NET WORK LISER

Discussion Paper Series

IZA DP No. 18941

September 2026

Individual-Level Inequality

Laurens Cherchye
KU Leuven

Bram De Rock
ECARES, Université Libre de Bruxelles
and KU Leuven

Frederic Vermeulen
KU Leuven
and IZA@LISER

IZA DP No. 18941 September 2026

Individual-Level Inequality*

Abstract

Standard household-level data hide substantial heterogeneity in how consumption, money, and time are allocated within households. This motivates the need for measuring welfare at the individual level and a structural approach on how household data can be used to capture these differences. After documenting what can be learned from direct survey evidence, we present the collective approach as an attractive framework for modeling intra-household decision-making with preference heterogeneity, public consumption, and home production. We then survey three identification strategies that recover individual preferences and resource shares from widely available datasets: differential, revealed preference, and preference-similarity methods. Applying these tools reveals that a sizable share of overall inequality originates within households and that individual poverty rates often diverge sharply from household-level measures. Overall, our review demonstrates that measuring welfare at the individual level fundamentally reshapes our understanding of inequality, and offers practical guidance for incorporating intra-household heterogeneity into empirical welfare evaluations.

JEL classification

D13, D31, I32, J22

Keywords

individual welfare measurement, intra-household inequality, collective household consumption, sharing rule, time use

Corresponding author

Frederic Vermeulen
frederic.vermeulen@kuleuven.be

* We thank Editor David Romer and four anonymous referees for their insightful comments and suggestions, which substantially improved this paper.

Laurens Cherchye gratefully acknowledges the Fund for Scientific Research - Flanders (FWO) and the Research Fund KU Leuven for financial support. He is also thankful to the Van Eeckhoudt-Ruttens family for supporting the 600 years of KU Leuven Chair on "Growing Up in Poverty".

Bram De Rock gratefully acknowledges the FWO and FNRS for their support.

Frederic Vermeulen gratefully acknowledges the Fund for Scientific Research - Flanders (FWO) and the Research Fund KU Leuven for financial support.

1 Introduction

This review focuses on inequality at the level of individuals. This immediately raises the question: inequality of what? Inequality measures used by international organizations such as the OECD and the World Bank generally assess inequality between households on the basis of income or total expenditure. Because households differ both in resources and demographic composition, equivalence scales were developed to adjust incomes and expenditures to make them comparable across households. As a result, equivalized household income or expenditure has become the standard measure employed by the abovementioned institutions. Beyond inequality measurement, equivalized income is routinely used as an indicator of household welfare, to evaluate the effects of policy reforms, and to identify households that fall below a predetermined poverty threshold.

A central feature of this approach is that it implicitly assumes that household resources are shared equally among household members or, equivalently, that all members attain the same level of welfare. Intra-household inequality is therefore ruled out by construction. The danger, however, is that substantial disparities within households remain hidden, potentially distorting assessments of inequality, poverty, and the distributional consequences of public policy. A growing empirical literature shows that this concern is far from merely theoretical, as discussed in greater detail in Sections 2 and 6. For example, Lise and Seitz (2011) document that individual consumption inequality in the United Kingdom is underestimated by roughly 50 percent when intra-household allocation is ignored. Similarly, Cherchye, De Rock, Lewbel, and Vermeulen (2015) and Lechene, Pendakur, and Wolf (2022) show that conventional household-level poverty rates substantially underestimate individual poverty rates in a variety of countries. These findings often reveal pronounced gender differences and connect to broader concerns regarding intra-household power and the distribution of resources, as highlighted in recent work such as Jayachandran and Voena (2026).

These concerns motivate the perspective adopted in this review: we take *individuals* as a starting point rather than households. Although this view is now widely shared by economists working on household behavior, following a research program that started about four decades ago and to which we return extensively below, it remains far less common in the broader profession. It is even less prevalent among policymakers and statistical agencies, whose measures of inequality and poverty continue to focus almost exclusively on households.

This focus on households rather than individuals has deep historical roots. Since the pioneering contributions of Lorenz (1905) and Gini (1912), inequality measurement has relied primarily on micro data on income, expenditure, and wealth. Earlier traditions in economics, including Ricardo (1817) and Marx (1867), emphasized the functional distribution of income across social groups and classes. As household-level micro data became increasingly available throughout the twentieth century, empirical work naturally shifted toward the measurement of inequality *between* households. This focus on households is also reflected in the nature of most available datasets. Household budget surveys typically record expenditures without distinguishing who consumes what, while many social surveys collect information on labor earnings at the individual level but aggregate most other income components at the household level. As a result, these data contain little direct information about how household resources are distributed among household members.

Theoretical models of consumer demand developed along similar lines. In the unitary

model described by Samuelson (1947), Deaton and Muellbauer (1980), and Mas-Colell, Whinston, and Green (1996), the household is treated as a single decision-making unit maximizing a household utility function that is assumed to represent household preferences. While analytically convenient, the unitary model cannot account for preference heterogeneity among household members. More importantly for our purposes, it provides no mechanism for recovering individual resource shares, defined as the fraction of total household resources allocated to the different household members. Yet such resource shares are essential if the objective is to measure inequality or poverty at the individual level.

This brings us to the heart of the matter. We argue that the move from households to individuals is not simply a matter of obtaining better data. It is fundamentally a theoretical issue. Meaningful analysis of intra-household inequality requires a *theoretical framework* that is specifically designed to address the allocation of resources within households.

To develop this argument, it is useful to return to the question posed at the outset: inequality of what? A natural starting point is individual welfare, which depends, among other things, on the goods and services available to individuals. Consider food consumption, drinks with friends, or a fitness subscription. These are examples of goods that are privately consumed within the household. Every apple purchased by the household can be consumed by only one person, and many fitness subscriptions are tied to a specific individual. Multiplying individual consumption quantities by market prices and aggregating them into an individual expenditure measure may thus provide a meaningful approximation of individual welfare.

At this point, it is important to emphasize that the relevant question is not who makes the purchases. Focusing on expenditure responsibility would be misleading. In a traditional household, for example, a working husband might appear poor simply because most shopping is undertaken by his spouse. What matters instead are the bundles of goods and services that are ultimately consumed by individuals. More fundamentally, welfare is ultimately experienced by individuals rather than households.

Yet household surveys typically provide information only on aggregate household consumption. It is precisely here that theory becomes indispensable. To infer how aggregate household resources are translated into individual consumption bundles, one requires a model of intra-household allocation. The unitary model offers one possibility, but only by imposing, explicitly or implicitly, the assumption that resources are allocated in a way that makes all household members equally well off. Given the growing body of evidence on intra-household inequality and differential poverty risks, this assumption appears overly restrictive. Measuring individual consumption bundles therefore requires models that take individuals, rather than households, as the starting point. In this respect, the analysis naturally aligns with the principle of methodological individualism (Weber, 1922), according to which social phenomena should ultimately be explained in terms of individual behavior rather than group behavior.

Let us add one more layer of complexity. Individuals in households do not only consume goods and services that benefit themselves. A substantial share of household resources is spent on goods and services that have a public character within the household. Rent, heating, and other forms of housing consumption are obvious examples. These goods are consumed jointly, even though they need to be purchased only once. Public consumption complicates welfare measurement because household members may value such goods differ-

ently. This creates an additional identification problem: without a theoretical framework, it is impossible to identify the potentially heterogeneous valuations that individuals attach to jointly consumed goods.

Thus far, we have focused on goods and services as sources of welfare. Yet time is equally important. Time allocation influences labor supply and earnings, shapes household production, and reflects constraints and preferences that vary across household members. Evidence from time-use surveys documents substantial gender differences in market work, childcare, household production, and leisure (Burda, Hamermesh, and Weil, 2013). These patterns suggest that understanding individual welfare requires looking beyond market expenditures alone. Moreover, time enters welfare not only through leisure but also through the production of goods and services within the household. Because domestic production relies heavily on time inputs, particularly in households with children, accounting for time use is necessary to understand both how resources are allocated within families and how household welfare is generated. Combined with expenditure information, time-use data make it possible to capture the trade-off between time and money and thereby provide a more complete picture of individual welfare based on a household's full income (Becker, 1965).

In sum, we need a theoretical framework capable of recovering individual consumption bundles from observed household-level choices, while also accounting for the individual valuation of publicly consumed goods and intra-household time allocations. This naturally raises the question of which framework is best suited for this task. Broadly speaking, two classes of models can be distinguished. The first assumes that household members behave non-cooperatively, whereas the second assumes that household decisions result in Pareto-efficient allocations. Although both approaches have received considerable attention in the literature, we adopt a clear perspective from the outset. Our focus is on the non-unitary *collective* approach developed by Chiappori (1988, 1992), which models the household as a group of individuals with potentially different preferences who make decisions jointly and reach Pareto-efficient intra-household allocations.

The appeal of the collective framework lies not only in its ability to recover individual consumption bundles, but also in its capacity to identify the preference structures that determine how private and publicly consumed goods translate into individual welfare. Importantly, this approach also enjoys substantial empirical support and offers a useful benchmark for the analysis of intra-household behavior, as discussed in Section 3.1. A defining feature of the collective model is the presence of Pareto weights attached to the utility functions of household members. These weights reflect each individual's influence over household decisions and are often interpreted as indicators of bargaining power. However, Pareto weights depend on the cardinalization of utility functions and are thus of limited direct interest from a welfare-economic perspective. Ultimately, welfare analysis should rest on ordinal preferences rather than cardinal utilities.

This naturally shifts attention to the *sharing rule*, which emerges as the key object of interest in the collective framework. It maps total household resources and other relevant determinants into the resource shares allocated to individual household members. As such, it describes how household resources are distributed across members and provides a tractable and theory-consistent link between observed household choices and the distribution of welfare within the household. Unlike Pareto weights, the sharing rule is invariant to the chosen cardinalization of utility and therefore offers a more robust basis for empirical welfare anal-

ysis. The objective is not to recover Pareto weights themselves, but rather to determine how household resources translate into well-being and, in turn, into measures of poverty and inequality.

A central research question then concerns how to identify the sharing rule in data environments that contain little direct information on the intra-household allocation of resources. More generally, this requires identifying both the allocation of resources across household members and the preferences that govern the welfare implications of this allocation. To address these issues, we review in Section 4 the basic essentials of three broad identification strategies that infer individual preferences and resource shares from widely available datasets: the differential approach, the revealed preference approach, and preference-similarity methods. Once preferences and resource shares have been identified, observed consumption choices and time allocations can be translated into measures of the distribution of resources within the household and, ultimately, into measures of individual welfare, inequality, and poverty.

An important implication is that welfare analysis requires more than identifying the allocation of resources within the household. In the presence of jointly consumed goods, welfare comparisons also depend on how the benefits of this public consumption are attributed to individual household members. As a result, the collective framework gives rise to several conceptually distinct measures of individual welfare, which we briefly introduce in Section 6. The differences between these measures primarily reflect alternative ways of incorporating public consumption into welfare assessment and, in particular, of accounting for potentially heterogeneous valuations of jointly consumed goods.

The objective of this review is therefore twofold. First, we review identification strategies based on the collective model that allow researchers to recover individual consumption bundles, resource shares, and preferences over both private and publicly consumed goods from widely used datasets. Second, we show how this information can be used to move beyond household-level analysis and conduct welfare, inequality, and poverty assessments at the level of individuals.

The remainder of the paper is organized as follows. Before turning to the theoretical and empirical strategies that recover intra-household allocations from standard household data, Section 2 discusses what can already be learned from the relatively rare datasets that contain direct evidence on individual consumption and time use. It examines both the insights provided by such direct survey evidence and its limitations, thereby motivating the methods reviewed in subsequent sections. Section 3 introduces the collective framework, highlights the role of the sharing rule, and shows how the model accommodates preference heterogeneity, public consumption, and time use. Section 4 reviews empirical strategies for identifying intra-household preferences and resource shares using conventional datasets, covering differential, revealed preference, and preference-similarity approaches. Section 5 extends these methods to richer empirical settings and organizes the discussion by identification strategy. Section 6 examines how identified resource shares can be translated into measures of individual welfare, poverty, and inequality. It also discusses some key empirical findings obtained through the discussed identification approaches applied to real-world data. Section 7 concludes and looks ahead.

2 Direct survey evidence

Although datasets containing both individual expenditure and individual time-use information remain uncommon, they provide the most direct evidence available on how resources and time are allocated within households. Examining such data offers a useful starting point for the analysis that follows. They reveal how consumption is divided between privately and publicly consumed goods, how private expenditures are distributed across household members, and how market work, home production, and leisure are allocated within couples. As such, they offer a unique window on dimensions of intra-household behavior that are largely invisible in standard datasets but are central to the measurement of individual welfare. At the same time, they highlight a broader theme of this review: even when unusually detailed individual-level information is available, translating observed allocations into measures of individual welfare ultimately requires a theoretical framework to reflect the heterogeneity across households.

To illustrate both the insights offered by such data and the limitations of direct survey evidence, we draw on three specific datasets that contain particularly rich information on intra-household behavior: the Dutch Longitudinal Internet Studies for the Social Sciences (LISS), the Belgian Measuring Equivalent Incomes (MEqIn) survey, and the Japanese Panel Survey of Consumers (JPSC).

We begin with the Dutch LISS data, which, among others, contain detailed information on both expenditures and time use in households with children. Table 1 shows that average private expenditures of husbands and wives are similar in levels, while the bulk of nondurable spending accrues to categories that are jointly consumed. It is important to emphasize once more that this type of information is generally unavailable in standard household budget surveys, which typically report only aggregate household expenditures on items such as food, drinks, and leisure activities. Although expenditures on goods such as children's clothing, diapers, or toys may sometimes be separately reported, they provide only partial information on resources devoted to children, since spending on food, drinks, and many leisure activities is usually contained in aggregate household expenditures and cannot be assigned directly to individual household members.

The same data also reveal a clear pattern of specialization in time use. Table 1 shows that husbands supply more market labor hours on average, whereas wives devote more time to child care and other domestic work. These differences are economically meaningful and consistent with the idea that household-level averages conceal substantial intra-household variation. At the same time, the reported standard deviations point to considerable variation across households, indicating significant inter-household heterogeneity.

Table 1 might suggest that spouses have private expenditures that are more or less equal to each other. And that the often made implicit assumption that resources are distributed equally within households is not too bad after all. This is clearly a wrong conclusion. We show this by means of another one of the few datasets that provide detailed information on the intra-household allocation of resources. More specifically, the Belgian MEqIn evidence complements the above picture based on Dutch data by documenting how private expenditure shares vary across households. Figure 1 plots women's shares in private expenditures by quartile. Moving from the first to the fourth quartile, the female share rises from roughly 35% to 61%, with intermediate quartiles around 45% and 52%. This dispersion contradicts

	Husband		Wife		Household	
	Mean	St.dev.	Mean	St.dev.	Mean	St.dev.
Expenditures (EUR/month)						
Private expenditures	311.08	217.10	302.64	182.41	—	—
Expenditures on children	—	—	—	—	479.15	493.21
Other public expenditures	—	—	—	—	1,827.77	875.32
Time use (hours/week)						
Market labor (incl. commuting)	47.71	11.60	28.54	11.32	—	—
Child care	9.04	8.52	15.43	13.79	—	—
Other domestic work	11.66	8.22	20.65	11.75	—	—

Table 1: LISS (2009), couples with children: expenditures and time use.
Source: Cherchye, De Rock, and Vermeulen (2012b).

equal sharing and highlights heterogeneity that is masked by household averages.

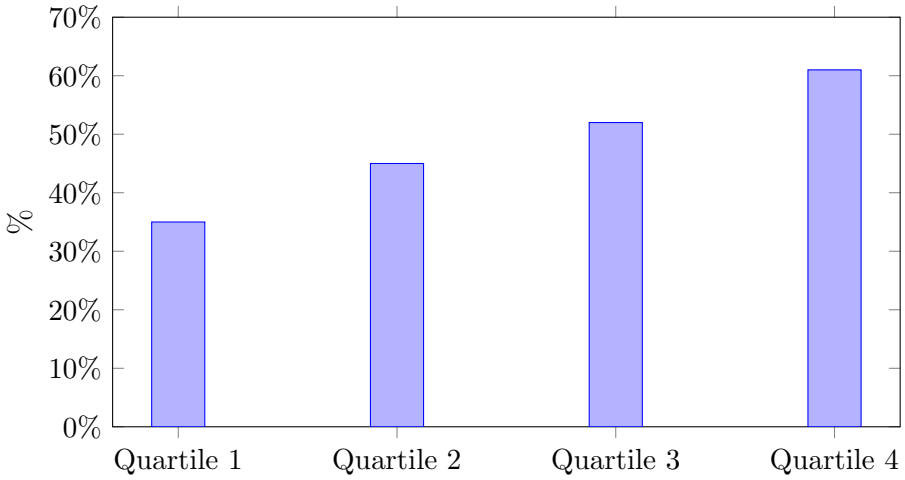

Figure 1: Women's share in private expenditures (quartiles), MEqIn.
Source: Capéau, Cherchye, Decancq, Decoster, Maniquet, Nys, Périlleux, Ramaekers, Rongé, Schokkaert, and Vermeulen (2020).

Finally, the Japanese JPSC data further show how spending and time allocations are organized within couples. Table 2 reports average private expenditures for men and women in couples, total household expenditures, and the share of total spending that is publicly consumed. The time use measures display pronounced specialization in market work and home production, while average weekly leisure is relatively similar across spouses. The distributions in Figure 2 reveal additional structure: women's relative shares of private consumption are skewed toward lower values, leisure shares concentrate near one half, home production shares are concentrated at high levels, while women's market hours shares tend to be low.

This reported evidence reveals substantial differences across countries. For example, it shows considerably more intra-household specialization and inequality in Japanese households than in Belgian and Dutch households. More importantly, these findings confirm that

household averages can be misleading for welfare assessment, reinforcing the case made in the Introduction for moving beyond household aggregates.

	Wife Mean [st.dev.]	Husband Mean [st.dev.]	Household Mean [st.dev.]
Expenditures (JPY/month)			
Private expenditures	36,711 [63,951]	77,650 [64,321]	442,640 [275,174]
Share of household total	6.5%	14.5%	79.0%
Time use (hours/week)			
Market work (incl. commuting)	29.7 [21.7]	62.7 [14.7]	—
Share of total	17.7%	37.3%	—
Home production (incl. child care)	44.0 [25.3]	7.4 [9.1]	—
Share of total	(26.2%)	(4.4%)	—
Leisure (incl. sleep)	94.2 [20.1]	97.9 [15.6]	—
Share of total	56.1%	58.3%	—
Other observables			
Age	36.5 [6.4]	38.9 [7.4]	—
Education (years)	13.2 [1.6]	13.5 [2.3]	—
Wage	925 [1,014]	1,676 [1,575]	—
Children aged 0–6	—	—	0.62 [0.80]
Household size	—	—	4.24 [1.49]

Table 2: Summary statistics, JPSC 1993-2013.
Source: Lise and Yamada (2019).

Direct surveys document how private spending is allocated across household members, how jointly consumed goods dominate household budgets, and how time use is specialized within couples, but they do not reveal how individuals value public goods or derive utility from private or shared consumption. Moreover, many prominent datasets report expenditures only at the household level. Because these data do not identify the intra-household sharing rule, which is the key link between household-level observations and individual welfare, descriptive patterns alone cannot support credible inference about individual well-being. This gap motivates the theory-based methods developed in the next sections: Section 3 introduces the cooperative collective framework, Sections 4 and 5 present identification strategies suited to standard datasets, and Section 6 shows how the recovered allocations inform the measurement of individual welfare, inequality, and poverty.

A final note concerns measurement. Although direct surveys provide essential information, collecting individual consumption and time use data is difficult and prone to error. Some surveys rely on a single respondent to report everyone's consumption and time use, or it is simply complicated to have an accurate view of the sharing of consumption. Alternatively, evidence also shows that household members may conceal income or spending from one another in ways tied to bargaining power (Zhang, 2024). These issues underscore the need for careful survey design, clear item definitions, appropriate respondent selection, and validation procedures. Section 7 revisits these concerns and discusses how richer linked expenditure and time use modules and improved data-collection protocols can strengthen the empirical basis for the tools presented in the following sections.

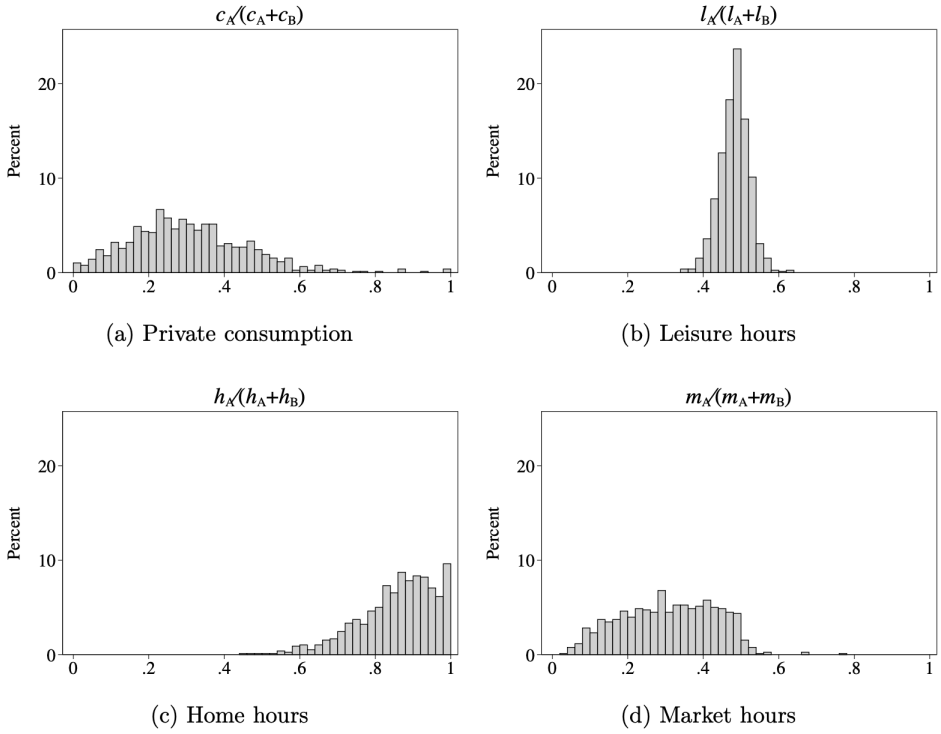

(a) Private consumption (b) Leisure hours

(c) Home hours (d) Market hours

Figure 2: Distributions of wives' relative shares: private consumption and time use (JPSC). *Source:* Lise and Yamada (2019).

3 Theoretical framework

The central objective is to explain how observed choices on private spending, shared consumption, and time use arise from underlying preferences and constraints, while recognizing that the decision process itself is part of what needs to be modeled. In short, to address questions about intra-household welfare, we require a framework that both respects individual heterogeneity and disciplines the way individual objectives aggregate into household choices.

3.1 Set-up

Two main approaches have been developed to model this intra-household allocation process. In the *noncooperative approach*, each spouse maximizes their own utility while taking the other spouse's decisions as given. Within the broader non-unitary framework of household consumption behavior, the leading noncooperative specification models intra-household interaction as a Nash equilibrium with voluntary contributions to public goods (see, e.g., Browning, Chiappori, and Lechene (2010) and Lechene and Preston (2011)). A well-known implication of this structure is that free-riding incentives may lead to underprovision of public goods, and thus to Pareto-inefficient outcomes.

By contrast, the *cooperative "collective" approach* assumes that spouses coordinate their

decisions so that allocations are Pareto-efficient (see Chiappori (1988, 1992); Browning, Chiappori, and Weiss (2014) provide an extensive review).[1] Although both perspectives have been examined, the collective framework has become the workhorse in household economics because it yields clear identification results even when only aggregate household choices are observed. Crucially, it also enjoys strong empirical support. For example, Browning and Chiappori (1998) and Cherchye and Vermeulen (2008) provide compelling evidence in favor of the collective model over the unitary model using Canadian and Dutch data; similarly, Rangel and Thomas (2019) make a strong case for Pareto-efficient consumption behavior using rich data on farm households in Burkina Faso. Direct empirical comparisons of cooperative and noncooperative models remain relatively scarce. A notable exception is Cherchye, Chiappori, De Rock, Ringdal, and Vermeulen (2026), who rely on detailed consumption data from Kenya and find strong support for the cooperative model relative to its noncooperative counterpart.

Although the cooperative approach offers a transparent and tractable benchmark, actual household behavior may exhibit departures from full efficiency due to frictions such as asymmetric information or limited commitment. As emphasized by Jayachandran and Voena (2026), models incorporating such frictions show that attempts by spouses to increase their control over resources can distort consumption and investment decisions, thereby reducing overall household welfare. Additional evidence of potentially inefficient intra-household decision-making comes from a number of empirical studies and experiments conducted in low- and middle-income countries. Examples include Udry (1996), who studied agricultural households in Ethiopia, and Ashraf (2009), who analyzed decision-making among married couples in the Philippines. Evidence for high-income countries remains relatively limited. One recent contribution is Choukhmane, Goodman, and O'Dea (2025), who show that roughly one-fifth of American couples could achieve Pareto improvements through alternative allocations of retirement-saving contributions. Their analysis further suggests that about half of these inefficiencies reflect deliberate choices by spouses, often associated with weaker marital commitment.

Jayachandran and Voena (2026) emphasize the importance of understanding departures from efficiency while noting that cooperative models, and in particular the collective framework, remain useful benchmarks for analyzing household behavior. Even in the high-stakes and inherently dynamic setting studied by Choukhmane, Goodman, and O'Dea (2025), the majority of couples appear to make decisions that are broadly consistent with Pareto efficiency. More generally, detailed household-level analyses point to substantial heterogeneity in behavior. For many couples, Pareto efficiency appears to provide a reasonable description of observed decision-making, while for others the cooperative framework may still offer a useful approximation despite some departures from efficiency. This interpretation is consistent with the findings of Cherchye, Chiappori, De Rock, Ringdal, and Vermeulen (2026), who simultaneously focused on both approaches.

We therefore use the cooperative framework as our primary reference point: it captures the dominant empirical patterns while providing the clarity and structure required for iden-

[1] A useful observation is that the distinction between cooperative and noncooperative behavior becomes irrelevant when there are no public goods and individual preferences depend only on own consumption (i.e., there are no externalities). In that case, each spouse maximizes their own utility function subject to their own budget constraint, independently of the other spouse. By construction, the resulting allocation is Pareto-efficient. See also Lechene and Preston (2011).

tification and individual welfare analysis.

Throughout the analysis, we focus on a static environment. By this we mean that household decisions are modeled as taking place within a single period, without explicitly modeling intertemporal dependencies. Most core results in the literature are derived in static settings, which provide a transparent mapping between assumptions, observables, and welfare conclusions. At the same time, important empirical applications demonstrate that intertemporal considerations can affect the household-level distribution of welfare; see, for example, Voena (2015) and Lise and Yamada (2019). This has motivated the use of dynamic household consumption models to better understand how observed intra-household inequality arises. Still, to the extent that the main role of dynamic decisions is to determine the amount of resources available in each period, a static collective model remains sufficient for identifying how those within-period resources are allocated across household members (Chiappori and Mazzocco, 2017). We therefore adopt a static benchmark for clarity and return briefly to dynamic aspects in Section 7.

Our theoretical framework jointly incorporates time allocation and consumption. In the baseline setting, individuals allocate time between market labor and leisure. Consistent with most empirical work, leisure is treated as privately consumed, although some models allow for externalities whereby one spouse's leisure enters the other spouse's utility (see, e.g., Chiappori (1988)). In richer variants, which we discuss in Section 5, time also enters as an input into home production alongside market purchases (see, e.g., Chiappori (1997) and Cherchye, De Rock, and Vermeulen (2012b)). This joint treatment of time and consumption is crucial for welfare analysis because it links labor supply, domestic production (including childcare), and the distribution of consumption benefits.

For clarity, we will mainly focus on two-decision maker households, although the framework generalizes to M-person households with $M \geq 2$ (see Browning and Chiappori (1998) and Chiappori and Ekeland (2009)). In practice, this means that we consider households with two adult decision makers. This is the simplest departure from the unitary model, which is appropriate for single-adult households but silent on intra-household allocation, and it encompasses couples with or without children, where children do not directly participate in the bargaining process. This two-person structure isolates the essential features of collective decision-making and provides the foundation for the identification results developed below.

Expenditures on children are treated as publicly consumed within the household for the purpose of analyzing adult welfare, so that children's welfare is reflected indirectly through these publicly consumed household expenditures. This interpretation is particularly natural when children do not actively participate in decision-making, yet their needs shape household-wide consumption of goods. In principle, children can be modeled as separate decision makers with their own decision power, allowing their welfare to be evaluated directly (see Dauphin, El Lahga, Fortin, and Lacroix (2011), Dunbar, Lewbel, and Pendakur (2013), Lechene, Pendakur, and Wolf (2022), and Aminjonov, Bargain, and Colacce (2025)). For expositional convenience, we do not explicitly consider this extension in what follows, although the generalization is relatively straightforward.

3.2 BCM model: collective labor supply with public consumption

We adopt the collective labor supply model with two decision makers and a public good introduced by Blundell, Chiappori, and Meghir (2005), which we refer to as the BCM model for short. This framework captures both the role of publicly consumed goods and the interaction between time allocation and consumption in determining individual welfare, while remaining tractable enough to illustrate how unobserved individual resource shares can be recovered from standard data. The BCM model assumes two decision makers, which we refer to as (household) members, irrespective of whether children are present. Likewise, we refer to the household as a couple.

There are three commodities in the setting we are focusing on: a Hicksian composite good and both household members' leisure. Given our static setting, the price of the Hicksian composite good has a price that is normalized to one. It has three uses: the private consumption of member 1, which is denoted by q^1, the private consumption of member 2, which is denoted by q^2, and public consumption, which is denoted by Q. Typical examples of private consumption include food or clothing, while public consumption may consist of spending on children or on housing. Given the nature of the consumption, both household members consume their own private consumption and the same level of the public good. Clearly, in household budget surveys and most other data sets, only the sum of both household members' private consumption is observed. We therefore denote this aggregate private consumption by q, where $q = q^1 + q^2$.

The model further assumes that the decision makers allocate their time exclusively between leisure, denoted by l^1 and l^2, and market work, denoted by m^1 and m^2. Each member i therefore faces the time constraint:

$$l^i + m^i = T, \tag{1}$$

where T is the total time available. In principle, the model could be developed without incorporating time and could focus solely on consumption. However, as argued above, time allocation is a crucial dimension for welfare analysis. Moreover, in our simple setting, the opportunity cost of time is the individual's wage rate. Because wage rates typically differ across spouses, they affect bargaining positions within the couple in ways that are highly informative for identification.

Let w^1 and w^2 denote the market wages of members 1 and 2, and let y denote the household's nonlabor income. Combining consumption and time yields the couple's full budget constraint:

$$q^1 + q^2 + Q + w^1 l^1 + w^2 l^2 \leq y + w^1 T + w^2 T. \tag{2}$$

Next, each member i is assumed to have rational preferences represented by a utility function of the form:

$$u^i = u^i(q^i, l^i, Q). \tag{3}$$

The utility functions u^1 and u^2 are strictly increasing, continuous, and concave. As the same quantity of the public good enters both utility functions, this specification reflects its non-rival nature. Note also that the specification excludes externalities across private goods or leisure: a member does not derive utility from the other member's private consumption or

leisure. These are therefore "egoistic" preferences. This assumption could, in principle, be relaxed. For example, leisure enjoyed by one member may affect the other member's utility, and private goods with externalities may be treated empirically as public goods. In this respect, see also our discussion in Section 5 on identification under "caring" preferences.

The utility functions of the two members need not be identical, reflecting the possibility of heterogeneous preferences. Nevertheless, the couple ultimately makes joint decisions. Following the collective approach, these decisions are assumed to be Pareto-efficient (Chiappori, 1988, 1992). Given the linear constraints and our assumptions on preferences, the couple's allocation of time and money can be represented as the solution to:

$$\max_{l^1,l^2,q^1,q^2,Q} \quad u^1(q^1,l^1,Q) + \mu(w^1,w^2,y,z)\, u^2(q^2,l^2,Q), \tag{4}$$

subject to

$$\begin{aligned} l^1 + m^1 &= T, \\ l^2 + m^2 &= T, \\ q^1 + q^2 + Q + w^1 l^1 + w^2 l^2 &\leq y + w^1 T + w^2 T. \end{aligned}$$

In this formulation, μ denotes member 2's relative Pareto-weight, which reflects the individual's relative bargaining position. Intuitively, a higher Pareto-weight implies that the couple's allocation of consumption and time more closely follows member 2's preferences, indicating a stronger influence on household decisions. In principle, the Pareto-weight may depend on the exogenous variables in the maximization problem, specifically the members' wages w^1 and w^2, the household's nonlabor income y, and *distribution factors*, which we denote by the vector z. Distribution factors affect household decisions solely through the Pareto-weight, but do not affect individual preferences or the full-income constraint (see Browning, Bourguignon, Chiappori, and Lechene (1994) for an early contribution, and Bourguignon, Browning, and Chiappori (2009) and Browning and Chiappori (1998) for more general results). Examples of distribution factors used in the literature include divorce laws and sex ratios in individuals' marriage markets (Chiappori, Fortin, and Lacroix, 2002), as well as the relative size and wealth of the husband's and wife's family networks (Attanasio and Lechene, 2014). See Browning, Chiappori, and Weiss (2014, Table 5.1) for a list of distribution factors that have been used in the empirical literature.

Summarizing, equation (4) stipulates the two key elements that need to be identified for analyzing individual welfare: individual preferences, represented by the members' utility functions and reflecting their specific trade-offs between the different forms of consumption and time use, and the Pareto weight, which summarizes the household decision process. As soon as the Pareto weight changes with exogenous variables, the differences in individual preferences become important because they may lead to different household decisions (see Browning, Chiappori, and Lechene (2006) for more discussion). Observed household data sets, including the ones documented in the previous section, reflect the combined influence of these two elements. Disentangling preferences from the intra-household allocation process therefore constitutes the central identification challenge that we address in Section 4.

Special cases. This simple collective model nests several well-known special cases. First, the standard unitary model is obtained when the Pareto-weight μ is constant. In this case,

the objective function in (4) becomes a household utility function. Because the private consumption, q^1 and q^2, and the public consumption, Q, have the same normalized price, Hicks' composite commodity theorem applies (see, e.g., Deaton and Muellbauer (1980)). The objective function can therefore be rewritten as $u(c, l^1, l^2)$, where $c = q^1 + q^2 + Q$. The maximization program then reduces to:

$$\max_{l^1, l^2, c} \quad u(c, l^1, l^2) \tag{5}$$

subject to

$$\begin{aligned} l^1 + m^1 &= T, \\ l^2 + m^2 &= T, \\ c + w^1 l^1 + w^2 l^2 &\leq y + w^1 T + w^2 T. \end{aligned}$$

In this case, utility maximization yields Marshallian demand functions for leisure and consumption, satisfying the standard Slutsky conditions. Observing both members' leisure (or labor supply) and household consumption as functions of wages and nonlabor income therefore allows identification of the household preferences, represented by $u(c, l^1, l^2)$, provided these Slutsky conditions hold (Deaton and Muellbauer, 1980). However, nothing can be learned about the allocation of c into q^1, q^2, and Q; the unitary model is thus silent on intra-household allocation.

A second special case is the original collective labor supply model of Chiappori (1988), which excludes public consumption. All monetary resources are allocated to private consumption q^1 and q^2, leading to:

$$\max_{l^1, l^2, q^1, q^2} \quad u^1(q^1, l^1) + \mu(w^1, w^2, y, z)\, u^2(q^2, l^2) \tag{6}$$

subject to

$$\begin{aligned} l^1 + m^1 &= T, \\ l^2 + m^2 &= T, \\ q^1 + q^2 + w^1 l^1 + w^2 l^2 &\leq y + w^1 T + w^2 T. \end{aligned}$$

This model will play a prominent role when outlining the differential approach to identifying the collective model in Section 4.1.

Finally, some identification approaches abstract from labor supply and focus exclusively on the intra-household allocation of consumption. In particular, the preference-similarity approach discussed in Section 4.3 often relies on the assumption that leisure is weakly separable from consumption. Under this assumption, attention can be restricted to consumption choices, leading to the following model:

$$\max_{q^1, q^2} \quad u^1(q^1) + \mu(y, z)\, u^2(q^2) \tag{7}$$

subject to

$$q^1 + q^2 \leq y,$$

where y denotes the couple's total expenditure.

3.3 The sharing rule

A particularly noteworthy feature of the maximization programs (4) and (6) is that they admit an equivalent representation based on decentralization results directly linked to the Second Fundamental Theorem of Welfare Economics (see, e.g., Mas-Colell, Whinston, and Green (1996)). This alternative representation proceeds in two stages. In the first stage, the household's resources are allocated across members, yielding the "sharing rule". In the second stage, each member maximizes their own utility subject to the resources allocated to them. The precise form of this two-stage representation depends on the presence and the treatment of public consumption within the household.

The simplest case is the collective labor supply model of Chiappori (1988), corresponding to maximization program (6). In the first stage, the household's nonlabor income y is divided into individual shares y^1 and y^2, with $y^1 + y^2 = y$. This *sharing rule* typically depends on the same variables that influence the Pareto-weight: the members' wages, total nonlabor income, and distribution factors. In the second stage, both members independently allocate their individual full income to private consumption and leisure, subject to their own budget constraint. This leads to the following maximization program for each member i:

$$\max_{l^i, q^i} \quad u^i(q^i, l^i) \tag{8}$$

subject to

$$l^i + m^i = T, \qquad q^i + w^i l^i \leq y^i(w^1, w^2, y, z) + w^i T.$$

The sharing rule (y^1, y^2) thus describes how household resources are allocated to members and contains essential information for measuring inequality and poverty within households. We return to these welfare questions in Section 6.

When a public good is present, two alternative forms of sharing rules arise. The first is the *conditional sharing rule*. In the BCM framework, the first stage allocates resources to the public good and to each member's share of the residual nonlabor income, defined as nonlabor income minus expenditures on the public good. In the second stage, each member maximizes their utility over private consumption and leisure, taking as given both the public good level and their share of residual nonlabor income. For each member i, this takes the form:

$$\max_{l^i, q^i} \quad u^i(q^i, l^i, \overline{Q}) \tag{9}$$

subject to

$$l^i + m^i = T, \qquad q^i + w^i l^i \leq y_c^i(w^1, w^2, y, z) + w^i T.$$

In this case, the conditional income shares satisfy

$$y_c^1(w^1, w^2, y, z) + y_c^2(w^1, w^2, y, z) = y - Q(w^1, w^2, y, z),$$

where $\overline{Q} = Q(w^1, w^2, y, z)$ is the Pareto-efficient public consumption level resulting from program (4).

An alternative decentralized representation of the collective model with public consumption involves the use of Lindahl prices; see Cherchye, De Rock, and Vermeulen (2011b) and

Cherchye, De Rock, Lewbel, and Vermeulen (2015) for extensive discussion. This formulation allows each member to face a personalized Lindahl price P^i for the public good. Under Pareto-efficiency, these Lindahl prices reflect each member's marginal willingness-to-pay for the public good and satisfy the adding-up condition (see Lindahl (1958); Samuelson (1954)):

$$P^1 + P^2 = 1.$$

This yields the following maximization problem for each member i:

$$\max_{l^i, q^i, Q} \quad u^i(q^i, l^i, Q) \tag{10}$$

subject to

$$l^i + m^i = T, \qquad q^i + P^i Q + w^i l^i \leq y_g^i(w^1, w^2, y, z),$$

where the *general sharing rule* defines individual income shares such that

$$y_g^1(w^1, w^2, y, z) + y_g^2(w^1, w^2, y, z) = y + w^1 T + w^2 T.$$

It is instructive to compare the three maximization programs (8), (9) and (10). The second-stage problems (8) and (9) are very similar: once individual income shares are determined, each member allocates their resources optimally between private consumption and leisure, facing the same market prices. The only difference is that in (9) the level of the public good Q is taken as given. Thus, conditional on income shares, the two models behave qualitatively alike: interaction between members arises solely through the first-stage sharing rule.

In contrast, in problem (10), each member also chooses Q and faces an individualized Lindahl price. Once shares in full income are determined, each member allocates resources optimally across private consumption, leisure, and the public good. Under efficiency, both members choose the same level of Q, and their Lindahl prices sum to the market price.

These comparisons highlight a fundamental distinction between private and public goods in a Pareto-efficient setting. For private goods, all members face the same market price but may consume different quantities. For public goods, all members consume the same quantity but generally face different Lindahl prices.

4 Identification

Identification refers to the conditions under which the structural components of the model (i.e., individual preferences for each household member and the intra-household resource allocation mechanism summarized by the sharing rule or the Pareto-weight) can be recovered from observed labor supply and household consumption behavior. A model is identified if no distinct structural specification can rationalize exactly the same observable behavior. This matters especially for welfare analysis: different structural models may imply markedly different welfare conclusions. We therefore emphasize identification arguments that do not hinge on functional form assumptions.

We next show how the BCM model can be identified, focusing on three complementary strategies: the differential approach, the revealed preference approach, and the approach

based on similarity of individual preferences. In principle, these strategies apply to any empirical specification nested within the structural model described above. This observation is particularly relevant for the differential and similarity-based approaches, which, in practice, require researchers to adopt functional forms for individual preferences, household demand, and labor supply. Consequently, applied identification results will generally depend on the imposed functional structure. By contrast, the revealed preference approach does not rely on empirical functional-form specification beyond the structural framework itself: identification follows from observed choices and revealed preference axioms alone, making it intrinsically nonparametric. Hence, in applications based on revealed preference, identification results do not hinge on functional assumptions about observed behavior.

Although these approaches differ in their assumptions and data requirements, they are best seen as complementary. The differential approach is most informative when there is substantial variation in prices or wages; under differentiability, it yields point identification, though its implementation still requires functional-form choices. The revealed preference approach avoids functional-form assumptions altogether but typically produces set rather than point identification. As we will explain, set identification means that we recover sets of feasible specifications for the unobserved structural components of collective consumption models (such as individual preferences and the sharing rule), which in turn allows us to place bounds on these unobserved elements. A distinctive feature of the preference-similarity approach is that it estimates a household consumption technology that determines the private versus public nature of goods, thereby avoiding any need to classify goods as private or public a priori. Like the differential approach, it yields point identification, but only under more restrictive behavioral assumptions regarding the similarity of individuals' preferences across different household situations.

4.1 Differential approach

The differential approach traces back to Chiappori (1988); Chiappori and Ekeland (2009) provide a comprehensive treatment of this identification strategy. In the specific BCM model under consideration, the approach treats individual labor supplies and expenditures on the public good as differentiable functions of wages and nonlabor income. The testable implications of the model then take the form of partial differential equations, analogous to the classical Slutsky conditions that apply to the unitary model. The identification it delivers is inherently *local*, relying on derivatives evaluated at observed data points together with integrability conditions that ensure these derivatives correspond to an underlying structural model. This local nature contrasts with the *global* reasoning in the next subsection, which is based on revealed preference methods and does not depend on differential properties of demand functions.[2]

[2]The distinction between our notions of local and global identification is closely related to the distinction between point and set identification discussed earlier. As we show in this subsection, the local differential approach obtains point identification: observed household behavior implies a unique value for the object of interest, such as the effect of a change in wages on individual resource shares. By contrast, the global revealed preference approach, discussed in the next subsection, entails set identification: a set of values for individual resource shares (rather than a unique one) is compatible with observed household behavior. See Lewbel (2019) for a more general discussion of point and set identification.

Preliminaries. For the BCM model in program (4), the empirical analyst observes each member's labor supply, m^1 and m^2, and expenditures on the public good, Q, as functions of wages w^1 and w^2 and nonlabor income y. These functions also depend on a distribution factor z, i.e., a variable that shifts the Pareto-weight without affecting preferences or the household's full income (see Section 3.2). Formally, we assume the observables:

$$\begin{aligned} m^1(w^1, w^2, y, z), \\ m^2(w^1, w^2, y, z), \\ Q(w^1, w^2, y, z). \end{aligned} \tag{11}$$

We note that retrieving these functions can be achieved with cross-sectional data under standard homogeneity assumptions across observationally similar households. It is equally important to emphasize what is not observed: in general, we do not observe the private consumption q^1 and q^2 separately, but only their aggregate $q = q^1 + q^2$. This data structure is common in many datasets, where expenditures on public goods or on children are recorded in detail, while many other items are aggregated.

The objective is then to separate individual preferences from the intra-household allocation process and thereby identify the structural components of the model from observed behavior. The key observation is that labor supplies are exclusive goods: each member's leisure enters only that member's utility. In the two-stage representation of Pareto-efficient allocations we study (see (9)), fixing the level of the public good Q implies that any change in a member's labor supply driven by the other member's wage, by nonlabor income, or by the distribution factor must operate through residual nonlabor income.

Unitary to collective. It is useful to briefly recall identification in the standard one-person labor supply model, which corresponds to the unitary program:

$$\max_{l,c} \quad u(c, l) \tag{12}$$

subject to

$$\begin{aligned} l + m = T, \\ c + wl \leq y + wT. \end{aligned}$$

In this setting, observing the labor supply function $m(w, y)$ identifies preferences represented by $u(q, l)$ under the usual Slutsky conditions; by adding up, the consumption demand function $c(w, y)$ adds no further information.

This unitary benchmark closely parallels the second stage of Chiappori's (1988) collective model without public consumption in program (8). The crucial difference is that, in the unitary case, observed nonlabor income y enters the budget directly, whereas in the collective case the relevant object is the unobserved income share $y^i(w^1, w^2, y)$. In the unitary model, one can vary w while holding y fixed (or vice versa) to achieve identification; in the collective model, varying w^i also changes $y^i(w^1, w^2, y)$, preventing such separation.

This is where the exclusivity of labor supplies becomes central. A change in w^1 affects member 2's choices only through $y^2(w^1, w^2, y)$; symmetrically, a change in w^2 affects

member 1 only through $y^1(w^1, w^2, y)$. Chiappori (1988) exploits this feature to identify the collective labor supply model in the absence of public consumption.[3] The observable labor supplies can be written as

$$m^1(w^1, w^2, y) = \tilde{m}^1\big(w^1,\, y^1(w^1, w^2, y)\big)\,, \tag{13}$$

$$m^2(w^1, w^2, y) = \tilde{m}^2\big(w^2,\, y - y^1(w^1, w^2, y)\big)\,. \tag{14}$$

From this structure it follows that

$$\frac{\partial m^1/\partial w^2}{\partial m^1/\partial y} = \frac{\partial y^1/\partial w^2}{\partial y^1/\partial y}, \tag{15}$$

$$\frac{\partial m^2/\partial w^1}{\partial m^2/\partial y} = \frac{-\,\partial y^1/\partial w^1}{1 - \partial y^1/\partial y}. \tag{16}$$

Hence, observable derivatives identify ratios of partial derivatives of the sharing rule, motivating the differential approach. The two equations relate two observable expressions (on the left side) to three unknown derivatives, i.e.,

$$\frac{\partial y^1}{\partial w^1}, \qquad \frac{\partial y^1}{\partial w^2}, \qquad \frac{\partial y^1}{\partial y}, \tag{17}$$

and the goal is to integrate these to recover $y^1(w^1, w^2, y)$. Standard results for partial differential equations imply this is feasible if and only if the symmetry (integrability) conditions hold:

$$\frac{\partial^2 y^1}{\partial y\,\partial w^1} = \frac{\partial^2 y^1}{\partial w^1\,\partial y}, \qquad \frac{\partial^2 y^1}{\partial y\,\partial w^2} = \frac{\partial^2 y^1}{\partial w^2\,\partial y}, \qquad \frac{\partial^2 y^1}{\partial w^1\,\partial w^2} = \frac{\partial^2 y^1}{\partial w^2\,\partial w^1}. \tag{18}$$

Imposing these restrictions on observable partial derivatives identifies the partial derivatives of the sharing rule. The sharing rule is then identified up to an integration constant, and the same holds for the private consumption of the Hicksian good, q^1 and q^2. For a given constant, the (given cardinalization of the) individual utility functions $u^1(\cdot)$ and $u^2(\cdot)$ are identified. In other words, individual utility functions are identified up to a translation. This contrasts with the unitary model, where the utility function is completely identified (for a given cardinalization of course, since any monotone increasing transformation of the utility function gives rise to exactly the same observed behavior).

Economically, one can trace the effect of a wage change on a member's nonlabor income share and private consumption. Although the levels of these shares are identified only up to an additive constant, this constant is welfare-irrelevant, as shown by Chiappori (1992).

More specifically, suppose the household members' utility functions representing their individual preferences, $u^1(c^1, l^1)$ and $u^2(c^2, l^2)$, along with their shares $y^1(w^1, w^2, y)$ and

[3]The current discussion assigns a key role to individual leisure (or labor supply), which is assumed to be exclusive. If leisure generates externalities (i.e., one household member's leisure enters the other member's utility function), then it can no longer serve as an exclusive good for identification. In that case, this identification strategy requires other exclusive goods with member-specific prices, such as male and female clothing. We return to this issue in Section 5.1.

$y^2(w^1, w^2, y)$ solve the above system of partial differential equations. Then, for any constant k, the transformed utility functions $u^1(c^1 - k, l^1)$ and $u^2(c^2 + k, l^2)$, combined with the transformed shares $y^1(w^1, w^2, y) + k$ and $y^2(w^1, w^2, y) - k$, also constitute a solution. The implication is that welfare conclusions are unchanged by this transformation. If a particular reform increases the welfare of household member 1 under $u^1(c^1, l^1)$ and $y^1(w^1, w^2, y)$, then it will also increase that member's welfare under $u^1(c^1 - k, l^1)$ and $y^1(w^1, w^2, y) + k$. In other words, it can always be assumed that household member 1 receives k more units of private consumption through the sharing rule, while household member 2 receives k units less. The reason is that the utility functions can be adjusted simultaneously, so that the additional consumption received through the sharing rule is exactly offset by a corresponding shift in the utility representation. As a result, both members are equally well off under $u^1(c^1, l^1)$ and $u^2(c^2, l^2)$ as under $u^1(c^1 - k, l^1)$ and $u^2(c^2 + k, l^2)$, provided the utility functions are adjusted accordingly.

We refer to Chiappori (1992) for an illustration and Chiappori and Ekeland (2009) for a more technical discussion. As shown above, this unidentified constant is directly related to the fact that one often only observes the aggregate private consumption $c = c^1 + c^2$ in standard data sets. With full information on intra-household allocations (see examples in Section 2), the constant itself can be identified; see Cherchye, De Rock, and Vermeulen (2012b).

Including public consumption. We now extend identification to the BCM model with public consumption, which is central to our analysis. As noted earlier, the second stages of the models without public goods (see (8)) and with public goods (see (9)) are analogous: conditional on a fixed Q, the objectives are qualitatively the same, so the differential strategy applies. However, when public consumption is present, changes in w^1, w^2, or y shift not only income shares but also the optimal $Q(w^1, w^2, y, z)$, so Q is not constant as wages or income vary. In this case, following Blundell, Chiappori, and Meghir (2005), one can use the distribution factor z to hold Q fixed.

By the implicit function theorem, the restriction $\overline{Q} = Q(w^1, w^2, y, z)$ (for fixed $\overline{Q}$) is equivalent to $z = \tilde{z}(w^1, w^2, y, \overline{Q})$. Substituting into (11) yields labor supply functions that depend only on (w^1, w^2, y), since $\overline{Q}$ is constant:

$$m^1\big(w^1, w^2, y, \tilde{z}(w^1, w^2, y, \overline{Q})\big) = \tilde{m}^1(w^1, w^2, y). \tag{19}$$

$$m^2\big(w^1, w^2, y, \tilde{z}(w^1, w^2, y, \overline{Q})\big) = \tilde{m}^2(w^1, w^2, y). \tag{20}$$

These observables identify the *conditional* sharing rule in (9) and conditional private consumption up to a constant. For a given constant, the (given cardinalization of the) conditional individual utilities are also identified (conditioning on Q). With additional information on individual consumption, the constant can be pinned down as well, yielding full identification of the conditional sharing rule and private consumption.

The remaining task concerns the first stage, where members allocate total resources between the public good and residual nonlabor incomes. As shown by Blundell, Chiappori,

and Meghir (2005), the optimal (y^1, y^2, Q) solves

$$\max_{y^1,y^2,Q} \; v^1(w^1, y^1, Q) + \mu\, v^2(w^2, y^2, Q) \tag{21}$$

subject to

$$y^1 + y^2 + Q \leq y,$$

where $v^i(\cdot)$ denotes member i's indirect utility given income share y^i and public consumption Q. The optimum satisfies the Bowen–Lindahl–Samuelson condition: the sum of the members' marginal rates of substitution between the public good and private consumption equals the public good's (normalized) price. Combined with the second-stage results, this delivers identification of preferences over private consumption, leisure, and the public good, as well as the Pareto-weight for the given cardinal representation of the individual preferences.

4.2 Revealed preference approach

The theory of revealed preference originates with the seminal contribution of Samuelson (1938), who sought to build consumer theory directly on observed choices rather than on unobservable utility constructs, aiming to "drop off the last vestiges of utility analysis" (Samuelson, 1938, p. 62). Samuelson's central proposition was that a consumer's behavior across different price regimes reveals preference orderings: comparing the expenditure required to purchase alternative consumption bundles allows one to infer which bundles are preferred. The framework relies on a minimal set of postulates and does not require any functional-form assumptions for demand. The key postulate, now known as the *Weak Axiom of Revealed Preference* (*WARP*), states that if a consumer chooses bundle one when bundle two is affordable, then the reverse choice cannot occur. Houthakker (1950) strengthened this requirement by adding transitivity, yielding the *Strong Axiom of Revealed Preference* (*SARP*). Somewhat ironically, demands that satisfy *SARP* are exactly those rationalizable by utility maximization, a complementarity already noted by Samuelson in his *Foundations of Economic Analysis*.

The revealed preference-based identification strategy for the BCM model builds on the approach of Afriat (1967), who made revealed preference theory particularly suitable for empirical analysis. This framework was further developed and popularized by Diewert (1973) and Varian (1982). Before turning to the collective setting, we first briefly review the revealed preference characterization of the unitary model. This serves both to clarify the basic identification logic and to prepare the groundwork for the collective model with public consumption. We also discuss how the resulting structure yields *set identification* of the sharing rule and individual preferences, in contrast to the *point identification* obtained under the differential approach. We restrict ourselves to presenting the main definitions and the intuition behind them. For a more detailed treatment, including worked numerical examples, we refer to Cherchye, De Rock, and Vermeulen (2011a).

Unitary model. A defining feature of the approach proposed by Afriat (1967) is that it begins with a finite set of observed price-quantity combinations rather than a full system of demand equations. These observations reflect choices made either by a single household

over multiple time periods or by different households assumed to share homogeneous preferences. In the unitary case, we observe a collection of consumption–leisure bundles and the corresponding wages (recall that the price of the Hicksian aggregate good is normalized to one). Formally, suppose the dataset consists of N observations:

$$S = \{(1, w_n; c_n, l_n),\ n = 1, \ldots, N\}. \tag{22}$$

We say that S can be rationalized by the unitary model if there exists a utility function u such that, for each observation n, the chosen bundle (c_n, l_n) yields at least as much utility as any alternative affordable bundle (c, l) satisfying $c + w_n l \leq c_n + w_n l_n$. In other words, each observed bundle must be utility maximizing subject to the relevant budget constraint (as specified in (12)).

Varian (1982) shows that such a utility function exists if and only if the dataset satisfies the *Generalized Axiom of Revealed Preference* (*GARP*). A dataset S satisfies *GARP* if there exist revealed preference relations R_0 and R such that, for any sequence of observations $(a, b, \ldots, f)$ contained in S:

$$\text{if } c_a + w_a l_a \geq c_b + w_a l_b, \text{ then } (c_a, l_a) R_0 (c_b, l_b), \tag{23}$$

$$\text{if } (c_a, l_a) R_0 (c_b, l_b),\ (c_b, l_b) R_0 (c_d, l_d), \ldots, (c_e, l_e) R_0 (c_f, l_f), \text{ then } (c_a, l_a) R (c_f, l_f), \tag{24}$$

$$\text{if } (c_a, l_a) R (c_b, l_b), \text{ then } c_b + w_b l_b \leq c_a + w_b l_a. \tag{25}$$

The three conditions have an intuitive interpretation. Condition (23) defines the *directly revealed preference* relation R_0. Specifically, if bundle (c_a, l_a) is chosen when bundle (c_b, l_b) is also affordable at the prices observed in situation a, then the consumer reveals a preference for (c_a, l_a) over (c_b, l_b). Condition (24) extends this direct revealed preference relation by imposing transitivity. For example, if bundle a is directly revealed preferred to bundle b, and bundle b is directly revealed preferred to bundle d, then bundle a is revealed preferred to bundle d, even if these two bundles are never directly compared in the data.

Condition (25) imposes the key consistency requirement. If bundle (c_a, l_a) is revealed preferred to bundle (c_b, l_b), then bundle (c_b, l_b) cannot be strictly more expensive than (c_a, l_a) at the prices observed in situation b. Otherwise, the consumer would choose (c_b, l_b) even though a revealed preferred bundle, namely (c_a, l_a), is affordable, which contradicts utility maximization. Thus, *GARP* requires that every observed choice be cost-minimizing relative to the set of bundles that are revealed preferred to it.

Beyond this axiomatic characterization, Afriat (1967) provides a constructive method for recovering a utility function whenever the data satisfy *GARP*. The *Afriat Theorem* shows that the dataset S is rationalizable by a utility function if and only if there exist numbers U_n and positive scalars λ_n for all observations $n \in \{1, ..., N\}$, that satisfy the *Afriat inequalities* for any pair of observations (n, m):

$$U_m \leq U_n + \lambda_n \big(c_m - c_n + w_n (l_m - l_n)\big). \tag{26}$$

These inequalities link the utility levels U_n at observed bundles to the corresponding expenditure differences, with the scalars λ_n representing marginal utilities of income. A

feasible specification of the pairs (U_n, λ_n) exists if and only if the dataset S is *GARP*-consistent, and Afriat provides an explicit rationalizing utility function:

$$u(c,l) = \min_n \left\{ U_n + \lambda_n \big(c - c_n + w_n(l - l_n) \big) \right\}.$$

Building on Afriat, Varian (1982) shows how these characterizations of rational consumer behavior can be used to construct upper and lower bounds on indifference curves through the observed bundles, as well as to forecast behavior at unobserved price–income combinations. Importantly, a rationalizing utility representation exists for any specification of the pairs (U_n, λ_n) that satisfy the linear inequalities, and because the observed dataset is finite, there are typically infinitely many such feasible specifications. In this sense, the revealed preference approach delivers *set identification*: many different utility functions, representing *different* preferences, can rationalize the same dataset based on the Afriat inequalities.

This feature contrasts sharply with the local differential approach discussed in Section 4.1, which yields point identification by relying on smoothness and derivative information. In the unitary model, identification via differentiability hinges on a labor supply function that maps each wage–income pair to a unique optimal bundle lying on a unique indifference curve. In contrast, the revealed preference approach observes only a finite menu of choices, implying that many indifference curves are consistent with each observation. Consequently, predictions for unobserved price–income combinations, as well as welfare measures such as Hicksian compensating and equivalent variation (see, e.g., Blundell, Browning, Cherchye, Crawford, De Rock, and Vermeulen (2015)), are set identified.

BCM model. Let us now turn to the set identification of the BCM model. Our discussion follows Cherchye, De Rock, and Vermeulen (2011b), who develop set identification results for a broad class of collective consumption models, including the one considered here.

The analogue of the dataset in (22) is:

$$S = \{(1, 1, w_n^1, w_n^2; q_n, Q_n, l_n^1, l_n^2),\ n = 1, \ldots, N\}, \tag{27}$$

where we explicitly include the normalized prices of the Hicksian private and public goods.

Then Pareto-efficiency implies the existence of two member-specific datasets, each of which must satisfy *GARP* (or, equivalently, the Afriat inequalities):

$$S^1 = \{(1, P_n^1, w_n^1, q_n^1, Q_n, l_n^1),\ n = 1, \ldots, N\}, \tag{28}$$

$$S^2 = \{(1, P_n^2, w_n^2, q_n^2, Q_n, l_n^2),\ n = 1, \ldots, N\}, \tag{29}$$

with $q_n^1 + q_n^2 = q_n$ and $P_n^1 + P_n^2 = 1$ for all n.

This last result merits further discussion. First, both S^1 and S^2 must satisfy *GARP*. This follows directly from the assumption that each household member behaves as a rational utility maximizer: if individual preferences are rational, then each member's observed choices must satisfy revealed preference consistency. Second, the structure of S^1 and S^2 reflects a core property of Pareto-efficient collective consumption. For private goods, household members face the same market price but may choose different quantities at that price. For public

goods, they consume the same quantity Q_n but generally face different individual Lindahl prices, which reflect heterogeneous willingness-to-pay. In equilibrium, these Lindahl prices must sum to the normalized market price of the public good, as shown by Samuelson (1954).

These observations yield a central implication of the revealed preference approach. If, for each observation n, one can assign private consumption q_n^1 and q_n^2 that sum to q_n, together with Lindahl prices P_n^1 and P_n^2 that sum to one, such that both S^1 and S^2 satisfy $GARP$, then the observed behavior is consistent with Pareto-efficiency. In general, many such allocations will satisfy these requirements, and the method therefore provides set identification of preferences, individual private consumption quantities, Lindahl prices, and the sharing rule. This applies to both the conditional sharing rule in (9) and the general sharing rule in (10), when Lindahl prices are used to evaluate the public good.

Cherchye, De Rock, and Vermeulen (2011b) show that the individual $GARP$ requirements for the datasets S^1 and S^2 can be reformulated as a system of linear constraints involving binary integer variables, which represent each member's binary revealed preference relations. This reformulation makes it possible to apply standard mixed integer linear programming (MILP) techniques to bring the revealed preference characterization of the collective model to empirical data. As a result, it provides an operational procedure for implementing set identification in practice, based directly on the revealed preference structure of the collective model.

4.3 Preference-similarity approach

A third identification approach builds on assumptions about the similarity of preferences across individuals, either across different household types or within households. Here, similarity should be understood as the idea that preferences remain close, though not necessarily identical, when individuals move between household environments. Over the past fifteen years, several such approaches have been proposed, each relying on the insight that suitably restricted variation in preferences can generate additional identifying power.

This approach is based on Browning, Chiappori, and Lewbel (2013). In their framework, a collective consumption model is characterized by a consumption technology that captures the extent to which goods are jointly consumed within the household. Their objective is to estimate economies of scale in consumption together with the sharing rule that describes the intra-household allocation of resources. This approach differs from those discussed earlier in two important respects. First, it does not impose any *a priori* classification of goods as private or public. Instead, the degree of joint consumption may vary across goods and is treated as an unknown parameter. Second, the identification strategy exploits data from three distinct groups: single women, single men, and couples. The model is fully identified under the assumption that the preferences of single women and single men are identical to those of their counterparts within couples. This assumption makes it possible to recover individual preferences from the observed behavior of singles. These identified preferences can then be combined with observed behavior of couples to identify both the consumption technology and the intra-household allocation rule.

At this point, we note that less restrictive versions of the identical preferences assumption can also deliver identification (for example, when parametrizing how preferences change when moving from singlehood to living in a couple), as discussed in Browning, Chiappori,

and Lewbel (2013). However, all such approaches share the central premise that individuals' preferences in different household configurations are not identical but sufficiently similar to enable identification. We discuss alternative relaxations of the preference-similarity assumption in Section 5.3, which can be employed when the specific focus is the identification of the sharing rule.

Adapted BCM model. As in the BCM model, we restrict attention to a three–commodity environment. However, because the private versus public nature of goods is not defined a priori in this approach, we no longer explicitly interpret the three goods as private consumption, leisure, and public consumption. Instead, we simply use q_j^i to denote the quantity of good $j \in \{1,2,3\}$ consumed by household member $i \in \{1,2\}$. Member i's direct utility function is then given by:

$$u^i = u^i(q_1^i, q_2^i, q_3^i). \tag{30}$$

The individual consumption q_1^i, q_2^i, and q_3^i are not observed, nor are the corresponding individual resource shares. We also do not observe the sum of the individuals' consumption: $q_1 = q_1^1 + q_1^2$, $q_2 = q_2^1 + q_2^2$, and $q_3 = q_3^1 + q_3^2$. What is observable is the couple's market purchase behavior: given prices (p_1, p_2, p_3) and total expenditures y, the couple purchases the bundle (z_1, z_2, z_3). Importantly, (q_1, q_2, q_3) generally differ from (z_1, z_2, z_3) because the commodities may be jointly consumed. The degree of joint consumption is captured by a household consumption technology, which in its simplest version is linear:

$$z_1 = \lambda_1 q_1, \tag{31}$$

$$z_2 = \lambda_2 q_2, \tag{32}$$

$$z_3 = \lambda_3 q_3, \tag{33}$$

where the unknown technology parameters λ_1, λ_2, and λ_3 lie between one half and one.

This specification is reminiscent of Becker (1965), who introduced a household production framework in which market goods are transformed into household goods. In the present setting, the couple purchases (z_1, z_2, z_3), which through the consumption technology yield larger household quantities (q_1, q_2, q_3). If good 1 is purely private, then $\lambda_1 = 1$ and $z_1 = q_1$. If good 2 is entirely public, then $\lambda_2 = 1/2$ and the household needs to buy only $z_2 = q_2/2$, with both members consuming $q_2^1 = q_2^2 = z_2/2$. Many goods will lie between these extremes, corresponding to $\lambda_j \in (1/2, 1)$. Browning, Chiappori, and Lewbel (2013) refer to the individual quantities q_j^i as *private-good equivalents*, which reflect the combination of the consumption technology and the intra-household allocation rule.

The couple's consumption problem can be written as:

$$\max_{q_1^1, q_1^2, q_2^1, q_2^2, q_3^1, q_3^2, z_1, z_2, z_3} u^1(q_1^1, q_2^1, q_3^1) + \mu\, u^2(q_1^2, q_2^2, q_3^2) \tag{34}$$

subject to

$$\begin{aligned} q_j &= q_j^1 + q_j^2, \qquad j = 1,2,3, \\ z_j &= \lambda_j q_j, \qquad j = 1,2,3, \\ p_1 z_1 + p_2 z_2 + p_3 z_3 &\le y. \end{aligned}$$

Identification under preference-similarity. Using the general results in Browning, Chiappori, and Lewbel (2013), this framework yields a two–stage representation that parallels the two–stage structure introduced earlier for the BCM model. In the second stage, each household member maximizes their direct utility function subject to an individual budget constraint. As before, this constraint is written in terms of personalized "Lindahl" prices rather than market prices and depends on the individual's share of household resources. Importantly, however, the definition and interpretation of these Lindahl prices differ from those in the earlier BCM setting. Under the linear consumption technology, both household members face the same Lindahl prices, given by $\lambda_j p_j$ for each good j. Because λ_j lies between one half and one, these Lindahl prices are always weakly below market prices. Intuitively, this reflects the fact that joint consumption effectively stretches household resources.

Using these Lindahl prices, we may write the maximization problem faced by each member i as:

$$\max_{q_1^i, q_2^i, q_3^i} u(q_1^i, q_2^i, q_3^i) \tag{35}$$

subject to

$$\lambda_1 p_1 q_1^i + \lambda_2 p_2 q_2^i + \lambda_3 p_3 q_3^i \leq y^i.$$

The optimal solution to this program yields the Marshallian demand of member i for good j:

$$g_j^i(\lambda_1 p_1, \lambda_2 p_2, \lambda_3 p_3, y^i).$$

Adapting Browning, Chiappori, and Lewbel (2013) to this simplified setting, the relationship between the couple's observed household demands and the individual demand functions is then given by

$$z_j(p_1, p_2, p_3, y) = \lambda_j \Big[g_j^1(\lambda_1 p_1, \lambda_2 p_2, \lambda_3 p_3, \eta(p_1, p_2, p_3, y)) + g_j^2(\lambda_1 p_1, \lambda_2 p_2, \lambda_3 p_3, y - \eta(p_1, p_2, p_3, y)) \Big], \tag{36}$$

where $z_j(\cdot)$ denotes household demand and

$$\eta(p_1, p_2, p_3, y) = \lambda_1 p_1 q_1^1 + \lambda_2 p_2 q_2^1 + \lambda_3 p_3 q_3^1$$

represents member 1's share of household resources.

In (36), the left–hand sides are observable, as they record the household's market purchases at given prices and total expenditure. By contrast, the individual Marshallian demands g_j^i are not observed and cannot be recovered from couples' data alone, because the individual consumption q_j^i and the resource shares y^i are inherently unobserved. Without further structure, the mapping from household purchases to individual demand functions is therefore underidentified. Preference-similarity offers a way forward. The most restrictive version assumes that an individual's preferences do not change upon entering a couple (i.e., preferences of single men and women are identical to those of their counterparts living in a couple). Under this assumption, the observed demand behavior of single men and single women directly identifies the preferences of men and women in couples.

Intuitively, identification proceeds as follows. Observed purchases by singles, expressed as functions of prices and income, identify singles' preferences, following the same logic used for

the unitary model in Section 4.1. Given these identified preference structures, the household demand system in (36) then pins down the remaining unknowns, namely the technology parameters λ_j and the individual resource shares $\eta(p_1, p_2, p_3, y)$ and $y - \eta(p_1, p_2, p_3, y)$, which effectively obtains identification of the collective model under study.

5 Extensions and refinements

The BCM framework is sufficiently general to cover most key concepts in the collective household literature, yet simple enough to present the three identification strategies in a transparent way accessible to a broad audience. In this section, we examine several important extensions and refinements of this baseline, including richer preference structures, and models that incorporate home production, nonparticipation in the labor market, or stability conditions in the marriage market. Naturally, the BCM model can be generalized in many directions, or simplified where appropriate. Moreover, there exists a variety of collective models that are neither straightforward specializations nor direct extensions of the baseline. Our aim is not to be exhaustive. Instead, we highlight several useful directions for researchers analyzing intra-household inequality, emphasizing intuition and providing references to detailed treatments. The choice of specification often hinges on the available data, for instance whether one has detailed time use information, and on the extent of price and income variation.

In what follows, we organize the discussion around the three identification approaches developed earlier, with each subsection presenting extensions that build directly on the corresponding framework. We begin with the differential approach, showing how its derivative-based logic extends to settings with caring preferences, multiple distribution factors, nonparticipation in the labor market, home production technologies, and the use of exclusive or assignable goods. We then turn to the revealed preference approach, outlining how its nonparametric structure accommodates environments with multiple private and public goods, how applying revealed preference restrictions to household demand systems can substantially tighten the bounds implied by set identification of intra-household allocations, and how marriage market stability conditions provide additional identifying power even when only a single observation per household is available while allowing for heterogeneous individual preferences. Finally, we examine the preference-similarity approach, highlighting how Engel curve methods enable sharing rule identification when price variation is limited, how weaker preference-similarity assumptions combined with assignable goods allow estimation using couple-level data alone, and how recent linear estimation procedures make the recovery of resource shares more accessible for applied work.

These extensions illustrate how the appropriate modeling choice depends on the empirical setting and objective, while preserving a coherent link from household-level observations to individual-level welfare. More broadly, they highlight the versatility and unifying nature of the collective approach.

5.1 Differential approach

The differential identification strategy introduced in Section 4.1 exploits smooth variation in wages, nonlabor income, and distribution factors to reveal the structure of intra-household allocations through systems of partial differential equations. The key insight underlying this approach is that exclusive goods, most notably leisure, provide a channel through which changes in one member's wage affect the other member's behavior only via the sharing rule. By tracing how observed labor supplies and public-good expenditures respond to marginal changes in exogenous variables, the differential method delivers point identification of individual preferences and the sharing rule up to a welfare-irrelevant constant. This local, derivative-based logic makes the approach particularly powerful in settings with sufficiently rich continuous variation in wages or distribution factors.

Caring preferences. While the BCM model can in principle accommodate caring preferences, it is most straightforward to introduce this concept in Chiappori (1988, 1992)'s original collective labor supply model. In this model, each household member i is characterized by a direct utility function of the form $u^i(q^i, l^i)$ $(i = 1, 2)$; see (8). This specification implies egoistic preferences: individuals derive utility only from their own private consumption and their own leisure. However, the identification of the sharing rule using the differential approach applies equally well in settings where members have "caring preferences". Formally, member i is said to have caring preferences if their direct utility function takes the form:

$$v^i\big(u^1(q^1, l^1),\, u^2(q^2, l^2)\big)\,. \tag{37}$$

With caring preferences, a household member derives utility not only from their own consumption and leisure, but also from the other member's consumption and leisure, though in a specific manner: the functions $u^1(q^1, l^1)$ and $u^2(q^2, l^2)$ are sub-utility functions in a weakly separable utility function $v^i(\cdot)$. A member's utility depends on the other member's sub-utility level, although the particular allocation generating that sub-utility (for example, high consumption and low leisure or vice versa) does not matter. As Chiappori (1988) shows, all identification results for a collective labor supply model with egoistic preferences carry over directly to a model with caring preferences. Without further information on the intra-household allocation of consumption, the sharing rule can be identified up to a constant (which is welfare-irrelevant), private consumption can be identified up to the same constant, and, for a given constant, individual preferences can be identified. Intuitively, the same Pareto-frontier applies under egoistic and caring preferences, except that in the latter case, extremely unequal allocations in favor of one member are ruled out.

Still, from a welfare point of view, the distinction between egoistic and caring preferences remains important. Since the functions $v^1(\cdot)$ and $v^2(\cdot)$ cannot be identified from observable choices, they cannot be used for welfare evaluation. This contrasts with the sub-utility functions $u^1(\cdot)$ and $u^2(\cdot)$, which *are* identifiable from observed choices. Assuming egoistic preferences, therefore, allows one to use these utilities directly in welfare analysis. Put differently, while caring preferences allow identification of the intra-household allocation of time and money, they do not allow evaluation of these allocations in utility terms. At the very least, such a welfare interpretation becomes open to debate.

Identification with distribution factors. In Section 4.1, we discussed how the presence of a distribution factor, which shifts the household members' Pareto-weights without affecting their preferences or the household budget constraint, helps identify the BCM model with public consumption. Earlier contributions such as Thomas (1990), Grossbard-Shechtman (1993), Angrist (2002), and Duflo (2003) document that variables like relative incomes and sex ratios can influence intra-household bargaining power and thereby alter the allocation of resources. Chiappori, Fortin, and Lacroix (2002) formalize these insights by incorporating distribution factors into a structural collective labor supply model, showing that such variables can be introduced coherently and that their inclusion leads to a simpler and more robust set of identification conditions than in Chiappori (1988). Technically, the main advantage is that the relevant identifying restrictions (analogous to those in (15) and (16)) require only second derivatives of the observed labor supply functions rather than third derivatives.

A complementary and more encompassing perspective on distribution factors is provided by Bourguignon, Browning, and Chiappori (2009), who offer a thorough analysis of their role in the testability and identification of collective models. They show that, under Pareto-efficient intra-household allocations, distribution factors generate specific proportionality restrictions on how demands respond to shifts in these factors (affecting individuals' bargaining power), and they characterize when these restrictions help identify the sharing rule and individual Engel curves across different observational schemes. Taken together, these insights clarify how exogenous variation in distribution factors can serve as a meaningful source of identification in collective consumption models, illustrating the broader potential of such variables for empirical work.

Nonparticipation and reservation wages. Labor supply choices provide crucial information for identifying intra-household allocations, since variation in labor supply responses to wages or distribution factors helps reveal how resources are shared within couples. Up to this point, we have abstracted from nonparticipation. This assumption is natural because the identification strategies discussed earlier rely on observable changes in labor supply, which may not arise when an individual does not work. In the unitary model, nonparticipation is typically explained by a market wage below the reservation wage, defined as the wage that leaves the individual indifferent between working and not working. In a collective setting, however, this reasoning does not apply directly, as labor supply decisions result from the joint optimization of two individuals rather than a single agent.

As shown by Donni (2003) and Blundell, Chiappori, Magnac, and Meghir (2007), reservation wages must therefore be modeled explicitly within the collective framework. They demonstrate that Pareto-efficiency requires that both partners are indifferent between one member working or not working, and they derive the corresponding restrictions on labor supply. Crucially, the identification results developed earlier continue to hold in this broader environment: the collective model remains identified even when one partner does not participate in the labor market. This extension substantially increases the empirical relevance of the model, as it allows researchers to analyze couples in which only one member works, a common situation in many datasets and demographic groups.

Home production. In the BCM model, household members allocate their time between leisure and market work only. As Becker (1965) emphasized, however, this is too restrictive: a substantial share of nonmarket time is devoted to home production activities such as child care, cleaning, and cooking. Although these activities may yield utility, it is unrealistic to treat them purely as leisure. Collective labor supply models that incorporate home production have therefore been developed; see Chiappori (1997) and Cherchye, De Rock, and Vermeulen (2012b). In the latter model, which is closest to the BCM framework, each member's utility depends not only on private consumption and leisure but also on domestically produced household goods. These domestic goods are produced by combining time and market inputs through a household production technology in the spirit of Becker (1965). One such domestic good is children's welfare, produced using parents' time inputs (child care) and monetary expenditures on children (for example, clothing, diapers, or toys). These domestic goods behave as public goods within the household, since their "quantity" enters both spouses' utility functions. Children's welfare, for instance, can be written as:

$$u_c(t_c^1, t_c^2, Q_c), \tag{38}$$

where t_c^i denotes the time spent by member i on child care and Q_c denotes expenditures on children. This set-up implies that each member's effective consumption depends on both their monetary share and their contribution of time to production. The shadow value of time, which is equal to the wage rate in simple settings or to the marginal product of time in more general environments, enters the full income constraint and therefore influences how resources are allocated within the household. As a result, variation in wages, labor supply, or time spent in home production shifts bargaining positions not only through earnings but also through their impact on the production of shared goods.

Variation in individual wages makes it possible to identify the subutility function u_c, by arguments similar to those used in our discussion following (5). When there is only one domestic good, the BCM model can be adapted by replacing the household-level public good Q with $u_c(t_c^1, t_c^2, Q_c)$ under constant returns to scale of the child-welfare production technology. Identification then proceeds as in the differential approach of Section 4.1, and the presence of a distribution factor allows identification of individual preferences, private consumption, the sharing rule, and the household production technology.

In models with multiple domestic goods, however, identification through the differential approach requires the use of "production shifters" (Cherchye, De Rock, and Vermeulen, 2012b). These variables operate on preferences only through the production technologies: they influence the production of domestic goods but do not affect preferences outside the technologies, bargaining positions, or the household budget constraint. Examples associated with the production of children's welfare include the number and ages of children. For identification purposes, production shifters play a role analogous to distribution factors: they allow the output of domestic goods to be held constant while wages and nonlabor income vary. Importantly, they are necessary for identification as soon as the collective model contains several domestic goods, because multiple distribution factors alone only shift the Pareto weight.

Exclusive and assignable private goods. In the BCM model, leisure is treated as a private good that generates utility only for the individual who consumes it. In this sense,

leisure is an *exclusive* good: it is privately consumed, priced at an individual-specific rate, and its consumption by one household member has no direct effect on the utility of the other. This distinction is useful for understanding the data requirements of collective consumption models that do not rely on time use information. Whenever the analyst observes an exclusive good for each household member, the differential identification strategy of Section 4.1 applies directly. Variation in the consumption of such exclusive goods as a function of prices and total expenditures then allows identification of individual preferences, private consumption, and the sharing rule, up to an additive welfare-irrelevant constant.

Another notable class of private goods in view of differential identification purposes consists of *assignable* goods. These goods, too, are privately consumed by a single household member, but unlike exclusive goods they share the same market price for both individuals; their distinguishing feature is that individual-level consumption is directly observed. When a distribution factor is available, the presence of even one assignable good is sufficient for identification, because it delivers individual-specific information that reveals how resources are allocated within the household. In this way, assignable goods play a role analogous to exclusive goods in environments where prices do not differ across members, provided that at least one distribution factor can be used for identification. For detailed analyses of identification with exclusive and assignable goods in collective models, see Chiappori and Ekeland (2009), Bourguignon, Browning, and Chiappori (2009), and Chiappori and Mazzocco (2017).

5.2 Revealed preference approach

The revealed preference framework introduced in Section 4.2 provides a basis for analyzing collective consumption behavior using only observed choices and the axioms of rationality. A central implication is that the approach delivers *set identification*: rather than pinpointing a unique structure, it characterizes all preference configurations, Lindahl prices, and sharing rules that can rationalize the data. This nonparametric perspective provides a flexible alternative to differential identification, particularly in settings where functional-form restrictions or smoothness assumptions are difficult to justify. By avoiding a parametric specification, revealed preference tests furnish a robust empirical benchmark for evaluating the plausibility of the collective model.

Multiple public and private goods. The identification approach outlined in Section 4.2 is readily generalizable. Cherchye, De Rock, and Vermeulen (2011b) show that individual consumption, Lindahl prices, and the sharing rule can be set identified in environments with multiple private goods per individual and multiple public goods. This generalization makes it possible to recover individual preferences in the spirit of Varian (1982) without requiring exclusive goods or distribution factors. In other words, even when only aggregate household consumption is observed, the revealed preference framework allows one to reconstruct the full set of individual preference structures and intra-household sharing patterns that are consistent with observed household behavior across different price regimes.

Relative to the basic setting discussed in Section 4.2, the added value of this extension lies in its capacity to handle richer commodity spaces and to exploit the informational content of multigood variation. The framework can readily accommodate (even partly) assignable goods when such information is available; price variation for these assignable goods further

increases identifying power by tightening the mapping from observed household purchases to feasible individual allocations. This makes the approach particularly useful for applications based on detailed expenditure surveys, where researchers seek informative bounds on resource shares despite not observing full individual-level consumption. The revealed preference conditions continue to deliver set identification, but with tighter and more empirically meaningful bounds that reflect the structure of multigood demand.

Demand systems and revealed preference restrictions. A notable modification of the revealed preference approach is proposed by Cherchye, De Rock, Lewbel, and Vermeulen (2015), who combine insights from the differential approach with those of the revealed preference framework. In particular, they show that set identification of the sharing rule, often with very tight bounds, can be achieved by applying revealed preference restrictions for the collective household consumption model to the observable household demand system, even in the absence of exclusive goods, assignable goods, or distribution factors. Moreover, when estimating the household demand system, one can incorporate differential restrictions that characterize collective household behavior, which serve as the collective analogues of the Slutsky conditions in the unitary model (see Browning and Chiappori (1998); Chiappori and Ekeland (2009)). This yields a powerful integration of the revealed preference and differential approaches, allowing researchers to recover informative and empirically relevant bounds on individual resource shares while maintaining full consistency with the rationality restrictions of the collective model of household consumption.

In the approach described in Section 4.2, the revealed preference restrictions are applied to a finite set of observed commodity bundles. Cherchye, De Rock, Lewbel, and Vermeulen (2015) instead apply them to the entire household demand system, which specifies a commodity bundle for any combination of market prices and total expenditures. By exploiting the full structure of the demand system, their approach incorporates substantially more information about how households adjust their consumption across a wide range of economic environments. This allows revealed preference inequalities to rule out a larger share of infeasible sharing rules and to isolate a much smaller set of allocations consistent with collective rationality. The method can generate markedly narrow bounds on individual resource shares, making it particularly valuable in applications where within-household allocation patterns are otherwise difficult to pin down.

Marriage market stability conditions. Another interesting extension of the revealed preference approach is developed by Cherchye, Demuynck, De Rock, and Vermeulen (2017). They combine the collective household model with the assumption of a stable marriage market in the sense of Gale and Shapley (1962) and Becker (1973) to set identify the sharing rule. The key idea is that marriage involves economies of scale that arise through the consumption of public goods, which may allow an individual to be materially better off in a couple than when living alone. This potential gain from marriage is realized only when the within-marriage allocation of resources ensures that both partners are at least as well off as they would be as singles, which is the "individual rationality" condition. Stability in the marriage market further requires that these within-marriage allocations are preferred to any feasible allocation that a partner could secure with an alternative spouse, taking equilibrium

conditions in the marriage market into account; this is the "no blocking pairs" condition.

This observation implies that comparing the budget sets that an individual faces as a single person or when matched with alternative potential partners yields additional information about the sharing rule that applies within the observed couple. Cherchye, Demuynck, De Rock, and Vermeulen (2017) show that this sharing rule can be identified using standard linear programming techniques applied to a cross-section of couples, that is, using only one observation per household while allowing for fully heterogeneous preferences across individuals. This approach differs fundamentally from earlier methods, which require multiple observations for each household or impose homogeneity assumptions across observationally similar couples. Here, identification is obtained by embedding household-level efficiency in a stable marriage market structure. Cherchye, De Rock, Surana, and Vermeulen (2020), Cherchye, De Rock, Vermeulen, and Walther (2021) and Cherchye, De Rock, and Surana (2025) provide various extensions and applications of this marriage-market-based approach, including the analysis of intra-household scale economies, marital instability of consumer-producer households in low income settings, and the welfare implications of non-working individuals in a collective framework.

5.3 Preference-similarity approach

The preference-similarity framework introduced in Section 4.3 models intra-household allocations by linking household-level consumption to the private-good equivalents consumed by individual members. It does so through a household consumption technology that maps household purchases into individual consumption goods. The central idea is that, under suitable restrictions on how preferences may vary across individuals or household types, one can recover both the household consumption technology and the sharing rule even when individual consumption is not directly observed. In contrast to revealed preference methods, which deliver set identification based solely on choice consistency, and to differential approaches, which rely on smooth responses to wages or prices, the preference similarity approach provides a structural route to point identification by exploiting similarities between the preferences of singles and those of individuals within couples.

The identification strategy based on preference-similarity assumptions turns out to be particularly valuable when relative price variation is limited or when estimating full demand systems is empirically demanding, and when the primary interest lies in identifying individual resource shares (i.e., the sharing rule) rather than recovering the full structural household consumption model (i.e., individual preferences and the household technology). The remainder of this section discusses how subsequent contributions have translated the preference-similarity framework into empirically tractable methods for estimating individual resource shares. A key development is the use of Engel curves to infer resource shares from variation in total household expenditure alone, thereby avoiding reliance on price variation.

Limited price variation and Engel curves. This line of research was initiated by Lewbel and Pendakur (2008), who combine the model of Browning, Chiappori, and Lewbel (2013) with assumptions on the structure of economies of scale within the household. In particular, they rely on Barten scales and assume that these scales do not vary with total household expenditure. Under these conditions, a systematic relationship emerges between

the household's Engel curve and the Engel curves of individual household members, which can be exploited to identify members' resource shares.

More formally, let $w_j(\tilde{x})$ (for $j = 1, \ldots, n$) denote the household's budget-share Engel curve for good j as a function of the logarithm of total household expenditure $\tilde{x}$. Let $w_j^i(\tilde{x})$ denote member i's Engel curve in budget-share form, and let η^i denote the resource share of member i. Further, let I^i denote the "indifference scale" of member i, defined as the income the individual would require as a single person to be as well off as in the current household, given the sharing rule. This indifference scale, a concept that will return in Section 6.1, reflects both the difference between the market prices faced by singles and the Lindahl prices faced by individuals within couples, as well as the intra-household allocation of resources. For two-person households, the relationship between the household-level and individual Engel curves is then given by ($j = 1, \ldots, n$):

$$w_j(\tilde{x}) = h_j + \eta^1 w_j^1(\tilde{x} - \ln(I^1)) + (1 - \eta^1) w_j^2(\tilde{x} - \ln(I^2)). \tag{39}$$

After specifying a functional form for the Engel curves, this model can be estimated using data on single men, single women, and couples. The approach proposed by Lewbel and Pendakur (2008) yields estimates of the members' relative resource shares η^1 and η^2, as well as their indifference scales I^1 and I^2. Since the method does not rely on variation in prices, it cannot identify individual preferences as the differential or revealed preference approaches. Nonetheless, it provides valuable estimates of resource shares in empirical settings where price variation is limited. Bargain and Donni (2012) extend this methodology to recover the share of resources allocated to children.

Weaker preference-similarity assumptions. The identification strategies of Lewbel and Pendakur (2008) and Bargain and Donni (2012) still require data on both singles and couples, and rely on preference-similarity assumptions across men and women in these groups. Additional assumptions can eliminate this requirement. This is shown by Dunbar, Lewbel, and Pendakur (2013), who propose a semiparametric identification strategy for the resource shares of men, women, and children in couples. Their method can be applied to data on couples only and requires the estimation of Engel curves without relying on variation in relative prices. A key assumption in their framework is the existence of an assignable good for men, women, and children in the household, meaning a privately consumed good for which individual consumption is observed.

In addition, their approach assumes that resource shares do not vary with total household expenditure. To close the model, they impose one of two types of preference restrictions on the consumption of the assignable good. The first, referred to as "similar across people" (SAP), requires that preferences for the assignable good are similar across individuals within a given household type. The second, "similar across types" (SAT), requires that preferences for the assignable good are similar across household types for a given individual. Other goods in the consumption bundle may be private, public, or partly shared, but identification relies on how the assignable good links individual behavior to observed household-level choices under these preference restrictions.

These identification ideas have proven highly influential, and subsequent empirical work has applied them in a wide range of settings. For example, Calvi (2020), Penglase (2021),

Bargain, Donni, and Hentati (2022a), and Calvi, Penglase, Tommasi, and Wolf (2023) employ variants of this approach to recover intra-household resource shares using data from India, Bangladesh, Côte d'Ivoire and Mexico. These applications illustrate how the Dunbar–Lewbel–Pendakur framework, under relatively weak data requirements, can be used to uncover detailed patterns of intra-household allocation in diverse economic environments.

Linear estimation procedures. Lechene, Pendakur, and Wolf (2022) seek to promote wider use of these methods by proposing a linear estimator (OLS, SUR, or 2SLS in the presence of endogenous variables) for recovering resource shares within households, including those accruing to children, that is straightforward to implement with standard datasets. Their approach adapts the model of Dunbar, Lewbel, and Pendakur (2013) so that resource shares can be inferred from linear Engel curves for one assignable good per person. This strategy requires additional behavioral assumptions, most notably that individual Engel curves, expressed in budget–share form, are linear in the logarithm of total household expenditure. Together with technical conditions on the household consumption technology, these assumptions yield the following linear household Engel curve for the assignable good of a person of type t in household h:

$$w_h^t \;=\; a^t \;+\; b^t \ln(y_h) \;+\; \epsilon_h^t, \tag{40}$$

where h indexes households, $t \in \{1, \ldots, \bar{t}\}$ indexes demographic types within the household (e.g., adult woman, adult man and child), and the parameters a^t and b^t depend on individual preference parameters and resource shares. Lechene, Pendakur, and Wolf (2022) show that the resource share for type t can then be estimated as

$$\hat{\eta}^t \;=\; \frac{\hat{b}^t}{\sum_{t=1}^{\bar{t}} \hat{b}^t}, \tag{41}$$

where $\bar{t}$ denotes the number of types. These specifications can be extended to include demographic covariates. Importantly, identification of individual resource shares is driven not by the levels of the assignable-good budget shares, but by how these budget shares respond to changes in total household expenditure.

As a concluding remark, it is worth emphasizing that the preference-similarity methods surveyed here are specifically designed to recover individual resource shares rather than individual preferences. This focus is well-suited to many empirical applications, since resource shares provide direct and policy relevant measures of how household resources are distributed among members. This brings us to the next section, which discusses various measures of individual welfare aimed at evaluating inequality and poverty at the individual level.

6 Empirical welfare measurement

Section 2 shows that direct surveys provide valuable descriptive evidence: they document how private spending is allocated across household members, how jointly consumed goods dominate budgets, and how time use is specialized within couples, with substantial cross-country variation. Yet these surveys do not reveal how individuals value public goods or

the utility they derive from private or public consumption, and many widely used datasets report expenditures only at the household level. As a result, direct evidence alone cannot recover the intra-household allocation of resources, and the impact of the latter on individual welfare.

This gap motivates the identification approaches developed in Sections 4 and 5. By imposing theoretically grounded structure on observed time use and household consumption, the collective framework makes it possible to infer individual preferences and resource shares even when person-specific data are unavailable. These recovered allocations form the basis for the individual welfare analysis. In what follows, we outline the welfare concepts implied by the theory and summarize the key empirical findings that emerge when applying them to real-world data.

6.1 Measuring individual welfare

As detailed in earlier sections, a central object that emerges across all identification strategies is the *sharing rule*, which describes how total household resources are allocated among members. Within the collective consumption framework, we distinguish three main types of sharing rules. When all consumption is private and preferences are egoistic, the sharing rule coincides with each member's disposable resources and therefore provides a natural basis for measuring individual welfare, inequality, and poverty.

In many households, however, public consumption is non-negligible. In such settings, multiple welfare concepts become relevant. As discussed in Section 3, one may first consider the *conditional sharing rule*, which captures the allocation to private goods while holding public consumption fixed. Although informative about the distribution of private resources, this measure alone is insufficient for welfare comparisons whenever household members attach different values to the public good. A broader concept is the *general sharing rule*, which incorporates personalized Lindahl prices to reflect each individual's valuation of public consumption. As emphasized by Chiappori and Meghir (2015), both the sharing rule and the associated Lindahl prices are needed to understand intra-household inequality, as together they determine access to resources and the implicit shadow prices individuals face for jointly consumed goods.

To evaluate welfare in the presence of public goods, several complementary measures have been proposed within the collective framework. In contrast to the sharing rule, which describes within-household allocation for a given household, these measures compare an individual's situation inside the household to the hypothetical situation of that same individual living alone. The existence of multiple welfare measures reflects the fact that public consumption complicates such comparisons. Once some consumption is shared, alternative welfare concepts need no longer coincide, as they differ in how the valuation of public consumption enters the welfare assessment.

The first such measure is the *money-metric welfare index* (MMWI), introduced by Chiappori and Meghir (2015) and further developed by Chiappori, Meghir, and Okuyama (2024). The MMWI measures the income an individual would require as a single person, evaluated at market prices for the public good, to attain the same utility level as within the household, where the public good is effectively priced at that individual's Lindahl price. It is therefore a Hicksian measure of welfare, grounded in individual utility and independent of the

preferences of other household members.

A closely related measure is the *indifference scale* proposed by Browning, Chiappori, and Lewbel (2013), which likewise asks how much income an individual would require as a single person to be as well off as within the household. The key difference relative to the MMWI is that indifference scales express cost-of-living comparisons in terms of private-good equivalents defined by the household's consumption technology, which characterizes the private versus public nature of goods; see Section 4.3. For a detailed comparison of MMWIs and indifference scales, see Chiappori, Meghir, and Okuyama (2024).

A third measure, the *relative individual cost of an equivalent bundle* (RICEB), introduced by Cherchye, De Rock, Surana, and Vermeulen (2020), complements the MMWI and indifference scale concepts by asking what share of total household resources an individual would need as a single person to reproduce exactly the same mix of private and public goods consumed within the household. While MMWI measures and indifference scales quantify the income required to maintain a given utility level, the RICEB evaluates the cost of replicating a given consumption bundle. It is therefore a Slutsky-type welfare measure, defined for a fixed consumption level instead of a fixed utility level, which marks a key conceptual distinction from the Hicks-type measures discussed above.

These distinctions matter only in the presence of public goods. Without public consumption, sharing rules, MMWIs, indifference scales, and RICEBs coincide; with public consumption, they offer complementary perspectives on individual welfare and inequality. Together, they constitute a coherent toolkit for translating the structural features of collective models into meaningful welfare assessments. In what follows, we illustrate how this toolkit has been applied empirically and summarize the most salient empirical findings from the literature. For compactness, we focus primarily on results for within-household allocations, as captured by the sharing rules.

6.2 Empirical findings

We synthesize empirical evidence on three key themes that emerge from the literature on intra-household allocation and welfare measurement. Our aim is not to provide an exhaustive survey, but rather present a selected number of studies to highlight several recurring findings that have proven robust across different empirical settings and methodological approaches. The available evidence shows that credible distributional analysis requires moving beyond household-level aggregates and toward individual welfare measures derived from a collective model of household consumption behavior. In doing so, it also connects with the wider literature on gender inequality within families, for which Jayachandran and Voena (2026) provide a recent comprehensive review.

The studies reviewed below also illustrate how the identification strategies discussed in Sections 4 and 5 are used in practice. Some contributions rely on the differential approach and exploit labor-supply responses to wages and nonlabor income to recover resource shares. Others use preference-similarity assumptions to infer intra-household allocations from household-level consumption data. Finally, Cherchye, De Rock, Lewbel, and Vermeulen (2015) combine revealed preference restrictions with estimated demand systems to recover bounds on individual resource shares. Although these approaches rely on different identifying assumptions and different data environments, they often lead to remarkably similar qual-

itative conclusions regarding inequality, poverty, and the determinants of intra-household allocations.

Before turning to these substantive findings, it is worth asking whether allocations recovered from standard household-level data resemble those obtained from unusually rich datasets containing direct information on intra-household resource allocation. Such comparisons provide a useful validation exercise for the identification strategies discussed in this review and for the empirical credibility of the results they generate. Evidence of this kind remains scarce. A notable exception is Bargain, Lacroix, and Tiberti (2022b), who exploit a unique dataset from Bangladesh containing individualized expenditure data for all household members. Using these data, they compare observed resource shares with those predicted by several collective-model identification approaches based on assignable goods and preference-similarity assumptions. Their results show that the collective model predicts intra-household resource allocation reasonably well, especially when clothing is used as the assignable good. More generally, their findings provide encouraging evidence for the empirical validity of the methods reviewed in this article and for the broader conclusions emerging from the literature.

Inequality and poverty. A first central insight is that a substantial share of overall inequality arises *within* households. Treating the household as a single decision-making welfare unit often masks considerable dispersion between spouses or between adults and children. This is nicely illustrated by Lise and Seitz (2011), who focus on data drawn from the U.K. Family Expenditure Survey. Using a structural collective model identified through the differential approach discussed in Section 4.1, they exploit information on wages, labor supply, and household consumption to recover intra-household allocations. They demonstrate that standard methods based on household-level consumption and equivalence scales miss pronounced differences in welfare across partners. Focusing on the period from 1968 to 2001, they show that conventional equivalence-scale adjustments understate the level of cross-sectional consumption inequality, measured by the variance of log consumption, by roughly 50% at the beginning of the sample period. Their decomposition of total consumption inequality into between-household and within-household components further reveals that within-household disparities were quantitatively important: in the late 1960s and early 1970s, about one half of total consumption inequality originated within households, while by 2000 this share had declined to about one quarter. In addition, they document that changes in marital sorting on wages and hours worked simultaneously lead to a substantial decline in within-household inequality and a substantial rise in between-household inequality.

A second insight concerns poverty measurement. When individual resource shares are estimated, rather than assuming equal resource sharing within households, poverty rates can shift substantially, often in ways that reveal gender- and age-specific vulnerabilities that are obscured by household-level statistics. Using data from the Integrated Household Survey (IHS) for Malawi, Dunbar, Lewbel, and Pendakur (2013) adopt a preference-similarity approach that we discussed in Section 5.3 and which is based on assignable goods to estimate resource shares. They estimate that the first child receives roughly 20% of household resources, with each subsequent child receiving an additional 5–10 percentage points. Under the World Bank's \$2/day per-capita poverty measure, which assumes equal resource shares across individuals, overall poverty is estimated at 91%. By contrast, allowing for unequal

intra-household sharing, using estimated individual resource shares, yields poverty rates of about 60% for men, 85% for women, and more than 95% for children, highlighting stark differences in individual poverty across household members.

Lechene, Pendakur, and Wolf (2022), which similarly rely on a preference-similarity method discussed in Section 5.3, use data drawn from twelve household surveys from, among others, the World Bank Living Standards Measurement Study (LSMS). They reject equal sharing and uncover systematic gender gaps. For instance, in Bangladesh, women's resource shares are about 5 percentage points lower than men's, implying roughly a 9-point higher poverty incidence; in Iraq, the corresponding gaps are around 4 points.

Similarly, using US PSID data, Cherchye, De Rock, Lewbel, and Vermeulen (2015) combine revealed preference methods with estimated demand systems, as discussed in Section 5.2, to recover bounds on individual resource shares. Their findings confirm that individual-level inequality can differ sharply from what household-level data imply. In particular, relative wages emerge as a key determinant of intra-household allocations, with higher individual earnings consistently associated with larger resource shares. They show that household-level poverty measures tend to understate individual poverty. For childless couples, for example, while 11% of households fall below the household-level poverty line, the corresponding individual-level poverty rate lies between 15% and 18%, with women substantially more frequently in poverty than men.

The studies reviewed above reject the view that household resources are shared equally among members and conclude that household-level measures tend to understate both inequality and poverty at the level of individuals. A striking feature of the evidence is the magnitude of the differences between household-level and individual-level measures. Depending on the country, dataset, and identification strategy, individual poverty rates can exceed household poverty rates by several percentage points, while measures of inequality based on household-level consumption may substantially understate the dispersion of welfare across individuals. Although the precise magnitudes differ across applications, the direction of the bias is remarkably consistent.

Determinants of resource sharing. An obvious question in then what the driving factors are for these unequal intra-household allocations. Empirical work shows that individual resource shares respond systematically to relative earnings, demographic composition, and institutional features such as marriage market conditions or divorce legislation.

Using the differential approach, Chiappori, Fortin, and Lacroix (2002) demonstrate, based on PSID data, that when the marriage market becomes more favorable to women (for instance, through a low female-to-male sex ratio or more protective divorce laws), women's bargaining power increases, which in turn raises their share of household resources and induces corresponding adjustments in household labor supply. In essence, shifts in distribution factors that strengthen women's outside options translate into a larger fraction of household income being allocated to them.

Using the preference-similarity approach of Browning, Chiappori, and Lewbel (2013), which we discussed in Section 4.3, and applying it to a series of Dutch household budget surveys, Cherchye, De Rock, and Vermeulen (2012a) find substantial economies of scale arising from joint consumption among elderly couples. This reflects the fact that individuals in

couples can maintain living standards at lower expenditure levels than singles. Their analysis further shows that wives' resource shares increase with total household expenditures, implying that in better-off elderly couples, women command a rising share of household resources. These findings indicate that intra-household allocations respond not only to relative earnings, but also to institutional factors, demographic characteristics, and the extent of economies of scale generated by joint consumption.

Public policy. The evidence reviewed in this section shows that household-level statistics provide an incomplete picture of the distribution of welfare within society. Particularly noteworthy is the consistency of the evidence across countries, datasets, and identification strategies. Policy interventions designed on the basis of household-level statistics may therefore fail to reach those most in need if resources are unevenly distributed within families. For example, a means-tested transfer based on per-capita household income may classify a household as non-poor, while women or children within that household face very low individual consumption levels. In contrast, using individual welfare measures derived from collective models allows policymakers to identify such hidden vulnerabilities, design better-targeted programs, and anticipate distributional effects within households, thereby improving both the efficiency and the equity of social policy.

A well-known illustration is provided by Lundberg, Pollak, and Wales (1997), who show that transferring child allowance specifically to mothers leads to a shift in spending toward women's and children's clothing relative to men's clothing, coinciding with the redistribution of income within the household. The collective framework provides a natural interpretation of this finding: by altering the intra-household distribution of resources, policies can affect both consumption patterns and individual welfare. More generally, the empirical evidence reviewed above suggests that policy evaluations based exclusively on household-level measures may overlook important distributional effects occurring within families. These concerns become even more pronounced once time use and domestic production are taken into account, which we discuss next.

6.3 Time allocation and home production

The welfare measures outlined above were originally developed for settings where individual welfare depends primarily on private and public consumption goods. Yet, as the substantial specialization in time use documented in Section 2 makes clear, time allocation is itself an essential component of individual welfare. A household member who benefits from their partner's childcare and meal preparation faces a higher cost of achieving the same welfare level when single, as they must either purchase market substitutes or invest their own time in these activities. Likewise, a lower monetary share may be partly offset by greater leisure. Ignoring time use can therefore misrepresent both the level and distribution of individual welfare.

Incorporating time use modifies welfare comparisons in important ways. First, it affects the measurement of resource shares: the sharing rule determines not only the division of money income but also the allocation of time-intensive production responsibilities. Second, it affects the mapping from resource shares to utility, since time use enters preferences. As

a result, poverty assessments based solely on consumption expenditures may diverge substantially from those based on full-income utility, a distinction that is especially relevant in households with children, where specialization in market and domestic work is most pronounced.

These observations also have direct implications for policy. Intra-household inequality may arise not only from unequal financial control but also from asymmetric time burdens. Consequently, evaluating reforms that influence labor supply, childcare provision, or parental leave requires a framework that jointly accounts for consumption, labor supply, and domestic production. Collective models with home production provide such a framework by allowing time inputs and market expenditures to interact through a household production technology.

As discussed in Section 5, the collective model naturally accommodates home production by treating domestically produced goods as household public goods generated from time and market inputs. Embedding these goods into sharing rules, MMWIs, indifference scales, and RICEBs yields an individual welfare analysis that captures the multiple channels through which time use and domestic production shape well-being. However, despite the conceptual feasibility of these extensions, empirical applications remain limited. The few notable examples reviewed in Section 5 demonstrate what is possible, but they also illustrate how much scope remains for further empirical work. Developing more applications that combine expenditure data, time use information, and domestic production represents a valuable and promising direction for future research.

7 Conclusion: looking ahead

Standard surveys typically record only aggregate household expenditures, partial information on time use, and no direct evidence on how public goods benefit each member. These limitations obscure the welfare of individuals who live inside households, especially when members differ in needs, preferences, or bargaining positions. Building on the descriptive evidence and the theoretical collective model, we have reviewed alternative approaches that recover unobserved intra-household allocation from widely used data. In particular, the collective approach provides the conceptual tools to interpret household choices through the lens of individual preferences and to recover the sharing rule that maps total resources into the shares received by each member.

The identification strategies surveyed in this article (i.e., differential, revealed preference and preference-similarity approaches) show that individual welfare can be recovered from broadly available datasets. We stipulated the assumptions needed for linking observed choices to unobserved allocations. These strategies clarify the conditions under which point identification is attainable and when the data and assumptions instead support only set identification, yielding informative bounds on individual welfare measures. Together, they complement the direct survey evidence reviewed earlier by showing precisely what can and cannot be inferred from household-level observations and by placing the sharing rule at the center of the empirical analysis.

A consistent empirical message is that a sizable part of inequality originates within households rather than between them. Gender differences in resource shares are also pervasive and have direct implications for inequality and for policy targeting. Moreover, the division

of resources within families responds systematically to individuals' relative incomes and distribution factors, so that policy reforms or labor market shifts can reallocate resources inside the household even when total household income is unchanged.

Notably, the collective framework is sufficiently flexible to accommodate rich household environments that feature time use, home production, and domestic production technologies. However, empirical applications that fully exploit these extensions remain scarce. This gap highlights the value of applying the identification strategies reviewed above to settings in which both money and time matter, an important direction for future research.

The analysis also underscores the importance of improving measurement. Collecting individual-level information on consumption and time use is inherently challenging and prone to error. In some surveys, a single household member reports the consumption shares of all others, which may generate inaccuracies if individuals assess their own consumption more precisely than that of other members. There is also evidence that income and consumption may be concealed within households in ways correlated with intra-household bargaining power, suggesting that misreporting may be systematic rather than random. These concerns underscore the need for careful survey design, appropriate respondent selection, clear item definitions, recall aids, and structured validation procedures to ensure that the data used for welfare measurement accurately reflect intra-household behavior.

While some newer surveys collect individual-level expenditures and detailed time use information, comprehensive datasets that track public goods, home-produced goods, and private consumption for all household members remain rare. Improvements in instruments and targeted modules can substantially reduce mismeasurement and strengthen identification. Even relatively modest improvements (such as collecting person-specific spending for a small number of assignable goods, distinguishing private from public consumption, or adding concise time use modules to budget surveys) can further enhance the application of the identification strategies reviewed in this article.

Better data strengthen empirical identification and naturally call for richer models and more powerful theoretical tools. Building on improved measurement, a natural next step is to develop structural frameworks that can explain the mechanisms underlying intra-household allocations and their heterogeneity across individuals and households. In particular, embedding collective household models within a broader context that features marriage market forces, partners' outside options, and labor supply choices provides a structural explanation for observed, and often unequal, intra-household allocations and how these translate into individual welfare outcomes. By explicitly modeling spouses' positions in the marriage market, researchers can obtain deeper insight into the determinants of household decision-making. This perspective is directly related to the marriage market-based identification approach discussed in Section 5.2 above and provides structural explanations for the determinants of matching patterns, bargaining positions, and intra-household allocations. See also Choo, Seitz, and Siow (2008), Chiappori (2017), Chiappori, Dias, and Meghir (2018), Galichon, Kominers, and Weber (2019), and Weber (2022) for structural models that offer motivating illustrations. For example, this marriage market approach helps explain matching patterns and the distribution of characteristics that individuals bring into the household. It does so by recognizing that equilibrium matches form where couples generate the highest joint surplus, given their traits and potential complementarities. At the same time, each individual's outside options in the wider marriage market shape their bargaining position within

the household. As a result, the distribution of characteristics and opportunities in the population naturally determines who matches with whom and how resources are subsequently allocated within households.

A closely related avenue is to move beyond the current focus on static models. Many elements of individual welfare, including the valuation of public goods (such as child related expenditures), time use, labor supply, wealth accumulation, and access to resources, evolve over the life cycle. Adopting a life cycle perspective also allows researchers to model key life defining decisions that occur only once or a few times in a lifetime, such as marriage, divorce, and fertility choices. Recognizing that allocations reflect shifting outside options, intra-household commitment to intertemporal decisions, policy interventions, and exposure to shocks underscores the value of dynamic collective models for connecting cross-sectional measurement to life cycle welfare analysis. We refer to Mazzocco (2007) and Chiappori and Mazzocco (2017) for broader discussions of methodological challenges, and to Voena (2015), Lafortune and Low (2023), Reynoso (2024), and Theloudis, Velilla, Chiappori, Giménez-Nadal, and Molina (2025) for recent applications. Dynamic collective models make it possible to examine how policies, outside options, commitment frictions, and life cycle shocks shape intertemporal choices and the distribution of resources within households. In doing so, they clarify how intra-household welfare evolves over time and how policy interventions translate into individual outcomes in a dynamic setting.

The overarching message is that measuring welfare at the level where it is experienced, the individual, fundamentally reshapes what we learn about inequality, poverty, and the distributional effects of policies. Once time use, home production, and shared consumption are taken seriously, the mapping from household resources to individual welfare becomes both richer and more accurate. By combining direct descriptive evidence with a coherent yet tractable theoretical framework and data-matched identification methods, researchers can recover who gains, who loses, and why within families. We hope this article serves as a practical guide for applied researchers who seek to incorporate intra-household heterogeneity into their work and as an invitation to further develop the conceptual and empirical tools needed to understand inequality at its most fundamental level.

References

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

Aminjonov, Ulugbek, Olivier Bargain, and Maira Colacce. 2025. "Global evidence on gender gaps and child poverty in consumption." *IZA Discussion Paper* (17767): 64p.

Angrist, Joshua. 2002. "How do sex ratios affect marriage and labor markets? Evidence from America's second generation." *Quarterly Journal of Economics* 117 (3): 997–1038.

Ashraf, Nava. 2009. "Spousal control and intra-household decision making: An experimental study in the Philippines." *American Economic Review* 99 1245–1277.

Attanasio, Orazio, and Valérie Lechene. 2014. "Efficient responses to targeted cash transfers." *Journal of Political Economy* 122 (1): 252–291.

Bargain, Olivier, and Olivier Donni. 2012. "Expenditure on Children: A Rothbarth-type Method Consistent with Scale Economies and Parents' Bargaining." *European Economic Review* 56 (4): 792–813.

Bargain, Olivier, Olivier Donni, and Imen Hentati. 2022a. "Resource sharing in households with children: A generalized model and empirical evidence from the UK." *Journal of the European Economic Association* 20 (6): 2468–2496.

Bargain, Olivier, Guy Lacroix, and Luca Tiberti. 2022b. "Intrahousehold resource allocation and individual poverty: Assessing collective model predictions using direct evidence on sharing." *The Economic Journal* 132 (643): 865–905.

Becker, Gary. 1965. "A theory of the allocation of time." *Economic Journal* 75 (299): 493–517.

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

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

Blundell, Richard, Pierre-André Chiappori, Thierry Magnac, and Costas Meghir. 2007. "Collective Labour Supply: Heterogeneity and Non-Participation." *Review of Economic Studies* 74 (2): 417–445.

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." *Review of Economic Studies* 76 503–528.

Browning, Martin, François Bourguignon, Pierre-André Chiappori, and Valérie Lechene. 1994. "Income and outcomes: A structural model of intrahousehold allocation." *Journal of Political Economy* 102 1067–1096.

Browning, Martin, Pierre-André Chiappori, and Valérie Lechene. 2006. "Collective and Unitary Models: A Clarification." *Review of Economics of the Household* 4 (1): 5–14.

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 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.

Browning, Martin, Pierre-André Chiappori, and Yoram Weiss. 2014. *Family Economics*. Cambridge University Press.

Burda, Michael, Daniel Hamermesh, and Philippe Weil. 2013. "Total work and gender: facts and possible explanations." *Journal of Population Economics* 26 (4): 239–261.

Calvi, Rossella. 2020. "Why are older women missing in India? The age profile of bargaining power and poverty." *Journal of Political Economy* 128 (7): 2453–2501.

Calvi, Rossella, Jacob Penglase, Denni Tommasi, and Alexander Wolf. 2023. "The more the poorer? Resource sharing and scale economies in large families." *Journal of Development Economics* 160(C).

Capéau, Bart, Laurens Cherchye, Koen Decancq, André Decoster, François Maniquet, Annemie Nys, Guillaume Périlleux, Eve Ramaekers, Zoé Rongé, Erik Schokkaert, and Frederic Vermeulen. 2020. *Well-being in Belgium: Beyond Happiness and Income*. Series in Economic Studies in Inequality, Social Exclusion and Well-Being, Springer.

Cherchye, Laurens, Pierre-André Chiappori, Bram De Rock, Charlotte Ringdal, and Frederic Vermeulen. 2026. "Feed the children." *The Economic Journal* 136 (678): 1983–2015.

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, and Khushboo Surana. 2025. "Labor market participation, marriage and individual welfare." *Journal of Political Economy: Mircoeconomics* forthcoming.

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. 2011a. "Collective Household Consumption Behavior: Revealed Preference Analysis." *Foundations and Trends in Econometrics* 4 (4): 225–312.

Cherchye, Laurens, Bram De Rock, and Frederic Vermeulen. 2011b. "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. "Economic well-being and poverty among the elderly: an analysis based on a collective consumption model." *European Economic Review* 56 (6): 985–1000.

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.

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

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

Cherchye, Laurens, and Frederic Vermeulen. 2008. "Nonparametric analysis of household labor supply: goodness of fit and power of the unitary and the collective model." *The Review of Economics and Statistics* 90 (2): 267–274.

Chiappori, Pierre-André, Monica Costa Dias, and Costas Meghir. 2018. "The marriage market, labor supply, and education choice." *Journal of Political Economy* 126 (S1): S26–S72.

Chiappori, Pierre-André, Costas Meghir, and Yoko Okuyama. 2024. "Intrahousehold welfare inequality and household public goods."Technical report, National Bureau of Economic Research.

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. 1997. "Introducing household production in collective models of labor supply." *Journal of Political Economy* 105 (1): 191–209.

Chiappori, Pierre-André. 2017. *Matching with Transfers: The Economics of Love and Marriage*. Princeton University Press.

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

Chiappori, Pierre-André, Bernard Fortin, and Guy Lacroix. 2002. "Marriage market, divorce legislation and household labor supply." *Journal of Political Economy* 110 37–72.

Chiappori, Pierre-André, and Maurizio Mazzocco. 2017. "Static and intertemporal household decisions." *Journal of Economic Literature* 55 (3): 985–1045.

Chiappori, Pierre-André, and Costas Meghir. 2015. *Intrahousehold inequality*. Volume 2. 1369–1418, Elsevier.

Choo, Eugene, Shannon Seitz, and Aloysius Siow. 2008. "The Collective Marriage Matching Model: Identification, Estimation and Testing." *Univ. of Toronto Working Paper* 340 1–65.

Choukhmane, Taha, Lucas Goodman, and Cormac O'Dea. 2025. "Efficiency in household decision-making: Evidence from the retirement savings of US couples." *American Economic Review* 115 1485–1519.

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.

Deaton, Angus, and John Muellbauer. 1980. *Economics and Consumer Behavior*. Cambridge University Press.

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

Donni, Olivier. 2003. "Collective household labor supply: Nonparticipation and income taxation." *Journal of Public Economics* 87 (5-6): 1179–1198.

Duflo, Esther. 2003. "Grandmothers and Granddaughters: Old-age Pensions and Intrahousehold Allocation in South Africa." *World Bank Economic Review* 17 (1): 1–25.

Dunbar, Geoffrey, Arthur Lewbel, and Krishna Pendakur. 2013. "Children's resources in collective households." *American Economic Review* 103 (1): 438–471.

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

Galichon, Alfred, Scott Duke Kominers, and Simon Weber. 2019. "Costly concessions: An empirical framework for matching with imperfectly transferable utility." *Journal of Political Economy* 127 (6): 2875–2925.

Gini, Corrado. 1912. "Variabilità e mutabilità: contributo allo studio delle distribuzioni e delle relazioni statistiche." *Studi Economico-Giuridici della Facoltà di Giurisprudenza della Regia Università di Cagliari* 3 (2): .

Grossbard-Shechtman, Shoshana. 1993. *On the Economics of Marriage. A Theory of Marriage, Labor, and Divorce*. Routledge (published originally by Westview Press).

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

Jayachandran, Seema, and Alessandra Voena. 2026. "Women's Power in the Household." *Journal of Economic Literature* forthcoming.

Lafortune, Jeanne, and Corinne Low. 2023. "Collateralized Marriage." *American Economic Journal: Applied Economics* 15 (4): 252–291.

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

Lechene, Valérie, Krishna Pendakur, and Alex Wolf. 2022. "Ordinary least squares estimation of the intrahousehold distribution of expenditure." *Journal of Political Economy* 130 (3): 681–731.

Lewbel, Arthur. 2019. "The identification zoo: meanings of identification in econometrics." *Journal of Economic Literature* 57 835–903.

Lewbel, Arthur, and Krishna Pendakur. 2008. "Estimation of collective household models with Engel curves." *Journal of Econometrics* 147 (2): 350–358.

Lindahl, Erik. 1958. "Just taxation—a positive solution." In *Classics in the theory of public finance*, 168–176, Springer.

Lise, Jeremy, and Shannon Seitz. 2011. "Consumption inequality and intra-household allocations." *Review of Economic Studies* 78 328–355.

Lise, Jeremy, and Ken Yamada. 2019. "Household sharing and commitment: Evidence from panel data on individual expenditures and time use." *Review of Economic Studies* 86 2184–2219.

Lorenz, Max Otto. 1905. "Methods of measuring the concentration of wealth." *Publications of the American Statistical Association* 9 (70): 209–2019.

Lundberg, Shelly J, Robert A Pollak, and Terence J Wales. 1997. "Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit." *Journal of Human Resources* 32 463–480.

Marx, Karl. 1867. *Das Kapital. Kritik der politischen Ökonomie. Volume I.* Otto Meisner.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1996. *Microeconomic Analysis*. Oxford University Press.

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

Penglase, Jacob. 2021. "Consumption inequality among children: evidence from child fostering in Malawi." *Economic Journal* 131 1000–1025.

Rangel, Marcos, and Duncan Thomas. 2019. "Decision-making in complex households."Technical report, National Bureau of Economic Research.

Reynoso, Ana. 2024. "The Impact of Divorce Laws on the Equilibrium in the Marriage Market." *Journal of Political Economy* 132 (12): 4155–4204. 10.1086/732532.

Ricardo, David. 1817. *On the Principles of Political Economy and Taxation*. John Murray.

Samuelson, Paul. 1954. "The Pure Theory of Public Expenditure." *Review of Economics and Statistics* 36 386–389.

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

Samuelson, Paul A. 1947. *Foundations of Economic Analysis*. Harvard University Press.

Theloudis, Alexandros, Jorge Velilla, Pierre-Andre Chiappori, José I. Giménez-Nadal, and José A. Molina. 2025. "Commitment and the Dynamics of Household Labour Supply." *The Economic Journal* 135 (665): 354–386. 10.1093/ej/ueae065.

Thomas, Duncan. 1990. "Intra-household Resource Allocation: An Inferential Approach." *Journal of Human Resources* 25 635–664.

Udry, Chris. 1996. "Gender, agricultural production, and the theory of the household." *Journal of Political Economy* 104 1010–1046.

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

Voena, Alessandra. 2015. "Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples?" *American Economic Review* 105 (8): 2295–2332.

Weber, Max. 1922. *Economy and Society. A New Translation [2019]*. Harvard University Press.

Weber, Simon. 2022. "Collective models and the marriage market." *arXiv preprint arXiv:2211.07416*.

Zhang, Sally. 2024. "Hidden in plain sight: Asymmetric information and hidden income within the household." *Available at SSRN 4883384*.