Matching to share risk*

Pierre-André Chiappori[†], Phil Reny[‡]
Spring 2005

Abstract

We consider a matching model in which individuals belonging to two populations ('males' and 'females', each of which can be ranked by risk aversion in the Arrow-Pratt sense) can marry to share some exogenous income risk. The model is characterized by imperfectly transferable utility, a context in which few general results have been derived so far. We show that in this framework (i) a stable matching always exists, (ii) it is unique, and (iii) married couples exhibit a negative assortative matching pattern (for any two married couples, the more risk averse man is always matched with the less risk averse woman). We discuss the implications of these results for the empirical analysis of household behavior under uncertainty.

*Paper presented at conferences at the New York Federal Reserve and in Chicago and Tel Aviv. Comments from the participants are gratefully acknowledged.

[†]Department of Economics, University of Chicago and Department of Economics, Columbia University. Email: pc2167@columbia.edu

[‡]Department of Economics, University of Chicago

1 Introduction

Consider a model of one-to-one matching in which individuals within each population (respectively 'males' and 'females') can be ranked according to their risk aversion; i.e., for any two males (or for any two females), one is more risk averse than the other in the sense of Arrow-Pratt. Assume, moreover, that each individual faces some iid income risk, and that the main motivation for matching is the ability to share risk within the couple. What would be the qualitative features of stable matches in this context ? Should one expect similar people to match? Would an individual be attracted to a partner with 'opposite' characteristics ? Or may several equilibrium with different qualitative properties coexist? These are the issues we investigate in this paper.

Besides its own interest, we believe these questions are important in two respects. On a technical side, the model we consider involves matching under imperfectly transferable utility. Indeed, while (compensating) transfers within couples are possible (and indeed essential) in our context, the Transferable Utility property (TU) does not hold in general; the non linearities implicit in risk aversion generally imply that the cost for a spouse of one 'expected utile' transfered to his/her partner is not constant along the Pareto frontier. The non linearity of the Pareto frontier raises interesting problems that have rarely been addressed in the literature; to the best of our knowledge, our paper is one of the first to prove existence and uniqueness, and to qualitatively characterize the properties of the equilibrium match in a context of imperfectly transferable utility.[1]

A second, and probably more compelling aspect of our work is its potential consequences for the empirical analysis of risk sharing within groups, a topic that has attracted considerable attention during the last decades. Most empirical tests of efficient risk sharing within a group rely on a basic property of ex ante efficient agreements, the so-called 'mutuality principle', whereby individual consumption should only depend on aggregate shocks. This prediction can be tested in different ways. One, which follows Townsend's (1994) seminal contribution, is to regress individual consumptions on aggregate consumption (or income) at the group level plus various individual-specific variables (individual incomes, prox-

[1] Our paper is related to a recent contribution by Legros and Newman (20$$, hereby LN) which provides sufficient conditions for assortative matching in a context of imperfectly transferable utilities. Indeed, a consequence of our main result is that (a modified version of) the LN conditions are satisfied. While LN conditions are complex and difficult to check in general, our result thus shows that they hold for a large class of models, namely all models in which matching is on risk aversion and agents can be ranked according to this characteristic.

ies for idiosyncratic shocks, etc.). The prediction is falsified if, controling for aggregate indicators, individual specific shocks are found to matter. An alternative approach, initiated by Cochrane (1991), uses an Euler equation formulation, which relates the marginal utility of consumption at date t to the expected marginal utility of consumption at date $t+1$ conditional on information available at date t. The key property, here, is that the ratio of marginal utilities of consumption between periods t and $t+1$ should be independent of any idiosyncratic shock at date t. Using, for instance, Constant Relative Risk Aversion (CRRA) preferences, this implies that growth in individual log consumption be orthogonal to idiosyncratic shocks, a property that can readily be tested provided (even short) panel data on consumption are available.

A potential problem affecting these approaches is that while the theory has been developped for *individuals*, available data are in general gathered at the *household* level. The existing literature almost uniformly solves this problem by assuming it away, i.e. by postulating (implicitly in general) that a household can be represented by a unique utility function. To see why this assumption plays a crucial role in the design of the empirical strategy, let us relax it for a moment; i.e., let us assume that the household is in fact a group of (at least) two individuals endowed with different preferences. Assume furthermore that household make Pareto efficient decisions; i.e., their behavior can be described as stemming from the maximization of a weighted sum of individual utilities. In general, individual Pareto weights depend on any economic variables (e.g., the characteristics of individual wage or income distributions, the distribution of prices, etc.) that may influence the intrahousehold decision process.[2] In practice, thus, one can still consider the weighted sum as a 'household utility', but the latter will in general be income-dependent. This fact, by itself, invalidates the tests usually performed. For instance, the tests of Euler equation described above are biased: if marginal utility of income, as computed at the household level, depends on idiosyncratic variables such as income distributions through the Pareto weights, then growth in individual log consumption cannot be orthogonal to these variables even in the presence of perfectly efficient risk sharing mechanisms.

This leads to a natural but crucial question, namely: When is it possible to summarize a group's attitude toward risk through a unique utility

[2] Note that efficient risk sharing *within* the couple requires that individual Pareto weights be independent of the *realization* of idiosyncratic shocks. However, thay may (and general will) depend on their *distribution*; i.e., the weight of an individual within the intrahousehold decision process is probably related to, say, the stochastic process followed by the individual's income.

function that *does not depend on Pareto weights* ? The answer to this question is by now well known.[3] A necessary and sufficient condition for such an aggregation result to obtain is the so-called ISHARA (for Identical Shape Harmonic Absolute Risk Aversion) property.[4] Namely, individual utilites u_i must be such that the index of absolute risk aversion is some harmonic function of income x - say. $\frac{1}{a_i+b_i x}$;furthermore, the shape parameter b must be identical accross individuals ($b_i = b\ \forall i$). This family includes as particular cases Constant Absolute Risk Aversion utilities with arbitrary coefficients (then $b_i = 0\ \forall i$) and Constant Relative Risk Aversion utilities with identical coefficients (then $a_i = 0\ \forall i$). Hence whenever preferences are assumed to belong to the Constant Relative Risk Aversion family (a very popular assumption in the litterature), the unitary representation requires in fact that individual preferences be identical. If they fail to be, then household 'preferences' are income-dependent, and testing Euler equations then requires a totally different strategy.[5]

It is thus fair to say that the vast majority of empirical work devoted to risk sharing implicitly assumes that individuals within couples have similar preferences ('similar' meaning ISHARA in general and identical in the CRRA case) - what can be called the Identical Risk Aversion (IRA) assumption. While the main justification of this assumption is its convenience, its foundations remain shaky. Basically all empirical studies aimed at estimating the distribution of risk aversion find that it is widely heterogenous, hence that the IRA assumption cannot even be seen as a partially acceptable approximation, at least population-wide. A milder version of IRA, which only postulates that differences in risk aversion are perfectly captured by observables (i.e., that observationnaly identical agents can be assumed to have the same risk aversion, so that the IRA assumption holds conditional on observables), is similarly rejected by empirical evidence: existing studies on decision under uncertainty find that unobserved heterogeneity is paramount.[6] This begs the questions of why preferences should be expected to be identical within couples, while they are found to be widely heterogenous in the population as whole.

One (and perhaps the only) possible reconciliation between the theoretical models and the empirical findings may stem from matching consideration. After all, couples are endogenously formed. Even if risk

[3] See Wilson ($$$), and Mazzocco ($$$) for a complete proof.

[4] Interestingly, ISHARA is also necessary and sufficient for expected utility to be transferable within the couple; see Schuhlhofer-Wohl (2004).

[5] Progresses have recently been made in this direction, See Mazzocco (2005) for an analysis based on an 'individualistic' approach, in which Euler equations are estimated at the *individual* level..

[6] See Chiappori (2005) for a brief presentation.

aversion is widely heterogenous in general, the dispersion may be much smaller or even nil within households, especially if risk sharing is an important purpose of household formation. In other words, if agents with similar characteristics are likely to marry to share risk, the IRA assumption, although rejected for the global population, may still make sense at the household level. If, on the contrary, the logic of risk sharing leads to *negative assortative* matching, then the IRA assumption is even less justified within an endogenously formed couples than within the population as a whole, which casts serious doubts on the validity of the considerable empirical literature relying on it.

Interestingly, our conclusions are quite clear-cut, and go against the standard practice. Indeed, we show that in our model: (i) a stable matching always exist, (ii) it is unique and (iii) it is negative assortative, in the sense that more risk averse males tend to be matched with less risk averse females and conversely. Our results thus imply that in a context in which matching is endogenous, risk aversions tend to systematically differ between matched agents - a conclusion that sounds particularly damaging for the identical risk aversion assumption.

2 The Model

The framework We consider a marriage model in which a set of n_M males are to be individually matched with a set of n_F females. Males and females are strictly risk averse; u_i denotes the VNM utility of male i, v_k that of female k. Both utilities are strictly increasing, strictly concave, and satisfy $u_i'(0) = +\infty$ ($v_k'(0) = +\infty$); we normalize $u_i(0)$ and $v_k(0)$ to be 0.

We assume that agents in each population can be ranked by risk aversion in the Arrow-Pratt sense: for two males (resp. females) i and j,either i is more risk averse than j or j is more risk averse than i. We can therefore rank each set of individuals by order of increasing risk aversion (so that Mr. i is more risk averse that Mr. j if and only if $i > j$, and similarly for women). It follows that for any two indices i and j, the strictly increasing functions f_{ij} and g_{ij} defined by $u_i = f_{ij}(u_j)$ and $v_i = g_{ij}(v_j)$ then f_{ij} and g_{ij} are strictly concave if and only if $i > j$.[7]

Each male i is endowed with some exogenous, random income $\tilde{y}_i$; male incomes are assumed iid. Similarly, each female j is endowed with some exogenous, random income $\tilde{z}_j$; female incomes are assumed iid and

[7]It should be noted that the property we impose is, in a sense, fairly weak: we only assume that if i and j receive the *same* income, then i's index of absolute risk aversion is larger than j's. Since, at a stable match, the (random) income received by an agent is endogenous, it may be that j's *actual* risk aversion is in fact larger than i's for some (or even most) income realizations.

independent from male incomes.

Agents match to form couples. Each couple (i, j) receives the (random) total income $\tilde{y} = \tilde{y}_i + \tilde{z}_j$; note that total income is identically distributed across couples. Consumption is one-dimensional; once married, the agents within a couple share risk in a Pareto efficient way. Efficiency, in turn, require that the allocation of consumption within the couple satisfies the mutuality principle (which states that each agent's consumption should depend only on total income $\tilde{y} = \tilde{y}_i + \tilde{z}_j$); hence, from now on, all consumptions are considered as functions of $\tilde{y}$. Note that in this context agents will always prefer to match than to remain single.[8] It follow that at any stable match, there may be single males (if $n_M > n_F$) or single females (if $n_M < n_F$) but not both.

When male i is matched with female k, his consumption (as a function of the income realization y) is denoted x_k^i; his spouse consumes $y - x_k^i$. The efficiency assumption implies that there exists a $\bar{v}_k$ such that x_k^i solves

$$\max_{x_k^i(y)} E\left[u_i\left(x_k^i(y)\right)\right] \quad \text{s.t.} \quad E\left[v_k\left(y - x_k^i(y)\right)\right] \geq \bar{v}_k$$

We then let $x_k^i(y, \bar{v}_k)$ denote the solution to this maximization problem.

3 Existence: a general argument

As a first step, the existence of a stable match is established by the following result:

Proposition 1 *There exist at least one stable match*

Proof. *A standard proof would rely on a generalization of the Gale-Shapley algorithm, applied to a discretized version of the model (i.e., imposing that individual expected utilities can only take finite a finite number of values). We propose here an alternative argument that exploits the specific nature of the problem.*

We may, without loss of generality, assume that $n_M \geq n_F$ *(i.e., there are no less males than females). Let* $V_{ij}(u)$ *denote male* i*'s expected payoff from being matched with female* j*, where* j *receives an expected payoff of at least* u*. Let* r_j *denote female* j*'s reservation utility and let* $\bar{u}_j$ *denote female* j*'s expected utility when she is matched and receives the entire joint income.*

Consider the following artificial game between the n_M *males,* $i = 1, ..., n_M$*, and an auctioneer, player* $i = 0$*. The auctioneer chooses a*

[8] Indeed, a possible way of sharing income is to mimic singlehood (i.e., Mr. i keeps $\tilde{y}_i$ and Mrs. j keeps $\tilde{z}_j$). Such an allocation will however never be efficient, since it fails to satisfy the mutuality principle; hence there exists allocations which are preferred over singlehood by both agents..

utility, u_j *for each female* $j = 1, ..., n_F \leq n_M$*. Male* i *chooses a probability* $x_{ij} \in [0,1]$ *of matching with each female* j*, where* $\sum_j x_{ij} \leq 1$ *permits male* i *to remain unmatched with positive probability. Let* $x_i = (x_{i1}, ..., x_{in_F})$ *denote male* i*'s strategy.*

Given the auctioneer's strategy $u = (u_1, ..., u_{n_F})$ *and the vector of male strategies* $x = (x_1, ..., x_{n_M})$*, male* i*'s payoff is*

$$\pi_i(x, u) = \sum_{j=1}^{n_M} x_{ij} V_{ij}(u_j),$$

and the auctioneer's payoff is

$$\pi_0(x, u) = \sum_{j=1}^{n_F} u_j \left[\left(\sum_{i=1}^{n_M} x_{ij} \right) - 1 \right].$$

Because each player's payoff is linear in his strategy, this game posseses a pure strategy equilibrium (x^*, u^*)*.*

We first claim that $\sum_i x^*_{ij} \leq 1$ *for each* $j = 1, ..., n$*. If not, then* $\sum_i x^*_{ij'} > 1$ *for some* j'*. Consequently, optimal play by the auctioneer requires* $u^*_{j'} = \bar{u}_{j'}$*. But because every male strictly prefers being unmatched than matching with* j' *at* $\bar{u}_{j'}$*, we must have* $\sum_i x^*_{ij'} = 0$*. This contradiction proves the claim.*

Next, we claim that $\sum_i x^*_{ij} = 1$ *for each* $j = 1, ..., n$*. If not, then* $\sum_i x^*_{ij'} < 1$ *for some* j'*. Consequently, optimal play by the auctioneer requires* $u^*_{j'} = r_{j'}$*. But because every male strictly prefers matching with* j' *at her reservation value to being unmatched, it must then be the case that* $\sum_j x^*_{ij} = 1$ *for every male* i*. But this implies* $\sum_i \sum_j x^*_{ij} = n_M \geq n$*, which implies that* $\sum_i x^*_{ij} = 1$ *for every* j*. This contradiction establishes the claim.*

Hence, letting X^* *denote the* $n_M \times n_F$ *matrix whose* i*th row is* x^*_i*, the columns add to unity and the rows add to no more than unity. Appending to* X^* $n_M - n_F$ *identical columns with* i*th entry* $(1 - \sum_j x^*_{ij})/(n_M - n_F)$ *results in an* $n_M \times n_M$ *matrix that is doubly stochastic, and hence a convex combination of permutation matrices. Any permutation matrix given positive weight in the convex combination is a stable match, where any male "matched" with an appended column is interpreted as unmatched.* ■

We now consider the most interesting part of the problem, namely the characterization of the stable matchs.

4 The 2 by 2 case

We start with the 2x2 case; we thus assume that there are exactly two males with VNM utilities u_1, u_2 and two females with VNM utilities are v_1, v_2, where, u_2 (resp. v_2) is strictly more concave than u_1 (resp. v_1). The key result is the following

Proposition 2 *In the 2x2 case, the negative-assortative match* $(1,4),(3,2)$ *is the unique stable match.*

Proof. *The Proof relies on the following Lemma:*

Lemma 3 *Assume that* $\bar{v}_1 > 0$ *and* $\bar{v}_2 > 0$ *are such that*

$$E\left[u_1\left(x_1^1(y,\bar{v}_1)\right)\right] \geq E\left[u_1\left(x_2^1(y,\bar{v}_2)\right)\right] > 0. \tag{1}$$

Then

$$E\left[u_2\left(x_1^2(y,\bar{v}_1)\right)\right] > E\left[u_2\left(x_2^2(y,\bar{v}_2)\right)\right] > 0 \tag{2}$$

The proof of the Lemma 3 is in Appendix. In words, Lemma 3 says the following. Suppose that neither side of the market receives the entire share of joint income. If $\bar{v}_1$ *and* $\bar{v}_2$ *are such that 1 weakly prefers being matched with 2 than with 4, then keeping* $\bar{v}_1$ *and* $\bar{v}_2$ *unchanged, 3 strictly prefers being matched with 2 than with 4. Or, equivalently, if one considers the two possible matchings,* $(1,2),(3,4)$ *on the one hand and* $(1,4),(3,2)$ *on the other, and if the surpluses are shared in such a way that both 2 and 4 are indifferent between the two matches while 1 weakly prefers the first, then 3 strictly prefers the second. Note that this property is purely ordinal; its statement does not require any particular cardinal representation of preferences.*

We now proceed to show that in the 2x2 case, when income shares are not extreme there is at most one stable matching, namely the negative-assortative one $(1,4),(3,2)$*. Indeed, Lemma 3 implies that the alternative matching* $(1,2),(3,4)$ *cannot be stable. For assume it is, and let* $\bar{v}_k$ *denote the utility of member k* ($k = 1,2$) *in this matching. Stability implies that* 1 *could not get a higher utility with 4 than what he gets with 2 without leaving 4 with less than* $\bar{v}_2$*; hence relation (1) must hold. By Lemma 3, we must have relation (2), which states that a match between 2 and 3 would be preferred (weekly) by 2 and (strongly) by 3 to their respective current matches, a contradiction.*

Finally, the previous result establishes the existence of (at least) one stable match, hence the conclusion. ∎

5 The general, finite case

5.1 Arbitrary, identical number of mates

We now consider the general case. Assume, first, that the populations of males and females have the *same* arbitrary size n. The main result is the following.

Proposition 4 *The only stable matching is the negative-assortative one, in which the k-th most risk averse man is matched with the k-th least risk averse woman for all k.*

Proof. *We know that there exist at least one stable match. Let us now show that any match different from the negative-assortative one cannot be stable. For any such match, define $\bar{k}$ as the smallest number k such that the k-th less risk averse male is* not *matched with the k-th more risk averse female (i.e., female of index $n_F - \bar{k}$). By definition of $\bar{k}$, male $\bar{k}$ must be matched with a female (say f) who is* less *risk averse than female $n_F - \bar{k}$, and female $n_F - \bar{k}$ must be matched with a male (say m) who is* more *risk averse than male $\bar{k}$. From Proposition 2, the submatching $\left(\bar{k}, f\right), \left(m, n_F - \bar{k}\right)$ cannot be stable: the matching $\left(\bar{k}, n_F - \bar{k}\right), (m, f)$ must be preferred either by $\bar{k}$ and $n_F - \bar{k}$ or by m and f. Hence the global matching cannot be stable.* ■

5.2 The general case

The main result Finally, we relax the assumption that the number of males and females are equal. As above, we can, without loss of generality, assume that the number of males, n_M, is larger than the number of females n_F . Then:

Proposition 5 *Any stable matching has the following features:*

- *all females are married, while $n_M - n_F$ males are single*
- *the matching is negative-assortative, in the sense that among married men, the k-th most risk averse is matched with the k-th least risk averse woman for all k.*

Proof. *Assume that a female is not married. Since there exist unmarried men, matching one of them with the single female increases both welfare, a contradiction. The second statement follows from Proposition 4.* ■

An ambiguity remains, however, regarding the identity of the men who are left single. Are they the most risk averse individuals in the population? The least risk averse? Or could they be located somewhere within the distribution of risk aversion? We now proceed to show, on a simple example, that any of these situations is indeed possible, depending on the parameters of the model.

An example The example we consider has the following features:

- Both $\tilde{y}_i$ and $\tilde{z}_j$ are uniformly distributed over $[0,1]$; it follows that their sum $\tilde{y}$ has the following distribution

$$P(\tilde{y} < a) = \int_0^1 (a-t)\,dt = \frac{1}{2}a^2 \text{ if } a \leq 1$$
$$P(\tilde{y} < a) = \int_0^{a-1} dt + \int_{a-1}^1 (a-t)\,dt = 2a - 1 - \frac{1}{2}a^2 \text{ if } 1 < a \leq 2$$
$$P(\tilde{y} < a) = 1 \text{ if } a > 2$$

- Individual preferences are CARA. In particular, for a match between a man and a woman with respective RA indices of σ and μ, the set of Pareto efficient contracts is characterized by

$$x(\tilde{y}) = \frac{\mu\tilde{y}}{\sigma+\mu} + k$$
$$\tilde{y} - x(\tilde{y}) = \frac{\sigma\tilde{y}}{\sigma+\mu} - k$$

  where k is an arbitrary constant.

We first compute the maximum level of expected utility a women with RA μ may expect when matched with a man with RA σ. This corresponds to an efficient contract in which the constant k is such that he is indifferent between being matched with μ or being single. After straightforward calaculations, this gives

$$\exp(-\sigma k) = \frac{(1-e^{-\sigma})\,\sigma\mu^2}{(\sigma+\mu)^2\left(1-e^{-\sigma\frac{\mu}{\sigma+\mu}}\right)2}$$

Let $\bar{U}(\sigma,\mu)$ denote the utility she derives from the match in that case (i.e. when she receives all the surplus). Then one can show that

$$\bar{U}(\sigma,\mu) = -\left(1-e^{-\sigma}\right)^{-\frac{\mu}{\sigma}} \sigma^{-\frac{\mu+2\sigma}{\sigma}} \mu^{-2\frac{\sigma+\mu}{\sigma}} (\sigma+\mu)^{2\frac{\sigma+\mu}{\sigma}} \left(1-e^{-\sigma\frac{\mu}{\sigma+\mu}}\right)^{2\frac{\sigma+\mu}{\sigma}}$$

A graph of this function is provided in Figure 1.

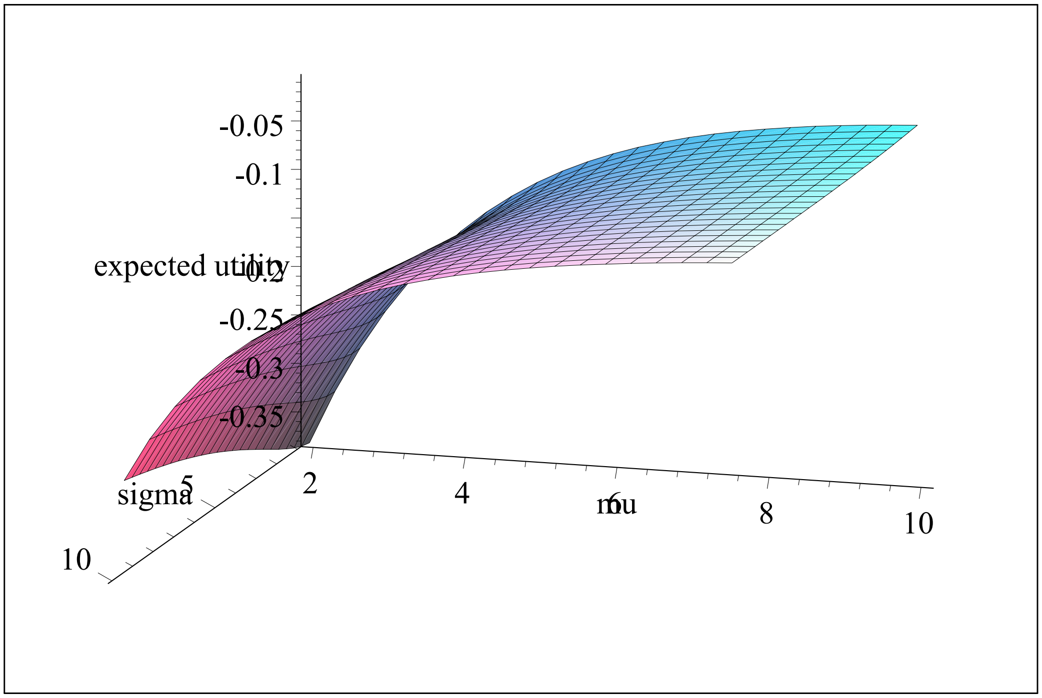

Figure 1 - Female maximum utility as a function of RA coefficients

In particular, Figure 2 provides a section of the corresponding surface by the plane $\mu = 5$:

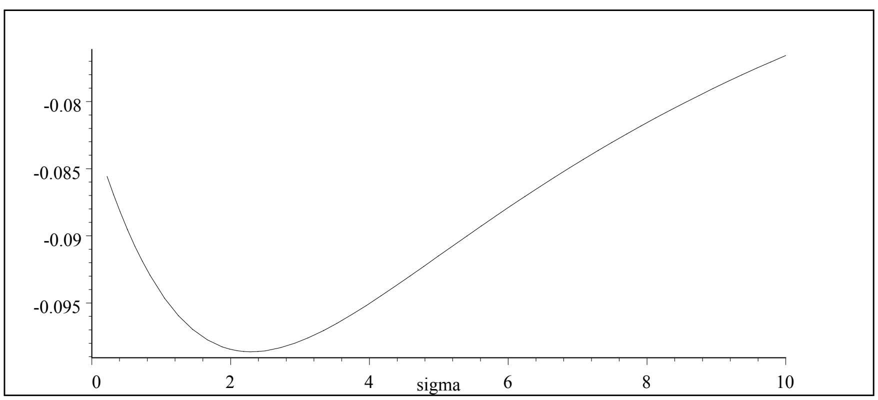

Figure 2: section for $\mu = 5$

The key remark, here, is that this section is non monotonic (and actually U-shaped). In words, when her RA parameter is 5, then her maximum utility $\bar{U}$, as a function of her mate's risk aversion, decreases then increases.

Consider, now, a situation in which three males (with respective RA σ_1, σ_2 and σ_3) are to be matched with two women with identical RA μ. We use the following Lemma:

Lemma 6 *Male i remains single if and only if*

$$\bar{U}(\sigma_i, \mu) = \min_j \bar{U}(\sigma_j, \mu)$$

Proof. *Assume not; then, i is single whereas $\bar{U}(\sigma_j, \mu) < \bar{U}(\sigma_1, \mu)$ for some j. But j's spouse cannot receive more than $\bar{U}(\sigma_j, \mu)$ in her match; a match with i in which she receives $\frac{\bar{U}(\sigma_j,\mu)+\bar{U}(\sigma_1,\mu)}{2}$ would be prefered by her and by i, which proves that the match cannot be stable.* ■

In other words, pick up any three values of σ on the horizontal axis of Figure 2; the coefficient giving the smallest value to the function defines the single male. Obviously, all outcomes are possible. For $\sigma_1 = 1, \sigma_2 = 1.5, \sigma_3 = 2$, the most risk averse agent 3 is single; for $\sigma_1 = 2, \sigma_2 = 4, \sigma_3 = 6$, the least risk averse agent 1 is single; finally, for $\sigma_1 = 1, \sigma_2 = 2, \sigma_3 = 4$, the intermediate agent 2 is single.

This result can be interpreted as follows. Think of a female trying to select a mate within a sample of single males. Two conflicting considerations should be taken into account. First, everything equal, a less risk averse partner would absord a larger share of the income risk, which for a risk averse aversion is a boon. However, since all males face the same income process when single, a less risk averse partner will have a higher reservation utility, which restricts the maximum utility his partner can receive. Which effect dominates depends on the values of the parameters. Figure 2 suggests that a risk averse woman will tend to prefer partners who are either almost risk neutral (so that they can provide almost perfect insurance) or very risk averse (so that she can extract a large rent), while her welfare is smaller for intermediate mates.

6 Conclusion

Clearly, the theoretical nature of our investigation has lead us to considering a very simple framework. In 'real life', heterogeneity is multidimensional; in particular, matching is not exclusively based on risk aversion, and other aspects (starting with incomes) also play a key role. An interesting extension would be one in which income processes are chosen by the agents, hence reflect their risk aversion. This is left for future work.

We believe, however, that our simple model has interesting consequences. The basic conclusion stemming from our work is that the logic of risk sharing matching is negative assortative. Very risk averse individuals are eager to match with less risk averse partners, who can provide the coverage they need at low cost; conversely, almost risk neutral agents exploit their comparative advantage by being matched with very risk averse spouses, who are willing to give up a large risk premium in exchange for coverage. All in all, to the extent that risk sharing plays a role in marital decisions, one should expect intrahousehold differences in risk aversion to be large - a conclusion that fits empirical evidence pretty well.[9]

This finding, in turn, has important consequences for the empirical literature on risk sharing. First, as discussed in introduction, it casts strong doubts on the 'unitary' approach to household behavior, whereby some aggregation property is assumed to hold within couples, resulting in household behavior being modeled as the maximization of a unique VNM utility. As it turns out, theory gives very precise conditions that are needed for such an aggregation result to hold true; namely, ISHARA is a necessary and sufficient requisite. If agents are matched randomly from an heterogeneous distibution (or, equivalently, if the matching criterion is orthogonal to risk sharing issues), ISHARA is very unlikely to hold. Our result shows that considering matching as endogenous and motivated by risk sharing does not help; actually, ISHARA is even less likely to hold true in this context. For instance, if agents are characterized by constant relative risk aversion, preferences can only be identical for *one* couple at most, and may be quite different for many or most of them. All in all, our conclusions provide a strong theoretical support for the alternative approach developped by Mazzocco ().

Secondly, the scope of the underlying intuition (that risk sharing motives tend to produce very heterogeneous groups) is probably much larger. A majority of empirical papers devoted to the empirical analysis of risk sharing assume not only that each household can be represented by a single VNM utility, but also that these utilities are identical across households. A major avantage of this assumption is that, for specific but widely used functional forms, it leads to very simple tests. Assume, indeed, in line with most of the literature, that agents have constant relative risk aversion (CRRA) utilities, and that the coefficient of risk aversion is the same over the entire population. This assumption leads to a remarkably tractable form; namely, individual consumptions are are linear functions of aggregate resources:

$$x_i^t(y) = \lambda_i y_t$$

where y_t denotes the group's aggregate income (or consumption) at date t, x_i^t denotes member i's consumption, and λ_i is a parameter depending on the Pareto weight of member i (note that assuming efficient risk sharing is *equivalent*, in this context, to assuming that λ_i is time-invariant).

[9] For instance, surveys

This specific form, whereby x_i^t obtains as the product of some observable, aggregate index y_t and an unobservable, time-invariant, individual-specific parameter λ_i, can easily be tested. For instance, one may, following Udry and Duflo (200$), take the log of both sides and differentiate out the time-invariant component, leading to:

$$\Delta \log x_i (y) = \Delta \log y$$

where the operator Δ denotes a one period difference. This prediction can readily tested from available data.

Models of efficient risk sharing which assume identical individual utilities include Altonji et al. (1992, 1996, 1997), Altug and Miller (19$$), $$$ and others.[10] Again, our results cast strong doubts on the validity of this approach. If groups are created randomly (or if the criterion implicit in the formation of the groups is orthogonal to risk sharing), there is no reason to expect that a widely heterogeneous population will produce homodenous groups. Assume, on the contrary, that risk sharing is an important motivation for the formation of these groups. Even though, technically, our results do not extend to the formation of larger groups,[11] they strongly suggest that homogenous groups are unlikely to be stable. Again, these arguments point to the need of more general models - a direction in which new results have recently been derived.[12]

[10] Interestingly enough, two of the earlier (and seminal) papers of the literature are an exception: while Townsend (1994) assumes CARA preferences and Cochrane (1991) considers CRRA utilities, both allow for risk aversions to vary across households. Note, however, that Townsend's preferences are still ISHARA, while in Cochrane's model the relationship between individual and aggregate consumptions admits no closed form solution.

[11] The literature has emphasized, in particular, that stable match may fail to exist in 'many to one homosexual' matching models of this type (see Roth and Sottomayor $$$, ch. $).

[12] See Chiappori, Schuhlhofer-Wohl and Townsend (2005).

APPENDIX

A Proof of Lemma 3

The following well-known fact will be helpful.

FACT: If a differentiable real-valued function defined on an interval is nonnegative at x_0 and its derivative is positive whenever the function vanishes, the function is positive at all $x > x_0$.[13]

Lemma 7 *Assume that* $\bar{v}_1 > 0$ *and* $\bar{v}_2 > 0$ *are such that*

$$E\left[u_1\left(x_1^1\left(y,\bar{v}_1\right)\right)\right] \geq E\left[u_1\left(x_2^1\left(y,\bar{v}_2\right)\right)\right] > 0.$$

Then

$$E\left[u_2\left(x_1^1\left(y,\bar{v}_1\right)\right)\right] > E\left[u_2\left(x_2^1\left(y,\bar{v}_2\right)\right)\right] > 0.$$

Proof of Lemma 7. Consider the difference $x_1^1\left(y,\bar{v}_1\right) - x_2^1\left(y,\bar{v}_2\right)$ as a function of y. A first claim is the following:

Claim 8 *There exists* $\bar{y} \geq 0$, *possibly infinite, such that the function* $x_1^1\left(y,\bar{v}_1\right) - x_2^1\left(y,\bar{v}_2\right)$ *is positive for* $0 < y < \bar{y}$ *and negative for* $y > \bar{y}$.

Proof of Claim 8. If $x_1^1\left(y,\bar{v}_1\right) - x_2^1\left(y,\bar{v}_2\right) > 0$ for all $y > 0$, then $\bar{y} = +\infty$ and we are done. So, assume that $x_1^1\left(\bar{y},\bar{v}_1\right) = x_2^1\left(\bar{y},\bar{v}_2\right) = \bar{x}$ for some $\bar{y} > 0$. Because $\bar{u}_k > 0$ and $E\left[u_1\left(x_k^1\left(y,\bar{u}_k\right)\right)\right] > 0$ for $k = 2, 4$, and because $u_k'(0) = +\infty$, for $k = 1, 2, 4$, we know that $0 < \bar{x} < \bar{y}$. Consequently, a straightforward and well known result by Wilson states that at $\bar{y}$,

$$\frac{dx_1^1\left(\bar{y},\bar{v}_1\right)}{dy} = \frac{\tau^1\left(\bar{x}\right)}{\tau^1\left(\bar{x}\right) + \theta^1\left(\bar{y}-\bar{x}\right)} \in (0,1)$$

and similarly

$$\frac{dx_2^1\left(\bar{y},\bar{v}_2\right)}{dy} = \frac{\tau^1\left(\bar{x}\right)}{\tau^1\left(\bar{x}\right) + \theta^2\left(\bar{y}-\bar{x}\right)} \in (0,1)$$

[13] A proof is as follows. If $N = \{x > x_0 : f(x) \leq 0\}$ is nonempty, it contains a smallest member, $\bar{x} > x_0$; otherwise $f(x_0) = 0$ and $f'(x_0) \leq 0$, a contradiciton. Consequently, $f(\bar{x}) = 0$ and f assumes a minimum at $\bar{x}$ on the interval $[x_0, \bar{x}]$, implying that $f'(\bar{x}) \leq 0$, a contradiction. Hence, N is empty.

where τ^i (resp. θ^i) denotes the risk tolerance associated with u_i (resp. v_i). Since $\theta^2(\bar{y}-\bar{x}) < \theta^1(\bar{y}-\bar{x})$ by assumption, we must have $\frac{dx_2^1(\bar{y})}{dy} > \frac{dx_1^1(\bar{y})}{dy}$. The uniqueness of such a $\bar{y}$ as well as the sign of $x_1^1(y,\bar{v}_1) - x_2^1(y,\bar{v}_2)$ follow from the FACT above. Finally, $E[u_1(x_2^1(y,\bar{v}_2))] > 0$ implies that $x_2^1(y,\bar{v}_2) > 0$ for a positive measure of y's (in fact for all y because $u_1'(0) = +\infty$) so that $E[u_2(x_2^1(y,\bar{v}_2))] > 0$.■

By Claim 8, the function $\phi(y) = u_1(x_1^1(y,\bar{v}_1)) - u_1(x_2^1(y,\bar{v}_2))$ is positive for $0 < y < \bar{y}$ and negative for $y > \bar{y}$. Also, define f by $u_2 = f(u_1)$; then f is increasing and concave. Because $x_1^1(y,\bar{v}_1)$ is strictly increasing in y and u_1 is strictly increasing, the function $g(y) = f'[u_1(x_1^1(y,\bar{v}_1))]$ is positive and strictly decreasing in y. Hence, if income is distributed according to the density $h(\cdot)$,

$$\begin{aligned} E[g(y)\phi(y)] &= \int_0^{\bar{y}} g(y)\phi(y)h(y)dy + \int_{\bar{y}}^{\infty} g(y)\phi(y)h(y)dy \\ &> \int_0^{\bar{y}} g(\bar{y})\phi(y)h(y)dy + \int_{\bar{y}}^{\infty} g(\bar{y})\phi(y)h(y)dy \\ &= g(\bar{y})E\left[u_1\left(x_1^1(y,\bar{v}_1)\right) - u_1 x_2^1(y,\bar{v}_2)\right] \geq 0, \end{aligned}$$

where the final inequality follows because the expectation is, by hypothesis, nonnegative. Hence, employing the definition of ϕ in the left-hand side above yields

$$E\left[g(y)\left(u_1\left(x_1^1(y,\bar{v}_1)\right) - u_1\left(x_2^1(y,\bar{v}_2)\right)\right)\right] > 0.$$

Finally, the concavity of f implies that

$$f\left[u_1\left(x_1^1(y,\bar{v}_1)\right)\right] - f\left[u_1\left(x_2^1(y,\bar{v}_2)\right)\right] \geq \left(u_1\left(x_1^1(y,\bar{v}_1)\right) - u_1\left(x_2^1(y,\bar{v}_2)\right)\right) f'\left[u_1\left(x_1^1(y,\bar{v}_1)\right)\right]$$

Since $u_2 = f_{31} \circ u_1$, we conclude that

$$E\left[u_2\left(x_1^1(y,\bar{v}_1)\right) - u_2\left(x_2^1(y,\bar{v}_2)\right)\right] > 0$$

which completes the proof of Lemma 7. ■

Lemma 7 states that if $\bar{v}_1$ and $\bar{v}_2$ are such that 1 prefers being matched with 2 than with 4, then 3 strictly prefers **1's** corresponding efficient allocation with 2 to **1's** corresponding efficient allocation with 4. Note that, so far, nothing is said about 3's efficient allocations.

We now consider 3's efficient allocations. Define, for $0 \leq \lambda \leq 1$,

$$\pi_k(\lambda) = \max_{x(y)} E\left[(1-\lambda)u_1(x(y)) + \lambda u_2(x(y))\right] \quad \text{s.t. } E\left[v_k(y - x(y))\right] \geq \bar{u}_k.$$

A key property of the value functions, $\pi_1(\lambda)$ and $\pi_2(\lambda)$, is the following.

Claim 9 *If* $\bar{\lambda}$ *is such that* $\pi_1\left(\bar{\lambda}\right) = \pi_2\left(\bar{\lambda}\right)$ *then* $\frac{d\pi_1\left(\bar{\lambda}\right)}{d\lambda} > \frac{d\pi_2\left(\bar{\lambda}\right)}{d\lambda}$.

Proof of Claim 9. Assume, first, that $\bar{\lambda} = 0$. Then $\pi_k(0) = E\left[u_1\left(x_k^1(y, \bar{u}_k)\right)\right]$ and, by the envelope theorem, $\frac{d\pi_k(0)}{d\lambda} = E\left[u_2\left(x_k^1(y, \bar{u}_k)\right)\right] - E\left[u_1\left(x_k^1(y, \bar{u}_k)\right)\right]$. Hence, the desired conclusion is a direct application of Lemma 7.

Hence, if $\pi_1(0) = \pi_2(0)$ then $\frac{d\pi_1(0)}{d\lambda} > \frac{d\pi_2(0)}{d\lambda}$. Of course, this conclusion holds for all utility functions u_k, $k = 1, ..., 4$ satisfying our hypotheses. Consequently, if instead $\bar{\lambda} > 0$, and we define $\tilde{u}_1 = \left(1 - \bar{\lambda}\right) u_1 + \bar{\lambda} u_2$, and define for $k = 1, 2$ the value function $\tilde{\pi}_k(\lambda)$ as before, but replacing the utility function u_1 with $\tilde{u}_1$, then, because $\tilde{u}_1$ is more convex than u_2, we may similarly conclude that if $\pi_1(0) = \tilde{\pi}_2(0)$ then $\frac{d\pi_1(0)}{d\lambda} > \frac{d\tilde{\pi}_2(0)}{d\lambda}$. But this is equivalent to the statement that if $\pi_1\left(\bar{\lambda}\right) = \pi_2\left(\bar{\lambda}\right)$ then $\frac{d\pi_1\left(\bar{\lambda}\right)}{d\lambda} > \frac{d\pi_2\left(\bar{\lambda}\right)}{d\lambda}$, as desired. ■

In view of the FACT above, a direct consequence of Claim 9 is that $\pi_1(0) \geq \pi_2(0)$ implies $\pi_1(\lambda) > \pi_2(\lambda)$ for all $\lambda \in (0, 1]$, and in particular that $\pi_1(1) > \pi_2(1)$. Hence, altogether we have shown that if $E\left[u_1\left(x_1^1(y, \bar{v}_1)\right)\right] \geq E\left[u_1\left(x_2^1(y, \bar{v}_2)\right)\right] > 0$, then $E\left[u_2\left(x_1^2(y, \bar{v}_1)\right)\right] > E\left[u_2\left(x_2^2(y, \bar{v}_2)\right)\right]$.

Finally, observe that $E\left[u_1\left(x_2^1(y, \bar{v}_2)\right)\right] > 0$ implies $E\left[u_2\left(x_2^1(y, \bar{v}_2)\right)\right] > 0$. Hence, $E\left[u_2\left(x_2^2(y, \bar{v}_2)\right)\right] \geq E\left[u_2\left(x_2^1(y, \bar{v}_2)\right)\right] > 0$, which completes the proof of Lemma 3. ■