Matching with Random Components: Simulations[*]

Pierre-André Chiappori[†] Dam Linh Nguyen[‡] Bernard Salanié[§]

Draft of December 10, 2019

Abstract

Several recent papers have analyzed matching markets under the dual assumption of perfectly transferable utility and a separable joint surplus. Separability rules out any contribution to the joint surplus of a match of interactions between characteristics of partners that are unobserved by the analyst. Since it may be unrealistic in some settings, we explore the consequences of mistakenly imposing it. We find 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.

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

[*]Columbia University. Email: pc2167@columbia.edu
[†]NERA Economic Consulting. Email: n.linh@columbia.edu
[‡]Columbia University. Email: bsalanie@columbia.edu

up with very different matching outcomes: some may be unmatched, 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 start from 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, 2019) study the properties of separable models in much more detail. 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. It has been 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 definitely a very useful assumption, 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 goal of this note 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 model.

Our first group of findings shows that non-separability impacts matching patterns and utilities in ways that, ex post, seem reasonable. It takes a rather large amount of non-separability to see a qualitative difference from the separable case, however. Second, 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 remainder of the paper is organized as follows. Section 1 introduces the set up, including the model and methods. Section 2 presents our simulation protocol and section 3 discusses the results. 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. 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.

1.1 The Model

To make things as simple and transparent as possible, each individual has only one, binary observed characteristic. To continue with our analogy to the marriage market, we refer to this observable variable as the education of the individual, which we code as 1 for "high school or less" and 2 for "at least some college". For simplicity, we will refer to them as high-school (HS) and college (CG).

Men are indexed by $i \in \mathcal{I} = \{1, 2\}$ and women by $j \in \mathcal{J} = \{1, 2\}$. Each man i (resp. woman j) has an education x_i (resp. y_j) in $\{1, 2\}$. This is observable to all men, all women, and to the analyst. When a man i and a woman j marry, their match generates a joint 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 $\tilde{\Phi}_{0j}$.

1.1.1 Separability

The separable model restricts the form of the surplus $\tilde{\Phi}_{ij}$. It imposes that there exist a matrix $\mathbf{\Phi} = (\Phi_{xy})_{x,y=1,2}$ and random variables $\boldsymbol{\varepsilon}^{\boldsymbol{i}} = (\varepsilon^i_0, \varepsilon^i_1, \varepsilon^i_2)$ and $\boldsymbol{\eta}^{\boldsymbol{j}} = (\eta^j_0, \eta^j_1, \eta^j_2)$ 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 $\boldsymbol{\varepsilon}^{\boldsymbol{i}}$ (resp. $\boldsymbol{\eta}^{\boldsymbol{j}}$) are drawn independently of each other, conditional on x_i (resp. y_j).

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. It does *not* rule out "matching on unobservables": a man i of education x and a woman j of education y are more likely to marry if his ε^i_y and her η^j_x take higher values. As shown by Chiappori, Salanié, and Weiss (2017), separability implies that if this man and this woman do marry in equilibrium, he gets utility

$$U_{xy} + \varepsilon^i_y$$

and she gets $V_{xy} + \eta^j_x$, where the split $U_{xy} + V_{xy} = \Phi_{xy}$ is determined in equilibrium. Therefore man i would be equally happy with any other woman of education y and woman j would be equally happy with any other man of type x. Galichon and Salanié (2017, 2019) study the class of separable models in more depth. They derive identification results and propose estimators.

As can be seen from this brief discussion, separability is more restrictive when the data contain little information on types and/or the unobserved types are relevant to the application under study. It is therefore important to evaluate the consequences of its failure, starting from the most popular separable model.

1.1.2 The Choo and Siow Specification

The best-known and most convenient separable model is 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; and that the number of individuals in each education category is very large. They showed that in this "large market" equilibrium, 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 4 equations in 8 unknowns; the system is completed by the scarcity constraints

$$I_x = \sum_{y=1,2} \mu_{xy} + \mu_{x0}$$
$$J_y = \sum_{x=1,2} \mu_{xy} + \mu_{0y}$$

for $x, y = 1, 2$, where I_x (resp. J_y) denotes the number of men of education x (resp. women of education y) and $I_1 + I_2 = I, J_1 + J_2 = J$.

Galichon and Salanié (2019) show how given the "margins" $(I_x), (J_y)$ and the surpluses $\tilde{\Phi}_{ij}, \tilde{\Phi}_{i0}, \tilde{\Phi}_{0j}$, this system of equations can be solved very efficiently using an Iterative Projection Fitting Procedure (IPFP).

The $U_{xy} + V_{xy} = \Phi_{xy}$ split referred to in the previous subsection takes the simple form

$$U_{xy} = \log \frac{\mu_{xy}}{\mu_{x0}}$$
$$V_{xy} = \log \frac{\mu_{xy}}{\mu_{0y}};$$

and the average utilities[1] of men of education x and women of education y are decreasing functions of their singlehood rates:

$$u_x = -\log \frac{\mu_{x0}}{n_x}$$
$$v_y = -\log \frac{\mu_{0y}}{m_y}.$$

The Choo and Siow (2006) specification has been criticized on a number of grounds; see for instance Galichon and Salanié (2017, 2019), as well as Mourifié (2019) and Mourifié and Siow (2017) for extensions that reach beyond separable models. Since it is so simple and has been widely used, it is a natural benchmark and we will adopt it as a starting point in this paper.

1.1.3 Beyond Separability

There are of course many ways to add non-separability to the Choo and Siow specification. The simplest one is to generate a matrix $\boldsymbol{\nu} = (\nu_{ij})_{i=1,\dots,I;j=1,\dots,J}$ of independent draws from some mean-zero distribution and to take ν_{ij} to represent a "pair-specific" preference shock. This obviously does not apply to singles, for whom we keep the same specification as Choo and Siow.

We will therefore simulate matching markets whose surplus is given by:

$$\tilde{\Phi}_{ij} \coloneqq \Phi_{x_i,y_j} + \tau \left(\varepsilon^i_{y_j} + \eta^j_{x_i} \right) + \sigma \nu_{ij}. \tag{2}$$

This would be exactly the separable surplus in (1) if we had $\tau = 1$ and $\sigma = 0$. The total variance of the surplus in the separable 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

[1]Here "average" means "in expectation over the unobserved type".

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 $R^2 = 1$.

1.2 The Methods

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 two conjectures.

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 $(\tilde{\boldsymbol{\Phi}}, \boldsymbol{I}, \boldsymbol{J})$, 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 compomnent 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[2]. 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 $O(\sigma^2)$. We conjecture that a less straightforward argument would yield the following, stronger result:

Conjecture 1: the effects of non-separability are at most in $O(\sigma^2)$ for small σ.

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

[2] We skip here on 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.

2 Monte Carlo Setup

2.1 Simulation Parameters

We simulate average-size samples of individuals with varying characteristics and matching preferences.

Our simulation framework takes in 1,000 individuals with a fifty-fifty split between men and women[3]. With respect to the allocation of 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 in 1.1.3 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 ($R^2 = 0.6, 0.8$).

2.2 Modularity and Distribution

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

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

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 supermodular 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

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

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 simulation, 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 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[4] for this purpose. The algorithm converges very robustly. Even

[4]Gurobi (2019).

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 (2019).

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 1.2 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 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 $M = 1, F = 2$ as with $M = 2, F = 1$ 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 Iterative Proportional Fitting Procedure (IPFP) algorithm proposed by Galichon and Salanié (2019), 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.

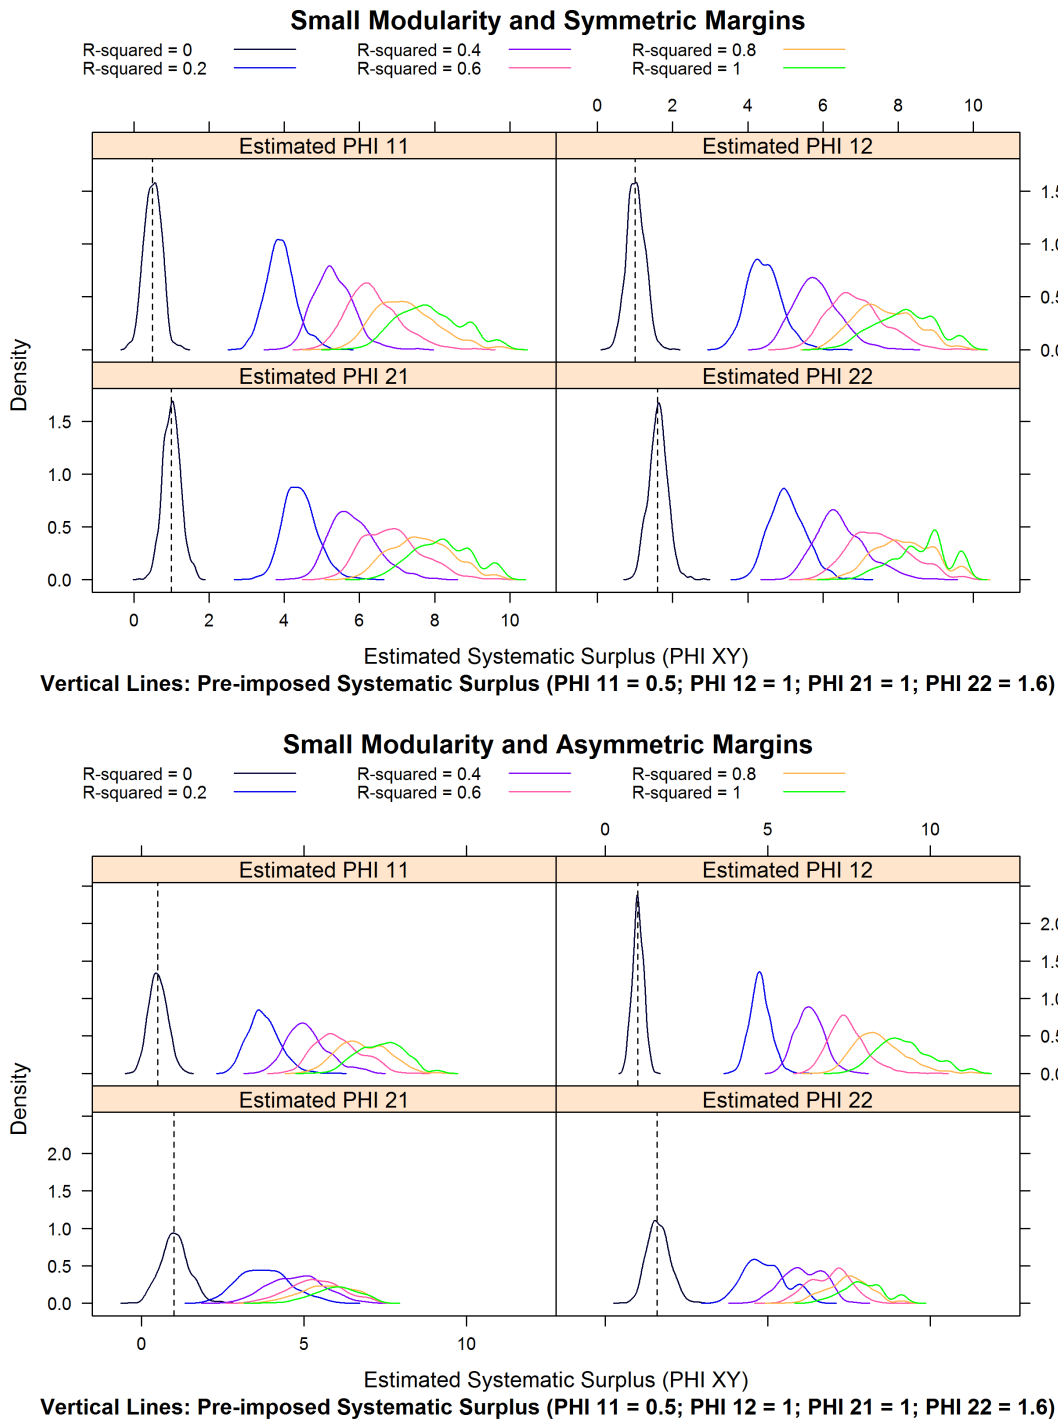

Figure 1: Estimates of Φ—small modularity

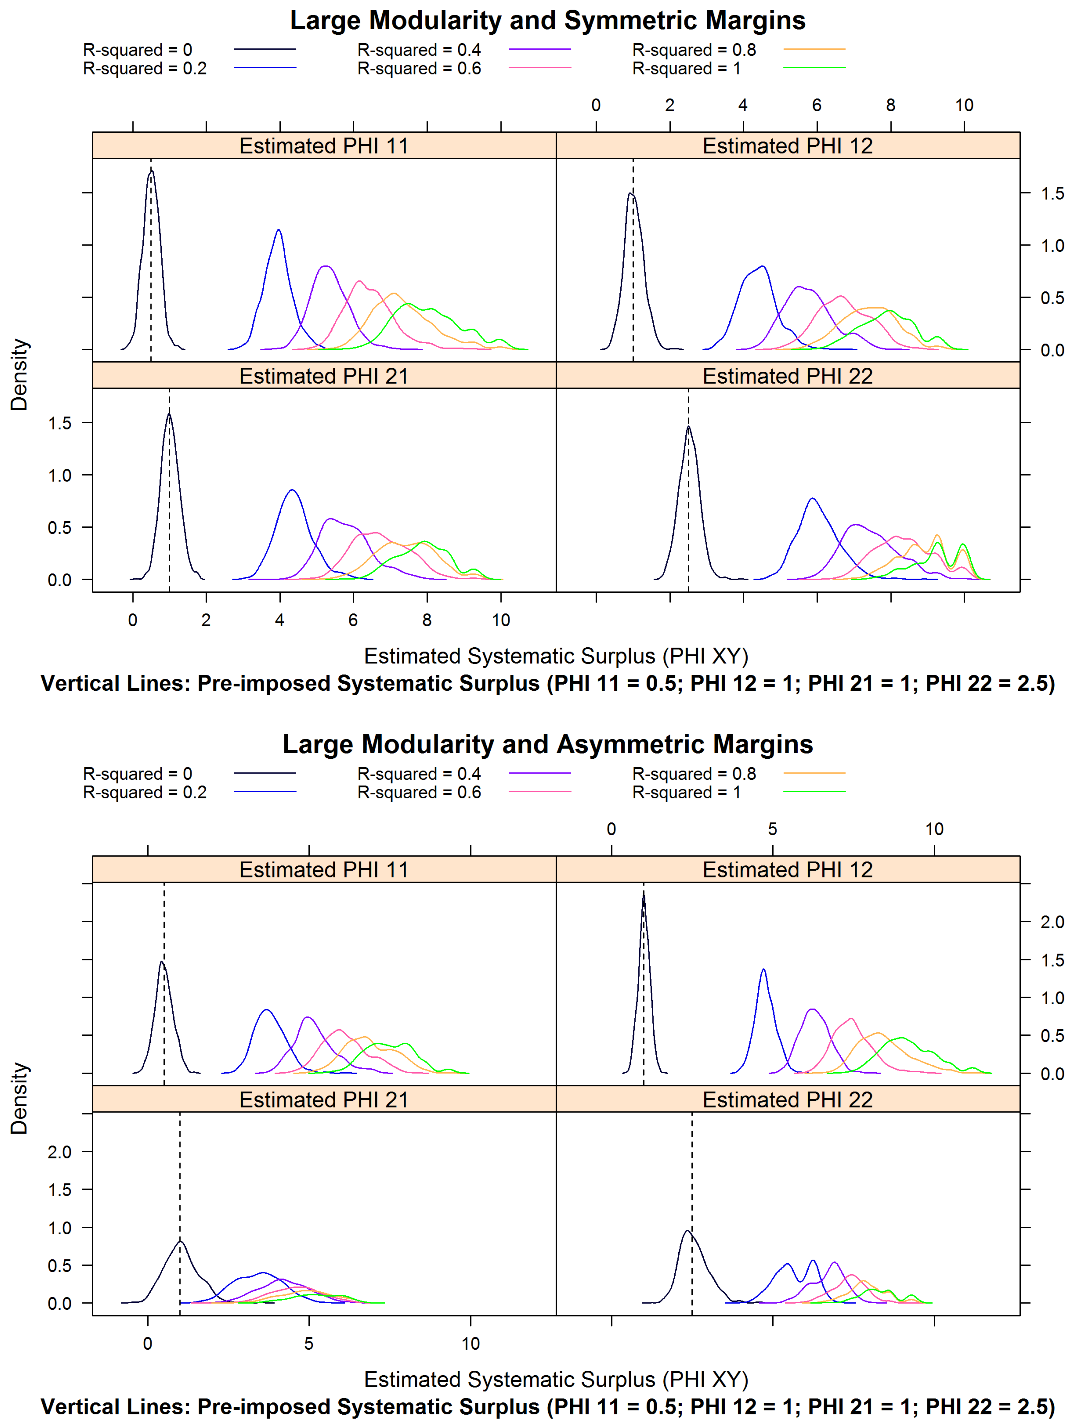

Figure 2: Estimates of Φ—large modularity

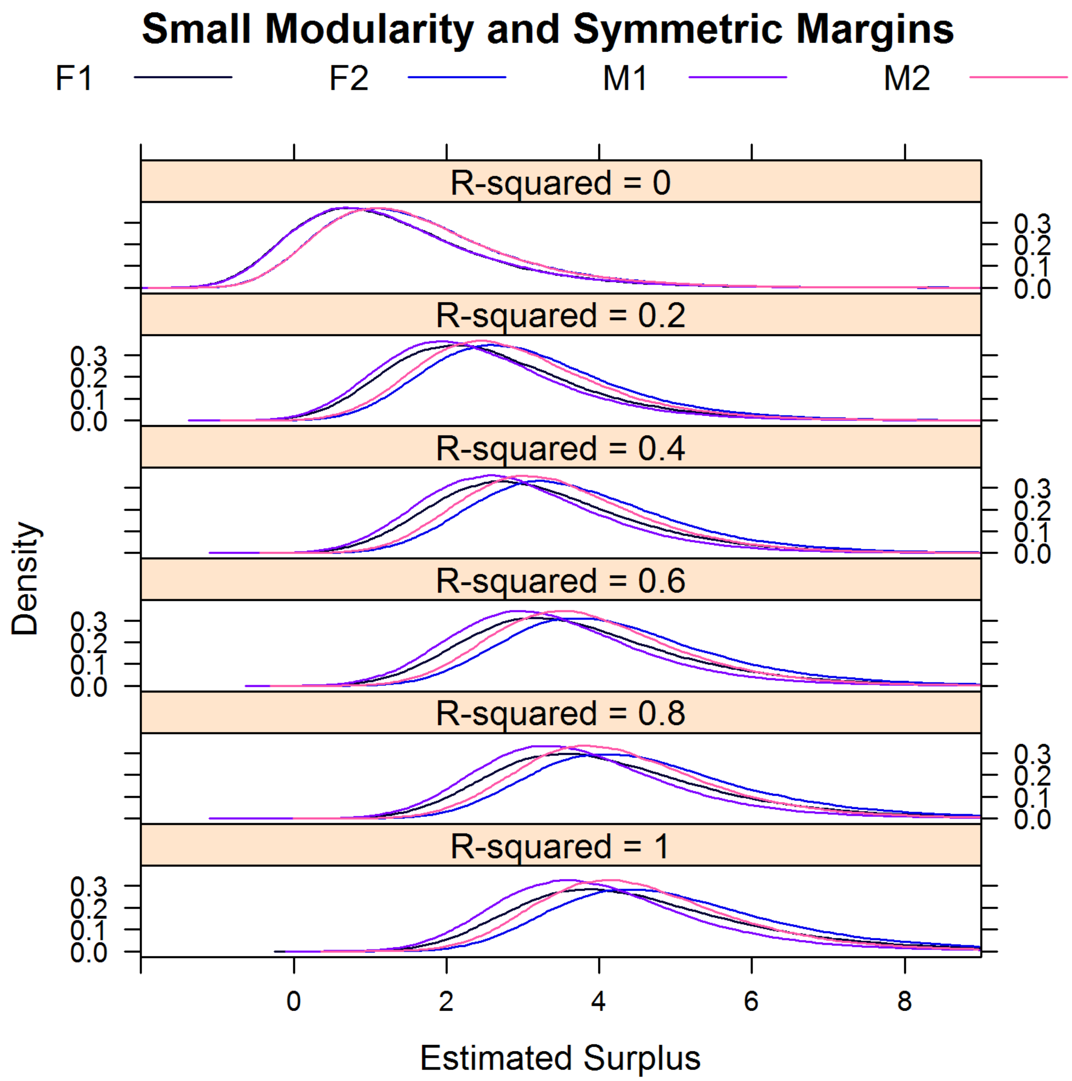

Figure 3: Estimates of u_x and v_y—symmetric, small modularity

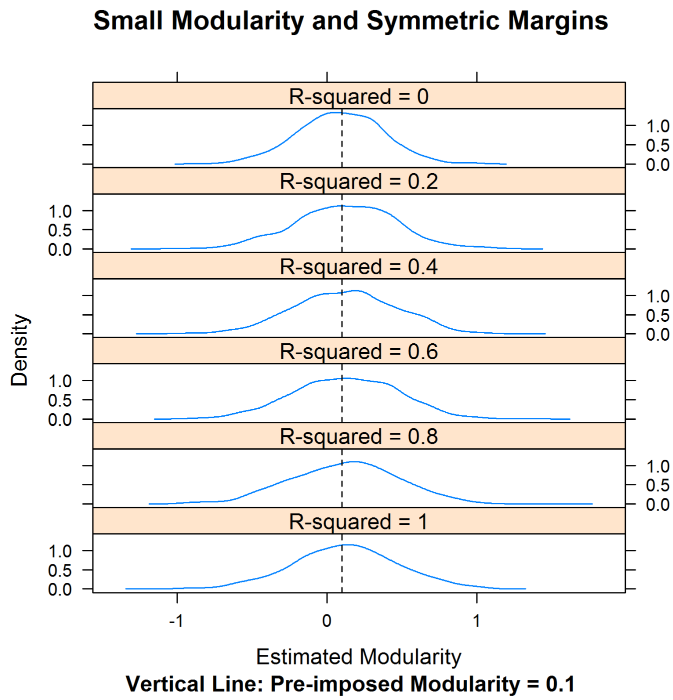

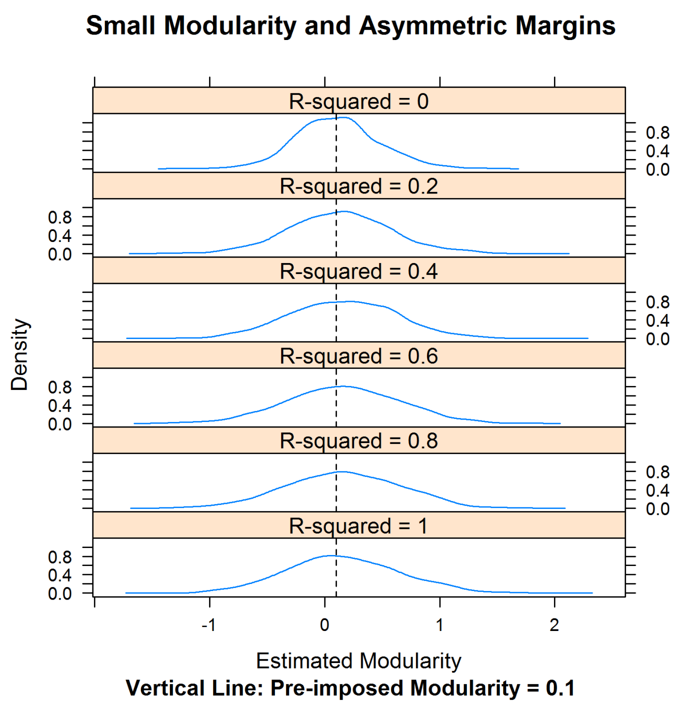

Figure 4: Estimates of $D_2\Phi$—small modularity

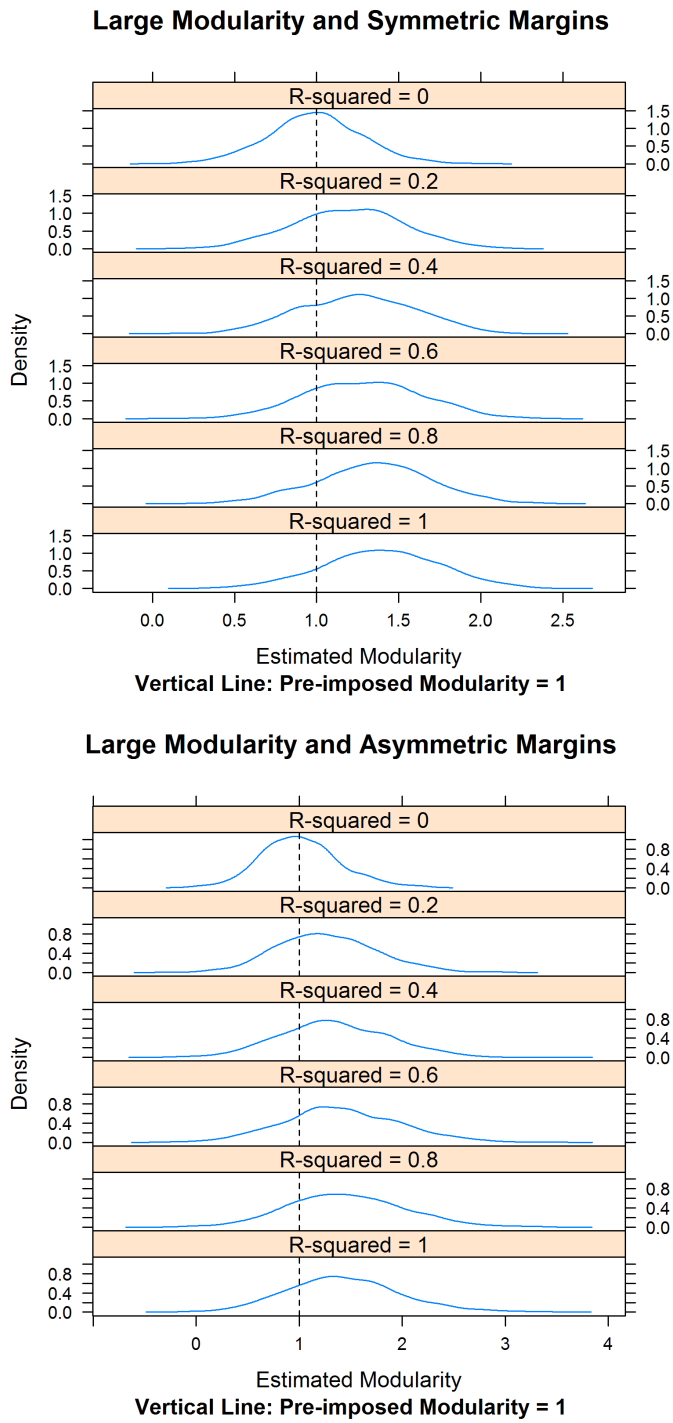

Figure 5: Estimates of $D_2\Phi$—large modularity

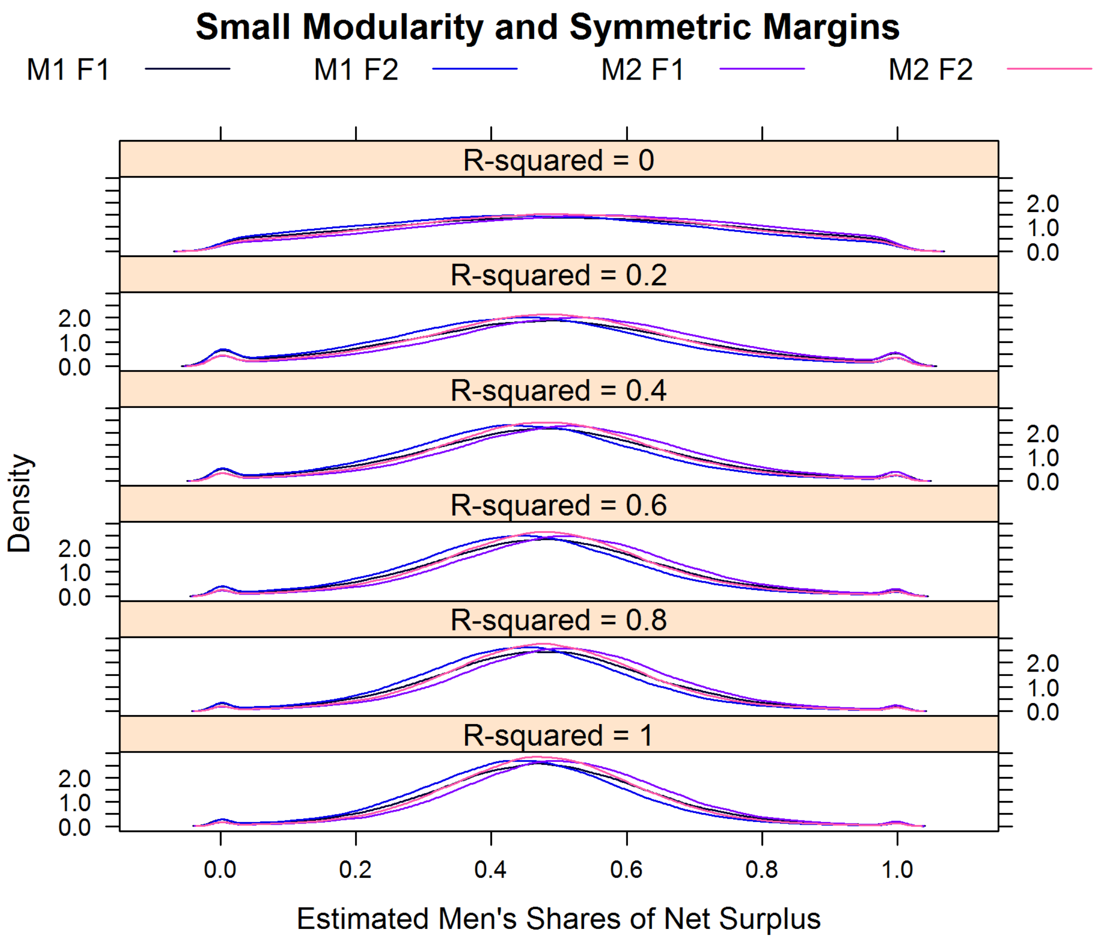

Figure 6: Estimates of Men's Shares—symmetric, small modularity

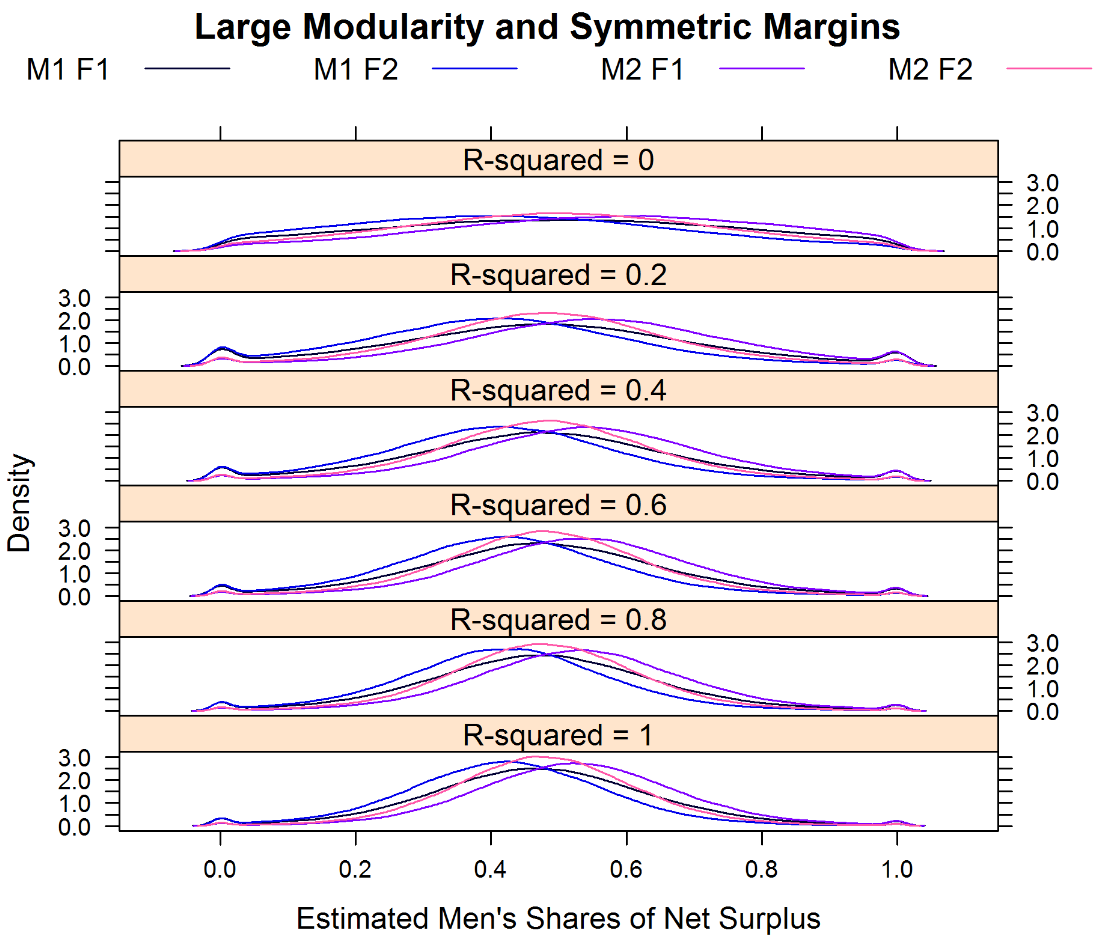

Figure 7: Estimates of Men's Shares—symmetric, large modularity

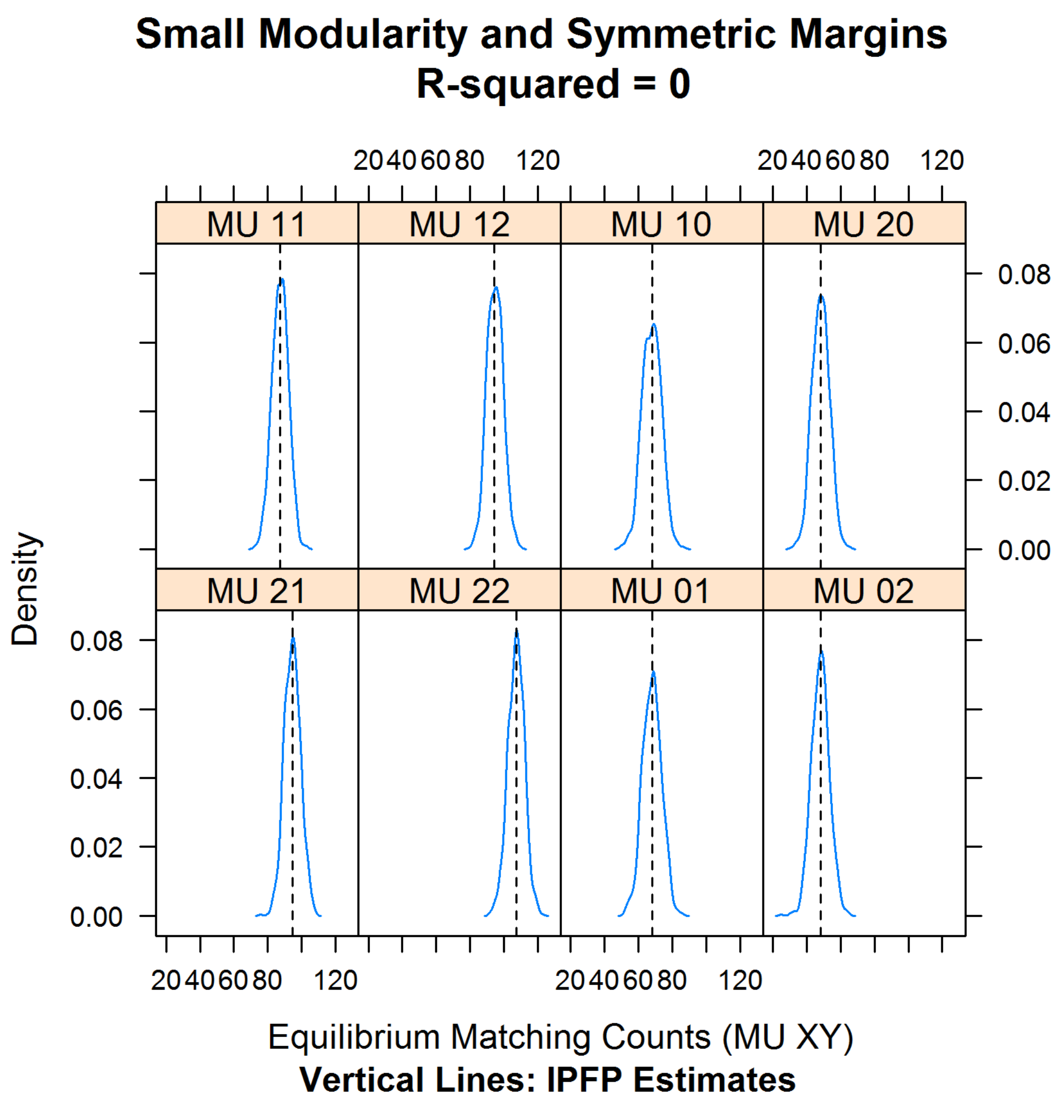

Figure 8: Estimates of μ_{xy}—symmetric, small modularity

Concluding Remarks

Even with our set of 24 scenarii, we have only explored a small area of the parameter space. Our conclusions are encouraging for separable models, however: even though they rule out interactions between unobservable characteristics, the induced misspecification biases 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 does little apparent harm to our ability to get reliable estimates of the most economically important statistics: the supermodular core and the surplus shares.

References

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

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*, Forthcoming.

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

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): "Cupid's Invisible Hand: Social Surplus and Identification in Matching Models," Columbia University mimeo.

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.