Household consumption when the marriage is stable*

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

September 2, 2014

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.

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 relevant 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,

*We thank Martin Browning, Jan Eeckhout, Patrick Legros and seminar participants in Leicester, Oxford, Tilburg, Toulouse and ESEM 2014 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.

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, again, Browning and Chiappori (1998), Cherchye and Vermeulen (2008) and Cherchye, De Rock and Vermeulen (2009).

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 gain 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 exactly 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.[1] 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.

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 approach typically (only) exploits the assumption that intrahousehold consumption is Pareto efficient

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

(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), Cherchye, De Rock, Lewbel and Vermeulen (2013) and Dunbar, Lewbel and Pendakur (2013) 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 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. But, 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. 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 giving 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.

[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.

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 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).[4]

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. The Appendix contains the proofs of our main results.

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. Before doing so, we first introduce some necessary notation.

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

[3] See, for example, Chiappori, Iyigun, Lafortune and Weiss (2013), Chiappori, Oreffice and Quintana-Domeque (2012), Choo and Siow (2006), Dupuy and Galichon (2012), Galichon and Salanié (2014) and Jacquemet and Robin (2013) 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 matchings for household consumption.

[4] We remark that our preference assumptions do 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, non-uniqueness does not interfere with the validity (and, thus, applicability) of the testable implications and (set) identification results that we derive below.

$\sigma : M \cup W \to M \cup W \cup \{\emptyset\}$. 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 women 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. 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. If the spouses' leisure is taken up in the analysis, then their wages will be elements of these price vectors. 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 women.[5] 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 women w. 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.

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 $S = \{(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m \in M}$ 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 continuous, strictly increasing 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 continuous, strictly increasing and concave utility function $u^w : \mathbb{R}^{n+k} \to \mathbb{R}$. We assume that males and females have complete information about each others' preferences.[6]

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 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

[5] 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.

[6] 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.

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 $S = \{(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m\in M}$ is* ***Pareto efficient*** *if, for all $m \in M$, there exists no feasible allocation (q^m, q^w, Q), i.e.*

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

such that

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

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 $V_{m,\emptyset}$ and $U_{\emptyset,w}$ represent the maximum utility levels that man m and, respectively, woman w could obtain by staying single, when faced with the prices $p_{m,\emptyset}$ and $p_{\emptyset,w}$ and incomes $y_{m,\emptyset}$ 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 $S = \{(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m\in M}$ 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 $p_{m,w}$ 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. We obtain the next condition, which is formally close to the Pareto efficiency condition that we defined above. The main difference is that, for any man m and woman w, we now consider all other potential partners on the marriage market (i.e. $w \neq \sigma(m)$).

Definition 3 *For a given matching σ, the matching allocation $S = \{(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m\in M}$ has* ***no blocking pairs*** *if, for all $m \in M$ and $w \in W$ with $w \neq \sigma(m)$, there exists no feasible allocation (q^m, q^w, Q), i.e.*

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

such that

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

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 $S = \{(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m\in M}$ is* ***stable*** *if it is Pareto optimal, individually rational and has no blocking pair.*

To conclude this section, we provide an alternative formulation of the no blocking pairs criterion in Definition 3.[7] While this formulation is somewhat less intuitive, it will be instrumental to introduce our revealed preference characterization in the next section. Specifically, for an unmatched pair (m, w), let $\overline{V}_{m,w}$ represent the utility level of male m if he were to spend the entire income $y_{m,w}$ (which means that m is the "dictator" in the pair (m, w)), i.e.

$$\overline{V}_{m,w} = \max_{\overline{q}^m, \overline{Q}} v^m(\overline{q}^m, \overline{Q}) \text{ s.t. } p_{m,w}\overline{q}^m + P_{m,w}\overline{Q} \leq y_{m,w}. \tag{3}$$

Then, if $\overline{V}_{m,w} \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)})$, we can define $U_{w,m}$ as the maximum utility for woman w under the constraint that m gets at least his utility under the matching σ, i.e.

$$\begin{aligned} U_{m,w} = &\max_{q^m, q^w, Q} u^w(q^w, Q) \qquad (4)\\ \text{s.t. } &v^m(q^m, Q) \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}),\\ \text{and } &p_{m,w}(q^m + q^w) + P_{m,w}Q \leq y_{m,w}. \end{aligned}$$

We note that the condition $\overline{V}_{m,w} \geq v^m(q^m_{m,\sigma(m)})$ effectively guarantees that the two constraints $v^m(q^m, Q) \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}$ and $p_{m,w}(q^m + q^w) + P_{m,w}Q \leq y_{m,w}$ can be satisfied simultaneously, i.e. $U_{m,w}$ is well defined.

Using these definitions, we conclude that the matching allocation S has no blocking pairs if, for any pair (m, w) with $w \neq \sigma(m)$, we have either

$$\overline{V}_{m,w} < v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}), \tag{5}$$

$$\text{or, if } \overline{V}_{m,w} \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)}), \text{ then } u^w(q^w_{\sigma(w),w}, Q_{\sigma(w),w}) \geq U_{m,w}. \tag{6}$$

The first condition states that a matching is always stable if the maximal utility that m can obtain with income $y_{m,w}$ (i.e. $\overline{V}_{m,w}$) is below his utility under the matching σ (i.e. $v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)})$). Indeed, in this case it is impossible for the male m to be better off in the pair (m, w) than under the matching σ. A fortiori, this implies that (m, w) is not a blocking pair. Next, the second condition pertains to the case with $\overline{V}_{m,w} \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)})$. In such a situation, we consider $U_{m,w}$, i.e. the maximal utility of w under the restriction that m must get the utility under the given matching σ. Then, for the pair (m, w) not to be a blocking one, we need that $U_{m,w}$ does not exceed w's utility under the matching σ.

As a final remark, we note that the criterion defined by conditions (5) and (6) is actually not exactly equivalent to the criterion in Definition 3. In some pathological cases, the conditions (5) and (6) might not be sufficient to exclude all blocking pairs. In particular, this occurs when Pareto frontiers for matched pairs $(m, \sigma(m))$ are not downward sloping. Thus, in what follows we will implicitly assume that individual utilities v^m and u^w are of a form that excludes such pathological situations.

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,

[7] Our alternative formulation will consider the dictator outcome for the male m (see (3) below). In fact, using the female dictator outcome w yields a criterion that is formally equivalent, except in some pathological cases (which we mention at the end of this section).

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. Finally, we conclude this section by introducing a method that allows us to evaluate deviations from "exactly" stable marriage. For example, such deviations may occur because there are costs of divorce (because of frictions on the marriage market or unobserved consumption in marriage). Interestingly, this extension will also enable us to analyze marriage behavior that is "close" to stable (instead of exactly stable), which will prove useful for our empirical application in Section 4.

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 $\{(q^m_{m,\sigma(m)}, q^w_{m,\sigma(m)}, Q_{m,\sigma(m)})\}_{m \in M}$ 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.[8] 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.

Revealed preference characterization. The next Theorem 1 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 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 Theorem 1.

Theorem 1 *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 $m \in M$, individual quantities $q^m_{m,\sigma(m)}$, $q^w_{m,\sigma(m)} \in \mathbb{R}^n_+$ that satisfy

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

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

b. for each unmatched pair $m \in M$ and $w \in M$ (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 $q^m_{m,\emptyset}$, $\overline{q}^m_{m,w} \in \mathbb{R}^n_+$ and public quantities $Q_{m,\emptyset}$, $\overline{Q}_{m,w} \in \mathbb{R}^k_+$ that satisfy

$$p_{m,\emptyset}q^m_{m,\emptyset} + P_{m,\emptyset}Q_{m,\emptyset} = y_{m,\emptyset} \text{ and } p_{m,w}\bar{q}^m_{m,w} + P_{m,w}\bar{Q}_{m,w} = y_{m,w},$$

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 (w, m) ($m \in M$, $w \in M$), 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 $V_{m,w}$, $V_{m,\emptyset}$, $\bar{V}_{m,w}$, $U_{m,w}$, $U_{\emptyset,w}$ and $\delta_{m,w}$, $\delta_{m,\emptyset}$, $\bar{\delta}_{m,w}$, $\lambda_{m,w}$, $\lambda_{\emptyset,w}$ (for any $m \in M$ and $w \in W$) that simultaneously meet the following constraints:

[8] 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.

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

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

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 $m \in M$ and $w \in M$ (with $\sigma(m) \neq w$), i.e.

$$
\bar{V}_{m,w} < V_{m,\sigma(m)} \text{ or } \left(V_{m,w} = V_{m,\sigma(m)} \text{ and } U_{\sigma(w),w} \geq U_{m,w}\right).
$$

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 when "dictator" in the pairs (m, w) (see (3)) 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 pairs). 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.

Theorem 1 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 $V_{m,w}, V_{m,\emptyset}, \bar{V}_{m,w}$ represent male utilities in alternative decision situations (respectively, in the pair (w, m), under single status, and as a "dictator" in the pair (m, w) (see again (3))). A directly similar interpretation applies to the numbers $U_{m,w}$ and $U_{\emptyset,w}$, which represent female utilities. Next, the numbers $\delta_{m,w}, \delta_{m,\emptyset}, \overline{\delta}_{m,w}$ (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 (continuous, strictly increasing and concave) utility functions v^m and u^w that explain the data. First, the inequalities ensure that these functions satisfy the Pareto efficiency criterion in Definition 1.[9] Next, they also guarantee that the Afriat numbers $V_{m,\emptyset}$, $U_{\emptyset,w}$, $\bar{V}_{m,w}$, $V_{m,w}$ and $U_{m,w}$ solve the problems (1), (2), (3) and (4) for the given specification of v^m and u^w. Given this, the conditions (iii) and (iv) impose consistency with the individual rationality criterion in Definition 2 and the no blocking pairs condition in Definition 3 (expressed in the form of (5)-(6)).

Linear conditions. The characterization of rationalizability in Theorem 1 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 Theorem 1. 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 1 *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 $m \in M$, individual quantities $q^m_{m,\sigma(m)}$, $q^w_{m,\sigma(m)} \in \mathbb{R}^n_+$ that satisfy

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

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

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

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

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 $w \in M$ (with $\sigma(m) \neq w$), 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 Theorem 1, 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 1 that are linear in the unknown quantities $q^m_{m,\sigma(m)}$, $q^w_{m,\sigma(m)}$ and prices $P^m_{m,w}$, $P^w_{m,w}$.[10]

[9] 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 conditions (1) and (2) of Theorem 1.

[10] It is interesting to observe that the linear conditions in Proposition 1 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.

Again, we can give a specific "revealed preference" interpretation to the different conditions in Proposition 1. The adding up restrictions (a) and (b) also appeared in Theorem 1. 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 1 are linear in unknown quantities and prices, they define testable implications for rationalizability that can 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 1 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 1, which effectively "set" identifies these unobservables (under the maintained assumption of a stable matching).[11]

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 $y^w_{m,\sigma(m)}$ as

$$y^m_{m,\sigma(m)} = p_{m,\sigma(m)}q^m_{m,\sigma(m)} + P^m_{m,\sigma(m)}Q_{m,\sigma(m)}, \tag{7}$$

$$y^w_{m,\sigma(m)} = p_{m,\sigma(m)}q^w_{m,\sigma(m)} + P^w_{m,\sigma(m)}Q_{m,\sigma(m)}. \tag{8}$$

We remark that $y^m_{m,\sigma(m)} + y^w_{m,\sigma(m)} = y_{m,\sigma(m)}$ 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 (2013) 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 (7) and (8) subject to the linear rationalizability restrictions in Proposition 1. 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, 2013) or observability of a discrete set of household consumption choices (see, for example, Cherchye, De Rock and Vermeulen, 2011).

[11] Recall that we focus on a necessary condition for rationalizability by a stable matching (see Proposition 1). This implies that the identified sets will be bigger than those that would result from the necessary and sufficient condition outlined in Theorem 1.

Deviations from "exact" stability. Proposition 1 defines "sharp" conditions for rationalizability by a stable matching, which only tell us whether or not the observed matching allocation is "exactly" stable. In practice, however, observed behavior that is inconsistent with exact stability may well be close to stable. Alternatively, a matching allocation can be 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).

Following this perspective, it is useful to quantify the 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 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 1, 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

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

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 1. 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 following empirical application, we will measure the degree of stability of our data set $\mathcal{D}$ by solving

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

subject to the feasibility constraints (a) and (b) in Proposition 1 and the linear constraints (9) and (10). It directly follows that the outcome of this linear programming problem gives us a measure of how close observed behavior (with original income values) is to rationalizable behavior. Next, if we carry out the (post-divorce) income decreases defined by the optimal values of $s^{IR}_{m,\emptyset}$, $s^{IR}_{\emptyset,w}$ and $s^{NBP}_{m,w}$, we 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 above.

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 a number of goods). We will first check consistency of our data with the rationalizability conditions in Proposition 1. Because our data will fail these sharp conditions (i.e. behavior is not exactly stable), we follow the procedure outlined at the end of the previous section to compute stability indices that rationalize the observed behavior in terms of divorce/remarriage costs. We will also consider the inter-household heterogeneity of these stability indices, which thus captures variation in divorce/remarriage costs across households. Using our stability indices, we can address sharing rule (set) identification, which we will do in the final step of our analysis.

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 variety of economic and socio-demographic variables.[12] Because we need to assume that individuals are active on the same marriage market, 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 40. We include both couples with and without children.[13] Next, we excluded the self-employed to avoid issues regarding 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), we obtained a sample of 62 households.[14]

Table 1 provides summary statistics on the data for 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 10 hours per day for sleeping and personal care (i.e. leisure = 168 - 70 - hours worked). Full income and (Hicksian) consumption are measured in euros per week. For completeness, Table 1 also reports on the male and female ages and the number of children in the households under consideration.

Our data set contains assignable consumption.[15] 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.[16] But the main part of the observed household consumption is nonassignable.[17] In our analysis, we assume 75% of this nonassignable consumption is privately consumed and 25% is publicly consumed within the household. This specific subdivision is based on an empirical goodness-of-fit criterion. In particular, it gives the best empirical fit of the stability conditions in Proposition 1 when using the stability index in (11) as our goodness-of-fit measure. We obtain the highest index value for the assumption that three quarters of the nonassignable consumption is private and the remaining quarter is public; the other subdivisions we checked yield lower index values.

Finally, our method requires prices and incomes that apply to the exit options from marriage (i.e. become 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.[18] Next, to reconstruct the potential full income in the unobserved outside options, we must

[12] 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.

[13] 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.

[14] We remark that our analysis does not need that each individual in our sample effectively knows all the individuals of the other gender. It suffices that (s)he knows at least one individual of the same type as each other observed individual (thus defining the associated exit option from marriage).

[15] 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 1 through appropriately defined linear constraints, which define feasibility bounds on the variables $q^m_{m,\sigma(m)}$ and $q^w_{\sigma(w),w}$.

[16] 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. To account for reporting error, we treat only 95% of these reported quantities as effectively assignable, and consider the remaining 5% as private but nonassignable.

[17] 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). It seems reasonable to assume that these goods are partly privately and partly publicly consumed. See also Browning, Chiappori and Lewbel (2013) for a collective consumption model that accounts for goods that are both privately and publicly consumed within the household.

[18] 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.

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.[19]

Table 1: Data summary statistics

	Full income	**Male wage**	**Female wage**
Mean	1406.66	12.05	11.98
St. dev.	421.12	3.73	3.39
Maximum	2614.10	26.54	26.96
Minimum	434.86	7.21	5.16
	Nonassignable	**Assignable private consumption**	
	Private consumption	**Male**	**Female**
Mean	467.70	75.82	81.43
St. dev.	176.34	46.97	50.95
Maximum	985.10	263.08	293.77
Minimum	29.71	8.33	12.06
	Public consumption	**Male leisure**	**Female leisure**
Mean	154.57	27.85	23.92
St. dev.	58.57	13.64	11.81
Maximum	326.539	66.50	53.17
Minimum	9.52	2	0
	Number of children	**Male age**	**Female age**
Mean	1.26	35.10	33.39
St. dev.	1.16	3.52	3.81
Maximum	4	40	40
Minimum	0	26	26

Note: full income and consumption are in euro per week,
wages in euro per hour, and leisure in hours per week.

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 1. As we discussed at the end of Section 3, a possible explanation is that the observed matching allocation is actually stable but we need to account for a cost of divorce (due to frictions on the marriage market and/or unobserved benefits from marriage). Therefore, we next compute the cost of divorce that is needed to obtain rationalizability. We do so by means of the stability measure that we defined in (11). This measure also indicates how close the observed behavior (with original income values) is to rationalizable behavior.

Table 2 provides summary statistics for our stability indices. We report results for all stability restrictions together and, for completeness, we also include separate results for the individual rationality and no blocking

[19] As compared to the alternative that fixes the intrahousehold distribution of nonlabor income (e.g. 50% for each individual), this solution 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.

pair restrictions. First, even though the observed behavior is not exactly stable, we find that the average index values are very close to unity. In words, when defined over all different exit options for the 62 households in our sample, we need only a rather small average income loss to rationalize the observed consumption and marriage behavior. However, when looking at the minimum index values, we also find for at least one individual and one (potentially blocking) pair that the required income loss amounts to almost 9%. In what follows (see our discussion of Table 3), we will use these stability indices for individual exit options to define household-specific stability measures that capture household-level divorce costs (under the assumption of stable marriage). Finally, Table 2 shows the fraction of individual rationality constraints and no blocking pair constraints with a stability index value below unity (see the rows "constraints < 1", which are based on (9) and (10)). When using this criterion, we conclude that we need a higher incidence of income loss associated with remarriage than with becoming single (i.e. 7.43% versus 2.72% of the corresponding exit options) if we want to rationalize the observed behavior.

Table 2: Stability index results

	All (individual rationality and no blocking pairs)
Mean index (st. dev.)	99.82% (0.83%)
Minimum index	91.21%
Constraints < 1	7.27% (284 out of 3906)
	Individual rationality
Mean index (st. dev.)	99.85% (1.06%)
Minimum index	91.21%
Constraints < 1	2.42% (3 out of 124)
	No blocking pairs
Mean index (st. dev.)	99.82% (0.82%)
Minimum index	91.46%
Constraints < 1	7.43% (281 out of 3782)

Inter-household heterogeneity. Our stability index also allows us to investigate inter-household heterogeneity in marriage stability. Specifically, we consider two measures of stability per individual household. Our first measure "stability violations" computes, for each different household, the number of individual rationality and no blocking pair restrictions with a stability index below unity, which means that these restrictions are violated in their "sharp" form. A higher value then indicates that more outside options dominate the current marriage matching if there were no cost of divorce (captured by the stability index). Our second measure "household stability index" calculates the severeness of these household-specific stability violations. In particular, we define it as the minimum stability index value over the violated restrictions. Following our above reasoning, this measure gives an indication of the cost of divorce that we need to rationalize the marriage behavior of a given household. Alternatively, we can also interpret it as quantifying how close a particular household is to an exactly stable marriage. In this sense, a higher value suggests a more stable couple.

Table 3 gives summary information on the distribution of these two measures for our sample of households. A first notable finding is that no couple is exactly stable. The minimum number of stability violations equals 2. In other words, if we assume that there is no cost of divorce, then any household has 2 or more exit options (remarry or become single) that dominate the current marriage. The maximum number in Table 3 reveals a situation with violations of no less than 62 stability restrictions (out of a total of 124), which means that half of the outside options is more attractive than the given marriage (in the absense of divorce costs). However, for most households the number of violated stability restrictions is fairly low, with a mode of 5 (i.e. about 4% of all restrictions).[20] To evaluate the severeness of these stability violations, we next consider our household-specific

[20] A closer inspection of our results reveals a very asymmetric distribution of stability violations for our sample. There are 5 households with at least 47 violations (i.e. 47, 59, 60 and two times 62 violations), while any of the remaining 57 households has no more than 7 violations.

stability index. As explained above, this measure indicates the cost of divorce that is needed to rationalize household behavior in terms of a stable marriage. We obtain that the average household stability index amounts to about 95%, which corresponds to a cost of divorce that equals approximately 5% of the income when exiting marriage. Importantly, we also observe quite some variation across households. For example, the maximum index value is almost 98% (i.e. a divorce cost of only 2%), while the minimum index value is close to 91% (i.e. a divorce cost of no less than 9%).

Table 3: Household-specific stability measures

	Stability violations	**Household stability index**
Mean (st. dev.)	9.11 (14.71)	94.85% (1.20%)
Minimum	2	91.21%
First quartile	5	94.44%
Median	5	95.08%
Third quartile	5	95.44%
Maximum	62	97.66%

Note: there are 2 individual rationality and 122 no blocking pair restrictions per household, so that the the maximum number of stability violations equals 124.

Our household stability index quantifies how close an observed marriage is to exact stability. From this perspective, it seems interesting to relate the inter-household heterogeneity that is revealed in Table 3 to observable household characteristics. This can provide insight into which household types are systematically more stable than others. To this end, we relate our household stability index to the intrahousehold wage and age structure, the number of children and the household's full income.

Table 4 reports the results of our regression exercise. Interestingly, even though our sample contains only 62 household observations, we do find a number of significant associations. First, we observe that a higher average wage, a larger number of children and a higher full income are associated with a couple's consumption and marriage behavior that is closer to exact stability. Intuitively, increasing the average wage, the number of children or the household income makes outside options less attractive. For children, we obtain this effect because children typically generate a larger gain from marriage.[21] In a similar vein, we have that higher wages and full income imply greater consumption possibilities.[22] In each case, it becomes more difficult to find an exit option that dominates the current marriage. Clearly, if outside options become less attractive, then we need a smaller cost of divorce to rationalize the observed marriage behavior, which obtains a higher household stability index.

Next, we find that a larger (absolute) wage difference between spouses is associated with household behavior that is further away from exact stability. Again, we can give this association an intuitive interpretation. If one household member has a substantially larger wage than his/her partner, then (s)he will have more attractive outside options. In turn, this can make it more difficult to rationalize the couples' marriage behavior, unless we account for a higher cost of divorce (i.e. frictions on the marriage market and/or unobserved benefits from marriage).

Finally, the age structure does not seem to have a significant impact on stability of marriage. We believe this result may at least partially follow from our sample construction. We only selected couples with both adults aged between 25 and 40 (because all individuals are assumed to be active on the same marriage market). Thus, it may well be that the age variation in our sample is too small to identify a significant age effect.[23]

[21] In our application, we treat expenditures on children as public consumption. It is through this channel that children give rise to a surplus of marriage in our consumption set-up.

[22] In our application, average wage changes can differ from full income changes because of non-labor income.

[23] In this respect, we note that the average age effect is close to significant (p-value of 0.14 for the null hypothesis that there is no effect). The negative sign may indicate that older couples experience a higher match quality (also because of "survival" effects), which implies a higher cost of divorce/remarriage (and, thus, a lower stability index). See, for example, Mazzocco, Ruiz and Yamaguchi (2013) and Adams, Cherchye, De Rock and Verriest (2014) for related empirical findings.

Table 4: Stability of marriage and household characteristics

Household stability index	coefficient	standard error
Intercept	0.82035***	0.02535
Average wage	0.00149***	0.00045
Absolute wage difference	-0.00086**	0.00038
Average age	-0.00059	0.00040
Absolute age difference	0.00053	0.00063
Number of children	0.00330***	0.00117
Full income (log)	0.01776***	0.00374
R^2	56.73%	

Note: average (w)age and (w)age differences defined over spouses;
** = significant at 5% level, *** = significant at 1% level.

Sharing rule identification. By using the stability index values that are summarized in Tables 2 and 3, we can construct a new data set (with divorce costs) that is rationalizable by a stable matching. Then, we can use the methodology outlined above to (set) identify the decision structure underlying the observed stable marriage behavior.[24] We illustrate this for the case of sharing rule identification. 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 allocated to the male (female).

The results of this exercise are summarized in Table 5. 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. The table reports on the percentage point difference between upper and lower bounds on the female relative income share. It does so for the naive bounds and the stable bounds. In addition, it gives summary statistics on the relative improvement of the stable bounds over the naive bounds. 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 44 percentage points, while this difference equals only about 24 percentage points for the stable bounds. Qualitatively similar results are obtained for the other quantiles reported in the Table 5. For example, in three quarters of our households the percentage point difference between the stable bounds is no larger than 30, compared to 52 percentage points for the naive bounds. In one case, we get a difference between the stable bounds that is as low as 1.71 percentage points, which comes close to point identification. Next, we find that the average improvement is close to 50 percentage points, and for three quarters of the households the bounds tightening amounts to at least 36 percentage points. In fact, for one household it equals no less than 77 percentage points.[25] Given all this, we may safely conclude that our methodology does exploit the marriage market implications in an effective way. It generates sharing rule bounds that are considerably tighter than

[24] 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. 10% of income) for all exit options. Such an alternative procedure yields results that are qualitatively similar to the ones that we report below.

[25] For 2 of our 62 households, we have that the stability bounds do not improve on the naive bounds, which explains the minimal improvement of 0 percentage points.

the naive ones. Notably, this conclusion holds for a fairly small sample (with only 62 households), with a single consumption observation per household, and without homogeneity of individual preferences.

We remark that, even though our stable bounds are considerably narrower than the naive bounds, they remain fairly wide in some cases. For example, the maximum difference between the relative lower and upper bounds still amounts to 56%. Different approaches can be used to further tighten the bounds, 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 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.[26]

Table 5: Identification of female relative income share

	Percentage points difference (upper minus lower = $\boldsymbol{\Delta}$)		**Relative improvement**
	Naive bounds $\boldsymbol{\Delta}$	**Stable bounds $\boldsymbol{\Delta}$**	$-\frac{\textbf{Stable }\Delta\textbf{ − Naive }\Delta}{\textbf{Naive }\Delta} \times 100$
Mean	44.41 p.p.	24.01 p.p.	46.45%
St. dev.	11.20 p.p.	10.28 p.p.	19.44%
Min.	7.37 p.p.	1.71 p.p.	0.00%
First quart.	38.03 p.p.	17.64 p.p.	35.96%
Median	44.96 p.p.	24.39 p.p.	47.92%
Third quart.	51.98 p.p.	29.66 p.p.	59.81%
Max.	70.94 p.p.	55.46 p.p.	76.77%

Note: p.p. stands for percentage points.

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 the intrahousehold wage ratio (i.e. female wage divided by male wage). We focus on this particular relationship because it received considerable attention in the literature on collective consumption models. That literature 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, and Oreffice, 2011). The underlying reasoning is that a higher wage improves the member's options outside marriage, which in turn yields a better bargaining position within marriage. From this perspective, it is interesting to see whether our analysis confirms the earlier findings, as it explicitly includes the structural implications of the outside options defined on the marriage market.

Figure 1 gives our results. We display the relation for the naive bounds (Panel A) as well as for the stable bounds (Panel B). Each • and + sign on 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 (i.e. nonparametric regressions) of these household-specific upper and lower bounds.

We find that the naive bounds are not really informative. They are generally wide. Actually, when looking at the upper and lower trendlines, we cannot exclude that (on average) the female income share is about 40% and independent of the wage ratio, i.e. there is no effect of the relative female wage on her relative income

[26] 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 1 to the estimated demands. See Cherchye, De Rock, Lewbel and Vermeulen (2013) for such an exercise in the context of collective consumption models. 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), Blundell, Kristensen and Matzkin (2014), Hoderlein and Stoye (2014) and Stoye and Kitamura (2013) address closely similar questions in a unitary context, also dealing with unobserved heterogeneity driving demand behavior.

share. By contrast, the stable bounds are substantially more informative. First, they are much narrower than the naive bounds, reflecting the results in Table 5. Next, and even more importantly, we now 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.

As a remark, one may be tempted to argue that this result is 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 our results for the naive bounds in Panel A of Figure 1. These bounds also exploit assignability of leisure but are not similarly upward sloping as the stable bounds in Panel B. Interestingly, our particular approach to sharing rule identification provides a structural interpretation of this difference between the naive bounds and the stable bounds, in terms of varying individual outside options defined by wage changes. In our opinion, this clearly demonstrates the empirical usefulness of endogenizing the marriage matching decisions in the household consumption analysis.[27]

More generally, we believe that the results in Table 5 and Figure 1 show the substantial potential of our framework to analyze the structural implications of the marriage market for household consumption patterns. It allows for an informative empirical analysis of intrahousehold decision processes, even for a fairly small sample of households and if we make minimal assumptions regarding the data at hand. As indicated above, a more powerful empirical analysis will result if we include a greater number of households and/or more (time-series) observations per household, or if we make stronger a priori assumptions (such as preference homogeneity). The next section will discuss alternative possible extensions of our framework.

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

Figure 1: Female income share (Y-axis) and the relative wage (female/male wage) (X-axis)

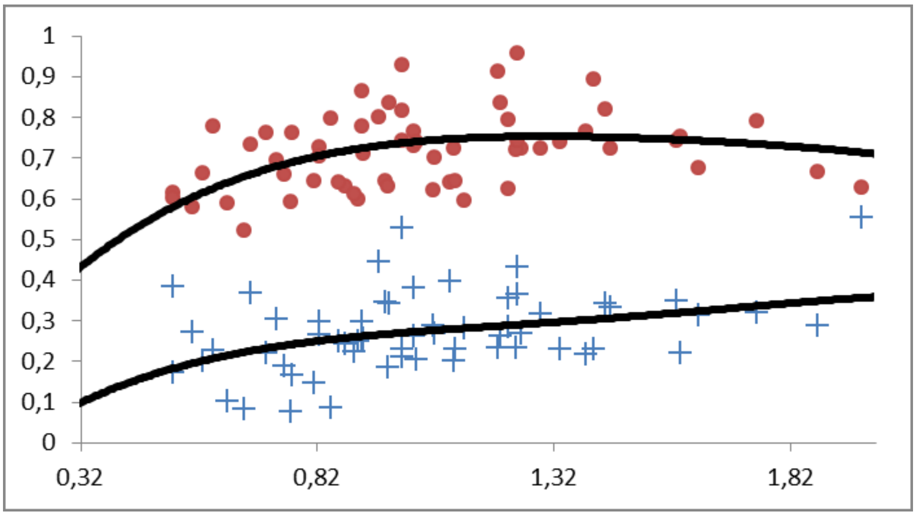

Panel A: Naive bounds

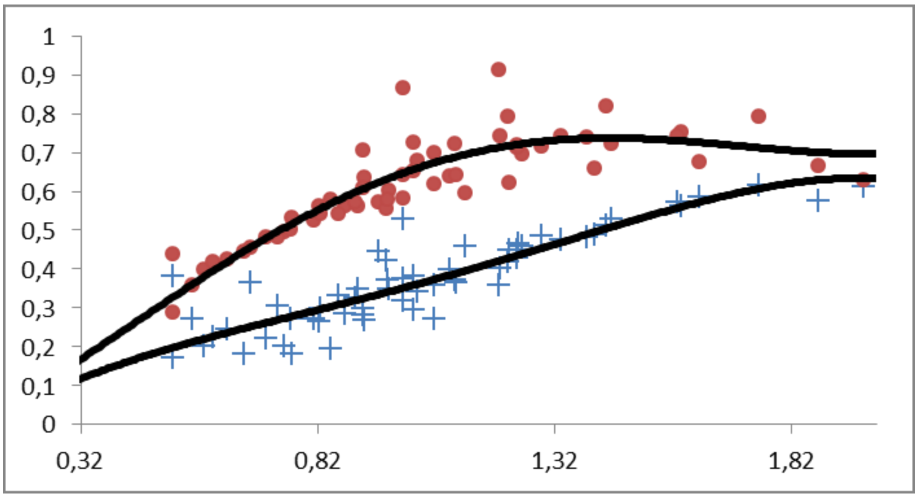

Panel B: Stable bounds

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 implications 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 (defined 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.[28]

Next, our empirical analysis has focused on the effect of relative wages on the sharing rule. We found that 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 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.[29]

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 (parametric) 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

[28] 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 identfication 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.

[29] 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.

cooperatively. However, it is sometimes argued that the assumption of Pareto efficiency is an overly strong one in a household context.[30] 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.

Appendix: proofs

Proof of Theorem 1

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:[31]

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

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

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

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

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

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

we get the conditions

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

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

[30] 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.

[31] Given the formal similarity between the Pareto efficiency criterion in Definition 1 and the no blocking pair condition in Definition 3, we can follow a directly analogous reasoning as under item 2. below to obtain the rationalizability conditions in Theorem 1 that pertain to our Pareto efficiency requirement (compare with Cherchye, De Rock and Vermeulen, 2011). For compactness, we do not include this reasoning here.

2. Let us then turn to the optimization problems for no blocking pairs. Here, a first optimization problem is defined as

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

i.e. $(\bar{q}^m_{m,w}, \bar{Q}_{m,w})$ represents the "dictatorial" allocation chosen by m if he could freely spend the entire income $y_{m,w}$. The corresponding first order conditions give

$$
\begin{aligned}
\frac{\partial v^m(\bar{q}^m_{m,w}, \bar{Q}_{m,w})}{\partial q^m} &\leq \overline{\delta}_{m,w} p_{m,w}, \\
\frac{\partial v^m(\bar{q}^m_{m,w}, \bar{Q}_{m,w})}{\partial Q} &\leq \overline{\delta}_{m,w} P_{m,w},
\end{aligned}
$$

where $\overline{\delta}_{m,w}$ is the Lagrange multiplier associated with the budget constraint.

Then, we consider the problem (for $\bar{V}_{m,w} \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)})$; see above)

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

In this case, we get the first order conditions

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

where $\mu_{m,w}$ and $\lambda_{m,w}$ are the Lagrange multipliers associated with the male's participation constraint and the budget constraint, respectively. The first four constraints represent the usual first order conditions for a Pareto optimal allocation. The last constraint requires the utility of the husband to equal exactly his utility under the matching σ, which corresponds to a situation where the constraint $v^m(q^m, Q) \geq v^m(q^m_{m,\sigma(m)}, Q_{m,\sigma(m)})$ is binding in the above optimization problem. In what follows, we use $\frac{\lambda_{m,w}}{\mu_{m,w}} = \delta_{m,w}$, $\frac{\partial u^w(q^w_{m,w}, Q_{m,w})}{\partial Q} = \lambda_{m,w} P^w_{m,w}$ 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 Theorem 1 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 Theorem 1 by using $v^m(q^m_{m,w}, Q_{m,w}) = V_{m,w}$, $u^w(q^w_{m,w}, Q_{m,w}) = U_{m,w}$ ($m \in M \cup \{\varnothing\}$ and $w \in W \cup \{\varnothing\}$) and $v^m(\bar{q}^m_{m,w}, \bar{Q}_{m,w}) = \overline{V}_{m,w}$.

Sufficiency. To obtain the sufficiency result, we consider

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

for

$$\begin{aligned} V^m(q^m, Q^m) &= \min_{w \in W \cup \{\varnothing\}} [V_{m,w} + \delta_{m,w}\left(p_{m,w}(q^m - q^m_{m,w}) + P^m_{m,w}(Q - Q_{m,w})\right)], \\ \overline{V}^m(q^m, Q^m) &= \overline{V}_{m,w} + \overline{\delta}_{m,w}\left(p_{m,w}(q^m - \overline{q}^m_{m,w}) + P^m_{m,w}(Q - \overline{Q}_{m,w})\right). \end{aligned}$$

Varian (1982) shows, in a unitary context, that $v^m(q^m_{m,w}, Q_{m,w}) = V_{m,w}$, $u^w(q^w_{m,w}, Q_{m,w}) = U_{m,w}$ ($m \in M \cup \{\varnothing\}$ and $w \in W \cup \{\varnothing\}$) and $v^m(\bar{q}^m_{m,w}, \bar{Q}_{m,w}) = \overline{V}_{m,w}$. 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 1

First, conditions (a) and (e) in Theorem 1 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 Theorem 1 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

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

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*, forthcoming.

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

[4] 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.

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

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

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

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

[9] Blundell, R., D. Kristensen, and R. Matzkin (2011), "Bounding quantile demand functions using revealed preference inequalities", *Journal of Econometrics*, forthcoming.

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

[11] 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.

[12] 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.

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

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

[15] 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.

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

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

[18] Cherchye, L., T. Demuynck, and B. De Rock (2014), "Is utility transferable? A revealed preference analysis", *Theoretical Economics*, forthcoming.

[19] Cherchye, L., B. De Rock, A. Lewbel, and F. Vermeulen, F. (2013), "Sharing rule identification for general collective consumption models", Working Paper, University of Leuven.

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

[21] 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.

[22] 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.

[23] 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.

[24] 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.

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

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

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

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

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

[30] 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.

[31] 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.

[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., and P. Reny (2006), "Matching to share risk", Working Paper, University of Chicago.

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

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

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

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

[38] 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.

[39] Dupuy, A., and A. Galichon (2012), "Personality traits and the marriage market", *Journal of Political Economy*, forthcoming.

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

[41] 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.

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

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

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

[45] Hoderlein, S., and J. Stoye (2014), "Revealed preferences in a heterogeneous population", *Review of Economics and Statistics*, forthcoming.

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

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

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

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

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

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

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

[53] 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.

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

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

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

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

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

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

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

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

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

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

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

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