

Matching with Trade-offs:

Revealed Preferences over Competing Characteristics

Alfred Galichon¹

Bernard Salanié ²

First version dated December 6, 2008. The present version is of October 11, 2009³.

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

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

³The authors are grateful to Guillaume Carlier, Pierre-André Chiappori, Piet Gauthier, Jim Heckman, Guy Laroque, Rob Shimer as well as seminar participants at Crest, Ecole Polytechnique, séminaire Roy, University of Chicago, and University of Alicante for useful comments and discussions. Much of this paper was written while Salanié was visiting the Toulouse School of Economics; he is grateful to the Georges Meyer endowment for its support.

Abstract

We investigate in this paper the theory and econometrics 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, as well as on characteristics that the analyst does not observe. 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. Given a flexible specification of the surplus function, we characterize under mild assumptions the properties of the set of feasible matchings and of the socially optimal matching. Then we show how data on the covariation of the types of the partners in observed matches can be used to test that the observed matches are socially optimal for this specification, and to estimate the parameters that define social preferences over matches.

Keywords: matching, logit, generalized linear models, revealed preferences, contingency tables.

JEL codes: C78, D61, C13.

Introduction

Louisa was naturally ill-tempered and cunning; but she had been taught to disguise her real disposition, under the appearance of insinuating sweetness, by a father who but too well knew that to be married would be the only chance she would have of not being starved, and who flattered himself that 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, if 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 corresponding 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 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 analysts 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 of optimal or equilibrium matches in such models. We show in this paper that these properties can be summed up in simple measures of covariation of types across partners; we analyze the set of values of such measures that can be rationalized by a matching model; and we show how to estimate this set from data and to test that the observed matching is socially optimal¹.

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—or bipartite matchings, to use the terminology of applied mathematics. In fact, we can even extend them to problems in which the sets of partners are determined endogenously—as with same-sex unions. This is investigated in Section 7, where we consider possible extensions of our setting.

We do require, however, that utility be transferable across partners. Our primitive function is indeed the surplus created by a match. We posit that it is an unknown function of the types of the partners only, plus preference shocks that are observed by all participants but not by the analyst—in the nature of unobserved heterogeneity. When utility is transferable, all optimal matchings must maximize the joint surplus; and so does the equilibrium of the assignment game.

As is well-known, this model is too general to be empirically testable: even without unobserved heterogeneity, any observed assignment can be rationalized by a well-chosen surplus function. This is a consequence of a more general theorem by Blair (1984). Echenique (2008) shows that on the other hand, some collections of matchings are not rationalizable: if the analyst can observe identical populations on several assignments, then these assign-

¹A word on terminology: like most of the literature, we call a “match” the pairing of two partners, and a “matching” the list of all realized matches.

ments must be consistent with each other in a sense that his paper makes precise. But we are unlikely to have such data at hand in general.

Relatedly, analysts sometimes observe several subpopulations which are matching independently and yet have the same surplus function. Fox (2009) shows that under a “rank-order condition” on the unobserved heterogeneity, it is then possible to identify several important features of the surplus function, and in particular how important complementarity is on various dimensions.

While analyzing complementarity is also one of our goals here, many of the applications we have in mind do not fit Fox’s assumption that there be enough variation across subpopulations with identical surplus. Marriage markets, for instance, seem to be either so disconnected that their surplus functions are unlikely to be similar, or too connected to make it possible to ignore matching across markets. In this paper, we will posit that we are only given data on one instance of a matching problem, such as the marriage market in the US in the 1980s, or the market for CEOs. Our data will consist of the values of the observable types of both partners in each realized match, and of the types of unmatched individuals. Since the optimal/equilibrium matching is determined on the basis of both the observable and the (to us) unobservable types, we will need to impose assumptions that allow us to integrate over the distribution of the unobservable types in a manageable way. Our aim is to start from the observable matching (the distribution of matches across observable types) and to recover information on the observable surplus function (the average surplus of matches for given observable types of both partners.)

To achieve this, we first resort to a separability assumption that rules out interactions between the unobservable types of the partners in the surplus from a match. This was used by Choo and Siow (2006), and then generalized by Chiappori, Salanié, Tillman, and Weiss (2008) who showed that the matching equilibrium then boils down to a series of parallel discrete choice models. While this is an important step on the way to a solution, the resulting model is still too rich to be taken to the data. We need to restrict the distribution of unobserved heterogeneity, and we do this by adopting again Choo and Siow (2006)’s

assumption that gives rise to multinomial logits. Under these assumptions, we prove that the cross-differences of the surplus function over observable types are nonparametrically identified from the data. In particular, we can test for complementarities between any two observable dimensions of the types of the partners, such as the education of the wife and the income of the husband. We can also identify the relative strengths of such complementarities across different dimensions.

If the analyst is lucky enough to have very rich data, then unobserved heterogeneity is almost irrelevant and the observable matching maximizes the observable surplus function. On the other hand, if data is so poor that unobserved heterogeneity dominates, then the analyst should observe something that, to him, looks like completely random matching. We show that under our assumptions, this amounts to maximizing the mutual information of the match distribution—a statistical object that here measures covariation of partner types. Moreover, for any intermediate amount of unobserved heterogeneity, the observable matching maximizes a straightforward linear combination of the observable surplus and of mutual information.

This observation suggests a strategy: approximate the observable surplus function with a linear expansion over some known basis functions, with unknown “assorting weights”. Then all relevant information can be expressed in terms of the average values of these basis functions across couples, and our results have a very neat geometrical interpretation. Take the abstract space where each point represents an hypothetical set of values for all the basis functions. All feasible matchings generate points within a polytope in this space. We first show that even with our restrictive assumptions, any point in this polytope is rationalizable: if the variance of unobserved heterogeneity is well-chosen, then there exist assorting weights for which the optimal matching generates exactly these covariations. Fortunately, this combination of heterogeneity and assorting weights is in general unique as we shall see: this allows us to introduce several consistent and asymptotically normal estimators of both the assorting weights and the variance of the unobserved heterogeneity. Moreover, models without any unobserved heterogeneity can only generate points on the boundary of the

polytope, and so the homogeneous model is testable.

This paper thus proves both a negative and a positive result. The negative part is that even if we assume separable heterogeneity with a multinomial logit structure, the model still cannot be rejected. The positive part is that given any theory about the way the observable types enter the surplus function (as embodied in a set of basis functions), we exhibit well-behaved estimators of the unknown parameters; and we can quantify how much heterogeneity is needed to rationalize the data. Moreover, our methods can be used heuristically, to explore ways to understand what goes on in matching markets—and how they change across time and space. Standard statistical techniques could for instance be put to work to find the basis functions that explain the largest share of the variation in the data. Such a methodological stance is reminiscent of revealed preferences in consumer theory; in fact the analogy is very sharp, as the underlying theoretical structure is the same.

Our depiction of matching markets of course abstracts from many features of real-world markets. We focus on static, frictionless markets, as in much of the literature on marriage markets. Models of matching on job markets, for instance, have on the whole adopted a much more dynamic perspective, in which job flows in fact provide a lot of information on the underlying parameters. In the applications we have in mind, the surplus function may involve many more dimensions and we do not want to restrict it too much a priori. The basic lack of identification mentioned at the beginning of this introduction would become even more severe if we introduced dynamics or frictions, unless these additional features are drastically simplified. We leave this for further research. The paper also currently focuses on discrete characteristics; we are exploring possible extensions to continuous types.

Section 1 sets up the matching model we study in the paper, along with the assumptions on the specification of the observable surplus and the process that drives unobserved heterogeneity. In section 2 we build on these assumptions to derive our main analytical results, and we give them a geometric interpretation on section 3. Section 4 gives our results on identification, and Section 5 introduces our tests and estimators and derives their asymptotic properties. In section 6 we show that the presence of heterogeneity in fact makes the

computation of the optimal matching much easier. We conclude by sketching extensions of our methods.

Since much of what we do uses convexity, we recall some definitions and basic results in Appendix A. All proofs are collected in Appendix B. Finally, we should note that there are close parallels between the analysis we develop in the present paper and familiar notions in thermodynamics and statistical physics. E.g the social utility of a matching evokes (minus) the internal energy of a physical system, and the standard error of unobservable heterogeneity parallels its physical temperature. Since the analogy may prove to be as useful to others as it was to us, we elaborate on it in Appendix C.

Summary of the notation used in the paper. For the reader's convenience, we regroup here the notation introduced in the text. We consider matches between N men and N women. $\mathfrak{S}_N$ is the set of permutations of $\{1, \dots, N\}$. A man has a full type $\tilde{x} = (x, \varepsilon)$, where the econometrician observes x but not ε ; we use $\tilde{y} = (y, \eta)$ for a woman. x is a random vector with distribution P , and $\tilde{x}$ is distributed according to $\tilde{P}$; we use Q and $\tilde{Q}$ for a woman. We denote $\mathcal{M}(P, Q)$ the set of probability distributions with margins P and Q ; we use $\mathcal{M}(\tilde{P}, \tilde{Q})$ for the full types. We denote $P \otimes Q$ the product measure which matches men and women randomly. A feasible matching generates a probability $\tilde{\Pi} \in \mathcal{M}(\tilde{P}, \tilde{Q})$, which assesses the odds that a man with full type $\tilde{x}$ is married to a woman with full type $\tilde{y}$. A man with full type $\tilde{x}$ and a woman with full type $\tilde{y}$ generate together a full surplus $\tilde{\Phi}(\tilde{x}, \tilde{y})$. We call $\Phi(x, y) = E\left(\tilde{\Phi}(\tilde{X}, \tilde{Y}) | X = x, Y = y\right)$ the observable surplus; in some of the paper we take it to be the structural quadratic form $\Phi(x, y) = x' \Lambda y$.

1 The Assignment Problem

Throughout the paper, we assume that two subpopulations M and W of equal size 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 model the determination of

the unmatched population (the singles) in this paper; we take it as data. We elaborate on this point in our concluding remarks. Note also that we assumed bipartite matching: the two subpopulations which define admissible partners are exogenously given. This assumption can also be relaxed; see Section 7.

Throughout the paper, we illustrate results on the education/income example sketched in the Introduction, which we denote (ER).

1.1 Population characteristics

Each man m has an r -dimensional type x_m of observable characteristics, and a vector of unobserved characteristics ε_m . Denote $\tilde{x}_m = (x_m, \varepsilon_m)$ the full description of man m 's characteristics, which we call his full type. Each woman w similarly has an s -dimensional type y_w of observed characteristics, and a full type $\tilde{y}_w = (y_w, \eta_w)$.

We denote $\tilde{P}$ (resp. $\tilde{Q}$) the distribution of *full* types $\tilde{x}$ (resp. $\tilde{y}$) in the subpopulation M (resp. W), and P (resp. Q) the distribution of *observable* types x (resp. y). Thus P is a probability distribution on $\mathbb{R}^r$ and Q is a distribution on $\mathbb{R}^s$. In observed datasets we will have a finite number N of men and women, so that P and Q are the empirical distributions over the characteristics samples of the men $\{x_1, \dots, x_N\}$ and the women $\{y_1, \dots, y_N\}$, respectively.

Take the education/income example: there $r = s = 2$, the first dimension of types is education $E \in \{D, G\}$ (dropout or graduate), and the second dimension is income class R , which takes values 1 to n_R . P describes both the number of graduates among men and the distributions of income among graduate men and among dropout men.

1.2 Matching

The intuitive definition of a matching is the specification of “who marries whom”: given a man of index $m \in \{1, \dots, N\}$, it is simply the index of the woman he marries, $w = \sigma(m) \in \{1, \dots, N\}$. Imposing that each man be married to one and only one woman at a given

time translates into the requirement that σ be a permutation of $\{1, \dots, N\}$, which we denote $\sigma \in \mathfrak{S}_N$. This definition is too restrictive in so far as we would like to allow for some randomization. This could arise because a given type is indifferent between several partner types; or because the analyst only observes a subset of relevant characteristics, and the unobserved heterogeneity induces apparent randomness.

A *feasible matching* (or *assignment*) is therefore defined in all generality as a joint distribution $\tilde{\Pi}$ over types of partners $\tilde{X}$ and $\tilde{Y}$, such that the marginal distribution of $\tilde{X}$ is $\tilde{P}$ and the marginal distribution of $\tilde{Y}$ is $\tilde{Q}$. We denote $\mathcal{M}(\tilde{P}, \tilde{Q})$ the set of such joint distributions. Note that when x and y are univariate, a feasible matching can be equivalently specified through a copula.

A matching is said to be *pure* if all conditional distributions $\tilde{\Pi}(\cdot|\tilde{x})$ and $\tilde{\Pi}(\cdot|\tilde{y})$ are point mass distributions. In a pure matching $\tilde{\Pi}$, there exists an invertible map $T(\tilde{x})$ such that a man with type $\tilde{x}$ almost surely marries a woman of type $\tilde{y} = T(\tilde{x})$, and conversely, a woman with type $\tilde{y}$ almost surely marries a man of type $\tilde{x} = T^{-1}(\tilde{y})$. (Of course, in the discrete case this map can be represented as a permutation on indices.)

In the education/income example (ER), a pure matching is described by $(2n_r)^2 - 1$ numbers; but the marginals impose $2(2n_r - 1)$ constraints, so that $(2n_r - 1)^2$ numbers are to be determined.

1.3 Surplus of a match

The basic assumption of the model is that matching man m of type $\tilde{x}_m$ and woman w of type $\tilde{y}_w$ generates a joint surplus $\tilde{\Phi}(\tilde{x}_m, \tilde{y}_w)$, where $\tilde{\Phi}$ is a deterministic function. Along with most of the matching literature, we assume that

Assumption (O): Observability. Each agent observes the full characteristics $\tilde{x}$ and $\tilde{y}$ of all men and all women, but the econometrician only observes the subvectors x and y .

Assumption (O) rules out asymmetric information between participants in the market, as the economics of matching with incomplete information is a subject of its own. On the other hand, we do not need to assume full information as the notation seems to imply: $\tilde{\Phi}$ could for instance be reinterpreted as the expectation of a random variable conditional on $\tilde{x}, \tilde{y}$, as long as all participants evaluate it in the same way.

Given Assumption (O), we need to define the *observable surplus* as the best predictor of $\tilde{\Phi}(\tilde{x}, \tilde{y})$ conditional on x and y , that is

$$\Phi(x, y) = E \left[\tilde{\Phi}(\tilde{X}, \tilde{Y}) | X = x, Y = y \right]$$

and we can write the decomposition

$$\tilde{\Phi}(\tilde{x}, \tilde{y}) = \Phi(x, y) + k(\tilde{x}, \tilde{y})$$

where $k(\tilde{x}, \tilde{y})$ is the *idiosyncratic surplus*.

Following the insight of Choo and Siow (2006), formalized by Chiappori, Salanié, Tillman, and Weiss (2008), we now assume:

Assumption (S): Separability. Let $\tilde{x}$ and $\tilde{x}'$ have the same observable type: $x = x'$. Similarly, let $\tilde{y}$ and $\tilde{y}'$ be such that $y = y'$. Then

$$\tilde{\Phi}(\tilde{x}, \tilde{y}) + \tilde{\Phi}(\tilde{x}', \tilde{y}') = \tilde{\Phi}(\tilde{x}, \tilde{y}') + \tilde{\Phi}(\tilde{x}', \tilde{y}).$$

While much of the literature on matching emphasizes complementarity, assumption (S) in fact requires that conditional on observable types, the surplus exhibit no complementarity across unobservable types.

It is easy to see that imposing assumption (S) is equivalent to requiring that the idiosyncratic surplus from a match must be additively separable, in the following sense:

$$k(\tilde{x}, \tilde{y}) = \chi(\tilde{x}, y) + \xi(\tilde{y}, x),$$

where χ and ξ are two deterministic functions and

$$E(\chi(\tilde{X}, Y) | X = x, Y = y) = E(\xi(\tilde{Y}, X) | X = x, Y = y) = 0.$$

Then the surplus function $\tilde{\Phi}$ can be rewritten as

$$\tilde{\Phi}(\tilde{x}, \tilde{y}) = \Phi(x, y) + \chi(\tilde{x}, y) + \xi(\tilde{y}, x).$$

Note that the model is invariant if one rescales the three terms on the right-hand side by the same positive constant. Later on we will normalize these three components.

As proved in Chiappori, Salanié, Tillman, and Weiss (2008), assumption (S) implies that at the optimum (or equilibrium), a given individual (say, a man $\tilde{x}$) has a preference $\xi(\tilde{x}, y)$ for a particular class of observable characteristics (say y), but he is indifferent between all partners which have the same y but a different η .

In fact, the optimal matching is characterized by two functions of observable characteristics $U(x, y)$ and $V(x, y)$ that sum up to $\Phi(x, y)$ such that if a man $\tilde{x} = (x, \varepsilon)$ is matched with a woman of characteristics $\tilde{y} = (y, \eta)$, he will get utility

$$U(x, y) + \chi(\tilde{x}, y)$$

while his match gets utility

$$V(x, y) + \xi(\tilde{y}, x).$$

Chiappori, Salanié, Tillman, and Weiss (2008) showed that given assumption (S), the matching problem boils down to a set of discrete choice models for each type of man and of woman: for instance, man $\tilde{x}$ is matched in equilibrium to a woman $\tilde{y}$ whose observable type y maximizes

$$U(x, y) + \xi(\tilde{x}, y)$$

over all values in the support of Q .

While this is already quite useful, we need to add more restrictions on the specification of the components of the idiosyncratic surplus $\chi(\tilde{x}, y)$ and $\xi(\tilde{y}, x)$.

1.4 Specifying the idiosyncratic surplus

Following Choo and Siow (2006) and Chiappori, Salanié, Tillman, and Weiss (2008), we introduce the following assumption²:

Assumption GUI: Gumbel Unobserved Interactions. It is assumed that:

- There are an infinite number of individuals with a given observable type in the population

- Fix the observable characteristics x of a man, and let $(y_1^*, \dots, y_{T_y}^*)$ be the possible values of the observable characteristics of women. Then the vector of preference shocks $\chi(x, \varepsilon, y_1^*), \dots, \chi(x, \varepsilon, y_{T_y}^*)$ are distributed as T_y independent and centered Gumbel random variables with scale factor σ_1 ;

similarly,

- Fix the observable characteristics y of a man, and let $(x_1^*, \dots, x_{T_x}^*)$ be the possible values of the observable characteristics of men. Then the vector of preference shocks $\xi(y, \eta, x_1^*), \dots, \xi(y, \eta, x_{T_x}^*)$ are distributed as T_x independent and centered Gumbel random variables with scale factor σ_2 .

■

In short: men of a given observable type have conditionally Gumbel iid draws of the χ 's for different individuals; and conversely for women of a given observable type.

We use (GUI) for the Independence of Irrelevant Alternatives property: without it, the odds ratio of the probability that a man with observable type x ends up in a match with a woman of observable type y rather than with z would also depend on the types of other women, and the model would become unmanageable.

(GUI) underlies the standard multinomial logit model of discrete choice. It has well-known limitations, one of which is that it does not extend directly to continuous choice. We are currently exploring alternative specifications that would allow us to deal with continuous

²We define the scale factor to be 1 for the standard Gumbel, which has variance $\pi^2/6$; thus e.g. χ has variance $\sigma_1^2 \pi^2/6$.

characteristics; but at this stage, we assume

Assumption (DD): The distributions of observed types P and Q are discrete, with probability mass functions $p(x)$ and $q(y)$. ■

In the (ER) example for instance, $p(D, 3)$ is the proportion of men who are dropouts and whose income lies in class 3. For simplicity, we now denote $i_P = 1, \dots, n_P$ the possible values of types of men, and $i_Q = 1, \dots, n_Q$ for women.

1.5 Specifying the observable surplus

We now introduce sets of assumptions on the observable surplus ranging from non-restrictive (NPOI below, suited for nonparametric identification) to more restrictive (SLOI below, convenient for a more concise analysis).

Let us first impose a normalization convention on the observable surplus. Notice that the optimal matching (but not the value of the social surplus) is left unchanged if we add an additively decomposable function $f(x) + g(y)$ to $\Phi(x, y)$. Therefore, without any loss of generality, we impose some identifying restriction on Φ , using the two-way ANOVA decomposition, according to which any vector $\Phi(x, y)$ admits the following orthogonal decomposition in $L^2(\pi)$ as

$$\Phi(x, y) = \bar{\Phi}(x, y) + f(x) + g(y) + c$$

where $E_p[f(X)] = E_q[g(Y)] = 0$ and $E[\bar{\Phi}(X, Y)|X] = E[\bar{\Phi}(X, Y)|Y] = 0$. We shall therefore often take the following convention when using a nonparametric approach:

Convention (ZMOI): Zero-mean Observable interactions. The observable surplus satisfies

$$E[\Phi(X, Y)|X] = E[\Phi(X, Y)|Y] = 0.$$

It will sometimes be useful to assume more structure on the function Φ (the observable joint surplus.) To do this, we consider K given *basis assorting functions* $\phi^1(x, y), \dots, \phi^K(x, y)$ whose values are interpreted as the utility benefit of interaction between type x and type y .

Given *assorting weights* $\Lambda \in \mathbb{R}^K$, we focus on *observable surplus functions* $\Phi_\Lambda(x, y)$ which are linear combinations of the basis assorting functions with weights Λ . That is,

Assumption (SLOI): Semilinear Observable Interactions. The observable surplus function can be written

$$\Phi_\Lambda(x, y) = \sum_{k=1}^K \Lambda_k \phi^k(x, y) \quad (1.1)$$

where the sign of each Λ_k is unrestricted. ■

Note that in the discrete case which we restrict to in this paper, this general form is absolutely *not* restrictive. Indeed, one can choose $K = T_x \times T_y$ and chose $\phi^{ij}(x, y) = 1_{\{x=x_i, y=y_j\}}$ so that $\phi^{ij}(x, y)$ captures interaction between observable man type x_i and observable woman type y_j . We shall refer to this specification as the:

Specification (NPOI): Nonparametric Observable Interactions. The observable surplus function is expanded in all generality

$$\Phi_\Lambda(x, y) = \sum_{i=1}^{T_x} \sum_{j=1}^{T_y} \Lambda_{ij} 1_{\{x=x_i, y=y_j\}}. \quad (1.2)$$

in which case social weight Λ_{ij} coincide with $\Phi(x_i, y_j)$. ■

However we favor parsimonious models for the sake of analysis, so in general we shall only assume (SLOI), unless explicitly stated.

To return to the education/income example (ER): we could for instance assume that a match between man m and woman w creates a surplus that depends on the similarity of the partners in both education and income dimensions. The corresponding specification would be (with education levels $E = (D, G)$ coded as $(0, 1)$):

$$\Phi(x_m, y_w) = \sum_{e_m=0,1; e_w=0,1} \Lambda_{e_m, e_w} \mathbf{1}(E_m = e_m, E_w = e_w) + \sum_{i=1, \dots, n_r; j=1, \dots, n_r} \Lambda_{ij} \mathbf{1}(R_m = i, R_w = j).$$

This specification only has $(n_r^2 + 4)$ parameters, while an unrestricted specification would have $4n_r^2$. Such an unrestricted specification would for instance allow the effect of matching partners in income class 3 to depend on both of their education levels.

An even more restrictive, “diagonal” specification would be

$$\Phi(x_m, y_w) = \sum_{e=0,1} \Lambda_e^E \mathbf{1}(E_m = E_w = e) + \sum_{i=1, \dots, n_r} \Lambda_i^R \mathbf{1}(R_m = R_w = i).$$

In this last form, it is clear that the relative importance of the Λ ’s reflects the relative importance of the criteria. Thus Λ_i^R measures the preference for matching partners who are both in income class i , while Λ_0^E measures the preference for matching dropouts. The relative values of these numbers indicate how social preferences value complementarity of incomes of partners more, relative to complementarity in educations. We will not need to assume such a diagonal structure in the following, although our results easily specialize to this case.

1.6 Summary: the model specification

Under assumptions (O), (S), (SLOI) and (GUI), the model is fully parametrized; its parameters can be collected in a vector

$$\theta = (\Lambda, \sigma_1, \sigma_2),$$

where Λ is the assorting weight matrix, and σ_1 (resp. σ_2) is the scale factor of the unobservable characteristics of the men (resp. of women). Without loss of generality, all components of θ can be multiplied by any positive number; hence we shall need to impose some normalization on θ .

Most of the results in the next section in fact only require assumptions (O), (S) and (GUI), with a general function $\Phi(x, y)$. In this case θ is just $(\Phi, \sigma_1, \sigma_2)$, and again it is defined up to a scale factor.

As we will see, the *total heterogeneity* ($\sigma_1 + \sigma_2$) plays a key role in our results; thus we introduce a specific notation for it:

$$\sigma = \sigma_1 + \sigma_2.$$

2 Solving for the Optimal Matching

In this section we only assume (O), (S), and (GUI), and we consider the problem of optimal matching:

$$\mathcal{W}(\theta) = \sup_{\tilde{\Pi} \in \mathcal{M}(\tilde{P}, \tilde{Q})} E_{\tilde{\Pi}} \left[\tilde{\Phi}(\tilde{X}, \tilde{Y}) \right]. \quad (2.1)$$

Our modeling strategy in this section and the next is to assume that the number of men and women in the population is large enough that averages can be replaced with expectations. When we describe our estimators in section 5, we of course take into account the fact that we only have a finite sample.

2.1 The heterogeneous model

Let us provide some intuition before we state a formal theorem. Under (O), (S) and (GUI), standard formulæ of the multinomial logit model give the expected utility of a man of observable type x at the optimal matching:

$$E \left[\max_y \left(U(x, y) + \chi(\tilde{X}, y) \right) | X = x \right] = \sigma_1 \log \sum_y \exp(U(x, y)/\sigma_1).$$

Therefore the expected social surplus from the optimal matching is simply³ (adding the equivalent formula for women of observable type y):

$$\sigma_1 E_P \log \sum_y \exp(U(X, y)/\sigma_1) + \sigma_2 E_Q \log \sum_x \exp(V(x, Y)/\sigma_2).$$

Now recall that $U(x, y)$ is the mean utility of a man with observable type x who ends up being matched to a woman with observable type y at the optimum. As in the general development of the theory of matching, U is the value of the multiplier of the population constraints; and as such, it (along with V) is the unknown function in the dual program

³Since this formula may not be entirely transparent, we develop one term below:

$$E_P \log \sum_y \exp(U(X, y)/\sigma_1) = \sum_x p(x) \log \sum_y \exp(U(x, y)/\sigma_1).$$

in which the expression for the social surplus above is minimized over all U, V such that $U + V \geq \Phi$. We now state this as a theorem (proved in Appendix B):

Theorem 1 (Social welfare-primal version) *Assume (O), (S), (GUI) and (DD). Then*

$$\mathcal{W}(\theta) = \inf_{(U,V) \in A} \left(\sigma_1 E_P \log \sum_y \exp(U(X, y)/\sigma_1) + \sigma_2 E_Q \log \sum_x \exp(V(x, Y)/\sigma_2) \right) \quad (2.2)$$

where the constraint set A is defined by the inequalities

$$\forall x, y, U(x, y) + V(x, y) \geq \Phi(x, y).$$

At an optimal matching, men with observable type x will be found in matches with women with observable types y such that $U(x, y) + V(x, y) = \Phi(x, y)$. The expected utility of men with observable type x matched with women of observable type y is $U(x, y)$.

This theorem also has a primal version, of course. While deriving it takes a bit more work (again, see Appendix B), the intuition is simple. First, if there were no unobserved heterogeneity (with σ close to zero) the optimal matching would coincide with the optimal observable matching Π , which solves

$$\mathcal{W}(\theta) = \sup_{\Pi \in \mathcal{M}(P, Q)} E_{\Pi} \Phi(X, Y).$$

Going to the polar opposite, in the limit when σ goes to infinity only unobserved heterogeneity would count; and since it is just noise, the optimal matching would simply assign partners randomly, yielding the product measure $P \otimes Q$.

As it turns out, when σ takes any intermediate value the optimal matching maximizes a weighted sum of these two extreme cases:

Theorem 2 (Social welfare-dual version) *Under the assumptions of Theorem 1*

$$\mathcal{W}(\theta) = \sup_{\Pi \in \mathcal{M}(P, Q)} \left(\sum_{x, y} \pi(x, y) \Phi(x, y) - \sigma I(\Pi) \right) + \sigma_1 S(Q) + \sigma_2 S(P), \quad (2.3)$$

where $S(P)$ and $S(Q)$ are the entropies of P and Q given by

$$S(P) = - \sum_x p(x) \log p(x); \quad \text{and} \quad S(Q) = - \sum_y p(y) \log p(y);$$

and $I(\Pi)$ is the mutual information of joint distribution Π , given by

$$I(\Pi) = \sum_{x,y} \pi(x,y) \log \frac{\pi(x,y)}{p(x)q(y)}.$$

The mutual information $I(\Pi)$ is nothing else than the Kullback-Leibler divergence of Π from the independent product $P \otimes Q$ to Π . Recall two important information-theoretic properties of I :

1. The map $\pi \rightarrow I(\pi)$ is strictly convex.
2. One has

$$\forall \Pi \in \mathcal{M}(P, Q), \quad S(P) + S(Q) \geq I(\Pi) \geq 0$$

the left handside becoming an equality in particular in the case of a pure matching, and the right handside inequality becoming an equality in the case where $\Pi = P \otimes Q$, as we shall see below.

Mutual information is a measure of the covariation of types x and y . Now $P \otimes Q$ is the independent product of P and Q , which corresponds to a completely random matching $\Pi = P \otimes Q$. Thus a large positive $I(\Pi)$ indicates that the matching Π induces strong correlation across types; $I(\Pi) = S(P) + S(Q)$ if and only if $\Pi = P \otimes Q$. If σ is very large then the Theorem suggests that $I(\Pi)$ should be minimized, which can only occur for the independent matching $\Pi = P \otimes Q$; whereas if σ is negligible then Π should be chosen so as to maximize the expected *observable* surplus $E_{\Pi} \Phi(X, Y)$. This corroborates the intuition given earlier.

Now the optimal matchings coincide with the solutions to this maximization problem. Since we only observe the realized Π over observable variables, Theorem 2 defines the empirical content of the model: a combination of the parameters $\theta = (\Phi, \sigma_1, \sigma_2)$ is identified if and only if the solution Π depends non-trivially on it.

We already knew that θ can be rescaled by any positive constant without altering the solution. We can now go one step further: while all components of θ figure in this theorem,

σ_1 and σ_2 only enter through their sum σ . Thus and as announced, σ_1 and σ_2 are not separately identified.

Accordingly, we redefine the parameter vector θ as

$$\theta = (\Phi, \sigma),$$

or $\theta = (\Lambda, \sigma)$ under (SLOI).

2.2 The homogeneous limit

In this section we consider the limit behavior of our model when σ goes to zero, so that unobservable heterogeneity vanishes. We denote

$$\mathcal{W}_0(\Phi) \equiv \mathcal{W}(\Phi, 0).$$

By taking the limit in Theorem 1, we obtain:

Theorem 3 (Homogeneous social welfare) *Assume (O) and (DD); then*

a) The value of the social optimum when $\theta = (\Phi, 0)$ is given both by

$$\mathcal{W}_0(\Phi) = \max_{\Pi \in \mathcal{M}(P, Q)} \sum_{x, y} \pi(x, y) \Phi(x, y), \quad (2.4)$$

and by

$$\mathcal{W}_0(\Phi) = \inf_{(u, v) \in A^0} \left(\sum_x p(x) u(x) + \sum_y q(y) v(y) \right) \quad (2.5)$$

where the constraint set A^0 is given by

$$\forall x, y, \quad u(x) + v(y) \geq \Phi(x, y);$$

A matching $(X, Y) \sim \Pi$ is optimal for Φ if and only if the equality

$$u(X) + v(Y) = \Phi(X, Y)$$

holds Π -almost surely, where u and v solve the optimization problem (2.5).

Thus in the limit we recover the standard primal and dual formulation of the matching problem; since all men with observable characteristics x have the same tastes, they all obtain the same utility at the optimum and $U(x, y)$ becomes a function of x only, which we denoted $u(x)$ above; and this is just the Lagrange multiplier on the population constraint

$$\sum_y \pi(x, y) = p(x)$$

which is implicit in the notation $\Pi \in \mathcal{M}(P, Q)$.

3 Qualitative properties of the optimum

In this section we first introduce the various statistics on which our analysis shall rest. We then provide comparative statics which help understanding the influences on the model parameters on these statistics; last, we study the influence on qualitative properties of the equilibria such as uniqueness and purity of the equilibria.

3.1 Matching summaries

Feasible summaries. Recall that under (SLOI), there exists an unknown vector Λ such that the observable surplus function takes the form

$$\Phi(x, y) = \sum_{k=1}^K \Lambda_k \phi^k(x, y)$$

with known basis functions ϕ^k . Now consider a hypothetical matching Π ; under this matching, the basis functions have expected values

$$C^k(\Pi) = \sum_{x,y} \pi(x, y) \phi^k(x, y).$$

We call each C^k a *covariation*. Take the (ER) example; then

- C^1 is the proportion of matches among graduate partners under Π

- C^2 is the expected income of a graduate man's wife multiplied by the proportion of graduate men; C^3 is defined similarly
- and C^4 is the expected product of the partners' incomes.

Random matching, as represented by $\Pi_\infty = P \otimes Q$, plays a special role in our analysis, as it obtains in the limit when heterogeneity becomes very large. We denote the corresponding covariations as C_∞^k . At the polar opposite is the matching Π_0 that obtains in the homogenous limit $\sigma = 0$; we denote the implied covariations $C_0^k(\Lambda)$. Note that C_∞ does not depend on Λ , but C_0 does.

We know from Theorem 2 that under (SLOI), the observable optimal matching Π maximizes

$$\Lambda \cdot C(\Pi) - \sigma I(\Pi).$$

Thus the vector $(C(\Pi), I(\Pi))$ summarizes all the relevant information about matching Π . We call each such vector a *matching summary*; matching summary vectors are set in *summary space*, which is a subset⁴ of $\mathbb{R}^{K+1}$.

Given an observed matching, it is of course very easy to estimate the associated covariation vector and mutual information. Again, the model is scale-invariant and we may impose an arbitrary normalization on $\theta = (\Lambda, \sigma)$. For that purpose we choose a vector C^* and we impose $\Lambda \cdot C^* = 1$. Later we make the choice of C^* more specific.

Given population distributions P and Q , we define the *set of feasible summaries* $\mathcal{F}$ as the set of summary vectors (C, I) that are generated by some feasible matching $\pi \in \mathcal{M}(P, Q)$, that is

$$\mathcal{F} = \left\{ (C, I) \in \mathbb{R}^K \times [0, S(P) + S(Q)] : \exists \Pi \in \mathcal{M}(P, Q), C^k = C^k(\Pi), I = I(\Pi) \right\}$$

Similarly, define the *covariogram* $\mathcal{F}_c$ as the set of covariations C that are implied by some feasible matching; that is,

$$\mathcal{F}_c = \left\{ C : \exists \Pi \in \mathcal{M}(P, Q), C^k = C^k(\Pi) \right\}.$$

⁴Remember that $I(\Pi) \geq 0$ for any feasible matching.

Covariograms provide us with a nice graphical representation of the properties of a matching. Figure 1 illustrates their relevant properties, and the reader should refer to it as we go along. To fit it within two dimensions, we assume that there are only two basis functions; e.g. in the (ER) example we could have

$$\Phi(E_m, E_w, R_m, R_w) = \Lambda_1 \mathbf{1}(E_m = E_w) + \Lambda_2 \mathbf{1}(R_m = R_w),$$

so that Λ_1 (resp. Λ_2) measures the preference for assortative matching on educations (resp. income classes.)

Proposition 1 *Under (O), (S), (GUI) and (SLOI), the sets $\mathcal{F}$ and $\mathcal{F}_c$ are nonempty closed convex sets, and their support functions are $\mathcal{W}(\Lambda, \sigma)$ and $\mathcal{W}(\Lambda, 0)$, respectively.*

As will soon become clear, the boundaries of the convex sets $\mathcal{F}$ and $\mathcal{F}_c$ have special significance in our analysis. For now, let us simply note that the boundary of $\mathcal{F}_c$ exhibits kinks when these distributions of characteristics are discrete—which is always the case in our setting. The reason for these kinks is that in the discrete case, the optimal matching for homogenous types is generically stable under a small perturbation of the assorting weights Λ ; starting from almost every Λ 's, a small change in Λ leaves covariations unchanged. Any such value of Λ generates a covariation vector on a vertex of the polytope. On the other hand, there exist a finite number of values of Λ where the optimal matching problem has multiple solutions, with corresponding multiple covariations; each such value of Λ generates a facet of the polytope. This is shown on Figure 1 with all $\lambda = \Lambda_2/\Lambda_1$ in an interval $[\lambda_i, \lambda_{i+1}]$ generating the same covariations in the homogeneous case. Remarkably, such kinks disappear as soon as there is enough positive amount of heterogeneity; we will come back to this in section 3.3.

3.2 Rationalizable boundary

The previous discussion suggests an intimate connection between the boundaries of the sets described above and optimal matchings. To make this clear, we now define the *set*




Figure 1: The covariogram and related objects

of *rationalizable summaries* $\mathcal{R}$ as the set of $(K + 1)$ -uples (C, I) such that C and I are covariations and mutual information corresponding to an optimal matching $\Pi \in \mathcal{M}(P, Q)$ for some parameter values (Λ, σ) :

$$\mathcal{R} = \{(C, I) \in \mathcal{F} : \exists (\Lambda, \sigma) \in \mathbb{R}^K \times [0, S(P) + S(Q)], \Lambda \cdot C - \sigma I = \mathcal{W}(\Lambda, \sigma)\}.$$

Obviously, rationalizable summaries are feasible and $\mathcal{R} \subset \mathcal{F}$. This definitions allow us to state that rationalizable summaries and extreme feasible summaries coincide. Or, to put it more formally:

Proposition 2 *$\mathcal{R}$ is the frontier of $\mathcal{F}$ in $\mathbb{R}^K \times [0, S(P) + S(Q)]$.*

Mutual information level sets. Now consider a covariation vector C in the covariogram $\mathcal{F}_c$, and define the *rationalizing mutual information*

$$I_r(C) := \sup \{I \in [0, S(P) + S(Q)] : (C, I) \in \mathcal{R}\};$$

clearly from the definition of $\mathcal{R}$ and positive homogeneity of $\mathcal{W}$, we see that the *implicit mutual information function*

$$I_r(C) = \sup_{\lambda} \{\lambda \cdot C - \mathcal{W}(\lambda, 1)\} \quad (3.1)$$

so $I_r(C)$ is the Legendre-Fenchel transform of $\mathcal{W}(\Lambda, 1)$ which is strictly convex; in particular $I_r(C)$ is a C^1 function. Conversely, for any mutual information $I \geq 0$ we define and the *set of rationalizable covariations* by

$$\mathcal{R}_c(I) = \{C \in \mathbb{R}^K : \exists I \in [0, S(P) + S(Q)], (C, I) \in \mathcal{R}\} = I_r^{-1}(\{I\}).$$

It follows directly from the convexity of I_r that $\mathcal{R}_c(I)$ is the boundary of the set $I_r^{-1}([0, I])$, which is convex and increasing (for inclusion) with respect to $I \in [0, S(P) + S(Q)]$.

Note the two limiting cases: when mutual information I is zero (corresponding to random matching), then $\mathcal{R}_c(0) = \{C_\infty\}$, where $C_\infty^k = E_{p \otimes q}[\phi^k(X, Y)]$. When $I = S(P) + S(Q)$, $\mathcal{R}_c(S(P) + S(Q))$ consists of the extreme points of the covariogram $\mathcal{F}_c$.

The following result combines the linearity embodied in (SLOI) and the convex structure of the problem:

Proposition 3 *Under (O), (S), (GUI) and (SLOI),*

a) The social welfare function $\mathcal{W}$ is positive homogeneous of degree one in $\theta = (\Lambda, \sigma)$. It is convex on $\mathbb{R}^K \times [0, +\infty)$ and strictly convex on its interior.

b) The subdifferential of $\mathcal{W}$ at (Λ, σ) is given by the set of $(K + 1)$ -uples

$$\partial\mathcal{W} = \{(C(\Pi), -I(\Pi))\}$$

generated by an optimal matching Π when parameter values $\theta = (\Lambda, \sigma)$ vary.

In particular, when the optimal matching Π is unique for some θ , then $\mathcal{W}$ is differentiable at θ and

$$C^k(\Pi) = \frac{\partial\mathcal{W}}{\partial\Lambda_k}(\theta), \quad I(\Pi) = -\frac{\partial\mathcal{W}}{\partial\sigma}(\theta),$$

in which case we define

$$C^k(\theta) := \frac{\partial\mathcal{W}}{\partial\Lambda_k}(\theta), \text{ and } I(\theta) := -\frac{\partial\mathcal{W}}{\partial\sigma}(\theta).$$

c) The function $I_r(C)$ is C^1 on the interior of $\mathcal{F}_c$, and one has

$$\frac{\partial I_r}{\partial C^k} = \frac{\Lambda_k}{\sigma}.$$

As a corollary, the limiting homogeneous case also has interesting comparative statics, which closely parallel the results above. When $\sigma = 0$ the mutual information does not play a role anymore in Theorem 2; so we focus on the covariogram $\mathcal{F}_C$; and we define $\mathcal{W}_0(\Lambda) = \mathcal{W}(\Lambda, 0)$. Note in particular that the boundary of the covariogram—which is the polygon in Figure 1—consists of the set of the covariation vectors $C_0(\Lambda)$ when Λ varies.

Corollary 1 (Homogeneous comparative statics) *Under (O), (S), (GUI) and (SLOI),*

a) The function $\mathcal{W}_0$ is convex and positive homogeneous of degree one in Λ .

b) The subdifferential of $\mathcal{W}_0$ at Λ is given by the set of K -uples

$$\partial\mathcal{W}_0 = \{C(\Pi)\}$$

generated by an optimal matching Π when Λ varies and $\sigma = 0$.

In particular, when the solution Π is unique for some Λ , then $\mathcal{W}_0$ is differentiable at Λ and

$$\frac{\partial\mathcal{W}_0}{\partial\Lambda_k}(\Lambda) = C_0(\Pi).$$

Our basic result is that any vector of covariations C that is feasible (that belongs to $\mathcal{F}_C$) can be rationalized for a well-chosen value of total heterogeneity. This is a byproduct of the following result, which sums up the relationships between the sets we introduced:

Proposition 4 *Under (O), (S), (GUI) and (SLOI),*

a) *The sets $\mathcal{R}_c(I)$ are the set of extreme points of nested closed and convex sets that expand from $\{C_\infty\}$ to $\mathcal{F}_c$ as mutual information I goes from 0 to $\bar{I}_0$.*

b) *Any point $\hat{C} \in \mathcal{F}_c$ belongs to exactly one frontier $\mathcal{R}_c(I)$, associated to the mutual information $I = I_r(\hat{C})$.*

c) *For a point C such that $I_r(C)$ is smooth, letting $\Lambda_k = \frac{\partial I_r(C)}{\partial C^k}$; then along $\partial\mathcal{R}_c(I)$*

$$\frac{dC^i}{dC^j} = -\frac{\Lambda_j}{\Lambda_i} \tag{3.2}$$

Proposition 4 is illustrated on figure 1. Note that when we fix Λ and increase σ from 0 to $+\infty$, the summary vector (C, I) for the optimal matching moves continuously from $(C_0(\Lambda), I_0(\Lambda))$ to $(C_\infty, 0)$; thus part a) tells us that increasing σ for given Λ moves us from a point on the boundary of $\mathcal{F}_C$ to C_∞ .

Proposition 4 may come as a surprise to the reader: we have imposed quite a few assumptions on the way, and yet it seems that our model still cannot rule out any feasible covariation of types across partners! (Observing a $\hat{C}$ that is outside of $\mathcal{F}$ is impossible by construction.) Proposition 4 tells us that observing $\hat{C}$ in the interior of $\mathcal{F}_C$ rejects the

homogeneous model; but that any such $\hat{C}$ can be rationalized by adding the right amount of unobserved heterogeneity.

The interpretation of part c) is simplest when the matrix Λ is diagonal. With several dimensions for types, the optimal matching must sacrifice some covariation in one dimension to the benefit of some covariation in another. The implied sacrifice ratio, quite naturally, is exactly the ratio of the assorting weights along these dimensions. Take for instance the homogeneous case with only two characteristics, and set $\Lambda_{11} = 1$ and $\Lambda_{22} = \varepsilon$. Then the function $\varepsilon \rightarrow C^{11}(1, \varepsilon)$ is decreasing, and the function $\varepsilon \rightarrow C^{22}(1, \varepsilon)$ is increasing. Therefore, when one puts more weight on the second dimension, the covariation of the characteristics in the second dimension increases, while the covariation on the first dimension decreases. Quite intuitively, in the limit where all the weights are put on one dimension, the classical Beckerian theory of positive assortative matching obtains.

More precisely, Carlier, Galichon, and Santambrogio (2008) have shown in an r -type homogeneous model that when $\Lambda_{11} = 1$ and $\Lambda_{jj} \rightarrow 0$ for $j \geq 2$, if $\Pi^*(\Lambda)$ is the Λ -optimal matching, and $(X, Y) \sim \Pi^*(\Lambda)$, then the joint distribution of the first characteristics (X^1, Y^1) converges towards the maximally correlated distribution. Equivalently, X^1 and Y^1 become comonotonic in the limit, just as in classical positive assortative matching.

3.3 Uniqueness and purity

Uniqueness. As mentioned earlier, the boundary of $\mathcal{F}_c$ has kinks when types are discrete. In the homogeneous model ($\sigma = 0$), the optimal matching is pure for almost all values of Λ . Start from such a value Λ_0 . A small change in the value of Λ will not change the optimal matching Π_0 , or the covariations it generates. Pick some direction in Λ -space and move further away from Λ_0 . At some point Λ_1 , the optimal matching will change to a different pure matching, say Π_1 ; but this new pure matching will vary with the direction we used to move away from Λ_0 . This is what generates kinks. Note also that in Λ_1 , any matching that is a convex combination of Π_0 and of Π_1 is also optimal. So kinks are related to non-uniqueness of the optimal matching. More formally:

Proposition 5 *Assume (DD): the distributions P and Q are discrete. Then feasible set $\mathcal{F}_c$ is a polytope with a finite number of vertices that correspond to pure matchings.*

When there is enough observed heterogeneity, the optimal matching is unique. Indeed, Decker, Stephens, and McCann (2009) have shown that when the total heterogeneity σ is large enough (so that I is small enough), the solution to Eq. (4.1) is unique.

Purity. A matching is pure if a given type of man cannot be matched to more than one type of women and conversely. Intuition suggests that given sufficient heterogeneity, the optimal matching will not be pure, and its probability weights will react to even small changes in Λ . In fact, we have an even stronger result: even tiny levels of heterogeneity will make the optimal matching impure. To see this, reason by contradiction: take a $\sigma > 0$ and any given Λ for which the optimal matching Π is pure. The objective function is

$$\sum_{x,y} \left(\pi(x,y) \Lambda \cdot \phi^k(x,y) - \sigma \pi(x,y) \log \frac{\pi(x,y)}{p(x)q(y)} \right).$$

Note that the derivative with respect to any $\pi(x,y)$ is infinite in $\pi(x,y) = 0$ and is finite anywhere else. Since Π is pure, for any x there is only one y for which $\pi(x,y)$ is nonzero. Subtract a positive ε from each such $\pi(x,y)$, and spread it over all zero elements. The new joint distribution is still a feasible matching, and the gain in social surplus from formerly zero probabilities outweighs the loss from other matches. Therefore a pure matching cannot be optimal.

4 Identification

The results in the previous sections give a very useful description of the optimal matchings, and they show that σ_1 and σ_2 cannot be identified separately. On the other hand, we have not provided a proof of identification of the remaining parameters yet. We now set out to do so make use for identification purposes of the geometrical interpretation of the matching problem when the observable surplus is a linear combination of known basis functions—this is assumption (SLOI), which we impose throughout this section.

4.1 Nonparametric identification

Remember that given assumptions (O) and (S), there exist two functions $U(x, y) + V(x, y) = \Phi(x, y)$ such that the optimal matching obtains when man $\tilde{x}$ maximizes $U(x, y) + \chi(\tilde{x}, y)$ over y and woman $\tilde{y}$ maximizes $V(x, y) + \xi(\tilde{y}, x)$ over x . Now if π is the observable component of an optimal matching, it was showed in Section 2 that given assumption (GUI),

$$\begin{aligned} U(x, y) &= \sigma_1 \log \pi(x, y) + \sigma_1 \log \frac{n_1(x)}{p(x)}; \text{ and similarly,} \\ V(x, y) &= \sigma_2 \log \pi(x, y) + \sigma_2 \log \frac{n_2(y)}{q(y)}. \end{aligned}$$

Now U and V depend on θ and are not easy to characterize as we will see; but we know that they sum up to Φ , so that

$$\Phi(x, y) = \sigma \log \pi(x, y) + \sigma_1 (\log n_1(x) - \log p(x)) + \sigma_2 (\log n_2(y) - \log q(y)).$$

In this formula n_1 and n_2 still depend on θ in a complex way; but they only appear in terms that depend only on characteristics of one partner. This means that the surplus function Φ is identified up to an additive function of the form $a(x) + b(y)$.

To state this more formally, define the *cross-difference operator* as

$$\Delta_2 F(x, y; x', y') = (F(x', y') - F(x', y)) - (F(x', y) - F(x, y)),$$

for any function F of (x, y) . Then we have:

Theorem 4 (Cross-differences are identified up to scale) *Assume (O), (S), (GUI) and (DD). For $\theta = (\Lambda, \sigma_1, \sigma_2)$ with $\sigma = \sigma_1 + \sigma_2 > 0$, one has:*

(i) *There exists a unique optimal observable matching π which maximizes the social welfare (2.3).*

(ii) *There exist three vectors $\pi(x, y)$, $u(x)$ and $v(y)$, and a constant c normalized by $E_p[u(X)] = E_q[v(Y)] = 0$, which are unique solutions to the following system*

$$\begin{cases} \pi(x, y) = p(x) q(y) \exp\left(\frac{\Phi(x, y) - u(x) - v(y) - c}{\sigma}\right), \\ \pi \in \mathcal{M}(P, Q). \end{cases} \quad (4.1)$$

Further, the constant c so defined coincides with the value of the social welfare $c = \mathcal{W}$.

(iii) The probability π defined in (ii) coincides with the optimal matching solution of (2.3).

This result expresses that by adjusting the functions u and v at the right level, one manages to satisfy the “budget constraint” that the matching has the right marginals distributions $\pi \in \mathcal{M}(P, Q)$: hence, these functions u and v can be interpreted as “shadow prices” of men and women’s observable characteristics.

Theorem 4 has another consequence: the complementarity of dimensions i of the observable types of the partners in (x, y) can be tested directly on $\log \pi$, since $\Delta_2 \log \pi$ and $\Delta_2 \Phi$ have the same sign. Moreover, the relative strengths of complementarities along dimensions i and j at a point (x, y) can be estimated by evaluating $\Delta_2 \log \pi$ for values of (x', y') that differ from (x, y) along these dimensions.

Theorem 4 immediately gives us an estimator of the observable joint surplus function Φ , up to additive functions of x and of y . But adding any combination $a(x) + b(y)$ to the joint surplus does not change the optimal matching, as long as we are determined not to have singles—as we assume throughout the paper; and the positive scale factor σ is irrelevant. So for instance $\log \hat{\pi}$ is a perfectly good estimator of Φ if $\hat{\pi}$ consistently estimates π .

When we add a parametric structure under (SLOI), Theorem 4 also gives us an estimator of the assorting weights Λ and the total heterogeneity σ^5 . In fact, the cross-difference operator is linear and so under (SLOI),

$$\Delta_2 \log \pi = \frac{\Delta_2 \Phi}{\sigma} = \sum_{k=1}^K \frac{\Lambda_k}{\sigma} \Delta_2 \phi^k;$$

if the cross-differences of the ϕ^k are linearly independent, then observing π gives us the Λ ’s (along with overidentifying restrictions.) This is a very weak requirement; having linearly dependent basis functions would indeed be a modelling mistake.

This can be very simple in practice; to illustrate, take the diagonal version of the (ER)

⁵Recall that σ_1 and σ_2 are not separately identified.

example. Then if in (x_1, y_1) man and woman are both dropouts, keeping their income classes unchanged and moving them to graduate level in (x'_1, y'_1) generates

$$\Delta_2 \Phi(x_1, y_1; x'_1, y'_1) = \Lambda_1^E - \Lambda_0^E.$$

On the other hand, taking man and woman to have different education levels in (x_2, y_2) and swapping their educations to create (x'_2, y'_2) (again keeping income classes fixed) generates

$$\Delta_2 \Phi(x_2, y_2; x'_2, y'_2) = -\Lambda_1^E - \Lambda_0^E.$$

Therefore we obtain for instance

$$\frac{\Lambda_1^E}{\Lambda_0^E} = \frac{\Delta_2 \log \pi(x_2, y_2; x'_2, y'_2) - \Delta_2 \log \pi(x_1, y_1; x'_1, y'_1)}{\Delta_2 \log \pi(x_1, y_1; x'_1, y'_1) + \Delta_2 \log \pi(x_2, y_2; x'_2, y'_2)},$$

which is readily estimated from the observed matching.

These results are reminiscent of those in Fox (2009), although we obtained them under quite a different set of assumptions: we do not use variation across subpopulations, neither does his rank-order condition apply to our model. Note also that when specialized to one-dimensional types, our result yields that of Siow (2009) on testing complementarity of the surplus function by examining log-supermodularity of the match distribution.

4.2 Parametric identification

Our parametric identification strategy will be either based on the knowledge of the matching summaries $(\hat{C}, \hat{I})$, which are the sufficient statistics for our model, or of just the covariation $\hat{C}$, with the assumption that $(\hat{C}, \hat{I})$ lies on the efficient frontier, that is $\hat{I} = I_r(\hat{C})$. In either cases, positive homogeneity imposes the need for a normalization of the parameter $(\hat{\Lambda}, \hat{\sigma})$.

4.2.1 The normalization rule

Once again, θ is only identified up to a positive scale factor. Take (SLOI) for instance: Λ was only used to specify the objective function, and so it can be multiplied by any positive

constant without any side-effect. In particular, $\mathcal{W}(t\Lambda, t\sigma) = t\mathcal{W}(\Lambda, \sigma)$ for $t \geq 0$. Therefore it is quite clear that (Λ, σ) cannot be identified without fixing some normalisation. So we normalize (Λ, σ) by the choice

$$\textbf{Normalization convention: } \sigma I(\Lambda, \sigma) = 1, \quad (4.2)$$

where as we recall, $I(\Lambda, \sigma) = -\frac{\partial \mathcal{W}(\Lambda, \sigma)}{\partial \sigma}$.

Our general approach will be to identify the parameter value $(\hat{\lambda}, 1)$, and then rescale

$$\hat{\Lambda} = \frac{\hat{\lambda}}{I(\hat{\lambda}, 1)}, \quad \hat{\sigma} = \frac{1}{I(\hat{\lambda}, 1)}.$$

4.2.2 Identification of σ

Note that $I_r(C(\lambda, 1)) = I(\lambda, 1)$, therefore, because of the normalization convention, σ is identified by

$$\hat{\sigma} = \frac{1}{I_r(\hat{C})}. \quad (4.3)$$

4.2.3 Identification of Λ

As described above, we look for identifying $\hat{\lambda}$ among the parameters of the form $(\lambda, 1)$. Remember

$$I_r(C) = \sup_{\lambda} \{\lambda \cdot C - \mathcal{W}(\lambda, 1)\}$$

so by the envelope theorem, $\hat{\lambda} = \frac{\partial I_r(\hat{C})}{\partial C}$ is such that $C(\hat{\lambda}, 1) = \frac{\partial \mathcal{W}(\hat{\lambda}, 1)}{\partial \lambda}$. Hence, Λ is identified by

$$\hat{\Lambda} = \frac{1}{I_r(\hat{C})} \frac{\partial I_r(\hat{C})}{\partial C}. \quad (4.4)$$

4.3 Comparative statics

We define the *best additive projector* $\mathcal{P}h$ of a vector $h(x, y)$ as

$$\mathcal{P}h(x, y) = f(x) + g(y)$$

where f and g minimize

$$E_{\pi} \left[(h(X, Y) - E[h(X, Y)] - f(X) - g(Y))^2 \right].$$

We have immediately that $E_P[f(X)] = 0$ and $E_Q[g(Y)] = 0$, and introducing the residue ε

$$\varepsilon(X, Y) = h(X, Y) - E[h(X, Y)] - f(X) - g(Y)$$

we get $E[\varepsilon(X, Y) | X] = 0$ and $E[\varepsilon(X, Y) | Y] = 0$. The decomposition

$$h(X, Y) = E[h(X, Y)] + f(X) + g(Y) + \varepsilon(X, Y)$$

is the *two-way ANOVA* decomposition of $h(X, Y)$. The following proposition will be the fundamental tool for inference. It expresses that the projection residue in the two-way ANOVA decomposition of ϕ^k is the score function $\sigma \frac{\partial \log \pi}{\partial \Lambda_k}$.

Proposition 6 (Score function) *Under (O), (S), (GUI), and (SLOI), the score function is given by*

$$\frac{\partial \log \pi}{\partial \Lambda_k}(x, y) = \frac{\phi^k(x, y) - \mathcal{P}\phi^k(x, y) - E[\phi^k(X, Y)]}{\sigma},$$

that is $\frac{\partial u(x)}{\partial \Lambda_k} + \frac{\partial v(y)}{\partial \Lambda_k} = \mathcal{P}\phi^k(x, y)$, where u and v are solution to Equation (4.1).

As a result, we get an expression for the computation of the Hessian of the social welfare function at fixed σ .

Proposition 7 (Fisher information matrix) *Under (O), (S), (GUI), and (SLOI), whenever $\mathcal{W}(\Lambda, \sigma)$ is derivable, we get*

$$\frac{\partial^2 \mathcal{W}(\Lambda, \sigma)}{\partial \Lambda_k \partial \Lambda_l} = \sigma \mathcal{I}^{kl}(\theta),$$

where

$$\mathcal{I}^{kl}(\theta) := E \left[\frac{\partial \log \pi}{\partial \Lambda_k}(X, Y) \frac{\partial \log \pi}{\partial \Lambda_l}(X, Y) \right]$$

is the Fisher information matrix. Further,

$$\mathcal{I}^{kl}(\theta) := \frac{\text{cov}(\phi^k(X, Y), \phi^l(X, Y)) - \text{cov}(\mathcal{P}\phi^k(X, Y), \mathcal{P}\phi^l(X, Y))}{\sigma^2}. \quad (4.5)$$

5 Inference

We now turn to the problem of inference. Our data will consist of matched characteristics of N pairs $\{(x_1, y_1), \dots, (x_N, y_N)\}$, and our null hypothesis is that they were generated by an optimal matching consistent with assumptions (O), (S), (GUI), and (SLOI). Given a proposed specification for the basis functions ϕ^k , and our estimates of the marginal distributions of types $\hat{P}_N$ and $\hat{Q}_N$, we would therefore like to infer the values of Λ and σ which come closest to rationalizing the observed matching. We use our theory to answer two questions:

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

The primary object of our investigation will be the empirical moments of ϕ^k ,

$$\hat{C}_N^k = \frac{1}{N} \sum_{n=1}^N \phi^k(x_n, y_n).$$

Let C^k denote the expectation of $\phi^k(X, Y)$ under the joint distribution Π of (X, Y) . Standard asymptotic theory of the empirical process (van der Vaart (1998)) implies the convergence in distribution

$$\sqrt{N} \left(\hat{C}_N^k - C^k \right) \Rightarrow \xi^k$$

where $\xi^k = \int \phi^k(x, y) dG(x, y)$, G being a Π -Brownian bridge. In particular,

$$\text{cov} \left(\sqrt{N} \left(\hat{C}_N^k - C^k \right), \sqrt{N} \left(\hat{C}_N^l - C^l \right) \right) = \text{cov}_\Pi \left(\phi^k(X, Y), \phi^l(X, Y) \right)$$

for all $1 \leq k, l \leq K$.

We shall call $\mathcal{W}_N(\theta)$ the value of the social surplus at parameter θ obtained with the empirical distributions of observable types P_N and Q_N .

Normalization. Recall that because of positive homogeneity, models $\theta = (\Lambda, \sigma)$ and $t\theta = (t\Lambda, t\sigma)$ are observationally indistinguishable. Just as in the previous section, we impose the normalization convention $\sigma I(\Lambda, \sigma) = 1$. When we describe estimators below, we first compute an estimator of the assorting weights Λ for total heterogeneity $\sigma = 1$; we denote it $\hat{\lambda}_N$. We then shall get an estimator of the mutual information $\hat{I}_N$. To obtain the normalized estimator in each case, the reader should divide the vector $(\hat{\lambda}_N, 1)$ by the scalar $\hat{I}$.

The results we obtained in sections 2 and 4 suggest two estimation strategies, which we will now define and compare.

5.1 Nonparametric inference

Theorem 4 and its corollary immediately suggest a very simple nonparametric approach. In this discrete case, a nonparametric estimator $\hat{\pi}_N(x, y)$ is readily obtained, by counting the proportion of matches between a man of characteristics x and a woman of characteristics y . We could pick arbitrary functions $a(x)$ and $b(y)$ and define

$$\hat{\Phi}_N(x, y) = \log \hat{\pi}_N(x, y) + a(x) + b(y),$$

without any reference to basis functions—imposing $\sigma = 1$ on the way. Then if we further assume (SLOI) with basis functions ϕ^k , we can apply minimum-distance techniques to recover an estimator $\hat{\lambda}_N^{SP}$, which minimizes some norm

$$\|\hat{\Phi}_N(x, y) - \lambda \cdot \phi(x, y)\|.$$

Note that as usual, the minimum value of the norm allows us to construct a test statistic for the hypothesis that Φ is a linear combination of the ϕ^k .

More generally, we know that under (O), (S) and (GUI) only,

$$\log \pi = \frac{\Phi}{\sigma};$$

thus a nonparametric estimate $\hat{\pi}_N$ can be used as a heuristic device to decide on a set of basis functions, and/or to test for the adequacy of such a set.

We now turn to parametric estimators.

5.2 Parametric inference: The Moment Matching Estimator

Our second estimator is based solely on the statistics of the matching covariations $\hat{C}$. It rests on identification of Λ provided by Eq. (4.4). Therefore $\hat{\lambda}$ is taken as a maximizer of

$$\Lambda \cdot \hat{C} - \mathcal{W}_N(\Lambda, 1) \quad (5.1)$$

over all possible Λ . This being a strictly concave function, its minimizer is unique; further efficient computation is available. Letting $\hat{I}$ the value of expression (5.1) at the optimal value $\hat{\lambda}$, we obtain the **Moment Matching (MM) estimator**, denoted $\hat{\Lambda}^{MM}$ and $\hat{\sigma}^{MM}$, by setting

$$\hat{\Lambda}^{MM} = \frac{\hat{\lambda}}{\hat{I}}, \quad \hat{\sigma}^{MM} = \frac{1}{\hat{I}}.$$

Now if our data was generated by an optimal matching Π for parameters $(\hat{\Lambda}^{MM}, \hat{\sigma}^{MM})$, the empirical covariations $\hat{C}_N$ would coincide with the optimal correlations $C(\hat{\Lambda}^{MM}, \hat{\sigma}^{MM})$. By construction, the MM estimator is the value of assorting weights λ such that the predicted covariations coincide with the observed covariations. The Moment matching estimator is consistent and asymptotically Gaussian, and

Theorem 5 *Under (O), (S), (GUI) and (SLOI),*

$$\sqrt{N}(\hat{\lambda}_N - \lambda) \implies \mathcal{I}^{-1}\xi$$

where ξ is the Brownian bridge characterized at the beginning of this section and the matrix $\mathcal{I}^{kl}$ is the Fisher information matrix expressed above in (4.5). In particular, the MM estimator is asymptotically efficient.

6 Computational issues

With the exception of the semiparametric estimator (SP), our inferential methods require solving for the optimal matching for potentially large populations, and a large number of

parameter vectors during optimization. This may seem to be a forbidding task: there exist well-known algorithms to find an optimal matching, and they are reasonably fast; but with large populations the required computer resources may still be large.

Fortunately, it turns out that introducing (our type of) heterogeneity actually makes computing optimal matchings much simpler; this is a boon for the ML and MM estimators⁶.

To see this, choose a parameter vector $\theta = (\Phi, \sigma)$ and return to the characterization of optimal matchings in equation 2.3, in the continuous case (CD) for simplicity. Dividing by σ and taking the logarithm, optimal matchings can also be obtained by solving the following minimization program:

$$\min_{\Pi \in \mathcal{M}(P, Q)} \sum_{x, y} \pi(x, y) \log \frac{\pi(x, y)}{p(x)q(y) \exp(\Phi(x, y)/\sigma)}.$$

Now define a set of probabilities r by

$$r(x, y) = \frac{p(x)q(y) \exp(\Phi(x, y)/\sigma)}{\sum_{x, y} p(x)q(y) \exp(\Phi(x, y)/\sigma)};$$

and note that given any choice of parameters θ and known marginals (p, q) , the probability r itself is known.

Determining the optimal matchings therefore boils down to finding the joint probabilities π with known marginals p and q which minimize the Kullback-Leibler distance to r :

$$\sum_{x, y} \pi(x, y) \log \frac{\pi(x, y)}{r(x, y)}. \quad (6.1)$$

Equivalently, we are looking for the Kullback-Leibler projection of r on $\mathcal{M}(P, Q)$.

This is a well-known problem in various fields, and algorithms to solve it have been around for a long time. National accountants, for instance, use RAS algorithms to fill cells of a two-dimensional table whose margins are known; here the choice of r reflects prior notions of the correlations of the two dimensions of the table. These RAS algorithms belong to a family called Iterative Projection Fitting Procedures (IPFP). They are very fast, and

⁶The BP estimator is designed for the homogeneous case and so the following does not apply to it.

are guaranteed to converge under weak conditions. We only describe the application of IPFP to our model here; we direct the reader to Rüschendorf (1995) for more information.

The intuition of equation 6.1 is quite clear: the random matching, which is optimal when σ is very large, has $\pi(x, y) = p(x)q(y)$. For smaller σ 's the probability of a match between x and y must increase with the surplus it creates, $\Phi(x, y)$; and given our assumption (GUI) on the distribution of unobserved heterogeneity, it should not come as a surprise that the corresponding factor is multiplicative and exponential.

To describe the algorithm, we split π into⁷

$$\pi(x, y) = r(x, y) \exp(-(u(x) + v(y))/\sigma).$$

The functions u and v of course will only be determined up to a common constant. The algorithm iterates over values (u^k, v^k) . We start from $u^0 \equiv -\sigma \log p$ and $v^0 \equiv 0$. Then at step $(k + 1)$ we compute

$$\exp(-v^{k+1}(y)/\sigma) = \frac{q(y)}{\sum_x r(x, y) \exp(-u^k(x)/\sigma)}$$

and

$$\exp(-u^{k+1}(x)/\sigma) = \frac{p(x)}{\sum_y r(x, y) \exp(-v^{k+1}(y)/\sigma)}.$$

Two remarks are in order here: first, we could just as well start from $u^0 \equiv 0$ and $v^0 = -\sigma \log q$ and modify the iteration formulæ accordingly. Second and just as in other Gauss-Seidel algorithms, it is important to update one component based on the other updated component: the right-hand sides have u^k and v^{k+1} .

If (u, v) is a fixed point of the algorithm, then

$$\frac{\pi(x, y)}{p(x)q(y)} = \exp\left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma}\right).$$

Comparing this formula to Theorem 4 shows the benefit of this reparameterization, since $u(x)$ and $v(y)$ have a simple interpretation: they represent (up to a common additive

⁷It can be shown that at the optimum $\pi(x, y) = 0$ where $r(x, y) = 0$.

constant) the expected utilities of a man of observable characteristics x and of a woman of observable characteristics y . This can be seen by checking, for instance, that

$$E(\max U(X, Y) | X = x) = \sigma_1 \log n_1(x).$$

Thus the IPFP algorithm gives us not only the optimal matching, but also these expected utilities.

The simplification does not stop there. In fact, given data on N couples, the marginal p assigns $1/N$ probability to each of $(x_1, \dots, x_N)$, and similarly for women. Define a matrix Ψ by $\Psi_{ij} = \exp(\Phi(x_i, y_j)/\sigma)$, and vectors $a_i^k = \exp(-u^k(x_i)/\sigma)$, $b_j^k = \exp(-v^k(y_j)/\sigma)$. Then we end up with the shockingly simple and inexpensive formulæ:

$$b^{k+1} = \frac{N}{\Psi' a^k} \quad \text{and} \quad a^{k+1} = \frac{N}{\Psi b^{k+1}}.$$

7 Possible extensions and concluding remarks

Our theory so far relies on several strong assumptions. Some of them are easy to relax; we discuss three of them, before turning to potential extensions.

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.

From the point of view of statistical inference, we may lose some efficiency in doing so; we note here that when the unobserved heterogeneity in preferences over partners is separable from the utility of marriage itself, our method does not incur any efficiency loss.

Non-bipartite matching. Bipartite matching refers to the fact that each individual

is exogenously assigned in one category—in our terminology, husband or wife. Our analysis in fact is very easy to extend so as to incorporate same-sex unions, and thus to rationalize endogamy in the gender dimension.

To do so, we just need to add one (observed) characteristic, in the form of gender. If for instance gender becomes the first dimension of the characteristics vector, then the observed surplus has an assorting weight $\Lambda_{11} < 0$ that reflects the more typical preference for the opposite sex; while heterogenous preferences χ and η will automatically take into account the dispersion of individual preference for same-sex unions.

Continuous distributions. While we have assumed discrete characteristics, we expect the main thrust of our arguments to carry over to the case where the distributions of the characteristics are continuous. We are working on such an extension; this will require adapting the (GUI) assumption to one that is better-suited to continuous choice.

Revealed Preferences. As mentioned in the section on the Boundary Projection estimator, the Lagrange multiplier e is known in the theory of revealed preferences as Afriat’s efficiency index. The analogy in fact goes deeper. Recall the basic theorem on revealed preferences:

Proposition 8 (Afriat) *The following conditions are equivalent:*

(i) *The observed quantity-price vectors $(x_k, p_k)_{k=1}^N$ are consistent with maximization of a single utility function;*

(ii) *There exist scalars $\lambda_k < 0$, $k = 1, \dots, N$ such that*

$$\sum_{k=1}^N \lambda_k p_k \cdot x_{\sigma(k)}$$

is maximized over $\sigma \in \mathfrak{S}_N$ when $\sigma(j) = j$ for all j .

This is reminiscent of a multidimensional matching problem in which prices p correspond to the characteristics x of men, consumptions q to those of women y , and there is no unobserved heterogeneity. We are currently exploring this analogy.

Screening. In the theory of screening, a “type” θ refers to a set of individual characteristics that are privately observed. Assume that utilities are additively separable in transfers, with

$$u(q, \theta) - t \text{ for an agent of type } \theta$$

and

$$W(q) + t \text{ for the principal.}$$

Then given quantity-transfer pairs $(q_k, t_k)_{k=1}^N$ that presumably correspond to different types, it can be shown that

$$\sum_{k=1}^N u(q_k, \theta_{\sigma(k)})$$

is maximized over $\sigma \in \mathfrak{S}_N$ when $\sigma(j) = j$ for all j .

This again suggests that our methods may help in estimating screening models.

A Facts from Convex Analysis

A.1 Basic results

We only sum up here the concepts we actually use in the paper; we refer the reader to Hiriart-Urrut and Lemaréchal (2001) for a thorough exposition of the topic.

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, denoted $\text{cch}(Y)$.

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 $\partial S_Y(0) = \text{cch}(Y)$.

A point in Y is an *extreme point* if it does not belong in any open line segment joining two points of 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

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

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

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

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

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

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 $\partial u(x) = \{\nabla u(x)\}$.

A.2 Generalized Convexity

In order to make the paper self-contained, we present basic results on the theory of *generalized convexity*, sometimes called the theory of *c-convex functions*. This theory extends many results from convex analysis and, in particular, duality results, to a much more general setting. We refer to Villani (2009), p. 54–57 (or Villani (2003), pp. 86–87) for a detailed account⁸.

Let ω be a function from the product of two sets $\mathcal{X} \times \mathcal{Y}$ to $[-\infty, +\infty)$.

Definition 1 *Consider any function $\psi : \mathcal{X} \rightarrow (-\infty, +\infty]$. Its generalized Legendre transform $\psi^\perp : \mathcal{X} \rightarrow [-\infty, +\infty)$ is defined by*

$$\psi^\perp(y) = \inf_{x \in \mathcal{X}} \{\psi(x) - \omega(x, y)\}.$$

Conversely, take any function $\zeta : \mathcal{Y} \rightarrow [-\infty, +\infty)$; then its generalized Legendre transform $\zeta^\top : \mathcal{X} \rightarrow (-\infty, +\infty]$ is defined by

$$\zeta^\top(x) = \sup_{y \in \mathcal{Y}} \{\zeta(y) + \omega(x, y)\}.$$

A function ψ is called ω -convex if it is not identically $+\infty$ and if there exists $\zeta : \mathcal{Y} \rightarrow [-\infty, +\infty]$ such that

$$\psi = \zeta^\top.$$

Recall that the usual Legendre transform is defined as

$$\psi^*(y) = \inf_{x \in \mathcal{X}} \{\psi(x) - x \cdot y\};$$

thus it coincides with the generalized Legendre transform when ω is bilinear, and then ω -convexity boils down to standard convexity.

Our analysis rests on the following fundamental result, which generalizes standard convex analysis.

⁸A cautionary remark is in order here: the sign conventions vary in the literature, so our own choices may differ from those of any given author.

Proposition 9 For every function $\psi : \mathcal{X} \rightarrow (-\infty, +\infty]$,

$$\psi^{\perp\top} \leq \psi$$

with equality if and only if ψ is ω -convex.

Proof Take any $x \in \mathcal{X}$; then

$$\psi^{\perp\top}(x) = \sup_{y \in \mathcal{Y}} \inf_{x' \in \mathcal{X}} \{ \psi(x') - \omega(x', y) + \omega(x, y) \};$$

taking $x' = x$ shows that $\psi^{\perp\top}(x) \leq \psi(x)$.

Conversely, if $\psi^{\perp\top} = \psi$ then $\psi(x) = \zeta^\top(x)$, with $\zeta = \psi^\perp$. But for any function ζ , the triple transform $\zeta^{\top\perp\top}$ coincides with $\zeta^\top$. To see this, write

$$\zeta^{\top\perp\top}(x) = \sup_{y \in \mathcal{Y}} \inf_{x' \in \mathcal{X}} \sup_{y' \in \mathcal{Y}} \{ \zeta(y') + \omega(x', y') - \omega(x', y) + \omega(x, y) \}.$$

Now for all x and y ,

$$\inf_{x' \in \mathcal{X}} \sup_{y' \in \mathcal{Y}} \{ \zeta(y') + \omega(x', y') - \omega(x', y) \} \geq \zeta(y)$$

as is easily seen by taking $y' = y$; therefore

$$\zeta^{\top\perp\top}(x) \geq \sup_{y \in \mathcal{Y}} \{ \zeta(y) + \omega(x, y) \} = \zeta^\top(x).$$

Applying this to the ζ such that $\psi = \zeta^\top$ concludes the proof.

QED.

B Proofs

B.1 Proof of Theorem 1

In order to prove Theorem 1, some preparation is needed. Remember our shorthand notation $\tilde{x} = (x, \varepsilon)$, and $\tilde{y} = (y, \eta)$. For any function $\tilde{u}(x, \varepsilon)$, fix x and use the theory of generalized convexity briefly recalled in Appendix (A.2) to define

$$\tilde{u}^\perp(x, y) = \inf_{\varepsilon} \{ \tilde{u}(x, \varepsilon) - \chi((x, \varepsilon), y) \}$$

the *generalized Legendre transform* of $\tilde{u}(x, \cdot)$ with respect to the partial surplus function $\chi((x, \cdot), \cdot)$. We define in the same manner

$$\tilde{v}^\perp(x, y) = \inf_{\eta} \{ \tilde{v}(y, \eta) - \xi(x, (y, \eta)) \}.$$

Similarly, for two functions $U(x, y)$ and $V(x, y)$, we define

$$\begin{aligned} U^\top(x, \varepsilon) &: = \sup_y \{ U(x, y) + \chi((x, \varepsilon), y) \} \\ V^\top(y, \eta) &: = \sup_x \{ V(x, y) + \xi(x, (y, \eta)) \}. \end{aligned}$$

Lemma 1 *Let A be the set of pairs of functions (U, V) such that*

$$\forall x, y, \quad U(x, y) + V(x, y) \geq \Phi(x, y).$$

Then

$$\mathcal{W} = \inf_{(U, V) \in A} \left\{ \int U^\top(\tilde{x}) d\tilde{P}(\tilde{x}) + \int V^\top(\tilde{y}) d\tilde{Q}(\tilde{y}) \right\}.$$

Proof of Lemma 1 By the Kantorovich duality theorem (Villani (2009) Theorem 5.10),

$$\mathcal{W} = \sup_{\tilde{\pi} \in \mathcal{M}(P, Q)} \int \tilde{\Phi}(\tilde{x}, \tilde{y}) d\tilde{\pi}(\tilde{x}, \tilde{y}) = \inf_{(\tilde{u}, \tilde{v}) \in \tilde{A}} \left\{ \int \tilde{u}(\tilde{x}) d\tilde{P}(\tilde{x}) + \int \tilde{v}(\tilde{y}) d\tilde{Q}(\tilde{y}) \right\}, \quad (\text{B.1})$$

where $\tilde{A}$ is the set of pairs of functions $(\tilde{u}, \tilde{v})$ such that

$$\forall \tilde{x}, \tilde{y}, \quad \tilde{u}(\tilde{x}) + \tilde{v}(\tilde{y}) \geq \tilde{\Phi}(\tilde{x}, \tilde{y}).$$

Note the following two facts about the right-hand side of this equality:

1. Since

$$\tilde{\Phi}(\tilde{x}, \tilde{y}) = \Phi(x, y) + \chi((x, \varepsilon), y) + \xi((y, \eta), x),$$

the infimum in (B.1) can be taken over the pair of functions $(\tilde{u}, \tilde{v})$ that satisfy

$$\tilde{u}(x, \varepsilon) \geq \sup_y \left\{ \Phi(x, y) + \chi((x, \varepsilon), y) + \sup_{\eta} [\xi((y, \eta), x) - \tilde{v}(y, \eta)] \right\},$$

or

$$\tilde{u}(\tilde{x}) \geq \sup_y \left\{ \Phi(x, y) + \chi((x, \varepsilon), y) - \tilde{v}^\perp(x, y) \right\}$$

At the optimum this must hold with equality. Going back to Definition 1, it follows that $\tilde{u}(x, \cdot)$ is $\chi((x, \cdot), \cdot)$ -convex for every x ; and using Proposition 9, we can substitute $\tilde{u}$ with $\tilde{u}^{\perp\top}$, that is:

$$\tilde{u}(x, \varepsilon) = \sup_y \left\{ \tilde{u}^{\perp}(x, y) + \chi((x, \varepsilon), y) \right\}.$$

Given a similar argument on $\tilde{v}$, the objective function can be rewritten as

$$\int \sup_y \left\{ \tilde{u}^{\perp}(x, y) + \chi((x, \varepsilon), y) \right\} d\tilde{P}(\tilde{x}) + \int \sup_x \left\{ \tilde{v}^{\perp}(x, y) + \xi(x, (y, \eta)) \right\} d\tilde{Q}(\tilde{y}).$$

2. Also note that the constraint of the minimization problem in (B.1) is also

$$\forall x, y, \quad \tilde{u}^{\top}(x, y) + \tilde{v}^{\perp}(x, y) \geq \Phi(x, y)$$

which follows directly from the fact that

$$\forall x, \varepsilon, y, \eta, \quad \tilde{u}(x, \varepsilon) - \chi((x, \varepsilon), y) + \tilde{v}(y, \eta) - \xi(x, (y, \eta)) \geq \Phi(x, y).$$

Now define

$$U(x, y) = \tilde{u}^{\perp}(x, y) \quad \text{and} \quad V(x, y) = \tilde{v}^{\perp}(x, y);$$

Given points 1. and 2. above, we can rewrite the value $\mathcal{W}$ as

$$\mathcal{W} = \inf_{(U, V) \in A} \left\{ \int U^{\top}(\tilde{x}) d\tilde{P}(\tilde{x}) + \int V^{\top}(\tilde{y}) d\tilde{Q}(\tilde{y}) \right\}.$$

QED.

We are now in a position to prove the theorem.

Proof of Theorem 1 Start by drawing two samples of size N of men and women from their population distributions P and Q ; we denote the corresponding values of the observed characteristics $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_n\}$. Call P_n and Q_n the corresponding sample distributions; e.g. P_n assigns a mass

$$p_{i,n} = \frac{1}{n} \sum_{j=1}^n \mathbf{1}(x_j = x_i^*)$$

to the value x_i^* of observable characteristics of men. The Law of Large Numbers implies that P_n and Q_n converge in distribution to P and Q , the population distributions of the observable types. Now we have for any possible x

$$\int U^\top(\tilde{x}) d\tilde{P}_n(\varepsilon|X=x) = \sum_{\substack{i=1,\dots,n \\ x_i=x}} \sup_{j=1,\dots,n} \{U(x_i, y_j) + \chi((x_i, \varepsilon_i), y_j)\} + o(1)$$

As N gets large enough, each of the possible values of observable characteristics of women y_t^* is included in the sample $\{y_1, \dots, y_n\}$; then the sup in the above expression runs over all such possible values $\{y_1^*, \dots, y_{T_y}^*\}$. But under (GUI), conditional on X the random variables $\chi((x, \varepsilon), y_t^*)$ are independent Gumbel random variables with scaling factor σ_1 , so we get

$$\frac{1}{\sigma_1} \int U^\top(\tilde{x}) d\tilde{P}_n(\varepsilon|X=x) = \log \sum_{t=1}^{T_y} \exp(U(x, y_t)/\sigma_1) + o_P(1)$$

hence, taking the limit and integrating over x ,

$$\int U^\top(\tilde{x}) d\tilde{P}(\tilde{x}) = \sigma_1 E_P \log \sum_y \exp(U(X, y)/\sigma_1)$$

and similarly

$$\int V^\top(\tilde{y}) d\tilde{Q}(\tilde{y}) = \sigma_2 E_Q \log \sum_x \exp(V(x, Y)/\sigma_2).$$

QED.

B.2 Proof of Theorem 2

Proof By theorem (1), we have

$$\mathcal{W}_N = \inf_{U(x,y)+V(x,y) \geq \Phi(x,y) \ \forall x,y} \left\{ \begin{array}{l} \sigma_1 \sum_x p(x) \log \left(\sum_y \exp(U(x, y)/\sigma_1) \right) \\ + \sigma_2 \sum_y q(y) \log \left(\sum_x \exp(V(x, y)/\sigma_2) \right) \end{array} \right\}$$

for which we form the Lagrangian

$$\begin{aligned}\mathcal{W}_N &= \inf_{U(x,y), V(x,y)} \sup_{\pi(x,y) \geq 0} \left\{ \begin{aligned} &\sigma_1 \sum_x p(x) \log \left(\sum_y \exp(U(x,y)/\sigma_1) \right) \\ &+ \sigma_2 \sum_y q(y) \log \left(\sum_x \exp(V(x,y)/\sigma_2) \right) \\ &+ \sum_{x,y} \pi(x,y) (\Phi(x,y) - U(x,y) - V(x,y)) \end{aligned} \right\} \\ &= \sup_{\pi(x,y) \geq 0} \left\{ \sum_{x,y} \pi(x,y) \Phi(x,y) + \inf_{U(\cdot,\cdot)} F(U) + \inf_{V(\cdot,\cdot)} G(V) \right\}\end{aligned}$$

where

$$\begin{aligned}F(U) &= \sigma_1 \sum_x p(x) \log \left(\sum_y \exp(U(x,y)/\sigma_1) \right) - \sum_{x,y} \pi(x,y) U(x,y) \\ G(V) &= \sigma_2 \sum_y q(y) \log \left(\sum_x \exp(V(x,y)/\sigma_2) \right) - \sum_{x,y} \pi(x,y) V(x,y).\end{aligned}$$

Clearly, $U(\cdot, \cdot)$ and $V(\cdot, \cdot)$ in the inner minimization problems satisfy

$$\pi(x,y) = \frac{p(x) \exp(U(x,y)/\sigma_1)}{\sum_y \exp(U(x,y)/\sigma_1)} = \frac{q(y) \exp(V(x,y)/\sigma_2)}{\sum_x \exp(V(x,y)/\sigma_2)}; \quad (\text{B.2})$$

note that these equations imply that $\sum_y \pi(x,y) = p(x)$ and $\sum_x \pi(x,y) = q(y)$, so that $\pi \in \mathcal{M}(P, Q)$. Rearranging terms,

$$\mathcal{W}_N = \sup_{\pi \in \mathcal{M}(P, Q)} \left\{ \begin{aligned} &\sum_{x,y} \pi(x,y) \Phi(x,y) - (\sigma_1 + \sigma_2) \sum_{x,y} \pi(x,y) \log \pi(x,y) \\ &+ \sigma_1 \sum_x p(x) \log p(x) + \sigma_2 \sum_y q(y) \log q(y) \end{aligned} \right\}$$

and noticing that $\sum_{x,y} \pi(x,y) \log \pi(x,y) = D(\pi) - S(P) - S(Q)$ gives the desired result.

B.3 Proof of Theorem 3

Proof The result follows directly from the Kantorovich duality (cf. Villani (2009), Ch. 2); it can also be obtained by letting σ_1 and σ_2 tend to zero in Theorem 1, and noting that, as $\sigma_1, \sigma_2 \rightarrow 0$,

$$\begin{aligned}\sigma_1 E_P \left[\log \sum_y [\exp(U(X,y)/\sigma_1)] \right] &\rightarrow E_P \left[\max_y U(X,y) \right], \\ \sigma_2 E_Q \left[\log \sum_x [\exp(V(x,Y)/\sigma_2)] \right] &\rightarrow E_Q \left[\max_x U(x,Y) \right].\end{aligned}$$

B.4 Proof of Theorem 4

Proof (i) For $\sigma > 0$, the map $\pi \rightarrow \sum_{x,y} \pi(x,y) \Phi(x,y) - \sigma I(\pi)$ is strictly concave and finite, on the convex domain $\mathcal{M}(P, Q)$; thus there exists a unique $\pi \in \mathcal{M}(P, Q)$ maximizing (2.3).

(ii) Let B be the set of pairs of functions $(u(x), v(y))$ such that $\sum_x u(x) p(x) = \sum_y v(y) q(y) = 0$, and for $(u, v) \in B$, let Z be the partition

$$Z(u, v) := \sum_{x,y} p(x) q(y) \exp \left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma} \right).$$

Introduce

$$\begin{aligned} p_{u,v}(x) &: = \frac{\partial \log Z(u, v)}{\partial u(x)} = \frac{\sum_y p(x) q(y) \exp \left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma} \right)}{\sum_{x,y} p(x) q(y) \exp \left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma} \right)} \\ q_{u,v}(y) &: = \frac{\partial \log Z(u, v)}{\partial v(y)} = \frac{\sum_x p(x) q(y) \exp \left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma} \right)}{\sum_{x,y} p(x) q(y) \exp \left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma} \right)} \end{aligned}$$

as a result $p_{u,v}$ and $q_{u,v}$ are probability vectors. By the strict concavity of $\log Z$, there exists a unique vector $(u, v) \in B$ such that

$$p = p_{u,v}$$

$$q = q_{u,v}$$

and $\pi(x, y) = p(x) q(y) \exp \left(\frac{\Phi(x, y) - u(x) - v(y)}{\sigma} \right) \in \mathcal{M}(P, Q)$.

(iii) Let $\pi \in \mathcal{M}(P, Q)$ be the solution of (2.3). From Expression (B.2) in the proof of Theorem 2, we have that

$$\begin{aligned} \sigma_1 \log \pi(x, y) &= U(x, y) + \sigma_1 \log p(x) - \sigma_1 \log \left(\sum_y \exp(U(x, y) / \sigma_1) \right) \\ \sigma_2 \log \pi(x, y) &= V(x, y) + \sigma_2 \log q(y) - \sigma_2 \log \left(\sum_x \exp(V(x, y) / \sigma_2) \right) \end{aligned}$$

thus, summing up

$$\sigma \log \frac{\pi(x, y)}{p(x) q(y)} = \Phi(x, y) - u(x) - v(y) - c$$

where

$$\begin{aligned} u(x) &= \sigma_2 \log p(x) + \sigma_1 \log \left(\sum_y \exp(U(x, y) / \sigma_1) \right) + c_1 \\ v(y) &= \sigma_1 \log q(y) + \sigma_2 \log \left(\sum_x \exp(V(x, y) / \sigma_2) \right) + c_2 \\ c &= c_1 + c_2 \end{aligned}$$

and c_1 and c_2 are constant adjusted so that $(u, v) \in B$. Hence π is solution of equation (4.1). It follows immediately that $c = \mathcal{W}$.

B.5 Proof of Proposition 3

Proof a) The convexity of $\mathcal{W}$ follows from the fact that it is the supremum of expression which are linear with respect to θ .

b) As a result, by the envelope theorem, the subdifferential of $\mathcal{W}$ at θ is the set of $\{C(\Pi), -I(\Pi)\}$ such that $\Lambda C(\Pi) - \sigma I(\Pi) = \mathcal{W}(\Lambda, \sigma)$. When this set consists of a single point, $\mathcal{W}$ is differentiable at θ and

$$\frac{\partial \mathcal{W}}{\partial \Lambda_k}(\theta) = C^k(\Pi), \quad \frac{\partial \mathcal{W}}{\partial \sigma}(\theta) = -I(\Pi).$$

B.6 Proof of Proposition 1

Proof Non-emptiness is obvious. Now $\mathcal{F}_c$ is convex: Let $\hat{C}$ and $\tilde{C}$ be two feasible cross-product matrices in $\mathcal{F}_c$. We first show that for any $\alpha \in [0, 1]$, $\alpha \hat{C} + (1 - \alpha) \tilde{C}$ is in $\mathcal{F}_c$. By definition of $\mathcal{F}_c$, 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 $\alpha \hat{C} + (1 - \alpha) \tilde{C} \in \mathcal{F}_c$. Now we prove that $\mathcal{F}_c$ is closed: Let C_n be a sequence in $\mathcal{F}_c$ 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 $C \in \mathcal{F}_c$. Finally, $\mathcal{F}$ is a closed convex set as it is the upper graph of the function $I_r(C)$ defined in Eq. (3.1).

B.7 Proof of Proposition 2

Proof $\mathcal{R}$ is the reunion of the subgradients of $\mathcal{W}$ which was seen in Prop. 1 to be the support function of $\mathcal{F}$: hence $\mathcal{R}$ is the frontier of $\mathcal{F}$.

B.8 Proof of Proposition 3

Proof a) Positive homogeneity and convexity of degree one follows from the fact that $\mathcal{W}$ is the support function of $\mathcal{F}$. Strict convexity for $\sigma > 0$ follows from the strict convexity of $I(\pi)$. Part b) follows directly from the envelope theorem. Part c) results of $I_r(C)$ being the Legendre transform of $\mathcal{W}(\lambda, 1)$ which is strictly convex, hence it convex on $\mathcal{F}_c$, differentiable on its interior, and by the envelope theorem, $\frac{\partial I_r}{\partial C^k} = \frac{\Lambda_k}{\sigma}$.

B.9 Proof of Proposition 4

Proof a) The sets $\mathcal{R}_c(I)$ are extreme points of the sets $I_r^{-1}([0, I])$ which are closed convex sets. One has $I_r^{-1}(\{0\}) = \{C_\infty\}$ which corresponds to $\Pi = P \otimes Q$, and $I_r^{-1}([0, S(P) + S(Q)]) = \mathcal{F}_c$. Clearly, one has $\hat{C} \in \mathcal{R}_c(I_r(\hat{C}))$. Finally, the form $\sum_k \Lambda_k dC^k$ vanishes along $\mathcal{R}_c(I)$, so one has $\sum_k \Lambda_k dC^k = 0$, hence the result.

B.10 Proof of Proposition 6

Proof By equation (4.1), we have

$$\log \frac{\pi(x, y)}{p(x)q(y)} = \frac{\Phi(x, y) - u(x) - v(y) - c}{\sigma}$$

hence $\sigma \frac{\partial \log \pi}{\partial \Lambda_k}(x, y) = \phi^k(x, y) - \frac{\partial u(x)}{\partial \Lambda_k} - \frac{\partial v(y)}{\partial \Lambda_k} - \frac{\partial c}{\partial \Lambda_k}$. But we have that

$$\sum_x \frac{\partial \log \pi_\Lambda(x, y)}{\partial \Lambda_k} \pi_\Lambda(x, y) = \sum_x \frac{\partial \pi_{\Lambda(x, y)}}{\partial \Lambda_k} = \frac{\partial}{\partial \Lambda_k} \sum_x \pi_{\Lambda(x, y)} = \frac{\partial q(y)}{\partial \Lambda_k} = 0,$$

thus for all x and y ,

$$E \left[\frac{\partial \log \pi_\Lambda(X, Y)}{\partial \Lambda_k} | X = x \right] = 0 \text{ and } E \left[\frac{\partial \log \pi_\Lambda(X, Y)}{\partial \Lambda_k} | Y = y \right] = 0$$

hence $\frac{\partial \log \pi}{\partial \Lambda_k}(x, y) \in V^\circ$, while $\frac{\partial u(x)}{\partial \Lambda_k} + \frac{\partial v(y)}{\partial \Lambda_k} \in V^+$, therefore

$$\phi^k(x, y) = \sigma \frac{\partial \log \pi}{\partial \Lambda_k}(x, y) + \frac{\partial u(x)}{\partial \Lambda_k} + \frac{\partial v(y)}{\partial \Lambda_k} + E \left[\phi^k(X, Y) \right]$$

is the orthogonal decomposition of $\phi^k(x, y)$ on $V^\circ \oplus V^+ \oplus \mathbb{R}$, hence $\frac{\partial u(x)}{\partial \Lambda_k} + \frac{\partial v(y)}{\partial \Lambda_k} = \mathcal{P}\phi^k(x, y)$.

B.11 Proof of Proposition 7

Proof We have

$$\frac{\partial \mathcal{W}(\Lambda, \sigma)}{\partial \Lambda_l} = E \left[\phi^l(X, Y) \right]$$

hence

$$\frac{\partial^2 \mathcal{W}(\Lambda, \sigma)}{\partial \Lambda_k \partial \Lambda_l} = E \left[\phi^l(X, Y) \frac{\partial \log \pi}{\partial \Lambda_k}(X, Y) \right] = \sigma E \left[\frac{\partial \log \pi}{\partial \Lambda_k}(X, Y) \frac{\partial \log \pi}{\partial \Lambda_l}(X, Y) \right].$$

Further, by the orthogonality of V° and V^+ ,

$$\begin{aligned} \text{cov} \left(\phi^k(X, Y), \phi^l(X, Y) \right) &= \sigma^2 \text{cov} \left(\frac{\partial \log \pi}{\partial \Lambda_k}(X, Y), \frac{\partial \log \pi}{\partial \Lambda_l}(X, Y) \right) \\ &\quad + \text{cov} \left(\mathcal{P}\phi^k(X, Y), \mathcal{P}\phi^l(X, Y) \right) \end{aligned}$$

QED.

B.12 Proof of Theorem 5

Proof We have $\hat{\lambda}_N = \frac{\partial I_r}{\partial C}$, hence at first order $\hat{\lambda}_N - \lambda = D^2 I_r \cdot (\hat{C}_N - C) + o_P(1/\sqrt{N})$. But as I_r is the Legendre transform of $\mathcal{W}(\cdot, 1)$, it results that $D^2 I_r = (D^2 \mathcal{W}(\cdot, 1))^{-1} = \mathcal{I}^{-1}$ by Proposition 7.

C Connections to Statistical physics

There is in fact, a very close parallel between our theory and Statistical physics and Thermodynamics. We refer to Parisi (1988) for more on Statistical physics, and to Mézard and Montanari (2009) for connection with Information theory. To give hints to the parallel, let

us just mention that the social welfare $\mathcal{W}$ is the analog of a *total energy*; the term $\sum \lambda_k C^k$ is the analog of an *internal energy*; $I(\pi)$ is the analog of an *entropy*; the parameter σ is the analog of a *temperature*. A pure matching is the equivalent of a *solid state*; the points of nondifferentiability of $\mathcal{W}$ are analog to *critical points*.

Note that equation 4.1 is known in the mathematical physics literature as the Schrödinger-Bernstein equation, cf. Rüschemdorf and Thomsen (1998) and references therein. It was first studied by Erwin Schrödinger as part of his research program in time irreversibility in Statistical Physics. Interestingly, it also bears some connections with the better-known “Schrödinger equation” in Quantum mechanics of the same inventor. In fact, as discovered by Zambrini, a dynamic formulation of this equation is the Euclidian Schrödinger equation which arises in Ed Nelson’s formulation of “Stochastic Mechanics,” an Euclidian analog of quantum mechanics. For more on this topic, see Parisi (1988), Chap. 19.

References

- BECKER, G. (1973): “A theory of marriage, part I,” *Journal of Political Economy*, 81, 813–846.
- BLAIR, C. (1984): “Every finite distributive lattice is a set of stable matchings,” *Journal of Combinatorial Theory, Series A*, 37, 353–356.
- 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.
- DECKER, C., B. STEPHENS, AND R. MCCANN (2009): “When do systematic gains uniquely determine the number of marriages between different types in the Choo-Siow matching model? Sufficient conditions for a unique equilibrium,” mimeo University of Toronto.
- DUDLEY, R. M. (2002): *Real Analysis and Probability*. Cambridge University Press.
- ECHENIQUE, F. (2008): “What matchings can be stable? The testable implications of matching theory,” *Mathematics of Operations Research*, 33, 757–768.
- FOX, J. (2009): “Identification in Matching Games,” Discussion paper, NBER.
- GALE, D., AND L. SHAPLEY (1962): “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 69, 9–14.
- HIRIART-URRUT, J.-B., AND C. LEMARÉCHAL (2001): *Fundamental of Convex Analysis*. Springer.

- MÉZARD, M., AND A. MONTANARI (2009): *Information, Physics, and Computation*. Oxford University Press.
- PARISI, G. (1988): *Statistical Field Theory*. Perseus Books.
- RÜSCHENDORF, L. (1995): “Convergence of the iterative proportional fitting procedure,” *Annals of Statistics*, 23, 1160–1174.
- RÜSCHENDORF, L., AND W. THOMSEN (1998): “Closedness of Sum Spaces and the Generalized Schrödinger Problem,” *Theory of Probability and its Applications*, 42, 483–494.
- SLOW, A. (2009): “Testing Becker’s Theory of Positive Assortative Matching,” Discussion paper, University of Toronto.
- VAN DER VAART, A. (1998): *Asymptotic Statistics*. Cambridge University Press.
- VILLANI, C. (2003): *Topics in Optimal Transportation*. American Mathematical Society.
- (2009): *Optimal Transport, Old and New*. Springer.