KU LEUVEN

CENTER FOR ECONOMIC STUDIES

DISCUSSION PAPER SERIES
DPS14.08

MARCH 2014

Household consumption when the marriage is stable

Laurens CHERCHYE, Thomas DEMUYNCK, Bram DE ROCK and Frederic VERMEULEN

Public Economics

Faculty of Economics
And Business

Household consumption when the marriage is stable*

Laurens Cherchye† Thomas Demuynck‡ Bram De Rock§
Frederic Vermeulen¶

July 6, 2015

Abstract

We develop a novel framework to analyze the structural implications of the marriage market for household consumption. We define a revealed preference characterization of efficient household consumption when the marriage is stable. Stability means that the marriage matching is individually rational and has no blocking pairs. We characterize stable marriage with intrahousehold (consumption) transfers but without assuming transferable utility. We show that our revealed preference characterization generates testable conditions even with a single consumption observation per household and heterogeneous individual preferences across households. The characterization also allows for identifying the intrahousehold decision structure (including the sharing rule) under the same minimalistic assumptions. An application to Dutch household data demonstrates the usefulness of our theoretical results. We find that the female gets a higher income share when her relative wage increases, which we can give a structural interpretation in terms of outside options from marriage that vary with individual wages.

*We thank Martin Browning, Jan Eeckhout, Patrick Legros, Arthur Lewbel, Jeremy Lise, Rika Ponnet and seminar participants in Antwerp, Barcelona, Bristol, Brussels, Ca'Foscari Venice, the Chinese University of Hong Kong, Ghent, Leicester, Leuven, Louvain-la-Neuve, Namur, Oxford, Tilburg, Toulouse, University College London, ESEM 2014, RES 2015 and the CEREC, MEQIN and THRED Workshops for useful suggestions.

†Center for Economic Studies, University of Leuven. E. Sabbelaan 53, B-8500 Kortrijk, Belgium. E-mail: laurens.cherchye@kuleuven.be. Laurens Cherchye gratefully acknowledges the European Research Council (ERC) for his Consolidator Grant 614221 and the Research Fund K.U.Leuven for the grant STRT1/08/004.

‡Maastricht University, Tongersestraat 53, 6370 Maastricht, Netherlands. Email: t.demuynck@maastrichtuniversity.nl.

§ECARES, Université Libre de Bruxelles. Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium. E-mail: bderock@ulb.ac.be. Bram De Rock gratefully acknowledges the European Research Council (ERC) for his Starting Grant 263707.

¶Department of Economics, University of Leuven, Naamsestraat 69, B-3000 Leuven, Belgium. E-mail: frederic.vermeulen@kuleuven.be. Frederic Vermeulen gratefully acknowledges financial support from the Research Fund K.U.Leuven through the grant STRT/12/001 and from the FWO through the grant G057314N.

JEL Classification: C14, D11, C78.
Keywords: marriage market, stable matching, Pareto efficient household consumption, testable implications, sharing rule identification, preference heterogeneity.

1 Introduction

We introduce a novel structural framework to study the implications of the marriage market for observed household consumption behavior. In particular, if we assume that a marriage matching is stable, does this generate testable implications for the observed consumption patterns? And, if so, can we use these testable implications to identify the within-household decision structure (including the so-called sharing rule) underlying this observed consumption? The remainder of this introductory section explains our research question in more detail, and positions our contribution in the literature.

Nonunitary household consumption and the sharing rule. This study fits within the nonunitary approach to modeling household consumption behavior. Nonunitary models of household consumption are to be contrasted with the more standard unitary model, which describes the household as if it were a single decision maker. Clearly, this unitary model is conceptually problematic in the case of multi-person households. Next, we also find that the unitary model does not provide a good empirical fit of multi-person household consumption behavior. In particular, the testable implications of the model are usually rejected when brought to data of multi-person households. Importantly, these conditions are typically not rejected for single-person households, which suggests that something is wrong with the implicit preference aggregation assumptions that underlie the unitary modeling of multi-person consumption behavior. See, for example, Browning and Chiappori (1998), Cherchye and Vermeulen (2008) and Cherchye, De Rock and Vermeulen (2009).

In response to these problems associated with the unitary model, Chiappori (1988, 1992) proposed the nonunitary "collective" model of household consumption. A distinguishing feature of this collective model is that it explicitly recognizes the multi-person nature of multi-person households. In particular, it assumes that multi-person households consist of multiple decision makers with their own rational preferences. Observed household consumption is then regarded as the outcome of a within-household bargaining process between these different decision makers. As for this interaction process, Chiappori's collective model (only) assumes that it yields Pareto efficient intrahousehold allocations. Attractively, the collective model does give a good fit of multi-person consumption data. See Browning and Chiappori (1998), Cherchye and Vermeulen (2008), Cherchye, De Rock and Vermeulen (2009) and Attanasio and Lechene (2014).

Our following analysis will assume that households behave in accordance with the collective consumption model (i.e. make Pareto efficient decisions). In particular, we assume a collective model that includes publicly as well as privately consumed goods.

Public consumption is particularly relevant in our context of marriage matching, as it generates gains from marriage. As for the privately consumed goods, we take the minimalistic prior that the empirical analyst only observes the aggregate household consumption and, so, does not know who consumes what within the household. Indeed, budget surveys typically do not contain information on the intrahousehold sharing of consumption quantities. As a matter of fact, an important issue in our following analysis will be to identify the intrahousehold sharing of resources that underlies the observed household consumption. Within the collective consumption literature, this sharing is summarized in terms of the so-called "sharing rule".

Formally, this sharing rule concept is intrinsic to the decentralized representation of rational consumption behavior in terms of a collective model. Essentially, this two-step representation is an application of the second fundamental theorem of welfare economics, which states that any Pareto efficient allocation can be represented as if it were the outcome of a two-step allocation process. In the first step, individual household members divide the household income among each other, which defines individual income shares. In the second step, each individual household member maximizes her/his utility subject to her/his individual budget constraint (using personalized "Lindahl" prices for evaluating the publicly consumed goods).

Within this representation, the sharing rule pertains to the first step, and defines the within-household sharing of resources. Typically, the sharing rule is not observed (i.e. individual shares of private goods or individual Lindahl prices for the public goods are unknown). Within the literature on collective consumption models, a main focus has been on identifying this sharing rule from observed household consumption behavior. If we can identify the sharing rule, then we can address a series of questions that are specific to the nonunitary modeling of household consumption behavior. For example, identifying individual incomes allows for welfare assessments (such as poverty and income inequality analysis) at the level of individuals within households, rather than aggregate households. Next, the sharing rule is often used as an indicator of individual bargaining power, i.e. a higher relative income share for a particular individual signals a better intrahousehold bargaining position. From this perspective, identifying individual income shares also provides insight into the within-household distribution of individual bargaining power.[1]

Sharing rule identification and the marriage market. In what follows, a main focus will be on sharing rule identification from observed (aggregate) household level consumption patterns. However, the approach that we follow is fundamentally different from the usual approach in the collective consumption literature. Basically, the usual

[1] See, for example, Browning, Bourguignon, Chiappori and Lechene (1994), Chiappori, Fortin and Lacroix (2002), Blundell, Chiappori and Meghir (2005), Lewbel and Pendakur (2008), Couprie, Peluso and Trannoy (2010), Lise and Seitz (2011), Bargain and Donni (2012), Cherchye, De Rock and Vermeulen (2012), Browning, Chiappori and Lewbel (2013), Dunbar, Lewbel and Pendakur (2013), and Cherchye, De Rock, Lewbel and Vermeulen (2015) for various applications of the collective consumption model that make use of the sharing rule concept.

approach (only) exploits the assumption that intrahousehold consumption is Pareto efficient (i.e. rational in terms of the collective model). It then shows that Pareto efficiency has testable implications as soon as one can use multiple consumption observations for one and the same household (e.g. a household demand function). If household demand satisfies these empirical restrictions of Pareto efficiency, we can use these restrictions to identify the within-household sharing of resources. Essentially, this obtains intrahousehold sharing rule identification under the maintained assumption of Pareto efficiency (i.e. collective rationality is the identifying hypothesis). See, for example, Chiappori and Ekeland (2006, 2009), Cherchye, De Rock and Vermeulen (2011), Dunbar, Lewbel and Pendakur (2013), and Cherchye, De Rock, Lewbel and Vermeulen (2015) for recent results that fit in this approach.

Our approach is very different from the usual one. Basically, we "endogenize" the marriage matching decisions in the household consumption analysis. Starting from a set of consumption observations for different households, we assume stable marriage in addition to Pareto efficient household consumption.[2] We will show that combining these two assumptions generates strong testable implications for household consumption patterns. In particular, these implications have empirical bite even in the limiting case with a cross-section containing (only) a single observation per household and when accounting for any heterogeneity across households (in terms of individual preferences and the within-household decision process). If these restrictions cannot be rejected, then they usefully allow for informative sharing rule identification under the same minimalistic conditions. Specifically, we will define bounds on individual income shares that are consistent with Pareto efficiency and stable marriage, which effectively "set" identifies the sharing rule. For ease of exposition, we will introduce our main theoretical results under the maintained assumption of frictionless matching, which means that divorce/remarriage is costless. Subsequently, we will also indicate how we can account for costs of divorce in practical applications (including our own application in Section 4). As we will explain, this cost of divorce may not only incorporate frictions on the marriage market but also unobserved benefits from marriage (including match-specific quality such as love).

The basic idea underlying our approach is that within-household bargaining positions (and, thus, individual income shares) are essentially defined by individuals' outside options, which pertain to the possibility to divorce (i.e. exit marriage) and stay single or remarry. Thus, if we put particular structure on marriage, we can actually incorporate these outside options within our model of household consumption. In this study, we assume that marriages are stable (i.e. no household member has an incentive to exit marriage), and show that this effectively does imply particular restrictions on observed household consumption. In turn, this allows us to identify the within-household decision structure underlying the observed household consumption.

[2] See the seminal papers of Gale and Shapley (1962), Shapley and Shubik (1972) and Becker (1973) for early contributions on the concept of stable marriage. Browning, Chiappori and Weiss (2014, Chapters 7 and 8) provide a recent account of the literature on stable matching on the marriage market. Chiappori and Salanié (2015) review the literature on the econometrics of matching models.

At this point, we emphasize that our framework can also be used to recover other fundamentals of the intrahousehold interaction process (such as individual preferences), in addition to the sharing rule. However, to focus our discussion, and given its prominent position in the literature on collective consumption models, our central focus here will be on identifying intrahousehold resource shares.

Outline. Before entering our analysis, we indicate two specific features of the approach we follow here. First, to address our central research question, we develop a characterization of efficient household consumption under stable marriage that follows the revealed preference tradition of Samuelson (1938, 1948), Houthakker (1950), Afriat (1967), Diewert (1973) and Varian (1982). An attractive feature of this revealed preference characterization is that it is intrinsically nonparametric: its empirical implementation does not require an (explicit or implicit) functional specification of individual utilities. This nonparametric orientation minimizes the risk of specification error, i.e. drawing erroneous empirical conclusions because of a wrongly specified functional form. We will show that, despite this fully nonparametric nature, our characterization does allow for a very informative empirical analysis. As such, the empirical methodology that we develop below significantly extends earlier work in Cherchye, De Rock and Vermeulen (2007, 2009, 2011), by explicitly integrating the marriage market in the analysis of Pareto efficient (or collectively rational) household consumption.

A second particular feature of our analysis implies an important difference with the existing literature on characterizing stable marriage. By construction, because we account for consumption sharing within the household, we consider intrahousehold transfers. However, in contrast to earlier studies, we do so without making the usual assumption that individual utilities are transferable.[3] Indeed, it is well-documented

[3]See, for example, Choo and Siow (2006), Chiappori, Oreffice and Quintana-Domeque (2012), Chiappori, Iyigun, Lafortune and Weiss (2013), Jacquemet and Robin (2013), Dupuy and Galichon (2014), Galichon and Salanié (2014) and Choo (2014) for theoretical and empirical analyses of stable marriage under the assumption that individual utilities are transferable. In this respect, another interesting study is the recent one of Echenique, Lee, Shum and Yenmez (2013; see also Echenique, 2008), who provide a revealed preference characterization of stable marriage that is close in spirit to the one that we develop below. However, these authors restrict attention to two polar cases, i.e. the case with transfers and transferable utility and the case without transfers and no transferable utility. By contrast, as we explained, our study considers stable marriage with transfers but no transferable utility. In this sense, it provides a useful complement to the one of Echenique, Lee, Shum and Yenmez (2013). Chiappori and Reny (2006), Legros and Newman (2007) and Browning, Chiappori and Weiss (2014, Chapter 7.3), for example, consider a setting that is formally close to ours (i.e. with transfers but no transferable utility). But these authors focus on theoretical conditions for monotone (assortative) matching patterns, whereas our interest is in the empirical implications of stable matching for household consumption. Choo and Seitz (2013) and Galichon, Kominers and Weber (2014) also consider the empirical analysis of stable marriage with transfers but no transferable utility. A main distinguishing feature of our approach is that we start from a revealed preference characterization that is intrinsically nonparametric, which -we show- enables an informative empirical analysis even with (only) a single consumption observation per household and heterogenous individual preferences across households.

that such transferable utilities imply substantial (and often unrealistic) structure for the individual preferences (i.e. they need to be of the generalized quasi-linear form; see, for example, Chiappori, 2010, and Cherchye, Demuynck and De Rock, 2014, for recent discussions). In what follows, we consider intrahousehold transfers but make no stronger assumptions for individual preferences than the standard ones in collective consumption analyses (i.e. we assume individual utility functions that are continuous, concave and increasing in their arguments).

The remainder of this study unfolds as follows. In Section 2, we introduce our notation and formally define our concept of stable marriage. Section 3 then provides the corresponding revealed preference characterization. Here, we also show that this characterization implies testable implications that are easy to operationalize for observational household consumption data. In addition, we will indicate that these testable implications provide a useful basis to address sharing rule (set) identification. Section 4 presents an empirical application to Dutch household consumption data, which demonstrates the empirical usefulness of our revealed preference methodology. In particular, this application shows that our testable conditions do have empirical bite even in the limiting scenario with only a single consumption observation per household and heterogeneous individual preferences across households. We also show that the conditions allow for meaningful sharing rule identification under the same minimalistic assumptions. Section 5 concludes and sets out some interesting avenues for follow-up research. Appendix A contains the proofs of our main results, and Appendices B and C provide additional information on our empirical application.

2 Stable marriage

In our following analysis, we will assume an empirical analyst who observes a set of matched/married households with (aggregate) consumption bundles that consist of publicly and privately consumed quantities. We assume that households make consumption decisions that are collectively rational, i.e. intrahousehold allocations are Pareto efficient. Next, we also assume that consumption patterns are such that marriages are stable, i.e. no individual wants to exit marriage. Formally, a marriage is stable if it is "individually rational" and has "no blocking pairs". Individual rationality means that no individual prefers becoming single over staying married. Similarly, no blocking pairs means that there are no two individuals who want to exit their current marriage to remarry each other. In what follows, we will formalize these assumptions, to subsequently define a "stable matching allocation" as one that meets Pareto efficiency, individual rationality and no blocking pairs. We will also demonstrate the existence of such a stable allocation under our set of assumptions. Before doing so, we first introduce some necessary notation.

2.1 Notation

We consider households that consist of males m and females w. In particular, we start from a finite set of men M and a finite set of women W. The marriage market is characterized by a matching function $\sigma : M \cup W \rightarrow M \cup W \cup \{\emptyset\}$. This function satisfies, for all $m \in M$ and $w \in W$,

$$\begin{aligned}
&\sigma(m) \in W \cup \{\emptyset\}, \\
&\sigma(w) \in M \cup \{\emptyset\}, \\
&\sigma(m) = w \in W \text{ if and only if } \sigma(w) = m \in M.
\end{aligned}$$

In words, the function σ assigns to every man or woman either a partner of the other gender (i.e. $\sigma(m) = w$ and $\sigma(w) = m$) or nobody (i.e. $\sigma(m) = \emptyset$ and $\sigma(w) = \emptyset$), which means that the man/woman remains single. If $\sigma(m) = w$, we say that man m is matched to woman w and vice versa, i.e. w and m form a married pair. Our analysis in Section 3 will assume data sets that only contain observations on married pairs, i.e. $\sigma(m) \neq \emptyset$ and $\sigma(w) \neq \emptyset$ for any m and w (which implies $|M| = |W|$). However, we emphasize that it is actually possible to extend our framework to incorporate single men and women. But this would substantially complicate the notation without adding substantial insights.

Married couples make consumption decisions. In particular, we assume that households consume a set of commodities, which may include the spouses' leisure (as in our application in Section 4). The set of commodities consists of both private and public goods. We denote by $q \in \mathbb{R}^n_+$ a (column) vector of n private goods and by $Q \in \mathbb{R}^k_+$ a (column) vector of k (intrahousehold) public goods. For any married pair $(m, \sigma(m))$, $(q_{m,\sigma(m)}, Q_{m,\sigma(m)})$ represents the observed aggregate consumption bundle of private and public goods.

Consumption decisions are made under budget constraints, which are defined by prices and incomes for any pair (m, w). We consider a (row) price vector $p_{m,w} \in \mathbb{R}^n_{++}$ for the private goods and a (row) price vector $P_{m,w} \in \mathbb{R}^k_{++}$ for the public goods. For leisure, prices equal the spouses' wages. The vectors $p_{m,\emptyset}$ and $P_{m,\emptyset}$ contain the private good and public good prices for a single man and, analogously, $p_{\emptyset,w}$ and $P_{\emptyset,w}$ contain the prices for a single woman.[4] Next, $y_{m,w} \in \mathbb{R}_{++}$ gives the potential income of the pair (m, w). Similarly, $y_{m,\emptyset}$ and $y_{\emptyset,w}$ are the incomes of a single man m and woman w. If leisure is considered in the analysis, then these incomes will be full incomes. We remark that we assume observed prices and incomes for (unobserved) pairs that are not matched and for (unobserved) singles. However, we only observe the actual consumption quantities for the matched pairs. We will return to these observational issues in Section 3, when we explain the type of data sets we consider, and in Section 4, when we present our empirical application.

[4] Admittedly, for singles the distinction between private and public consumption becomes artificial. Still, we choose to maintain the distinction here to ease our exposition and to avoid an overload of notation.

For a given pair (m, w), the private consumption bundle $q_{m,w}$ is shared between the male and the female. This obtains the male quantities $q^m_{m,w} \in \mathbb{R}^n_+$ and female quantities $q^w_{m,w} \in \mathbb{R}^n_+$ that satisfy the adding up condition $q^m_{m,w} + q^w_{m,w} = q_{m,w}$. For a bundle $(q_{m,w}, Q_{m,w})$, this defines the household allocation $(q^m_{m,w}, q^w_{m,w}, Q_{m,w})$. Then, for given σ the matching allocation $S = \{(q^m_{m,\sigma(m)},\ q^{\sigma(m)}_{m,\sigma(m)},\ Q_{m,\sigma(m)})\}_{m\in M}$ is the collection of household allocations defined over all matched pairs. We exclude externalities for the privately consumed goods. We note, though, that we can easily account for such externalities by formally treating private goods with externalities as public goods. As such, the above approach does not entail any loss of generality.

Finally, every man m is endowed with a non-negative, increasing, continuous and concave utility function $v^m : \mathbb{R}^{n+k}_+ \to \mathbb{R}_+$, which associates a certain level of utility with every bundle (q^m, Q). Analogously, each woman w has a non-negative, increasing, continuous and concave utility function $u^w : \mathbb{R}^{n+k} \to \mathbb{R}_+$. We assume that $v^m(q^m, Q)$ and $u^m(q^w, Q)$ are strictly increasing in, respectively, q^m and q^w. Next, as a regularity condition, we use that $v^m(0, Q) = u^w(0, Q) = 0$ for any level of public goods Q. In words, both women and men must consume at least some private goods to have positive utility (for example, if food is a private good, individuals need to consume at least some food to experience positive utility). Finally, we assume that males and females have complete information about each others' preferences.[5]

2.2 Stable matching allocation

A matching allocation is stable if it is Pareto efficient, individually rational and has no blocking pair. First, Pareto efficiency requires that no Pareto improvement is possible for any matched pair $(m,\ \sigma(m))$. That is, for the given prices $p_{m,\sigma(m)}$ and $P_{m,\sigma(m)}$ and income $y_{m,\sigma(m)}$, there does not exist another intrahousehold allocation over the consumption goods that makes at least one member better off without making the other member worse off. As explained before, Pareto efficiency means that observed consumption behavior is consistent with the collective model of household consumption.

Definition 1 *For a given matching σ, the matching allocation $S = \{(q^m_{m,\sigma(m)},\ q^{\sigma(m)}_{m,\sigma(m)},$ $Q_{m,\sigma(m)})\}_{m\in M}$ is **Pareto efficient** if, for all $m \in M$, there exists no feasible allocation (q^m, q^w, Q), i.e.*

$$p_{m,\sigma(m)}(q^m + q^w) + P_{m,\sigma(m)}Q \leq y_{m,\sigma(m)},$$

[5] Incomplete preference information may result in deviations from the "exact" stability conditions that we formulate in Section 3. In what follows, we introduce stability indices to account for such deviations in empirical applications. In this respect, see also Liu, Malaith, Postlewaith and Samuelson (2014) for a recent discussion on stable matching with incomplete information, and its relation to stable matching with complete information.

such that

$$v^m(q^m, Q) \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}),$$
$$u^{\sigma(m)}(q^w, Q) \geq u^{\sigma(m)}(q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)}),$$

with at least one strict inequality.

Next, individual rationality requires that no individual is better off as a single than under the matching σ. To define this concept formally, we let $V_{m,\emptyset}$ ($U_{\emptyset,w}$) represent the maximum utility level that man m (woman w) could obtain by staying single, when faced with the prices $p_{m,\emptyset}$ and $P_{m,\emptyset}$, and income $y_{m,\emptyset}$ ($p_{\emptyset,w}$ and $P_{\emptyset,w}$, and $y_{\emptyset,w}$), i.e.

$$V_{m,\emptyset} = \max_{q^m, Q} v^m(q^m, Q) \text{ s.t. } p_{m,\emptyset}q^m + P_{m,\emptyset}Q \leq y_{m,\emptyset}, \tag{1}$$
$$U_{\emptyset,w} = \max_{q^w, Q} u^w(q^w, Q) \text{ s.t. } p_{\emptyset,w}q^w + P_{\emptyset,w}Q \leq y_{\emptyset,w}. \tag{2}$$

Then, we have the following definition.

Definition 2 *For a given matching σ, the matching allocation $S = \{(q^m_{m,\sigma(m)}, q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m\in M}$ is **individually rational** if, for all $m \in M$ and $w \in W$, we have*

$$u^w(q^w_{\sigma(w),w}, Q_{\sigma(w),w}) \geq U_{\emptyset,w},$$
$$v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}) \geq V_{m,\emptyset}.$$

Finally, we say that an (unmatched) pair (m, w) is a blocking one if the associated prices $p_{m,w}, P_{m,w}$ and income $y_{m,w}$ admit an allocation such that, when compared to the matching σ, at least one member of the unmatched pair is better off while the other member is not worse off. A stable matching requires that no such blocking pairs exist. To formalize this idea, we consider, for any man m and woman w,

$$\begin{aligned} \psi_{m,w}(u) = &\max_{q^m_{m,w}, q^w_{m,w}, Q_{m,w}} v^m(q^m_{m,w}, Q_{m,w}) \\ \text{s.t. } &p_{m,w}(q^m_{m,w} + q^w_{m,w}) + P_{m,w}Q_{m,w} \leq y_{m,w}, \\ &u^w(q^w_{m,w}, Q_{m,w}) \geq u. \end{aligned} \tag{3}$$

In words, $\psi_{m,w}(u)$ gives the maximal utility that the man m can obtain when he is married to woman w, under the condition that w's utility equals at least u. Let $\overline{V}_{m,w}$ and $\overline{U}_{m,w}$ represent the maximum attainable utility of the male m and female w in the couple (m, w). Then, if we restrict $u \in [0, \overline{U}_{m,w}]$, the function $\psi_{m,w} : [0, \overline{U}_{m,w}] \to [0, \overline{V}_{m,w}]$ traces out the Pareto frontier of the couple (m, w). Using this, we get the following definition of our no blocking pairs requirement.

Definition 3 *For a given matching σ, the matching allocation $S = \{(q^m_{m,\sigma(m)}, q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m \in fM}$ has* ***no blocking pairs*** *if, for all $m \in M$ and $w \in W$ with $w \neq \sigma(m)$, there exist no utility levels $V_{m,w}$ and $U_{m,w}$ such that*

$$\begin{aligned} V_{m,w} &= \psi_{m,w}(U_{m,w}), \\ V_{m,w} &\geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}), \\ U_{m,w} &\geq u^w(q^w_{\sigma(w),w}, Q_{\sigma(w),w}), \end{aligned}$$

with at least one strict inequality.

We can now define our concept of a stable matching allocation.

Definition 4 *For a given matching σ, a matching allocation $S = \{(q^m_{m,\sigma(m)}, q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m \in M}$ is* ***stable*** *if it is Pareto optimal, individually rational and has no blocking pair.*

We can prove the following lemma.

Lemma 1 *The function $\psi_{m,w}(u)$ is strictly decreasing and continuous over the interval $[0, \overline{U}_{m,w}]$.*

Given this lemma, we can apply a general result of Alkan and Gale (1990) to show existence of a stable matching allocation.[6]

Proposition 1 *For any set of women W and men M with utility functions u^w and v^m, and given the incomes $y_{m,w}, y_{m,\emptyset}, y_{\emptyset,w}$ and prices for private and public goods $p_{m,w}$, $p_{m,\emptyset}$, $p_{\emptyset,w}$, $P_{m,w}$, $P_{m,\emptyset}$, $P_{\emptyset,w}$, there exists at least one matching σ that defines a stable matching allocation.*

To conclude this section, we provide an alternative formulation of the no blocking pairs criterion in Definition 3, which will be instrumental for our revealed preference characterization in the next section. Specifically, given that the Pareto frontier is continuous and strictly decreasing (by Lemma 1), it is easy to see that the no blocking pair condition in Definition 3 is equivalent to the requirement that, for any man m and woman w, there must exist at least one combination of $V_{m,w}$ and $U_{m,w}$ such that

$$\begin{gathered} V_{m,w} = \psi_{m,w}(U_{m,w}), \\ V_{m,w} \leq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}) \text{ and } U_{m,w} \leq u^w(q^w_{\sigma(w),w}, Q_{\sigma(w),w}). \end{gathered} \tag{4}$$

[6] Because the result follows directly from Theorem 1 of Alkan and Gale (1990), we need not include a formal proof. We remark that this existence result does not necessarily imply a unique stable marriage matching. See, for example, Eeckhout (2000), Clark (2006) and Legros and Newman (2010) for conditions that guarantee uniqueness in a non-transferable utility setting that is similar to ours. Importantly, however, non-uniqueness does not interfere with the validity (and, thus, applicability) of the testable implications and (set) identification results that we derive below.

3 Testable implications and identification

After defining the type of data sets that we consider, we will introduce our revealed preference conditions for rationalizability by a stable matching. These conditions are nonlinear in unknowns, which makes them difficult to implement. Therefore, in a following step we will define testable implications that are linear in unknowns and, thus, easy to operationalize. Importantly, these linear conditions will have an intuitive interpretation in terms of the stability criteria that we outlined in the previous section. In addition, as we will indicate, they provide a useful basis for (set) identifying the decision structure (including the sharing rule) underlying household consumption behavior if this behavior is found consistent with stable marriage.

3.1 Rationalizability

We assume that the empirical analyst only has consumption observations on married pairs, i.e. there are no singles. For a given set of males M and females W (with $|M| = |W|$), we assume a data set $\mathcal{D}$ that contains the following information:

- the matching function σ,
- the consumption bundles $(q_{m,\sigma(m)}, Q_{m,\sigma(m)})$ of all matched couples $(m, \sigma(m))$ with $m \in M$,
- the prices $p_{m,w}, P_{m,w}$ for all $m \in M \cup \emptyset$ and $w \in W \cup \emptyset$,
- the incomes $y_{m,w}$ for all $m \in M \cup \emptyset$ and $w \in W \cup \emptyset$.

Obviously, the empirical analyst needs to observe who matches whom (i.e. the function σ) to check stability of marriages. Next, we observe the (aggregate) consumption bundles $q_{m,\sigma(m)}$ and $Q_{m,\sigma(m)}$ only for pairs $(m, \sigma(m))$ that are effectively matched. By contrast, we do not observe any consumption if there is no match (i.e. a pair (m, w) with $w \neq \sigma(m)$). In that case, the vectors $q_{w,m}$ and $Q_{w,m}$ represent possible consumption outcomes of (w, m) if the pair had been matched, and $q^w_{w,m}$ and $q^m_{w,m}$ give the corresponding private consumption shares. The underlying idea is that individuals anticipate this consumption when evaluating alternative possible matches. Finally, we do assume that the empirical analyst can reconstruct the budget conditions (i.e. prices $p_{m,w}, P_{m,w}$ and income $y_{m,w}$) for any $m \in M \cup \emptyset$ and $w \in W \cup \emptyset$, which also includes unobserved decision situations pertaining to unmatched pairs and single status. As a specific example, take a standard labor supply setting where couples have to choose a leisure-consumption bundle. Then, the price vectors $p_{w,m}$ and $P_{w,m}$ contain exogenously defined individual wages, and the income $y_{w,m}$ stands for the corresponding full income, which can be reconstructed from observed individual wages and nonlabor income. We will consider such a labor supply setting in our empirical application in Section 4.

Referring to Definition 4, we can now state our condition for a data set $\mathcal{D}$ to be rationalizable.

Definition 5 *For a given matching σ, the data set $\mathcal{D}$ is* ***rationalizable by a stable matching*** *if, for any $m \in M$ and $w \in W$, there exist utility functions v^m and u^w and individual quantities $q^m_{m,\sigma(m)}$, $q^w_{m,\sigma(m)} \in \mathbb{R}^n_+$, with*

$$q^m_{m,\sigma(m)} + q^w_{m,\sigma(m)} = q_{m,\sigma(m)},$$

such that the matching allocation $\{(q^m_{m,\sigma(m)}, q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m \in M}$ is stable.

At this point, it is useful to emphasize the minimalistic nature of our assumptions. Specifically, our rationalizability criterion requires only a single consumption observation per married pair. In addition, we account for heterogeneous preferences for all individuals (females and males) that are observed. A main conclusion of this study will be that we can meaningfully analyze stable marriages even under these minimalistic priors. In particular, in what follows we will introduce an easy to implement (linear) methodology for testing the empirical validity of stability, and for identifying the intrahousehold decision structure if stability cannot be rejected. Our empirical application will show the empirical usefulness of this methodology.

In this respect, we also recall that our concept of a stable matching allocation actually requires both Pareto efficient (or collectively rational) household consumption decisions and stable marriage matching (i.e. individual rationality and no blocking pairs). Notably, Pareto efficiency alone generates no testable implications for observed consumption if we can use only a single observation per household.[7] Therefore, the empirical bite of our methodology stems essentially from the assumption of stable marriage. Because our central focus is precisely on the testable implications of this stability assumption, this also directly motivates us concentrating on data sets with only a single consumption observation per household. However, we want to point out that it is actually fairly easy to extend our framework to settings with multiple household-specific observations (albeit at the cost of notational complexity). We briefly return to this extension in the concluding Section 5.

3.2 Revealed preference characterization

The next Proposition 2 gives a revealed preference characterization of a data set $\mathcal{D}$ that is rationalizable in the sense of Definition 5. As explained in the Introduction, such a revealed preference characterization is intrinsically nonparametric. It does not imply an explicit reference to individual utility functions, and so its verification does not need a specific parametric/functional structure for these utilities. It is directly

[7] See Cherchye, De Rock and Vermeulen (2007) for a detailed analysis of the minimal data requirements (including the number of observations) that are needed for Pareto efficiency (or collective rationality) to generate testable implications.

expressed in terms of the information that is contained by the actual data set $\mathcal{D}$; no additional (possibly confounding) structure is to be imposed.

Usually, revealed preference characterizations are expressed in terms of so-called "Afriat inequalities" (after Afriat, 1967). In our particular case, these Afriat inequalities are defined in unknown (individual, private and public) quantities as well as "personalized prices" and "Afriat numbers". We will explain the interpretation of these prices and Afriat numbers directly after Proposition 2.

Proposition 2 *For a given matching σ, the data set $\mathcal{D}$ is* ***rationalizable by a stable matching*** *if and only if there exist,*

a. *for each matched pair $m \in M$ and $\sigma(m) \in W$, individual quantities $q^m_{m,\sigma(m)}$, $q^{\sigma(m)}_{m,\sigma(m)} \in \mathbb{R}^n_+$ that satisfy*

$$q^m_{m,\sigma(m)} + q^{\sigma(m)}_{m,\sigma(m)} = q_{m,\sigma(m)},$$

which define a matching allocation $\{q^m_{m,\sigma(m)}, q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)}\}_{m\in M}$,

b. *for each unmatched pair $m \in M$ and $w \in W$ (with $\sigma(m) \neq w$), individual quantities $q^m_{m,w}$, $q^w_{m,w} \in \mathbb{R}^n_+$ and public quantities $Q_{m,w} \in \mathbb{R}^k_+$ that satisfy*

$$p_{m,w}(q^m_{m,w} + q^w_{m,w}) + P_{m,w}Q_{m,w} = y_{m,w},$$

c. *for each male $m \in M$, private quantities $q^m_{m,\emptyset} \in \mathbb{R}^n_+$, and public quantities $Q_{m,\emptyset} \in \mathbb{R}^k_+$, that satisfy*

$$p_{m,\emptyset}q^m_{m,\emptyset} + P_{m,\emptyset}Q_{m,\emptyset} = y_{m,\emptyset},$$

d. *for each female $w \in W$, private quantities $q^w_{\emptyset,w} \in \mathbb{R}^n_+$ and public quantities $Q_{\emptyset,w} \in \mathbb{R}^k_+$ that satisfy*

$$p_{\emptyset,w}q^w_{\emptyset,w} + P_{\emptyset,w}Q_{\emptyset,w} = y_{\emptyset,w},$$

e. *for each pair (m, w) ($m \in M$, $w \in M$), personalized prices $P^m_{m,w}$, $P^w_{m,w} \in \mathbb{R}^k_{++}$ that satisfy*

$$P^m_{m,w} + P^w_{m,w} = P_{m,w},$$

as well as strictly positive Afriat numbers $V_{m,w}$, $V_{m,\emptyset}$, $U_{m,w}$, $U_{\emptyset,w}$ and $\delta_{m,w}$, $\delta_{m,\emptyset}$, $\lambda_{m,w}$, $\lambda_{\emptyset,w}$ (for any $m \in M$ and $w \in W$) that simultaneously meet the following constraints:

i. *Afriat inequalities for all males $m \in M$, i.e. (for any $w, w' \in W$)*

$$\begin{aligned}
V_{m,w} - V_{m,w'} &\leq \delta_{m,w'}\left(p_{m,w'}(q^m_{m,w} - q^m_{m,w'}) + P^m_{m,w'}(Q_{m,w} - Q_{m,w'})\right), \\
V_{m,w} - V_{m,\emptyset} &\leq \delta_{m,\emptyset}\left(p_{m,\emptyset}(q^m_{m,w} - q^m_{m,\emptyset}) + P_{m,\emptyset}(Q_{m,w} - Q_{m,\emptyset})\right), \\
V_{m,\emptyset} - V_{m,w'} &\leq \delta_{m,w'}\left(p_{m,w'}(q^m_{m,\emptyset} - q^m_{m,w'}) + P^m_{m,w'}(Q_{m,\emptyset} - Q_{m,w'})\right),
\end{aligned}$$

ii. Afriat inequalities for all females $w \in W$, i.e. (for any $m, m' \in M$)

$$
\begin{aligned}
U_{m,w} - U_{m',w} &\leq \lambda_{m',w} \left(p_{m',w}(q^w_{m,w} - q^w_{m',w}) + P^w_{m',w}(Q_{m,w} - Q_{m',w}) \right), \\
U_{m,w} - U_{\emptyset,w} &\leq \lambda_{\emptyset,w} \left(p_{\emptyset,w}(q^w_{m,w} - q^w_{\emptyset,w}) + P_{\emptyset,w}(Q_{m,w} - Q_{\emptyset,w}) \right), \\
U_{\emptyset,w} - U_{m,w} &\leq \lambda_{m,w} \left(p_{m,w}(q^w_{\emptyset,w} - q^w_{m,w}) + P^w_{m,w}(Q_{\emptyset,w} - Q_{m,w}) \right),
\end{aligned}
$$

iii. individual rationality restrictions for all males $m \in M$ and females $w \in W$, i.e.

$$
\begin{aligned}
V_{m,\sigma(m)} &\geq V_{m,\emptyset}, \\
U_{\sigma(w),w} &\geq U_{\emptyset,w},
\end{aligned}
$$

iv. no blocking pair restrictions for all males $m \in M$ and females $w \in M$, i.e.

$$
\begin{aligned}
V_{m,\sigma(m)} &\geq V_{m,w}, \\
U_{\sigma(w),w} &\geq U_{m,w}.
\end{aligned}
$$

Thus, a necessary and sufficient condition for a data set $\mathcal{D}$ to be rationalizable by a stable matching is that it simultaneously satisfies the conditions (a)-(e) and (i)-(iv). Interestingly, the different conditions can be given a specific interpretation. First, the adding up constraints in (a)-(d) specify feasibility restrictions on the unknown quantities. In particular, condition (a) pertains to individual quantities for matched pairs $(m, \sigma(m))$, condition (b) to individual quantities and public quantities for unmatched pairs (m, w), condition (c) to private and public quantities of males m when single and, finally, condition (d) to private and public quantities of females w when single.

Next, condition (e) defines a formally similar feasibility constraint on the personalized prices $P^m_{m,w}$ and $P^w_{m,w}$ (for any matched or unmatched pair). Intuitively, these personalized prices represent the willingness-to-pay of individual members for the public consumption. Because they must add up to the actual prices $P_{m,w}$, they can actually be interpreted as Lindahl prices that correspond to a Pareto optimal provision of public goods.

Proposition 2 requires the existence of feasible quantities and prices that simultaneously meet the rationalizability conditions (i)-(iv). These rationalizability conditions are defined in terms of Afriat numbers. First, the numbers $V_{m,w}, V_{m,\emptyset}$ represent male m's utilities in alternative decision situations (respectively, in the pair (m, w) and as a single). A directly similar interpretation applies to the numbers $U_{m,w}$ and $U_{\emptyset,w}$, which represent female w's utilities. Next, the numbers $\delta_{m,w}, \delta_{m,\emptyset}$ (for male m) and $\lambda_{m,w}$, $\lambda_{\emptyset,w}$ (for female w) can be interpreted as marginal utilities of individual expenditures (or Lagrange multipliers) in the respective decision scenarios (using, for a given pair (m, w), the personalized prices $P^m_{m,w}$ and $P^w_{m,w}$ to allocate public good expenditures to the individuals m and w).

Then, the Afriat inequalities in conditions (i) and (ii) make sure that there exist (non-negative, continuous, strictly increasing and concave) utility functions v^m and u^w

that explain the data. First, the inequalities ensure that, for all matched couples, these functions satisfy the Pareto efficiency criterion in Definition 1.[8] Next, they also guarantee that the Afriat numbers $V_{m,\emptyset}$, $U_{\emptyset,w}$ solve the individual maximization problems (1) and (2), and that the numbers $V_{m,w}$ and $U_{m,w}$ solve the maximization problem (3) (so that $V_{m,w} = \psi_{m,w}(U_{m,w})$), i.e. $V_{m,w}$ and $U_{m,w}$ represent utilities that are situated on the Pareto frontier of the couple (m, w). Given this, the conditions (iii) and (iv) impose consistency with the individual rationality criterion in Definition 2 and the no blocking pairs criterion in Definition 3 (expressed in the form of (4)).

3.3 Linear conditions

The characterization of rationalizability in Proposition 2 is not directly useful in practice, because the Afriat inequalities in conditions (i) and (ii) are nonlinear in unknowns. In what follows, we will define testable conditions of rationalizability that are linear in unknowns, which makes them easy to apply. While these conditions are necessary for rationalizability, they are, in general, no longer sufficient. That is, we can conclude that a data set $\mathcal{D}$ is not rationalizable if it does not meet the conditions, but there may well exist data sets that pass these (linear) conditions but not the (nonlinear) conditions in Proposition 2. However, as we will explain, our linear conditions do have several attractive features. First, they have an intuitive interpretation in terms of our criteria for stable marriage that we introduced in Section 2. Next, they easily allow for identifying the intrahousehold decision structure (including the sharing rule) if the data satisfy the rationalizability constraints. Finally, and importantly, the necessary conditions do have sufficient empirical bite for an informative empirical analysis, which we will show in Section 4.

Our linear conditions are summarized in the following result.

Proposition 3 *For a given matching σ, the data set $\mathcal{D}$ is* ***rationalizable by a stable matching*** *only if there exist,*

a. *for each matched pair $m \in M$ and $\sigma(m) \in W$, individual quantities $q^m_{m,\sigma(m)}$, $q^{\sigma(m)}_{m,\sigma(m)} \in \mathbb{R}^n_+$ that satisfy*

$$q^m_{m,\sigma(m)} + q^{\sigma(m)}_{m,\sigma(m)} = q_{m,\sigma(m)},$$

which define a matching allocation $\{q^m_{m,\sigma(m)}, q^{\sigma(m)}_{m,\sigma(m)}, Q_{m,\sigma(m)}\}_{m \in M}$,

b. *for each pair (m, w) ($m \in M$, $w \in M$), personalized prices $P^m_{m,w}$, $P^w_{m,w} \in \mathbb{R}^k_{++}$ that satisfy*

$$P^m_{m,w} + P^w_{m,w} = P_{m,w},$$

[8] See in particular Cherchye, De Rock and Vermeulen (2011), who present a revealed preference characterization of Pareto efficient (or collectively rational) household consumption in a setting that is formally similar to ours. The Afriat inequalities in their Proposition 1 are contained in the constraints (i) and (ii) of Proposition 2.

that simultaneously meet the following constraints:

i. individual rationality restrictions for all males $m \in M$ and females $w \in W$, i.e.

$$y_{m,\emptyset} \leq p_{m,\emptyset} q^{m}_{m,\sigma(m)} + P_{m,\emptyset} Q_{m,\sigma(m)},$$
$$y_{\emptyset,w} \leq p_{\emptyset,w} q^{w}_{\sigma(w),w} + P_{\emptyset,w} Q_{\sigma(w),w},$$

ii. no blocking pair restrictions for all $m \in M$ and $w \in M$, i.e.

$$y_{m,w} \leq \left(p_{m,w} q^{m}_{m,\sigma(m)} + P^{m}_{m,w} Q_{m,\sigma(m)}\right) + \left(p_{m,w} q^{w}_{\sigma(w),w} + P^{w}_{m,w} Q_{\sigma(w),w}\right).$$

Technically, we obtain our linear conditions in this result by dropping the Afriat numbers in our earlier characterization. In particular, referring to Proposition 2, we combine the Afriat inequalities (i) and (ii) with the individual rationality and no blocking pairs restrictions (iii) and (iv). This obtains the (necessary) conditions (i) and (ii) in Proposition 3 that are linear in the unknown quantities $q^{m}_{m,\sigma(m)}$, $q^{\sigma(m)}_{m,\sigma(m)}$ and prices $P^{m}_{m,w}$, $P^{w}_{m,w}$.[9]

Again, we can give a specific "revealed preference" interpretation to the different conditions in Proposition 3. The adding up restrictions (a) and (b) also appeared in Proposition 2. Next, the rationalizability restrictions (i) and (ii) bear an intuitive meaning in terms of the stability conditions that we defined in Section 2. First, condition (i) requires, for each individual male and female, that incomes and prices under single status (i.e. $y_{m,\emptyset}, p_{m,\emptyset}$, $P_{m,\emptyset}$ for male m and $y_{\emptyset,w}, p_{\emptyset,w}$, $P_{\emptyset,w}$ for female w) do not allow buying a bundle that is strictly more expensive than the one consumed under the current marriage (i.e. $\left(q^{m}_{m,\sigma(m)}, Q_{m,\sigma(m)}\right)$ for male m and $\left(q^{w}_{\sigma(w),w}, Q_{\sigma(w),w}\right)$ for female w). Indeed, if these conditions are not met, then at least one man or woman is better off (i.e. can attain a strictly better bundle) as a single, which means that the marriage allocation is not stable. In a similar vein, the right hand side of the inequality in condition (ii) gives the sum value of the bundles within marriage for male m (i.e. $p_{m,w} q^{m}_{m,\sigma(m)} + P^{m}_{m,w} Q_{m,\sigma(m)}$) and female w (i.e. $p_{m,w} q^{w}_{\sigma(w),w} + P^{w}_{m,w} Q_{\sigma(w),w}$), evaluated at the prices that pertain to the pair (w, m) (and using personalized prices to evaluate the public quantities). Condition (ii) then requires that the pair's income $y_{m,w}$ must not exceed this sum value. Intuitively, if this condition is not met, then man m and woman w can allocate their income so that both of them are better off (with at least one strictly better off) than with their current matches $\sigma(m)$ and $\sigma(w)$, which makes (w, m) a blocking pair.

Because the conditions (a)-(b) and (i)-(ii) in Proposition 3 are linear in unknown quantities and prices, they define testable implications for rationalizability that can

[9] It is interesting to observe that the linear conditions in Proposition 3 bear some formal similarity to the ones derived by Browning, Chiappori and Weiss (2014, Chapter 7.2) for the model of Shapley and Shubik (1972) and Becker (1973). However, a crucial difference is that Browning, Chiappori and Weiss's conditions assume that individual utilities are transferable, whereas our conditions apply to more general utility structures.

be verified through simple linear programming, which is particularly convenient from a practical point of view. Interestingly, for a data set that satisfies these conditions, Proposition 3 also implies an operational way to identify the intrahousehold decision structure that underlies the rationalizable consumption behavior. It allows for recovering individual quantities and personalized prices that represent the observed behavior in terms of a stable matching. Specifically, it defines feasible sets of these quantities and prices as (non-empty) feasible sets characterized by the linear constraints in Proposition 3, which effectively "set" identifies these unobservables (under the maintained assumption of a stable matching).

Importantly, our linear conditions also allow for recovering the sharing rule that corresponds to rationalizable household consumption. In the collective model, this sharing rule defines the individual incomes that are allocated to the male m and female w. For a matched pair $(m, \sigma(w))$, we can define the male income share $y^m_{m,\sigma(m)}$ and female income share $y^{\sigma(m)}_{m,\sigma(m)}$ as

$$y^m_{m,\sigma(m)} = p_{m,\sigma(m)} q^m_{m,\sigma(m)} + P^m_{m,\sigma(m)} Q_{m,\sigma(m)}, \tag{5}$$
$$y^{\sigma(m)}_{m,\sigma(m)} = p_{m,\sigma(m)} q^{\sigma(m)}_{m,\sigma(m)} + P^{\sigma(m)}_{m,\sigma(m)} Q_{m,\sigma(m)}. \tag{6}$$

We remark that $y^m_{m,\sigma(m)} + y^{\sigma(m)}_{m,\sigma(m)} = y_{m,\sigma(m)}$ by construction, i.e. every share exhaustively assigns a part of total household expenditures to each individual member. Actually, this particular definition of individual income shares (with personalized "Lindahl" prices to evaluate the public quantities) directly corresponds to the two-step representation of collectively rational behavior that we explained in the Introduction. It can be shown that, in the case of public goods, these are the income shares required in the first step to obtain that representation. See, for example, Chiappori and Ekeland (2009) and Cherchye, De Rock, Lewbel and Vermeulen (2015) for a formal argument.

Similar to before, we can set identify the individual income shares through linear programming. In particular, we obtain upper/lower bounds on these shares by maximizing/minimizing the linear functions (5) and (6) subject to the linear rationalizability restrictions in Proposition 3. As we emphasized before, this obtains sharing rule identification even with only a single observation per household and heterogeneous individual preferences across households. This is in stark contrast with the usual identification approach, which assumes either observability of household demand as a function of prices and income (see, for example, Chiappori, 1988, 1992, Chiappori and Ekeland, 2009, and Cherchye, De Rock, Lewbel and Vermeulen, 2015) or observability of a discrete set of household consumption choices (see, for example, Cherchye, De Rock and Vermeulen, 2011).

In our following empirical application, we will refer to $y^m_{m,\sigma(m)}$ and $y^{\sigma(m)}_{m,\sigma(m)}$ as the "total" income shares of male m and female $\sigma(m)$, as they capture all (private and public) consumption that is allocated to the individuals m and $\sigma(m)$. We will distinguish these total shares from "conditional" shares, which are defined as

$$z^m_{m,\sigma(m)} = p_{m,\sigma(m)} q^m_{m,\sigma(m)}, \tag{7}$$

$$z^{\sigma(m)}_{m,\sigma(m)} = p_{m,\sigma(m)} q^{\sigma(m)}_{m,\sigma(m)}. \tag{8}$$

In words, the conditional sharing rule defines the individuals' private consumption ($q^m_{m,\sigma(m)}$ and $q^w_{m,\sigma(m)}$) conditional upon the household's public consumption ($Q_{m,\sigma(m)}$), i.e. the private parts of the total income shares defined in (5) and (6). This conditional sharing rule concept is often used in applications of collective consumption models. See, for example, Browning, Chiappori and Weiss (2014) for more discussion.

4 Empirical application

We consider a nonunitary labor supply setting in which households allocate their full income (i.e. the sum of both spouses' maximum labor income and total non-labor income) to spouses' leisure and remaining consumption (captured by Hicksian aggregate commodities). We subdivide the non-leisure consumption in a private and public part. For our particular data set, private consumption is partly assignable to individual household members (i.e. we observe who consumes what for some goods) and partly nonassignable. We will first check consistency of our data with the rationalizability conditions in Proposition 3. Because our data will fail these sharp conditions (i.e. behavior is not "exactly" stable), we will introduce a procedure that can rationalize the observed behavior in terms of divorce/remarriage costs. By using this procedure, we will be able to address (total and conditional) sharing rule (set) identification. Attractively, it will turn out that the linear programming approach that we outlined in the previous section generates sharing rule bounds that are informatively tight. To illustrate this last feature, we investigate how individuals' income shares vary with spouses' relative wages and households' full income levels.

4.1 Data

We apply our method to a sample of Dutch households drawn from the 2012 wave of the Dutch LISS (Longitudinal Internet Studies for the Social sciences) panel that is gathered by CentERdata. This survey, which is representative for the Dutch population, contains a rich set of economic and socio-demographic variables.[10] We will first discuss the sample selection of our base data set. Next, we will explain our construction of the marriage markets that we consider in our empirical analysis. In particular, for

[10] Households without any Internet access are provided with a basic computer (a 'SimPC') that enables them to connect to the Internet and thereby participate in the survey. See Cherchye, De Rock and Vermeulen (2012) for a collective consumption analysis that is based on the same LISS panel (2009 wave). These authors provide more details on the characteristics of the panel and the data collection procedure.

each different couple in our sample, we will define a specific marriage market on the basis of observable characteristics, and we will use this marriage market to check the rationalizability conditions that we defined above and to address sharing rule identification.

Sample selection. The set of households used for this study was subject to the following sample selection rules. First, we only consider couples with both adults working at least 10 hours per week, and aged between 25 and 65. We include both couples with and without children.[11] Next, we excluded the self-employed to avoid issues regarding the imputation of wages and the separation of consumption from work-related expenditures. After deleting the households with important missing information (mostly, incomplete information on one of the spouses) and some obvious outliers, we obtained a sample of 264 households.

Table 1 provides summary statistics on the sample at hand. Wages are net hourly wages. Leisure is measured in hours per week. To compute leisure hours we assume that an individual needs 8 hours per day for sleeping and personal care (i.e. leisure = 168 - 56 - hours worked). Full income and (Hicksian) consumption are measured in euros per week. For completeness, Table 1 also reports on some important background information of the households under consideration.

Our data set contains assignable consumption.[12] In what follows, we will treat leisure as an assignable private good. Next, the LISS data set also allows us to assign part of the remaining consumption to individual household members.[13] But the main part of the observed household consumption is nonassignable.[14] In our analysis, we assume that 50% of this nonassignable consumption is privately consumed and 50% is publicly consumed within the household. We can motivate this choice by referring to a method proposed by Browning, Chiappori and Lewbel (2013) to compute household-specific economies of scale in a collective consumption setting. For a given household, the method computes a relative measure of scale economies as the ratio of the (sum of) the expenditures that the male and female would need as singles to buy their consumption bundles within marriage (i.e. public and private quantities evaluated at

[11] We implicitly assume that expenditures on children are internalized in the parents' preferences through individual or public consumption. See Bargain and Donni (2012), Cherchye, De Rock and Vermeulen (2012) and Dunbar, Lewbel and Pendakur (2013) for alternative approaches to dealing with children in collective consumption models.

[12] Using our notation of the previous sections, this means that part of the privately consumed quantities $q^m_{m,\sigma(m)}$ and $q^w_{\sigma(w),w}$ is effectively observed. Clearly, such information is easily included in the linear characterization in Proposition 3 through appropriately defined linear constraints, which define feasibility bounds on the variables $q^m_{m,\sigma(m)}$ and $q^w_{\sigma(w),w}$.

[13] The assignable good categories are food at home and outside home, tobacco, clothing, personal care products and services, medical care and health costs not covered by insurance, leisure time expenditures, (further) schooling expenditures, donations and gifts, and other personal expenditures.

[14] The non-assignable consumption includes mortgage, rent, utilities, transport, insurance, daycare, alimony, debt, holiday expenditures, housing expenditures, other public expenditures, and child expenditures (i.e. expenditures on assignable private goods for children).

the observed market prices), divided by the actual (observed) outlay of the household. Clearly, more public consumption implies greater economies of scale. We compute this scale economies measure for consumption without leisure. For the above specification of private and public consumption (with 50% of the nonassignable consumption treated as public consumption), this obtains on average economies of scale of 1.372 for our sample of households, with a standard deviation of 0.056. These figures fall in line with other estimates of scale economies that have been reported in the empirical literature on collective models (see, for example, Browning, Chiappori and Lewbel, 2013).[15]

Finally, our method requires prices and incomes that apply to the exit options from marriage (i.e. becoming single or remarry). For our labor supply application, prices correspond to individual wages. We assume that wages outside marriage are the same as inside marriage (i.e. exiting marriage does not affect labor productivity). Given that we consider the same individuals in and (potentially) outside marriage, this seems not a particularly strong assumption.[16] Next, to reconstruct the potential full income in the unobserved outside options, we must define the individual nonlabor incomes after divorce. For the observed households, we use a consumption-based measure of total nonlabor income, i.e. nonlabor income equals full income minus reported consumption expenditures. Then, in our linear programming method we treat individual nonlabor incomes as unknowns (similar to the individual quantities $q^m_{m,\sigma(m)}$, $q^w_{m,\sigma(m)}$ and personalized prices $P^m_{m,w}$, $P^w_{m,w}$) that are subject to the restriction that they must add up to the observed (consumption-based) total nonlabor income. Basically, given that the actual nonlabor incomes of individual males and females are unobserved, this checks whether there exists at least one feasible specification of these nonlabor incomes that rationalizes the observed behavior by a stable matching.[17]

Marriage markets. It can hardly be assumed that all individuals in our base data set operate on the same marriage market, given that there are big age differences between them. Although celebrity marriages between spouses with a considerable age gap get quite some media attention, they are rather rare. The question remains what is an individual's marriage market and what observable characteristics are correlated

[15] As a robustness check, we redid all our following analyses for two alternative specifications, i.e. 25% and 75% of the nonassignable consumption treated as public consumption. Of course, these alternative choices resulted in different scale economies measures (average scale economies of, respectively, 1.186 and 1.558 for our sample of households). However, our main qualitative conclusions regarding households' sharing rules remained unaffected.

[16] In fact, our framework could easily integrate alternative assumptions regarding changes in individuals' labor productivity (e.g. resulting from post-divorce adjustments of individual labor supply). For simplicity, our empirical analysis will abstract from such more sophisticated wage effects.

[17] As compared to the alternative that fixes the intrahousehold distribution of nonlabor income (e.g. 50% for each individual), this procedure to endogenously define the individual nonlabor incomes effectively gives the "benefit-of-the-doubt" to our assumption of stable matching. In that sense, we treat the (unknown) individual nonlabor incomes the same as the (unknown) individual quantities and personalized prices. However, to exclude unrealistic scenarios, we impose that individual nonlabor incomes after divorce must lie between 40% and 60% of the total nonlabor income under marriage.

	Mean	St. dev.	Min.	Max.
Male wage	13.52	4.25	5.77	36.06
Female wage	11.98	3.31	1.80	26.96
Full income	2855.85	663.64	1256.41	5854.15
Total private consumption	473.72	156.82	79.04	1148.08
Assignable male private consumption	90.03	53.16	8.77	375
Assignable female private consumption	93.67	48.77	4.62	309.23
Public consumption	290.02	121.16	14.42	830.19
Male leisure	71.36	10.33	22	102
Female leisure	83.74	10.77	39	102
Male age	46.67	8.57	26	65
Female age	44.66	8.69	26	62
Number of children	1.30	1.1	0	4
Male dummy for college degree	0.34	0.47	0	1
Female dummy for college degree	0.33	0.47	0	1
Dummy for mixed marriage	0.14	0.35	0	1

Table 1: Summary statistics. Full income and consumption are in euros per week, wages in euros per hour and leisure in hours per week.

with stable marriages. To obtain further insights into these intriguing questions, we consulted a marriage counselor and sexuologist, who has been running a marriage counseling agency for two decades. It turns out that stable marriages are characterized by rather small age and educational differences between the spouses, and that non-mixed marriages have a higher success rate.

To practically implement these insights, we have allocated the households in our base data set to sixteen mutually exclusive marriage markets. In principle, spouses do not need to operate on the same marriage market. It is sufficient that there is some overlap in their respective markets. Although allowing for this may be more realistic, it would be computationally cumbersome. As a pragmatic choice, we constructed marriage markets on the basis of an indicator variable that captures the age band of the husband (between 25 and 35 years old, between 35 and 45 years old, between 45 and 55 years old and between 55 and 65 years old), a dummy variable that captures whether the husband has a college degree, and a dummy variable that captures whether the couple is a mixed marriage. This resulted in sixteen different marriage markets (i.e. 4 (age classes) $\times$ 2 (college degree or not) $\times$ 2 (mixed versus non-mixed marriages)). It appears that there is quite some variation in the size of the different marriage markets for our sample of households. There are on average 16.5 couples in a marriage market, while the smallest and largest marriage markets contain respectively 1 couple and 58 couples. We refer to Appendix B for more details. In what follows, our core results will come from the application of our methodology to these sixteen markets. However, as a robustness check, we will also discuss results for differently defined marriage markets (see Section 4.5).

4.2 Rationalizability

We begin by checking whether and to what extent the observed consumption and marriage behavior satisfies the rationalizability conditions that we outlined above. Here, a first result is that our data set does not satisfy the sharp conditions in Proposition 3. More specifically, only 7 of the 16 different marriage markets turn out to be stable in terms of these requirements. A possible explanation is that the observed matching allocation is stable, but only if we account for a cost associated with exiting marriage, which lowers the available income after divorce (as a single or when newly married). Such a cost of divorce may also result from (e.g. search) frictions on the marriage market, which make it costly to match a new partner. Or, we may want to account for unobserved (material or immaterial) benefits from marriage (e.g. love), which similarly imply a divorce cost (e.g. the monetary value of love).

As these examples demonstrate, a structural modeling of the cost associated with exiting marriage is not straightforward. Therefore, in the current paper we limit ourselves to quantifying the minimal cost of divorce that we must account for to rationalize the observed behavior by a stable matching. Actually, this will also reveal how close the observed behavior (with the original income levels) is to exactly stable behavior. We operationalize this idea by introducing "stability indices", which represent income losses associated with exiting marriage.

Formally, starting from our characterization in Proposition 3, we include a stability index in each restriction of individual rationality ($s^{IR}_{m,\emptyset}$ for the male m and $s^{IR}_{\emptyset,w}$ for the female w) and no blocking pair ($s^{NBP}_{m,w}$ for the pair (m,w)). Specifically, we replace the inequalities in condition (i) by

$$\begin{aligned} \left(s^{IR}_{m,\emptyset} * y_{m,\emptyset}\right) &\leq p_{m,\emptyset} q^{m}_{m,\sigma(m)} + P_{m,\emptyset} Q_{m,\sigma(m)} \text{ and} \\ \left(s^{IR}_{\emptyset,w} * y_{\emptyset,w}\right) &\leq p_{\emptyset,w} q^{w}_{\sigma(w),w} + P_{\emptyset,w} Q_{\sigma(w),w}, \end{aligned} \tag{9}$$

and the inequality in condition (ii) by

$$\left(s^{NBP}_{m,w} * y_{m,w}\right) \leq \left(p_{m,w} q^{m}_{m,\sigma(m)} + P^{m}_{m,w} Q_{m,\sigma(m)}\right) + \left(p_{m,w} q^{w}_{\sigma(w),w} + P^{w}_{m,w} Q_{\sigma(w),w}\right), \tag{10}$$

and we add the restriction $0 \leq s^{IR}_{m,\emptyset}, s^{IR}_{\emptyset,w}, s^{NBP}_{m,w} \leq 1$. Clearly, imposing $s^{IR}_{m,\emptyset}, s^{IR}_{\emptyset,w}, s^{NBP}_{m,w} = 1$ obtains the original (sharp) conditions in Proposition 3. A lower stability index corresponds to a greater income loss associated with a particular exit option (i.e. become single or remarry). As an extreme scenario, $s^{IR}_{m,\emptyset}, s^{IR}_{\emptyset,w}, s^{NBP}_{m,w} = 0$ means that income after divorce is zero, which implies that the individual rationality and no blocking pair restrictions lose their empirical bite.

In our application, we measure the degree of stability of our data set $\mathcal{D}$ by computing

$$\max_{s^{IR}_{m,\emptyset}, s^{IR}_{\emptyset,w}, s^{NBP}_{m,w}} \sum_{m} s^{IR}_{m,\emptyset} + \sum_{w} s^{IR}_{\emptyset,w} + \sum_{m} \sum_{w} s^{NBP}_{m,w}, \tag{11}$$

subject to the feasibility constraints (a) and (b) in Proposition 3 and the linear constraints (9) and (10).[18] By using the thus computed values of $s^{IR}_{m,\emptyset}$, $s^{IR}_{\emptyset,w}$ and $s^{NBP}_{m,w}$, we can construct an adjusted data set that is rationalizable by a stable matching. Then, for this new data set, we can set identify household-specific sharing rules by using the linear programming method that we introduced in Section 3.[19]

By solving (11), we compute a different stability index for every individual rationality constraint ($s^{IR}_{m,\emptyset}$ and $s^{IR}_{\emptyset,w}$; see (9)) and no blocking pair constraint ($s^{NBP}_{m,w}$; see (10)). Correspondingly, for each exit option, we can define a divorce cost that obtains rationalizability of the observed marriage and consumption behavior. Formally, we define these divorce costs as a fraction of the post-divorce income, i.e. $(1 - s^{IR}_{m,\emptyset}) \times 100$ and $(1 - s^{IR}_{\emptyset,w}) \times 100$ for the individual rationality constraints (for the males m and females w) and $(1 - s^{NBP}_{m,w}) \times 100$ for the no blocking pair constraints (for the pairs (m, w)).

Table 2 provides summary statistics on these cost of divorce for our data set. The second and third columns report on the distribution for the individual rationality constraints of males and females, respectively. The fourth and fifth columns pertain the no blocking pair constraints. For a matched pair $(m, \sigma(m))$, Maximum (fourth column) refers to the highest divorce cost defined over all possible remarriages (i.e. $\max_{m',w'}[(1-s^{NBP}_{m,w'}) \times 100, (1-s^{NBP}_{m',\sigma(m)}) \times 100]$), and Mean (fifth column) to the average divorce cost (i.e. the mean of the values $(1 - s^{NBP}_{m,w'}) \times 100$ and $(1 - s^{NBP}_{m',\sigma(m)}) \times 100$ for all w' and m').

The results in Table 2 reveal that we need only small divorce costs to obtain consistency with the individual rationality constraints: divorce costs are zero for all the females in the sample, and only moderately positive for a small fraction of the males (with an average of 0.22 percent). This suggests that only few individuals have an incentive to become single. Intuitively, this can be explained by the presence of public consumption in our model of stable marriage. Public consumption generates economic gains from marriage, which makes it less attractive to become single. Next, we generally need higher divorce costs to rationalize behavior in terms of the no blocking pair constraints. Still, also here the divorce costs are fairly low in most cases. For example, the mean value in the fourth column (maximum divorce cost per couple) equals only 2.14 percent, and the mean value in the fifth column (average divorce cost) amounts to no more than 0.12 percent.

All these results suggest that we need only mildly adjust the post-divorce incomes

[18] In this respect, we recall that we define the individual nonlabor incomes endogenously in our application. Therefore, in our empirical analysis we only use the stability indices $s^{IR}_{m,\emptyset}$, $s^{IR}_{\emptyset,w}$ and $s^{NBP}_{m,w}$ to rescale potential (post-divorce) labor incomes (i.e. wages multiplied by total available time), and not nonlabor incomes, in our analogues of (9) and (10). This ensures that our conditions are linear in unknowns, so that we can use linear programming to compute (11). It also makes that our following concept of divorce cost is to be interpreted in terms of labor income rather than full income.

[19] This procedure effectively accounts for a varying divorce/remarriage cost for different individuals and exit options. An alternative consists of using the same (a priori fixed) divorce cost (e.g. 2.50 percent of income) for all exit options. Such an alternative procedure yields results that are qualitatively similar to the ones that we report below.

	Individual rationality		No blocking pairs	
	Males	Females	Maximum	Mean
Mean	0.22	0.00	2.14	0.12
St. dev.	1.29	0.00	2.35	0.42
Minimum	0.00	0.00	0.00	0.00
First quartile	0.00	0.00	0.02	0.00
Median	0.00	0.00	1.71	0.04
Third quartile	0.00	0.00	3.25	0.08
Maximum	15.96	0.00	12.38	5.98

Table 2: Cost of divorce as a fraction of post-divorce full income (in percentage).

to rationalize the observed consumption behavior in terms of a stable matching allocation. Nevertheless, some couples do require a substantial cost of divorce to obtain rationalizability. For example, the maximum values in the second and fourth columns of Table 2 equal 15.96 percent and 12.38 percent, respectively. It may be argued that these divorce costs are too high to be realistic, and so these marriages are effectively unstable. Therefore, as a robustness check, we will also conduct a sharing rule identification analysis in which we omit couples that require a "too high" divorce cost (characterized as maximum cost above 2.50 percent; see Section 4.5).

4.3 Identification of the total sharing rule

As indicated above, by using the divorce costs that are summarized in Table 2 we can construct a new data set that is rationalizable by a stable matching. Subsequently, this allows us to (set) identify the decision structure underlying the observed stable marriage behavior. We will first do this for the "total" sharing rule (see (5) and (6)). Because this total sharing rule allocates the household's full income to the individual household members by accounting for both private and public consumption (evaluated at the individuals' Lindahl prices), it can be regarded as an indicator of individual well-being. The main distinguishing feature of our framework is that it explicitly includes the marriage market implications for household consumption patterns. As such, it effectively identifies the sharing rule through a structural modeling of the individual's outside options on the marriage market.

As a first exercise, we compare the bounds on female income shares (male shares are one minus the female shares) that are obtained by our revealed preference methodology with "naive" bounds. These naive bounds do not make use of the (theoretical) restrictions associated with a stable matching allocation, and are defined as follows: the lower bound for a female in a particular household equals the share of the value of her assignable consumption (including leisure) in this household's full income; the corresponding upper bound adds the share of nonassignable consumption in the household's full income to this lower bound. In other words, the lower (upper) bound corresponds to an (extreme) scenario where all the household's nonassignable consumption is allo-

	Naive bounds	Stable bounds	
		Only IR constraints	All constraints
Mean	21.10	17.25	13.24
St. dev.	6.90	7.03	5.91
Minimum	1.59	1.59	1.59
First quartile	16.08	12.44	9.23
Median	20.72	17.48	12.08
Third quartile	25.70	21.66	15.58
Maximum	43.64	43.64	43.64

Table 3: Total sharing rule identification. Percentage point differences between upper and lower bounds on female relative income shares.

cated to the male (female).

The results are summarized in Table 3. In that table, we call the bounds that we obtain by our methodology "stable" bounds, as they correspond to a stable matching allocation on the marriage market. We compute two types of stable bounds: the first type is defined by solely using the empirical implications of the individual rationality (IR) constraints, while the second type uses all (i.e. both the individual rationality and no blocking pair) constraints. Comparing the results of these two analyses will provide insight into the identifying power of each category of stability restrictions. In fact, we can give our two types of stable bounds a particular interpretation in terms of two specific models of stable marriage behavior. The first type corresponds to a model with extreme search frictions, i.e. individuals do not meet other potential partners and, thus, the only exit option from marriage is to become single. By contrast, the second type of bounds corresponds to the other limiting case in which individuals have complete information on all potential new partners on the marriage market (that are contained in our sample).

From Table 3, we find that our stable bounds provide a substantial gain in precision compared to the naive bounds. The average difference between the upper and lower naive bounds is 21.10 percentage points, while this difference equals only 17.25 and 13.24 percentage points for our stable bounds. Further, when comparing our two types of stable bounds, we observe that both the individual rationality restrictions and no blocking pair restrictions have considerable identifying power: our stable bounds of the first type (which only use the individual rationality constraints) are much tighter than the naive bounds and, similarly, our stable bounds of the second type (which use all stability constraints) are systematically narrower than our bounds of the first type. Intuitively, this indicates that adding information on outside options (to become single or to remarry) will generally enhance the precision of our identification method. Interestingly, the results also demonstrate that the stable marriage model with only the assumption of individual rationality (i.e. no information on potential new partners) can obtain informative sharing rule bounds.

Importantly, despite our minimalistic set-up, the bounds that we obtain are informatively tight. We illustrate this feature for the relation between female resource shares and, respectively, households' full incomes and intrahousehold wage ratios (i.e. female wage divided by male wage). We focus on these particular relationships because they received considerable attention in the literature on collective consumption models. It is frequently assumed in the empirical literature that bargaining power is independent of total household income; see, for instance, Lewbel and Pendakur (2008), Bargain and Donni (2012) and Dunbar, Lewbel and Pendakur (2013). These authors use this assumption to obtain point identification for resource shares. Next, the literature also provided systematic evidence that a household member's bargaining power generally increases with her/his wage; see, for example, Chiappori, Fortin and Lacroix (2002), Blundell, Chiappori, Magnac and Meghir (2007), Oreffice (2011) and Cherchye, De Rock, Lewbel and Vermeulen (2015). The underlying reasoning is that a higher wage improves the member's outside options, which in turn yields a better bargaining position within marriage.

Figure 1 presents our results (for the stable bounds that use all stability constraints in Table 3). Each $\bullet$ and $+$ sign in the figure represents the upper and lower bound for a given household in our sample. To help visualize the results, we included trendlines showing local sample averages of these household-specific upper and lower bounds. The trendlines on the figure in the left panel suggest a slightly decreasing trend, but they do not allow us to reject the hypothesis that the female's relative income share is independent of full income. By contrast, in the right panel of Figure 1 we do observe a significant upward sloping pattern. The stable bounds clearly suggest that a higher relative wage for the female does give her a better bargaining position and, via this channel, a larger resource share.

Summarizing, the bounds in Figure 1 show the potential of our framework to analyze the structural implications of the marriage market for household consumption patterns. It allows for an informative analysis of the intrahousehold decision process, even if we make minimal assumptions regarding the data at hand. Interestingly, the conclusions that we obtain regarding income and wage effects fall in line with existing results in the literature. Importantly, because of the specific set-up of our analysis, we can give these findings a structural interpretation by explaining a higher relative share in terms of better outside options on the marriage market. In Appendix C we consider the same relations as in Figure 1 for subgroups of our sample, which we split up on the basis of age, education, ethnicity and having children or not. This shows robustness of our main qualitative findings at these subgroup levels.

4.4 Identification of the conditional sharing rule

We next conduct a similar analysis as above but now for the "conditional" sharing rule (see (7) and (8)). As explained in Section 3, this sharing rule defines the intrahousehold distribution of private consumption (for the given level of public consumption). It provides information on individuals' well-being that is complementary to the information revealed by the total sharing rule. An individual's total resource share may vary

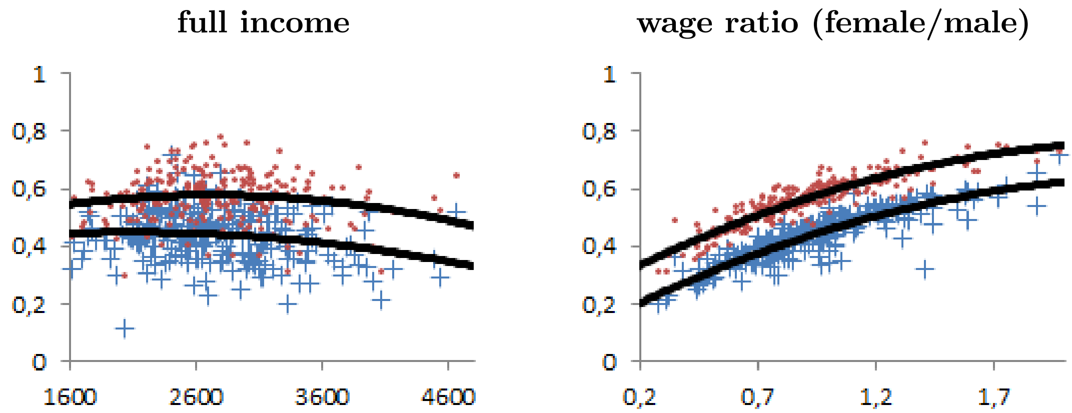

Figure 1: Total sharing rule (Y-axis: female relative share), full income (X-axis in left panel) and wage ratio (X-axis in right panel).

because of changes in his/her private consumption as well as changes in the household's public consumption. In what follows, we disentangle these two channels. In particular, we separately consider the impact of a household's full income and the intrahousehold wage ratio on, respectively, the budget share of public consumption and the female's (conditional) share of private consumption.

Figure 2 sets out the budget share of public consumption as a function of full incomes and wage ratios in our data set. We find that the public share varies considerably across the households in our sample. However, it is fair to say that, on average, we cannot discern an obvious income or wage effect. As a following exercise, Figure 3 presents the female's relative share of private consumption (defined in (8)) and the budget share of public consumption. Again, we do not observe an obvious relation between these two shares. Summarizing, from Figures 2 and 3 we may conclude that the share of public consumption is, on average, independent of the household's full income, the intrahousehold wage ratio and the individuals' private consumption shares. As an implication, it must be that the patterns in Figure 1 (for the total sharing rule) are mainly driven by variation in private consumption shares (i.e. the conditional sharing rule). We will explore this in more detail below.

Before doing so, we first discuss our results in Table 4, which is similar to Table 3, except that it pertains to the conditional sharing rule instead of the total sharing rule. Like before, we find that our stable bounds are substantially tighter than the naive bounds, and that both the individual rationality and no blocking pair constraints have considerable identifying power. Further, comparing the bounds in Tables 3 and 4 reveals that our bounds for the conditional sharing rule are much narrower than for the total sharing rule. In the last column of Table 4, the average difference between upper and lower bounds on the female's conditional share is as small as 3.09 percentage points. Moreover, for three quarters of the households in our sample this difference is not above 4.21 percentage points. The smallest difference is zero to two decimal places,

Figure 2: Budget share of public consumption (Y-axis), full income (X-axis in left panel) and wage ratio (X-axis in right panel).

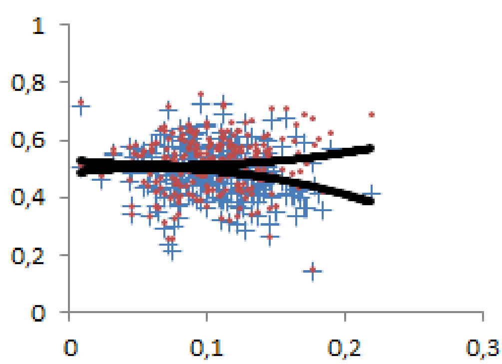

Figure 3: Conditional sharing rule (Y-axis: relative female share) and budget share of public consumption (X-axis).

	Naive bounds	Stable bounds	
		Only IR constraints	All constraints
Mean	11.96	7.63	3.09
St. dev.	4.35	5.02	3.98
Minimum	0.80	0.00	0.00
First quartile	8.74	3.30	0.51
Median	11.56	7.64	1.76
Third quartile	14.74	10.86	4.21
Maximum	27.94	27.94	27.94

Table 4: Conditional sharing rule identification. Percentage point differences between upper and lower bounds on female relative income shares.

showing that for some households we come extremely close to point identification.

From Table 4, we learn that our bounds for the conditional sharing are very informative. This is further confirmed by Figure 4, which depicts wage and income effects for the female's private consumption share. We find that the average trendlines are very close to each other, reflecting the results in Table 4. The left panel of Figure 4 suggests that, on average, full income leads to a lower share of private consumption for females, but this effect is rather weak. By contrast, we find a significantly positive effect of the wage ratio on the private's income share. In combination with the result in the right panel of Figure 2, this leads us to conclude that the wage effect for the total sharing rule (in Figure 1) mainly runs through the individuals' conditional share (and not the share of public consumption). Appendix C shows that this conclusion remains valid if we condition on age, education, ethnicity and having children or not.

As a final remark, one may be tempted to argue that our results are an artefact of our set-up, which assumes that leisure is privately assignable and priced at the individual's own wage level. Indeed, if leisure demands were not responsive to their prices (i.e. individual wages), then by construction this would obtain higher relative income shares for higher relative wages. However, this alternative explanation is contradicted by the results that we present in Figure 5, which depicts the female's private share without leisure as a function of full income and the wage ratio. Admittedly, we now observe more heterogeneity across households. However, the average income and wage effects are similar to before.

4.5 Robustness checks

To end our empirical application, we check robustness of our above results with respect to three specific features of our analysis: we considered sixteen fairly restricted marriage markets, we did not use information on singles in the analysis, and we included couples with high divorce costs in our identification exercises. We will only report on the tightness of the (total and conditional) sharing rule bounds. Our main qualitative conclusions regarding income and wages effects remained unaffected for each of our

Figure 4: Conditional sharing rule (Y-axis: relative female share), full income (X-axis in left panel) and wage ratio (X-axis in right panel).

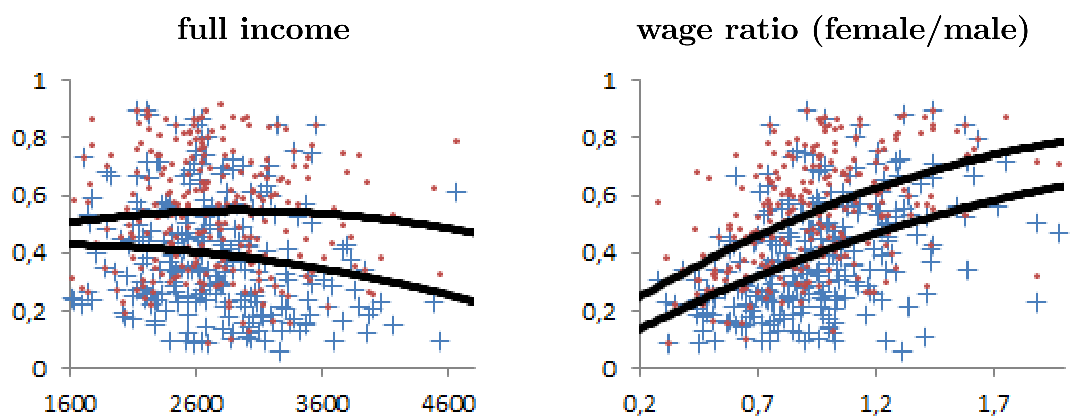

Figure 5: Private share without leisure (Y-axis: relative female share), full income (X-axis in left panel) and wage ratio (X-axis in right panel).

three robustness exercises. We omit these results for compactness.

Our first robustness exercise considers a different construction of marriage markets. In particular, while in our main analysis we used sixteen relatively small marriage markets, we now want to investigate the impact of enlarging these marriage markets. More precisely, we consider a different marriage market for each different age group of the husband (between 25 and 35 years old, between 35 and 45 years old, between 45 and 55 years old and between 55 and 65 years old). This resulted in four markets, with 28, 79, 100 and 57 couples, respectively. Table 5 summarizes our findings. Comparing this table to Tables 3 and 4, we learn that enlarging the marriage market generally enhances the identifying power of our method (i.e. narrower sharing rule bounds). Actually, this should not be very surprising as larger marriage markets imply more outside options for the married individuals, which can only improve the identification analysis. However, it is also fair to say that this improvement is rather limited. For example, the average difference between upper and lower bounds narrows down from 13.24 percent (in Table 3) to 12.68 percent (in Table 5) for the total sharing rule, and from 3.09 percent (in Table 4) to 2.45 percent (in Table 5) for the conditional sharing rule.

As a second robustness check, we include singles in our identification analysis. In particular, we account for the possibility that a married individual may consider remarrying a single of the other gender, i.e. singles are used to construct potentially blocking pairs.[20] For our data set, we have consumption information on 198 female singles and 170 male singles, which we added to our original analysis with sixteen marriage markets. Table 6 summarizes the results of this exercise. We get narrower sharing rule bounds than in our previous exercise (see Tables 3 and 4). Like before, this could be expected a priori, as we added outside options for the married individuals. However, we again conclude that the bounds tightening is rather modest: the average difference in Table 6 is 12.68 percent for the total sharing rule and 2.11 percent for the conditional sharing rule.

Our final robustness analysis excludes couples with a "too high" divorce cost when identifying intrahousehold sharing rules. This acounts for the argument that high divorce costs effectively signal unstable marriages, which thus cannot be used to learn about intrahousehold decision processes (by using stability of marriage as an identifying assumption). In our robustness check, we exclude couples that require a divorce cost of at least 2.50 percent for at least one exit option (as a single or with a new partner). This criterion led us to drop 116 couples (i.e. about 44 percent of our sample), and we redid our original analysis for the remaining 148 "stable" households. The results are given in Table 7. We obtain that sharing rule bounds are wider than before, because we now use less information in our identification analysis. Once more, however, the impact is moderate: average differences are 15.15 percentage points for the total sharing rule

[20] We remark that we do not include rationalizability conditions for singles' behavior. Unless there is a shortage on one side of the marriage market, rationalizing this behavior requires an explicit model for frictions on the marriage market or marriage costs. To focus our discussion, we abstract from such extensions in the current study.

	Total sharing rule	Conditional sharing rule
Mean	12.68	2.45
St. dev.	5.37	3.38
Minimum	1.59	0.00
First quartile	9.44	0.41
Median	11.71	1.35
Third quartile	14.58	2.96
Maximum	42.38	26.29

Table 5: Sharing rule identification with larger marriage markets. Percentage point differences between upper and lower bounds on female relative income shares.

	Total sharing rule	Conditional sharing rule
Mean	12.39	2.11
St. dev.	5.17	3.16
Minimum	1.59	0.00
First quartile	9.23	0.00
Median	11.67	1.07
Third quartile	14.46	2.89
Maximum	43.67	27.94

Table 6: Sharing rule identification with singles included. Percentage point differences between upper and lower bounds on female relative income shares.

and 4.24 points for the conditional sharing rule.

As an overall conclusion, our above results clearly demonstrate the empirical usefulness of endogenizing the marriage matching decisions in the household consumption analysis.[21] Moreover, they also neatly show the potential of our framework to analyze the structural implications of the marriage market for household consumption patterns. It generates sharing rule bounds that are considerably tighter than the naive ones. Notably, this conclusion is based on a fairly small sample selection, with a single consumption observation per household, and without homogeneity of individual preferences.

Our robustness exercises illustrate that enlarging marriage markets or including singles' information in the analysis can further tighten the bounds (albeit that the impact was fairly limited for our sample). Different approaches can be used to obtain additional improvements, by expanding the minimalistic set-up of the application we consider here. Obviously, tighter bounds can be obtained by including more households. Additional households imply that a larger range of outside options is incorporated in the sharing rule identification analysis. Or, one can use panel data that contain a time-series of consumption observations for individual households. As we explain in

[21] In this respect, see also Ackerberg and Botticini (2002), who emphasize the importance of endogenizing matching decisions in the empirical analysis of contract forms.

	Total sharing rule	Conditional sharing rule
Mean	15.15	4.24
St. dev.	6.48	4.59
Minimum	1.59	0.00
First quartile	10.63	0.90
Median	14.29	2.86
Third quartile	18.44	5.80
Maximum	43.67	27.94

Table 7: Sharing rule identification without unstable couples. Percentage point differences between upper and lower bounds on female relative income shares.

the concluding section, this can strengthen the analysis by combining the empirical restrictions of the Pareto efficiency assumption with the stable marriage implications that we have developed. Finally, and naturally, narrower sharing rule bounds are also obtained by making stronger assumptions, such as preference homogeneity across individuals.[22]

5 Concluding discussion

We have defined testable (revealed preference) restrictions of stable marriage under the maintained assumption of Pareto efficient household consumption. Importantly, our characterization allows for intrahousehold consumption transfers but does not require individual utilities to be transferable. We have shown that this characterization provides a useful basis for identifying the intrahousehold decision structure (including the sharing rule) that underlies stable marriage behavior. Interestingly, the application of our testability and identification results merely requires standard linear programming, which is particularly attractive from a practical point of view. We also conducted an empirical application to Dutch household data, which shows that this linear programming methodology has substantial empirical bite (i.e. yields informative results) even in the limiting case with only a single consumption observation per household and without assuming any preference homogeneity across households.

Basically, we have developed a novel framework to analyze the structural impli-

[22] To correct for heterogeneous observable characteristics of households and individuals, one can use the observed consumption behavior to (parametrically or nonparametrically) estimate household demand while conditioning on these characteristics, and subsequently apply the revealed preference restrictions in Proposition 3 to the estimated demands. See Cherchye, De Rock, Lewbel and Vermeulen (2015) for such an exercise in the context of collective consumption models (without explicit marriage market restrictions). They show that this combination of estimated demand functions with revealed preference restrictions obtains a particularly powerful sharing rule (set) identification analysis. Blundell, Browning and Crawford (2008), Stoye and Kitamura (2013), Blundell, Kristensen and Matzkin (2014), Hoderlein and Stoye (2014) address similar questions in a unitary context, also dealing with unobserved heterogeneity driving demand behavior.

cations of the marriage market for household consumption behavior. It endogenizes the marriage matching decisions in the household consumption analysis. Because it explicitly incorporates individuals' outside options on the marriage market, the framework allows us to further open the "black box" of intrahousehold decision making. We strongly believe that this paves the way for many interesting new developments.

For example, in our empirical application, we have used stability indices to account for deviations of observed behavior from exactly stable behavior. These indices capture the cost of divorce, which is caused by frictions on the marriage market and/or unobserved benefits from marriage (such as love). From this perspective, a first interesting extension of our framework consists of explicitly modeling (e.g. search) frictions related to marriage and remarriage. Similarly, one can specifically include unobserved characteristics that drive marriage decisions (e.g. the unobserved consumption of love). Such unobserved characteristics can also capture preference shifts (e.g. single versus married). Generally, a structural modeling of these different aspects can help to disentangle the different aspects that we aggregated in our stability indices.[23]

Next, our empirical analysis has focused on the effects of households' full income levels and relative wages on total and conditional sharing rules. On average, individuals' resource shares appear to be independent of the household's full income. By contrast, a higher relative wage of the female gives her a higher income share under stable marriage. The underlying mechanism is that a higher wage defines better outside options on the marriage market, which we explicitly model in our framework. Following applications can focus on other determinants of individuals' outside options (and, through this channel, income shares). In particular, they may consider alternative characteristics of the individuals (e.g. differences in age, education, ...) or the marriage market itself (e.g. sex ratio, divorce laws, ...). In the literature on collective consumption models, these defining characteristics are usually referred to as "distribution factors" (see, for example, McElroy, 1990, Browning, Bourguignon, Chiappori and Lechene, 1994, and Bourguignon, Browning and Chiappori, 2009). By integrating individuals' outside options in the household consumption analysis, our methodology allows for a structural investigation of the effect of these distribution factors, which should provide a deeper insight into the specific (matching) mechanics that are at play.

Other useful extensions pertain to the basic set-up that we adopted in the current study. For example, because our central focus was on the testable and identifying implications of stable marriage, we have concentrated on data sets with only a single

[23] At this point, if we do not impose specific structure on them, frictions or unobserved characteristics will lead to vacuous rationalizability conditions (i.e. stable marriage loses its testable implications and identification power). This negative result is close in spirit to the one of Varian (1988) in a formally similar revealed preference context. For frictions, we obtain the negative result if we assume the extreme case in which the only person one meets is her/his partner and, in addition, no individual has the option to become single. For unobservable characteristics, we can rationalize any matching by assuming that the match-specific quality (e.g. love) is high enough to outweigh any outside option. As for this last case, identifying structure may be, for example, to assume that all potential partners rank a person-specific attribute (e.g. "amiability") in the same way. We thank Martin Browning for pointing this out to us.

consumption observation per household. In practice, however, time-series of observations for one and the same household are increasingly available. As indicated in the Introduction, the assumption of Pareto efficiency generates specific testable implications as soon as one can use multiple household-specific consumption observations. Extending our framework to a panel data setting (containing time-series for a sample of households) can combine these implications with the stable marriage restrictions that we developed above. Clearly, such a combination can only enrich the empirical investigation. Interestingly, it also enables a structural analysis of dynamic aspects related to intrahousehold consumption and marriage decisions.[24]

Another interesting development consists of explicitly including household production in the consumption model (see, for example, Jacquemet and Robin, 2013, who consider a similar marriage matching context). Our above analysis incorporates expenditures on public goods to model gains from marriage. By modeling the household production technology, we could identify how these household inputs lead to household outputs (that enter the individual utilities). By extending our methodology to also identify the within-household production structure, we obtain a revealed preference toolkit that can empirically address research questions related to, for example, marriage matching on productivity and specialization in marriage. By the very nature of our framework, it could do so while minimizing the assumptions needed for this empirical analysis.

Finally, by adopting the widely used collective consumption model, we have maintained the assumption that households make Pareto efficient consumption decisions, which essentially means that household members act cooperatively. However, it is sometimes argued that the assumption of Pareto efficiency is an overly strong one in a household context.[25] As an alternative, the noncooperative model assumes Nash equilibrium allocations within the household (see, for example, Browning, Chiappori and Lechene, 2010, Lechene and Preston, 2011, and Cherchye, Demuynck and De Rock, 2011). In terms of the resulting within-household allocations, the main difference between the two models is that the noncooperative alternative allows for free riding behavior regarding the consumption of public goods. In our opinion, it would be interesting to extend our framework towards investigating the implications of the marriage market in the case of noncooperative household consumption.

[24] See, for example, the recent study of Mazzocco, Ruiz and Yamaguchi (2013) on the relationship between household consumption decisions (on labor supply and savings behavior) and marital choices. Mazzocco (2007) and, more recently, Lise and Yamada (2014) consider dynamic versions of the collective consumption model. Adams, Cherchye, De Rock and Verriest (2014) analyze such dynamic collective consumption behavior by following a revealed preference approach that is formally similar to ours. Choo (2014) presents an empirical framework for dynamic marriage matching with transferable utility.

[25] See, for example, Lundberg and Pollak (2003) for a discussion on the implicit assumptions underlying the Pareto efficiency assumption in the specific context of married couples. Del Boca and Flinn (2014) recently provided an empirical analysis of efficient (or cooperative) versus inefficient (or noncooperative) household behavior on the basis of observed sorting patterns on the marriage market.

References

[1] Ackerberg, D., and M. Botticini (2002), "Endogenous matching and the empirical determinants of contract forms", *Journal of Political Economy*, 110, 564-591.

[2] Adams, A., L. Cherchye, B. De Rock and E. Verriest (2014), "Consume now or later? Time inconsistency, collective choice and revealed preference", *American Economic Review*, 104, 4147-4183.

[3] Afriat, S. (1967), "The construction of utility functions from expenditure data", *International Economic Review*, 8, 67-77.

[4] Alkan, A. and D. Gale (1990), "The core of the matching game", *Games and Economic Behavior*, 2, 203-212.

[5] Attanasio, O. and V. Lechene (2014), "Efficient responses to targeted cash transfers", *Journal of Political Economy*, 122, 178-222.

[6] Bargain, O., and O. Donni (2012), "The measurement of child costs: A Rothbarth-type method consistent with scale economies and parents' bargaining", *European Economic Review*, 56, 792-813.

[7] Becker, G. (1973), "A Theory of Marriage: Part I", *Journal of Political Economy,* 81, 813–846.

[8] Blundell, R., M. Browning, and I. Crawford (2008), "Best nonparametric bounds on demand responses", *Econometrica*, 76, 1227-1262.

[9] Blundell, R., P.-A. Chiappori, T. Magnac, and C. Meghir (2007), "Collective labor supply: heterogeneity and nonparticipation", *Review of Economic Studies*, 74, 417-445.

[10] Blundell, R., P.-A. Chiappori, and C. Meghir (2005), "Collective labor supply with children", *Journal of Political Economy*, 113, 1277-1306.

[11] Blundell, R., D. Kristensen, and R. Matzkin (2014), "Bounding quantile demand functions using revealed preference inequalities", *Journal of Econometrics*, 179, 112-127.

[12] Bourguignon, F., M. Browning, and P.-A. Chiappori (2009), "Efficient intra-household allocations and distribution factors: implications and identification", *Review of Economic Studies*, 76, 503–528.

[13] Browning, M, F. Bourguignon, P.-A. Chiappori, and V. Lechene (1994), "Income and outcomes: a structural model of intrahousehold allocations", *Journal of Political Economy*, 102, 1067-1096.

[14] Browning, M., and P.-A. Chiappori (1998), "Efficient intra-household allocations: a general characterization and empirical tests", *Econometrica,* 66, 1241-1278.

[15] Browning, M., P.-A. Chiappori, and V. Lechene (2010), "Distributional effects in household models: separate spheres and income pooling", *Economic Journal*, 120, 786-799.

[16] Browning, M., P.-A. Chiappori, and A. Lewbel (2013), "Estimating consumption economies of scale, adult equivalence scales, and household bargaining power", *Review of Economic Studies* 80, 1267-1303.

[17] Browning, M., P.-A. Chiappori, and Y. Weiss (2014), *Family Economics*, Cambridge University Press.

[18] Cherchye, L., T. Demuynck, and B. De Rock (2011), "Revealed preference analysis of noncooperative household consumption", *Economic Journal*, 121, 1073–1096.

[19] Cherchye, L., T. Demuynck, and B. De Rock (2014), "Is utility transferable? A revealed preference analysis", *Theoretical Economics*, 10, 51-65.

[20] Cherchye, L., B. De Rock, A. Lewbel, and F. Vermeulen, F. (2015), "Sharing rule identification for general collective consumption models", *Econometrica*, forthcoming.

[21] Cherchye, L., B. De Rock, and F. Vermeulen (2007), "The collective model of household consumption: a nonparametric characterization", *Econometrica*, 75, 553-574.

[22] Cherchye, L., B. De Rock, and F. Vermeulen (2009), "Opening the black box of intra-household decision-making: theory and non-parametric empirical tests of general collective consumption models", *Journal of Political Economy*, 117, 1074-1104.

[23] Cherchye, L., B. De Rock, and F. Vermeulen (2011), "The revealed preference approach to collective consumption behavior: testing and sharing rule recovery", *Review of Economic Studies*, 78, 176-198.

[24] Cherchye, L., B. De Rock and F. Vermeulen (2012), "Married with children: a collective labor supply model with detailed time use and intrahousehold expenditure information", *American Economic Review*, 102, 3377-3405.

[25] Cherchye, L., and F. Vermeulen (2008), "Nonparametric analysis of household labor supply: goodness-of-fit and power of the unitary and the collective model", *Review of Economics and Statistics*, 90, 267-274.

[26] Chiappori, P.A. (1988), "Rational household labor supply", *Econometrica*, 56, 63-89.

[27] Chiappori, P.A. (1992), "Collective labor supply and welfare", *Journal of Political Economy*, 100, 437-467.

[28] Chiappori, P.A. (2010), "Testable implications of transferable utility", *Journal of Economic Theory* , 145, 1302-1317.

[29] Chiappori, P.A., and I. Ekeland (2006), "The micro economics of group behavior: general characterization", *Journal of Economic Theory*, 130, 1-26.

[30] Chiappori, P.A., and I. Ekeland (2009), "The micro economics of efficient group behavior: identification", *Econometrica*, 77, 763-799.

[31] Chiappori, P.-A., M. Iyigun, J. Lafortune, and Y. Weiss (2013), "Changing the rules midway: the impact of granting alimony rights on existing and newly-formed partnerships", Working paper, Columbia University.

[32] Chiappori, P.-A., B. Fortin, and G. Lacroix (2002), "Marriage market, divorce legislation and household labor supply", *Journal of Political Economy*, 110, 37-72.

[33] Chiappori, P.-A., S. Oreffice, and C. Quintana-Domeque (2012), "Fatter attraction: Anthropometric and socioeconomic matching on the marriage market", *Journal of Political Economy*, 120, 659–695.

[34] Chiappori, P.-A., and P. Reny (2006), "Matching to share risk", Working Paper, University of Chicago.

[35] Chiappori, P.-A. and B. Salanié (2015), "The Econometrics of Matching Models", *Journal of Economic Literature*, forthcoming.

[36] Choo, E. (2014), "Dynamic marriage matching: and empirical framework", *Econometrica*, forthcoming.

[37] Choo, E. and S. Seitz (2013), "The Collective Matching Model: Identification, Estimation and Testing", in *Advances in Econometrics, Volume 32: Structural Econometric Models*, ed. E. Choo and M. Shum, Emerald Publishing Group, forthcoming.

[38] Choo, E., and A. Siow (2006), "Who marries whom and why", *Journal of Political Economy*, 114, 175-201.

[39] Clark, S. (2006), "Uniqueness of stable matchings", *Contributions to Theoretical Economics*, 6, Article 8, Berkeley Electronic Press.

[40] Couprie, H., E. Peluso, and A. Trannoy (2010), "Is power more evenly balanced in poor households?", *Journal of Public Economics*, 94, 493-507.

[41] Del Boca, D., and C. Flinn (2014), "Household behavior and the marriage market", *Journal of Economic Theory*, 150, 515-550.

[42] Diewert, E. (1973), "Afriat and revealed preference theory", *Review of Economic Studies*, 40, 419-426.

[43] Dunbar, G., A. Lewbel, and K. Pendakur (2013), "Children's resources in collective households: identification, estimation and an application to child poverty in Malawi", *American Economic Review*, 103, 438–471.

[44] Dupuy, A., and A. Galichon (2014), "Personality traits and the marriage market", *Journal of Political Economy*, 122, 1271-1319.

[45] Echenique, F. (2008), "What matchings can be stable? The testable implications of matching theory", *Mathematics of Operations Research*, 33, 757-768.

[46] Echenique, F., S. Lee, M. Shum, M., and B. Yenmez (2013), "The revealed preference theory of stable and extremal stable matchings", *Econometrica*, 81, 153-172.

[47] Eeckhout, J. (2000), "On the uniquess of stable marriage matchings", *Economics Letters*, 69, 1-8.

[48] Gale, L.S., and D. Shapley (1962), "College admissions and the stability of marriage", *American Mathematical Monthly*, 69, 9–15.

[49] Galichon, A., and B. Salanié (2014), "Cupid's invisible hand: social surplus and identification in matching models", Working Paper, Columbia University.

[50] Galichon, A., S.D. Kominers and S. Weber (2015), "An empirical framework for matching with imperfectly transferable utility", Working Paper, Sciences Po.

[51] Hoderlein, S., and J. Stoye (2014), "Revealed preferences in a heterogeneous population", *Review of Economics and Statistics*, 96, 197-213.

[52] Houthakker, H. (1950), "Revealed preference and the utility function", *Economica*, 17, 159-174.

[53] Jacquemet, N., and J.-M. Robin (2013), "Assortative matching and search with labor supply and home production", CeMMAP working paper CWP07/13.

[54] Lechene, V., and I. Preston (2011), "Noncooperative household demand", *Journal of Economic Theory*, 146, 504-527.

[55] Legros, P, and A.F. Newman (2007), "Beauty is a beast, frog is a prince: assortative matching with nontransferabilities", *Econometrica* 75 1073-1102.

[56] Legros, P., and A.F. Newman (2010), "Co-ranking mates: assortative matching in marriage markets", *Economics Letters*, 106, 177-179.

[57] Lewbel, A., and K. Pendakur (2008), "Estimation of collective household models with Engel curves", *Journal of Econometrics*, 147, 350-358.

[58] Lise, J., and S. Seitz (2011), "Consumption inequality and intra-household allocations", *Review of Economic Studies*, 78, 328-355.

[59] Lise, J., and K. Yamada (2014), "Household sharing and commitment: evidence from panel data on individual expenditures and time use", Working paper, University College London.

[60] Liu, Q., G.J. Mailath, A. Postlewaite and L. Samuelson (2014), "Stable matching with incomplete information", *Econometrica* 82, 541-587.

[61] Lundberg, S., and R. Pollak (2003), "Efficiency in marriage", *Review of Economics of the Household*, 1, 153-167.

[62] Mazzocco, M. (2007), "Household intertemporal behaviour: a collective characterization and a test of commitment", *Review of Economic Studies* 74, 857-895.

[63] Mazzocco, M., C. Ruiz, and S. Yamaguchi (2013), "Labor supply, wealth dynamics and marriage decisions", Working Paper, University of California at Los Angeles.

[64] McElroy, M. (1990), "The empirical content of Nash-bargained household behaviour", *Journal of Human Resources*, 25, 559-583.

[65] Oreffice, S. (2011), "Sexual orientation and household decision making. Same-sex couples balance of power and labor supply choices", *Labour Economics*, 18, 145-158.

[66] Samuelson, P. (1938), "A note on the pure theory of consumers behaviour", *Economica*, 5, 61-71.

[67] Samuelson, P. (1948), "Consumption theory in terms of revealed preference", *Economica*, 15, 243-253.

[68] Shapley, L. S., and M. Shubik, (1972), "The assignment game I: the core", *International Journal of Game Theory*, 1, 111–130.

[69] Stoye, J., and Y. Kitamura (2013), "Nonparametric analysis of random utility models: testing", CeMMAP Working Papers CWP36/13.

[70] Varian, H. (1982), "The nonparametric approach to demand analysis", *Econometrica*, 50, 945-972.

[71] Varian, H. (1988), "Revealed preference with a subset of goods", *Journal of Economic Theory*, 46, 179-185.

Appendix A: proofs

Proof of Lemma 1

Consider two utility levels u, u' with $\overline{U}_{m,w} \geq u > u' \geq 0$. Given that every solution to (3) with utility level u is also feasible with the utility level u' (by strict monotonicity of the utility functions), we have that $\psi_{m,w}(u') \geq \psi_{m,w}(u)$.

Let $q^w_{m,w}$ be part of the optimal solution for utility level u, then $q^w_{m,w}$ is strictly positive for at least one good. Indeed, otherwise we have $u^w(0, Q_{m,w}) = 0$ by assumption, which contradicts the inequality $u^w(0, Q_{m,w}) \geq u > u' \geq 0$. Given the continuity and strict monotonicity of the utility function u^w, we can take a tiny bit of these private goods from w and give them to m in such a way that w still receives utility level u' and the budget constraint is still satisfied. From this redistribution, we see that the utility level of m strictly increases, which means that the optimal solution must also strictly increase. This shows that $\psi_{m,w}(u') > \psi_{m,w}(u)$.

To show continuity, we consider

$$\begin{aligned} \phi_{m,w}(v) = & \max_{q^m_{m,w}, q^w_{m,w}, Q_{m,w}} u^w(q^w_{m,w}, Q_{m,w}) \\ \text{s.t. } & p_{m,w}(q^m_{m,w} + q^w_{m,w}) + P_{m,w} Q_{m,w} \leq y_{m,w}, \\ & v^m(q^m_{m,w}, Q_{m,w}) \geq v. \end{aligned}$$

The functions $\phi_{m,w}$ and $\psi_{m,w}$ are each other's inverse. To see this, assume that $u = \phi_{m,w}(v)$ and let $(q^m_{m,w}, q^w_{m,w}, Q_{m,w})$ be the solution to the woman's optimisation problem given v. Clearly, this bundle satisfies all restrictions for the man's optimisation problem, so $\psi_{m,w}(u) \geq v$.

We can prove $\psi_{m,w}(u) = v$ by contradiction. Assume that $\psi_{m,w}(u) > v$ and let $(q^m_{m,w}, q^w_{m,w}, Q_{m,w})$ be the optimal solution to the man's optimisation problem for u. This allocation is also feasible for the woman's optimisation problem given v. Moreover, $q^m_{m,w}$ is strictly positive for at least one good. Then, consider reallocating a tiny bit of these private goods from m and give them to w in such a way that m still receives utility level v and the budget constraint is still satisfied. This allows the woman to reach a utility level strictly above u. Thus, we obtain that $\phi_{m,w}(v) > u$, which gives the wanted contradiction.

We conclude that the function $\psi_{m,w}$ is a strictly monotone (invertible) function from an interval to an interval. As such, it must be continuous.

Proof of Proposition 2

Necessity. As a first step to deriving our revealed preference characterization, we define the first order conditions that are used to formulate this characterization. In particular, we consider these conditions for the optimization models that underlie our criteria of individual rationality and no blocking pairs:[26]

1. We begin with the two optimization problems for individual rationality. First, we consider the problem

$$\begin{aligned}(q^m_{m,\emptyset}, Q_{m,\emptyset}) = &\arg\max_{q^m,Q} v^m(q^m, Q)\\ &\text{s.t. } p_{m,\emptyset}q^m + P_{m,\emptyset}Q \leq y_{m,\emptyset},\end{aligned}$$

   i.e. $(q^m_{m,\emptyset}, Q_{m,\emptyset})$ represents the optimal allocation for m if he spends the income $y_{m,\emptyset}$. The first order conditions yield

$$\begin{aligned}\frac{\partial v^m(q^m_{m,\emptyset}, Q_{m,\emptyset})}{\partial q^m} &\leq \delta_{m,\emptyset}p_{m,\emptyset},\\ \frac{\partial v^m(q^m_{m,\emptyset}, Q_{m,\emptyset})}{\partial Q} &\leq \delta_{m,\emptyset}P_{m,\emptyset},\end{aligned}$$

   where $\delta_{m,\emptyset}$ is the Lagrange multiplier associated with the budget constraint and the expressions on the left hand side of the inequalities represent subdifferentials of the utility function v^m.

   Similarly, for the problem

$$\begin{aligned}(q^w_{\emptyset,w}, Q_{\emptyset,w}) = &\arg\max_{q^w,Q} u^w(q^w, Q)\\ &\text{s.t. } p_{\emptyset,w}q^w + P_{\emptyset,w}Q \leq y_{\emptyset,w},\end{aligned}$$

   we get the conditions

$$\begin{aligned}\frac{\partial u^w(q^w_{\emptyset,w}, Q_{\emptyset,w})}{\partial q^w} &\leq \lambda_{\emptyset,w}p_{\emptyset,w},\\ \frac{\partial u^w(q^w_{\emptyset,w}, Q_{\emptyset,w})}{\partial Q} &\leq \lambda_{\emptyset,w}P_{\emptyset,w},\end{aligned}$$

   where $\lambda_{\emptyset,w}$ is the Lagrange multiplier associated with the budget constraint.

2. Let us then turn to the optimization problems for no blocking pairs. Here, the

[26] We remark the formal similarity between the Pareto efficiency criterion in Definition 1 and the no blocking pair condition in Definition 3. Essentially, the condition in Definition 3 reduces to the one in Definition 1 for (m, w) with $w = \sigma(m)$. Therefore, we can follow a directly analogous reasoning as under item 2. below to obtain the rationalizability conditions in Proposition 2 that pertain to our Pareto efficiency requirement (compare with Cherchye, De Rock and Vermeulen, 2011). For compactness, we do not include this reasoning here.

optimization problem is defined as

$$\begin{aligned}(q^m_{m,w}, q^w_{m,w}, Q_{m,w}) = \arg\max_{q^m,Q} & \, v^m(q^m, Q) \\ \text{s.t. } & \, p_{m,w}q^m + P_{m,w}Q \leq y_{m,w}, \\ & \, u^w(q^w, Q) \geq U_{m,w}.\end{aligned}$$

i.e. $(q^m_{m,w}, q^w_{m,w}, Q_{m,w})$ represents the allocation chosen by m if he could freely spend the entire income $y_{m,w}$ given that w should receive utility level $U_{m,w}$. The corresponding first order conditions give

$$\begin{aligned}\frac{\partial v^m(q^m_{m,w}, Q_{m,w})}{\partial q^m} &\leq \delta_{m,w}p_{m,w}, \\ \mu_{m,w}\frac{\partial u^w(q^w_{m,w}, Q_{m,w})}{\partial q^w} &\leq \delta_{m,w}p_{m,w}, \\ \frac{\partial v^m(q^m_{m,w}, Q_{m,w})}{\partial Q} + \mu_{m,w}\frac{\partial u^w(q^m_{m,w}, Q_{m,w})}{\partial Q} &\leq \delta_{m,w}P_{m,w},\end{aligned}$$

where $\delta_{m,w}$ is the Lagrange multiplier associated with the budget constraint and $\mu_{m,w}$ is the Lagrange multiplier associated with the utility constraint.

In what follows, we use $\lambda_{m,w} = \delta_{m,w}/\mu_{m,w}$, $\frac{\partial u^w(q^w_{m,w},Q_{m,w})}{\partial Q} = \lambda_{m,w}P^w_{m,w}$ and $P^m_{m,w} = P_{m,w} - P^w_{m,w}$ (which implies $\frac{\partial v^m(q^m_{m,w},Q_{m,w})}{\partial Q} \leq \delta_{m,w}P^m_{m,w}$).

In a final step, we can define the characterization in Proposition 2 by combining the above first order conditions with the postulated concavity property of the utility functions u^w and v^m. In particular, concavity implies (for any $q^{m\prime}$, $q^{w\prime}$, $q^{m\prime\prime}$, $q^{w\prime\prime} \in \mathbb{R}^n_+$ and Q', $Q'' \in \mathbb{R}^k_+$)

$$\begin{aligned}v^m(q^{m\prime}, Q') - v^m(q^{m\prime\prime}, Q'') &\leq \frac{\partial v^m(q^{m\prime\prime}, Q'')}{\partial q^m}(q^{m\prime} - q^{m\prime\prime}) + \frac{\partial v^m(q^{m\prime\prime}, Q^{m\prime\prime})}{\partial Q}(Q' - Q''), \\ u^w(q^{w\prime}, Q') - u^w(q^{w\prime\prime}, Q'') &\leq \frac{\partial u^w(q^{w\prime\prime}, Q'')}{\partial q^w}(q^{w\prime} - q^{w\prime\prime}) + \frac{\partial u^m(q^{w\prime\prime}, Q^{w\prime\prime})}{\partial Q}(Q' - Q'').\end{aligned}$$

Then, we obtain the rationalizability conditions in Proposition 2 by using $v^m(q^m_{m,w}, Q_{m,w}) = V_{m,w}$, and $u^w(q^w_{m,w}, Q_{m,w}) = U_{m,w}$ ($m \in M \cup \{\varnothing\}$ and $w \in W \cup \{\varnothing\}$).

Sufficiency. To obtain the sufficiency result, we consider

$$\begin{aligned}v^m(q^m, Q) &= \min_{w\in W\cup\{\varnothing\}}[V_{m,w} + \delta_{m,w}\left(p_{m,w}(q^m - q^m_{m,w}) + P^m_{m,w}(Q - Q_{m,w})\right)], \\ u^w(q^w, Q) &= \min_{m\in M\cup\{\varnothing\}}[U_{m,w} + \lambda_{m,w}\left(p_{m,w}(q^w - q^w_{m,w}) + P^w_{m,w}(Q - Q_{m,w})\right)].\end{aligned}$$

Varian (1982) shows, in a unitary context, that $v^m(q^m_{m,w}, Q_{m,w}) = V_{m,w}$ and $u^w(q^w_{m,w}, Q_{m,w}) = U_{m,w}$ ($m \in M \cup \{\varnothing\}$ and $w \in W \cup \{\varnothing\}$). Using this, we can use a readily similar argument as in Varian (1982) (for the unitary consumption model) and Cherchye, De Rock and Vermeulen (2011) (for the collective consumption model) to show that the utility functions v^m and u^w defined above rationalize the data set $\mathcal{D}$ by a stable matching (i.e. the data solve the optimization problems underlying our stability criteria for these functions v^m and u^w).

Proof of Proposition 3

First, conditions (a) and (e) in Proposition 2 define the constraints

$$q^m_{m,\sigma(m)} + q^w_{m,\sigma(m)} = q_{m,\sigma(m)} \text{ and } P^m_{m,w} + P^w_{m,w} = P_{m,w}.$$

Next, the individual rationality constraints (iii) together with the Afriat inequalities (i) (for male m) and (ii) (for female w) in Proposition 2 give

$$\begin{aligned} 0 &\leq \left[p_{m,\emptyset}\left(q^m_{m,\sigma(m)} - q^m_{m,\emptyset}\right) + P_{m,\emptyset}\left(Q_{m,\sigma(m)} - Q_{m,\emptyset}\right)\right], \\ 0 &\leq \left[p_{\emptyset,w}\left(q^w_{\sigma(w),w} - q^w_{\emptyset,w}\right) + P_{\emptyset,w}\left(Q_{\sigma(w),w} - Q_{\emptyset,w}\right)\right]. \end{aligned}$$

In turn, this obtains

$$\begin{aligned} y_{m,\emptyset} &\leq p_{m,\emptyset} q^m_{m,\sigma(m)} + P_{m,\emptyset} Q_{m,\sigma(m)}, \\ y_{\emptyset,w} &\leq p_{\emptyset,w} q^w_{\sigma(w),w} + P_{\emptyset,w} Q_{\sigma(w),w}. \end{aligned}$$

Similarly, the no blocking pairs condition (iv) together with the Afriat inequalities (i) and (ii) give that, for all m, w such that $\sigma(m) \neq w$,

$$\begin{aligned} 0 &\leq \left[p_{m,w}\left(q^m_{m,\sigma(m)} - q^m_{m,w}\right) + P^m_{m,w}\left(Q_{m,\sigma(m)} - Q_{m,w}\right)\right], \\ 0 &\leq \left[p_{m,w}\left(q^w_{\sigma(m),w} - q^w_{m,w}\right) + P^w_{m,w}\left(Q_{\sigma(w),w} - Q_{m,w}\right)\right]. \end{aligned}$$

The first inequality states that the man m should not prefer his allocation in (m, w) over his matching allocation (in revealed preference terms). The second condition does the same for woman w. Now, adding these two equations together yields

$$y_{m,w} \leq p_{m,w} q^m_{m,\sigma(m)} + p_{m,w} q^w_{\sigma(w),w} + P^m_{m,w} Q_{m,\sigma(m)} + P^w_{m,w} Q_{\sigma(w),w}.$$

Appendix B: information on marriage markets

Marriage market	Number of couples	Couples' characteristics
1	16	Husband between 25 and 35 years old, no college degree, in a nonmixed couple
2	1	Husband between 25 and 35 years old, no college degree, in a mixed couple
3	8	Husband between 25 and 35 years old, college degree, in a nonmixed couple
4	3	Husband between 25 and 35 years old, college degree, in a mixed couple
5	47	Husband between 35 and 45 years old, no college degree, in a nonmixed couple
6	6	Husband between 35 and 45 years old, no college degree, in a mixed couple
7	22	Husband between 35 and 45 years old, college degree, in a nonmixed couple
8	4	Husband between 35 and 45 years old, college degree, in a mixed couple
9	58	Husband between 45 and 55 years old, no college degree, in a nonmixed couple
10	8	Husband between 45 and 55 years old, no college degree, in a mixed couple
11	28	Husband between 45 and 55 years old, college degree, in a nonmixed couple
12	6	Husband between 45 and 55 years old, college degree, in a mixed couple
13	31	Husband between 55 and 65 years old, no college degree, in a nonmixed couple
14	7	Husband between 55 and 65 years old, no college degree, in a mixed couple
15	17	Husband between 55 and 65 years old, college degree, in a nonmixed couple
16	2	Husband between 55 and 65 years old, college degree, in a mixed couple

Appendix C: sharing rule and wage ratio for subgroups

Starting from our original data set, we created subsamples based on age, education, ethnicity, and having children or not. More precisely, we considered (i) four husband age groups (i.e. husband age within 25-35, 35-45, 45-55 and 55-65), (ii) two husband education groups (i.e. low and high education), (iii) mixed versus non-mixed couples, and (iv) having children or not. Figures 6-9 report on wage effects for the total sharing rule, and Figures 10-13 on wage effects for the conditional sharing rule. These figures also give an idea of differences in levels and changes for the different subgroups. We find that patterns may be slightly different across subgroups, but the main qualitative conclusions are the same as for our core analysis in the main text. The same applies to the other exercises that we conducted in Section 4 (on income effects and the budget share of public consumption). For compactness, we do not include these last results for the subgroups.

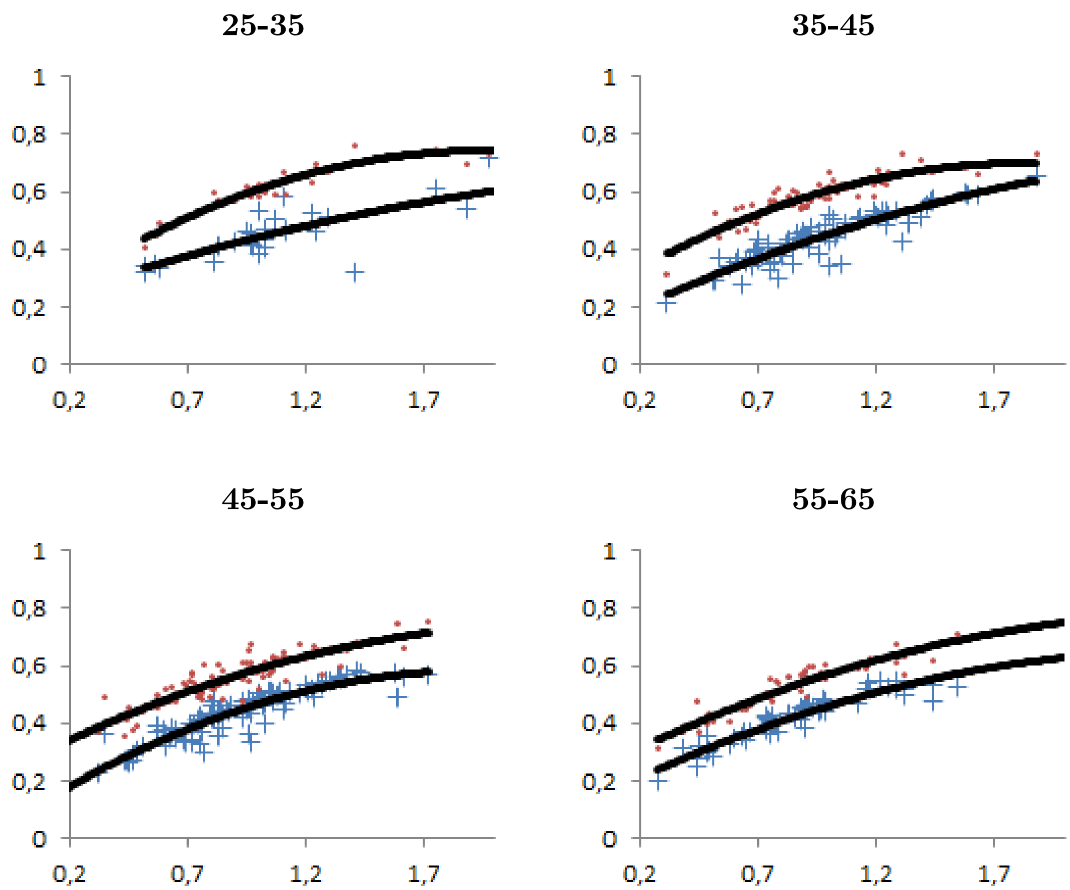

Figure 6: Total sharing rule and wage ratio per husband age group.

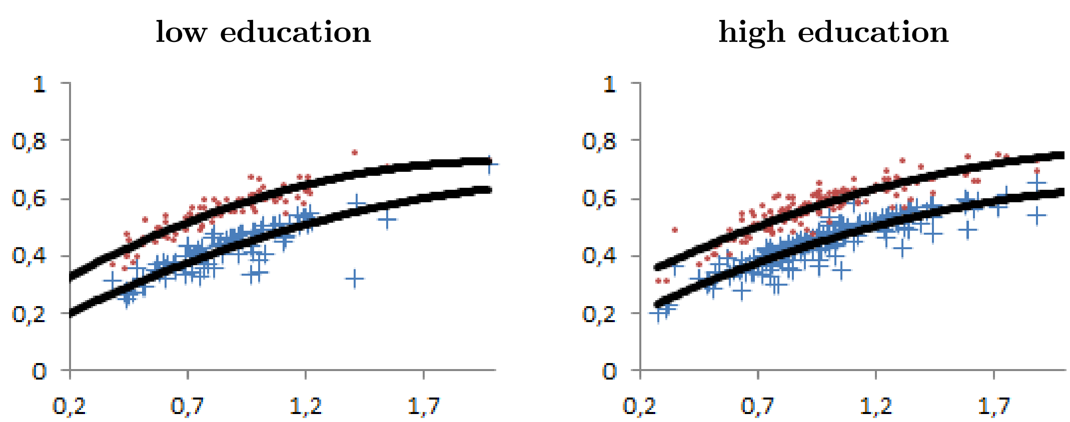

Figure 7: Total sharing rule and wage ratio per husband education group.

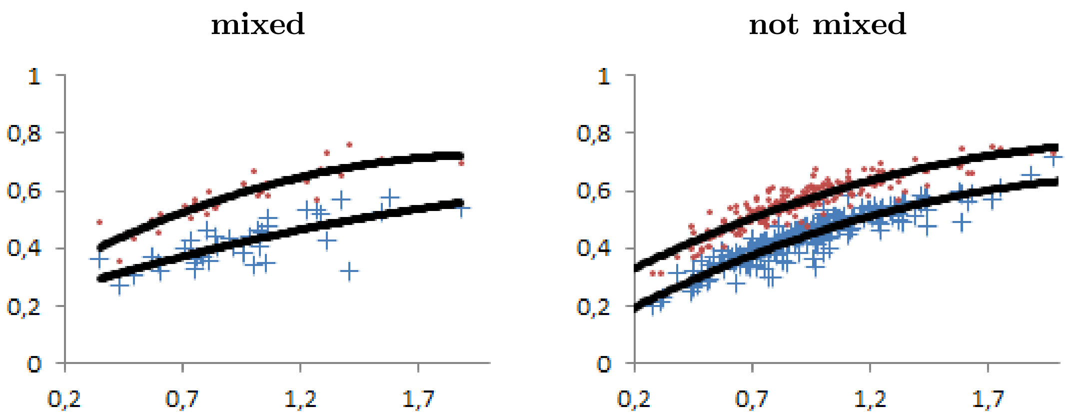

Figure 8: Total sharing rule and wage ratio, mixed verus non-mixed couples.

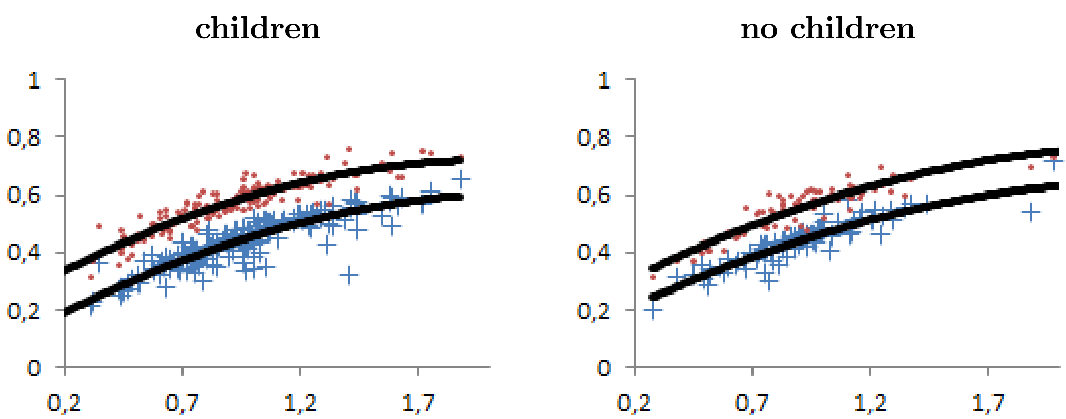

Figure 9: Total sharing rule and wage ratio, children versus no children.

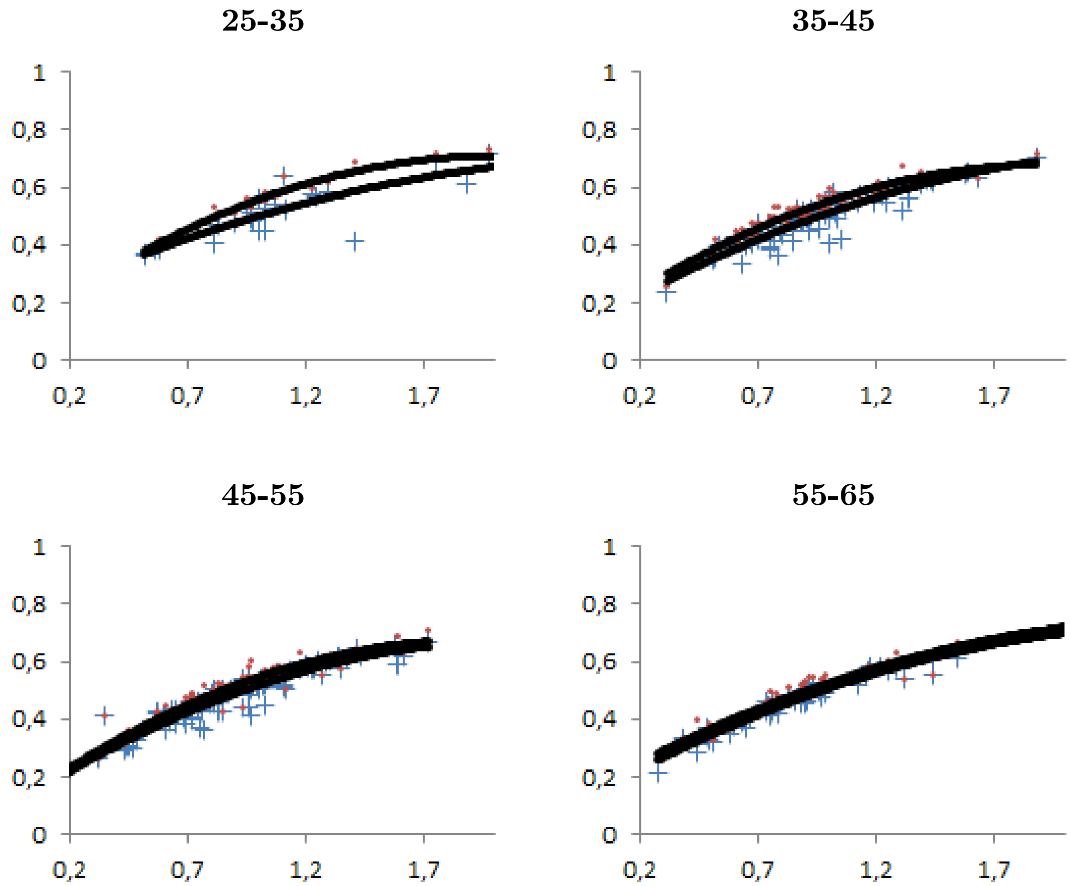

Figure 10: Conditional sharing rule and wage ratio per husband age group.

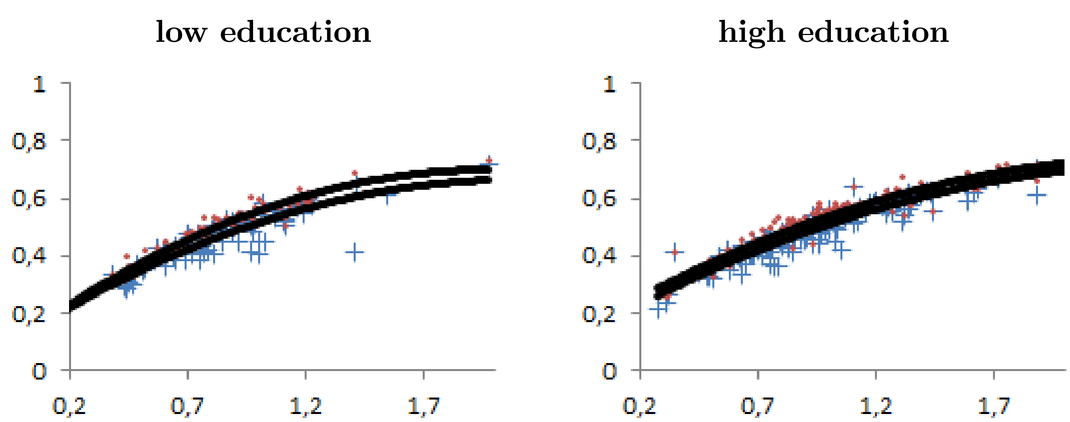

Figure 11: Conditional sharing rule and wage ratio per education group.

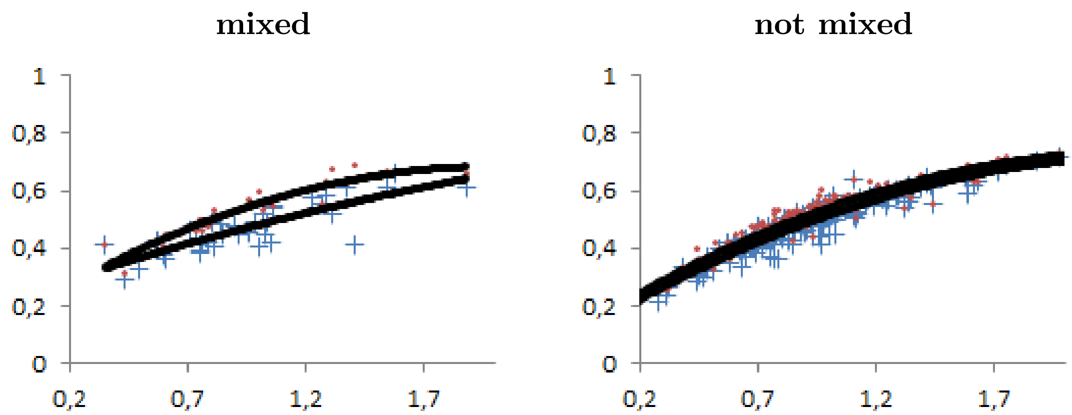

Figure 12: Conditional sharing rule and wage ratio, mixed versus non-mixed couples.

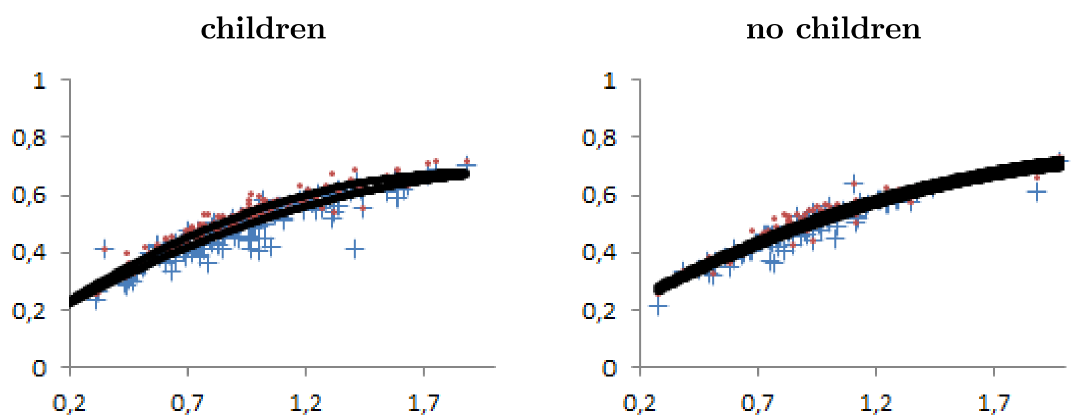

Figure 13: Conditional sharing rule and wage ratio, children versus no children.