Matching Markets with N-Dimensional Preferences

Flanders, Sam

University of North Carolina, Chapel Hill

February 2014

MATCHING MARKETS WITH N-DIMENSIONAL HORIZONTAL PREFERENCES

SAM FLANDERS

University of North Carolina at Chapel Hill

December 2013

Abstract. This paper analyzes matching markets where agent types are n-vectors of characteristics–i.e. points in $\mathbb{R}^n$–and agent prefer matches that are closer to them according to a distance metric on this set (horizontal preferences). First, given a few assumptions, I show that in the Gale-Shapley stable matching in this environment, agents match to a linear function of their own type. I show that restrictions on preferences are not as onerous as they may seem, as a rich variety of preference structures can be mapped into the horizontal framework. With these results in hand, I develop a highly stylized model of an online dating platform that helps consumers find and contact potential matches, where consumers have preferences over many characteristics (e.g. height, income, age, etc.) and have the option to pay to join the platform or look for a match off the platform. I characterize the firm's optimal pricing strategy and the concomitant market outcomes for consumers. Finally, I address an unanswered question in the matching literature–can multidimensional preferences be aggregated (e.g. into a univariate measure of quality) without changing the salient features of the model? I find that, in the dating platform model I introduced, consumer preferences can be aggregated without any change to firm strategy or market outcomes, providing some justification for the univariate-type matching models prevalent in the theoretical matching literature.

1. Introduction

Matching models are of great interest in economics, and, for matching problems relating to dating and marriage, nontransferable utility is a common assumption. That is, when agents in a dating market match, it is often assumed that one agent cannot offer a transfer to a potential match to entice them, perhaps because such an offer may be socially unacceptable or a matching will involve a pooling of resources regardless. Further, it might be assumed that preferences are are over a single parameter which is vertical, where all agents share a preference ordering over types, e.g. wealth; or horizontal, where agents prefer their own type, e.g. religion, race, etc. For both vertical (Becker[2]) and horizontal models (Clark (2003)[5], Clark (2007)[7], and Klumpp[10]), simple matching algorithms have been derived for continuous and discrete cases. However, it would be desirable to have multiple dimensions representing all the characteristics we believe agents have preferences over.

In this paper I 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, I 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, I test it's implications under both modest and moderate deviations from the symmetry assumption and find that the theoretical results well approximate nearly symmetric matching markets and even under significant deviations from symmetry the stable matching outcome closely corresponds to the theory. I also use the results to construct a theoretical model of an online dating market where agents have preferences over n characteristics, and I'm able to characterize the market structure for an arbitrary number of characteristics.

The results derived for horizontal preferences can easily be extended to vertical preferences, categorical horizontal preferences (e.g. there are several ethnic categories and agents prefer their own category), and even more general single peaked preferences (e.g. women most prefer men who are 80% their height plus 18 inches, with preference decreasing in distance from this ideal), 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 characteristics. Because of this, the assumptions are not terribly onerous and these results may plausibly be directly applied to real world matching markets.

This paper follows a rich literature on stable matching problems, starting with the seminal paper by Gale and Shapley [8], which introduced an algorithm for deriving stable matches given arbitrary preferences and a finite set of agents. While this algorithm is highly general, it utilizes an iterative process that must be run before matches can be derived. Thus, while it is quite useful for empirical analysis, it does not lend itself to being embedded in a theoretical model. Becker [2] found that positive assortative matching occurs when there is a continuum of types and the utility of a match is increasing in types and nontransferable–that is, when the two matching agents can't bargain over the apportionment of the utility of the match–and that assortative matching 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 pair up agents, so it is suitable for use in theoretical models, but the fairly onerous assumption of vertical preferences–higher types are universally preferred to lower types–limits its application. Legros and Newman [9] extended positive and negative assortative matching 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. Assuming horizontal preferences where agents want to match to their own type rather than the vertical preferences of Becker, Clark (2003) [5] gives an algorithm for finding stable matchings in a market with a finite set of agents. Clark (2007) [7] then treats the horizontal case with an infinite set of agents, finding a very simple matching result. Clark (2006) [6] also gives a condition guaranteeing a unique stable matching.Klumpp [10] derives a very simple "inside-out" algorithm for horizontal matching with finitely many agents.

The remainder of this paper is organized as follows: Section 2.1 explores the possibility of reducing an n-dimensional matching problem to a single composite type and finds serious shortcomings to aggregating characteristics. Section 2.2 shows how the n-dimension horizontal model can be used to represent a variety of characteristics. Section 3.1 derives the matching function for n-dimensional horizontal matching with symmetric agents and characterizes the matching outcome. Section 3.2 contrasts the symmetric and asymmetric cases and gives some conjectures about matching in asymmetric markets. Section 4.1 outlines the empirical model that is used to analyze the asymmetric case. Section 4.2 gives the details of the empirical analysis. Section 5 reviews the results. Section 6 develops and solves a model of a monopolist online dating platform using this matching result, and Section 7 summarizes the paper and suggests avenues for further research. Section 8, the appendix, provides several proofs not included in the main body of the article and background information on the Gale-Shapley algorithm.

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. We'll denote specific agents in A as a, and specific agents in B as b. These sides could be interpreted as, for example, 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 $b_1 \underset{a}{\succ} b_2$ we'll say a strictly prefers b_1 to b_2, and if $b_1 \underset{a}{\succcurlyeq} b_2$ we'll say a prefers b_1 to b_2 or is indifferent between them.

The goal of our analysis will be to find stable matchings in this environment. In this environment, a matching is a function $\mu : A \cup B \to A \cup B$ such that, for each agent $x \in A$ or $x \in B$, $\mu(x)$ is an agent on the opposite side or the empty set (no match), and $\mu$is a bijection. A stable matching is one in which there is no a and b such that $b \underset{a}{\succ} \mu(a)$ and $a \underset{b}{\succ} \mu(b)$. Such an (a, b) is a 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, whereas with stable matchings, so long as preferences remain the same, the matching should remain unchanged over time.

Preferences can be very general in the framework outlined so far, but we'll restrict them to horizontal preferences–that is, agents prefer matches closer to them. Specifically, we look at an environment where agents of each side S have n characteristics–that is, their type is an n-vector, $\theta_s \in \mathbb{R}^n$. For example. 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.

Figure 2.1. Mapping vertical preferences to a horizontal model

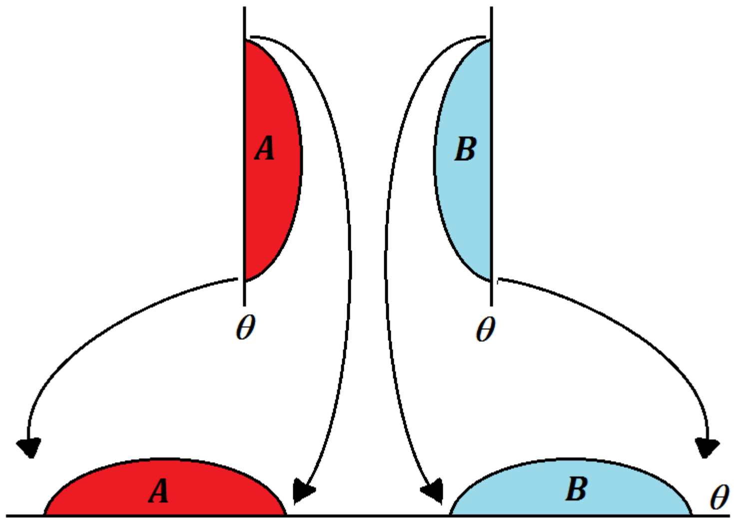

We'll also specify utility functions corresponding to these preferences. We'll assume nontransferable utility in this matching problem so that agents cannot offer some of their matching utility to a potential mate to induce them to match. Thus, the only criterion agents have when matching is the utility they get from their potential mate's type–agents can't strategically change their attractiveness to others.

2.2. **Modeling Various Preference Types in a Horizontal Framework.** As mentioned previously, we're considering agents with horizontal preferences over n characteristics, such that they want matches closer to their own type along these n dimensions. However, preferences over many characteristics are manifestly not horizontal for most individuals. For example, people generally prefer more attractive partners, not a partner of their own level of attractiveness. Luckily, while the horizontal preference assumption is restrictive, requiring preferences corresponding to a shared distance function over all n characteristics, it still allows considerable flexibility, and many types of preferences can be mapped into this framework.

In the attractiveness example, we assume people 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 type with vertical preferences (the higher the type, the more desirable to all agents on the other side), we can map the two distributions to the real line with preferences based on least-distance. A agent types unchanged and new B agent types equal to a constant minus the original b agent type. Assuming distributions with finite support and a large enough constant, the best B agent will have a type greater than the best A agent, and higher type A agents (lower type B agents) will be preferred by all agents on the opposing side. For example, we could have attractiveness for A and B, α_A and α_B, range from 0 to 1. Then we can map to the new A attractiveness using the identity function $\alpha'_A = \alpha_A$ and the new B attractivenss using $\alpha'_B = 2 - \alpha_B$. Since higher type A agents (lower 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 characteristic we'd like to include in the model, we can preserve that preference structure in the horizontal model we've developed. Note that we need the best A agent to be below the best B agent in the horizontal mapping, else the overlap agents will prefer their own type to the best type.

Using similar mapping techniques, we can even represent a larger class of single peaked preferences of which vertical and horizontal are special cases. Single peaked preferences are preferences where agents have an ideal type and prefer matches closer to that ideal type to matches further away from the ideal type. Given the Euclidean distance and utilizing a vertical (θ_1) and horizontal (θ_2) dimension for the single peaked characteristic, we we can see in figure 2.2 that the preferences in the 2-dimensional model are neither horizontal nor vertical, but that b most prefers a (b's indifference curve is shown in green, and no other

FIGURE 2.2. Mapping a single peaked characteristic to two horizontal dimensions

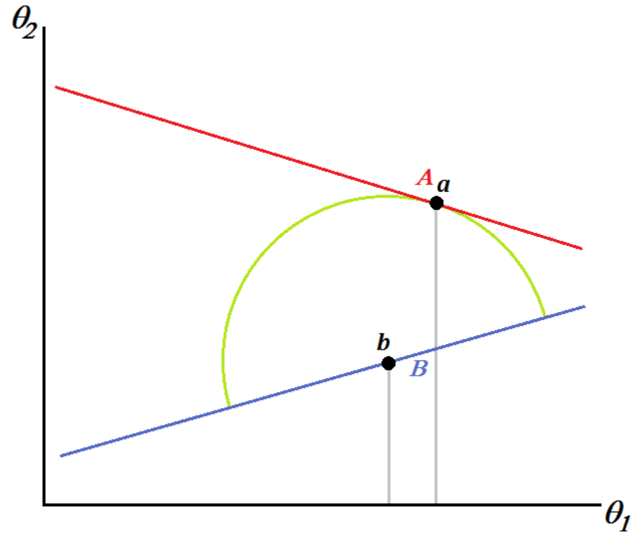

FIGURE 2.3. Mapping categorical characteristics into several horizontal dimensions

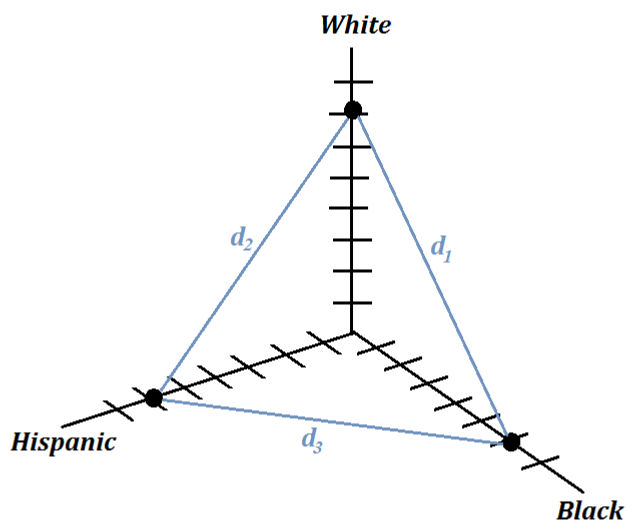

agent in A is within that radius), and a has a higher horizontal type than b, but b does not most prefer the maximal element of A. For the horizontal dimension, we simple use the value of the characteristic–say, height–so that θ_1 is height and agents prefer heights close to theirs. In dimension 2, however, we use a mapping like the one in the vertical case above. In this case, let the second height characteristic be equal to half the original height value for B agents. That is, $b_{\theta_2} = b_{\theta_1}/2$, where dividing by 2 means that the disutility of distance in the vertical dimension is half that of the horizontal dimension. For A agents, set the height in meters to $a_{\theta_2} = (2.5 - a_{\theta_1})/2$. Then the distribution of agents on each side will have support along lines as in figure 2.2, and the resulting preferences in the 2-dimensional space will be single peaked.

We can also represent preferences over non-ordinal categorical characteristics. For example, many individuals have strong preferences over race, preferring matches of their own race to others. We can also map this into a horizontal framework. For a categorical characteristic with k categories, we simply add k dimensions to our n-dimensional environment, with each dimension representing one of the categories, and agents of a given race having a positive value in their own race's dimension and zeroes in all others, where the magnitude of the positive component will depend on the relative strength of the preferences over various race pairs. In figure 2.3 we see that, using 3 horizontal dimensions, we can represent a mutually exclusive categorical racial characteristic, where in each dimension the value is 0 if you are not that race and some positive value if you are. Then every individual will have a positive value along one axis and zero along all others, and the distance between these points, d_i, will correspond to the disutility of matching for that race pair. Note that, because of the horizontal assumption, disutility is the same for both agents in a pair, but we have substantial leeway to vary the disutility of different pairs.

Thus, while our results only extend to matching problems where agents have horizontal preferences, we see that we can solve matching problems with a much richer set of preferences by mapping non-horizontal characteristics into preference-equivalent horizontal characteristics.

2.3. **Aggregation to a Single Dimension.** An obvious question when dealing with multiple preference dimensions is the following: can we reduce these preferences into a single variable? Specifically, we'll consider the following problem: we have a n-dimensional matching problem as described in 2.1, with either horizontal preferences in terms of distance or where agents are assumed to have vertical preferences over each characteristic. 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 1-dimensional model, where the 1-dimensional model is also of the type described in 2.1. Ideally, we'd like to preserve preference orderings for each agent and matching outcomes for each agent.

One route we might consider would be to simply look at one characteristic in isolation, but this presents clear difficulties. An example will be instructive: often, univariate matching models of dating and marriage are motivated with a type corresponding to wealth for the vertical case. However, in any practical application to dating or marriage markets, agents will have preferences over characteristics other than wealth. One such factor is attractiveness. Suppose then that attractiveness is perfectly negatively correlated with wealth in women and perfectly positively correlated with wealth in men. Further, suppose that that wealth and attractiveness are distributed in a uniform continuum over $[0, 1]$ on both sides, and that the utility of a match is the partner's wealth plus twice their attractiveness. Then the utility of a match is deceasing in wealth on the women's side and increase in wealth on the men's side. The result is that the matching will exhibit negative assortation in wealth–the higher the man's wealth, the lower the wealth of his match, and vice versa. On the other hand, solving the univariate model including only wealth gives positive assortation–higher wealth means a higher wealth match. This means that abstracting away from the multivariate case gave us the opposite conclusion! If we assume that all vertical parameters are perfectly positively or negatively correlated on both sides we can avoid this complication, but this is a very special case, and in general cherry-picking a single trait and matching on that will give us an inaccurate matching outcome.

This suggests the following question: can we somehow aggregate multiple type dimensions into a composite type and retain the same matching outcome? In fact, a single aggregated type is used in a number of papers in the literature, such as "pizazz" in Burdett and Coles (1997)[4]. With an aggregate univariate type, we could use a univariate model without the problems outlined above. However, this can also be problematic.

If we restrict attention to n-dimensional vertical preferences and further assume that the underlying assumption that all agents have the same preference ordering over each type parameter extends to type n-vectors–that is, that everyone agrees on the ordering for any pair of type profiles–then we have a very nice result, as we can preserve the matching outcome in a univariate matching model where type is strictly increasing in composite quality. If we are interested in cardinal preferences as well, we can impose some modest structure–making the utility of a match a function only of the composite qualities of the two agents. Then the set of possible stable matchings will be preserved as well as the matching utility. On the other hand, if one's own utility of a match with a fixed opposing agent differs for two own-type vectors with the same composite quality level (same utility of matching for any opposing agent) then the matching utility in the univariate model will not correspond to the multivariate problem, but the matching function will still be preserved and therefore true utility can easily be recovered.

However, there are many characteristics that are manifestly not vertical and are highly relevant to matching. With demographic characteristics like ethnicity and religion and other characteristics like life goals and values we would instead expect to see many people preferring their own type. That is, they should exhibit horizontal matching. As before, ignoring these characteristics will almost certainly yield inaccurate matching predictions, so being able to relate a multivariate horizontal matching problem to a univariate matching problem would be desirable. In fact, when horizontal preferences are included there even more reason to expect more problems with matching models considering preferences over only one characteristic. While many traits that are modeled vertically are positively correlated, yielding qualitatively similar matching outcomes for the univariate and multivariate matching problems, horizontal types and vertical types need not be. For example, women from cultures with many short individuals may prefer men from their culture (same horizontal type) and also prefer taller men, which is negatively correlated with being a member of their culture. Matching only on height would match the women almost exclusively to men of other cultures, which would be contrary to what is observed. Again, we must turn to aggregation. Mapping from n to 1 dimensions is more difficult than the vertical case, and the result we get will be weaker. Before that, however, we need some preliminaries.

First, we need an algorithm to solve the original (n-dimensional) matching problem in a way that will allow us to map it into one dimension. The most commonly used algorithm for solving matching problems in the environment of section 2.1 is the Gale-Shapley or Deferred Acceptance Algorithm [1]. However, this algorithm does not lend itself to mapping into one dimension. On the other hand, Klumpp derives an "inside out" algorithm to solve univariate horizontal matching problems [?], which simply iteratively matches the closest agents to one another and removes them from the problem. More concretely, consider the environment in Section 2.1 with a finite set of agents, and, in the first step, choose $(a_1, b_1) \in \underset{a \in \mathrm{A}, \mathrm{b} \in \mathrm{B}}{argmin}\ d((a,b))$. This is our first pairing in the stable matching we're constructing. Suppose $k-1$ pairs have been matched out. Define $A_{k-1} = \{a_1, a_2, ..., a_{k-1}\}$ and $B_{k-1} = \{b_1, b_2, ..., b_{k-1}\}$. In the kth step, choose $(a_k, b_k) \in \underset{a \in A \setminus A_{k-1}, b \in B \setminus B_{k-1}}{argmin}\ d((a,b))$ Proceeding this way, we'll match out agents until one or both sides are fully matched out.

Since, in horizontal matching, agents prefer closer matches, and distance is symmetric, a distance minimal pair of A and B agents must be the mutually most preferred match by the pair among currently unmatched agents, and already matched agents weakly preferred their match to the current matches, so the matches are stable. This algorithm can easily be extended to n-dimensions given some additional structure on preferences. Specifically, we will assume that preferences on both sides correspond to a distance metric d over $\mathbb{R}^n$. Note that the proof sketch for Klumpp's algorithm above relies only on the symmetry of the distance metric, so the algorithm can be applied to the n-dimensional case without any adjustment so long as preferences are dictated by a shared distance metric over the n-vector types.

With this method in hand, we can derive a simple n to 1 dimensional mapping. First, consider a discrete n-dimensional distribution f_n and the solution supplied by Klumpp's algorithm. Take all matches where $d > 0$ and assign them a univariate type equal to half the distance from their match in the stable matching, multiplied by -1 for B agents. Now, for each agent with a match of distance 0 and where some other agent of identical type on the same side has a match with nonzero distance, assign them half that nonzero distance as their univariate type, multiplied by -1 for B agents, and assign their identical match on the other side to that type as well. Finally, for agents whose type only receives perfect matches, assign the univariate type 0. Call this univariate distribution f_1.

To show that this preserves the multivariate matches as possible univariate matches, take a in A such that $\mu_n(a) = b$. Suppose $\theta_{1b} = \theta_{1a}$. Then $d_n(a,b) = 0$ and either the one-dimensional type of both is 0, such that they are a valid match in the univariate analogue, or a's type is $\theta_{na} \neq 0$. Then b must be type θ_{na} as well, since it was a's match in the n-dimensional case. Finally, suppose $\theta_{nb} \neq \theta_{na}$. Let $d_n(a,b) = D$. Then $\forall a' \in A$ where $d_1(a', b) < D$, then $|\theta_{1a'}| < D/2$, so $d_1(a', \mu_1(a')) = 2|\theta_{1a'}| < |\theta_{1b}| + |\theta_{1a'}| = d_1(a', b)$, so b cannot deviate to a more desirable a'. By a symmetric argument, a has no more preferred available match than b, so there is no blocking pair and they are valid matches in the one-dimensional case.

This algorithm does not perfectly preserve the structure of the n-dimensional matching. In particular, this matching only preserves the quality of matches, not the types that match together. While no matching over preferences based on type will go beyond preserving the types that match, as in the matching problem two agents with the same type are interchangeable and their matches can be traded amongst themselves without violating stability, in this horizontal mapping we can also stably match agents that cannot match in the n-dimensional case. For example, consider a matching problem where each side is composed of two agents and all four form the vertices of a rectangle, as in Figure 2.4. Then there will be two matching pairs,

[1]see appendix 8.3 for a summary of the algorithm

FIGURE 2.4. Mapping a multidimensional matching problem to one dimension creates extraneous matches

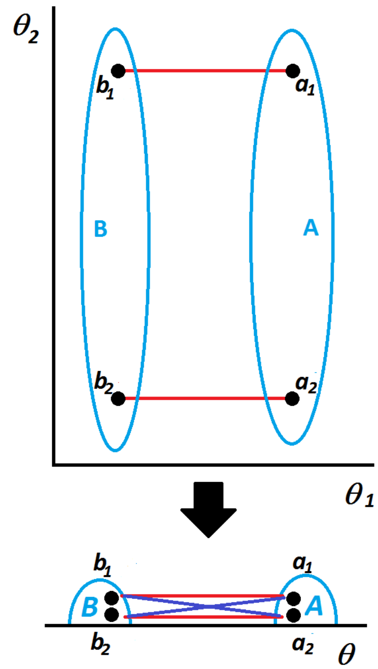

with each agent matching to the adjacent agent on the other side. When this is mapped to one dimension, there will be a mass point on each side with both of that side's agents, and either pairing would be stable, while in the n-dimensional case the alternative pairing would leave all agents worse off and thus be unstable. Also, agents of identical types that match to types of different nonzero distances will have different types in the one-dimensional matching, since their 1-type is derived from the distance, which is different even though their n-type is the same, as seen in Figure 2.5. However, if we are only interested in the preserving the match quality that agents receive, this mapping can still justify the use of a univariate composite type in theoretical work.

Because this mapping requires the n-dimensional matching outcome in order to generate the one-dimensional analogue, it cannot be used solve n-dimensional matching problems by turning them into more tractable one-dimensional problems. However, we may still find use for it as a justification for using one dimensional theoretical models. In this case, we don't need to match our one dimensional matching market to a particular n-dimensional equivalent, we just want to know that n-dimensional models can be represented in one dimension while retaining characteristics we consider salient. If all that is required is that the original matching outcome remain stable, then clearly this algorithm is sufficient. However, the one dimensional model only retains the matching outcome, not the underlying n-dimensional distributions. Nor does it preserve the full preference orderings, as in the one dimensional problem your preferences for opposing agents depend on the distance from their match, not from you. For example, in Figure 2.2 $a_3$prefers b_2 to b_1, but the distance betweena_1 andb_1 and a_3 and b_3 are the same, so the distance between a_3 and b_1 in the one dimensional mapping is the same as the distance between a_1 and b_1, while the distance between a_3 and b_2 is greater, since b_2 is further from a_2. In general, if we want to e.g. model the effect of changes in the distributions of agents or compare outcomes under different matching processes (for example, matching with search frictions), the one dimensional model will yield incorrect results matching outcomes.

Thus, we see that this mapping method of one dimensional representation of n-dimensional horizontal matching problems has limited applications to theoretical matching models. However, while the matching outcomes may differ substantially, there still may be some use for univariate analogues to multidimensional matching problems. This topic is explored in Section 6.

FIGURE 2.5. Mapping a multidimensional matching problem to one dimension changes preference orderings

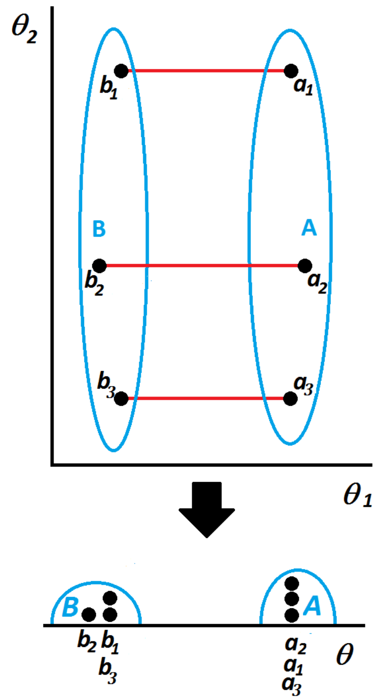

3. Theoretical Results for Symmetric Distributions

3.1. **Nontransferable Utility Matching with Symmetric Distributions.** If we can't reduce an n-dimensional horizontal problem to a one-dimensional one, we must then 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 and even characterize the type of one's match as a linear function of one's own characteristics. The model here is still the one outline in section 2.1, with disutility of distance given by the Euclidean distance metric $d(a, b)$, where d is shared by all agents on both sides. Agent a has type $\theta_a \in \mathbb{R}^n$, and agent b has type $\theta_b \in \mathbb{R}^n$. Denote the unit norm to h with origin at zero $\eta(h)$. Denote the normal to a hyperplane h beginning in h and terminating at a point a $\eta(h, a)$ Denote the reflection or Householder matrix of h $R(h) = I - 2\eta(h)\eta(h)^T$. Denote the vector with all zero components except a value of 1 at component i e_i. Define $d_\eta(a, b)$ as the distance between a and b along vector η and $d_{h_i}(a, b)$ as the distance between a and b along the ith 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 also need a result from linear algebra: $d(a, R(h).a) = a - 2\eta(h, a)$. We'll need to make several assumptions to get a simple matching function:

- **Assumption 1 (SEP)** : $\exists h = \{x : ax = k\}$ for some x 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. Note that, if there is at least one vertical characteristic, we can satisfy this condition, 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. If the distance between the distributions is large enough, we can find a hyperplane that satisfies SEP as well as assumption 2.

- **Assumption 2 (REF)** : the distribution A is the reflection about h of the distribution B. That is, $\forall a \in A, \exists b \in B$ such that $R(h).a = b$, and $\forall b \in B, \exists a \in A$ such that $R(h).b = a$.

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.

- **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 between the points along the normal to a hyperplane and the distance between the points along that hyperplane, which will be crucial both for proving the that agents most prefer matching to their reflections and for actually guaranteeing that that result is true (it won't generally be 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, whereas for the, say, 1-norm or supnorm, the indifference curve will be a hypercube, which is not invariant to rotation.

What we want to do, essentially, is what is seen in figure 4.2. We want to take matching problems with distributions A and B reflected by a hyperplane in any sort of orientation in the typespace, and solve an equivalent problem where they are reflected by a hyperplane normal to one of the basis vectors, so that an agent and its reflection differ only along one dimension, and that distance can be additively separated from the others under the root in the norm. This will be critical in the proof, but such a rotation, or equivalently a change of basis in the norm, will change the matching problem unless the norm has rotational symmetry.

Now we can state the result:

Proposition 1. *(Continuous Symmetric NTU Matching) Given a two sided NTU matching market with sides A and B, suppose $\exists$ a hyperplane $h \subset \mathbb{R}^n$ satisfying SEP and REF. 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. For a contradiction, consider the matching outcome of proposition 1 and suppose there is a blocking pair (a_1, b_2) such that $b_2 \underset{a_1}{\succ} \mu(a_1) = b_1$ and $a_1 \underset{b_2}{\succ} \mu(b_2) = a_2$. 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 $d(a_1, b_1) = 2d(a_1, h) = d_\eta(a_1, h)$, $d(a_2, b_2) = 2d(b_2, h) = d_\eta(b_2, h)$. But, since d is the Euclidean distance,

$$d(a_1, b_2) = \sqrt{\sum \begin{matrix} n \\ i = 1 \end{matrix} \; d_i(a_1, b_2)^2}$$

and equivalently, we have

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

Contradiction. $\square$

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, we can immediately derive a characterization of the matching function from this result:

Corollary 2. $\dfrac{\partial \mu(a)}{\partial a_1} = -e_1$, $\dfrac{\partial \mu(a)}{\partial a_i} = e_i$ *for* $i > 1$

Proof. Note that $\mu(a) = a - 2\eta(h,a) = a - 2d(h,a)e_1$. Then $\dfrac{\partial \mu_i(a)}{\partial a_i} = e_i - 2\dfrac{\partial d(h,a)e_1}{\partial a_i}$ and $\dfrac{\partial d(h,a)e_1}{\partial a_i}$ is 0 if $i > 1$ and e_i if $i = 1$, since, as one varies a_i, a moves parallel to the hyperplane. $\square$

Corollary 2 has a simple interpretation–along the normal to the hyperplane dividing the two distributions, that is, along the dimension across which the distributions face each other, the matching exhibits negative assortation. Along vectors orthogonal to the first, the matching is perfectly assortative. 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. Imagine moving left and right or up and down in front of a mirror (moving parallel to the mirror). Your reflection moves exactly as you do, that is, the change in your reflections position varies positively and one-to-one with your movement. However, when you move towards the mirror, your reflection moves towards the mirror too, its position moving in the opposite direction as yours, and if you move away from the mirror your reflection similarly moves in the opposite direction.

3.2. **Transferable Utility Matching with Symmetric Distributions.** We now move on to an analogues matching 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 is, convexity of the disutility of distance. This will ensure that the marginal cost of a closer pairs being moved further apart is greater than the marginal cost of further pairs being seperated further. Convexity is important in the TU framework because, since agents are free to bargain with each other over that division of match surplus, TU-stability requires that the sum of match utilities be maximized. Thus, as with Becker and our previous result, along the vertical dimension closer agents will match with closer, further agents will match with further. 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 ideal.

Generally, explicitly solving for TU-stable matchings is very difficult, since one must find not just the matching but also show there are allocations that support that matching as stable. Finding those allocations can be very difficult, but here we can simply assume the allocations are even splits of the match surplus and then show that this is stable. Generally, the allocations can be thought of as a shadow price for the agent's presence in the matching market[3], and 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 identical, and in the stable maching agents match to their mirror type, so 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 3. *(Continuous Symmetric TU Matching) Given a two sided TU matching market with sides A and B, suppose $\exists$ a hyperplane $h \subset \mathbb{R}^n$ satisfying SEP and REF. Suppose agents prefer closer matches in the Euclidean distance metric and the match utility is convex and decreasing in distance. Then all agents matching to their reflection is stable. That is,$\mu(a) = a - 2\eta(h, a) = R(h).a$. Further, every pair splitting the match surplus equally is an allocation consistent with the matching being stable.*

Proof. For a contradiction, consider the matching outcome of proposition 3 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

$$d(a_1, b_2) = \sqrt{\sum_{i=1}^{n-1} d_{h_i}(a_1, b_2)^2 + d_\eta(a_1, b_2)^2}$$

$$\geq \sqrt{d_\eta(a_1, b_2)^2}$$

FIGURE 3.1. Distributions will generally not be symmetric

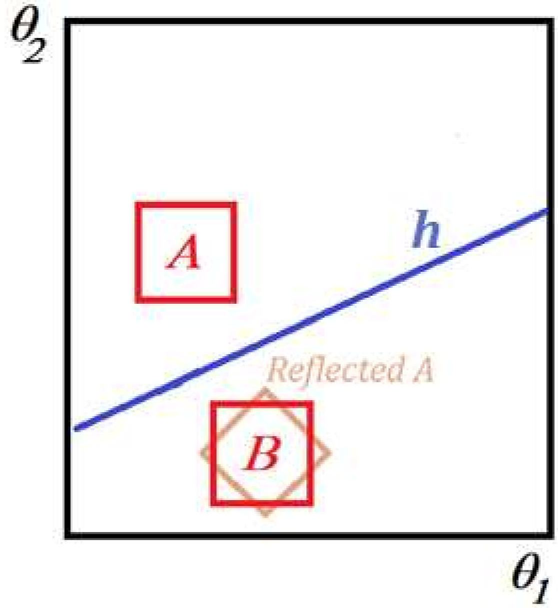

$$= d_\eta(a_1, b_2)$$
$$= d_\eta(a_1, h) + d_\eta(b_2, h)$$
$$= (d(a_1, b_1) + d(a_2, b_2))/2$$

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. □

3.3. **NTU Matching with Asymmetric Distributions.** Proposition 1 gives us an easily derived and interpreted matching function. However, in practice, we're very unlikely to encounter perfectly symmetric sides. That is, we can't expect the n-dimensional distribution of men to be the exact reflection of the n-dimensional distribution of women about a separating hyperplane. The natural question to ask, then, is whether the sort of approximate symmetry we might reasonably hope to see in the data exhibits approximately the same matching structure. Unfortunately, deriving general matching outcomes for n-dimensional horizontal matching markets is extremely difficult.

We can make some conjectures, though. The factors that ensure the assortation 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) 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 isn't a symmetric match for all agents, one side or another will be in shortage at 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 characteristics on corresponding match characteristics and possibly some modest effect of own characteristics on non-corresponding match characteristics. Notice, however, 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, say, 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, and if the separation of the two sides in the vertical characteristic is large enough, or there are enough vertical characteristics with smaller separations, we should have enough space between the distributions to fit a hyperplane that both separates and reflects

FIGURE 4.1. Rotating the typespace

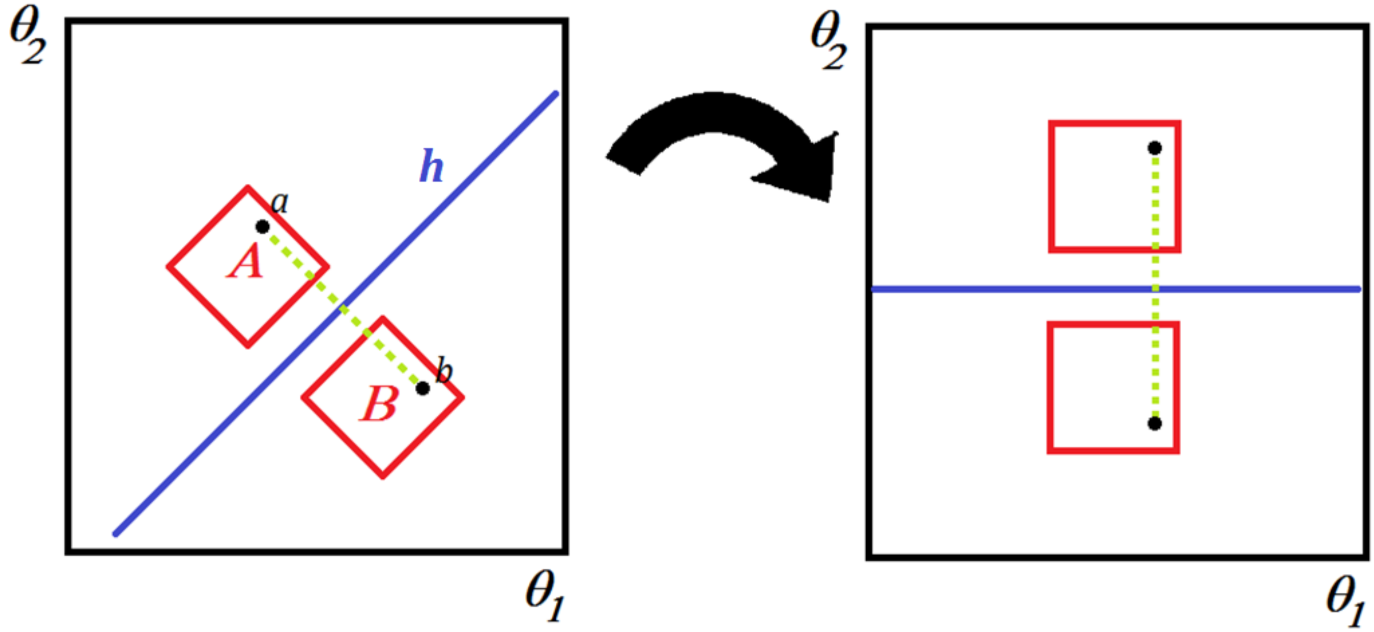

between them. Thus, while a lack of symmetry may change the matching outcome, getting an approximate reflection should not generally be a problem in an empirical setting.

4. EMPIRICAL FRAMEWORK

4.1. **Empirical Setup.** We now consider how to test the validity of applying the symmetric distributions 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 don't 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-characteristic effects significantly below one and other effects may be nonzero.

To simplify the analysis and facilitate visualization of the model, we use a two-dimensional typespace. In both cases, the observations on both sides are drawn from square bivariate uniform distributions. In the first case, they are stacked vertically as in the righthand portion of figure 4.1, such that they are symmetric about h, h is horizontal, and the expected own-characteristic effects should be along $\theta_1$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, and 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. Additionally, it will be necessary to estimate the relative preference weights of each characteristic for each side of the market, as those generally will not be known a priori. Until now, relative weightings have been implicitly assumed to be 1. Note that if the relative weights differ between the two sides of the market, we will not have the symmetry in preferences necessary for the inside out algorithm to function, and the results derived earlier need not hold.

4.2. **Simulation Model Specifications.** In both specifications, 200 agents are simulated for each side, drawn from a independent bivariate uniform distribution with support from 0 to 5. In the symmetric case, the two distributions are offset by 10 along the second characteristic (thus, along the second characteristic, A agents range from -5 to 0 and B agents range from 5 to 10), while in the asymmetric case they're offset by 10

FIGURE 5.1

Relative Preference for Characteristic 2

	Symmetric	Asymmetric
Side A	1.076	1.011
Side B	1.022	0.976

along the second characteristic and 2.5 along the first so that they are not symmetric about the hyperplane that approximately mirrors them.

With the simulated agents in hand and their preferences specified, we can simply run the Gale-Shapley algorithm[2] to find the stable matching outcome. Given a separating hyperplane along the horizontal axis, we can then run regressions with one b characteristic as the dependent variable and both $a = \mu(b)$ characteristics as the independent variables, where the resulting coefficients estimate the effects of of a change in each on the b characteristic. In the idealized symmetric case we would expect the coefficient to be 1 for characteristic 1 on characteristic 1, -1 for characteristic 2 on characteristic 2, and zero otherwise.

There are some additional steps we must take. Specifically, we must find the relative weight for the second characteristic for both sides in order to find the matching outcome, and then derive a rotation matrix and rotate the typespace to one where the hyperplane is horizontal before running the regression. The method I use to find the optimal weightings is minimizing the sum of squared errors (SSE) of the matching function $\mu(b)$. Since I am simulating the data, I choose equal weights as the "true" weighting for both sides, then compare the matching outcome with equal weights to the matching outcome with the guessed weights to find the SSE.

Unfortunately, the SSE is a stepwise function of the guessed relative weights and is not single peaked, so finding a global minimum is difficult. To overcome this issue, I estimate a differentiable approximation with a cubic interpolation fitting using 50 SSE's from randomly drawn weight-pairs, then randomly draw 50 initial guesses from a distribution centered on (1,1), running Newton-Raphson for each and taking the minimum over the 50 trials. Given these weightings, I 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. I then rotate the typespace, creating new "synthetic" characteristics 1 and 2 which should correspond to the vectors along which matching is assortative or negatively assortative. Finally, I run the regression as before using the synthetic characteristics.

5. RESULTS

First, we'll look at the estimated relative preference weights. In both the symmetric and asymmetric case we have relatively close approximations to the true relative weight of 1. These estimates could be improved by using more points in the SSE fitting process and by drawing more initial points for optimization.

Now we'll look at characteristic 1. Recall that this is the characteristic where both sides have the same distribution in the symmetric case before rotation and in both cases the distributions will overlap after rotation. It's essentially the horizontal characteristic after rotation, while characteristic 2 is vertical after rotation. We see that the same-characteristic coefficients are significantly less than 1, but also highly significantly greater than zero. Some attenuation is observed, as was conjectured, but the same-characteristic effect is still very strong and relatively close to 1, even in the asymmetric case. The effect of characteristic 2 after rotation, on the other hand, is extremely close to zero in both cases, though it is statistically significantly less than zero in the symmetric case. Note that in the symmetric case the rotation is extremely minor as the underlying distributions require no rotation and any rotation is due solely to random variation in the draws. The primary change from symmetry to asymmetry is that the standard errors rise and the R^2 decreases.

[2]See Appendix 8.3 for a summary of the algorithm

FIGURE 5.2

Predicting $\mu_1(a)$ by a

	Symmetric		Asymmetric	
	Original	Rotated	Original	Rotated
a_1	0.81	0.8112	0.714	0.825
	(0.0273)	(0.0273)	(0.0411)	(0.0432)
a_2	0.066	-0.05	-0.46	0.0418
	(0.0288)	(0.0288)	(0.0437)	(0.0432)
R^2	0.82	0.81	0.66	0.65

FIGURE 5.3

Predicting $\mu_2(a)$ by a

	Symmetric		Asymmetric	
	Original	Rotated	Original	Rotated
a_1	0.026	-0.01	-0.413	-0.006
	(0.0055)	(0.0055)	(0.0105)	(0.007)
a_2	-0.942	-0.943	-0.88	-0.991
	(0.0058)	(0.0058)	(0.0111)	(0.007)
R^2	0.99	0.99	0.98	0.99

Thus there are other factors influencing matching in the asymmetric case, but linear same-characteristic effects still explain most of the variation and their average effect doesn't change.

Finally, we consider characteristic 2, the vertical characteristic. We see that in this case the linear same-characteristic effects explain virtually all of the variation in the vertical characteristic, and the coefficient is very close to -1, showing only minor attenuation, and in fact is even closer in the asymmetric case where it is statistically indistinguishable from -1. The opposite-characteristic effects are also quite close to zero, although statistically different in the symmetric case. Unlike the horizontal parameter, there is no noticeable loss of explanatory power when we drop the symmetric distribution assumption.

These results are quite positive for the theory. Even in very small, coarse matching markets of a few hundred agents, the idealized result well approximates the actual outcome. The primary question we're left with is why the horizontal characteristic's effect is significantly weaker and why it explains less of the horizontal variation. The most obvious possibility is that, because the Euclidean distance is used and, under the ideal symmetric outcome, agents differe from their match only along the vertical parameter, and thus 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. 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 5.4 we see that a large horizontal deviation is on the same indifference curve as a tiny vertical deviation.

FIGURE 5.4

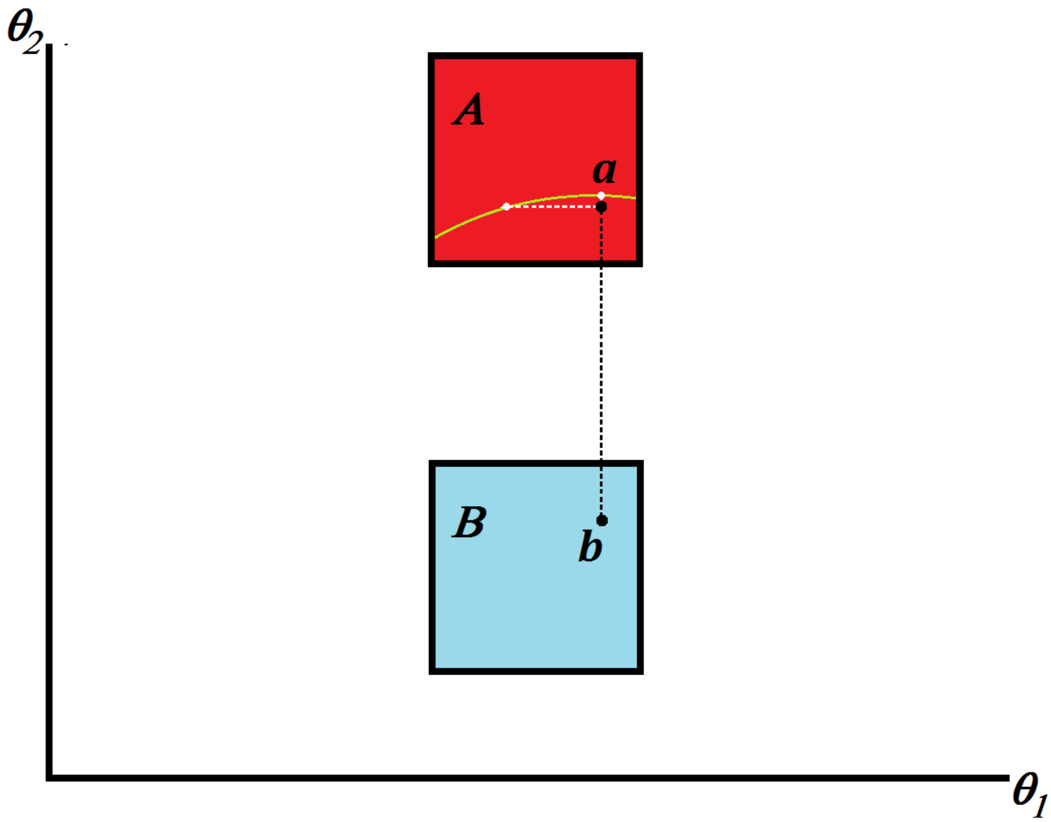

6. THEORETICAL APPLICATION

6.1. **Model.** So far, we've looked at the empirical applications of this result, but have only briefly mentioned how it could be applied to theory work. We will now work through a simple model of a monopolist firm[3] selling access to a platform at single price for all agents, where consumers on the two sides can observe the n-dimensional types of all other agents on the platform and propose matches to any agent in the other side–that is, the environment delineated in Section 2.1–leading the stable matching outcome derived in section 3. Agents have the outside option of searching for a match on their own, simplified in this model to receiving a single random draw from the overall distribution. Essentially, this is a stylized model of a platform that provides agents with information on potential matches and helps them make contact, e.g. an online dating website.

For concreteness, suppose the mass of agents are distributed over two n-dimensional continuous jointly uniform distributions, A and B, with support on $[0,1] \times [0, \frac{1}{n-1}]^{n-1}$ and $[2,3] \times [0, \frac{1}{n-1}]^{n-1}$, respectively. The first dimension is the dimension normal to the separating hyperplane and the remaining n-1 dimensions are along the separating hyperplane. The choice of support here appears peculiar, but will be explained once the model is solved.

Despite focusing on the Euclidean distance earlier due to its rotational symmetry, we will switch to the 1-norm here for tractability.[4] The expected distance of a random draw in Euclidean distances will give us an expression that will make computing the exact form of the cutoff for joining the platform intractable, so

[3]We can easily model a competitive market using the same framework. The only difference will be a lower price and a correspondingly shifted cutoff region where more consumers join the platform.

[4]Note that the proofs of the symmetric matching propositions assuming the Euclidean distance apply to the 1-norm as well, so long as the distributions don't require a change of basis to set the reflecting hyperplane orthogonal to the first dimension.

we look at the analogous case for the 1-norm, and later discuss how the Euclidean norm case relates. The useful property of the 1-norm that we take advantage of is that the distance is additively separable.

We can interpret this model in relation to an aggregate variable "pizazz" model like Burdett and Coles[4], where agents are modelled as having preferences over a single aggregate characteristic, as well as the investigation of aggregation pursued in Section 2.3. Specifically, we have a vertical component and $n-1$ horizontal components, where we scale down the lengths of the horizontal components so that the average distance between agents remains the same for each n, with that distance split amongst progressively more dimensions. We can then compare higher aggregation models (lower n) versus lower aggregation models (higher n).

Individuals have the option of joining the platform or getting a random match from the entire distribution, but cannot do both. This model essentially assumes the platform is "small" so that the agents who join the platform don't decrease the mass of agents of their type off the platform.[5] As with all two sided market models, we may have situations where no one believes anyone else will join and, consequently, no one joins. To avoid these coordination problems, assume the platform allows everyone to join for free and they do. The platform then offers them their on platform match at the equilibrium price, and they can either accept it or reject it and get a random off platform draw. Then all agents will join, and those whose individual rationality constraint (IR) is satisfied will accept their on platform match, knowing that, since the model is symmetric between sides A and B, their match's IR is also satisfied.

Our first goal is to find the cutoff between joining the platform and searching off platform. Define a consumer's type vector $x = (x_1, x_2, ..., x_n)$ and Define $x_{\neg 1} = (x_2, ..., x_n)$. Define the cutoff function $x_1 = f(x_{\neg 1})$ for side A (The cutoff for B will be 3-$f(x_{\neg 1})$). Given the set up, we know that on platform (full information) matches will be reflections along the hyperplane–that is, that they will have the same type along the 2nd to nth dimensions, and will be mirrored along the first. Thus the on platform utility for an agent on the cutoff with type $(f(x_{\neg 1}), x_{\neg 1})$ will be $2f(x_{\neg 1}) - 3 - p$.[6] Off platform, an A agent of type $(f(x^*_{\neg 1}), x^*_{\neg 1})$ gets a random draw with expected disutility (negative distance)

$$-\int_0^{\frac{1}{n-1}}\int_0^{\frac{1}{n-1}}...\int_0^{\frac{1}{n-1}}(n-1)^{n-1}(\int_2^3|x_1 - f(x^*_{\neg 1})| + \sum_{i=2}^{n}|x_i - x^*_i|)dx_1dx_2...dx_n$$

Evaluating the on and off platform utility and equating them (see appendix 8.2), we have

$$2f(x^*_{\neg 1}) - 3 - p = -3 + f(x^*_{\neg 1}) - \sum_{i=2}^{n}((n-1)x^*_i - 1)x^*_i$$

$$f(x^*_{\neg 1}) = p + \sum_{i=2}^{n}(1-(n-1)x^*_i)x^*_i$$

As derived in the appendix, the optimal price is 5/12 and the cutoff is

$$f(x^*_{\neg 1}) = 5/12 + \sum_{i=2}^{n}(1-(n-1)x^*_i)x^*_i$$

when the firm maximizes profit. Notice that the maximum and minimum values the cutoff takes stay within the bounds of [0,1] when $n \geq 2$, so we can simply integrate over the entire support along the $n-1$ horizontal dimensions without the complication that the cutoff goes outside the support on the vertical dimension. We keep the vertical dimension support fixed at $[0,1]$ to avoid this problem–if the vertical dimension also scaled down, we'd run into situations where the cutoff is below or above $[0, \frac{1}{n-1}]$ for some values of $x_{\neg 1}$, and addressing that complication in an arbitrary number of dimensions is too difficult. We can justify this by assuming $n-1$ vertical dimensions with the additional assumption that all agents agree on a preference ordering over $n-1$-vectors of vertical characteristics, regardless of the horizontal characteristics. We can then aggregate them into a single vertical dimension as discussed in section 2.3. This vertical dimension will then have support $[0, (n-1)\frac{1}{n-1}] = [0,1]$, as we've assumed.

[5] A model where all agents of a given type can join the platform, leaving off platform types as the only potential matches off platform, is interesting, but finding the cutoff becomes intractable.

[6] We could add a benefit term to make this utility positive, but it would add nothing to the analysis as we won't be considering the choice to remain unmatched

FIGURE 6.1. Two-dimensional example with risk aversion and income preferences–1-Norm

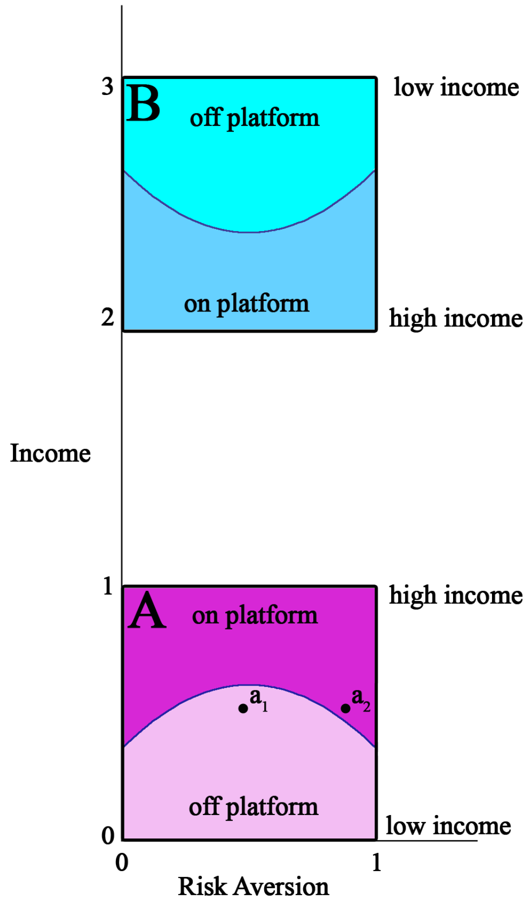

6.2. **Interpretation.** We can derive a number of implications from this analysis. First, closer agents along the separating dimension will join the platform while further agents will find a match off the platform, as in a one dimensional vertical preference model. However, we can see that, in the n-dimensional model, agents closer to the middle of the distribution along the $n-1$ horizontal dimensions will have a higher cutoff in the vertical dimension–that is, the middle agents will find the platform less attractive relative to their outside option than agents near the edges of the distribution along the n-1 horizontal dimensions, and thus will require a better match on the platform to induce them to join.

The reason for this is quite simple: for A agents on the platform, the expected utility of a match increases in ones own vertical characteristic by twice the increase in that characteristic. That is, $\dfrac{\partial E[-d(x,\mu_P(x))]}{\partial \mathrm{x}_1} = 2$, and there is no change in the quality of ones match with changes in own horizontal type. That is, $\dfrac{\partial E[-d(x,\mu_P(x))]}{\partial \mathrm{x}_i} = 0$, $i > 1$. Off the platform, agents get a random match, so expected utility increases only 1 for 1 with own vertical characteristic. That is, $\dfrac{\partial E[-d(x,\mu_O(x))]}{\partial \mathrm{x}_1} = 1$, and, $i > 1$. Also, agents with horizontal types further from the center of the distribution are further away from potential matches on the other side, on average, so their off platform utility from a random draw is lower.

FIGURE 6.2. Two-dimensional example with risk aversion and income preferences-Euclidean distance

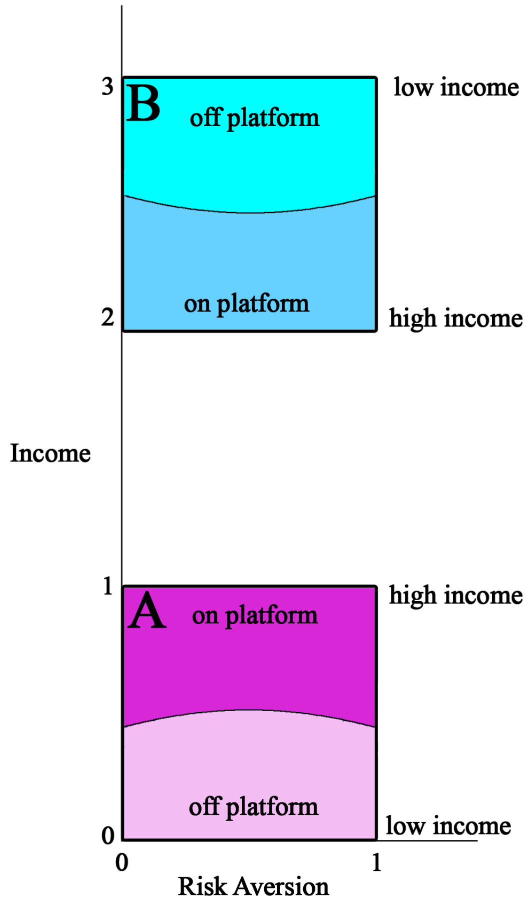

It may be helpful here to consider a concrete example. Suppose agents have horizontal preferences over risk aversion, because, once matched, pairs will make many joint decisions, and even decisions they make unilaterally will tend to have externalities for their partner. Thus, agents will prefer matches with risk aversion closer to them, so that they won't have to choose suboptimal options in compromise. For example, a compulsive gambler and a highly risk averse individual will likely have many conflicts over how to spend money.

Suppose agents also have vertical preferences over income. As in figure 6.1, suppose we've mapped risk aversion to [0,1] and income to [0,1] for A and [2,3] for B. Then the cutoff regions will be parabolas facing away from each other, where higher income individuals join the platform to get the good match they're guaranteed, rather than a random draw, and agents with high or low risk aversion tend to prefer the platform. For example, agent a_1 in figure 6.1 is willing to take a chance with a random draw, since her income is only moderate, not guaranteeing her a high income match, and her risk aversion is moderate, so that her expected distance from a random draw is not too great. Agent a_2, on the other hand, is very risk averse, so while her income is the same, the expected difference between her risk aversion and a random draw is greater, thus the certainty of getting a perfect risk aversion match on the platform is more attractive to her and she joins the platform.

The qualitative structure we see in the 2 dimensional case is preserved in higher dimensions, with higher dimensional paraboloids describing the cutoff and agents finding the platform more attractive the further they are away from the center along each horizontal dimension. We find that, as the number of characteristics increases while average distance between agents is held constant, price, quantity joining, and firm profits all stay constant. This is a positive result for the application of aggregation models to pricing questions in matching markets, as it suggests that, if our interest is solely in prices, profits, and quantities, rather than thye matching function itself, we can aggregate without changing the salient characteristics of the model. However, this result is likely to be particular to the 1-norm, and norms that are very different from the 1-norm may respond differently to aggregation. Also, in markets where the cutoffs go outside the support of the distributions, a complication intentially avoided in this exercise, then aggregation may change optimal prices, etc.

As noted earlier, we've been looking at a model using the 1-norm. This is fine if that is in fact the preference structure consumers have, and if the distributions exhibit symmetry about a hyperplane normal to one of the characteristic dimensions. However, the hyperplane will generally not have this nice property, so we need the Euclidean norm and its rotational symmetry to get the matching result. While we can't solve this model for an arbitrary number of dimensions, we can check a two-dimensional example to see if the matching structure is qualitatively similar. In fact it is, as we can see in figure 6.2, with almost (but not quite) parabolic cutoffs. Consumers with horizontal types near the edges find the platform more attractive for the same reason as in the 1-norm case, though the effect is attenuated with the Euclidean distance because the large vertical separation of the distributions makes the horizontal type relatively less important in determining distance. If the separating hyperplane isn't normal to any one characteristic, we can expect to see the same structure of matching and the same sort of cutoff, but the cutoff will now be rotated as well, and thus its relationship to any particular characteristic will depend on how it was rotated.

We might also see qualitatively similar results from higher dimensional Euclidean distance cases, except that, in order to make the mean distance between agents invariant as we increase dimensions, we'll need to shrink the supports by a factor of $\frac{1}{\sqrt{n-1}}$, rather than $\frac{1}{n-1}$, due to the square root of sums structure of the Euclidean distance.

We can also relate these results to the literature on search models of matching. Above, we've explored a model where agents get perfect information, no frictions matches on platform and get random matches off platform. We can interpret this as the limit of a search model where agents have a slow search technology as their outside option and a much faster search technology is offered by the platform. More concretely, agents draw potential matches each period and if both accept they will match. On the platform, draws come very quickly from the set of other agents on the platform, so that there is little time discounting between periods. Off platform, however, draws come from the whole distribution and agents get draws much more slowly, so there is significant time discounting between periods. Taking the limit of this model where time between draws goes to infinity off platform and to 0 on platform, we have an outside option of getting a single random draw (after which you'd have to wait forever to get another), and on platform the matching outcomes will converge to the Gale-Shapley stable matching, as shown by Adachi[1]. Thus, while solving a multidimensional search model as outlined above directly would be intractable, for sufficiently low discount factor off platform, and sufficiently high discount factor on platform, the search model should approximate–and have the same qualitative characteristics as–the extremely simple model derived in this section.

7. Conclusions

The results we see strongly support the symmetric result's applicability to modestly asymmetric markets. While, as expected, there is some attenuation of the anticipated same-characteristic effects on matches, the coefficient is relatively close to 1 or -1 even in small matching markets of a few hundred, and for the symmetric distribution case the predicted effects explain almost all of the variation in matching outcomes. Even for the asymmetric case, nearly all of the negatively assortative characteristic matching variation and most of the positively assortative characteristic matching variation are explained by the theoretically predicted same-characteristic effects.

There are many avenues by which to extend this research. The most obvious work yet to be done on the empirical side is Monte Carlo simulations to derive confidence intervals for the match characteristic coefficients and the relative weights. This paper only considers one pair of asymmetric distributions, but there is no guarantee that the same results will hold for different types of distributions with different types of asymmetries, so evaluating a variety of distribution pairs could help confirm or reject the results in general. Also, while this paper considered asymmetric deviations from the symmetric case, another assumption which may not hold in practice is the shared preference weightings. It would be worthwhile to test the robustness of these results to significant deviations from shared relative preferences. It would also be enlightening to see whether the results hold in higher dimension models. Finally, applying the techniques of this paper to real world data is the clear next step for this research program, and this will be the ultimate arbiter of the result's usefulness.

8. Appendix

8.1. **Uniqueness of Symmetric stable matching with finitely many agents.** While we've proven in Section 3 that the symmetric matching outcome is stable, we haven't proven that it is unique. While the following proof technique doesn't 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 4. *(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) isn't distance minimal, a contradiction. Suppose$d(a, b) \leq d(a, b')$ 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) isn't 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 characteristic 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.[10] □

8.2. **Theoretical Model Derivations.** Off platform, an A agent of type $(f(x^*_{\neg 1}), x^*_{\neg 1})$ gets a random draw with expected disutility (negative distance) $-\int_0^{\frac{1}{n-1}}\int_0^{\frac{1}{n-1}} \ldots \int_0^{\frac{1}{n-1}}\int_2^3 (n-1)^{n-1}(|x_1 - f(x^*_{\neg 1})| + \sum_{i=2}^n |x_i - x^*_i|)dx_1 dx_2 \ldots dx_n$ Since the 1-norm is additively separable, we can rewrite this as

$$(-\int_2^3 |x_1 - f(x^*_{\neg 1})|dx_1 - \sum_{i=2}^n \int_0^{\frac{1}{n-1}} (n-1)|x_i - x^*_i|dx_i)/\int_2^3 (n-1)^{n-1} f(x^*_{\neg 1})dx_1$$

$$= -5/2 + f(x^*_{\neg 1}) - \sum_{i=2}^n \frac{1}{2(n-1)} - x^*_i(1-(n-1)x^*_i)$$

$$= -5/2 - \frac{1}{2} + f(x^*_{\neg 1}) - \sum_{i=2}^n ((n-1)x^*_i - 1)x^*_i$$

Then at the cutoff we have

$$2f(x^*_{\neg 1}) - 3 - p = -3 + f(x^*_{\neg 1}) - \sum_{i=2}^n ((n-1)x^*_i - 1)x^*_i$$

so $f(x^*_{\neg 1}) = p + \sum_{i=2}^n (1-(n-1)x^*_i)x^*_i$. Note that the cutoff region will take the form of an n-dimensional paraboloid, with the center region having the highest cutoff and cutoffs dropping as one moves away from the center. We can now compute the quantity of agents that enter the platform by finding the proportion of A that is above the cutoff. However, it may be that the cutoff is below the support of A. We we see, however, that for the model specified that won't be a problem. The quantity joining the platform is then

$$Q(p,n) = \int_0^{\frac{1}{n-1}} \ldots \int_0^{\frac{1}{n-1}} (n-1)^{n-1}(1 - f(x_{\neg 1}))dx_2 \ldots dx_n$$

$$= \int_0^{\frac{1}{n-1}} \ldots \int_0^{\frac{1}{n-1}} (n-1)^{n-1}(1 - (p + \sum_{i=2}^n (1-(n-1)x^*_i)x^*_i)dx_2 \ldots dx_n$$

$$= 5/6 - p$$

It is now easy to compute the firm's profit, $\Pi(p,n) = p(5/6 - p)$, and the optimal price is 5/12. The cutoff is $f(x^*_{\neg 1}) = 5/12 + \sum_{i=2}^n (1-(n-1)x^*_i)x^*_i$. Note that, for all values of x^*, the cutoff remains between 0 and 1.

8.3. **The Gale-Shapley or Deferred Acceptance Algorithm.** The Gale-Shapley Algorithm is a widely used iterative algorithm for solving matching problems where there are no search frictions or limitations to information, where utility is nontransferable, and where agents on each of two sides of the market can match only to agents of the other side, for example men and women. Given a complete preference ordering over potential matches for each agent, it is guaranteed to give a stable matching that is optimal for one side of the market amongst the set of stable matchings. The algorithm proceeds as follows, this formulation due to Roth:

- Step 0. If some preferences are not strict, break ties arbitrarily, assign one side of the market as "men" and the other as "women"
- Step 1
  - a. Each man proposes to his most desired match, if any matches are acceptable.
  - b. Each woman rejects all but her most preferred proposal, if any are acceptable.
- Step k
  - a. Any man rejected at step $k-1$ makes a new proposal to its most preferred acceptable mate who hasn't yet rejected him. (If no acceptable choices remain, he makes no proposal.)
  - b. Each woman rejects all but her most preferred proposal to date.
- STOP: when no further proposals are made, and match each woman to the man (if any) whose proposal she is holding.

References

[1] Adachi, Hiroyuki. "A search model of two-sided matching under nontransferable utility." Journal of Economic Theory 113, no. 2 (2003): 182-198.

[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): 141-168.

[5] Clark, S. (2003), "Matching and Sorting with Horizontal Heterogeneity", ESE Discussion Paper 98,

[6] Clark, S. (2006). The uniqueness of stable matchings. Contributions in Theoretical Economics, 6(1), 1-28.

[7] Clark, S. (2007). Matching and sorting when like attracts like.

[8] 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

[9] "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

[10] "Two-Sided Matching with Spatially Differentiated Agents," Klumpp, Tilman, Journal of Mathematical Economics 45, 376-390, 2009.