

MATCHING MARKETS WITH N-DIMENSIONAL PREFERENCES

SAM FLANDERS

University of North Carolina at Chapel Hill¹

August 2015²

ABSTRACT. This paper analyzes matching markets where agent types are points in \mathbb{R}^n and agents prefer matches that are closer to them according to a distance metric over this space (horizontal preferences). Given a few additional assumptions (most notably a form of symmetry between the two sides of the market), we show that in the Gale-Shapley stable matching in this environment agents match to a linear function of their own type. We show this result holds both when the division of match surplus is fixed (nontransferable utility), and when agents can bargain over match surplus (transferable utility) given a convexity assumption on the utility function. We also show that the restriction to horizontal preferences is not as onerous as it may seem, as a rich variety of preference structures can be mapped into the horizontal framework, including vertical preferences where all agents agree on the preference ordering of a trait. This result can be interpreted as a generalization of Becker's univariate assortative matching to multiple preference dimensions and more general preference structures.

Keywords: Transferable Utility, Non-transferable Utility, Multidimensional Heterogeneity, Assortative Matching, Closed Form

JEL codes: C78, D1

¹Email: sflander@email.unc.edu. Phone: (919) 636-9852.

²I'd like to thank Gary Biglaiser, Atila Abdulkadiroglu, Melati Nungsari, and the UNC micro theory workshop group for their indispensable advice and suggestions.

1. INTRODUCTION

In theoretical models of matching markets it is often assumed that agents have preferences over a single parameter which is either *vertical*, where all agents share a preference ordering over types; or *horizontal*, where agents prefer their own type. For both vertical (Becker (1973)) and horizontal models (Clark (2003), Clark (2007), and Klumpp (2009)), simple matching functions have been derived for continuous and discrete cases. However, it would be desirable to model multiple preference dimensions representing all the traits we believe agents have preferences over. This would allow us to make predictions about how an agent's own multivariate type will correspond to the multivariate type of their match in a real matching market. It would also allow us to explore the qualitative structure of matching over many traits, which cannot be studied in a univariate model.

In this paper we derive a simple matching function for a special case of n -dimensional horizontal preferences, where agent types are points in \mathbb{R}^n and agents prefer matches that are closer to them in terms of distance. Specifically, we consider the case where the set of agents on each side are symmetric about a separating hyperplane. Because this assumption is implausible in real world applications, we simulate matching markets with both modest and moderate deviations from the symmetry assumption and find that the theoretical results for symmetric markets well approximate the stable matching assignments observed in markets with moderate deviations from symmetry. We treat both the case where utility is nontransferable (NTU) and the case where utility is transferable (TU). In the NTU case the two matching agents cannot bargain over the apportionment of the utility of the match, while in the TU case agents can divide the match payoff between one another in any way they choose.

While this model assumes horizontal preferences, the results can easily be extended to a preference structure that includes vertical preferences, categorical horizontal preferences³,

³For example, there are several ethnic categories and agents prefer their own category.

and even more general single peaked preferences⁴, as all these preference types can be represented by horizontal preferences. Thus, these results can be applied to a matching problem where the economist observes an arbitrary number of horizontal, vertical, categorical horizontal, or single peaked preference traits. Because of this, the horizontal preference assumption is not terribly restrictive, and these results may plausibly be directly applied to real world matching markets.

These results make two primary contributions to the literature. First, they contribute to the wide empirical literature on *assortative* matching.⁵ While research on assortment has generally been confined to single traits, such as whether rich individuals marry rich individuals and poor individuals marry poor individuals, our result yields testable predictions for the structure of assortment among all traits simultaneously. Under a special case of symmetry, it predicts *positively assortative matching*⁶ (PAM) along all but one trait, and *negatively assortative matching*⁷ (NAM) along the remaining trait. More generally, it predicts that an agent's match's type will be a linear function of their own type. Lindenlaub (2013) recovers matching functions and studies assortment in a similar n-dimensional matching environment, but focuses exclusively on vertical preferences and TU.

The second contribution is to the theoretical matching literature. Univariate models are the norm in the literature because with theoretical models it is easier to work with closed form solutions, and these are much easier to obtain in a univariate model. By providing closed form matching functions for multivariate matching problems, our results open up new possibilities for analyzing matching models with multiple preference traits in a theoretical setting. There are already matching theory results for more general preference structures, such as the famous Gale-Shapley algorithm. However, while these algorithms can solve matching problems with arbitrary preferences, including multiple preference dimensions, they are iterative algorithms that do not give closed form solutions. Thus, while analyzing multivariate matching problems

⁴For example, women most prefer men who are 80% their height plus 18 inches, with preference decreasing in distance from this ideal.

⁵In one dimension, assortment is a matching structure where the type of an agent's match is monotonic in the agent's own type.

⁶PAM means an agent's match's type is monotonically increasing in the agent's own type.

⁷NAM means agent's match's type is monotonically decreasing in the agent's own type.

has been tractable in empirical settings for some time, theoretically tractable n-dimensional matching models have only begun to be studied, and our results provide extremely simple matching functions for a wide range of preferences in both TU and NTU environments.

Additionally, we find that the NTU and TU matching assignments are identical in our environment given a common assumption on match utility. A major implication of this is that the NTU assignment maximizes total match surplus and internalizes any externalities. Also, the equilibria of finite NTU matching markets with search frictions must approach surplus maximization as frictions go to zero (low search costs or high patience) in many environments⁸. The assumptions of our paper may frequently be satisfied in future theoretical work due to the need for tractability, so we can expect surplus maximization to be a common feature in tractable multivariate NTU models. However, the strong distributional assumption of symmetry needed to get this result suggests that this absence of externalities is a property of a special class of matching markets and cannot be expected to hold generally. Externalities in matching markets may drive rationales for intervention and provide opportunities for matching platform owners to profitably manipulate user matching behavior, so this is of practical interest.

This paper follows a rich literature on stable matching problems, starting with the seminal paper by Gale and Shapley (1964) mentioned above. Becker (1973) found that PAM occurs when there is a continuum of types and the utility of a match is increasing in types and nontransferable and that PAM also occurs when utility is transferable and the total utility of a match exhibits increasing differences in the two agents' types. Unlike Gale-Shapley, this requires no iterative process to find agent pairs in the stable matching, so it is suitable for use in theoretical models. However, it imposes the fairly onerous assumption of univariate vertical preferences—higher types are universally preferred to lower types, and agents only care about one trait. Legros and Newman (2007) extended PAM and NAM results to a class of partially nontransferable utility problems, where there are limitations on the ability of some or all agents to transfer utility to their match.

⁸Environments that satisfy the assumptions of this paper and those of e.g. Lauer mann and Nöldeke (2014).

Assuming horizontal preferences over a single trait where agents want to match to their own type, Clark (2003) gives an algorithm for finding stable matchings in a market with a finite set of agents. Clark (2007) then treats the univariate horizontal case with an infinite set of agents, finding a very simple matching result, which, like Becker’s result for vertical preferences, is well suited to a theoretical model. Clark (2006) also gives a condition guaranteeing a unique stable matching. Finally, Klumpp (2009) derives a very simple “inside-out” algorithm for horizontal matching with finitely many agents.

Multivariate matching has been studied empirically for some time. Choo and Siow (2003) develop an empirical model of TU marriage matching on age and education. Hitsch, Hortacsu, and Ariely (2010) study online dating, recovering preferences over many traits using a multivariate NTU model with horizontal and vertical preference dimensions. Chiaporri et al. (2012) study multivariate marriage matching empirically and recover a simple matching function by assuming that preferences can be aggregated to a single index of quality. Theoretical treatments of multidimensional matching include Chiaporri et al. (2010), which applies optimal transport theory to multidimensional TU matching problems and finds a very general (but not closed form) characterization of TU matching functions. Lindenlaub (2013) finds closed form solutions to multivariate matching problems in a very similar environment to this paper, and studies the effects of varying complementarities between traits. The analysis, however, is restricted to vertical preferences and TU. This paper extends the literature by finding closed form matching functions for TU and NTU in a framework where agents have preferences over multiple traits and where they can have a wider variety of preferences over each trait.

The remainder of this paper is organized as follows: Section 2 lays out the basic theoretical framework this paper uses and explores issues surrounding the modeling of agent preferences that shape the paper. Section 3 derives the main propositions of the paper, characterizing the matching functions for various symmetric n-dimensional horizontal matching problems. It also includes discussion of how these results might extend to asymmetric matching problems. Section 4 outlines the simulation model that is used to analyze the asymmetric case. Section

5 reviews the results. Section 6 summarizes the paper and suggests avenues for further research. Section 7, the in-text appendix, provides many Monte Carlo simulations to test the robustness of the theoretical results to deviations from the theoretical assumptions. It also includes several proofs not included in the main body of the article. An online appendix (located at http://sflanders.web.unc.edu/files/2013/09/ndimmatching_online_appendixc.pdf) provides additional Monte Carlo simulations, an extension of the results of this paper to the Roommates Problem, and background information on the various matching algorithms used and referenced in this paper.

2. THEORETICAL PRELIMINARIES

2.1. The Model. The environment we'll be considering is a matching market with two sides, or sets of agents, A and B ⁹. We'll denote specific agents in A as a , and specific agents in B as b . These sides could be interpreted as men and women in a heterosexual dating market. Agents of each side seek exclusive matches with agents of the other side. These agents can costlessly and perfectly observe every other agent in the market and costlessly propose and accept or reject any number of matches. Time is not modeled in this environment; everything happens simultaneously and with no time discounting. Agents have preferences over potential matches, and if $b \succ_a b'$ we'll say a strictly prefers b to b' , and if $b \succsim_a b'$ we'll say a prefers b to b' or is indifferent between them.

The goal of our analysis will be to find stable matchings in this environment. In this environment, a *matching* or *assignment* is a function $\mu : A \cup B \rightarrow A \cup B$ such that, for each agent $x \in A \cup B$, $\mu(x)$ is an agent on the opposite side or the empty set (no match), and μ is a bijection. This tells us what a match is, but our real goal is to predict how they will form. To do this, we need to specify agent preferences over matches. We will assume that, when agents a and b match, they produce a match surplus $u(a, b)$ which will be split between the two agents. In the nontransferable utility (NTU) environment, we will assume that agents cannot bargain over the apportionment of $u(a, b)$. For example, if a and b match, each agent will get

⁹We will sometimes abuse notation by denoting the type distributions associated with these sets by A and B as well.

$u(a, b)/2$. Agents want to maximize their own utility, so an agent a will prefer agents b that yield a higher $u(a, b)$. In the NTU environment, a *stable matching* is a matching in which there is no a and b such that $b \succ_a \mu(a)$ and $a \succ_b \mu(b)$. Such an (a, b) is called a *blocking pair*. In the transferable utility (TU) environment, match surplus can be apportioned between the two agents in any way. Because of this, the utility an agent gets from a match is not entirely determined by the agent she matches to—the transfers between agents must also be accounted for. Thus, a *stable matching* with TU is a matching μ such that there exists an allocation rule $v : A \cup B \rightarrow \mathbb{R}$ giving the utility for each matched agent such that is *feasible*: $v(a) + v(\mu(a)) \leq u(a, \mu(a)) \forall a \in A$ and $v(b) + v(\mu(b)) \leq u(b, \mu(b)) \forall b \in B$, and under which the match is *stable*: there is no a and b such that $u(a, b) > v(a) + v(b)$. Such an (a, b) is called a *blocking pair*.

We focus on matchings because we want to find out how agents pair up in this environment, and we restrict our consideration to stable matchings as we assume that, if agents are matched in an unstable way, it's likely that some matches will dissolve as poorly matched agents pursue better matches that also prefer them. In stable matchings, by contrast, the matching should remain unchanged over time so long as preferences and the distribution of agents remain the same.

Preferences can be very general in the framework outlined so far, but we'll restrict them to *horizontal* preferences, where agents prefer matches with types closer to their own. Specifically, we look at an environment where agents s of each side S have n traits, and their type is an n -vector, $\theta_s \in \mathbb{R}^n$. When unambiguous, we'll use s to denote an agent's type vector to conserve notation. The i th trait in this environment is θ_i . These could be income, height, BMI, risk aversion, etc. The horizontal preference assumption means that agents prefer matches whose n -dimensional type is closer to their own n -dimensional type in a given distance metric on \mathbb{R}^n . Typically we'll use the Euclidean distance.

We'll also specify utility functions corresponding to these preferences, where utility is decreasing in distance. In our first case, we'll assume nontransferable utility in this matching problem so that agents cannot offer some of their matching utility to a potential mate to

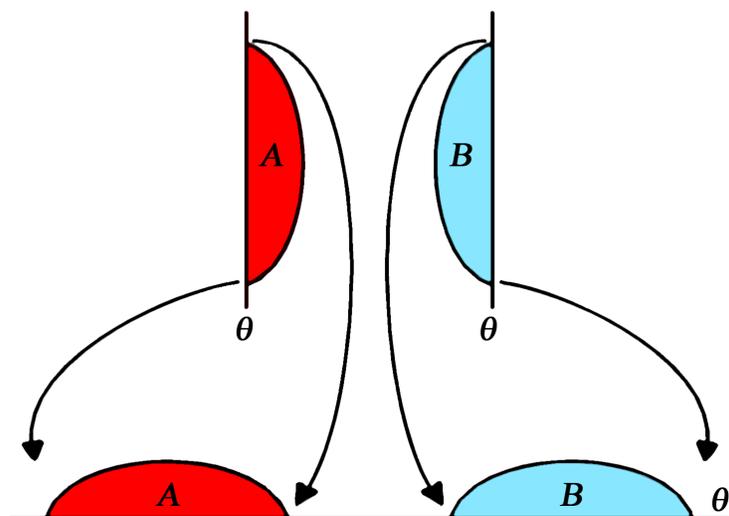


FIGURE 2.1. Mapping vertical preferences to a horizontal model: side A 's type remains the same in θ' , but B 's type is multiplied by -1 and translated by a constant, so that the highest A type is below the lowest B type. Thus the highest A and lowest B types in θ' are the highest θ type agents and are mutually most desirable in the horizontal framework, while the lowest A and highest B types are the low types in the original vertical framework and are mutually least desirable in the horizontal mapping.

induce them to match. We'll then assume transferable utility with an additional assumption that the utility function is convex in distance.

2.2. Modeling Various Preference Types in a Horizontal Framework. As mentioned previously, we are considering agents with horizontal preferences over n traits. However, there are many traits where preferences are manifestly not horizontal for most individuals. For example, people generally prefer more attractive partners, not a partner of their own level of attractiveness.¹⁰ Luckily, while the horizontal preference assumption requires preferences to correspond to a shared distance function over all n traits, it still allows considerable flexibility. Many types of preferences can be mapped into this framework. We'll show how vertical preferences can be mapped into this framework below. More general single peaked preferences, and certain types of categorical preferences can be also mapped into the horizontal framework, as shown in Online Appendix 1.1.1.

¹⁰A common result in matching models with vertical preferences is that agents match to mates of their own quality, but this is a characteristic of the equilibrium, not of agents' own preferences.

In the attractiveness example, we assume agents prefer more attractive individuals. If everyone can agree on the relative attractiveness of any two individuals, and everyone prefers more attractive to less attractive individuals, we call this a *vertical* preference. Vertical preferences can be represented in a horizontal framework, as shown in Figure 2.3. Given two distributions over a single trait θ with vertical preferences (the higher the type, the more desirable to all agents on the other side), we can generate a new trait θ' by mapping the two distributions to the real line with preferences based on least-distance. For example, we could have attractiveness for A and B , θ , range from 0 to 1. Then we can map to the new A attractiveness using the identity function $\theta'_A = \theta_A$ and the new B attractiveness using $\theta'_B = 2 - \theta_B$. Since higher θ' type A agents (lower θ' type B agents) had better vertical types, and are also closer to and thus more preferred by all B (A) agents, we preserve the preference orderings of all agents. Thus, if we expect agents to have vertical preferences over a trait we'd like to include in the model, we can preserve that preference structure in the horizontal model we've developed.¹¹

2.3. Aggregation to a Single Dimension. An obvious question we might ask is the following: can we reduce multivariate preferences to a single variable? Many theoretical models use univariate preferences, and existing research has already discovered closed form matching functions for many types of preferences over a single trait. Thus, if we could transform a multidimensional problem into a one-dimensional problem, that would be an attractive way to proceed. Specifically, we'll consider the following problem: we have an n -dimensional matching problem as described in Section 2.1, with either horizontal or vertical preferences over each trait. Our goal is to construct a univariate type and corresponding values along that one parameter for each agent such that the salient features of the n -dimensional matching model are preserved in the one-dimensional model. Ideally, we'd like to preserve preference orderings over potential matches for each agent and stable matching outcomes for each agent. In fact, a single aggregated type is used in a number of papers in

¹¹Note that we need the best A agent to be below the best B agent in the horizontal mapping. If not, there will be overlap in the support of A and B , and overlap agents will most prefer their own θ' type, rather than the "best" agent of the other side.

the literature, such as “pizzazz” in Burdett and Coles (1997). More recently, the ability to aggregate type vectors into a univariate index has been used as an identifying assumption in empirical work such as Chiappori et al. (2012).

As shown in Online Appendix 1.1.2, aggregation that preserves the set of stable matchings is possible with vertical preferences under an additional assumption. However, when horizontal traits are introduced, mapping from n dimensions to one dimension will generally lead to a larger set of stable matches. Most importantly, it cannot preserve the full preference orderings of each agent (Online Appendix 1.1.2). Because the full preference orderings are not preserved, if we change some parameter of the matching problem and cause the stable matching to change, we cannot expect those changes to be the same in the univariate model as in the underlying multivariate model. For example, if we want to examine a model where a market designer is optimizing over some parameters¹² that change the structure of the matchings, aggregation will render this optimization invalid with respect to the original n -dimensional problem. Generally, we cannot assume theoretical economic models involving univariate matching problems are valid stand-ins for those same models with multivariate matching problems unless we have a specific reason to believe such aggregation preserves the characteristics of the model we consider salient¹³. Thus, we’ll now consider the problem of explicitly solving multivariate matching problems.

3. THEORETICAL RESULTS FOR SYMMETRIC DISTRIBUTIONS

3.1. Nontransferable Utility Matching with Symmetric Distributions. If we cannot reduce an n -dimensional horizontal problem to a one-dimensional one, we must consider how to directly solve an n -dimensional problem. We will see that, given a form of symmetry between the distributions of each side and the condition that the distributions are separated, we can solve the matching problem. We can even characterize the type of one’s match as a linear function of one’s own traits. The model here is still the one outlined in Section 2.1,

¹²e.g. a price of entry into the matching platform.

¹³In a companion paper (Flanders (2014)), we find just such an environment where aggregation does not change the salient characteristics of the model.

with disutility of distance given by an increasing function f of the negation of the Euclidean distance metric: $u(a, b) \equiv f(-d(a, b))$, where d is shared by all agents on both sides.

Define the unit normal to a hyperplane h with origin at zero as $\eta(h)$. Denote the normal to a hyperplane h beginning in h and terminating at a point a as $\eta(h, a)$. Define $d_\eta(a, b)$ as the distance between a and b along vector $\eta(h, a)$ and $d_{h_i}(a, b)$ as the distance between a and b along the i th basis vector of h . Define $d(a, h) \equiv \|\eta(h, a)\| = d_\eta(a, h)$. Note that this is the minimal distance between a and the hyperplane h , and also the distance between a and a 's projection onto h . We'll need to make several assumptions to get a simple matching function:

- **Assumption 1 (SEP)** : $\exists h = \{x : ax = k\}$ for some a and k such h separates A and B . That is, $ay < k < az \forall y \in A, z \in B$.

Separation of the two distributions ensures that no one can get their own type as a match, which they would always accept. We could eliminate overlap by matching out identical agents and using the proposition to be proved on the remaining agents, but typically these remainder distributions will still not satisfy the separation criterion, as this condition is stronger than a requirement that the sets be disjoint. While this assumption appears strong, it will often be quite easy to satisfy: if there is at least one vertical trait, this condition is automatically satisfied, since vertical preferences require that the distributions of the two sides be separated along the vertical dimension, and by constructing a hyperplane with a normal along the vertical dimension, we can separate the entire n -dimensional distributions. Thus, if there are any traits that can be assumed to be quality based, such as income in a dating/marriage application, this condition imposes no further restriction. Alternatively, if more general preferences are decomposed into vertical and horizontal components, as done in Hitsch, Hortaçsu, and Ariely (2010) and discussed in Online Appendix 1.1.1, then the vertical component will be sufficient for Assumption 1. If the distance between the distributions is large enough, we can find a hyperplane that satisfies SEP as well as assumption 2. To state assumption 2, we must first define the reflection or Householder matrix of h as $R(h) \equiv I - 2\eta(h)\eta(h)^T$.

- **Assumption 2 (REF)** : the set A is the reflection about h of the distribution B .
That is, $R(h) \cdot A = B$.

We'll need this assumption to ensure that every agent has a reflected agent, which, combined with the shared distance metric, will ensure that we can match every agent to their reflection stably. Note that this assumption is a generalization of univariate symmetry assumptions (e.g. $F_M(x) = F_W(x)$ for type distributions F_S and sides M and W) common to many matching papers (e.g. Burdett and Coles (1997), Bloch and Ryder (2000), etc.) that also allows for more general forms of symmetry.

- **Assumption 3 (EUC)** : the distance metric on which preferences are based is the Euclidean distance.

Using the Euclidean distance, we'll be able to restate the distance between two points in terms of distance along the normal to a hyperplane and the distance along the basis vectors of that hyperplane, which will be crucial for proving that agents stably match to their reflections (this will not generally be true for other norms). The important characteristic of the Euclidean norm is that it is rotationally symmetric. That is, the indifference curve of any agent with distance preferences based on the Euclidean norm is a hypersphere, which has rotational symmetry. In contrast, for the 1-norm or sup-norm the indifference curve will be a hypercube, which is not invariant to rotation.

Essentially, we want to do what is seen in Figure 4.1. Given a matching problem with distributions A and B that are symmetric about some hyperplane which may have any arbitrary orientation in the typespace, we want to solve an equivalent problem where the typespace is redefined through a change of basis such that the reflecting hyperplane is now normal to one of the new basis vectors. Then an agent and their reflection will differ only along one dimension—along the vector normal to the hyperplane. This will be critical in the proof, and rotational symmetry of the Euclidean norm ensures that the rotation of the typespace due to the change of basis will not change the matching problem.

Now we can state the result. Recall the property of reflection matrices that $R(h) \cdot a = a - 2\eta(h, a)$:

Proposition 1. (*Continuous Symmetric NTU Matching*) Given a two sided NTU matching market with sides A and B , suppose there exists a hyperplane $h \subset \mathbb{R}^n$ satisfying SEP and REF. Suppose agents prefer closer matches in the Euclidean distance metric (EUC). Then all agents matching to their reflection is stable. That is, $\mu(a) = a - 2\eta(h, a) = R(h) \cdot a$.

Proof. For a contradiction, consider the matching outcome of Proposition 1 and suppose there is a blocking pair (a_1, b_2) such that $b_2 \succ_{a_1} \mu(a_1) = b_1$ and $a_1 \succ_{b_2} \mu(b_2) = a_2$. Then $d(a_1, b_2) < \min\{d(a_1, b_1), d(a_2, b_2)\}$. Since the agents in pairs (a_1, b_1) and (a_2, b_2) are each reflections of their respective matches, we know $d(a_1, b_1) = 2d(a_1, h) = d_\eta(a_1, h)$, $d(a_2, b_2) = 2d(b_2, h) = d_\eta(b_2, h)$. Since d is the Euclidean distance,

$$d(a_1, b_2) = \sqrt{\sum_{i=1}^n d_i(a_1, b_2)^2}$$

and equivalently, we have

$$\begin{aligned} d(a_1, b_2) &= \sqrt{\sum_{i=1}^{n-1} d_{h_i}(a_1, b_2)^2 + d_\eta(a_1, b_2)^2} \\ &\geq \sqrt{d_\eta(a_1, b_2)^2} \\ &= d_\eta(a_1, b_2) \\ &= d_\eta(a_1, h) + d_\eta(b_2, h) \\ &= (d(a_1, b_1) + d(a_2, b_2))/2 \\ &\geq \min\{d(a_1, b_1), d(a_2, b_2)\} \end{aligned}$$

Contradiction. □

Since an agent's a 's match is $R(h) \cdot a$ where R is a matrix, the matching function is linear. To interpret this, we'll introduce a new definition of assortment for multiple dimensions. Define a_i as the value of the i th trait of agent a , and a_{-i} as the vector of a 's traits excluding i .

Definition 2. (*Unconditional PAM(NAM)*) We'll say a matching μ satisfies unconditional PAM (NAM) in trait i if $a_i > a'_i$ implies $\mu_i(a) > \mu_i(a')$ ($\mu_i(a) < \mu_i(a')$) $\forall a, a' \in A$.

This extends the univariate definition to multiple dimensions by ensuring the one dimensional assortment holds throughout the typespace. We could imagine a weaker definition requiring only, say, $a_i > a'_i$ implies $\mu_i(a) > \mu_i(a')$ ($a_i > a'_i$ implies $\mu_i(a) < \mu_i(a')$) for a given a_{-i} vector ($a_{-i} = a'_{-i}$). Then we could have, for example, PAM in income for low education individuals and NAM in income for high education individuals. Definition 2, by contrast, requires a much stronger notion of assortment.

Given this definition and the fact that μ is linear in own type, the following holds:

Corollary 3. (*n-Dimensional Assortation*) for each i , μ_i either satisfies unconditional PAM, satisfies unconditional NAM, or, for each a_{-i} , all (\cdot, a_{-i}) type agents match to B agents of type (b_i, \cdot) for some b_i .

While the direction of assortment for each trait depends on the orientation of the two distributions, the linear matching pattern ensures that if matching on one trait is PAM (NAM) for one vector of other traits, it is PAM (NAM) for every vector of other traits. This yields strong testable implications about the structure of matching, but to fully characterize the qualitative structure of matching, it will be helpful to normalize the typespace. Defining a rotated typespace with the normal to h as the first dimension and $n - 1$ orthogonal spanning vectors of h as the remaining $n - 1$ dimensions and denoting the vector with all zero components except a value of one at component i as e_i , we can immediately derive a characterization of the matching function from Proposition 1:

Corollary 4. (*Normal n-Dimensional Assortation*) Where defined, $\frac{\partial \mu(a)}{\partial a_1} = -e_1$, $\frac{\partial \mu(a)}{\partial a_i} = e_i$ for $i > 1$

Proof. Note that $\mu(a) = a - 2\eta(h, a) = a - 2d(h, a)e_1$. Then $\frac{\partial \mu_i(a)}{\partial a_i} = e_i - 2\frac{\partial d(h, a)e_1}{\partial a_i}$ and $\frac{\partial d(h, a)e_1}{\partial a_i}$ is 0 if $i > 1$ and e_i if $i = 1$. □

Corollary 4 has a simple interpretation: along the normal to the hyperplane dividing the two distributions, the matching exhibits NAM. Along vectors orthogonal to the first, the matching exhibits PAM. Additionally, match type along one dimension depends *only* on own type along that same dimension. This is very intuitive given the fact that matches are reflections of one another along the hyperplane.

Note that Corollary 4 hold only for the synthetic traits of the rotated typespace, which are vectors in the original typespace. Further, the original typespace itself may be composed of synthetic traits generated from the original traits in order to map non-horizontal preferences into the horizontal preference framework, as shown in Section 2.2. Thus, if we want to interpret the assortment results with respect to the original traits, we'll need to map the rotated, synthetic traits back to the original set of traits. For a simple example, consider an n -dimensional matching problem with $n - 1$ horizontal dimensions and 1 vertical dimension, where no rotation is required. Then there will be PAM in all horizontal traits. Horizontal traits do not need to be mapped into the horizontal framework and we assumed they are unrotated, so no no mapping—or equivalently the identity mapping—is required. Thus the PAM of Corollary 4 applies directly to the horizontal traits. The vertical trait is still unrotated, but one side's values have been multiplied by -1. Thus, the NAM of Corollary 4 corresponds to PAM in the original vertical trait. Thus in this example we have PAM along all traits, and every agent matches to their own type.

3.2. Transferable Utility Matching with Symmetric Distributions. We now move on to an analogous matching problem for transferable utility. Just as Becker showed that TU and NTU-stable matchings coincide for univariate vertical preferences when match utility is supermodular in types, we find that the NTU-stable matching derived above is also TU-stable given the appropriate analogue for supermodularity in this framework. That analogue is convexity of the disutility of distance. This will ensure that the marginal cost of a closer pair being moved further apart is greater than the marginal cost of further pairs being separated further. Convexity is important in the TU framework since agents are free to

bargain with each other over that division of match surplus. Due to that bargaining, TU-stability requires that the sum of match surpluses be maximized, and the aggregate surplus maximizing allocation depends on convexity. As with our previous result, we will see NAM along the vertical dimension. Convexity ensures that a distant pairing and a close pairing has a higher total match surplus than two mediocre pairings, so in order to maximize match surplus the closest pairs will be preferentially matched together. Along all other dimensions agents will match to their own type (after the typespace has been rotated), as this is their ideal match.¹⁴

Generally, explicitly solving for TU-stable matchings is more difficult, since one must find not just the matching but also show there are surplus allocations that support that matching as stable. Finding those surplus allocations can be very difficult in general, but the REF assumption ensures that an even split of the match surplus for every pair will admit a stable matching. Generally, the allocations can be thought of as a shadow price for the agent's presence in the matching market (Browning et al. (2014)). As such, stable allocations vary widely depending on the outside options of each agent in the match—colloquially, whether they are in shortage or surplus. However, we've assumed that the two distributions are symmetric, and in the stable matching agents will turn out to match to their mirror type, as before. Thus, every agent's decision problem is mirrored by the decision problem of their mirror match, and neither has any sort of advantage or disadvantage relative to the other in bargaining over the split, so an even split is supportable.

Proposition 5. (*Continuous Symmetric TU Matching*) *Given a two sided TU matching market with sides A and B , suppose there exists a hyperplane $h \subset \mathbb{R}^n$ satisfying SEP and*

¹⁴Becker also found a result for TU stability with submodular match utility, and we conjecture that this result too can be generalized to the n-dimensional framework, where the disutility of distance is instead concave. In this framework, instead of the two distributions being reflections of one another, they must instead be translations of one another—the same distributions up to an offset. This is actually a less restrictive assumption than the reflective symmetry assumption, as it does not require any sort of rotation, or that the distributions be separated. Given this framework, we conjecture (and Monte Carlo simulations support) that agents will match to their translated twin in the opposing distribution. However, proving this result will be more difficult than in the convex disutility case, and, as with Becker's result for submodular utility, the result is of less interest since it implies NAM, which is typically not observed empirically. Thus, the proof is not pursued here.

REF. Suppose agents prefer closer matches in the Euclidean distance metric (EUC) and the match utility is weakly convex and decreasing in distance. Then all agents matching to their reflection is a stable matching assignment. That is, $\mu(a) = a - 2\eta(h, a) = R(h) \cdot a$. Further, every pair splitting the match surplus equally is an allocation consistent with stability. If match utility is strictly convex in distance and there are finitely many agents, the stable assignment is unique.

Proof. See Appendix 7.1.2 □

We can extend this result beyond two-sided matching problems as well. Online Appendix 1.2 gives an analogous result for the one-sided matching or “Stable Roommates” problem—a result that is in some ways more robust, as it does not require rotation of the typespace.

3.3. Matching with Asymmetric Distributions. Propositions 1 and 5 gives us an easily derived and interpreted matching function. However, we are very unlikely to encounter perfectly symmetric sides empirically; we cannot expect the n -dimensional distribution of men to be the exact reflection of the n -dimensional distribution of women about a separating hyperplane. However, we can easily find an approximate reflection. For example, we can choose a hyperplane that reflects the center of mass of distribution A to the center of mass of distribution B . The natural question to ask, then, is whether the sort of approximate symmetry we might see in the data corresponds to approximately the same matching structure. Unfortunately, deriving more general closed form matching functions for n -dimensional horizontal matching markets is extremely difficult.

However, we can make some conjectures. The factors that ensure the assortment in the symmetric case are still at work in an asymmetric market. In an asymmetric case like Figure 3.1, agents still want closer matches, which means A agents closer to B (more desirable to all B agents) will match to agents in B that are themselves close to A (also more desirable). Similarly, agents on the top right of B are likely to match to agents on the right side of A , who they prefer and to whom they’re among the more attractive options. However, because there is not a symmetric match for all agents, one side or another will be in shortage at

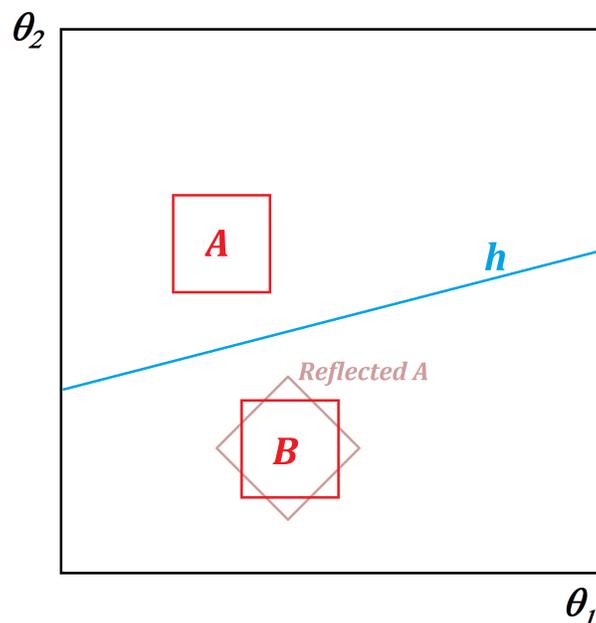


FIGURE 3.1. In any empirical application, the symmetry assumption will never be fully satisfied. In this example, distributions A and B are uniform square distributions, identical up to translation. Because they are offset from each other along the horizontal dimension, no hyperplane can perfectly reflect one onto the other, so the symmetry assumption fails.

various times in the inside-out algorithm, so the matching outcome will be distorted from the ideal symmetric case. Thus we would expect some attenuation in the effect of own traits on corresponding match traits and possibly some modest effect of own traits on non-corresponding match traits. Notice that, while the reflection may not be a perfect match, as long as there is sufficient separation between the two distributions we will be able to find a separating hyperplane that maps the center of mass of A onto the center of mass of B , giving an approximate reflection. At least one vertical dimension will guarantee that the two distributions are separated, as seen in Section 2.2. If the separation of the two sides is large enough, we should have enough space between the distributions to fit a hyperplane that both separates them and reflects the center of mass of one onto the center of mass of the other. Thus, while a lack of symmetry may change the matching outcome, getting an approximate reflection should not generally be a problem in an empirical setting.

4. SIMULATION FRAMEWORK

4.1. Simulation Setup. We'll now develop a framework to test the validity of symmetric distribution results in situations with asymmetric distributions. We'll consider two cases: first, we assume a best case scenario where the underlying distributions for A and B are symmetric but the realized observations are drawn randomly and thus do not exhibit perfect symmetry. Note that this will completely eliminate the matching structure we relied on for Proposition 1, since agents no longer have mirror matches. However, the overall distribution should be approximately the same, so we can hope that the results will be almost identical. Second, we consider a less optimistic scenario where the underlying distributions are not perfectly symmetric, but exhibit moderate asymmetry when reflected onto one another, as in Figure 3.1. In this case we can expect same-trait effects significantly below one and other effects may be nonzero.

To simplify the analysis and facilitate visualization of the model, we'll primarily focus on a two-dimensional typespace, and later look at how increasing the number of dimensions changes the outcome. In both the symmetric and asymmetric cases, the observations on both sides are drawn from square bivariate uniform distributions. In the first case, they are stacked vertically as in the right-hand portion of Figure 4.1. In this case A and B are symmetric about h , h is horizontal, and the assortment should be along θ_1 and θ_2 . For the second case, the distributions are offset along θ_1 , yielding a market like that seen in Figure 3.1. In this case, h is not horizontal. To match the predicted effects to the axes of the model, it will be necessary to rotate the typespace such that h becomes horizontal as seen in Figure 4.1.

4.2. Simulation Model Specifications. In the simulation model, the two dimensional matching market is as described in section 4.1. For higher dimensions, only the asymmetric case is simulated: the distribution is uniform over an n -cube, offset greatly along one dimension, and slightly offset along all others. With the simulated agents in hand and their preferences specified, we can simply run a version of Klumpp's inside-out algorithm in the

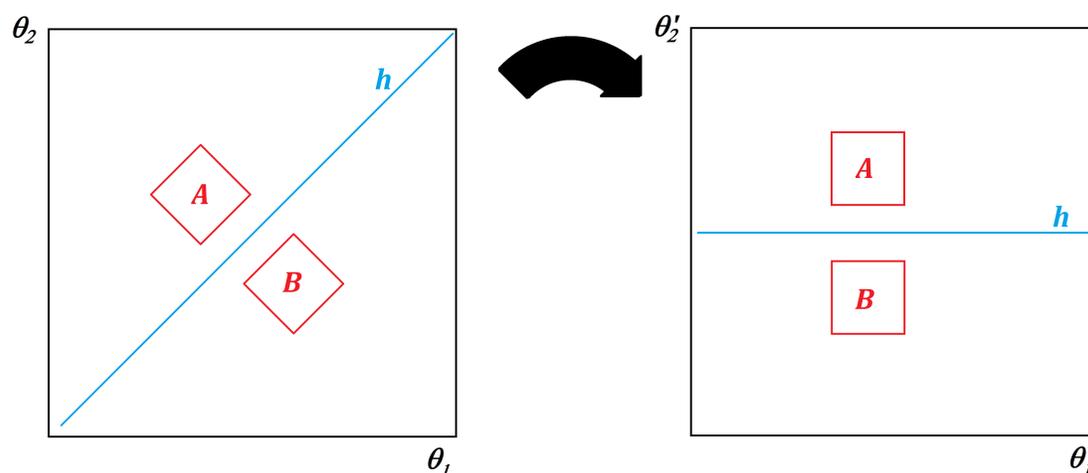


FIGURE 4.1. An example where A and B are symmetric about the separating hyperplane h , but the hyperplane is not normal to any of the basis vectors of the type-space. To observe the predicted assortment, we must rotate the typespace so that h is normal to one of the newly created synthetic traits.

NTU case to find the stable matching outcome.¹⁵ In the TU case, we can solve for the stable assignment and transfers pair by formulating the matching problem as a linear program (Shapley and Shubik (1972)).¹⁶ Given a separating hyperplane along the horizontal axis, we can then run regressions with one b trait as the dependent variable and both $a = \mu(b)$ traits as the independent variables, where the resulting coefficients estimate the effects of a change in each on the b trait. In the idealized symmetric case we would expect the coefficients to be one for trait one on trait one, -1 for trait two on trait two, and zero otherwise. We run each specification many times and find the mean of recovered coefficient values, as well as a 90% confidence interval. More detailed model specifications can be found in Appendix 7.2.1.

While a few simulations are shown in Section 5 below, many more simulations are included in Appendix 7.2 and Online Appendix 1.3. The environments simulated include alternate distributions for A and B , non-uniform distributions, an example with a categorical variable, deviations from the 2-norm assumption, simulations with correlated traits, different levels of convexity with TU, and NTU and TU simulations for various market sizes and numbers of preference traits.

¹⁵See Online Appendix 1.4.3 for a summary of the algorithm.

¹⁶See Online Appendix 1.4.2 for a summary of the algorithm.

Predicting $\mu_1(a)$ (vertical characteristic) by a			
Symmetric, n=100			
	a_1	a_2	R^2
95th %ile	-0.89	0.05	0.99
mean	-0.98	0.00	0.97
5th %ile	-1.09	-0.04	0.95
Asymmetric, n=100			
	a_1	a_2	R^2
95th %ile	-0.87	0.05	0.98
mean	-0.98	0.00	0.97
5th %ile	-1.08	-0.04	0.94

FIGURE 5.1. Monte Carlo Simulation Results for the Vertical trait: NTU, 200 iterations, baseline specification.

5. RESULTS

First, we'll look at the two dimensional case with NTU and a very coarse market of 100 agents on each side (Figure 5.1). We start with the 100 agent case so that we can compare the NTU results to the TU results, which cannot easily be simulated for larger markets. We'll initially look at the match's first trait, the "vertical" or separating trait. We see that in this case the linear same-trait effects explain virtually all of the variation in one's match's vertical trait, and the coefficient is very close to -1 for both the symmetric and asymmetric distributions. The opposite-trait effects are quite close to zero, as predicted. The R^2 is very close to 1, showing that almost all the variation in your match's vertical trait is explained by your own vertical trait, with no noticeable drop off for the asymmetric case. The range of coefficient estimates is fairly tight around both traits, showing that we can expect the estimated coefficients to be consistently close to their predicted values in this environment.

Now we'll look at the match's second trait, the "horizontal" trait (Figure 5.2). In the symmetric case, we see that the linear same-trait effects explain much of the variation in one's match's horizontal trait, and the coefficient is fairly close to 1, as predicted. However, there is much more attenuation than with the vertical coefficient for both the symmetric and the asymmetric case. The opposite-trait effects are quite close to zero, as before. The R^2 is significantly lower than for the vertical coefficient, and drops off more significantly in the asymmetric case. The range of coefficient estimates is still fairly tight around both traits.

Predicting $\mu_2(a)$ (horizontal characteristic) by a			
Symmetric, n=100			
	a_1	a_2	R^2
95th %ile	0.15	0.96	0.77
mean	-0.02	0.83	0.67
5th %ile	-0.21	0.72	0.55
Asymmetric, n=100			
	a_1	a_2	R^2
95th %ile	0.15	0.85	0.67
mean	-0.01	0.75	0.56
5th %ile	-0.20	0.65	0.45

FIGURE 5.2. Monte Carlo Simulation Results for the Horizontal trait: NTU, 200 iterations, baseline specification.

Predicting $\mu_1(a)$ (vertical characteristic) by a			
Symmetric, n=100			
	a_1	a_2	R^2
95th %ile	-0.87	0.16	0.94
mean	-0.96	-0.02	0.92
5th %ile	-1.08	-0.16	0.88
Asymmetric, n=100			
	a_1	a_2	R^2
95th %ile	-0.84	0.13	0.92
mean	-0.96	-0.02	0.88
5th %ile	-1.10	-0.21	0.84

FIGURE 5.3. Monte Carlo Simulation Results for the Vertical trait: TU, 60 iterations, baseline specification.

As we'll see later, the results much better approximate the ideal symmetric case as the size of the market increases.

Now, we'll look at the analogous two dimensional case with TU (Figure 5.3). We'll initially look at the match's first trait, the "vertical" or separating trait. We see that the linear same-trait effects still explain virtually all of the variation in one's match's vertical trait, and the coefficient is very close to -1 for both the symmetric and asymmetric distributions—though not quite as close as in the NTU case. The opposite-trait effects are also quite close to zero, as predicted. The R^2 is fairly close to 1, showing that most of the variation in your match's vertical trait is explained by your own vertical trait, with a small drop off for the asymmetric case. The range of coefficient estimates is fairly tight around both traits, showing that we

Predicting $\mu_2(\mathbf{a})$ (horizontal characteristic) by \mathbf{a}			
Symmetric, n=100			
	a_1	a_2	R^2
95th %ile	0.09	1.07	0.97
mean	0.01	1.00	0.96
5th %ile	-0.09	0.88	0.93
Asymmetric, n=100			
	a_1	a_2	R^2
95th %ile	0.03	1.07	0.96
mean	-0.04	0.95	0.94
5th %ile	-0.12	0.87	0.90

FIGURE 5.4. Monte Carlo Simulation Results for the Horizontal trait: TU, 60 iterations, baseline specification.

can expect the estimated coefficients to be consistently close to their predicted values in this environment.

Finally, we'll look at the match's horizontal trait in the TU case (Figure 5.4). In the symmetric case, we see that in this case the linear same-trait effects explain virtually all of the variation in one's match's horizontal trait, and the coefficient is fairly close to 1, as predicted, with modest attenuation in the asymmetric case. The opposite-trait effects are quite close to zero, as before. The R^2 is quite close to 1 and the range of coefficient estimates is still fairly tight around both traits.

These results are quite auspicious for applications of the theoretical result to empirical data. Even in very small, coarse matching markets of two hundred agents, the idealized result well approximates the actual outcome. The only exception to this is the weaker horizontal assortment results in the NTU case. We'll see that the NTU case's horizontal results improve markedly as the number of agents on each side grows. The primary question we are left with is why the horizontal trait's effect is significantly weaker and why it explains less of the horizontal variation in the NTU case. This question is treated in Appendix 7.2.2.

So far, we've considered very small, coarse matching markets. How do the results we've seen change when there are more agents on each side? Do the results more closely mirror the theoretical predictions? In particular, do the less than ideal results we saw in the horizontal NTU case improve with more agents on each side? Also, so far we've looked at

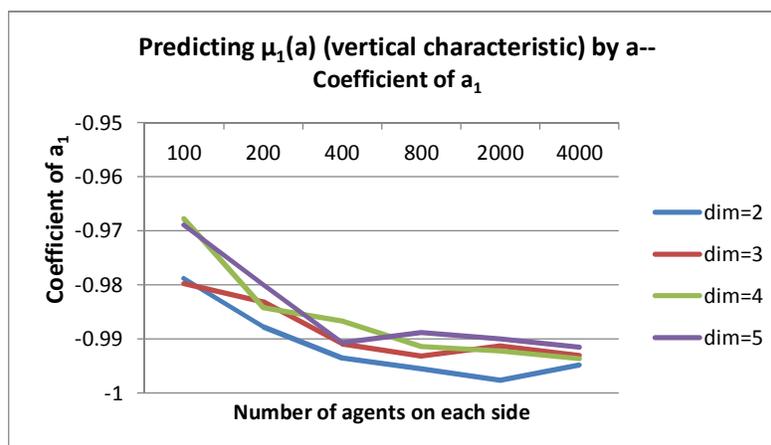


FIGURE 5.5. NTU with baseline specifications: predicting match’s vertical trait from own vertical trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

models with just two traits over which agents have preferences. What happens if there are more traits? We now address these questions.¹⁷

In Figure 5.5 we see that, for an agent’s match’s vertical trait, the average coefficient on the corresponding trait of the agent appears to asymptotically approach 1 as the number of agents on each side increases. We see some slight attenuation from this result as we increase the number of preference dimensions, but even with five dimensions the result is quite strong. To the extent that the coefficient values do not monotonically decrease as the number of agents increases and as the number of traits decreases, we can attribute this to the finite sample size for the Monte Carlo simulations, which introduces some noise into the mean coefficient estimates. In Figure 5.6 we see that, for an agent’s match’s vertical trait, the

¹⁷Before we proceed, a note on some differences in the TU and NTU simulations: the inside-out algorithm is faster than Gale-Shapley, which is already an extremely fast algorithm. However, the method used to solve for TU stable matches is extremely slow. Thus TU simulations in this paper are for markets with between 30 and 100 agents per side, while NTU simulations go up to 4,000 agents per side. This is because, for TU simulations, the linear program to be solved has 10,200 constraints with just 100 agents to a side. If we were to attempt to solve the model with 4,000 agents to a side, there would be 16,008,000 constraints. Thus, the computation time increases very quickly with larger markets. In the 100 agents per side case, the inside-out algorithm takes about 0.001 seconds, while the TU algorithm takes 2-4 minutes per simulation. Even increasing the market size to 150 per side requires at least an hour of computation, if not more. Thus, it is not possible to run Monte Carlo simulations for large markets in the TU case, and the number of iterations must be lower than in the NTU case.

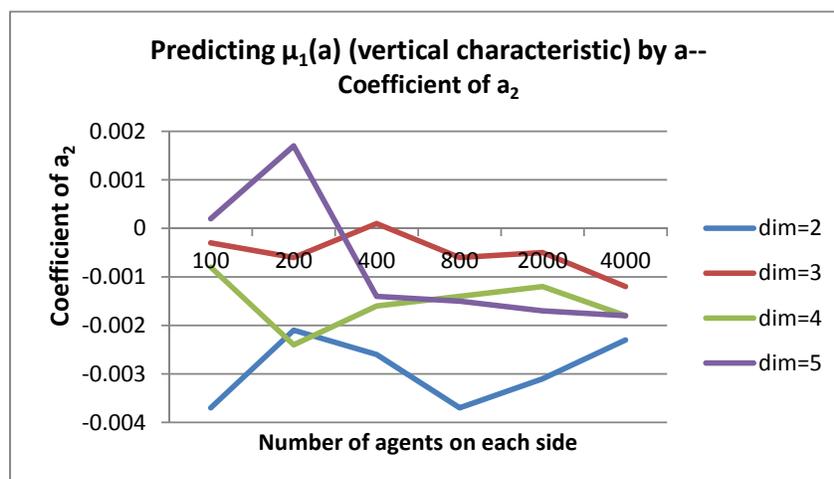


FIGURE 5.6. NTU with baseline specifications: predicting match's vertical trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

average coefficient on the horizontal trait of the agent appears to asymptotically approach a value just slightly below 0. To the extent that the coefficient values are not monotonic in sample size, we can probably attribute this to the finite sample size for the Monte Carlo simulations, noting the extremely small region of the y-axis that's being graphed. The fact that the coefficient seems slightly biased from the predicted coefficient of 0 should not be surprising—our baseline specification includes asymmetry in both the draws and the underlying distributions, and while asymmetry in the draws should asymptotically approach zero as the size of the market increases, the asymmetry of the underlying distributions will not. If anything, it is quite impressive that there is so little bias, given the significant deviation from symmetry we've specified.

In Figure 5.7 we see that, for an agent's match's vertical trait, the average R^2 appears to asymptotically approach 1 as the number of agents on each side increases. We also see that the R^2 is attenuated as we increase the number of preference traits. Note that, to recover the R^2 , we regress only the same trait coefficient so that the R^2 gives us the variation in one's match trait solely attributable to the predicted same own trait effect.

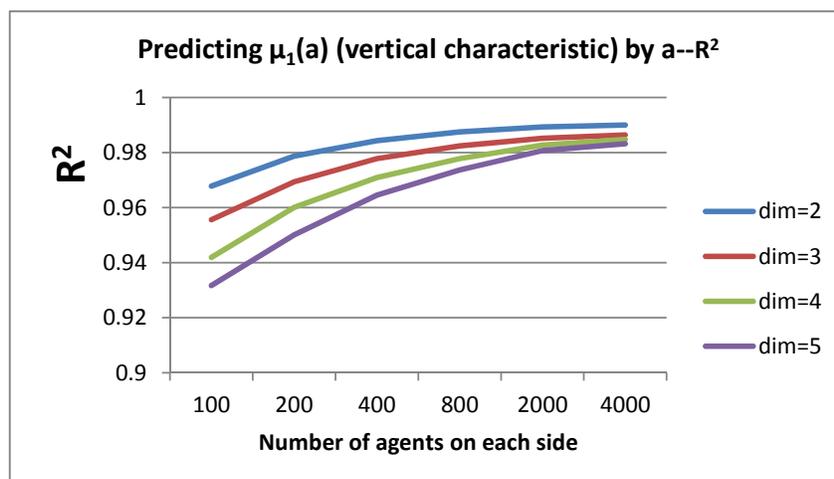


FIGURE 5.7. NTU with baseline specifications: average predicted R^2 for the regression of own vertical trait on match's vertical trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

For the NTU case, we've looked at simulations where both the individual draws and underlying distributions have significant asymmetry, and we generally see that the predictions from the theoretical results for symmetric distribution are a good approximation for the actual stable assignments. As the size of the market increases, the stable assignments appear to asymptotically approach the predicted assortment in the vertical trait. As seen in Appendix 7.2.3, we also come fairly close to the predicted assortment along horizontal traits, though the results are not quite as strong. Additionally, this asymptotic behavior appears quite quickly as the size of the market increases—just a few thousand agents to a side gives us extremely strong fidelity to the predicted result. In many applications such as online dating, the market will be more than large enough to see asymptotic behavior. We do see that we need larger and larger markets to get the same level of fidelity to the predicted results as we increase the number of traits. This is not terribly surprising, as increasing the number of traits while holding the number of agents on each side constant effectively makes the distributions on each side sparser, since they vary over a larger type-space. A concern here

Predicting $\mu_1(a)$ (vertical characteristic) by a		
NTU, Asymmetric, n=1000		
	a_1	R^2
Equal weights	-0.99	0.9758
Higher weight for second characteristic	-0.99	0.4779
Predicting $\mu_2(a)$ (horizontal characteristic) by a		
	a_2	R^2
Equal weights	0.83	0.6879
Higher weight for second characteristic	1	0.9211
Predicting $\mu_3(a)$ (horizontal characteristic) by a		
	a_3	R^2
Equal weights	0.83	0.6932
Higher weight for second characteristic	0.71	0.505
Predicting $\mu_4(a)$ (horizontal characteristic) by a		
	a_4	R^2
Equal weights	0.83	0.6921
Higher weight for second characteristic	0.71	0.507
Predicting $\mu_5(a)$ (horizontal characteristic) by a		
	a_5	R^2
Equal weights	0.78	0.6921
Higher weight for second characteristic	0.65	0.507

FIGURE 5.8. Here we compare the fidelity of a five dimensional NTU stable matching to the predicted results for two cases. In the equal weights case, each trait is given equal weight in the distance metric. In the unequal weights case, trait two is given ten times as much weight as the others. The reported values are the means of 200 stable matchings with 1000 agents on each side. The model used is the asymmetric baseline for five dimensions.

is that we could presumably enumerate dozens or hundreds of traits over which agents have preferences. The simulations above suggest we might need millions or billions of agents on each side to get a good approximation to the symmetric result if there are too many traits. Luckily, not all preferences are created equal, and some traits will be of great importance, while others are of little. Then, as we see in Figure 5.8, the traits that are very important to agents should have good fidelity to the predicted results at reasonable market sizes even if there are many more less important traits. The less important traits, conversely, will have very poor fidelity to the predicted result. However, traits which agents do not care much about are probably not of great interest to begin with.

As mentioned before, we can only simulate very small markets for the TU case. Thus, we cannot observe the asymptotic TU trends the way we did in with NTU. The TU simulations

for various market sizes and numbers of traits appear quite similar to those for NTU. The main difference is that, as seen before, there is a relatively better fit for the horizontal traits and a relatively worse fit for the vertical trait compared to NTU. However, the inability to simulate to markets of many hundreds of agents and the very small number of iterations that were possible for the Monte Carlo simulations makes the interpretation much more difficult. Generally, the TU results seem consistent with the NTU results in the range we can examine them in, and we assume that they would continue to mirror the NTU results in larger markets. The TU simulations are presented in Appendix 7.2.4.

6. CONCLUSIONS

In this paper, we found a simple, closed form matching function for a special case of frictionless two-sided matching where agents have preferences over multi-dimensional types. To get this result, we needed to make strong assumptions on the distributions of agents and the structure of preferences, most notably that the distribution of agents on each side was the reflection of the distribution of agents on the other. However, the simulations in Section 5 and in the appendix strongly support the symmetric mirror-matching result's applicability to modestly asymmetric markets. While, as expected, there is some attenuation of the anticipated same-trait effects on matches, the coefficients are relatively close to their predicted values even in small matching markets of a few hundred, and improve as the size of the market increases. Thus, these results may plausibly be applied to empirical matching data. These results also have relevance for theory work. We can embed the closed form matching functions into more complex economic models, such as models of online dating markets. This allows for the theoretical study of matching phenomena involving multiple traits, such as how agents tradeoff between various match traits. It also allows one to compare the highly aggregated, univariate theoretical matching models that are typical in the literature to multivariate models, in order to see whether the qualitative characteristics of a multivariate model are preserved in a more stylized univariate model.

One major implication of our result is that the NTU assignment maximizes total match surplus when the match utility function is convex. In an analogue to the Coase Theorem, the frictionless TU assignment maximizes total surplus (Shapley and Shubik (1972)) and internalizes externalities via transfers, and the two assignments coincide, so the NTU assignment has the same properties. Also, Adachi (2003), Eeckhout (1999), and Lauermaun and Nöldeke (2014) show that, in many search environments with NTU, search equilibria approach the frictionless stable assignment as frictions go to zero, provided that assignment is unique. Thus, there are a wide range of environments¹⁸ where finite¹⁹ NTU search markets must have equilibria that approach surplus maximization as frictions go to zero. The setting of this paper encompasses many environments with univariate vertical or horizontal preferences and symmetric distributions, and future theoretical work in both search and frictionless multi-dimensional matching will likely often satisfy the assumptions imposed in this paper due to the tractability issues with asymmetric distributions. Thus, tractable models with NTU, especially with multiple preference dimensions, will likely have these qualities. This implies, for instance, that externality issues with NTU search markets of this type can be resolved by simply improving the search technology. However, the strong assumptions needed in this paper—and the divergence of TU and NTU assignments in simulations that relax these assumptions—illustrate how special these environments are and suggest that these efficiency results cannot be expected to hold generally in the broader universe of possible matching markets. Failure of NTU matching to maximize social surplus due to externalities may provide a justification for market intervention or for platform owners like online dating websites to influence consumer matching behavior through contracts, platform structure, etc., so this issue has practical importance. Indeed, a recent survey of the search and matching literature (Chade, Eeckhout, and Smith (2015)) identifies the role of externalities in matching markets as one of the most important open questions in the field.

¹⁸e.g. those satisfying the assumptions of one of the above papers and this paper.

¹⁹The NTU uniqueness result in this paper is only proved for finite markets.

7. APPENDIX

7.1. Proofs.

7.1.1. *Uniqueness of Symmetric stable matching with finitely many agents.* While we've proven in Section 3 that the symmetric matching outcome is stable for NTU, we have not proven that it is unique. While the following proof technique does not work in the infinite case. In the finite case, we can construct the only possible type of stable matching and show that, under certain conditions, the set of stable matchings is a singleton.

Proposition 6. (*NTU Finite Symmetric Matching*) *Suppose \exists a hyperplane $h \subset \mathbb{R}^n$ such that the finite set of agents A is the reflection of the finite set of agents B about h and h separates A and B . Suppose agents prefer closer matches in the Euclidean distance metric. Then all agents match to their reflection. That is, $\mu(a) = a - 2\eta(h, a) = R(h).a$.*

Proof. Consider the first step of Klumpp's inside-out algorithm and a pair (a, b) such that $d(a, b)$ is distance minimal among all $a \in A$ and $b \in B$. We will show that a and b are reflections of each other. Without loss of generality, consider a 's matching problem. Suppose $d(a, b) > d(a, b')$, where $b' = R(h).a$. Then (a, b) is not distance minimal, a contradiction. Supposed $d(a, b) \leq d(a, b')$ where $b' = R(h).a$ and $b \neq b'$. Since Euclidean distance is rotation invariant, we can find $n-1$ orthogonal vectors spanning h and decompose $d(a, b)$ into distance

along the normal and distance along h , $\sqrt{\sum_{i=1}^{n-1} d_{h_i}(a, b)^2 + d_\eta(a, b)^2}$. Then

$$\sqrt{\sum_{i=1}^{n-1} d_{h_i}(a, b)^2 + d_\eta(a, b)^2} \leq \sqrt{\sum_{i=1}^{n-1} d_{h_i}(a, b')^2 + d_\eta(a, b')^2}$$

Since b' is the reflection of a about h , $d_{h_i}(a, b') = 0$. Therefore we have

$$\sqrt{\sum_{i=1}^{n-1} d_{h_i}(a, b)^2 + d_\eta(a, b)^2} \leq \sqrt{d_\eta(a, b')^2}$$

$$\sum_{i=1}^{n-1} d_{h_i}(a, b)^2 + d_\eta(a, b)^2 \leq d_\eta(a, b')^2$$

$$d_\eta(a, b)^2 < d_\eta(a, b')^2$$

$$d_\eta(a, b) < d_\eta(a, b')$$

Note that distance $d_\eta(a, b) = d(a, h) + d(b, h)$, so we have

$$d(a, h) + d(b, h) < 2d(a, h)$$

$$d(b, h) < d(a, h)$$

But we know that $a' = R(h).b$ is an agent in A since A is the reflection of B about h, and $d(a', b) = 2d(b, h) < d(a, h) + d(b, h) = d_\eta(a, b) < \sqrt{\sum_{i=1}^{n-1} d_{h_i}(a, b)^2 + d_\eta(a, b)^2} = d(a, b)$, so (a,b) is not distance minimal, a contradiction. Continuing inductively, if all previous steps in the inside-out algorithm have resulted in mirror pairs matching out, every agent remaining unmatched has a mirror pair still unmatched, and the result just proved applies. Thus all agents getting a mirror match is a stable matching. Note that having at least one vertical trait will ensure the separation condition, as along the vertical axis, all agents in A will be above (below) all agents in B. Since this is just a special case of the inside out algorithm, the properties of that algorithm's matching outcome are preserved, most importantly the uniqueness of the stable match given strict preferences (Klumpp (2009)). \square

7.1.2. Continuous Symmetric TU Matching.

Proof. First, we'll show stability. For a contradiction, consider the matching outcome of Proposition 5 and suppose there is a blocking pair (a_1, b_2) such that $u(d(a_1, b_2)) > u(d(a_1, b_1))/2 + u(d(a_2, b_2))/2$. But by convexity, we know $u(d(a_1, b_1))/2 + u(d(a_2, b_2))/2 \geq$

$u((d(a_1, b_1) + d(a_2, b_2))/2)$. Since the agents in pairs (a_1, b_1) and (a_2, b_2) are each reflections of their respective matches, we know $d(a_1, b_1) = 2d(a_1, h) = d_\eta(a_1, h)$, $d(a_2, b_2) = 2d(b_2, h) = d_\eta(b_2, h)$. Since d is the Euclidean distance,

$$d(a_1, b_2) = \sqrt{\sum_{i=1}^n d_i(a_1, b_2)^2}$$

and equivalently, we have

$$\begin{aligned} d(a_1, b_2) &= \sqrt{\sum_{i=1}^{n-1} d_{h_i}(a_1, b_2)^2 + d_\eta(a_1, b_2)^2} \\ &\geq \sqrt{d_\eta(a_1, b_2)^2} \\ &= d_\eta(a_1, b_2) \\ &= d_\eta(a_1, h) + d_\eta(b_2, h) \\ &= (d(a_1, b_1) + d(a_2, b_2))/2 \end{aligned}$$

Thus

$$u(d(a_1, b_1))/2 + u(d(a_2, b_2))/2 \geq u((d(a_1, b_1) + d(a_2, b_2))/2) \geq u(d(a_1, b_2))$$

Contradiction.

Now, we show uniqueness. Recall that a stable allocation must maximize the aggregate match surplus. Suppose there are k agents on each side and that mirror agents have identical indices. Suppose that each agent's type is unique.²⁰ Then for any potential stable allocation μ , we must have $\sum_{i=1}^k u(d(a_i, \mu(a_i))) = \sum_{i=1}^k u(d(a_i, b_i))$. As before, we can decompose the

²⁰If there is more than one agent of a given type, we can easily amend the proof to account for this. The stable allocation will remain unique up to agent type, though not up to individuals, since two identical agents can have their matches switched without changing total surplus.

distances along the basis vectors, noting again that $d(a_i, b_i) = d_\eta(a_i, b_i)$. Then we have

$$\sum_{i=1}^k u(d_\eta(a_i, b_i)) = \sum_{i=1}^k u\left(\sqrt{\sum_{i=1}^{n-1} d_{h_i}(a_i, \mu(a_i))^2 + d_\eta(a_i, \mu(a_i))^2}\right)$$

$$\sum_{i=1}^k u(d_\eta(a_i, b_i)) \leq \sum_{i=1}^k u(d_\eta(a_i, \mu(a_i)))$$

Note that, having removed all horizontal components, we have a condition on a single vertical component. This condition is analogous to the optimality condition in the standard Becker TU matching problem. For any distinct pairs i and j , we have that $d_\eta(a_i, b_i) + d_\eta(a_j, b_j) = d_\eta(a_i, b_j) + d_\eta(a_j, b_i)$ and $d_\eta(a_i, b_j) = d_\eta(a_j, b_i)$. Then convexity ensures that $u(d_\eta(a_i, b_j)) = u(d_\eta(a_i, b_j)/2 + d_\eta(a_j, b_i)/2) = u(d_\eta(a_i, b_i)/2 + d_\eta(a_j, b_j)/2) < u(d_\eta(a_i, b_i))/2 + u(d_\eta(a_j, b_j))/2$. For any $\mu(a_i)$, define $b_{j(i)} \equiv \mu(a_i)$. Then $\sum_{i=1}^k u(d_\eta(a_i, \mu(a_i))) = \sum_{i=1}^k u(d_\eta(a_i, b_{j(i)})) \leq \sum_{i=1}^k u(d_\eta(a_i, b_i))/2 + u(d_\eta(a_{j(i)}, b_{j(i)}))/2 = \sum_{i=1}^k u(d_\eta(a_i, b_i))$ since μ is a bijection, and the inequality is strict if $j(i) \neq i$ for some i . In fact, if μ is some matching other than the stable matching described above, it must be that $j(i) \neq i$ for some i . Then this alternate matching cannot be stable and the stable assignment is in fact unique. ²¹ \square

7.2. Monte Carlo Simulations.

7.2.1. Detailed Simulation Model Specifications. In the NTU case we simulate 100, 200, 400, 800, 2,000, or 4,000 agents for each side for both the symmetric and asymmetric specifications. The agents are drawn from an independent bivariate uniform distribution with support from 0 to 5. In the symmetric case, the two distributions are offset by 10 along the second trait (thus, along the second trait, A agents range from -5 to 0 and B agents range from 5 to 10). They are not offset along the first trait. In the asymmetric case A and B are offset by 10 along the second trait and 2.5 along the first. Therefore, they are not symmetric about the hyperplane that approximately mirrors them. For higher dimensions, the distribution is uniform over an n -cube whose edges are of length 5, and the offsets are 10 along one trait,

²¹This uniqueness result is proved in much greater generality in Theorem 4.11 of Chiappori et al., but mapping the match surplus functions from this environment to their framework and showing their conditions are satisfied to demonstrate that it is a special case of their result would be too long-winded for this paper.

and 2.5 along all others. For TU simulations, the matching disutility is the square root of the two-norm distance. In the TU case 30, 50, 70, or 100 agents are drawn for each side. For future reference, we'll call the above class of models our *baseline specification*. We'll run models with 2, 3, 4, or 5 traits over which agents have preferences. Larger numbers of traits were not simulated because the formulae defining the general form of the rotation matrix increase in both number and length as n increases, becoming unmanageable with more than a few dimensions.

With the simulated agents in hand and their preferences specified, we can simply run a version of Klumpp's inside-out algorithm in the NTU case to find the stable matching outcome.²² In the TU case, we can solve for the stable assignment and transfers pair by formulating the matching problem as a linear program (Shapley and Shubik (1972)).²³ Given a separating hyperplane along the horizontal axis, we can then run regressions with one b trait as the dependent variable and both $a = \mu(b)$ traits as the independent variables, where the resulting coefficients estimate the effects of a change in each on the b trait. In the idealized symmetric case we would expect the coefficients to be one for trait one on trait one, -1 for trait two on trait two, and zero otherwise. We run each specification a number of times—200 times for NTU specifications, and either 60 or 20 for TU depending on the model specification—and find the mean of recovered coefficient values, as well as a 90% confidence interval.

Before we can run the regressions in the asymmetric case, we must derive a rotation matrix and rotate the typespace to one where the hyperplane is horizontal. We find the vector from the center of mass (mean) of A to the center of mass of B , and construct the rotation matrix that maps that vector to the vertical axis. We then rotate the typespace, creating new “synthetic” traits 1 and 2 which should correspond to the vectors along which matching is positively assortative or negatively assortative. Finally, we run the regression using the synthetic traits.

²²See Online Appendix 1.4.3 for a summary of the algorithm.

²³See Online Appendix 1.4.2 for a summary of the algorithm.

7.2.2. *Understanding Differences in TU and NTU Simulation Results.* In Section 5, we saw that, in the NTU case, we had better fidelity to the predicted assortment results along the vertical trait. In the TU case we saw the opposite, though in this case the difference between the vertical and horizontal results was smaller. We'll now try to find some intuition as to why we'd see these results, starting with the NTU case. Recall that one's predicted match—their reflection about the hyperplane h —generally will not exist in these simulations. Thus, agents will have to match to some substitute with a different trait vector. Because the Euclidean distance is used and agents differ from their predicted match only along the vertical parameter, small changes from this outcome are much more costly in utility terms if they result in vertical change than if they result in horizontal change. In fact, the utility effect of a horizontal deviation per unit distance asymptotically approaches zero for small horizontal deviations. Thus, deviations from this ideal matching due to shortage or surplus of agents on a given side are likely to be realized primarily via horizontal deviations which are less costly. For example, in Figure 7.1 we see that a large horizontal deviation is on the same indifference curve as a tiny vertical deviation.

This explains the results in the NTU case, but what about TU? Why do TU stable matchings seem to better approximate the theoretical assortment prediction along the horizontal dimension, and approximate the vertical assortment prediction relatively less well? First, let's consider the simplified example in Figure 7.2. In this example the symmetry assumption does not hold, and we see how the unique stable matchings under the two transferability assumptions differ. In both the NTU and TU cases, the stable matches are shown by lines between agents, and the reflection matches for B agents predicted by propositions 1 and 3 are shown as $\mu(b_i)$. Notice that, in the NTU case, the actual stable matching $\{a_1b_2, a_2b_1\}$ has matches for B agents with the same θ_2 values (vertical types) as the predicted matches for those B agents. However, the θ_1 or horizontal types of the actual matches are very different than the symmetric matching prediction. In the TU case, by contrast, stable matches for the B agents have the same horizontal types as their predicted matches, but the vertical

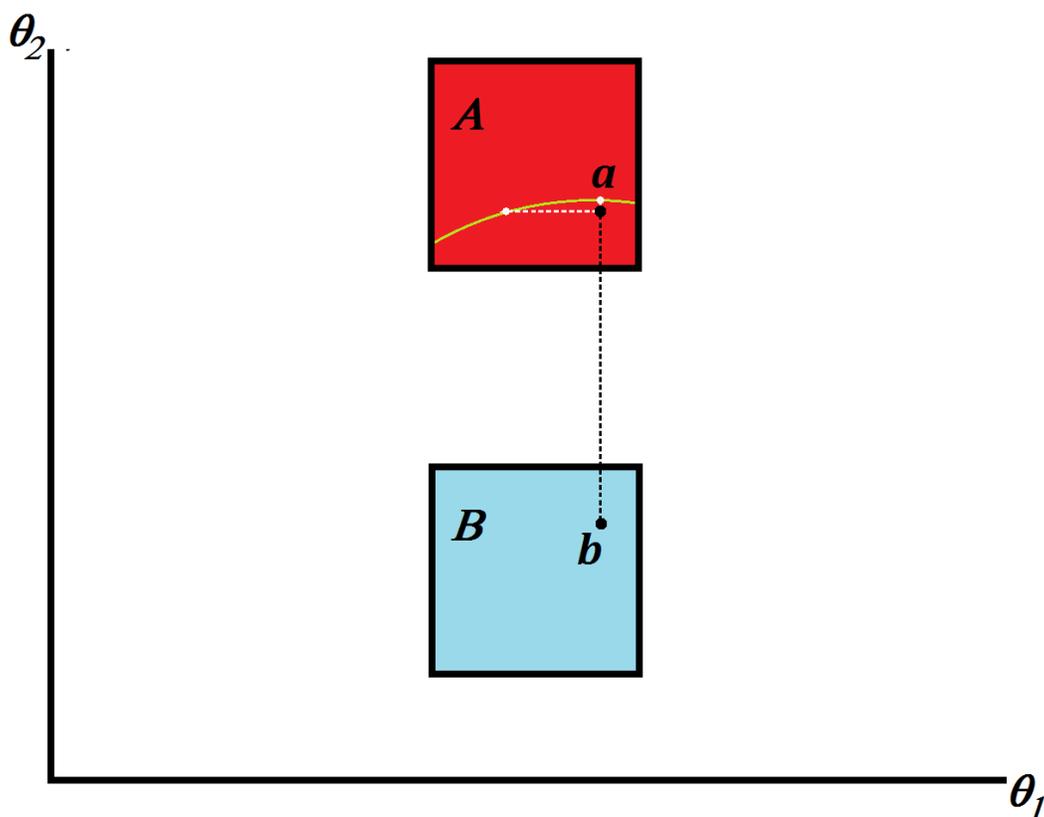


FIGURE 7.1. Horizontal deviations from the theoretical matching outcome have small utility effects relative to vertical deviations: an agent b is predicted to match to their reflection, a . Because a and b differ only along the “vertical” dimension, the distance between b and an agent near a has a much larger vertical component than horizontal component, so in the Euclidean distance horizontal deviations cause a much smaller change in distance and thus utility.

types differ. In the NTU case we have the predicted vertical assortment, but not the horizontal assortment, and in the TU case we have the predicted horizontal assortment, but not the vertical assortment. Note also that we cannot have both at the same time—the predicted matches do not actually exist, and the two configurations shown in 7.2 are the only possible matchings. While matching problems with more agents and less carefully chosen trait values will not be as extreme as this example, Figure 7.2 distills an important quality of n -dimensional matching problems of the sort we’ve been studying: there will generally be a tradeoff between assortment along one dimension and assortment along another. In

the two dimensional case, we can come up with matchings (not necessarily stable) that get closer to the horizontal assortment prediction, but this will often result in less fidelity to the vertical assortment prediction, and vice versa. We will see below that TU stable matching puts a greater premium on horizontal assortment than vertical assortment relative to NTU, so the stable matches in the TU framework will exhibit better horizontal assortment and commensurately worse vertical assortment.

Why is this? Whenever an A agent matches to a B agent, they create a negative externality for any A agents that would have liked to match to that B agent, since that B agent is now removed from the set of possible matches. In the NTU case, however, agents only care about their own match. It is irrelevant to an A agent whether matching to a particular B agent leaves another A agent with a far worse outcome. We can see exactly this in the NTU case of Figure 7.2. First, the closest agents, b_1 and a_2 , match. This leaves a_1 and b_2 with an extremely bad pairing, but since transfers are impossible, they have no recourse. In the TU case, by contrast, agents can freely trade their match surplus with their match in order to entice potential mates. In the TU case of Figure 7.2, we see that the stable matching is a_1 matching with b_1 and a_2 matching with b_2 . Notice that these matches are almost as good (close) for b_1 and a_2 as the a_2b_1 match, and that they are vastly better for a_1 and b_2 than the a_1b_2 match. Essentially, a_1 and b_2 are able to offer more of their surplus to a_2 and b_1 in order to attract them, and they greatly prefer this to getting a terrible match. We can think of the TU stable matching structure in terms analogous to the Coase Theorem—agents cause externalities by removing mates from the pool of potential matches, but they internalize those externalities since those affected can offer transfers embodying the cost that has been imposed on them.

How does this relate to horizontal and vertical preferences? Consider how vertical assortment affects individual utility: everyone agrees on the rank ordering of potential matches, so if one A agent matches to a B agent they find more desirable along this dimension, another A agent must match to someone they find less desirable. That is, one agent's gain is another agent's loss—this is the negative externality discussed before. In light of this, we see that

vertical assortment will make some agents better off, but it will make other agents worse off. Along the horizontal dimension, however, different agents prefer different matches, since they prefer their own type. In fact, horizontal assortment will make all agents better off, since it will give all agents their ideal type along that dimension. Thus, when we move from an NTU framework where externalities are ignored to a TU framework where they are accounted for, we can expect to see a shift towards horizontal assortment at the expense of vertical assortment.

Note that we've been running simulations where the support and variance of the distribution along each dimension is identical. If we were to increase the support of the distribution along the i th trait, or equivalently we were to increase the weight on the i th trait in the distance metric, the relative fidelity of the simulations to the horizontal and vertical assortment predictions would change. Specifically, putting more weight on a trait will generally improve the assortment along that trait, while worsening it along all others. An illustration of this can be seen in Figure 5.8 in Section 5, where the coefficient and R^2 dramatically improve for predicting a match's second trait by own second trait, while all other assortment results worsen. Intuitively, we can think about this in much the same way we looked at the better vertical assortment results in the NTU case. As in that case, the stable matching in the approximate symmetry environment is going to be a close approximation to the perfect symmetry stable matching, which has perfect assortment. However, it will not be possible to match each agent to their exact mirror-type because of the lack of perfect symmetry, so agents must deviate from their predicted matches. When a given trait is assigned more weight in agents' utility functions, they will be relatively more sensitive to deviations along this axis, and relatively less sensitive to deviations along other axes, so the assortment result will be stronger along the higher-weight axis and weaker along all others.

7.2.3. Additional NTU simulations for various market sizes and numbers of traits. Below we have the remainder of the baseline NTU simulations from Section 5. These are the horizontal same trait coefficients and the corresponding R^2 .

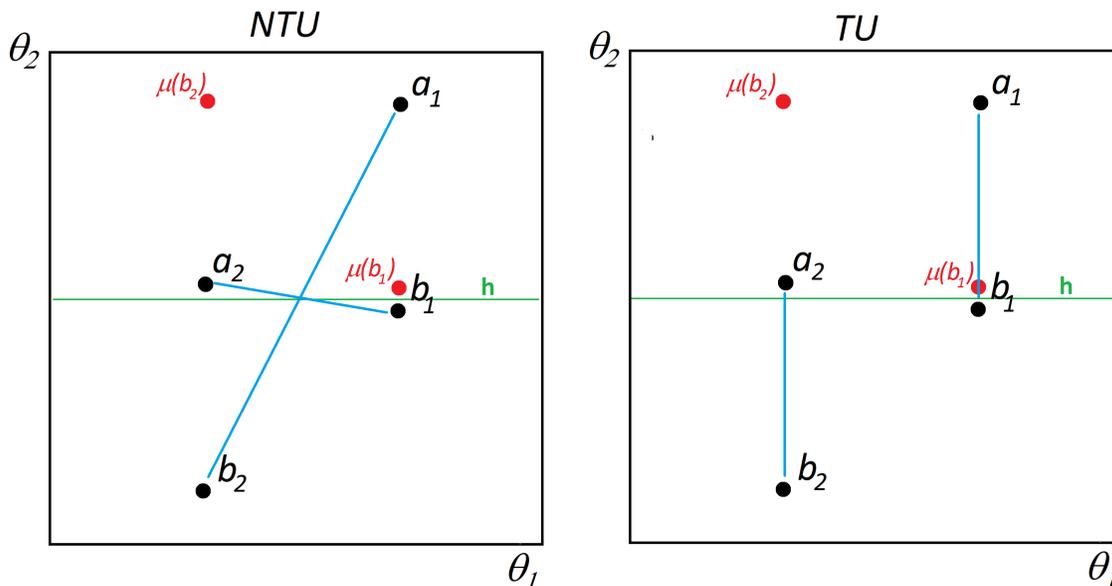


FIGURE 7.2. TU and NTU stable matching outcomes differ when the symmetry assumption fails. Here we consider the simplest nontrivial matching problem—two agents on each side—and assume mild convex preferences for distance for the TU case. Assume $u_{TU}(a) = -\sqrt{d(a, \mu(a))}$. The predicted but nonexistent symmetric matches for B agents are shown as $\mu(b_i)$.

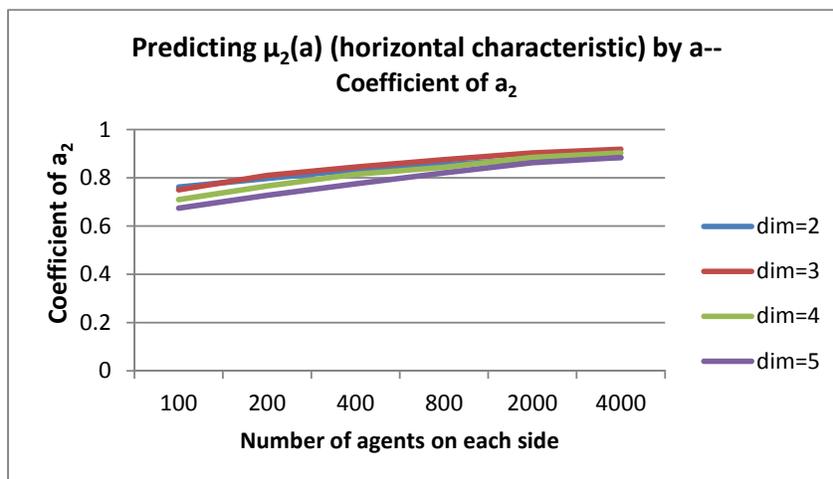


FIGURE 7.3. NTU with baseline specifications: predicting match’s horizontal trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

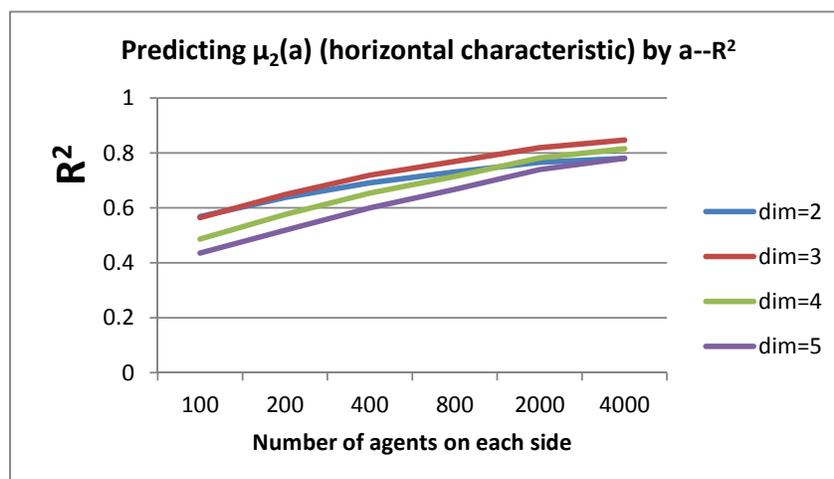


FIGURE 7.4. NTU with baseline specifications: average predicted R^2 for the regression of own horizontal trait on match's horizontal trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

In Figure 7.3 we see that, for an agent's match's horizontal trait, the average value for the agent's horizontal coefficient increases as the number of agents on each side increases. Recall that this is the case where the fidelity of the average coefficient values to the predicted result of 1 was poorest in the small market simulations. Here we see that increasing the number of agents on each side significantly improves the result. We also see that adding more traits causes more attenuation from the predicted result. It may be the case that, even with an arbitrarily large market, the average coefficient will remain below 1. Again, the baseline specification has asymmetric distributions, so there may be some deviation from the predicted symmetric result even with very large markets.

In Figure 7.4 we see that, for an agent's match's horizontal trait, the average R^2 increases as the number of agents on each side increases. Recall that this is the case where the fidelity of the average coefficient values to the predicted result of 1 was poorest in the small market simulations. Here we see that increasing the number of agents on each side significantly improves the result. We also see that adding more traits causes more attenuation from

the predicted result, except that the two dimension case improves more slowly with more agents on each side. It may be the case that, even with an arbitrarily large market, the average coefficient will remain below 1. Again, the baseline specification has asymmetric distributions, so there may be some deviation from the predicted result even with very large markets. The nature of the asymmetry also changes slightly as the number of dimensions changes, since the n-cube distributions are offset along each dimension. This could be the source of the strange behavior for the two dimensional case. Note that, to recover the R^2 , we regress only the same trait coefficient so that the R^2 gives us the variation in one's match trait solely attributable to the predicted same own trait effect.

7.2.4. *TU simulations for various market sizes and numbers of traits.*

Below, we have simulations for the TU environment analogous to the NTU simulations for various numbers of traits and different market sizes presented in Section 5. As discussed in Section 5, TU simulations are much more computationally intensive, so only small markets are simulated here, and the number of iterations in the Monte Carlo process is small—just 20 per specification. While these limitations make evaluation of the asymmetric stable matching outcomes much more difficult in than in the NTU case, we see that the results are qualitatively similar to the NTU results. Including more preference traits worsen cause attenuation in the same-trait coefficients from the predicted value of 1 or -1, and also worsen the R^2 , while the different-trait coefficients remain around the predicted value of zero. Increasing the size of the matching market improves the fit of the asymmetric model to the predictions of the symmetric model, though the inability to simulate markets with many hundreds or thousands of agents prevents us from seeing the limit behavior that we saw with the NTU simulations. The primary difference is a relatively better fit for the horizontal traits, and a relatively worse fit for the vertical traits, as compared the the NTU case. This is the same behavior we saw in the 2 dimensional, 100 agent per side TU vs. NTU comparison in Section 5. Given the similarity of the TU and NTU results in the region in which we can compare them (small markets), we can conjecture that the fidelity of the asymmetric stable assignments to the symmetric model predictions should drastically improve as the number of

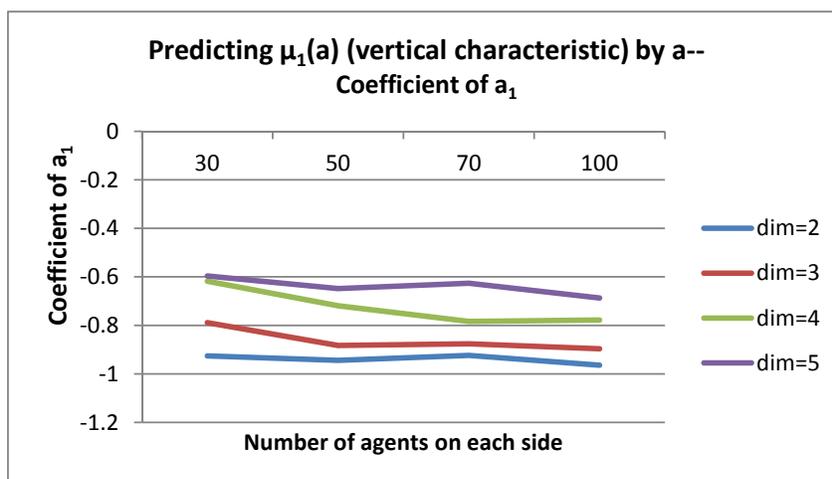


FIGURE 7.5. TU with baseline specifications: predicting match's vertical trait from own vertical trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

agents increases, as in the NTU case. It should become quite good even for several preference traits as the number of agents on each side reaches several thousand.

REFERENCES

- [1] Adachi, Hiroyuki. "A search model of two-sided matching under nontransferable utility." *Journal of Economic Theory* 113, no. 2 (2003): 182d more general preference structures.
 - [2] Becker, Gary S, 1973. "A Theory of Marriage: Part I," *Journal of Political Economy*, University of Chicago Press, vol. 81(4), pages 813-46, July-Aug.
 - [3] "Economics of the Family," Martin Browning, Pierre-André Chiappori, and Yoram Weiss, forthcoming.
 - [4] Burdett, Ken, and Melvyn G. Coles. "Marriage and class." *The Quarterly Journal of Economics* 112, no. 1 (1997): 141-167.
- [Chade, Eeckhout, and Smith] Chade, Hector, Jan Eeckhout, and Lones Smith. Search and matching models in macroeconomics. working paper, 2015.

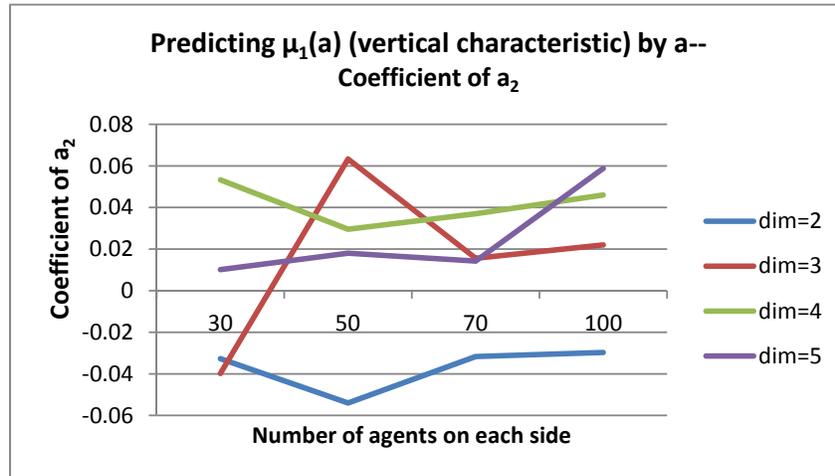


FIGURE 7.6. TU with baseline specifications: predicting match's vertical trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

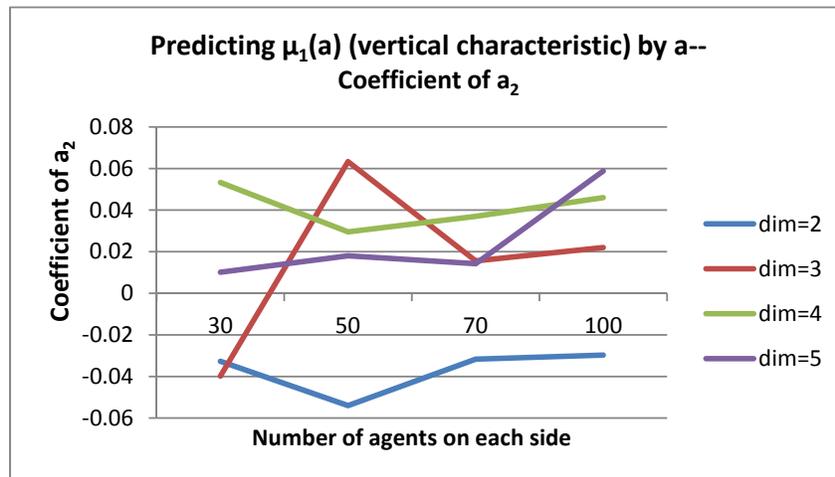


FIGURE 7.7. TU with baseline specifications: average predicted R^2 for the regression of own vertical trait on match's vertical trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

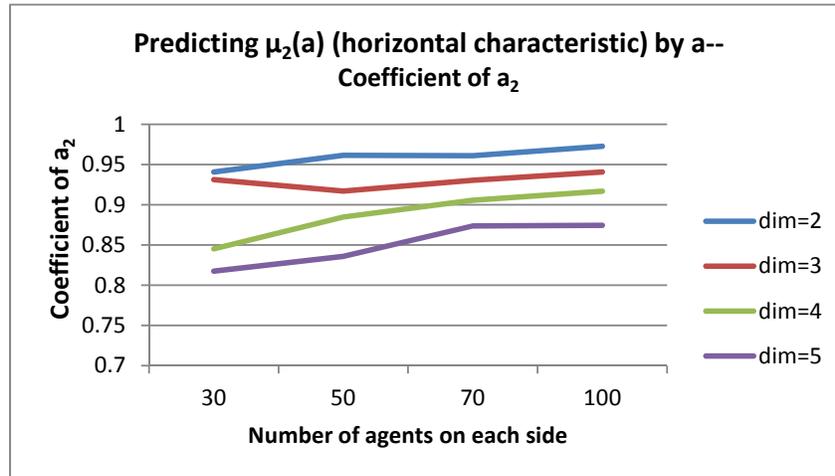


FIGURE 7.8. TU with baseline specifications: predicting match's horizontal trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

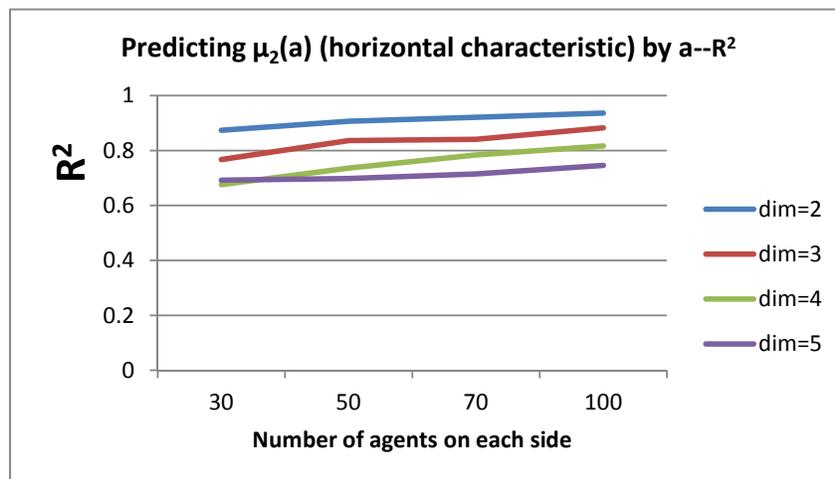


FIGURE 7.9. TU with baseline specifications: average predicted R^2 for the regression of own horizontal trait on match's horizontal trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

- [5] Chiappori, Pierre-André, Robert J. McCann, and Lars P. Nesheim. "Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness." *Economic Theory* 42, no. 2 (2010): 317-354.
- [6] Chiappori, Pierre-André, Sonia Oreffice, and Climent Quintana-Domeque. "Fatter attraction: anthropometric and socioeconomic matching on the marriage market." *Journal of Political Economy* 120, no. 4 (2012): 659-695.
- [7] Clark, S. (2003), "Matching and Sorting with Horizontal Heterogeneity", ESE Discussion Paper 98,
- [8] Clark, S. (2006). The uniqueness of stable matchings. *Contributions in Theoretical Economics*, 6(1), 1-27.
- [9] Clark, S. (2007). Matching and sorting when like attracts like. Working paper.
- [10] J. Eeckhout, Bilateral search and vertical heterogeneity, *International Economic Review* 40 (1999) 869–887.
- [11] Jan Eeckhout, 2006. "Local Supermodularity and Unique Assortative Matching," 2006 Meeting Papers 127, Society for Economic Dynamics.
- [12] Flanders, S. (2014). Characteristic Aggregation in Models of Matching Markets. Working paper.
- [13] College Admissions and the Stability of Marriage D. Gale and L. S. Shapley *The American Mathematical Monthly* , Vol. 69, No. 1 (Jan., 1962), pp. 9-15
Published by: Mathematical Association of America
- [14] Hitsch, Günter J., Ali Hortaçsu, and Dan Ariely. "Matching and sorting in online dating." *The American Economic Review* (2010): 130-163
- [15] Lauer mann, Stephan, and Georg Nöldeke. "Stable marriages and search frictions." *Journal of Economic Theory* 151 (2014): 163-195.
- [16] "Beauty Is a Beast, Frog Is a Prince: Assortative Matching with Nontransferabilities," Patrick Legros and Andrew F. Newman, *Econometrica*, Vol. 75, No. 4 (Jul., 2007), pp. 1073-1102
- [17] Lindenlaub, I. (2013). "Sorting Multidimensional Types: Theory and Application ", Job Market Paper.
- [18] "Two-Sided Matching with Spatially Differentiated Agents," Klumpp, Tilman, *Journal of Mathematical Economics* 45, 376-390, 2009.

- [19] Niederle, Muriel, Alvin E. Roth and Tayfun Sönmez. "matching and market design." The New Palgrave Dictionary of Economics. Second Edition. Eds. Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan, 2007. The New Palgrave Dictionary of Economics Online. Palgrave Macmillan. 13 April 2014 <http://www.dictionaryofeconomics.com/article?id=pde2008_M000356> doi:10.1057/9780230226203.1058
- [20] Shapley, Lloyd and Martin Shubik, 'The Assignment Game 1: The Core', International Journal of Game Theory, 1 (1972), 111-130.