Identification in Separable Matching with Observed Transfers

Bernard Salanié*

June 5, 2015

Abstract

Imposing a separability assumption on the joint surplus in tranferable utility matching models has proved very useful in empirical work. Yet when only "who matches whom" is observed, the distributions of unobserved heterogeneity cannot be identified separately. This note derives the distribution of equiilibrium transfers and shows that if the distribution of transfers within cells is observed, the distribution of heterogeneity can often be recovered, separability can be tested, and complementarities in surplus inferred.

1 Matching with Separable Surplus

We impose throughout this note the central assumption introduced by Choo and Siow (2006), which Chiappori-Salanié-Weiss (2014, hereafter CSW) and Galichon-Salanié (2014, hereafter GS) call "separability". To define it, assume a population of men $m = 1, \ldots, M$ and a population of women $w = 1, \ldots, W$ (the terms "men" and "women" are only for concreteness; they should be adapted to fit other one-to-one matching contexts.)

A hypothetical match between man m and woman w generates a joint surplus $\tilde{\Phi}_{mw}$. While these are known to all participants in the market, the analyst only observes that n_x men have characteristic x, m_y women have characteristics y, and μ_{xy} marriages between x-men and y-women took place. By subtraction, $n_x - \sum_y \mu_{xy}$ remain single; we use the standard convention of matching them with 0. In the rest of this note summations over y typically also include the 0 term. Similar notation is used for single women.

*Columbia University. Email: bsalanie@columbia.edu. I thank Nikhil Agarwal, Matthieu Bellon, Pierre-André Chiappori and Alfred Galichon for discussions on this topic.

Our problem is to identify as much as we can of the matrix $\tilde{\Phi}_{mw}$ while observing only the numbers n_x, m_y, and μ_{xy}; and possibly some data on transfers. As it is, this is clearly a hopeless task: there are many more unknown numbers than observations. To reduce the degree of underidentification, we impose:

Assumption S (separability) for all $m \in x$ and $w \in y$,

$$\tilde{\Phi}_{mw} = \Phi_{xy} + \varepsilon_{my} + \eta_{wx}.$$

Assumption S requires that conditional on observables, the partners' unobservables do not interact in creating joint surplus.

2 Identification without transfers

As proved in CSW and in GS, assumption S implies that there exist two matrices

$$U_{xy} + V_{xy} = \Phi_{xy}$$

such that man m chooses the characteristic of his partner y by solving

$$\max_y \left(U_{xy} + \varepsilon_{my}\right)$$

and woman w chooses the characteristic of her partner x by solving

$$\max_x \left(V_{xy} + \eta_{wx}\right)$$

Moreover, Galichon and Salanié prove that in large markets[1], the matrices U and V are just identified *if* the analyst knows the distributions of the unobserved heterogeneities.

More precisely, denote

- $\mathbf{P}_x$ the distribution of the vector $(\varepsilon_{my})_y$ conditional on $m \in x$
- and $\mathbf{Q}_y$ the distribution of the vector $(\eta_{wx})_x$ conditional on $w \in y$;

and for any matrix $U_{x\cdot} = (U_{xy})_y$,

$$G_x\left(U_{x\cdot}\right) = E_{\mathbf{P}_x} \max_y \left(U_{xy} + \varepsilon_{my}\right)$$

[1] That is, if the numbers n_x and m_y are very large for every x and y—see GS for a precise statement.

the expected maximized utility of men of characteristic x under $U_{x\cdot}$. Now define the Legendre–Fenchel transform over the marriage patterns $\mu_{\cdot|x} = (\mu_{y|x})_y$ of men of characteristic x:

$$G_x^*(\mu_{\cdot|x}) = \max_{U_{x\cdot}} \left(\sum_y \mu_{y|x} U_{xy} - G_x(U_{x\cdot}) \right).$$

GS proves that in equilibrium,

$$U_{xy} = \frac{\partial G_x^*}{\partial \mu_{y|x}}.$$

A similar equality obtains for V, replacing $\mathbf{P}_x$ with $\mathbf{Q}_y$, G_x with H_y, etc.

Therefore observing the marriage patterns μ identifies U, V and their sum $\Phi = U + V$, provided that surplus is separable and we know the distributions $\mathbf{P}_x$ and $\mathbf{Q}_y$. The latter assumption is clearly very strong. It was imposed in Choo and Siow 2006, who specified these distributions as products of iid standard type I extreme value distributions. Relaxing it has been a priority of more recent literaure.

A first possibility is to use data on several "similar" submarkets. CSW used data on thirty cohorts of men and women in the US to identify the heteroskedasticity of $\mathbf{P}_x$ and $\mathbf{Q}_y$, while still maintaining the type I EV assumption and independence[2].

This note proposes a second approach, which relies on observing data on transfers. We will explore two extreme cases:

- **Case 1:** the analyst only observes the mean value of the transfer $\tilde{t}_{mw}$ in each "cell" (x, y)—that is, she observes

  $$t_{xy} = E\left(\tilde{t}_{mw}|x, y\right)$$

  where the expectation is over realized matches of $m \in x$ and $w \in y$.

- **Case 2:** for each realized match, the analyst observes the value of the transfer $\tilde{t}_{mw}$.

We will always assume that pre-transfer utilities are themselves separable, so that the post-transfer utility of man $m \in x$ in a match with woman $w \in y$ is

$$a_{xy} + \varepsilon_{my}^a + \eta_{wx}^a - \tilde{t}_{mw}$$

[2]Fox and Yang (2013) show that one can sometimes identify the distribution of unobserved heterogeneity using data on many markets, even without separability. To do so, they impose restrictions on what we called Φ_{xy}.

and that of his partner is

$$b_{xy} + \varepsilon^b_{my} + \eta^b_{wx} + \tilde{t}_{mw},$$

with $a + b \equiv \Phi$, $\varepsilon^a + \varepsilon^b \equiv \varepsilon$, and $\eta^a + \eta^b \equiv \eta$.

3 Transfers are Separable

Given the results in GS, post-transfer utilities in a match $(m \in x, w \in y)$ equal

$$\frac{\partial G^*_x}{\partial \mu_{y|x}} + \varepsilon_{my}$$

and

$$\frac{\partial H^*_y}{\partial \mu_{x|y}} + \eta_{wx}.$$

Equating these expressions to those in section 2 shows that

$$\eta^a_{wx} - \varepsilon^b_{my} - \tilde{t}_{mw}$$

can only be a function of (x, y); that is, it cannot depend on the identity of $m \in x$ and $w \in y$. It follows that we must have

$$\tilde{t}_{mw} = T_{xy} + \eta^a_{wx} - \varepsilon^b_{my} \tag{1}$$

for some matrix T_{xy}. This proves that the separability of pre-transfer utilities implies the separability of equilibrium transfers.

In equilibrium, a man $m \in x$ who marries a woman $w \in y$ faces competition from all other men of characteristic x who value woman w; given prefect competition, this implies that he will have to pay her a salary that is shifted up by η^a_{my}. By a similar argument, the salary woman w receives is shifted down by the value of ε^b_{my}.

Moreover, we have

$$U_{xy} = a_{xy} - T_{xy} = \frac{\partial G^*_x}{\partial \mu_{y|x}} \tag{2}$$

and

$$V_{xy} = b_{xy} + T_{xy} = \frac{\partial H^*_y}{\partial \mu_{x|y}} \tag{3}$$

so that

$$T_{xy} = a_{xy} - \frac{\partial G^*_x}{\partial \mu_{y|x}} = \frac{\partial H^*_y}{\partial \mu_{x|y}} - b_{xy}. \tag{4}$$

4 The Equilibrium Value of Transfers

Given a specified model and the observed (or simulated) matching patterns, formulas (1) and (4) can be used to compute equilibrium transfers.

4.1 An Example

As a simple example, take the TU analog of Agarwal's (2014) matching market, in which one side of the market agrees on the ranking of all individuals on the other side. In my notation, say that all men agree on how to rank all women. This can only hold if pre-transfer utilities of men have $\varepsilon^a_{my} \equiv 0$ and $\eta^a_{wx} \equiv \eta_w$. Then in equilibrium in large markets, the results in GS imply that

$$\tilde{t}_{mw} = a_{xy} - \frac{\partial}{\partial \mu_{y|x}} \max_{U_{x\cdot}} \left(\sum_z \mu_{z|x} U_{xz} - E_{\boldsymbol{P}_x} \max_z \left(U_{xz} + \varepsilon^b_{mz} \right) \right) - \varepsilon^b_{my}.$$

Why this may seem complicated, it is often easy to evaluate in closed form. If for instance the ε^b_{my} are type I extreme value $G(-\gamma, \sigma^2_x)$, then

$$\tilde{t}_{mw} = a_{xy} - \sigma_x \log \mu_{y|x} - \varepsilon^b_{my}.$$

4.2 A Test for Separability

In addition, these formulas generate a simple prediction of separability (when conjoined with perfect competition, the absence of frictions and large markets): if m and m' share the same x and women w and w' share the same y, then

$$t_{mw} + t_{m'w'} = t_{mw'} + t_{m'w}.$$

In Case 1, this is non-testable since we only observe average transfers for each (x, y)-cell. But in case 2, this generates a testable prediction. As a deterministic equality on observables, it will of course be rejected by any data set. It is more useful to consider (1) as an approximation, and to evaluate its quality. A simple way to do so is to regress transfers on interacted dummies:

$$\tilde{t}_{mw} = d_{xy} + d_{my} + d_{wx} + \xi_{mw}.$$

The R^2 of this regression is a measure of how well separability fits the data.

5 The Identifying Power of Transfers

Now let us go beyond Choo and Siow by letting the distributions of $\varepsilon^a_{my}, \varepsilon^b_{my}, \eta^a_{wx}$ and η^b_{wx} depend on unknown parameters θ. Our identifying equations are (1), (2), and (3), and the RHS in the latter two now depend on θ. We wish to identify θ, a, b, and T.

First note that from the results in GS, given any value of θ,

$$a_{xy} - T_{xy} = \frac{\partial G_x^*}{\partial \mu_{y|x}}(\mu_{\cdot|x}; \theta)$$

and

$$b_{xy} + T_{xy} = \frac{\partial H_y^*}{\partial \mu_{x|y}}(\mu_{\cdot|y}; \theta)$$

are just identified from the matching patterns. In addition, we now know observe the realized transfers

$$\tilde{t}_{mw} = T_{xy} + \eta_{wx}^a - \varepsilon_{my}^b$$

in Case 2, and the average transfers

$$t_{xy} = T_{xy} + E\left(\eta_{wx}^a - \varepsilon_{my}^b | x, y\right)$$

in Case 1.

5.1 Restrictions on Pre-transfer Utilities

In some applications, the analyst will be able to assume that the matrices a and b are restricted in some way; this would help her identify θ. This is already true for models in which we do not observe transfers, and we do not pursue it further.

5.2 Restrictions on Unobserved Components

A particularly simple case is that in which the components ε^b and η^a are known to be identically zero. Then $\tilde{t}_{mw}$ is simply T_{xy}; so that T_{xy} is just-identified in Case 1, and massively over-identified (and testable) in Case 2.

In Case 1 as in Case 2, the knowledge of T also gives us a and b for any given θ; on the other hand, we cannot use the observability of transfers to learn about θ.

Note that $\varepsilon^b \equiv 0$, for instance, implies that all women of characteristics y are indifferent (pre-transfer) as to which man of characteristics x they may end up marrying: if they were offered a different x-husband with the same transfer, they would be equally happy. This sounds rather implausible.

The polar assumption, on which we focus from now on, would have ε^a and η^b be identically zero: all men of characteristics x agree in the way they rank (pre-transfer) the women of type y, and all women of characteristics y agree in the way they rank (pre-transfer) the men of type x. Then the distribution of $\tilde{t}_{mw}$ conditional on a match $(m \in x, w \in y)$ is the distribution of

$$T_{xy} + \eta^a_{wx} - \varepsilon^b_{my}$$

conditional on

$$\frac{\partial G^*_x}{\partial \mu_{z|x}}(\mu_{\cdot|x}; \theta) + \varepsilon^b_{mz} \tag{5}$$

being maximal in $z = y$ and

$$\frac{\partial H^*_y}{\partial \mu_{z|y}}(\mu_{\cdot|y}; \theta) + \eta^a_{wz}$$

being maximal in $z = x$. In Case 2, we observe this entire conditional distribution; and in Case 1 we only observe its mean.

5.2.1 An Illustrative Example

Let us take a simple example that minimally generalizes Choo and Siow 2006. As in CSW, we assume that the ε^b_{my} (resp. the η^a_{wx}) are type I EV iid with scale parameter σ_x (resp. τ_y.) Therefore $\theta = ((\sigma_x)_x, (\tau_y)_y)$.

The special properties of the type I EV distribution[3] imply that the distribution of the maximum in (5) is the original type I EV distribution translated by the expected maximum utility,

$$\sigma_x \log \sum_z \exp(U_{xz}/\sigma_x) = -\sigma_x \log \mu_{0|x}.$$

In addition, the results in GS show that

$$\frac{\partial G^*_x}{\partial \mu_{y|x}}(\mu_{\cdot|x}; \theta) = \sigma_x \log \left(\mu_{y|x}/\mu_{0|x}\right).$$

Putting this together, $\tilde{t}_{mw}$ should be distributed as the difference of T_{xy} and two independent type I EV distributions,

- one with location parameter $-\tau_y \log \mu_{x|y}$ and scale parameter τ_y
- one with location parameter $-\sigma_x \log \mu_{y|x}$ and scale parameter σ_x.

[3] See de Palma-Kilani 2007.

In Case 1, we obtain equivalently

$$\begin{aligned} t_{xy} &= T_{xy} - \tau_y \log \mu_{x|y} + \sigma_x \log \mu_{y|x} \\ &= a_{xy} - \tau_y \log \mu_{x|y} + \sigma_x \log \mu_{0|x} \\ &= \sigma_x \log \mu_{y|x} - \tau_y \log \mu_{0|y} - b_{xy}. \end{aligned}$$

in which the terms T_{xy}, a_{xy}, b_{xy}, σ_x and τ_y are unknown. Jointly with

$$a_{xy} + b_{xy} = \sigma_x \log \left(\mu_{y|x} / \mu_{0|x} \right) + \tau_y \log \left(\mu_{x|y} / \mu_{0|y} \right),$$

this is the empirical content of the model. Given the structure of these equations, it is clear that we cannot identify θ.

On the other hand, in Case 2 we can identify the scale parameters σ_x and τ_y. They are in fact just identified from the second- and third-order moments of $\tilde{t}$; and they are overidentified using higher-order moments, so that the specification is testable. It would even be possible to allow for some mismeasurement of transfers.

While this is only an illustrative example, it is interesting to note that it is exactly the model in CSW. Observing the distribution of transfers in every cell therefore substitutes for the repeated cross-sections they used[4].

5.2.2 A More General Treatment

For any non-negative integer k, define

$$c_{ky} = E \left(\varepsilon_{my}^k \mathbb{1} \left(U_{xy} + \varepsilon_{my} \geq U_{xz} + \varepsilon_{mz} \text{ for all } z \right) \right).$$

Note that each c_{ky} is a random function of the differences of utilities $(U_{xy} - U_{xz})$ and of θ. Moreover, c_{0y} is simply the probability that y is chosen by $m \in x$, which is $\mu_{y|x}$. Therefore the equation

$$c_{0y} = \mu_{y|x}$$

simply rewrites equation (2). Solving these equations for all values of y gives us the differences of utilities as a function of θ. Moreover, the conditional moments of ε_{my} given that the maximum is achieved in y are simply the c_{ky}/c_{0y} for $k \geq 1$; and since we already solved for the differences of utilities, these are functions of θ only.

In Case 1, we only observe t_{xy}, which is the sum of T_{xy}, of c_{1y}/c_{0y} and of the equivalent term for η. This is essentially uninformative on θ.

[4]It does a little more, in fact, since it identifies the general scale of utilities.

In Case 2, we know the distribution of transfers for (x, y) matches. In particular, we know its moments of all orders; but these are functions of the unknown T_{xy}, of the conditional moments c_{ky}/c_{0y} of ε, and of those of η. Therefore we have an infinity of equations to solve for θ, and a tempting conjecture is that any finite-dimensional θ is overidentified.

While we do not have a proof for this conjecture, some more special results are easily obtained. Consider for instance the case in which the shape of the distribution for each gender is known but its scale parameter is not; then $\theta = ((\sigma_x), (\tau_y))$. It is easy to see that the c_{ky} scale as σ_x^k with the difference utilities divided by σ_x:

$$E\left((\sigma_x \varepsilon_{my})^k \mathbb{1}\left(U_{xy} + \sigma_x \varepsilon_{my} \geq U_{xz} + \sigma_x \varepsilon_{mz} \text{ for all } z\right)\right)$$
$$= \sigma_x^k E\left(\varepsilon_{my}^k \mathbb{1}\left(\frac{U_{xy}}{\sigma_x} + \varepsilon_{my} \geq \frac{U_{xz}}{\sigma_x} + \varepsilon_{mz} \text{ for all } z\right)\right).$$

As a consequence, the conditional moments c_{ky}/c_{0y} scale as σ_x^k, and knowledge of the moments of $\tilde{t}_{mw}$ identifies the θ parameter. Thus the findings reported in 5.2.1 for the type I EV distribution apply much more generally, and scale parameters are identified.

6 Testing for Assortative Matching

Much attention has been devoted to "sorting" in matching markets. Eeckhout and Kircher (2011) for instance make the point that even in a frictionless market like the one we consider, a given set of transfers may be the product of very different surplus functions, only some of which involve positive complementarities. In this section we are concerned about inferring complementarities in surplus from the observed data. "Complementarities in surplus" have a straightforward definition: we require that for the given orderings on the sets of men characteristics x and of women characteristics y, the joint surplus be supermodular in (x, y).

To put it more formally, there are complementarities in surplus iff for all men characteristics x and z and women characteristics y and t,

$$\Phi_{x \vee z, y \vee t} + \Phi_{x \wedge z, y \wedge t} \geq \Phi_{xt} + \Phi_{zy}.$$

Complementarity in surplus is an ordinal property: it is invariant if x and y are subjected to increasing transformations.

Take the CSW example of section 5.2.1. When the σ_x and τ_y parameters all equal one (which defines the model of Choo and Siow 2006), then

$$\Phi_{xy} = \log \frac{\mu_{y|x}}{\mu_{0|x}} + \log \frac{\mu_{x|y}}{\mu_{0|y}}.$$

Eliminating the terms that do not interact x and y, it is clear that Φ is supermodular iff $\log \mu$ is—a point already made by Siow (2009). Observing transfers is not necessary to infer complementarities; the matching patterns μ are all that is needed. Their log-supermodularity is equivalent to the supermodularity of the joint surplus.

If we reintroduce the parameters σ_x and τ_y, then

$$\Phi_{xy} = \sigma_x \log \frac{\mu_{y|x}}{\mu_{0|x}} + \tau_y \log \frac{\mu_{x|y}}{\mu_{0|y}},$$

and the supermodularity of Φ is equivalent to that of the matrix

$$M_{xy} = (\sigma_x + \tau_y) \log \mu_{xy},$$

The supermodularity of M in turn depends on the logsupermodularity of μ, but also on its monotonicity and on that of σ_x and τ_y. While the first two are readily observed from the data, the latter is not; and inferring complementarities requires learning about the additional parameters σ_x and τ_y.

To illustrate this, suppose that we observe aggregate transfers as in case 1, and that (a, b, σ, τ) rationalizes the data (t, μ), with

$$\begin{aligned} a_{xy} + b_{xy} &= \sigma_x \log \frac{\mu_{y|x}}{\mu_{0|x}} + \tau_y \log \frac{\mu_{x|y}}{\mu_{0|y}} \\ t_{xy} &= a_{xy} - \tau_y \log \mu_{x|y} + \sigma_x \log \mu_{0|x} \\ &= \sigma_x \log \mu_{y|x} - \tau_y \log \mu_{0|y} - b_{xy}. \end{aligned}$$

Now pick any $(\delta\sigma_x, \delta\tau_y)$, and modify a and b as follows:

$$\begin{aligned} \delta a_{xy} &= \delta\sigma_x \ \log \mu_{xy} \\ \delta b_{xy} &= \delta\tau_y \ \log \mu_{xy}. \end{aligned}$$

It is easy to see that $(a + \delta a, b + \delta b, \sigma + \delta\sigma, \tau + \delta\tau)$ rationalizes the same data; and with these degrees of freedom we can easily make M supermodular, submodular, or neither. The only exception is the intermediate case in which we assume that σ_x is independent of x and τ_y is independent of y. Then we can still conclude that Φ is supermodular iff $\log \mu$ is..

The lesson from this exercise is that unless we restrict the pre-transfer utilities a or b, their separate and unknown complementarities will defeat our attempts to infer those of the joint surplus. Restrictions on a and b do not yield that much, however. Assume that there are no interactions between x and y in a (so that a_{xy} only depends on x, or on y.) Then t is supermodular iff $\tau_y \log \mu_{xy}$ is; and M is supermodular if and only if

$$\sigma_x \log \mu_{xy} + t_{xy}$$

is supermodular. If we knew the value of σ_x, then we would identify the complementarities in surplus; but even an additional assumption that σ_x is a constant σ will not allow us to conclude unless for instance $\log \mu$ and t are both supermodular.

Finally, remember from section 5.2.1 that in case 2, σ_x and τ_y are identified from the individual transfers; and therefore we can test for complementarities once we have recovered estimates of these parameters.

Conclusion

The results in this note are mixed. When the only information on transfers is available at the "cell" level as in Case 1, then it will not allow us to identify any parameter of the distribution of unobserved heterogeneity. On the other hand, information on the distribution of transfers within each cell has the potential to be very informative on this distribution under separability. It also allows the analyst to test for separability without resorting to specific assumptions on the distributions of heterogeneity.

While this note focused on the identifying power of transfers, in some situations the analyst observes measurements of outcomes; these will be informational on the joint surplus $\tilde{\Phi}_{mw}$, or on the way it is shared. In marriage markets, such measurements could include divorces, children outcomes, or allocation of household expenditure. In labor markets one could have measurements of the productivity of a match. The survey by Graham (2011) discusses this at greater length.

References

Agarwal, N. (2014), "An Empirical Model of the Medical Match", *American Economic Review*, forthcoming.

Chiappori, P.-A., B. Salanié and Y. Weiss (2014), "Partner Choice and the Marital College Premium", mimeo, Columbia University.

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

de Palma, A. and K. Kilani (2007), "Invariance of Conditional Maximum Utility", *Journal of Economic Theory*, 132, 137–146.

Eeckhout, J. and P. Kircher (2011), "Identifying Sorting: In Theory", *Review of Economic Studies*, 78, 873–906.

Fox, J. and C. Yang (2013), "Unobserved Heterogeneity in Matching Games", mimeo University of Wisconsin.

Galichon, A., and B. Salanié (2014), "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models ", mimeo.

Graham, B. (2011), "Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers", in J. Benhabib, M. Jackson and A. Bisin eds: *Handbook of Social Economics*, North-Holland, Vol. 1B, 965–1052.

Siow, A. (2009), "Testing Becker's Theory of Positive Assortative Matching," University of Toronto, mimeo.