulations and algorithms." *Operations Research* 64 (6): 1197–1216.

Theloudis, Alexandros, Jorge Velilla, Pierre-André Chiappori, J. Ignacio Giménez-Nadal, and José Alberto Molina. 2024. "Commitment and the Dynamics of Household Labor Supply."Technical report. 10.48550/arXiv.2307.10983, Version 2, February 2024.

Varian, Hal R. 1982. "The nonparametric approach to demand analysis." *Econometrica* 50 (4): 945–972.

Warshall, Stephen. 1962. "A theorem on boolean matrices." *Journal of the ACM (JACM)* 9 (1): 11–12.

Wolsey, LA, and GL Nemhauser. 2014. *Integer and Combinatorial Optimization (John-Wiley & Sons, New York).*