The Econometrics of Matching with Transferable Utility: A Progress Report[1]

Pierre-André Chiappori[2] Dam Linh Nguyen[3]
Bernard Salanié[4]

October 21, 2025

[1]We thank Thierry Magnac and Shruti Sinha for helpful comments. A previous version of this paper was circulated under the title "Matching with Random Components: Simulations".

[2]Department of Economics, Columbia University. Email: pc2167@columbia.edu
[3]Department of Economics, New York University. Email: n.linh@nyu.edu
[4]Department of Economics, Columbia University. Email: bsalanie@columbia.edu

Abstract

Since Choo and Siow (2006), a burgeoning literature has analyzed matching markets when utility is perfectly transferable and the joint surplus is separable. We take stock of recent methodological developments in this area. Combining theoretical arguments and simulations, we show that the separable approach is reasonably robust to omitted variables and/or non-separabilities. We conclude with a caveat on data requirements and imbalanced datasets.

Introduction

The empirical analysis of matching markets has made considerable progress in recent years. We will focus here on markets where partners exchange transfers freely. In the usual terminology, we study matching markets with perfectly transferable utility (hereafter "TU"). In most matching markets, observationally identical agents end up with very different matching outcomes: some may be unmatched, and others will be matched to partners with variable observable characteristics. This can be rationalized by introducing frictions, unobserved heterogeneity, or a mixture of both. As explained in Chiappori and Salanié (2016), these two groups of rationalizations have essentially identical predictions for cross-sectional data; they can only be distinguished with data featuring transitions, which are often unavailable to the analyst. We choose here to model the dispersion of outcomes using unobserved heterogeneity. More precisely, we work with the class of models pioneered by Choo and Siow (2006), which incorporates a quasi-additive error structure called *separability* by Chiappori, Salanié, and Weiss (2017). Galichon and Salanié (2017, 2022) study the properties of separable models in much more detail. Galichon and Salanié (2024a) show that this class of models can be easily estimated using minimum distance, which often boils down to generalized least squares; and that the Choo and Siow model can also be estimated with Poisson GLM.

Beyond Choo and Siow (2006), this framework has been applied by Chiappori, Salanié, and Weiss (2017) to changes in the marital college premium in the US. Ahn (2023) used it to analyze cross-border marriages in South-East Asia, and Chiappori, Florio, Galichon, and Verzillo (2024) to describe matching on income levels in the marriage market. It was also applied to the labor market (in a many-to-one version) by Corblet (2022), and to mergers and acquisitions by Guadalupe, Rappoport, Salanié, and Thomas (2024). It was extended to continuous types by Dupuy and Galichon (2014) to study the contribution of personality traits to marital surplus; by Galichon, Kominers, and Weber (2019) to imperfectly transferable utilities; and by Ciscato, Galichon, and Goussé (2019) to non-bipartite matching in order to compare same-sex and different-sex marriages.

By definition, separability rules out the interaction between partners' unobserved heterogeneity in the production of joint surplus. While it is a very useful assumption for keeping the model tractable, it is also restrictive, especially if the set of characteristics available to the analyst is small. A standard intuition would suggest that this may not matter much if unobserved characteristics are drawn independently of observed characteristics. However, we are dealing here with a two-sided market where we cannot just transpose this intuition. The first goal of this paper is to explore the consequences of relaxing separability on equilibrium matching patterns, utilities, and division of surplus. We will also quantify the misspecifications that result from mistakenly assuming separability when estimating a non-separable model.

Theoretical arguments and a Monte Carlo simulation suggest that non-separability impacts matching patterns and utilities in ways that, ex post, seem intuitive. It takes a rather large amount of non-separability to see a qualitative difference from the separable case, however. We find that when we estimate a non-separable model as if it were separable, the estimated complementarities in surplus are surprisingly robust. This is reassuring since we have known since Becker (1973) that complementarities play a crucial role in TU matching markets. Our conclusions are strongest when the distributions of the observable characteristics of the partners are similar ("symmetric margins"). As we will see, non-separabilities matter more when markets are very unbalanced.

The second contribution of the paper is to point out that estimates may be fragile when the data contain matching cells of very different size. We give a simple, two-by-two example to show that the asymptotic variance of the estimator may be dominated by the inverse of the size of the smallest cell, unless precautions are taken to minimize its influence.

The remainder of the paper is organized as follows. Section 1 introduces the class of separable matching models with perfectly transferable utilities. Section 2 discusses the separability assumption and its implications. Section 3 explores the consequences of omitted variables and of mistakenly imposing separability. Section 4 discusses the effects of imbalanced data, and how to remedy them. We conclude with some directions for further work.

1 Model and Methods

We focus throughout on a bipartite one-to-one model of matching with unobserved heterogeneity and perfectly transferable utilities[1]. Since it is similar to Becker's "marriage market", we will use that terminology and refer to potential partners as "men" and "women". Our framework easily accommodates other interpretations. We maintain some of the standard assumptions: matching is frictionless and all potential partners have the same information. The analyst only observes a subset of individual characteristics.

Men are indexed by $i \in \mathcal{I}$ and women by $j \in \mathcal{J}$. Each man i (resp. woman j) has a "type" $x_i \in \mathcal{X}$ (resp. $y_j \in \mathcal{Y}$). This is observable to all men, all women, and to the

[1] Galichon and Salanié (2024a) extend the analysis to more complex separable models.

analyst. When a man i and a woman j marry, their match generates a surplus $\tilde{\Phi}_{ij}$ which they can share freely. Singles match to "0": a single man i attains a utility of $\tilde{\Phi}_{i0}$ and a single woman j attains $\tilde{\Phi}_{0j}$.

This paper will assume that the sets of types are finite: $\mathcal{X} = \{1, \ldots, X\}$ and $\mathcal{Y} = \{1, \ldots, Y\}$. As mentioned in the introduction, Dupuy and Galichon (2014) extend the model to continuous types. We will denote by $\tilde{\boldsymbol{\Phi}} = (\tilde{\Phi}_{ij})_{i=1,\ldots,I;j=1,\ldots,J}$ the matrix of surpluses.

A matching is a list of numbers $(\mu_{ij}, \mu_{i0}, \mu_{0j})_{i=1,\ldots,I;j=1,\ldots,J}$, with $\mu_{ij} = 1$ if i and j match, $\mu_{i0} = 1$ if i is single, etc. We focus on the stable matching (it is generically unique). Since only types are observed by the analyst, the data consist of *matching patterns*: the numbers of men and women of each type, which we denote n_x and m_y, and the numbers of marriages μ_{xy} between men of each type x and women of each type y. By construction, $\mu_{x0} = n_x - \sum_{y=1}^{Y} \mu_{xy}$ and $\mu_{0y} = m_y - \sum_{x=1}^{X} \mu_{xy}$ are the numbers of single men of type x and of single women of type y. Note that in some empirical applications, the analyst may not observe singles—only the numbers of marriages μ_{xy}.

We denote by $\boldsymbol{\Phi}$ and $\boldsymbol{\mu}$ the matrix of surpluses and the matching patterns, respectively; and $\boldsymbol{r} = (\boldsymbol{n}, \boldsymbol{m})$ the *margins*, i.e., the numbers of men and women of each type. Like most of the literature, we work with the *large market limit* in which the vectors $\boldsymbol{n}$ and $\boldsymbol{m}$ grow proportionately.

1.1 Separability

The separable model restricts the form of the surplus $\tilde{\Phi}_{ij}$. It imposes that there exist a matrix $\boldsymbol{\Phi} = (\Phi_{xy})_{x \in \mathcal{X}, y \in \mathcal{Y}}$, random variables $\boldsymbol{\varepsilon^i} = (\varepsilon^i_0, \varepsilon^i_1, \ldots, \varepsilon^i_Y)$, and $\boldsymbol{\eta^j} = (\eta^j_0, \eta^j_1, \ldots, \eta^j_X)$ such that

$$\tilde{\Phi}_{ij} \coloneqq \Phi_{x_i, y_j} + \varepsilon^i_{y_j} + \eta^j_{x_i} \tag{1}$$

within any match, and that single men and women get $\tilde{\Phi}_{i0} = \varepsilon^i_0$ and $\tilde{\Phi}_{0j} = \eta^j_0$, respectively. Moreover, the random vectors $\boldsymbol{\varepsilon^i}$ (resp. $\boldsymbol{\eta^j}$) are assumed to be drawn independently of each other, conditional on x_i (resp. y_j).

Chiappori, Salanié, and Weiss (2017) and Galichon and Salanié (2017, 2022, 2024a) studied the class of separable models in depth. They derived some of their properties, proved identification results, and proposed estimators. We summarize them briefly here and refer the reader to recent surveys by Salanié (2024) and Galichon and Salanié (2024b) for more information.

1.1.1 Identification

We assume in this paper that the available data only describe "who matches with whom": that is, they consist of the matching patterns $\boldsymbol{\mu}$. Data on transfers between partners and/or proxies for the outcomes of matches, if available, may provide more identifying power[2].

Galichon and Salanié (2022) show that the stable matching patterns maximize the social welfare (the sum of the surpluses of all matches)

$$\sum_{x=1}^{X}\sum_{y=1}^{Y}\mu_{xy}\Phi_{xy} + \mathcal{E}(\boldsymbol{\mu})$$

where the *generalized entropy* $\mathcal{E}$ depends on the margins $\boldsymbol{r}$ and on the distributions of the ε^i and η^j terms. This is a globally strictly convex problem. When (as we will assume) these terms have full support, all matching patterns must be strictly positive and the set of first-order conditions

$$\Phi_{xy} = -\frac{\partial \mathcal{E}}{\partial \mu_{xy}}(\boldsymbol{\mu}) \tag{2}$$

defines what can be identified from a single matching market.

For any choice of the distributions of the ε^i and η^j terms, the joint surplus matrix $\boldsymbol{\Phi}$ is just identified. If for instance the analyst wants to identify the standard errors of these terms, restrictions on the joint surplus must be imposed. The alternative is to pool data from several markets and impose cross-market restrictions. This was the approach of Chiappori, Salanié, and Weiss (2017), for instance.

When the analyst only observes realized matches (no data on singles), a similar set of equations apply; but they only identify the joint surplus matrix up to arbitrary additive transformations $\Phi_{xy} \to \Phi_{xy} + a_x + b_y$. This simply translates the fact that the value of marriage cannot be identified without data on the proportion of singles.

1.1.2 Inference

Now suppose that the distributions of the unobserved heterogeneity terms (resp. the joint surplus matrix $\boldsymbol{\Phi}$) are known up to a parameter vector $\boldsymbol{\alpha}$ (resp. $\boldsymbol{\beta}$), and that the model is identified from data on a single market[3]. For any value of $\boldsymbol{\alpha}$, the generalized entropy $\mathcal{E}^{\boldsymbol{\alpha}}$ can be evaluated easily—often in closed form. Given observed matching

[2] See Salanié (2015) for the case when transfers are observed.

[3] This is typically the case if the dimensions of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ add up to $X \times Y$ or less.

patterns $\hat{\boldsymbol{\mu}}$, the parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ can be estimated by minimizing the norm of the $X \times Y$ vector with components

$$\Phi^{\boldsymbol{\beta}}_{xy} + \frac{\partial \mathcal{E}^{\boldsymbol{\alpha}}}{\partial \mu_{xy}}(\hat{\boldsymbol{\mu}}).$$

The resulting minimum distance estimator has the usual properties: it is consistent and asymptotically normal; there is a choice of the norm that minimizes its variance-covariance matrix; and for that choice of the norm, the minimized value of the objective function provides a χ^2 test of the specification.

While we focus here on parametric models, we note that Gualdani and Sinha (2023) have explored nonparametric inference in this context.

1.2 The Choo and Siow Specification

The best-known and most convenient separable model has the "multinomial logit" form popularized by Choo and Siow (2006). They assumed that each component of the vectors $\boldsymbol{\varepsilon^i}$ and $\boldsymbol{\eta^j}$ is drawn independently from a centered standard type I extreme value distribution. They showed that the stable matching patterns $\boldsymbol{\mu}$ obey very simple formulæ: the number μ_{xy} of marriages between men of education x and women of education y can be written as

$$\mu_{xy} = \sqrt{\mu_{x0}\mu_{0y}} \exp(\Phi_{xy}/2)$$

where μ_{x0} (resp. μ_{0y}) is the number of men of education x (resp. women of education y) who are single in equilibrium.

This gives $X \times Y$ equations in $X \times Y + X + Y$ unknowns; the system is completed by the scarcity constraints

$$n_x = \sum_{y=1}^{Y} \mu_{xy} + \mu_{x0}$$
$$m_y = \sum_{x=1}^{X} \mu_{xy} + \mu_{0y}.$$

Galichon and Salanié (2022) show how given the "margins" $(n_x), (m_y)$ and the joint surplus matrix $\boldsymbol{\Phi}$, this system of equations can be solved very efficiently using an Iterative Projection Fitting Procedure (IPFP). Identification is straightforward: since the distributions of the unobserved heterogeneity terms are parameter-free, the joint surplus matrix $\mathbf{\Phi}$ is just identified from the matching patterns $\boldsymbol{\mu}$. In fact, the generalized entropy in this case is the usual entropy, and the first-order conditions (2) give

$$\Phi_{xy} = \log \frac{\mu_{xy}^2}{\mu_{x0}\mu_{0y}}.$$

When the numbers of singles are not observed, $\Phi_{xy} = 2\log\mu_{xy} + a_x + b_y$, with arbitrary $\boldsymbol{a}$ and $\boldsymbol{b}$. In addition to the minimum distance procedure described earlier, a Poisson model with two-way fixed effects can be used to estimate the Choo and Siow model (see Galichon and Salanié 2024a).

Since the Choo and Siow (2006) specification is so simple and has been widely used, it is a natural benchmark and we will use it in our examples. It has been criticized on a number of grounds, however; see for instance Galichon and Salanié (2017, 2019, 2022). Since the components of $\boldsymbol{\varepsilon^i}$ are independent of each other, the specification rules out "local" correlations. In addition, the Choo and Siow matching model inherits a two-sided version of the Independence of Irrelevant Alternatives property of the one-sided multinomial logit. As in the one-sided case, this has some advantages and many drawbacks[4].

2 Discussion of Separability

By definition, separability rules out any contribution to the joint surplus of a match of interactions between characteristics of partners that are unobserved by the analyst. Separability is best viewed through the lens of an ANOVA decomposition: it requires that conditional on the observed types of the partners, interactions between their unobserved types do not contribute to the variation in the surplus.

Separability does *not* rule out "matching on unobservables": a man i of education x and a woman j of type y are more likely to marry if his ε_y^i and her η_x^j take higher values. It does have strong consequences, however. Suppose i and j do marry at the stable matching. Since this man gets utility $U_{xy} + \varepsilon_y^i$, he would be equally happy with any other woman of type y. If for instance the observed types only record education, and 1,000 couples form between college-educated partners, one could shuffle these 1,000 men and 1,000 women without changing anyone's utility level.

As can be seen from this example, separability is more restrictive when the data contain little information on types and/or unobserved heterogeneity is highly relevant for the application under study. In a sense, assuming separability acknowledges that

[4]See Mourifié (2019) and Mourifié and Siow (2017) for extensions of the Choo and Siow framework that reach beyond separable models.

in these two-sided problems, the data are unlikely to give us much information about interactions between unobserved components[5]. Still, it is worth investigating whether separable models yield robust conclusions.

Separability has very useful properties that greatly simplify analysis and inference. It implies that at the stable matching, there exist matrices $\boldsymbol{U}$ and $\boldsymbol{V} = \boldsymbol{\Phi} - \boldsymbol{U}$ such that if a man i of type x and a woman j of type y marry in equilibrium, he gets utility

$$u_i = U_{xy} + \varepsilon^i_y = \max(\varepsilon^i_0, \max_{t=1}^{Y}(U_{xt} + \varepsilon^i_t))$$

and she gets $v_j = V_{xy} + \eta^j_x = \max(\eta^j_0, \max_{z=1}^{Y}(V_{zy} + \eta^j_z))$.

The stable matching $\boldsymbol{\mu}$ is homogeneous of degree 1 in the vector $(\boldsymbol{n}, \boldsymbol{m})$: if the numbers of men and women of each type are multiplied by a common positive integer k, then all numbers of marriages μ_{xy} and of singles μ_{x0} and μ_{0y} are multiplied by k as well. On the other hand, utilities are unchanged; in that sense, there is no scale effect in a separable matching market.

As the matching patterns maximize the social welfare, which is globally convex, its Hessian must be semi-definite positive at the optimum. This entails testable implications that (depending on the available data) may allow for tests of the separability assumption (see Galichon and Salanié 2022).

3 Misspecifying Separability

We explore here the consequences of mistakenly imposing separability. There are of course many ways to add non-separable components to a separable matching model. We could for instance add a simple interaction term to Equation (1):

$$\tilde{\Phi}_{ij} \coloneqq \Phi_{x_i,y_j} + \tau(\varepsilon^i_{y_j} + \eta^j_{x_i}) + \sigma\, \boldsymbol{\xi}_i \cdot \boldsymbol{\zeta}_j \tag{3}$$

with τ and σ two positive parameters that control the relative strength of the non-separable component. We call this a case of "missing interaction".

Alternatively, we could use a "missing shock" model. Generate a matrix $\boldsymbol{\nu} = (\nu_{ij})_{i=1,\dots,I; j=1,\dots,J}$ of independent draws from some mean-zero distribution and take ν_{ij} to represent a pair-specific unobserved preference shock:

$$\tilde{\Phi}_{ij} \coloneqq \Phi_{x_i,y_j} + \tau(\varepsilon^i_{y_j} + \eta^j_{x_i}) + \sigma\nu_{ij}. \tag{4}$$

In both cases, we recover the separable surplus in (1) if we set $\tau = 1$ and $\sigma = 0$.

[5] But see Fox, Yang, and Hsu (2018) for some results along these lines.

3.1 Missing Interactions

Missing interactions may not create major difficulties, as long as they are only weakly correlated to the included interactions. Suppose that the true model is as in Choo and Siow, with one difference: the types $\boldsymbol{x}$ and $\boldsymbol{y}$ are continuous variables distributed as two zero-mean Gaussian random variables in $\mathbb{R}^{K+1}$. We take the joint surplus to have the following quadratic form:

$$\Phi(x, y) = \sum_{k,l=1}^{K+1} A_{kl} x_k y_l.$$

This is a version of the quadratic-normal Tinbergen (1956) model. It follows from Bojilov and Galichon (2016) that the stable matching is described by an affine mapping between men and women types: $\boldsymbol{y} = \boldsymbol{T}\boldsymbol{x} + \boldsymbol{\xi}$ where conditional on $\boldsymbol{x}$, $\boldsymbol{\xi}$ has mean zero. The matrix $\boldsymbol{T}$ depends on the matrix $\boldsymbol{A}$ and the variance-covariance matrices $\boldsymbol{\Sigma_x}$ and $\boldsymbol{\Sigma_y}$ of $\boldsymbol{x}$ and $\boldsymbol{y}$. The matrix $\boldsymbol{T}$ can be estimated simply by regressing $\boldsymbol{y}$ on $\boldsymbol{x}$ over observed matches. Moreover, given $\boldsymbol{\Sigma_x}$ and $\boldsymbol{\Sigma_y}$ (which are easily estimated), the mapping from $\boldsymbol{A}$ to $\boldsymbol{T}$ can be inverted to recover an estimator of $\boldsymbol{a}$.

Now suppose that there are no interactions between x_{K+1} and $(y_1, \ldots, y_K)$, as well as between y_{K+1} and $(x_1, \ldots, x_K)$, so that the matrix $\boldsymbol{A}$ has the same block-diagonal structure:

$$\Phi(x, y) = \sum_{k,l=1}^{K} A_{kl} x_k y_l + A_{K+1,K+1} x_{K+1} y_{K+1}.$$

If we only observe $(x_1, \ldots, x_K)$ and $(y_1, \ldots, y_K)$, can we still estimate consistently the coefficients $(A_{11}, \ldots, A_{KK})$? This obviously won't work in general: if for instance $E(x_{K+1}|x_1, \ldots, x_K) = x_1$ and $E(y_{K+1}|y_1, \ldots, y_K) = y_1$, then omitting x_{K+1} and y_{K+1} will bias the estimate of the coefficient A_{11}. However, if we further assume that the $(K+1)$-th components of $\boldsymbol{x}$ and $\boldsymbol{y}$ are independent of the first K components, the answer is positive[6]. This is easy to see from the formulæ in Bojilov and Galichon (2016, Theorem 1): under our assumptions, the matrix $\boldsymbol{T}$ inherits the block-diagonal structure of $\boldsymbol{\Sigma}$, and its top-left (K, K) block does not depend on the value of $A_{K+1,K+1}$ or of the variances of x_{K+1} and y_{K+1}. As a consequence, the coefficients $(A_{11}, \ldots, A_{KK})$ can be estimated by regressing the vector $(y_1, \ldots, y_K)$ on the vector $(x_1, \ldots, x_K)$.

It follows from this result that if we do not observe x_{K+1} or y_{K+1}, we can still get consistent estimates of $(A_{11}, \ldots, A_{KK})$. This finding obviously extends to more general block structures. It is superficially similar to classic results on (the absence

[6] If some omitted variables are *not* independent of the included variables, one could conceivably use instrumental variables, as is done in a linear model. We leave this for further research.

of) omitted variable bias in the linear model: it is much less general, however. It only obtains here because with a quadratic joint surplus and Gaussian types, the "optimal transport map" $\boldsymbol{y} = \boldsymbol{T}(\boldsymbol{x})$ yields a linear statistical model. Still, this suggests that the separable model has some degree of robustness to omitted variables that are functionally and statistically independent of those that are included.

3.2 Missing Shocks

Let us now turn to the model of (4), where the true joint surplus includes pair-specific unobserved preference shocks that are excluded from the estimated model. We explored this case extensively via Monte Carlo simulations. We assumed that the ε and η terms are drawn independently from a standard type I extreme value distribution, as in the Choo and Siow (2006) model. The total variance of the surplus in the separable ($\sigma = 0, \tau = 1$) model is 2; we ensure that it is the same in the non-separable model by imposing $\tau^2 = 1 - \sigma^2/2$. This implies that σ must take values in $[0, \sqrt{2}]$. To put it differently, it is tempting to define a "coefficient of determination" by $R^2 = \sigma^2/2$. It represents the fraction of the total variation in the joint surplus that is explained by the non-separable component (the pair-specific individual preferences ν_{ij}) for given x_i and y_j. Our simulations will cover the whole range from $R^2 = 0$ to $R^2 = 1$.

We found that the biases that result from this misspecification grow slowly with the magnitude of the contribution of the interaction terms. In particular, the estimated complementarities in the Choo and Siow (2006) model are remarkably robust to the inclusion of interaction terms.

3.2.1 Some Theory

We do not know of any study of the properties of matching under non-separability, even in the limit case where $R^2 = 1$. However, we state below two results and a conjecture.

Our first result builds on standard properties of linear programs in Euclidean spaces. As is well-known, the solution to any such program is generically unique and robust to small changes in the parameters. To put it more formally, define a general linear program as

$$\max_{\boldsymbol{x} \in \mathbb{R}^n} \boldsymbol{c}'\boldsymbol{x} \ \text{ s.t. } \ \boldsymbol{A}\boldsymbol{x} \leqslant \boldsymbol{b}$$

and assume that $\boldsymbol{A}$ and $\boldsymbol{b}$ define a non-empty constrained set. Then the problem has a solution set $\boldsymbol{x}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$. For a generic choice of $(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$, the solution is a singleton;

and if $(\boldsymbol{A}', \boldsymbol{b}', \boldsymbol{c}')$ is close enough to $(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c})$,

$$\boldsymbol{x}(\boldsymbol{A}', \boldsymbol{b}', \boldsymbol{c}') = \boldsymbol{x}(\boldsymbol{A}, \boldsymbol{b}, \boldsymbol{c}).$$

Whether it is separable or not, TU matching is an instance of a linear program; and it is finite-dimensional if the number of potential partners is finite. Therefore for a generic draw of the parameters, the optimal matching is a piecewise constant function of σ. In particular, it is generically the same for fully separable and for almost separable models:

Result 1: in finite markets, generically (for almost all draws), there is a $\bar{\sigma}$ such that the optimal matching is the same for all $0 \leqslant \sigma < \bar{\sigma}$.

Result 1 exploits the fact that in finite markets, any statistic of the model that is a continuous function of σ is locally constant. This does not extend to the "large markets" limit that is usually assumed in the empirical applications of separable matching models: then such statistics become smooth functions of σ. Our second result states that for small σ, these statistics (and in particular any misspecification bias) are of order $O(\sigma^2)$ for small σ when the non-separable component is drawn from a distribution that is symmetric around zero.

Result 2: in large markets, if the distribution of ν is symmetric around zero then the effects of non-separability are in σ^2 for small σ.

The proof of result 2 relies on the fact that in large markets, the empirical distribution of all draws of the ν_{ij} component converges to its generating distribution. If the latter is symmetric, then changing all $\sigma\nu_{ij}$ to $(-\sigma)\nu_{ij}$ cannot affect the equilibrium in the large market limit[7]. A fortiori, all statistics computed on the equilibrium must be locally even functions of σ. The equilibrium matching is a smooth function of σ in the limit; therefore its difference with the equilibrium matching of the separable model must be at most $O(\sigma^2)$.

In general, adding a non-separable mean-zero term in a separable model is akin to increasing the opportunities for successful matches. By making the market "thicker", it increases the probabilities of all kinds of matches: we expect (and our simulations will confirm) $\mu_{11}, \mu_{12}, \mu_{21}$, and μ_{22} to be larger, and $\mu_{10}, \mu_{01}, \mu_{20}$ and μ_{02} to be smaller, than in the fully separable model. On the other hand, since we are adding random terms that are distributed independently of partners' observed types, there is no obvious reason why the probabilities of non-diagonal matches (μ_{12} and μ_{21}) should

[7] We skip here over the normalizations that are needed to rescale the non-separable model as the number of individuals grows without bounds—see Menzel (2015) for a thorough analysis.

increase much differently than those of diagonal matches (μ_{11} and μ_{22}). In particular, the log-odds-ratio

$$\log \frac{\mu_{11}\mu_{22}}{\mu_{12}\mu_{21}}$$

may not be so different in non-separable models. In the Choo and Siow (2006) specification that we use to estimate the parameters of the model, this ratio identifies with the "supermodular core" defined in Chiappori, Salanié, and Weiss (2017), which in this simple case is simply the double difference

$$D_2\Phi \equiv \Phi_{11} + \Phi_{22} - \Phi_{12} - \Phi_{21}.$$

From this admittedly imprecise argument we derive the following conjecture:

Conjecture: mistakenly assuming separability generates "small" misspecification biases on the estimated supermodular core $D_2\Phi$ of the joint surplus.

3.2.2 Monte Carlo Setup

To test our conjecture, we simulated average-size samples of individuals with varying characteristics and matching preferences. Our simulation framework takes in 1,000 individuals, equally split between men and women[8]. We assign to each one of two educational levels; we examine both symmetric and asymmetric margins. In the symmetric case, there are 250 men and 250 women in each of the two groups (HS and CG). The asymmetric scenario assumes a larger number of college educated women (375) than men (125); conversely, it has a smaller number of high-school educated women (125) than men (375).

Recall that we defined a "non-separability" R^2. We calibrate our scenarii so that this R^2 takes values of $0, 0.2, 0.4, 0.6, 0.8,$ and 1. This covers the range from the fully separable Choo and Siow (2006) model to a "fully random" surplus (drawn independently of the observed types). The intermediate cases represent a mildly non-separable model ($R^2 = 0.2, 0.4$) and a strongly non-separable one ($R^2 = 0.6, 0.8$).

All of our simulations pre-impose a supermodular and symmetric systematic surplus function Φ. We distinguish two cases: "small modularity", with

$$\Phi = \begin{pmatrix} 0.5 & 1.0 \\ 1.0 & 1.6 \end{pmatrix}$$

and "large modularity", where we use

$$\Phi = \begin{pmatrix} 0.5 & 1.0 \\ 1.0 & 2.5 \end{pmatrix}.$$

[8] We will also report more briefly on a scaled-down population of 200 individuals.

Note that the only difference lies in the surplus generated by a couple of college-educated partners, which is larger in the latter case. As a consequence, the super-modular core

$$D_2\Phi \equiv \Phi_{11} + \Phi_{22} - \Phi_{12} - \Phi_{21}$$

equals 0.1 with small modularity and 1.0 with large modularity.

As a result of selecting two different specifications of modularity and two contrasting population margins, we arrive at four distinct calibrations of the variance of the systematic surplus Φ. In the "large market" limit of the separable Choo and Siow (2006) model when $R^2 = 0$, the estimate of the variance of Φ in the case of small modularity is about 0.357 with symmetric margins and roughly 0.336 with asymmetric margins. In the case of large modularity, they are more than twice as high: about 0.856 with symmetric margins and roughly 0.716 with asymmetric margins.

To complete the description of our simulations, we need to describe the specification of the ε, η, and (for the non-fully separable cases $R^2 > 0$) also ν. We choose to draw all $\varepsilon_y^i, \eta_x^j$, and ν_{ij} independently from the centered standard type I extreme value distribution such that when $R^2 = 0$, this is just the Choo and Siow (2006) model; scenarii with positive R^2 explore its robustness to deviations from separability.

For each simulation scenario, we generate 1,000 datasets. Table 1 summarizes the simulation scenarii.

Population:	1,000			(Separable)	$R^2 = 0$
Draws:	1,000			(Non-Separable)	$R^2 = 0.2, 0.4, 0.6, 0.8, 1$
Modularity:	Small or Large				
		Symmetric Margins		**Asymmetric Margins**	
		Count	Share of Population	Count	Share of Population
Men	HS $(x = 1)$	250	25%	375	37.5%
Men	CG $(x = 2)$	250	25%	125	12.5%
Women	HS $(y = 1)$	250	25%	125	12.5%
Women	CG $(y = 2)$	250	25%	375	37.5%

Table 1: Simulation Parameters

3.2.3 Monte Carlo Results

The combination of six values of R^2, symmetric or asymmetric populations, and two modularity subcases generates 24 different scenarii. As going through all of our results would quickly bore the reader, we focus on the most striking ones. Sometimes we only show plots for the small modularity/symmetric margins case to save space.

With non-separable surplus, the only way we know of solving for the equilibrium matching is to solve the linear programming problem associated with the primal (maximizing the total surplus) or dual (minimizing the sum of utilities under the stability constraints). We used the `R` interface of the free academic version of the Gurobi software[9] for this purpose. The algorithm converges very robustly. Even with our relatively small populations of 1,000 individuals, each run requires several gigabytes of memory. This stands in contrast with the separable case, for which very efficient methods exist—most notably the IPFP algorithm of Galichon and Salanié (2022).

The $R^2 = 0$ simulation serves as our benchmark, since in that case the model is well-specified. Biases and inefficiencies introduced by the misspecifications for $R^2 > 0$ will show up in a translation and a spreading out of the estimated density of our estimators of the four elements of the Φ matrix, and of the supermodular core $D_2\Phi$.

Figures 1 and 2 plot the distributions of the estimated $\hat{\Phi}_{xy}$ for respectively $(x, y) = (1, 1), (1, 2), (2, 1)$, and $(2, 2)$. As our discussion in section 3.2.1 suggested, the estimators have a positive bias that grows with the extent of the non-separability. This reflects the growing thickness of the market and the resulting higher probability of finding a suitable partner. This is apparent in Figure 3: as the joint surplus becomes more non-separable, each group gets better outcomes.

More interestingly, Figures 4 and 5 show that our estimate of the supermodulaar core $D_2\Phi$ has very little bias, even when R^2 becomes as large as 0.6. Recall that for this value of R^2, the non-separable term ν_{ij} contributes half more variance than the sum of the separable terms ε^i_y and η^j_x. This finding is important as we have known since Becker (1973) that the supermodularity of the joint surplus drives the essential properties of the matching. It points to a remarkable robustness of the Choo and Siow (2006) estimator in the face of rather large deviations from separability.

Another important property of the matching equilibrium are its "equilibrium prices": how it shares the joint surplus between the two partners in any couple that forms in equilibrium. Figures 6 and 7 show the distribution of the share that goes to the man on average, that is

$$\frac{u_x}{u_x + v_y}$$

for all types of couples (depending whether either partner is college-educated). We focus on the symmetric case, where the shares should be close to 0.5 when the partners are equally educated. Both figures confirm this; again, the lack of bias in

[9] Gurobi (2019).

the estimated surplus shares is striking, even with $R^2 = 1$ when the basic structure of the Choo and Siow (2006) model vanishes.

With symmetric populations and supermodularity, one expects the couples in which one partner is better-educated than the other to attribute a smaller share of the surplus to the less-educated partner; and the difference should grow with the supermodular core. Comparing Figures 6 and 7 shows that while men do almost as well with $x = 1, y = 2$ as with $x = 2, y = 1$ when the modularity is small, the difference becomes more noticeable with large modularity.

Since solving the linear programming problem is much more expensive than using the IPFP algorithm proposed by Galichon and Salanié (2022), it is also interesting to compare their performances in our simulation runs. While we already know that IPFP is much faster (and uses very little memory), it is only rigorously valid in the "large market" limit when the number of individuals goes to infinity. Figure 8 shows that even with our population size of 1,000, IPFP performs remarkably well: it yields equilibrium matching probabilities μ_{xy} that are very close to the mode of the distribution of the Gurobi results for the finite-population case.

We also ran a set of simulations with only 200 individuals. The results are basically similar as with 1,000 individuals, with more variation across runs as expected.

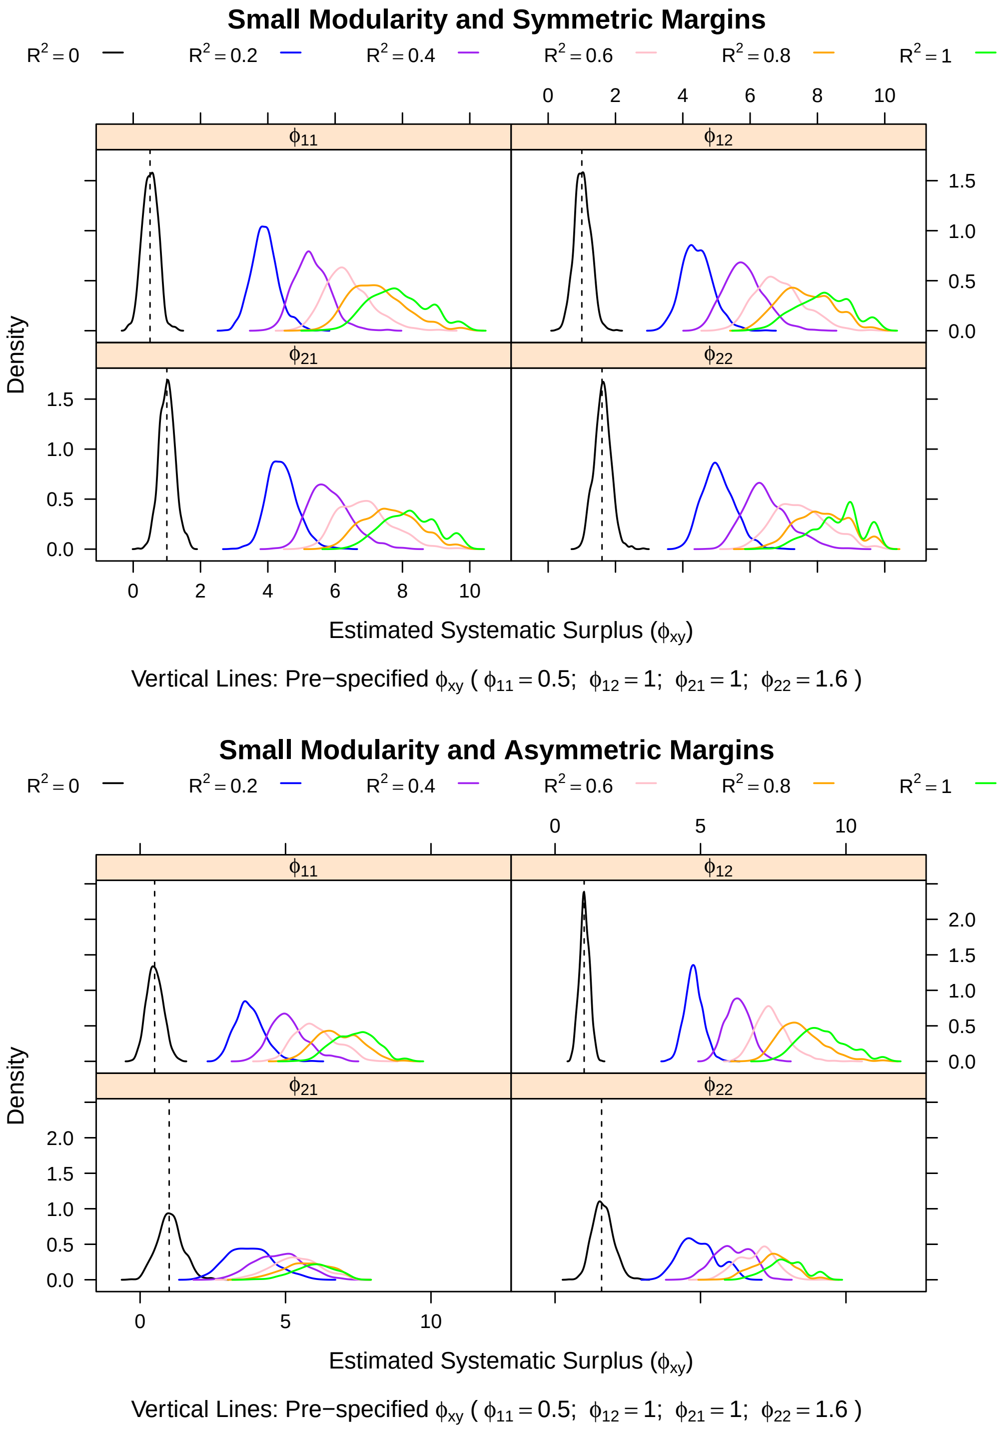

Figure 1: Estimates of Φ—small modularity

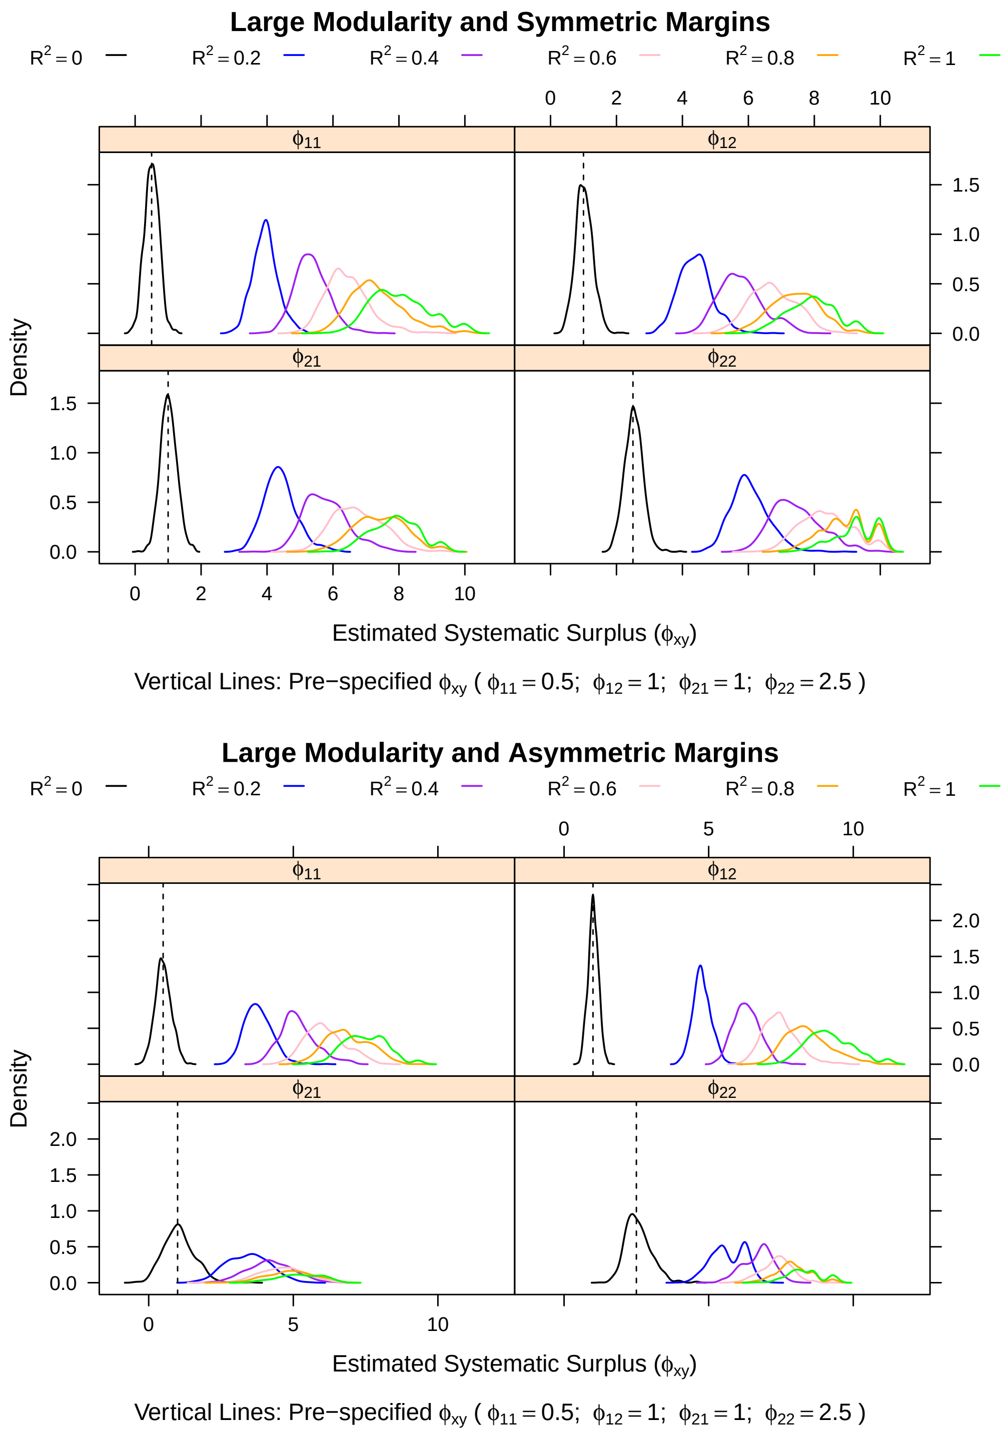

Figure 2: Estimates of Φ—large modularity

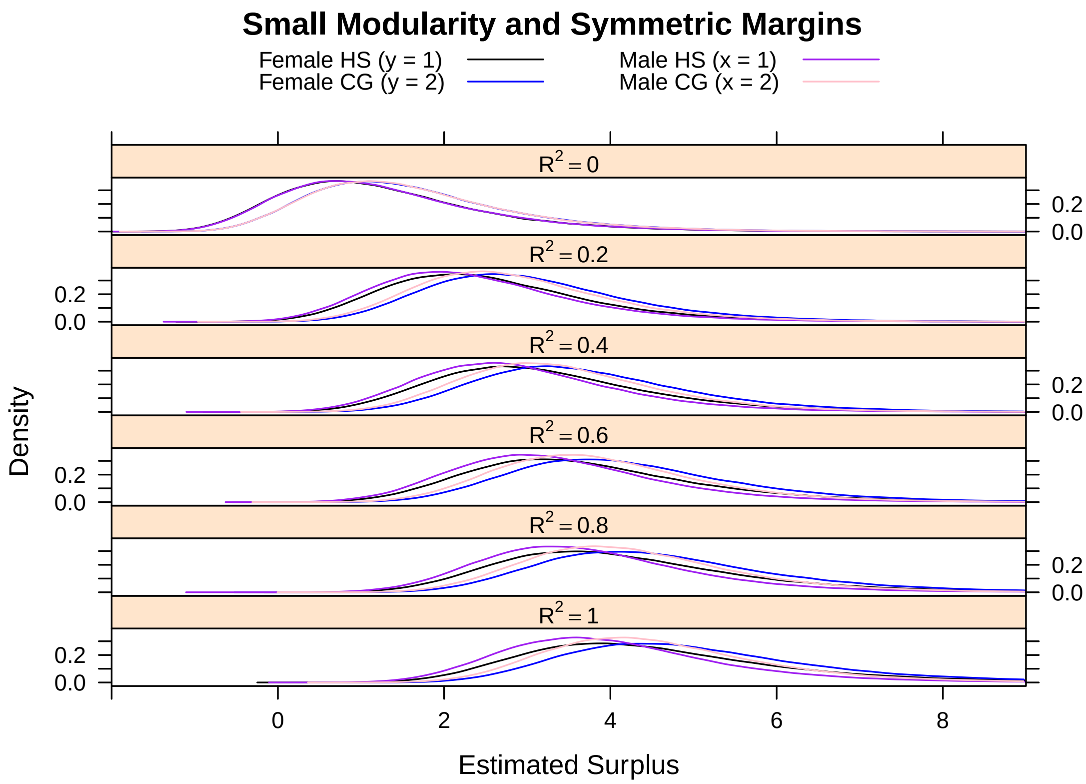

Figure 3: Estimates of u_x and v_y—symmetric, small modularity

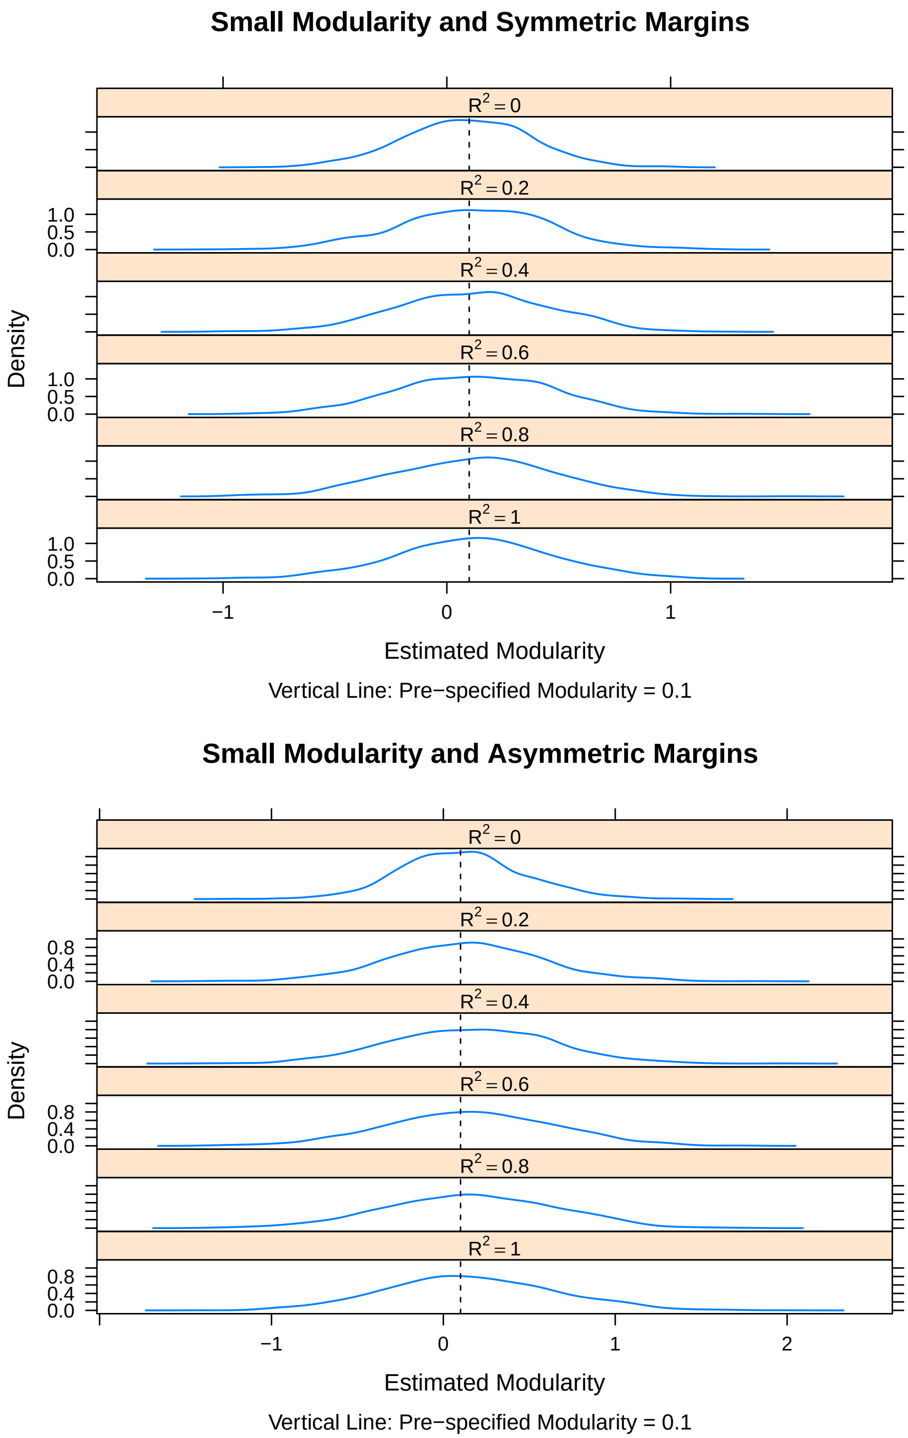

Figure 4: Estimates of $D_2\Phi$—small modularity

Large Modularity and Symmetric Margins

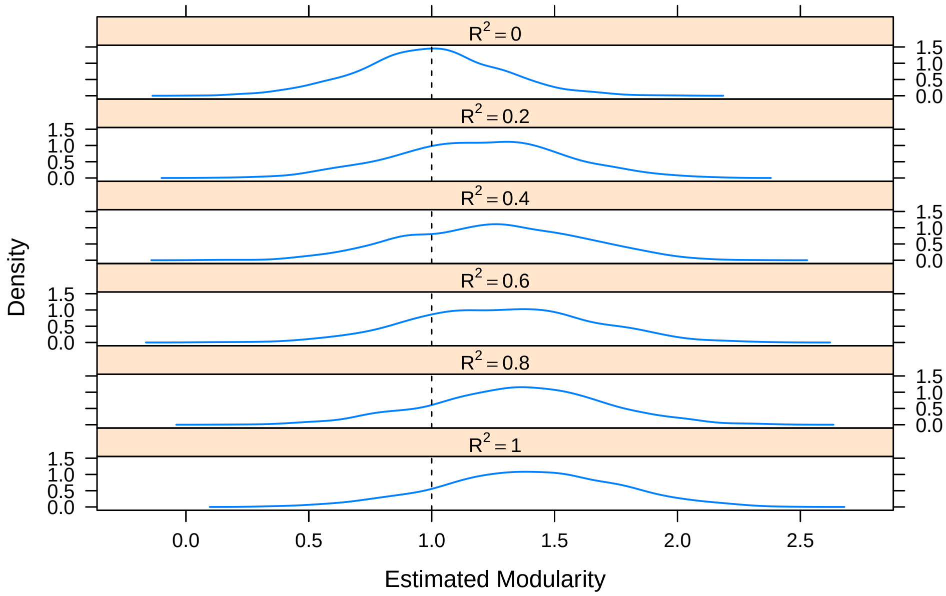

Vertical Line: Pre−specified Modularity = 1

Large Modularity and Asymmetric Margins

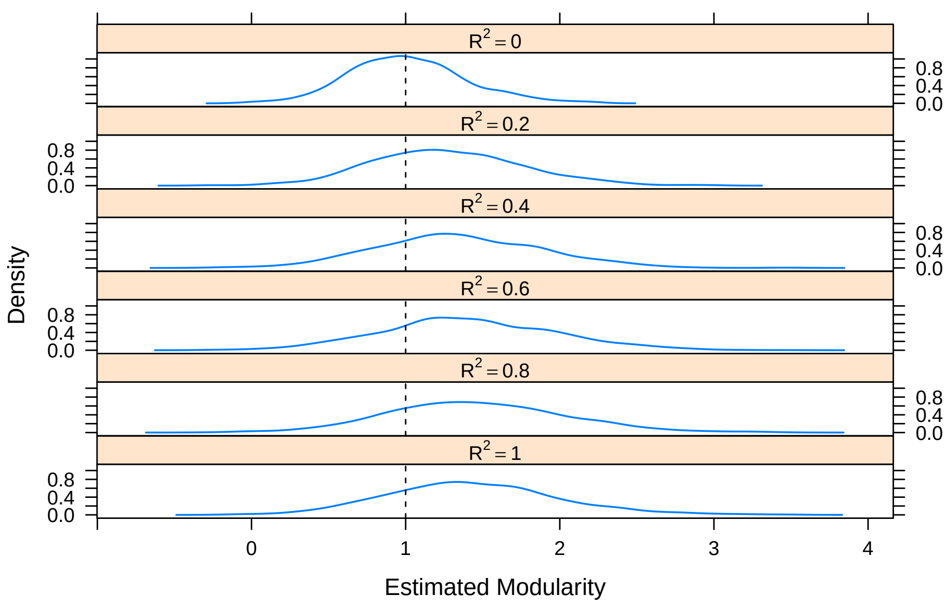

Vertical Line: Pre−specified Modularity = 1

Figure 5: Estimates of $D_2\Phi$—large modularity

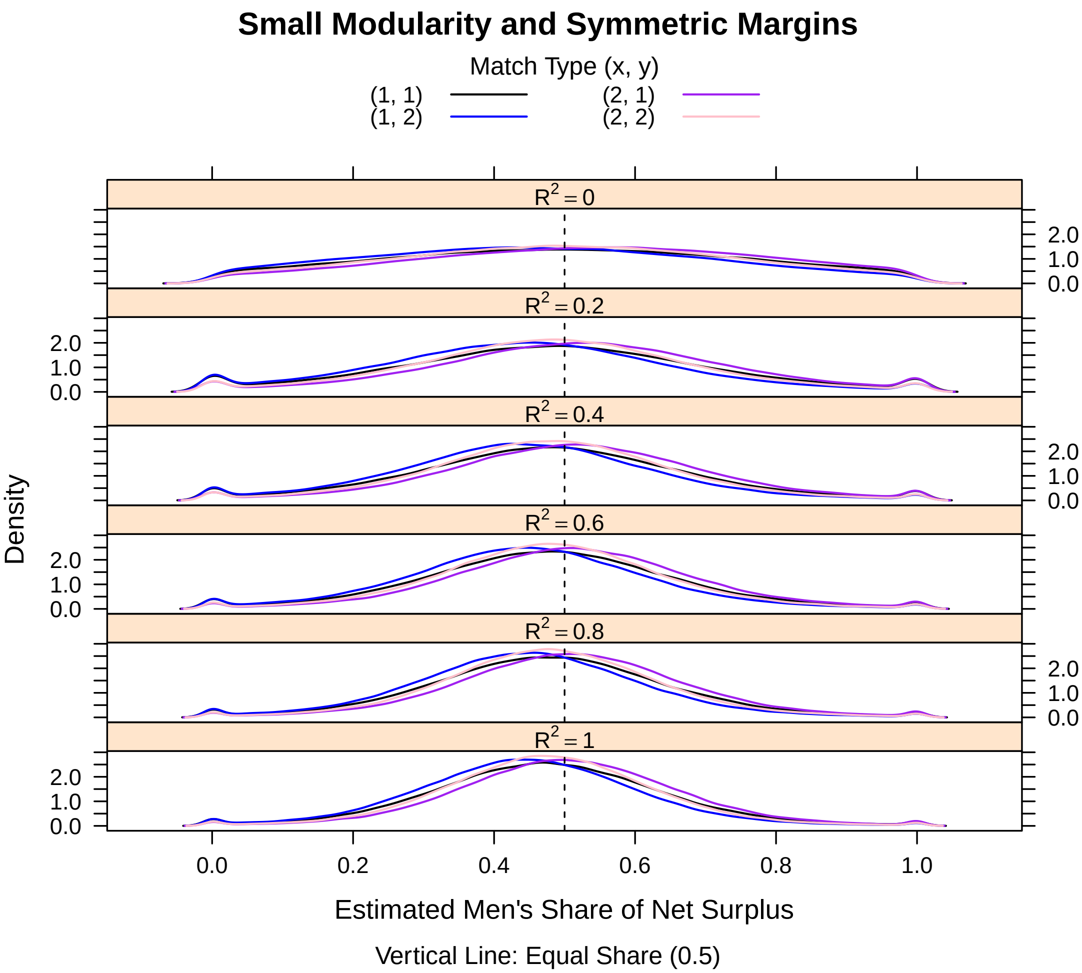

Figure 6: Estimates of Men's Shares—symmetric, small modularity

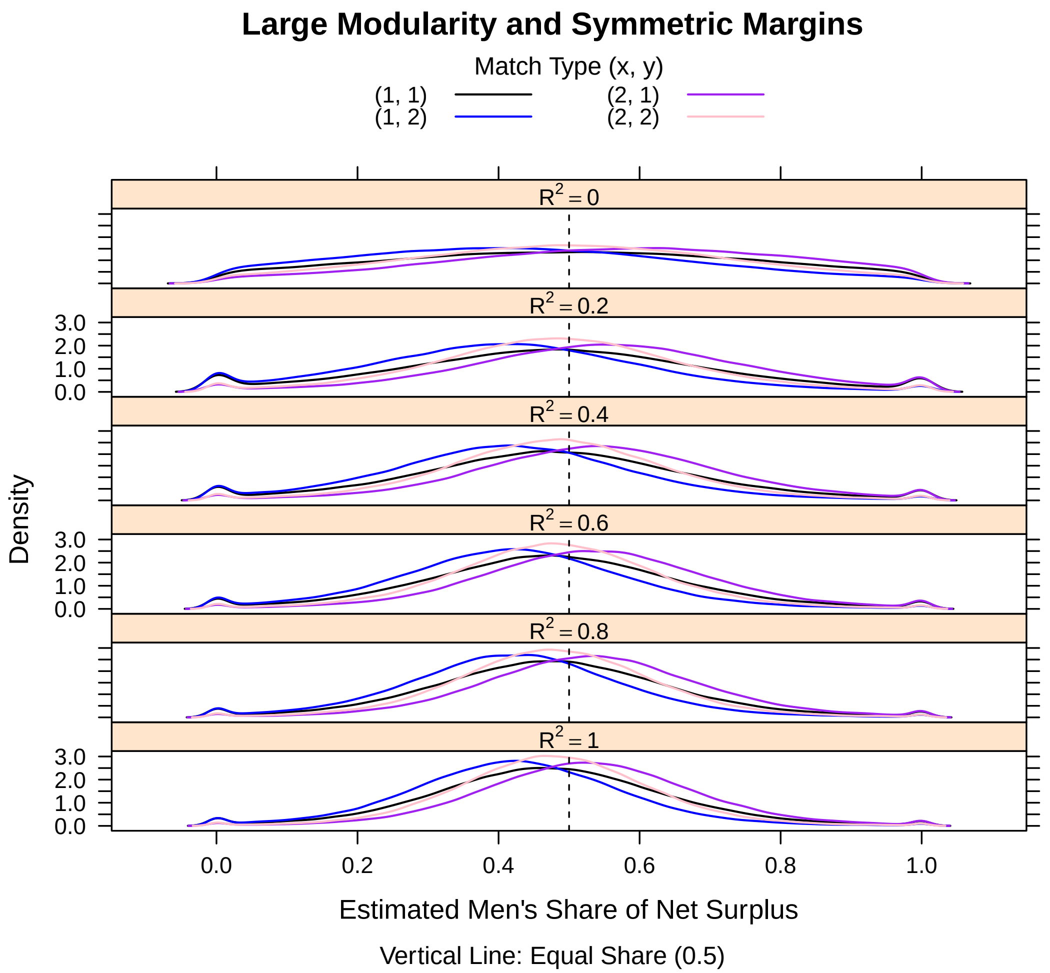

Figure 7: Estimates of Men's Shares—symmetric, large modularity

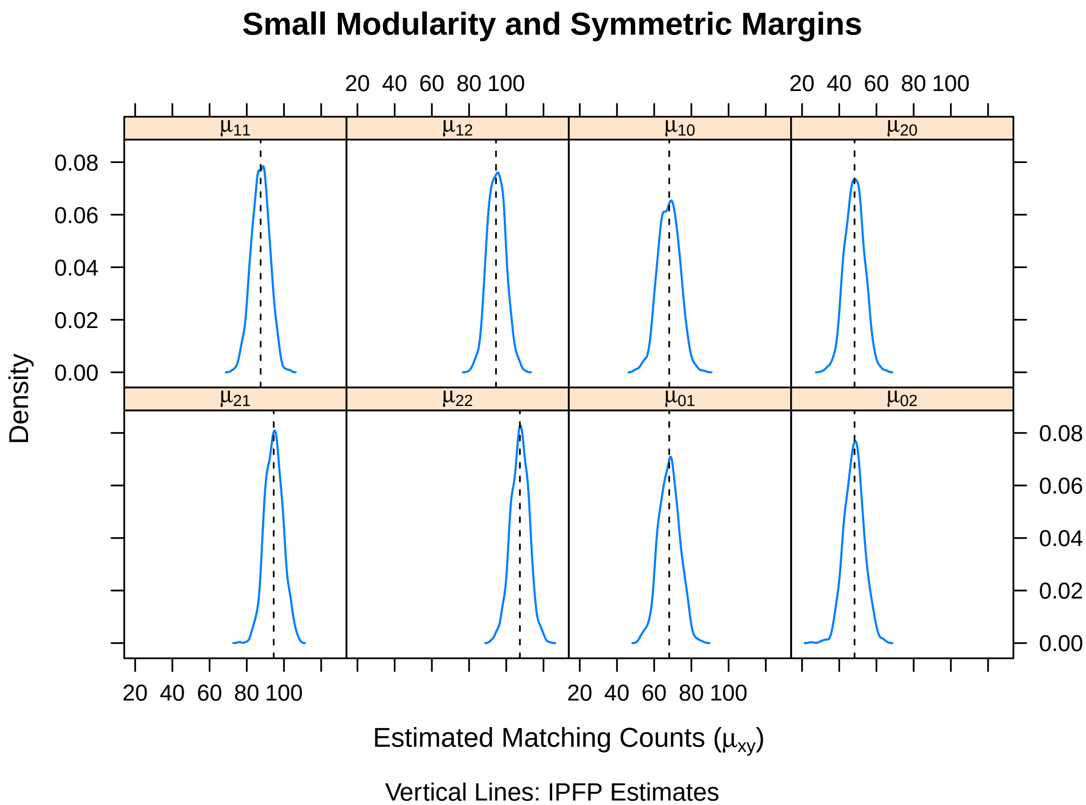

Figure 8: Estimates of μ_{xy}—symmetric, small modularity

4 A Caveat on Imbalanced Data

The estimation of matching models, separable or not, is very sensitive to the presence of large variation in cell sizes. To take a very bare-bones example, consider a Choo and Siow model with $X = Y = 2$ observed types and suppose that we have no data on singles: we only observe the proportion of marriages $0 < \hat{\mu}_{xy} < 1$ in each cell $x, y = 1, 2$. Then we can only estimate the double difference

$$\phi_0 = \Phi_{11} + \Phi_{22} - \Phi_{12} - \Phi_{21}.$$

The maximum likelihood estimator of ϕ_0 is simply

$$\hat{\phi} = 2 \log \frac{\hat{\mu}_{11}\hat{\mu}_{22}}{\hat{\mu}_{12}\hat{\mu}_{21}}.$$

It is easy to check that given a sample of n marriages, the asymptotic approximation to its variance is

$$V\hat{\phi}_n = \frac{4}{n}\left(\frac{1}{\mu_{11}} + \frac{1}{\mu_{12}} + \frac{1}{\mu_{21}} + \frac{1}{\mu_{22}}\right),$$

where the μ_{ij} are the asymptotic limits of the equilibrium proportions. It is clear in this form that since $\mu_{11} + \mu_{12} + \mu_{21} + \mu_{22} = 1$, the variance of the estimator is always at least $16/n$. To get an estimate of ϕ_0 with a standard error of say 0.1, we need $1,600$ marriages in the best of cases (that is, when $\phi_0 = 0$ so that all μ_{ij} equal 1/4.). It it not hard to show that this bound also holds with X equally probable types on each side (rather than just two)[10].

Even more importantly, a larger dispersion in cell sizes increases the variance of the estimator. In fact,

$$V\hat{\phi}_n > \frac{4}{n} \frac{1}{\min(\mu_{11}, \mu_{12}, \mu_{21}, \mu_{22})}:$$

the presence of *any* small cell will make the estimator imprecise.

It may not be so surprising, a posteriori, that estimating a multinomial two-sided choice model requires fairly large samples; and we know that even in one-sided models, small cells create problems for estimation. This difficulty is more salient in matching models, as cells with small numbers of observations are often found in the data. Marriage markets are a good illustration: e.g. Choo and Siow (2006) used data with individuals ages 16 to 75, with non-zero cells clustered near the diagonal.

[10]To be precise, the variance is $4/(nD_0(1 - D_0))$, where D_0 is the proportion of marriages with $x = y$.

Galichon and Salanié (2024a, Section 5.1) is a first stab at some possible strategies to deal with zero cells. Non-zero but small cells may be easier to deal with: they could be trimmed out of the data, cautiously inflated, or weighted down in the estimation. This is an important topic that clearly requires more work.

Concluding Remarks

This paper provides a balanced assessment of the empirical strengths and limitations of separable matching models. Our theoretical arguments and our Monte Carlo simulations suggest (cautiously) that omitted variables and omitted non-separabilities induce misspecification biases that do not seem to be severe. Estimating a separable, Choo and Siow (2006) model on data that may have been generated by a non-separable model (or a separable model with some unobserved variables) does little apparent harm to our ability to get reliable estimates of the most economically important statistics: the supermodular core and the surplus shares.

Our main caveat is that when the data contain cells that vary widely in size, estimators are likely to be very imprecise. It may in fact be best to exclude some small cells from the estimation, to inflate their size, or to smooth cell sizes in some way. As such data are a common occurrence, more work is needed to explore which of these strategies are most useful.

References

Ahn, S. Y. (2023): "Matching Across Markets: An Economic Analysis of Cross-Border Marriage," *Journal of Labor Economics*, 43.

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

Bojilov, R., and A. Galichon (2016): "Matching in closed-form: equilibrium, identification, and comparative statics," *Economic Theory*, 61, 587–609.

Chiappori, P.-A., C. Florio, A. Galichon, and S. Verzillo (2024): "Assortative Matching on Income," *Econometrica*, forthcoming.

Chiappori, P.-A., and B. Salanié (2016): "The Econometrics of Matching Models," *Journal of Economic Literature*, 54, 832–861.

Chiappori, P.-A., B. Salanié, and Y. Weiss (2017): "Partner Choice, Investment in Children, and the Marital College Premium," *American Economic Review*, 107, 2109–67.

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

Ciscato, E., A. Galichon, and M. Goussé (2019): "Like Attract Like: A Structural Comparison of Homogamy Across Same-Sex and Different-Sex Households," *Journal of Political Economy*, 128, 740–781.

Corblet, P. (2022): "Education Expansion, Sorting, and the Decreasing Education Wage Premium," Discussion paper, Sciences Po Paris, mimeo.

Dupuy, A., and A. Galichon (2014): "Personality traits and the marriage market," *Journal of Political Economy*, 122, 1271–1319.

Fox, J., C. Yang, and D. Hsu (2018): "Unobserved Heterogeneity in Matching Games with an Appplication to Venture Capital," *Journal of Political Economy*, 126, 1339–1373.

Galichon, A., S. Kominers, and S. Weber (2019): "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility," *Journal of Political Economy*, Forthcoming.

Galichon, A., and B. Salanié (2017): "The Econometrics and Some Properties of Separable Matching Models," *American Economic Review Papers and Proceedings*, 107, 251–255.

Galichon, A., and B. Salanié (2019): "Labeling Dependence in Separable Matching Markets," Columbia University mimeo.

——— (2022): "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," *Review of Economic Studies*, 89, 2600–29.

——— (2024a): "Estimating Separable Matching Models," *Journal of Applied Econometrics*, 39, 1021–44.

——— (2024b): "Hedonic and Matching Models," in *prepared for the Handbook of Econometrics, vol. 7B*, ed. by L. Hansen, J. Heckman, and R. Matzkin. North Holland.

Guadalupe, M., V. Rappoport, B. Salanié, and C. Thomas (2024): "The Perfect Match: Assortative Matching in Mergers and Acquisitions," Columbia mimeo.

Gualdani, C., and S. Sinha (2023): "Partial identification in matching models for the marriage market," *Journal of Political Economy*, 131, 1109–1171.

Gurobi (2019): "Gurobi Optimizer Reference Manual," Discussion paper, Gurobi Optimization, Inc.

Menzel, K. (2015): "Large Matching Markets as Two-Sided Demand Systems," *Econometrica*, 83, 897–941.

Mourifié, I. (2019): "A Marriage Matching Function with Flexible Spillover and Substitution Patterns," *Economic Theory*, 67, 421–461.

Mourifié, I., and A. Siow (2017): "The Cobb Douglas Marriage Matching function: Marriage Matching with Peer and Scale Effects," University of Toronto mimeo.

Salanié, B. (2015): "Identification in Separable Matching with Observed Transfers," Columbia University mimeo.

Salanié, B. (2024): "Matching with Transfers: Applications," in *Handbook of the Economics of Matching, vol. 1*, ed. by Y.-K. Che, P.-A. Chiappori, and B. Salanié. North Holland.

Tinbergen, J. (1956): "On the Theory of Income Distribution," *Weltwirtschaftliches Archiv*, 77, 155–73.