Matching with Trade-offs: Revealed Preferences over Competing Characteristics

Alfred Galichon[1] Bernard Salanié [2]

First version dated December 6, 2008. The present version is of February 22, 2009[3].

[1]Economics Department, École polytechnique; e-mail: alfred.galichon@polytechnique.edu

[2]Department of Economics, Columbia University; e-mail: bsalanie@columbia.edu.

[3]The authors are grateful to Guillaume Carlier, Pierre-André Chiappori, Ivar Ekeland and to seminar participants at Ecole polytechnique and séminaire Roy for useful comments and discussions.

Abstract

We investigate in this paper the structure of optimal matchings with competing criteria. The surplus from a marriage match, for instance, may depend both on the incomes and on the educations of the partners. Even if the surplus is complementary in incomes, and complementary in educations, imperfect correlation between income and education at the individual level implies that the social optimum must trade off matching on incomes and matching on educations. We characterize, under mild assumptions, the properties of the set of feasible matches, and of the socially optimal match. Then we show how data on correlations of the types of the partners in observed matches can be used to test that the observed matches are socially optimal. Under optimality, our procedure also provides an estimator of the parameters that define social preferences over matches, thereby introducing the Matching Revealed Preferences (MaRP) estimator. We illustrate our approach on data from the June 1995 CPS.

JEL codes: C78, D61, C13.

Introduction

> Louisa Burton was naturally ill-tempered and cunning; but (...) with such an extraordinary share of personal beauty, joined to a gentleness of manners, and an engaging address, she might stand a good chance of pleasing some young Man who might afford to marry a Girl without a Shilling.
>
> Jane Austen, *Lesley Castle* (1792).

Starting with Becker (1973), most of the economic theory of one-to-one matching has focused on the case when the surplus created by a match is a function of just two numbers: the one-dimensional types of the two partners. As is well-known, when the types of the partners are one-dimensional and are complementary in producing surplus, then the socially optimal matches exhibit positive assortative matching. Moreover, the resulting configuration is stable: it is in the core of the matching game; and it can be implemented by the celebrated Gale and Shapley (1962) deferred acceptance algorithm.

While this result is both simple and powerful, its implications are also quite unrealistic. If we focus on marriage and type is education for instance, then positive assortative matching has the most educated woman marrying the most educated man, then the second most educated woman marrying marrying the second most educated man, and so on. In practice the most educated woman would weigh several criteria in deciding upon a match; even in the frictionless world studied by theory, the social surplus her match creates may be higher if she marries a man with less education but, say, a similar income. Since income and education are only imperfectly correlated, the optimal match must trade off assortative matching along these two dimensions. This point is quite general: with multiple types, the stark predictions of the one-dimensional case break down.

Empirical models of matching have long felt the need to accommodate the imperfect assortative matching observed in the data, of course. This can be done by introducing noise, in the form of heterogeneity in creation of surplus that is unobserved by the analyst (see Choo and Siow (2006).) Models with multidimensional types can also be estimated from the data, as in Chiappori, Salanié, Tillman, and Weiss (2008). But as far as we know, there has been little theoretical work exploring the properties optimal or equilibrium matches in such models. This is the task we set ourselves in the theoretical (and for now the only) part of this paper: we analyze the set of cross-products across partners that can be rationalized by such a model, and we show how to estimate this set from data and to test that the observed matching is socially optimal.

We first focus on a model without noise: the surplus created by a match is an unknown function of the types of the partners only. With one-dimensional types, positive assortative matching of course implies that the correlation between the types of partners is maximized. We show that with multidimensional types, the predictions of the matching model can be phrased in terms of the "cross-products", that is the set of expectations of functions of type dimensions of both partners across matches. In the education/income example, these could boil down to the three proportions of matches with given education levels of the partners, along with the average income of men matched with an educated woman, and the average income of women matched with an educated man.

Our main result shows that the cross-products implied by the socially optimal match must be on the frontier of the convex set of all feasible cross-products. We also show how knowing the observed cross-products in a dataset is enough to recover the match surplus function if the observed match is socially optimal; and to reject this hypothesis otherwise. The computational burden required for estimation is testing is much alleviated by the existence of very fast algorithms for matching.

We abstract away in this version from the determination of individuals who are unmatched in the optimal assignment, e.g. singles on the marriage market. This is without loss of generality when there is no unobserved heterogeneity, or when the unobserved heterogeneity is "separable", in a sense that we make clear. Then our analysis is just conditional on the optimal set of individuals who have a match at the optimum.

While we use the language of the economic theory of marriage in our illustrations, nothing we do actually depends on it. The methods proposed in this paper apply just as well to any one-to-one matching problem[1].

In work in progress, we implement our method to study the marriage market using data from the Current Population Survey in the US. We plan to use six waves of the June CPS, between 1971 and 1995[2]. This version of the paper only reports illustrative results on the 1995 wave.

1 The Assignment Problem

In all of the paper, we assume that two subpopulations M and W must be matched; each man (as we will call the members of M) must be matched with one and only one member of W (we will call them women.) Thus we do not allow for unmatched individuals—we return to this assumption later.

Each man m has an r-dimensional type x^m, and each woman w has a s-dimensional type y^w. Matching man m and woman w produces a social surplus $V(x^m, y^w)$, which we assume to be known by the social planner. For now we also assume away unobserved heterogeneity: V is a deterministic function of types.

A word on terminology: like most of the literature, we call a "match" the pairing of a man and a woman, and a "matching" or a complete set of matches, in which every man is paired with a woman, and vice versa.

1.1 Basic Assumptions

We consider K given *basis social benefit functions* $v_1(x, y), ..., v_K(x, y)$ whose values are interpreted as the social benefit of interaction between type x and type y. Given a *social weight* $\lambda \in \mathbb{R}^K$, we consider the *social benefit function*

$$V_\lambda(x, y) = \sum_{k=1}^{K} \lambda_k v_k(x, y) \tag{1.1}$$

[1]Or bipartite matchings in the terminology of applied mathematics.

[2]Later surveys unfortunately lack information that is required to analyze matches.

where the sign of each λ_k is unrestricted.

To return to the example in the introduction, which we denote (ER): there $r = s = 2$, the first dimension of types is education $E \in \{D, G\}$ (dropout or graduate), and the second dimension is income R, a continuous variable. Then a match between man m and woman w creates a surplus

$$V(x^m, y^w) = \lambda_{EE} v_{EE}(E^m, E^w) + \lambda_{ER} v_{ER}(E^m, R^w) + \lambda_{RE} v_{RE}(R^m, E^w) + \lambda_{RR} v_{RR}(R^m, R^w).$$

We shall make the simplifying restriction to restrict each $v_{ij}(x_i, y_j)$ to be the cross-product $x_i y_j$ of the types. With only one dimension ($r = s = 1$), this would give us the simplest form of complementarity between types—and depending on the sign of λ_{11}, positive or negative assortative matching would be optimal. Thus cross-products provide us with a useful starting point. We explain later how our inference procedure can be adapted to actually estimate the set of (v_{ij}) functions that rationalize an observed matching as optimal.

In our (ER) example, after obvious normalizations that do not change the structure of the assignment problem we can now write, with $E = (D, G)$ coded as $(0, 1)$:

$$V(x^m, y^w) = \mathbb{1}(E^m = E^w) + \lambda_{ER} E^m R^w + \lambda_{RE} R^m R E^w + \lambda_{RR} R^m R^w.$$

We now introduce a class of basis social benefit functions which will be of particular interest in the sequel.

Quadratic Interactions (QI) will denominate the case where the basis social benefit functions are polynomial functions of degree 2, namely (changing the indexation) $v_{ij}(x, y) = x_i y_j$. Denoting Λ the matrix of λ_{ij}, we get that

$$V_\Lambda(x, y) = \sum_{\substack{1 \le i \le r \\ 1 \le j \le s}} \lambda_{ij} x_i y_j = x' \Lambda y \qquad (1.2)$$

As moments involving only x or y depend on the fixed marginal distributions of the types and not on the matching, the terms of degree zero and one do not matter in the expression, and can be omitted, as well as the terms in x_i^2 and y_j^2. Thus under (QI), we could have equivalently chosen to define the surplus as

$$\hat{V}_\Lambda(x,y) = -\sum_{\substack{1\le i\le r\\ 1\le j\le s}} \lambda_{ij}(x_i - y_j)^2,$$

as maximizing $E[V_\Lambda(X,Y)]$ and maximizing $E\left[\hat{V}_\Lambda(X,Y)\right]$ over the feasible matchings are clearly equivalent. Also note that the quadratic restriction on the basis functions implies the notable restriction that for any x and y $1 \le i \le r$ and $1 \le j \le s$, the cross-derivative $\frac{\partial^2 V_\Lambda}{\partial x_i \partial y_j}$ does not depend on the other dimensions x_{-i} and y_{-j}.

We will sometimes make the stronger assumption that characteristics do not cross-interact, eg. there is no variation of surplus of matching a rich man with a more or less educated woman, all other things being equal:

Diagonal Quadratic Interactions (DQI) will describe the case $r = s$ and $\Lambda_{ij} = 0$ for $i \neq j$, so that

$$V_\lambda(x,y) = \sum_{i=1}^{r} \lambda_i x_i y_i. \tag{1.3}$$

In this form, it is clear that the relative importance of the λ_i's reflects the relative importance of the criteria. Thus λ_i also measures the social concern for matching partners who are similar in dimension i.

In the (ER) example, this would lead to the following specification

$$V(x^m, y^w) = -\lambda_E(E^m - E^w)^2 - \lambda_R(R^m - R^w)^2,$$

and λ_R would measure how much more it matters for incomes of partners to be similar[3] than it matters for educations.

[3]Or dissimilar, if $\lambda_R < 0$.

1.2 Notation and Definitions

We denote P (resp. Q) the distribution of types x (resp. y) in the subpopulation M (resp. W.) Thus P is a probability distribution on $\mathbb{R}^r$ and Q is a distribution on $\mathbb{R}^s$. In actual datasets we will have a finite number N of men and women, so that

$$P = \frac{1}{N} \sum_{m=1}^{N} \delta_{x_m}, \text{ and } Q = \frac{1}{N} \sum_{w=1}^{N} \delta_{y_w}.$$

Choosing a set of matches is essentially equivalent to choosing a correlation structure, or a copula, between types of partners. We denote $\Gamma(P, Q)$ the set of random vectors (x, y) such that the marginal distribution of x is P and that of y is Q; and $\mathcal{M}(P, Q)$ is the set of distributions of elements of $\Gamma(P, Q)$.

The formal definition of a *matching* (or *assignment*) follows immediately: it is defined as a probability measure π on whose marginals coincide with P and Q, so that the set of matchings is $\mathcal{M}(P, Q)$.

All of our analysis is easily adapted to the case where P and Q are absolutely continuous with respect to the Lebesgue measure; this case can be thought as an approximation to the previous one, with a large number of individuals and a distribution of characteristics which is sufficiently spread.

Alternatively (and equivalently), a matching π is described by two functions:

- one (π^m) mapping the set of types in the support of P to probability distributions over the support of Q;
- and one (π^w) mapping the set of types in the support of Q to probability distributions over the support of P.

Then in a matching π, a man of type x will be matched randomly with a woman of type y drawn according to the conditional distribution

$$\pi^w(y; x) = \frac{\pi(x, y)}{\int \pi(x, z)\, dz},$$

and similarly for a woman of type y.

A matching is said to be *pure* if both π^m and π^w are Dirac masses. Then there exists an invertible function $T = \pi^m$ such that π assigns zero outer probability to the set $\{(x, y) : y \neq T(x)\}$; and $T^{-1} = \pi^w$.

In a pure matching π, a man with type x almost surely is matched with a woman of type $y = T(x)$, and conversely, a woman with type y almost surely is matched with a man of type $x = T^{-1}(y)$.

2 Solving for the Optimal Matching

We focus here on the matching that maximizes the total surplus. When utilities are transferable across partners, we shall see below that the social optimum can be decentralized through a competitive matching market. In any case, the socially optimal matching solves

$$K(\Lambda) := \max_{\pi \in \mathcal{M}(P,Q)} E_\pi \left[V_\Lambda (x, y) \right]. \tag{2.1}$$

Unless explicitly mentioned otherwise, we will make the Quadratic Interactions (QI) assumption throughout the rest of the paper, thus V_Λ has the form (1.2). Note however that the results and the methodology can be directly generalized to the more general restriction for V given by (1.1). We feel however that such a generalization would come at cost of an extra expositional burden, so we chose to restricts the exposition to the quadratic case.

In order to study the structure of optimal matchings, we introduce some terminology. For any matching π, we denote $C(\pi)$ the (r, s) matrix of expectations of cross-products of types within matches:

$$C_{ij}(\pi) = E_\pi x_i y_j.$$

We call such a matrix C_π a *cross-product matrix*.

The *feasible set of cross-products*, denoted $\mathcal{F}$, is the set of cross-product matrices C for which there exists a matching π, not necessarily pure, such that $C = C(\pi)$.

For any given matrix Λ, the *efficient set of cross-products* $\mathcal{E}(\Lambda)$ is the set of feasible cross-product matrices $C(\pi)$ such that π solves (2.1).

Finally, the *rationalisable set of cross-products*, denoted $\mathcal{R}$, is the union of the efficient sets of cross-products $\mathcal{E}(\Lambda)$ when the matrix Λ varies in $\mathbb{R}^{rs}$. When C is in $\mathcal{R}$, there exists a matrix Λ and a matching π such that

- π solves (2.1) for this value of Λ;
- and $C = C(\pi)$.

Then we say that Λ *rationalizes* the cross-product matrix C.

Finally, we call the graph of the feasible set of cross-products in $\mathbb{R}^{rs}$ the *covariogram.*

To sum up, the assignment problem can be viewed as follows: given distributions P and Q on sets of types, a feasible matching π is a joint distribution that has P and Q as marginals; and any feasible matching π implies a matrix of cross-products C. The resulting feasible C define the set $\mathcal{F}$ in $\mathbb{R}^{rs}$. Among all feasible matchings, the optimal matchings for a matrix of social weights Λ solve (2.1) and generate cross-products in $\mathcal{E}(\Lambda)$. Finally, the union of all sets $\mathcal{E}(\Lambda)$ is the set of rationonalizable cross-products $\mathcal{R}$.

Note that these sets are unchanged if the surplus $V_\Lambda(x, y)$ is changed into $V_\Lambda(x, y) - u(x) - v(y)$, for any functions u and v. This remark will be important when considering the competitive equilibrium.

2.1 The feasible set

We begin by studying the set of feasible cross-products. Since much of what we do uses convexity, we first recall some definitions[4].

Take any set $Y \subset \mathbb{R}^d$; then the *convex hull* of Y is the set of points in $\mathbb{R}^d$ that are convex combinations of points in Y. We usually focus on its closure, the closed convex hull.

[4]We refer the reader to Ekeland and Témam (1976) for more.

The *support function* S_Y of Y is defined as

$$S_Y(x) = \sup_{y \in Y} x \cdot y$$

for any x in Y. It is a convex function, and it is homogeneous of degree one. Moreover, $S_Y = S_{\text{cch}(Y)}$ where $\text{cch}(Y)$ is the closed convex hull of Y, and $\partial S_Y(0) = \text{cch}(Y)$.

Now let u be a convex, continuous function defined on $\mathbb{R}^d$. Then the gradient ∇u of u is well-defined almost everywhere and locally bounded. If u is differentiable at x, then

$$u\left(x'\right) \geq u(x) + \nabla u(x) \cdot (x' - x)$$

for all $x' \in \mathbb{R}^d$. Moreover, if u is also differentiable at x', then

$$\left(\nabla u(x) - \nabla u\left(x'\right)\right) \cdot \left(x - x'\right) \geq 0.$$

When u is not differentiable in x, it is still *subdifferentiable* in the following sense. We define $\partial u(x)$ as

$$\partial u(x) = \left\{ y \in \mathbb{R}^d : \forall x' \in \mathbb{R}^d, u\left(x'\right) \geq u(x) + y \cdot (x' - x) \right\}.$$

Then $\partial u(x)$ is not empty, and it reduces to a single element if and only if u is differentiable at x; in that case $\partial u(x) = \{\nabla u(x)\}$.

Our first result is as follows:

Proposition 1 *a) The feasible set $\mathcal{F}$ is a non-empty closed convex set in $\mathbb{R}^{rs}$.*

b) The function $K(\Lambda)$ is convex and homogeneous of degree one, and it is the support function of the set of feasible cross-products $\mathcal{F}$.

Proof a) Non-emptiness is obvious. Now $\mathcal{F}$ is convex: Let $\hat{C}$ and $\tilde{C}$ be two feasible cross-product matrices in $\mathcal{F}$. We first show that for any $\alpha \in [0, 1]$, $\alpha\hat{C} + (1-\alpha)\tilde{C}$ is in $\mathcal{F}$.

By definition of $\mathcal{F}$, there exist $\hat{\pi}$ and $\tilde{\pi}$ in $\mathcal{M}(P, Q)$ such that $\hat{C}_{ij} = E_{\hat{\pi}}[X_{ij}Y_{ij}]$ and $\tilde{C}_{ij} = E_{\tilde{\pi}}[X_{ij}Y_{ij}]$.

Let $\bar{\pi} = \alpha\hat{\pi} + (1-\alpha)\tilde{\pi}$. Then $\alpha\hat{C}_{ij} + (1-\alpha)\tilde{C}_{ij} = E_{\bar{\pi}}[X_{ij}Y_{ij}]$, and $\bar{\pi} \in \mathcal{M}(P, Q)$, thus $\alpha\hat{C} + (1-\alpha)\tilde{C} \in \mathcal{F}$.

Now we prove that $\mathcal{F}$ is closed: Let C_n be a sequence in $\mathcal{F}$ converging to $C \in \mathbb{R}^{rs}$, and let π_n be the associated matching. By Theorem 11.5.4 in Dudley (2002), as $\mathcal{M}(P, Q)$ is uniformly tight, π_n has a weakly converging subsequence in $\mathcal{M}(P, Q)$; call π its limit. Then C is the cross-product associated to π, so that $C \in \mathcal{F}$.

b) By definition of $\mathcal{F}$, we can rewrite K as

$$K(\Lambda) = \sup_{C \in \mathcal{F}} \sum_{i=1}^{r} \sum_{j=1}^{s} \lambda_{ij} C_{ij} \tag{2.2}$$

which shows that K is the support function of $\mathcal{F}$, and in particular that it is convex and homogeneous of degree one.

2.2 The rationalisable set

We have the following important characterization of the rationalisable set as the frontier of the feasible set of cross-products.

Proposition 2 *a) The frontier of $\mathcal{F}$ is the rationalisable set $\mathcal{R}$.*

b) The efficient set $\mathcal{E}(\Lambda)$ is the subgradient of K at Λ, $\mathcal{E}(\Lambda) = \partial K(\Lambda)$.

c) In particular, when K is differentiable at Λ, $\mathcal{E}(\Lambda)$ is a singleton $\{C^*(\Lambda)\}$; there is a unique optimal matching $\pi^*(\Lambda)$, and it is pure; finally,

$$C^*_{ij}(\Lambda) = E_{\pi^*(\Lambda)}[x_i y_j] = \frac{\partial K(\Lambda)}{\partial \lambda_{ij}}.$$

Proof a) The frontier of $\mathcal{F}$ is given by $\mathrm{Fr}(\mathcal{F}) = \bigcup_{\Lambda \in \mathbb{R}^{rs}_*} \partial K(\Lambda)$; but as we saw, $\mathcal{E}(\Lambda) = \partial K(\Lambda)$, and thus the frontier of $\mathcal{F}$ is precisely $\mathcal{R}$.

b) By definition, $\mathcal{E}(\Lambda) = \arg\max \left\{ C \in \mathcal{F} : \sum_{i=1}^{d} \Lambda_{ij} C_i \right\}$. From the envelope theorem, it follows that $\mathcal{E}(\Lambda) = \partial K(\Lambda)$.

c) follows as an immediate corollary.

Any value of the social weights Λ where $K(\Lambda)$ is not differentiable generates a kink in the frontier $\mathcal{R}$ of the covariogram. These kinks are not theoretical curiosa; in fact, with finite datasets the rationalizable frontier always has a finite number of kinks. We shall mention two important special cases concerning the regularity of K.

Proposition 3 *a) When the distributions P and Q are absolutely continuous with respect to the Lebesgue measure, the function K is everywhere differentiable.*

b) When the distributions P and Q are discrete with N atom points of mass $1/N$, the function K is piecewise linear. The space $\mathbb{R}^{rs}$ of Λ matrices is partitioned into a finite number p of convex cones, in the interior of which the efficient cross-product matrix $C^*(\Lambda)$ is constant. Let $\{C_1, ..., C_p\}$ be the values of these p constants. The rationalisable set $\mathcal{R}$ is a polytope whose vertices are the C_k's.

Proof a) By Brenier's theorem (cf. Villani (2003), pp. 66-67), when the distributions P and Q are absolutely continuous with respect to the Lebesgue measure, then the optimal transportation plan $\pi \in \mathcal{M}(P, Q)$ which is the distribution of $(X, Y) \in \Gamma(P, Q)$ solution to (2.1) is unique for any (nonzero) Λ.

In particular, $\partial K(\Lambda) = \{C^*(\Lambda)\}$ where $C^*_{ij}(\Lambda) = E[x_i y_j]$ at the optimal matching. Therefore K is differentiable at Λ.

b) Denote $P = \frac{1}{N}\sum_{m=1}^N \delta_{x_m}$, and $Q = \frac{1}{N}\sum_{w=1}^N \delta_{y_w}$. Let $\mathfrak{S}_N$ denotes the set of permutations of $\{1, ..., N\}$. Then $\mathcal{F}$ is by definition the convex hull of $C(\sigma)$ where $C(\sigma)_{ij} = \frac{1}{N}\sum_{m=1}^N x_{ij}^m y_{ij}^{\sigma(m)}$. Since $K(\lambda) = \max_{\sigma \in \mathfrak{S}_N} \frac{1}{N}\sum_{m=1}^N x^{m\prime}\Lambda y^{\sigma(m)}$, K is piecewise linear.

As $\mathcal{F} = conv\,(C(\sigma) : \sigma \in \mathfrak{S}_N)$, it follows that $\mathcal{F}$ is a polytope, and there is a subset of $\{C(\sigma) : \sigma \in \mathfrak{S}_N\}$ which is the set of extreme points of $\mathcal{F}$, which are also the vertices of $\mathcal{F}$.

Figure 1 illustrates the case of more empirical interest in which the distributions P and Q are discrete. The "ideal" or limiting case of continuous distributions is illustrated in Figure 2. The details of the construction of these figures can be found in the appendix.

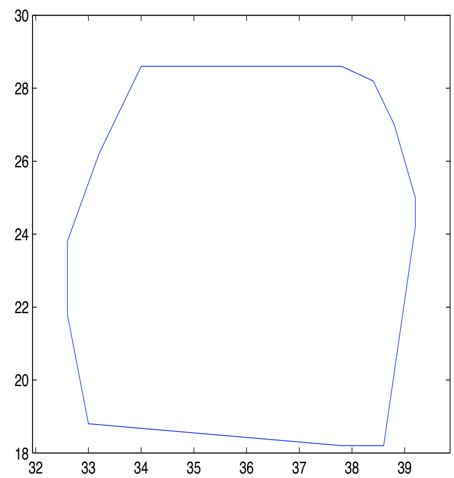

Figure 1: Covariogram when both distributions are discrete. The x (resp. y) axis measures the cross-product of the first (resp. second) characteristics.

2.3 Comparative statics

First, note that $\mathcal{R}$ can be defined in an alternative manner, which directly leads to an interpretation in terms of *sacrifice ratios*:

Proposition 4 *a) Take any* Λ *and any cross-product matrix* C *that is efficient for* Λ*. Reorder and redefine dimensions so that* $\lambda_{11} > 0$*. Then* C_{11} *is the value of the constrained optimization problem* (CO)*:*

$$\max_{(X,Y)\in\Gamma(P,Q)} E\left[X_1 Y_1\right] \;:\; E\left[X_i Y_j\right] \geq C_{ij} \;\; forall(i,j) \neq (1,1).$$

b) When K *is differentiable at* Λ*, then on the rationalisable frontier*

$$\frac{dC_{ij}}{dC_{kl}} = -\frac{\lambda_{kl}}{\lambda_{ij}}. \tag{2.3}$$

Proof a) Assume that C_{11} is the value of (CO). Write the Lagrangian associated to the problem: there exist some real numbers l_{ij}, $(i,j) \neq (1,1)$ such that C is the solution of

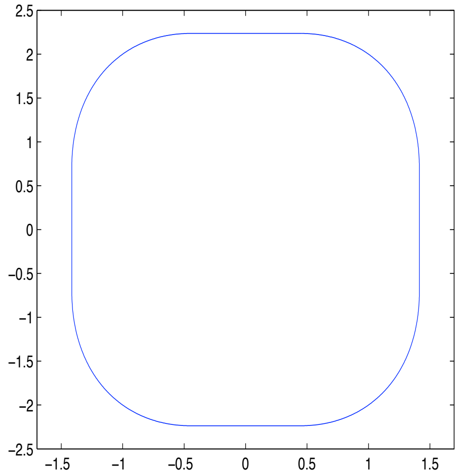

Figure 2: Covariogram for two bivariate Gaussian distributions.

$\max_{c\in\mathcal{F}} c_{11} + \sum_{(i,j)\neq(1,1)} l_{ij}C_{ij}$. Therefore, $C \in \mathcal{E}(\Lambda)$ if we define Λ by $\lambda_{11} = 1$ and $\lambda_{ij} = l_{ij}$ for $(i,j) \neq (1,1)$.

Conversely, assume that C_{11} is not the value of (CO). Then either $C \notin \mathcal{F}$, in which case C cannot be in $\mathcal{E}(\Lambda)$, or $C \in \mathcal{F}$ and $C_{11} < \max_{(X,Y)\in\Gamma(P,Q)} E[X_1Y_1] \ : \ E[X_iY_j] \geq C_{ij}$, $\forall (i,j) \neq (1,1)$.

But as $\mathcal{F}$ is a convex closed set, this would imply that C is in the strict interior of $\mathcal{F}$, hence by Theorem 1 in this case as well C cannot be in $\mathcal{E}(\lambda)$. This completes the proof.

b) Recall that $K(\Lambda)$ is positive homogeneous of degree 1. Then by Euler's theorem, if K is differentiable at Λ then

$$\sum_{i=1}^{r}\sum_{j=1}^{s} \lambda_{ij} \frac{\partial K}{\partial \lambda_{ij}}(\Lambda) = K(\Lambda).$$

Differentiating this w.r.t. some λ_{kl}, we get

$$\sum_{i=1}^{r}\sum_{j=1}^{s} \lambda_{ij} \frac{\partial^2 K}{\partial \lambda_{ij} \partial \lambda_{kl}}(\Lambda) = 0.$$

But we know from Theorem 1 that

$$\frac{\partial K}{\partial \lambda_{ij}}(\Lambda) = C^*_{ij}(\Lambda);$$

it follows that

$$\sum_{i=1}^{r}\sum_{j=1}^{s} \lambda_{ij} \frac{\partial C_{ij}}{\partial \lambda_{kl}}(\Lambda) = 0 \qquad (L)$$

Now Euler's equation can be rewritten as

$$\sum_{i=1}^{r}\sum_{j=1}^{s} \lambda_{ij} C^*_{ij}(\Lambda) = 0;$$

a fortiori,

$$\left(\sum_{a=1}^{r}\sum_{b=1}^{s} \lambda_{ab} \frac{\partial C^*_{ab}}{\partial \lambda_{ij}}\right) d\lambda_{ij} + \left(\sum_{a=1}^{r}\sum_{b=1}^{s} \lambda_{ab} \frac{\partial C^*_{ab}}{\partial \lambda_{kl}}\right) d\lambda_{kl} + C^*_{ij} d\lambda_{ij} + C^*_{kl} d\lambda_{kl} = 0.$$

Given (L), the first two terms are zero and we obtain

$$C^*_{ij} d\lambda_{ij} + C^*_{kl} d\lambda_{kl} = 0.$$

Thus on the rationalisable frontier $\frac{dC_{ij}}{dC_{kl}} = -\frac{\lambda_{kl}}{\lambda_{ij}}$.

The interpretation of part b) of this result is clearest under (DQI). With several dimensions for types, the optimal matching must sacrifice some cross-product in one dimension to the benefit of some cross-product in another. The implied sacrifice ratio, quite naturally, is exactly the ratio of the social weights along these dimensions. Note that in particular, when there are only two observed characteristics, if we set $\lambda_1 = 1$ and $\lambda_2 = \varepsilon$, then the function $\varepsilon \to C^*_1(1, \varepsilon)$ is decreasing, and the function $\varepsilon \to C_2*(1, \varepsilon)$ is increasing. Therefore, when one puts more weight on the second dimension, the cross-product of the characteristics in the second dimension increases, while that of the first dimension decreases.

2.4 Market equilibria and the structure of rationalisable matchings

2.4.1 Market equilibria

We now turn to the case where utilities are transferable. In this case, it is known that the socially optimal assignment can be decentralized though a competitive matching market. Given a matching $\pi \in \mathcal{M}(P,Q)$, define two *payoff functions* $\mathcal{U} : \mathcal{X} \to \mathbb{R}$ and $\mathcal{V} : \mathcal{Y} \to \mathbb{R}$ such that

$$\mathcal{U}(x) + \mathcal{V}(y) \leq V(x,y), \ \forall (x,y) \in \operatorname{supp}(\pi) \tag{2.4}$$

We say that payoff functions $\mathcal{U}$ and $\mathcal{V}$ *stabilize* a matching π if

$$\mathcal{U}(x) + \mathcal{V}(y) \geq V(x,y), \ \forall (x,y) \in \mathcal{X} \times \mathcal{Y} \tag{2.5}$$

Gretsky, Ostroy, and Zame (1999) proved that

Proposition 5 (Gretsky, Ostroy, and Zame) *A matching π is stabilized by two payoff functions $\mathcal{U}$ and $\mathcal{V}$ if and only if it is optimal in the sense of equation (2.1).*

The matching market can be seen as a cooperative game where a matching should not be blocked by any coalition of players. A matching which is not blocked by any coalition is said to be stable. Two players (a man and a woman) might decide to block a given matching if together they generate a larger surplus than the sum of their payoffs in the current matching, hence the stability condition (2.5). The set of stable matchings is the core of this cooperative game where utilities are supposed transferable, and in a stable matching where x and y are matched, the equation $V(x,y) = \mathcal{U}(x) + \mathcal{V}(y)$ indicates that $\mathcal{U}(x)$ is the share of the surplus assigned to the man, and $\mathcal{V}(y)$ is the share of the surplus assigned to the woman.

2.4.2 Structure of the rationalisable matchings

Proposition 6 *a) When the distributions P and Q are absolutely continuous with respect to the Lebesgue measure, then a matching is efficient for parameter value Λ, if and only* *if there exists a convex function $\mathcal{U}$ such that x gets almost surely matched with an y that satisfies*

$$\Lambda y \in \partial \mathcal{U}(x).$$

b) When the distributions $P = \frac{1}{N}\sum_{m=1}^{N} \delta_{x_m}$ and $Q = \frac{1}{N}\sum_{w=1}^{N} \delta_{y_w}$ are discrete with N atom points, let π be any matching that is optimal for some parameter value Λ. Then there exists two utility vectors $(u_m)_{m=1,...,N}$ and $(v_w)_{w=1,...,N}$ such that for any m and w,

$$u_m + v_w \geq x^{m\prime} \Lambda y^w,$$

with equality if and only if man m and woman w are matched.

Proof a) Suppose that Λ is invertible (by density there is no loss of generality in doing so). Assume that matching π is optimal for parameters Λ. Then by Proposition 5, there exist payoff functions $\mathcal{U}$ and $\mathcal{V}$ that stabilize the matching. Denote $y' = \Lambda y$, and $V'(y') = V(\Lambda^{-1} y')$, and let π' be the measure image of π by $(Id \times \Lambda)$. Then the matching π' is stabilized by the payoff functions $\mathcal{U}$ and $\mathcal{V}'$, that is

$$\mathcal{U}(x) + \mathcal{V}'(y') \geq x \cdot y'$$

with equality if and only if $(x, y') \in \text{supp}(\pi')$. Note that payoffs $\mathcal{U}$ and $\mathcal{V}'$ can be taken convex without lack of generality, in which case the equality condition translates into:

$$(x, y') \in \text{supp}(\pi')$$

if and only if $y' \in \partial \mathcal{U}(x)$, that is $\Lambda y \in \partial \mathcal{U}(x)$. The reverse implication works similarly.

b) This follows directly from Proposition 5.

We now state a proposition that shows that when all weights but one tend to zero, the classical one-dimensional assortative matching obtains in the limit. We refer the reader to Carlier, Galichon, and Santambrogio (2008) for a proof and for a more detailed investigation of the structure of the limiting matchings.

Proposition 7 *When $\lambda_1 = 1$ and $\lambda_j \to 0$ for $j \geq 2$, let letting $\pi^*(\Lambda)$ be the Λ-optimal matching, and $(X, Y) \sim \pi^*(\Lambda)$. Then the distribution of the first characteristics (X_1, Y_1) converges towards the distribution π_1 such that $(X_1, Y_1) \sim \pi_1$ is maximally correlated, or equivalently such that X_1 and Y_1 are comonotone, thus recovering classical Positive Assortative Matching.*

3 Empirical Estimation

3.1 The Matching Revealed Preferences (MaRP) estimator

Our empirical strategy starts with an observed matching of N matches, on which we first straightforwardly estimate the matrix of the cross-products of the observed types $\hat{c}_N$, as well as the marginal distributions of types $\hat{P}_N$ and $\hat{Q}_N$. Then we use our theory to answer two questions:

1. is the observed matching optimal?
2. which parameter Λ best rationalizes the observed matching (exactly if the observed matching is optimal, approximately if it is not)?

To address these two questions, we also need to construct the covariogram for the observed marginal distributions $\hat{P}_N$ and $\hat{Q}_N$, maximizing the empirical K_N objective function. Then we denote $\mathcal{F}_N$ the feasible cross-products set, and $\mathcal{R}_N$ its frontier. Finally, the cross-products that results from matching men and women randomly (so that π is the independent product of $\hat{P}_N$ and $\hat{Q}_N$) plays a special role, so we denote it c_N^0. Note that without loss of generality we can assume that the distribution of the men and the women's types are centered, so we will take $c_N^0 = 0$.

Given this data, we use the following simple proposition:

Proposition 8 *There exists a unique $\hat{t}_N \geq 1$ such that $\hat{t}_N \hat{c}_N \in \mathcal{R}_N$; and $\hat{t}_N = 1$, if and only if $\hat{c}_N \in \mathcal{R}_N$.*

There exists a convex cone $\hat{L}_N \subset \mathbb{R}^{rs}$ *such that* $\hat{t}_N \hat{c}_N \in \partial K_N(\Lambda)$ *if and only if* $\Lambda \in \hat{L}_N$.

Before we give the proof of the proposition, note that in the discrete case the covariogram is a polyhedrom, and the kinks of this polyhedron are in fact typical, in the sense that the set of Λ's which do not map to a vertex of this polyhedron are of measure zero: they coincide with the points of nondifferentiability of K. When N gets large, these vertices get closer and closer and (if the distribution in the population are absolutely continuous with respect to the Lebesgue measure), the covariogram converges to a smooth convex set in the Hausdorff topology.

Proof Note that $c_N^0 = 0 \in \mathcal{F}_N$ as this is the cross-product matrix obtained by the independent matching (again, we have supposed that the types distributions were centered). (X, Y) where $X \sim \hat{P}_N$, $Y \sim \hat{Q}_N$ and (X, Y) independent. The existence and unicity of $\hat{t}_N$ follows from the compacity of $\mathcal{F}_N$.

Therefore $\hat{t}_N = 1$ if and only if $\hat{c}_N$ is rationalizable, and the matrices $\Lambda \in \hat{L}_N$ are the estimators of the social weights we are looking for. We now turn to a characterization of $\hat{t}_N$ and $\hat{L}_N$. We now denote $\Lambda \cdot c$ the matrix scalar product $tr(\Lambda' c)$.

Theorem 1 *Take a number* t *and a matrix* Λ. *Then* $t = \hat{t}_N$ *and* $\Lambda \in \hat{L_N}$ *if and only if* (t, Λ) *solves*

$$\min_{\Lambda} \max_{t \geq 0} K_N(\Lambda) + t(1 - \Lambda \cdot \hat{c}_N);$$

or equivalently, if Λ *solves*

$$\min_{\Lambda} K_N(\Lambda) \ \ s.t. \ \ \Lambda \cdot \hat{c}_N = 1$$

and t *is the Lagrange multiplier associated to the constraint.*

Proof The first problem is the Lagrangian associated to the second problem. The first order condition associated to the first problem is $\hat{t}_N \hat{c}_N \in \partial K_N(\Lambda)$, QED.

Note that $t > 0$, hence if $\hat{c}_{rs} \neq 0$, one can rewrite the constrained optimization problem as an unconstrained optimization problem on $\mathbb{R}^{rs-1}$ as

$$\min_{(\lambda_{1,1},\ldots,\lambda_{r,s-1})} \Phi\left(\lambda_{1,1},...,\lambda_{r,s-1}\right)$$
$$\Phi\left(\lambda_{1,1},...,\lambda_{r,s-1}\right) \quad : \quad = K\left(\lambda_{1,1},...,\lambda_{r,s-1},\frac{1-\hat{c}^{1,1}\lambda_{1,1}-...-\hat{c}^{r,s-1}\lambda_{r,s-1}}{\hat{c}^{r,s}}\right),$$

which is the form we are going to use to determine Λ using a gradient algorithm. For a choice of $\epsilon > 0$:

Algorithm 1 (MaRP) *Take some initial choice of* $(\lambda_{1,1},...,\lambda_{r,s-1}) \neq 0_{\mathbb{R}^{rs-1}}$.

1. Compute $\lambda_{r,s} = \frac{1-\hat{c}^{1,1}\lambda_{1,1}-...-\hat{c}^{r,s-1}\lambda_{r,s-1}}{\hat{c}^{r,s}}$. *By running a standard assignment algorithm, (cf. Bertsekas (1981)), compute the optimal matching with weights* Λ *and the associated cross-product* C. *Compute* $\nabla\Phi = \left(C^{ij} - \frac{\hat{c}^{ij}}{\hat{c}^{rs}}C^{rs}\right)_{ij}$, *and* $\lambda'_{ij} := \lambda_{ij} - \epsilon\left(C^{ij} - \frac{\hat{c}^{ij}}{\hat{c}^{rs}}C^{rs}\right)$.

2. Replace Λ *by* Λ', *and iterate step* 1 *until* $\|\Lambda - \Lambda'\|$ *is small enough. One has* $C = t\hat{c}$, *with* $t \geq 1$. *Return* $t = C^{11}/\hat{c}^{11}$ *and* Λ.

3.2 Asymptotic analysis

It is possible to show that when the distribution of men and women's types in the population are absolutely continuous with respect to the Lebesgue measure, then our estimator (t_N, Λ_N) has an asymptotic Gaussian distribution with standard $n^{1/2}$ rate of convergence. Exposition is left for future work.

4 An Empirical Illustration

To explore the fruitfulness of our proposed approach, we extracted data on married couples in the US from the June 1995 Current Population Survey. We chose to focus on the trade-off between matching on education levels and matching on earnings. Thus by necessity we exclude couples in which one at least of the partner is not working at the time of the survey. To reduce the heterogeneity in the sample, we only keep couples whose members are both

Table 1: Descriptive statistics

	Age of husband	Age of wife	Education, H	Education, W	Earnings, H	Earnings, W
Q1	35	34	1	1	0.473	0.249
Median	40	38	2	2	0.653	0.385
Q3	44	43	2	2	0.898	0.575

white, non-Hispanic, were born in the US from US-born parents. We also eliminate couples where one of the partners was younger than 30 or older than 50 in 1995. This selection leaves us with 1,133 couples.

We recoded the education variable in the CPS so that it takes three values:

0. high school dropout

1. high school graduate

2. college graduate.

Table 1 describes the data (earnings are measured in thousands of dollars per week.)

Denote E_m and Y_m (resp. E_w and Y_w) the education level and earnings of the husband (resp. the wife.) We specify the social surplus function to take into account educational endogamy, earnings endogamy, and the interaction between both:

$$V(E_m, Y_m, E_w, Y_w) = \sum_{e=0,1,2} \lambda_e \mathbb{1}(E_m = E_w = e) - \lambda_y (Y_m - Y_w)^2 - \sum_{e=0,1,2} \frac{\lambda_{ey}}{2} \mathbb{1}(E_m = E_w = e)(Y_m - Y_w)^2.$$

We expect all λ's to be positive, but we do not impose it at priori. Thus we end up with a set of 7 basis functions, two for each value $e = 0, 1, 2$:

$$\mathbb{1}(E_m = E_w = e) \text{ and } -\mathbb{1}(E_m = E_w = e)(Y_m - Y_w)^2,$$

and the function $-(Y_m - Y_w)^2$.

We gather the corresponding coefficients λ_e, λ_y and λ_{ey}, in a vector Λ.

This specification embodies three restrictions that may be too severe. First, it imposes that when partners have different education levels, the social surplus is the same whatever these levels are. Second, it restricts the earnings endogamy term to be quadratic. Third, it neglects other factors such as age, or at least it assumes the kind of orthogonality discussed in our subsection on heterogeneity. We will work to remove these restrictions in further research.

We first compute the efficient frontier of the feasible polytope by drawing 500 values of the vector Λ on the unit sphere S_5. For each value, we use the Munkres algorithm to determine the optimal matching. Given that we have to do it a large number of times, we reduced the sample size to 200 randomly drawn couples. The 500 optimal matchings trace the 6-dimensional efficient frontier.

With seven dimensions, the feasible polytope is not easy to describe. Recall that the seven components of a point in a polytope measure the expectations of the seven basis functions for the corresponding efficient matching. The first three are just the probabilities of a match in which both partners have the education level $e = 0, 1, 2$. The absolute value of the fourth basis function is minus the average square difference of earnings between partners; and the last three modulate it when partners have the same education level.

First, we note that three of our estimated $\hat{\lambda}_k$'s are actually negative: the λ_e for high-school dropouts, λ_y, and the λ_{ey} for high-school dropouts. Thus the social benefit function that comes closest to rationalizing the observed matching values marrying high-school dropouts to more educated and higher-earnings partners, as well as marrying partners with dissimilar incomes when they have different education levels. It is easy to play with the estimated parameters to answer "what if" questions. For instance, start from a couple whose members have the same education e and the same earnings; then changing the earnings of one partner by d thousand dollars a week (and keeping all matches as they are!) reduces the social benefit by $(\lambda_y + \lambda_{ey})d^2$, while changing the education level of one partner reduces it by λ_e. This gives an "income equivalent of education endogamy". We compute it as a weekly \$560 for high-school dropouts, \$1,450 for high-school graduates, and a much smaller \$460 for college graduates.

Rather than elaborate further on these results, we go directly to the test that the observed matching is efficient, and to our estimator for the social benefit function.

Easy algebra shows that $\hat{\lambda}$ and $\hat{t}$ can be obtained by the following steps:

1. find $(\hat{\lambda}_1, \ldots, \hat{\lambda}_{K-1})$ to minimize
$$K\left(\lambda_1, \ldots, \lambda_{K-1}, \frac{1 - \sum_{k=1}^{K-1} \lambda_k \hat{c}_k}{\hat{c}_K};\right)$$
2. define
$$\hat{\lambda}_N = \frac{1 - \sum_{k=1}^{K-1} \hat{\lambda}_k \hat{c}_k}{\hat{c}_K};$$
3. compute the cross-products $C(\hat{\lambda})$ for the efficient matching corresponding to $\hat{\lambda}$ (this is a direct by-product of the minimization in step 1);
4. regress $\hat{c}$ on $C(\hat{\lambda})$ by OLS, without a constant term;
5. take $\hat{t}$ to be the estimator in this regression.

The latter is justified by the fact that under the null of efficiency, $\hat{c} = C(\hat{\lambda})$ and so
$$\sum_{k=1}^{K} (\hat{c}_k - tC_k(\hat{\lambda}))^2$$
is minimized in $t = 1$.

When running this, we find a multiplier $\hat{t} = 0.83$. We are not ready to give a standard error yet, and so we refrain from overinterpreting this number; but it does seem reasonably close to one. Computing the social benefit from the observed matching confirms that given the estimated $\hat{\lambda}$, the observed matching is about 15% less efficient than the optimal matching. A social planner with these preferences would for instance regret that high-school dropouts are too endogamous, and other categories not enough. An alternative conclusion, which should wait for a formal test, is that the observed matching is inefficient; then the estimator $\hat{\lambda}$ has no particular claim to represent “social preferences”, and such statements are unwarranted.

5 Remarks and extensions

Our theory so far makes some assumptions that we intend to relax.

5.1 Single households

So far we have not allowed for unmatched individuals. In an optimal matching, some men and/or women may remain single, as of course some must if there are more individuals on one side of the market. The choice of the socially optimal matching can be broken down into the choice of the set of individuals who participate in matches and the choice of actual matches between the selected men and women. Our theory applies without any change to the second subproblem; that is, all of our results extend to M and W as selected in the first subproblem.

5.2 Unobserved heterogeneity

Of course, any empirical application must take unobserved heterogeneity into account. There are several ways to do this here. Let us focus on the situation where agents are matched efficiently relative to a vector of types of dimension $r = s$, but only the first $r_0 < r$ dimensions are observed by the econometrician. The true empirical cross-product $\hat{C} = \left(\hat{C}^1, ...\hat{C}^r\right)$ is on the efficient set $\mathcal{R}$ in $\mathbb{R}^r$; but the observed cross-product $\bar{C} = \left(\hat{C}^1, ...\hat{C}^{r_0}, 0, ..., 0\right)$ is in the intersection of the feasible set $\mathcal{F}$ and the space

$$V_{\text{obs}} = \{c \in \mathbb{R}^r : c_{r_0+1} = ... = c_r = 0\}.$$

Similarly, define

$$V_{\text{lat}} = \{c \in \mathbb{R}^r : c_1 = ... = c_{r_0} = 0\}.$$

For $c \in \mathbb{R}^r$, denote π_{obs} and π_{lat} the orthogonal projections of c on respectively V_{obs} and V_{lat}. In general, $\bar{C} = \pi_{\text{obs}}\left(\hat{C}\right)$ is not in the efficient set $\mathcal{R}$, and estimators based only on the observed cross-products will be biased. However, there is one important case in which the bias vanishes:

Proposition 9 *Assume that observed and latent characteristics are distributed independently within each subpopulation; that is,* $P = P_{obs} \otimes P_{lat}$ *where* P_{obs} *is the distribution on the* r_0 *first characteristics and* P_{lat} *is the distribution of the* $r - r_0$ *last characteristics, and with similar notation* $Q = Q_{obs} \otimes Q_{lat}$.

Then for any $C \in \mathcal{F}$,

$$C \in \mathcal{R} \text{ if and only if } \pi_{obs}(C) \in \mathcal{R} \text{ and } \pi_{obs}(C) \in \mathcal{R},$$

hence $\bar{C} = \pi_{obs}\left(\hat{C}\right)$ *allows unbiased estimation of* $\bar{\lambda} = \pi_{obs}(\lambda)$.

Proof The proof is similar to the argument in the Appendix A.3.

5.3 Rationalizing other functional forms

We would like to go beyond cross-products. Let us return to the more general form of the social surplus

$$V_{ij}(.,.) \equiv \lambda_{ij} v_{ij}(.,.),$$

where v_{ij} is a function known to the analyst and the scalar weights λ_{ij} are unknown to him (up to scale.)

For any specification of the v_{ij}'s , we can redefine "generalized cross-products" to be expectations of $E_\pi v_{ij}(x_i, y_j)$, and our results again apply, with a generalized covariogram that depends on the specification. A natural question is whether we can *estimate* the v_{ij} functions, assuming that the observed matching is socially optimal for one particular specification. This would amount to using the (generalized) algorithm in section 3.1 to determine the value of t for any given choice of the v's and then to solve the equation $t = 1$. While this is an equation with only one unknown and so it seems unlikely to give much identifying power, recall that it allowed us to estimate the Λ matrix; so the task may not be as hopeless at it looks.

A related question is whether there always exists a specification of the v's that makes the observed matching optimal. Or, to rephrase it, does assumption (TWI)—along with optimality—have any testable implication? We conjecture that the answer is positive, but we have yet to prove it.

Conclusion

It is a bit early to conclude, but we would like here to point out the strong links between our theory and the Revealed Preferences principle in utility theory, stemming from Afriat (1967) and Varian (1982). The key notion is that of *consistency with utility maximization*, which we recall here.

Definition 1 (Consistency with utility maximization) *We define consumptions* $x_k \in \mathbb{R}^d$*,* $k = 1...n$ *and prices* $\pi_k \in \mathbb{R}^d_+$ *to be consistent with utility maximization if and only if there exists a concave utility function* u *and wealths* w_k *such that*

$$x_k \in \arg\max u(x) : x \cdot \pi_k = w_k$$

Equivalently, this holds if and only if there exists scalars λ_k, $k = 1, ..., n$ and a concave function v such that $\lambda_k \pi_k = \nabla v(x_k)$, where $\lambda_k > 0$ is to be interpreted as a Lagrange multiplier in the convex optimization problem above. Introducing $U_k = u(x_k)$, we see that if the latter condition holds true then we have $U_l \leq U_k + \lambda_k \pi_k (x_l - x_k)$ for all k and l. Conversely, if this holds, then we can introduce $u(x) = \min_k \{U_k - \lambda_k \pi_k x_k\}$ and it follows immediately that (x_k, π_k) are consistent with utility maximization. Hence Afriat's theorem, and also Varian's theorem which can be deduced using the notion of cyclical monotonicity (cf. e.g. Villani (2003), p. 79). Note that we changed λ_k into $-\lambda_k$ to work with convex functions instead of concave ones.

Proposition 10 (Afriat; Varian) *The following conditions are equivalent:*

(i) Quantity-prices (x_k, π_k) *are consistent with utility maximization;*

(ii) There are scalars $\lambda_k < 0$*,* $k = 1, ..., n$ *such that*

$$\sum_{k=1}^{n} \sum_{i=1}^{d} \lambda_k x^i_{\sigma(k)} \pi^i_k$$

is maximized over $\sigma \in \mathfrak{S}_N$ *for* $\sigma = id$*;*

(iii) (Afriat's theorem) There are scalars $\lambda_k < 0$ *and* U_k*,* $k = 1, ..., n$ *such that*

$$U_l \leq U_k - \lambda_k \pi_k \left(x_l - x_k \right)$$

for all k *and* l*;*

(iv) (Varian's theorem) The (x_k, π_k)*'s satisfy the Generalized Axiom of Revealed preferences (GARP): for any any cyclical relabelling* $k_1, ..., k_p$ *(with* $k_{p+1} = k_p$*), all the quantities* $\pi_{k_i} \left(x_{k_{i+1}} - x_{k_i} \right)$ *are not* ≥ 0*.*

This formulation makes the link with our theory precise: while testing for consistency with utility maximization consists in looking for $\lambda_k < 0$, $k = 1, ..., n$ such that $\sum_{k=1}^{n} \sum_{i=1}^{d} \lambda_k x_{\sigma(k)}^i \pi_k^i$ is maximized over $\sigma \in \mathfrak{S}_N$ for $\sigma = id$, in our problem (with diagonal social weights) we look for λ_i, $i = 1, ..., d$ such that $\sum_{k=1}^{n} \sum_{i=1}^{d} \lambda_i x_{\sigma(k)}^i y_k^i$ is maximized over $\sigma \in \mathfrak{S}_N$ for $\sigma = id$.

This analogy implies that our approach also may provide a new way to test for consistency of preferences; we plan to explore this further.

References

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

Becker, G. (1973): "A theory of marriage, part I," *Journal of Political Economy*, 81, 813–846.

Bertsekas, D. (1981): "A new algorithm for the assignment problem," *Mathematical Programming*, 21, 152–171.

Carlier, G., A. Galichon, and F. Santambrogio (2008): "From Knothe's transport to Brenier's map and a continuation method for optimal transport," preprint available on http://arxiv.org/abs/0810.4153.

Chiappori, P.-A., B. Salanié, A. Tillman, and Y. Weiss (2008): "Assortative Matching on the Marriage Market: A Structural Investigation," mimeo Columbia University.

Choo, E., and A. Siow (2006): "Who Marries Whom and Why," *Journal of Political Economy*, 114, 175–201.

Dudley, R. M. (2002): *Real Analysis and Probability*. Cambridge University Press.

Ekeland, I., and R. Témam (1976): *Convex analysis and variational problems*. North-Holland.

Gale, D., and L. Shapley (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–14.

Gretsky, N., J. Ostroy, and W. Zame (1999): "Perfect competition in the continuous assignment model," *Journal of Economic Theory*, 88, 60–118.

Rachev, S., and L. Ruschendorf (1998): *Mass Transportation Problems. Volume I: Theory and Volume II: Applications*. Springer.

Varian, H. (1982): "The Nonparametric Approach to Demand Analysis," *Econometrica*, 50, 945—973.

Villani, C. (2003): *Topics in Optimal Transportation.* American Mathematical Society.

A Examples and particular cases of interest

A.1 Gaussian distributions

When $P = N(0, \Sigma_1)$ and $Q = N(0, \Sigma_2)$, one has

$$K(\lambda) = tr\left(\sqrt{\sqrt{\Sigma_1}\bar{\Sigma}_2(\lambda)\sqrt{\Sigma_1}}\right)$$

where $\bar{\Sigma}_2(\lambda) = D_\lambda \Sigma_2 D_\lambda$ with $D_\lambda = diag(\lambda)$. This follows from a calculation based on Rachev and Ruschendorf (1998) I, Ex. 3.2.12. In particular, when $d = 2$ and $P = N(0, I_2)$, and $Q = N\left(0, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}\right)$, this becomes $K(\lambda) = tr\left(\sqrt{\bar{\Sigma}_2(\lambda)}\right)$, thus making use of $tr\left(\sqrt{S}\right) = \sqrt{tr(S) + 2\sqrt{\det S}}$, we get

$$K(\lambda) = \sqrt{\sigma_1^2\lambda_1^2 + \sigma_2^2\lambda_2^2 + 2\sigma_1\sigma_2\lambda_1\lambda_2\sqrt{(1-\rho^2)}}.$$

In particular when $\sigma_1 = \sigma_2$ and $\rho = 1$, the covariogram is a circle.

We illustrated the Gaussian covariogram in Figure 2 in the case where $P_1 = N(0, I_2)$ and $P_2 = N\left(0, \begin{pmatrix} 2 & 3 \\ 3 & 5 \end{pmatrix}\right)$.

A.2 Multinomial distributions

When the distributions are discrete, as it is the case with a multinomial distribution, we have seen that the covariogram is a polytope. This was illustrated in Figure 1 in the text.

A.3 Independent marginals

For simplicity take $d = 2$, and suppose that $P = P_1 \otimes P_2$, and $Q = Q_1 \otimes Q_2$, that is the distributions of the various individuals' characteristics (male or female) are independent. Let F_{P_i} (resp. F_{Q_i}) be the cumulative distribution functions associated to distribution P_i (resp. Q_i).

Take $\lambda \in \mathbb{R}^2_{+*}$. Then by 4, the λ-rationalisable matching are the solutions of constrained optimization problem

$$\max_{(X,Y)\in\Gamma(P,Q)} E\left[X^1Y^1\right] \ : \ E\left[X^2Y^2\right] \geq c^2,$$

but by independence between the two dimensions, in the optimal coupling the pair $\left(X^1, Y^1\right)$ is independent from $\left(X^2, Y^2\right)$, thus the solution to the constrained problem coincides with the unconstrained one. Thus

$$\begin{aligned} C^1 &= \max_{(X^1,Y^1)\in\Gamma(P_1,Q_1)} E\left[X^1Y^1\right] \\ &= \int_0^1 F_{P_1}^{-1}(u)\, F_{Q_1}^{-1}(u)\, du := C^1_{\max} \end{aligned}$$

and a similar result holds for C^2. Therefore for any value of $\lambda \in \mathbb{R}^2_{+*}$, the optimal cross-product is given by $C(\lambda) = \left(C^1_{\max}, C^2_{\max}\right)$. For $\lambda_1 > 0$ and $\lambda_2 < 0$, $C(\lambda) = \left(C^1_{\max}, C^2_{\min}\right)$, and so on. This is illustrated in Figure 3.

A.4 Maximally dependent marginals

Again, take $d = 2$, and suppose that P is the distribution of $\left(X^1 = f_1(U), X^2 = f_2(U)\right)$ with $U \sim \mathcal{U}([0,1])$ and f_1 and f_2 two increasing and continuous functions, and similarly, that Q is the distribution of $\left(Y^1 = g_1(V), Y^2 = g_2(V)\right)$ with $V \sim \mathcal{U}([0,1])$ and g_1 and g_2 two increasing and continuous functions. Then the cross-product between X^1 and Y^1 determines the cross-product between the cross-product between X^2 and Y^2, thus

$$C(\lambda) = \max\left(\lambda_1 C^1_{\min} + \lambda_2 C^2_{\min}, \lambda_1 C^1_{\max} + \lambda_2 C^2_{\max}\right)$$

which is illustrated in in Figure 4.

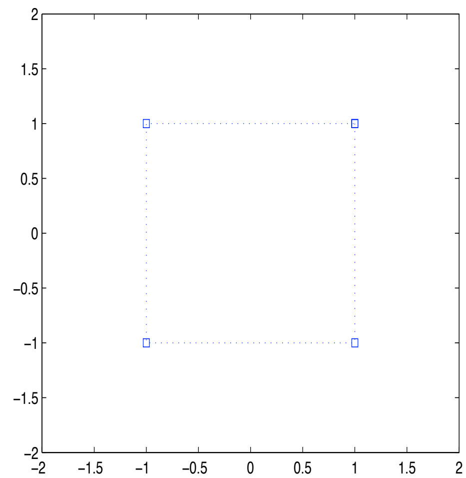

Figure 3: covariogram when both distributions have independent components. The x (resp. y) axis measures the cross-product of the first (resp. second) characteristics.

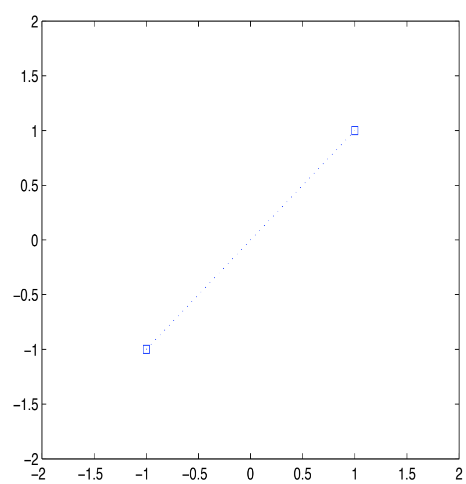

Figure 4: covariogram when both distributions have maximally dependent components. The x (resp. y) axis measures the cross-product of the first (resp. second) characteristics.