ESSAYS ON MICROECONOMICS

Sam Flanders

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics.

Chapel Hill
2016

Approved by:

Gary Biglaiser

Fei Li

Peter Norman

Ju Hyun Kim

Brian McManus

Ju Hyun Kim

Atila Abdulkadiroglu

ABSTRACT

SAM FLANDERS: Essays on Microeconomics
(Under the direction of Gary Biglaiser.)

This dissertation consists of three essays on applied microeconomic theory, focusing on matching markets. The first two essays focus on finding closed form solutions to matching problems. Both generalize Gary Becker's well-known assortative matching results to more general environments. Becker studied a frictionless one-to-one matching environment with vertical or quality-based univariate preferences and found that, the higher one's type, the higher the type of their match will be, an implication that has been extremely influential in empirical work. The first paper generalizes this analysis to an environment where agents care about many traits instead of just one and have more general preferences, rather than restricting attention to vertical preferences. The latter generalizes it to a univariate environment where agents can have any ideal type and preference is decoupled from type, such that different agents of the same type can have different preferences. Both papers provide closed-form matching functions and make empirical predictions about the structure of matching. Theoretical and empirical interest in matching is currently shifting to richer settings where agents have varied preferences and must make tradeoffs between various traits they care about, so foundational work on the qualitative structure of sorting in these settings is necessary to provide intuition for researchers and to direct empirical questions, and closed-form matching functions can make theoretical models with embedded matching problems tractable.

In the third essay, I study a search model of online dating with nontransferable utility where agents are vertically differentiated, self-report quality, and must go on costly dates to verify a match's quality. We show that these per-date costs induce some agents to over-report their type, consistent with the stylized facts of online dating platforms where users frequently over-report characteristics like height and income, a phenomenon known as catfishing. This make agents less picky by preventing high types from rejecting some low types, and since externalities in matching markets without transfers can make agents inefficiently picky, these

costs can improve total market surplus. A monopolist platform owner may also have an incentive to increase per-date costs in order to increase profits. Thus, inducing lying amongst users can actually be optimal for a platform.

To my parents.

ACKNOWLEDGMENTS

There are countless individuals who have helped me along the way, most of whom will not be mentioned by name here. Needless to say, I am grateful to every single person who has aided me. Thank you from the bottom of my heart.

I would like to first and foremost thank my advisor Gary Biglaiser, whose help and guidance throughout the years have been invaluable. I could not have asked for a better mentor. I would also like to thank my dissertation committee members, Peter Norman, Fei Li, Ju Hyun Kim, Brian McManus, and Atila Abdulkadiroglu, for their tutelage. They have aided me from start to finish, from drafts to presentations, and from ideas to realizations. Thank you also to all of the members of the UNC Microeconomic Theory Workshop, namely Matthew Horne, Yiyi Liu, and Ben Labe, for their insightful comments that have helped shape the essays in this dissertation. I would also like to thank the Royster Society of Fellows for supporting my final year of research and the UNC Childcare Scholarship for enabling the graduate student parents of a feisty little boy, Erhan, to actually get work done during the day.

Last but not least, I would like to thank my family and friends. Thank you to my parents for always supporting me, and my son Erhan for giving me a reason to push forward. Finally, the biggest thanks of all goes to my amazing and long suffering wife, Melati Nungsari. Without her incredible love, effort, and sacrifice I could never have made it through this.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

CATFISH: LYING IN MATCHING MARKETS WITH CHEAP TALK

1.1 Introduction

This paper studies matching markets with "catfish", a neologism for someone who attempts to attract matches on an online dating platform by lying about themselves. Online dating has become a common component of dating and partnership formation, and is a fast growing market, taking in $1.08 billion in revenue for dating sites and $572 million for dating apps in 2014.[1] Lying is an important factor in search on these platforms; the distributions of reported types for traits like income and male height tend to be shifted right on online dating platforms relative to the broader population, suggesting misreporting,[2] and one industry study found that 20% of women and 33% of men admit to lying on their online dating profiles.[3] This report also offers advice on how to account for such misreporting, advising women to "assume the men you meet might not be quite as tall, as successful or as connected as they say they are, and then decide whether you'd still consider dating them regardless", suggesting that misreporting is an important factor in agent search strategies. We study a stylized model motivated by this issue, with a two sided (men and women) platform allowing agents to search for matches on the other side[4] where agents can misreport their type.

Specifically, we model a one-to-one nontransferable utility (NTU) matching market with

---

[1]http://www.wsj.com/articles/the-dating-business-love-on-the-rocks-1433980637

[2]http://blog.okcupid.com/index.php/the-biggest-lies-in-online-dating/

[3]http://www.ayi.com/dating-blog/ayi-top-online-dating-profile-lies/

[4]This analysis can easily be extended to a same-sex dating market.

random search[5] and time discounting where agents pursue long term (permanent) matches. Agents are vertically differentiated,[6] are distributed over a continuum of types, and self-report type. When agents meet, they see one another's reports and choose whether to enter a type verification phase (date) or to part and return to search. If they go on a date, they see each other's true type. They then decide whether or not to match permanently.

Going on a date is costly, allowing some agents to profitably lie. Before a date, agents weigh the expected payoff from continued searching against the benefit of matching today less the cost of the date, accepting a match if the latter is higher. However, after the date this per-date cost is sunk and drops from the agent's decision, making them less selective. That is, once you've already made the effort to meet and learn about someone you should be less picky, since if you start your search over again you'll have to make a costly investment in learning about the next person. Thus, some agents will be rejected for a date if they truthfully report their type, but, if they get a date by pooling with more attractive peers, they will be accepted after the date when standards are lower.

We focus on an equilibrium with a minimal amount of pooling, which we term the "Limited Pooling" equilibrium. In this equilibrium, agents self partition into "classes", where each agent only matches within their class. Classes will typically be larger–and thus agents less choosy–when per-date costs are higher, since these costs allow lower quality agents to pool with and match to higher types who would like to reject them ex-ante.

We study how a strategic platform charging a single fixed fee[7] can utilize per-date costs to improve firm profits or maximize surplus as a social planner. We first consider the case where per-date costs are prices (for example, a price to communicate with an individual you're interested in). Externalities endemic to NTU search markets make agents too picky. In particular, agents don't care about their match's payoff, or the payoff of other agents in the market, only their own. Without transfers, other agents have no way of making them

[5] Over time, agents receive random draws from the set of agents on the other side of the market.

[6] Each type is characterized by a level of quality, and all agents prefer higher quality matches to lower quality matches.

[7] We discuss a schedule of fixed fees in Appendix 1.6.1.

internalize these costs. Thus agents will chase after high quality matches, ignoring the fact that if they get a high quality match, someone else must get a low quality match, so their gain is another's loss. Thus the social planner will want agents to be less picky, since choosy behavior results in surplus loss due to time discounting and benefits the social planner much less than it does the individual. Thus, per-date costs can counter this excessive pickiness, and we find that a social planner will utilize positive per-date prices to make agents less picky and increase total surplus.

A profit seeking monopolist [8]platform will also charge positive per-date prices.[9] In this environment, the platform fully extracts surplus from the lowest quality agent who joins, while leaving rents to all higher class agents due to the lack of price discrimination. By forcing higher types to match to lower types they'd like to reject, per-date costs can be interpreted as inducing a transfer from high types to low types, which allows the platform to extract more surplus from agents by making the indifferent agent better off (and willing to pay more) at a cost to the rents of high type agents, which the platform does not value.

We also consider the case where per-date costs are frictions, and a platform has access to technology that can lower per-date frictions at a cost to the platform. Some online dating platforms in fact offer technologies such as video-chat and ID verification to lower informational frictions. Even in this case, we find that a monopolist platform often prefers high per-date costs, despite their great cost to agents. Per-date costs can again be interpreted as inducing a transfer from high types to low types, increasing the fixed fee the platform can charge. With appropriately chosen per-date costs, this benefit offsets the direct surplus losses due to per-date frictions. Thus, in this environment platforms often do not have an incentive to induce truth-telling or informative reporting, and in fact may benefit from inducing lying if lowering frictions is even slightly costly. A social planner will generally prefer to lower

[8]A platform with significant market power may be a reasonable approximation of the online dating market as the market is fairly concentrated and the fact and much of the market outside of the two largest firms, IAC and eHarmony, is highly differentiated niche platforms like JDate and ChristianMingle.

[9]Explicit per-date prices are currently uncommon in dating platforms, though some platforms, like Ashley Madison and It's Just Lunch, have utilized them. More common are contracts that limit the number of contacts one can make in a given period of time and charge a premium for unlimited contacts, which has qualitatively similar effects but is less tractable to model.

per-date frictions to zero if the cost to the platform of this technology is sufficiently low, but there are special cases where even a social planner will prefer higher per-date frictions.

In the broader set of equilibria, the ability to freely report type allows for many forms of pooling, including highly non-monotonic reporting where, for example, agents in a high type class and a low type class make the same report while agents in an intermediate quality class make reports that uniquely identify their class. However, we find that, regardless of reporting strategies, equilibria are characterized by a class partition where, after going on a date, agents will only match to draws whose true type is within their class. Hence, this paper contributes to the literature by extending the coarse, class based form of positively assortative matching (PAM)[10] found previously in NTU search models like Macnamara and Collins (1990) and Burdett and Coles (1997) to an environment with cheap talk and costly type verification.

It also contributes to the literature on two-sided platforms–specifically, the literature on strategic matching platforms. In particular, we show how informational frictions can be profitably used by a strategic platform to counteract externalities in matching markets. A recent survey of the search and matching literature by Chade, Eeckhout, and Smith (2015) identified the nature and implications of externalities in matching markets as one of the major open questions in the field, so this analysis addresses an important hole in the literature.

We'll now describe some additional salient features of the online dating market. It features significant concentration, with eHarmony taking in \$310 million in revenue in 2014 and platforms owned by the IAC (Match, Tinder, OkCupid, etc.) taking in \$601 million, almost all of which was from dating websites, where total dating site revenue was \$1.08 billion. [11]. Membership is slightly less concentrated, with IAC platforms serving 21% of users in 2014 and eHarmony serving 13% [12]. The remainder of the market is composed of small platforms, many of which are niche dating sites, and recent dating app entries. Platforms typically either

[10] Higher types match with higher types, lower types match with lower types.

[11] http://www.washingtonpost.com/news/business/wp/2015/04/06/online-datings-age-wars-inside-tinder-and-eharmonys-fight-for-our-love-lives/

[12] http://www.wsj.com/articles/the-dating-business-love-on-the-rocks-1433980637

offer a single service at a positive price, or a free, ad-supported service and a premium paid service with additional amenities, such as unlimited messaging, better search options, and video chat. Many platforms engage in second degree price discrimination, offering significant discounts for longer contracts. Typical contract lengths are 1, 3, 6, or 12 months. Few major platforms engage in overt demographic based price discrimination, although Tinder prices based on age.

While this paper focuses on an application to online dating markets, matching markets with search and costly type verification–and thus an incentive for lower quality agents to pool with higher quality agents–appear in a variety of contexts, notably in job search. While the NTU assumption is less palatable for job search, as wages are often bargained over and can be set flexibly by the employer, there are also often limits to the degree of transferability. Among other things, firms may have their wage setting abilities constrained by regulations like minimum wages, and progressive taxation makes larger transfers more costly. Hence, the efficiency and profitability of positive per-date frictions we find assuming NTU may extend–perhaps with some attenuation–to partially transferable utility (PTU) environments that may more credibly model applications like job search. Additionally, in some job search applications like the market for medical residents, wage offers are extremely compressed due to the market structure, making NTU a reasonable assumption. While the National Resident Matching Program ends with a one-shot assignment game based on rankings reported by each side, hospitals and students meet in time consuming and costly interviews before reporting their preferences, suggesting a search model like the present may capture some stylized characteristics of such markets.

The remainder of this paper is organized as follows: Section 1.2 discusses the related literature and this paper's place within it. Section 1.3 lays out the basic theoretical framework this paper uses, characterizes the set of equilibria in this environment, and provides equilibrium selection arguments. Section 1.4 incorporates a strategic platform that can change the magnitude of the per-date costs. Section 1.5 concludes. Section 1.6 is the appendix, which includes proofs for many of the propositions in the paper, as well as analysis of the model with alternative assumptions, such as different matching technologies.

1.2 Related Literature

This paper follows a rich literature on search and matching. In particular, it fits within the literature on search and matching with NTU. McNamara and Collins (1990) first studied the NTU search environment with a continuum of types and found the distinctive partition or class equilibrium common in this literature, where agents in each class only match within their class. Burdett and Coles (1997) extended this analysis to a steady state environment with exogenous inflows of agents and endogenous outflows. The present paper is closely related to this strand of the literature. In particular, the constant returns to matching model can be interpreted as Burdett and Coles with an additional reporting stage. Eeckhout (1999) extends this result to multiplicatively separable preferences, and Chade (2001) extends it to fixed search costs. Smith (2006) looks at even more general preferences, situating the partition result in a larger class of equilibria where partitioning does not necessarily hold. There is a parallel literature for the transferable utility (TU) assumption, with Shimer and Smith (2000) studying equilibria in the analogous TU environment. Generally, TU makes characterizing equilibria, payoffs and agent behavior more difficult.

There is also a small literature on strategic matching platforms with search. Bloch and Ryder (2000) study a monopolist platform environment that can offer frictionless NTU matches for a fixed fee or a fixed proportion of match surplus, with an outside option of NTU search. Damiano and Li (2007) study vertically differentiated agents and a monopolist that creates a continuum of platforms and sets prices to induce agents to join their assigned platform and match with identical agents. Given the simplicity of observed contracts in this market and the potential frictional costs of partitioning agents into many small platforms, our paper instead considers what a less ambitious platform can do in an environment where draw rates proportional to the mass of agents on a platform make infinite partition of the space into measure zero platforms inefficient. Damiano and Li (2008) study competition between matching markets.

This paper also relates to the literature on cheap talk and information transmission. The cheap talk literature was pioneered in Crawford and Sobel (1982). Applications of cheap talk, signaling, and information transmission to matching markets include Hoppe et al. (2009) and

Hopkins (2012), who study matching with signaling. Bilancini and Boncinelli (2013) address a similar NTU environment but assume only one side has unobservable types and consider a binary choice between type certification and full information matching and hiding one's type and matching randomly. The present paper differs from these works by focusing on a cheap talk environment. Ko and Konishi (2010) study a profit-maximizing platform matching firms and workers in a many-to-one environment, where firms and workers report match-specific wage offers and desired wages, respectively. They find that manipulating reporting by curtailing the message space can improve profits. A current working paper, Hagenbach et al. (2015), studies an environment very similar to ours, with an initial reporting stage and a costly type verification stage before permanent matching in a search environment with vertical differentiation and NTU. However, they consider a two point type distribution, while our analysis focuses on the class structure that only appears non-degenerately in a model with a continuum of types. We also study strategic platforms, while they focus on a nonstrategic platform.

This paper also relates to the literature on two-sided platforms, pioneered in Rochet and Tirole (2003, 2006) and Armstrong (2006). More recent work includes Weyl (2010), Bedre-Defolie and Calvano (2013), and Lee (2013). In contrast to the majority of this literature, which takes advantage of simple, exogenous specifications for network externalities, this paper explicitly models the special case of network effects induced by a search model of matching.

1.3 Model

1.3.1 Preliminaries

We model a heterosexual market on an online dating platform and denote the two sides men (m) and women (w). Agents are characterized by a single vertical characteristic representing quality or attractiveness, where every agent strictly prefers higher quality matches. Quality for side j is distributed over $[\underline{q}_j, \overline{q}_j]$, $\underline{q}_j > 0$.[13] When agents join the platform, the platform

[13] If $\underline{q} = 0$ there can be infinitely many classes, which poses difficulties for certain aspects of the analysis.

solicits a report on their true type $\hat{q} \in [\underline{q}_j, \overline{q}_j]$, representing their online dating profile. On the platform, agents engage in bilateral search for partners in continuous time, with a discount rate of $r \in (0, \infty)$. Agents receive random draws from from the endogenous distribution of agents on the other side $G_j(q,t)$ according to a Poisson process with arrival rate α, where G_j is continuous[14]

When they meet, each agent observes the other's report, and they make an ex-ante decision whether or not to propose a date. If both propose a date, they pay a per-date cost,[15] learn each other's true type, and make an ex-post decision whether to propose a match. If both propose a match they marry forever. If either rejects in the first stage they part costlessly, while if either rejects after the date they part having paid the per date cost. Inflow into the platform is exogenous and time invariant, with the cumulative distribution given by $F_j(q)$, where F_j is twice differentiable and has full support on $[\underline{q}_j, \overline{q}_j]$. The corresponding density is given by $f_j(q)$. The total mass of inflow is equal for both sides and normalized to 1. Outflow is determined endogenously by the rate of acceptances and the mass on the platform. The mass of agents on the platform is given by N.

Total match surplus is given by a function $u(q_m, q_w) \equiv \psi_m(q_m)\phi_m(q_w) + \psi_w(q_w)\phi_w(q_m)$, where each agent's payoff is multiplicatively separable into an own type component ψ and a match's type component ϕ. Both are weakly increasing twice differentiable positive functions, with ϕ is strictly increasing. We assume non-transferable utility (NTU), where $u_w(q_m, q_w) \equiv \psi_w(q_w)\phi_w(q_m)$ and $u_m(q_m, q_w) \equiv \psi_m(q_m)\phi_m(q_w)$. This means that agents cannot bargain over the apportionment of surplus, perhaps due to social norms, which may be plausible in some matching markets such as dating markets. NTU, along with the multiplicative separability of each agent's own type and match type in their payoff, ensures the very simple and tractable class structure common to this literature. Per-date cost for an agent of type q_j on side j is $\psi_j(q_j)c$. This is a strong assumption, but it is also necessary to preserve the

[14]We'll prove this later.

[15]We assume agents cannot match sight-unseen. This seems consistent with most marriage/partnership formation, where some amount of quality verification precedes commitment. Even arranged marriages typically involve reconnaissance by relatives.

class structure of the equilibrium.[16] When $\psi = 1$, as in the commonly assumed case where utility is simply match's type, per-date costs can be thought of as either a price imposed by the platform or the opportunity or effort cost of going on a date. When $\psi \neq 1$, per date costs should be thought of as opportunity costs that are increasing in type. [17] Unless otherwise noted, agents have an outside option of zero, such that every possible match is preferable to remaining unmatched.

We'll focus on stationary equilibria where

- **Assumption 1 (STN)** : Each agent believes $G_j(q,t) = G_j(q)$.

Further assume stationary agent strategies, where a strategy is an agent's ex-ante date decision for each reported type and ex-post match decision for each true type. Let $\mu_i(\hat{q})$ be agent i's belief about the distribution of q given a report $\hat{q}$. Following Burdett and Coles (1997) and extending the definition to a game of incomplete information, we utilize the following definition:

Definition 1 *Given (G_m, G_w), a Bayesian perfect partial equilibrium (BPPE) is a strategy profile and beliefs μ where STN is satisfied, agents maximize utility subject to their belief about other agents' types and actions and follow sequential rationality, and beliefs are consistent with Bayes rule wherever possible.*[18]

This definition identifies a set of candidates for a steady state equilibrium, which we will later winnow down by requiring that inflows equal outflows.

We can now establish several useful properties of the agents' strategies and payoffs. First, we'll establish that agents follow cutoff strategies:

---

[16]We'll relax this assumption in numerical simulations in Appendix 1.6.1.

[17]Note that, if $\phi_w = \phi_m$, we can define a new trait $x' = \phi(x)$ and find a distribution H such that $F(\phi^{-1}(x')) = H(x')$. Thus, when we assume symmetric payoffs it will suffice to consider $\phi(x) = x$. However, when we make distributional assumptions we must note that they will change when types are mapped back into the original distribution.

[18]Also assume agents reject ex-post matches when indifferent and accept ex-post matches they strictly prefer to continued search when they believe the probability of beign accepted is zero.

Lemma 1 *In a BPPE each agent x accepts all draws above some cutoff* $\underline{q}_x$ *ex-post and rejects all draws below. Each agent x accepts all draws with expected discounted match quality above a cutoff ex-ante and rejects all draws below. Each agent accepts a strictly positive measure of ex-post matches in equilibrium.*

Proof. After true type is revealed, an agent chooses between a continuation value independent of current draw and the value of the draw. Accepting is costless, so they must accept if the utility of the draw exceeds the continuation value and reject otherwise. Before true type is revealed, agents choose between a continuation value independent of the current report and the expected discounted match quality associated with that report.

Accepting a zero measure of dates or matches yields a continuation value of zero. Accepting any agent yields a strictly positive payoff ex-post, thus agents must have a cutoff below the maximal type on the other side. ■

Let $U_w(q|\hat{q})$ denote woman $q's$ expected discounted lifetime utility when reporting $\hat{q}$. $q_{wE}(q|\hat{q}) \equiv U_w(q|\hat{q})/\psi_w(q)$ is then the expected discounted match quality. Given symmetric definitions for men, we can easily show that higher type agents get matches of weakly better discounted quality (and thus higher utility) and that agents with higher quality matches must have weakly higher types:

Lemma 2 *If an agent is of type* $x > x'$, $q_{jE}(x|\hat{q}_x) \geq q_{jE}(x'|\hat{q}_{x'})$. *If* $q_{jE}(x|\hat{q}_x) > q_{jE}(x'|\hat{q}_{x'})$, $x > x'$.

Proof. Given Lemma 1, any agent that will accept x' ex-post must accept x and x can always mimic the x' strategy. Thus an x agent can always obtain at least as high expected match quality as an x' agent. ■

We will make one of two assumptions about the rate of draws agents face. Specifically, we'll assume they face *linear returns to matching* (LRM) in the main body of the paper, and consider the case of *constant returns to matching* (CRM)

[19] in Appendix 1.6.2.

[19]Note that LRM is sometimes referred to as a quadratic search technology, owing to the quadratic nature of the total number of draws in the market as a function of the number of agents, and CRM is sometimes referred to as a linear search technology based on the linear rate of total draws in the market as a function of the total number of agents.

- **Assumption 2A (LRM)** : Agents receive a rate of draws α proportional to the mass of agents on the platform, normalized to N.

- **Assumption 2B (CRM)** : Agents receive a constant rate of draws α, normalized to 1.

Linear returns to matching means that the frequency of draws is proportional to the mass of agents on the platform and that thick markets make search faster. Most past work on search with NTU has focused on the constant returns to matching environment, and this may be more appropriate for traditional forms of partner search where finding potential matches is time consuming and these frictions put an upper bound on the number of draws an agent can consider, regardless of the size of the market. However, on an online dating platform, we'll argue that linear returns may be more realistic. With easy search, filtering, and detailed information available with a single click, it's plausible that more agents on the platform means more draws, since one may quickly exhaust a small list of potential matches by paring it down to a handful of likely matches.

In fact, the linear returns environment is significantly more tractable than the constant returns one: since there are no per-draw costs in this model, having to eliminate more agents outside of your acceptance region imposes no cost, and thus changes in the mass of agents outside your class has no effect on your optimization problem. This can be motivated by the nearly costless filtering out of undesired matches that may be achieved on a search platform. With CRM, by contrast, more agents in other classes means it will take longer to get a draw from your class, making behavior in each class dependent on behavior in every other class.

Define $\lambda \equiv Pr[match|q,\hat{q}]$ and $\gamma \equiv Pr[date|q,\hat{q}]$. Given that the continuation value of a woman of type q reporting $\hat{q}$ (and the case for men is symmetric) is their lifetime expected utility, $U_w(q|\hat{q})$, and is also equal to the ex-post cutoff draw $q_{wl}(q,\hat{q})$, the dynamic program for this environment gives us the following optimization condition ex-post for a small time period dt:

$$U_w(q|\hat{q}) = \frac{U_w(q|\hat{q})(1-\lambda\alpha dt) + \lambda\alpha dt E[\phi_w(q')\psi_w(q)|q,\hat{q},match] - \gamma\alpha dt\psi_w(q)c}{(1+rdt)} \tag{1.3.1}$$

Where c does not appear on the left-hand side (LHS) because the per-date cost has already been paid when the ex-post match decision is being made. Taking the limit as $dt \to 0$, we have

$$U_w(q|\hat{q}) = \frac{(\alpha\lambda E[\phi_w(q')|q, \hat{q}, match] - \alpha\gamma c)\psi_w(q)}{\lambda\alpha + r} \tag{1.3.2}$$

and applying the equality between continuation and cutoff acceptance utility,

$$\phi_w(q_{wl}(q,\hat{q}))\psi_w(q) = \frac{(\alpha\lambda E[\phi_w(q')|q, \hat{q}, match] - \alpha\gamma c)\psi_w(q)}{\lambda\alpha + r} \tag{1.3.3}$$

Notice that this specification for the cutoff means that the way one's own type enters the utility function doesn't affect agent behavior, and thus only matters when one considers welfare or adds prices to the model.

Because the continuation value is the same for both the ex-post and ex-ante decisions, but ex-ante the per-date cost is not sunk, the ex-ante cutoff is simply $\phi_w^{-1}(\phi_w(q_{wl}(q,\hat{q})) + c)$ if type is certain or expected discounted match quality equal to the same.

Lemma 3 *Ex-post cutoffs and optimal strategies are independent of ψ_j.*

Proof. Direct inspection of 1.3.3. ■

1.3.2 The Set of Equilibria

Ex-Post matching structure

We can now analyze ex-post matching behavior. While the structure of reporting can be quite complex, ex-post matching behavior is simple and highly consistent with previous research on NTU search models with observable types. In particular, Proposition 1 shows that agents will partition themselves into classes in equilibrium. First, we'll formally define the terminology:

Definition 2 *We'll call an interval of types a* class *if every agent with a type in that class accepts any type from that interval ex-post and only forms ex-post matches with types within the class. Define the lower bound of a class n on side j $q_j(n)$.*

Note that this is a condition on the second stage where type has been revealed. Agents in a given class may reject agents within their class ex-ante, and accept dates from agents outside their class depending on the reporting structure.

Proposition 1 *The distribution of agents on each side is partitioned by intervals (or classes) of agents where, for each class n, men(women) in class n will accept any woman(man) in their corresponding class n ex-post, will reject any woman(man) below class n, and will be rejected by any woman(man) above class n.*

Proof. Consider a $\overline{q}_m$ man. $\overline{q}_m$ men are accepted by every woman ex-post and must accept women above a cutoff strictly below $\overline{q}_w$ by Lemma 1. Thus, there is a nontrivial interval over which every $\overline{q}_m$ man accepts matches ex-post. Denote the lower bound $q_w(1)$ for the highest type man's cutoff type (in the distribution of women) and $q_m(1)$ for women. Then by Lemmas 1 and 2, every agent must accept this interval ex-post as they have lower types. Since women will accept men above $q_m(1)$ ex-post, a man above that type can mimic any man's strategy. Thus, every man above $q_m(1)$ will get the same payoff in expectation and thus the same cutoff $q_w(1)$. We'll call the interval of women $(q_w(1), \overline{q}_w]$ class 1 of side w and denote the nth class class n. A symmetric analysis yields $(q_m(1), \overline{q}_m]$, class 1 of side m. We can proceed inductively from here. Given that every woman above $q_w(n)$ rejects any man at or below $q_m(n)$, $q_m(n)$ type agents face a problem analogous to $\overline{q}_m$ men, and accept every woman in an interval whose lower limit is defined as $q_w(n+1)$. Similarly, every lower type man accepts all women in $(q_w(n+1), q_w(n)]$. Thus every woman in class $n+1$ is accepted by the same set of men and must have the same payoff in expectation and thus the same cutoff $q_m(n+1)$. A symmetric analysis shows that every man in $(q_m(n+1), q_m(n)]$ only accepts women above $q_w(n+1)$.

■

Figure 1.3.1 shows the class structure of an equilibrium. Agents in each class only match to agents in the same class on the opposite side. An agent in class 3, for example will accept anyone in a classes 1, 2, or 3 ex-post, and will be rejected ex-post by every agent in classes 1 and 2. Note that the length of classes can vary based on density and the class cutoffs need not be symmetric if the distributions are not. This class structure ensures that agents in

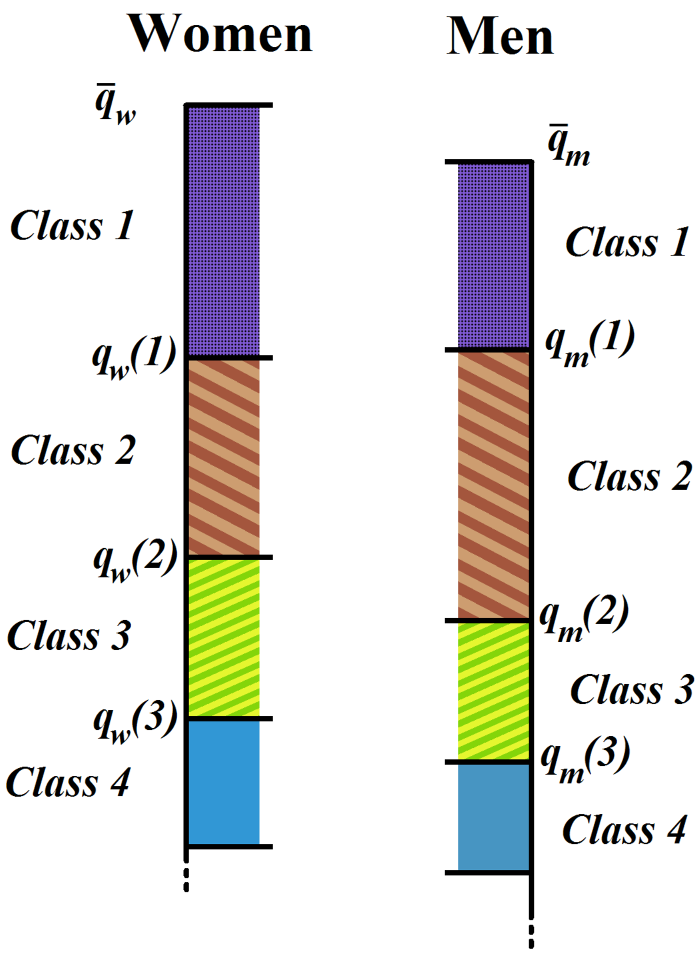

Figure 1.3.1: The class structure of an equilibrium.

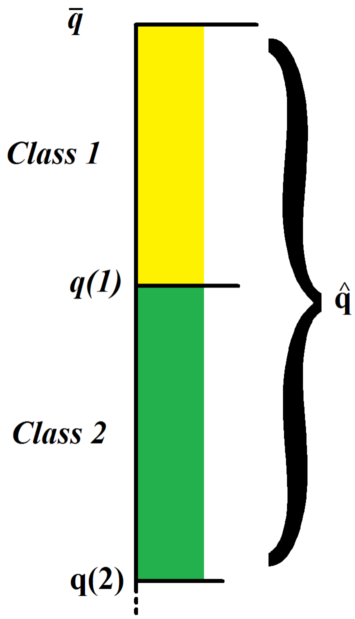

Figure 1.3.2: This diagram shows inter-class pooling, as in Proposition 2.

each class must have a higher payoff than agents receive in any lower class.

Lemma 4 *Suppose type q is in class m and type q' is in class n. If $m>n$, $q_{jE}(q) > q_{jE}(q')$.*

Proof. Suppose not. Then the q and q' agents receive the same payoff and thus must have the same ex-post cutoff. But then they are in the same class. Contradiction. ■

Reporting Structure

We can now address the reporting stage of the game. While the ex-post class structure was simple, there are a wide range of possible reporting structures, including pooling over multiple classes. Because of the bilateral nature of the signaling, it's important to note the possibility of asymmetric reporting where, for example, every woman pools and every man reports their class. The redundancy of the two sided reporting and acceptance decisions means that, so long as one side reports their class and accepts everyone ex-ante, the other side can simply accept only the the class they will match to ex-post.

Definition 3 *On a given side, we'll call an interval a* contiguous pool *if there is a report $\hat{q}$ such that, for every type in the interval, some agent of that type reports $\hat{q}$ or another report* *that is accepted by the same set of agents ex-ante.*

Pools, then, are defined based on reporting and the ex-ante dating stage. This is in contrast to classes, which are defined based on ex-post acceptance decisions. The following proposition shows that, for any any contiguous set of classes n through $n+k$, an equilibrium exists where all agents in each of these classes pool on a single report $\hat{q}$. That is, any pooling structure that nests contiguous sets of classes within pools can be supported in equilibrium. Figure 1.3.2 shows an example with two classes pooling on a single report: all agents in classes 1 and 2 make the same report $\hat{q}$. Note that the cutoffs for these classes are endogenous in the reporting structure.

Proposition 2 *Any contiguous pool $(q_j(n+k), q_j(n-1)]$ can be supported in a BPPE.*

Proof. Suppose that, for each side j, all agents in classes n through $n+k$ pool on a single report $\hat{q}_j$, no agent outside $(q_j(n+k), q_j(n-1)]$ reports $\hat{q}_j$ and every agent rejects any report made only by agents outside of their class. Then if a pooling agent makes any report other than $\hat{q}_j$ they will never receive a match, yielding a non-positive payoff. Thus there is no profitable deviation for pooling agents. If each other class i forms a pool where all agents report $\hat{q}_{ij}$, this is an equilibrium, so a BPPE exists. ■

We will now provide a lemma that establishes the range of characteristics pools can have in equilibrium. In particular, it shows that pool cutoffs can coincide with class cutoffs as above or appear within classes, with the possibility of multiple pools within a class and mutual rejection by agents within the same class but in different pools, even though they would like to match ex-post.

Lemma 5 *For any contiguous pool in a BPPE, on at least one side of the market, one of the following holds for the lower(upper) bound of the pool:*

i) The lower(upper) bound is also the lower(upper) cutoff of the lowest(highest) class with reports in the pool.

ii) the non-endogenous pool cutoff induces indifference between reporting within the pool and giving any report given by agents in the lowest(highest) pooling class but outside the pool.

Proof. Without loss of generality, consider the lower bound case. Suppose i) and ii) are violated and the pool is consistent with equilibrium. Then there is a report made in equilibrium by an agent in the lowest pooling class with a different payoff than some other report made by an agent outside pool but in the class. Then agents in that class have a strict incentive to give the higher payoff report and the assumed reporting is not optimal. Contradiction. ■

The above lemma could be presented more tersely by giving a more general form of ii., but this formulation provides more intuition about the range of possible reporting strategies. It suggests several different sorts of equilibria, and we will describe examples for both cases. If reports satisfy i) for both bounds, classes are nested within pools and agents can't profitably leave the pool by reporting above, where they'll get rejected, or reporting below, where they'll get a lower payoff. If reports satisfy ii) for both bounds, agents will reject one another ex-ante even though they are in the same class. For example, suppose a report $\hat{q}_m$ is made by some men in class n and some outside of it, and $\hat{q}_w$ is made by some women in class n and some outside. If men reporting $\hat{q}_m$ always reject women reporting $\hat{q}_w$ and vice versa, deviating to accepting may yield a costly date with no match and will never yield a match, so it is strictly better to reject. This can be supported if the payoff for each report on a given side in class n is equal, e.g. half of n men of each type in the class report $\hat{q}_{m1}$ and only accept $\hat{q}_{w1}$ and the other half report $\hat{q}_{m2}$ and only accept $\hat{q}_{w2}$, and women behave symmetrically. Note that reports need not satisfy the same case for both the top and the bottom bounds– the upper bound could satisfy i) while the lower bound satisfies ii)

This can analysis can trivially be extended to pools that are discontiguous but can be represented as a finite union of contiguous pools. In fact, we can quite easily find equilibria where, for example, classes 1 and 3 pool on a single report despite rejecting one another ex-post while class 2 agents make a report that uniquely identifies their class.

Corollary 1 *Suppose the support of a pool can be expressed as the union of a finite set of intervals. Then each interval must satisfy Lemma 5.*

Thus we see that the structure of reporting can be highly non-monotonic, and agents within a given class can even reject each other ex-ante. However, the underlying ex-post matching structure retains the coarsely assortative class structure found in previous research in this environment with observable types.

The Limited Pooling Equilibrium

Having established some characteristics of the set of possible equilibria, we will now introduce the equilibrium[20] of interest for the remainder of this paper. Unlike many other equilibria in this environment, this equilibrium looks very similar to those found in Macnamara and Collins (1990) and Burdett and Coles (1997), with agents only dating within their class. The primary difference from equilibria with observable types is that there is a region between the ex-ante and ex-post (class) cutoffs for each class that, with observable types, would be rejected by the agents in the class. However, since per-date costs induce a lower ex-post cutoff, agents in this interval can pool with those above them and get accepted due to the laxer ex-post standards.

Definition 4 (Limited Pooling Partial Equilibrium (LPPE)) *We'll call a BPPE a*n LPPE *if each agent makes a report only made by agents in their class and accepts every report made by agents in their class.*[21]

We can now easily characterize the equilibrium pooling structure and the relationship between the ex-ante and ex-post acceptance decisions:

Lemma 6 *(1) In any LPPE, agents between the ex-ante and ex-post cutoffs must pool with agents above the ex-ante cutoff such that the expected quality of that report exceeds the ex-ante cutoff.*

(2) In any LPPE, agents will always accept after a date.

Proof. Consider an agent $m(w)$ in class n. Suppose a report $\hat{q}$ is never made by women(men) above the ex-ante cutoff in class n. Then $m(w)$'s continuation value is higher than the expected payoff of dating and matching to the $\hat{q}$ woman(man) and he must reject. Thus women(men) below the ex-ante cutoff must pool with agents above it to gain acceptance in their class. It is immediate that any agent with a type below the ex-ante cutoff must pool

[20] Formally, a set of equilibria with reporting strategies that lead to equivalent payoffs for all agents.

[21] Note that there is a larger set of equivalent equilibria where every agent accepts every agent in their class ex-ante and only goes on dates with agents in their class, but where, for example, men make informative reports and accept all matches and women make uninformative reports and are selective. It is without loss to consider the special case of Definition 4.

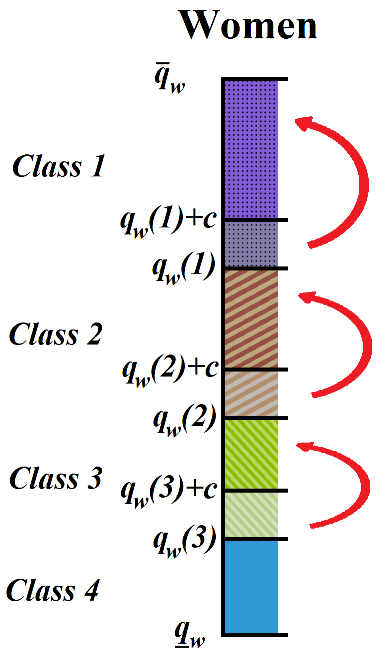

Figure 1.3.3: The pooling structure of an LPPE.

with an agent of a type above the ex-ante cutoff for the average quality of an agent in that pool to exceed the ex-ante cutoff. Finally, since agents only accept dates within their class, they always accept a match after a date.

■

Figure 1.3.3 shows the pooling structure of an LPPE. For women(men) in each class, agents the region from $q_w(n)$+c to $q_w(n)$ are below the ex-ante class cutoff but above the ex-post class cutoff, and thus will be rejected ex-ante if they reveal their true type, but will be accepted ex-post. Thus, they must pool with agents above the ex-ante cutoff (classes will always have length greater than c, so this is always possible). Note that, in this example, $\underline{q}_w > q_w(4) + c$, so the fourth class ends before the ex-ante cutoff. While the first J-1 classes must end with the class's endogenous ex-post cutoff, the last class J may end with the lower limit of the support of G, which may be above $q_w(J)$ or even above $q_w(J) + c$.

Lemma 6 shows that agents between the two cutoffs must pool in order to maximize their payoffs. Comparing an LPPE to the larger set of equilibria, we see two primary ways other equilibria differ: we can find equilibria where multiple classes pool reports and we have between-class date acceptance, and we can find equilibria where agents in a given class reject dates with others within that class due to coordination failures.

The lack of a complex reporting structure independent of ex-post classes make LPPE very tractable, and thus very attractive. There are reasons beyond convenience to focus on LPPE, though. Given linear returns to matching (LRM) and a steady-state assumption these equilibria must exhibit the finest class partition in the set of steady-state equilibria in terms of the number of classes.[22]

Additionally, in any other equilibrium some agents must pool and accept dates they will later reject, or reject dates they'd like to accept. This means that every such agent would strictly prefer to leave the pool and report their true class, and they are only prevented from doing so by the absence of messages that can reveal their class.[23] In particular, any man(woman) who goes on dates outside their class or rejects dates inside their class would prefer to unilaterally reveal their true class and commit to accepting agents in their class ex-post, since this would allow them to match to every good draw and to reject every bad date. This does not imply a simple equilibrium dominance argument, unfortunately–everyone switching to truthful reporting of class creates higher expected payoffs, which induces higher class cutoffs. Thus, agents between the old and new cutoffs could be made worse off by being bumped from class n to class $n+1$.

However, given LRM, consider a non-LPPE equilibrium. There must be some class where men(women) reject dates within the class or accept dates outside the class. For the first such class, n, class n men(women) who would remain in the class under the limited pooling class cutoff would all strictly prefer to form a coalition and report their true class, if that were possible. Importantly, if they did, no man(woman) outside their class would have an incentive to re-pool with their new report, since they would simply be rejected ex-post, and any coalition attempting to repool would necessarily include agents who would be worse off under the coalitional deviation. Thus, non-LPPE seem fragile in the sense that all agents will deviate from the multi-class pools that define them if they can find a way to report their

[22]Depending on the parameterization, there may be other equilibria with the same number of classes, but none will have fewer.

[23]We'll proceed informally here. Formally modeling this game with coalitions, repeated reports, and deviations that induce off equilibrium path beliefs and play outside of steady-state would be intractable.

true type, and when agents deviate from these pools there is no incentive for the agents in their former pool to unilaterally follow them to a new report.

Could this same argument also exclude the LPPE, where pooling also occurs? We'll provide a heuristic argument to the contrary. In an LPPE, high type agents in a class would like to reveal their true types and thus escape pooling with agents between the ex-ante and ex-post cutoffs. However, all of these low type agents strictly prefer to pool with higher type agents in the class. If we assume that agents can change reports at intervals and coalitional deviations are significantly more costly than unilateral deviations due to coordination costs, we'd expect to generally see pooling as in the LPPE since it will be too costly for high types to repeatedly coordinate on new reports only to have low types unilaterally follow them. The finest partition and equilibrium preference arguments are formalized in Section 1.3.3.

We can now write down the explicit form of the agents optimization problem. We'll focus on the case where the distributions of men and women and their utility functions are symmetric for tractability. Define the number of classes as J, the proportion of agents in class n $\lambda_n \equiv G(q(n-1)) - G(q(n))$, and define q_l and q_u as the lower and upper cutoffs for a class, respectively. Given that agents accept any agent in their class and reject all others ex-ante, every date results in a match and the probability of accepting a draw is λ_n. Then equation 1.3.3 can be rewritten as

$$q_l = \frac{\alpha\lambda}{\lambda\alpha + r}\int_{q_l}^{q_u}(x-c)\frac{g(x)}{\lambda}dx \tag{1.3.4}$$

Rearranging and applying integration by parts, we have:

$$q_l = \frac{\alpha}{r}(\int_{q_l}^{q_u} G(q_u) - G(x)dx - \lambda c) \tag{1.3.5}$$

We can now characterize the class structure explicitly:

Proposition 3 *Given G, a LPPE implies sequence of cutoffs for men and women* $\{q(n)\}_{n=0}^{J}$ *satisfying* $q(n) = \frac{\alpha}{r}(\int_{q(n)}^{q(n-1)} G(q(n-1)) - G(x)dx - \lambda_n c)$, where $q(0) = \overline{q}$, *and* $q(J) <= \underline{q}$.

Proof. The first and third claims follow directly from Proposition 1 and equation 1.3.5, and if the fourth were not true there would be another class $J+1$. ■

It also follows that agents accepting measure zero masses of agents outside their class or rejecting measure zero masses of agents inside their class generates the same cutoffs and payoffs.

1.3.3 Steady State in the Limited Pooling Equilibrium.

Linear Returns to Matching

The linear returns to matching equilibrium analysis closely follows Burdett and Coles (1997). Some proofs go through nearly unchanged, but others must be amended to account for per-date costs and the differing assumptions on returns to matching. Define the distribution of agents leaving the platform as $H(q)$ and the mass of agents leaving the platform by O. We can now define our complete equilibrium concept by combining the partial equilibrium of the LPPE, which ensures all behavior and beliefs are rational and assumes steady-state, with a balanced flow condition that ensures steady-state holds by equating the endogenous outflows with the exogenous inflows, closing the model.

Definition 5 *(LPPE Steady-State Equilibrium (LSSE)): given exogenous inflows* (F)*, a steady state equilibrium is pair* (G, N) *satisfying LPPE and balanced flow: for every interval* $[q_1, q_2) \in [\underline{q}, \overline{q}]$*,* $O(H(q_2) - H(q_1)) = F(q_2) - F(q_1)$*.*

The cutoff equation is now

$$q(n) = \frac{N}{r}(\int_{q(n)}^{q(n-1)} G(q(n-1)) - G(x)dx - \lambda_n c) \tag{1.3.6}$$

Within a given class, we can get a simple characterization of outflow. Outflow in a class is given by the number of agents on the platform, N, times the proportion of agents in the class, λ_n, times the rate of draws of an agent in that class $N\lambda_n$. Then outflow from class n is $\lambda_n^2 N^2$. Then, in an LSSE,

$$\lambda_n = \sqrt{(F(q(n-1)) - F(q(n)))}/N \tag{1.3.7}$$

We also have that, for any $[q_1, q_2)$ in class n, $\lambda_n(G(q_2) - G(q_1))N^2 = F(q_2) - F(q_1)$ and thus, with the differentiability of F,

$$g(q) = \frac{f(q)}{\lambda_n N^2} \tag{1.3.8}$$

Thus the density of agents on the platform in a given class is inflow density times a scalar. Combining (1.3.6) and balanced flow, we can get eliminate G terms, yielding class cutoffs solely in terms of inflows and c.

$$q(n) = \frac{1}{r}(\int_{q(n)}^{q(n-1)} \frac{F(q(n-1)) - F(x)}{F(q(n-1)) - F(q(n))} dx - c)\sqrt{(F(q(n-1)) - F(q(n)))} \tag{1.3.9}$$

Note that in the linear returns environment, the N's cancel out, and we have cutoffs that depend on the previous cutoff. We can now explicitly characterize the LSSE in this environment:

Proposition 4 *Given F, then (G, N) defines a LSSE if and only if G satisfies (1.3.8) and λ_n satisfies (1.3.6), (1.3.7), $q(0) = \overline{q}$, $q(J) \leq \underline{q}$, and $\sum_n \lambda_n = 1$.*

Proof. $\sum_n \lambda_n = 1$, the boundary conditions, and (1.3.6)-(1.3.8) are necessary in an LSSE by construction. Conversely, the assumptions guarantee $G(\overline{q}) = 1$, $G(\underline{q}) = 0$ and G increasing, so G is a well defined steady state distribution and any G and N satisfying them form a valid LSSE. ■

Thus, equilibrium requires that each class cutoff is the solution to the agents' optimal stopping problem and the density on the platform is consistent with balanced flow. $\sum_n \lambda_n = 1$ ensures that g is well defined.

To ensure existence of an LSSE, we'll need to make some distributional assumptions. An increasing hazard function will ensure that the class structure is unique. Note that, while Burdett and Coles need this assumption to deal with a multiplicity of cutoffs due to N, that channel is shut down in the linear returns environment. However, this assumption also constrains multiplicity induced by per-date costs, so it is still necessary in this environment.

- **Assumption 3 (HAZ)** : The hazard function $f(q)/(1 - F(q))$ is increasing in q.

Proposition 5 *Given F, the partition satisfying (1.3.6)-(1.3.8) and the boundary conditions is unique.*

Proof. See Appendix 1.6.2. ■

The intuition for this proposition is that the RHS of (1.3.9) is decreasing in $q(n)$, while the LHS $q(n)$ is obviously increasing, yielding a single crossing. The RHS can be interpreted as the expected surplus quality of an accepted match over the cutoff quality, multiplied by the probability of acceptance. Generally, we'd expect this to be decreasing in cutoff type $q(n)$, since a higher cutoff lowers the surplus over cutoff for any given draw, and, were G exogenous, a higher cutoff would lower the probability of accepting a draw. However, due to the endogenous nature of G, it's possible for the density to rise as $q(n)$ increases, swamping the aforementioned effects. The hazard rate assumption ensures that the density can't rise too fast, excluding this possibility.

In this environment, existence and uniqueness of the equilibrium follow directly. The class cutoffs are unique and these cutoffs imply a unique steady-state mass on the platform, N. This in turn ensures a unique density $g(q) = \frac{f(q)}{\sqrt{(F(q(n-1))-F(q(n)))N}}$ and thus a unique distribution on the platform G.

Proposition 6 *A unique LSSE exists.*

Proof. The class summation condition and (1.3.7) yield

$$N = \sum_n \sqrt{(F(q(n-1)) - F(q(n)))} \tag{1.3.10}$$

and N does not enter the cutoff equation so $q(n)$ is not a function of N and uniqueness is ensured. Existence is similarly direct. ■

We can also show that the cutoffs are decreasing in c–that is, increasing per-date costs generally makes classes coarser, and for sufficiently high c the class structure completely unravels as every man(woman) accepts every woman(man) due to the high costs of continuing their search and paying the per-date cost again.

Proposition 7 *Given LRM, $q(n)$ is decreasing and continuous in c in an LSSE for all n and $q(n) \to 0$ for c sufficiently high.*

Proof. Suppose c increases. Consider the first endogenous cutoff, $q(1)$. The RHS of (1.3.9) is decreasing in c, and the RHS is decreasing in $q(1)$ by Lemma 11, so the lower RHS from c

must be compensated with a lower $q(1)$ for equality to hold. We can now proceed inductively. Suppose $q(n-1)$ is decreasing in c. Lemma 11 also shows that the RHS increasing in $q(n-1)$, so to maintain equality, $q(n)$ must increase in $q(n-1)$. Additionally the argument in the base case ensures that, fixing $q(n-1)$, $q(n)$ is decreasing in c. Thus, $q(n)$ is increasing in the argument that decreases, $q(n-1)$, and decreasing in the argument that increases, c, and must decrease on net. Direct inspection shows continuity given the continuity of F. Lemma 11 shows that, for sufficiently high per-date costs, the RHS goes to zero and thus the cutoff goes to zero. ∎

The intuition for this result is that per-date costs lower expected match utility, and in the optimal stopping problem cutoff utility must be equal to expected match utility, so higher per-date costs should yield lower cutoffs.

We're now ready to formalize the justifications for our equilibrium selection. Lemma 7 shows that LPPE classes with the same upper bound generate higher payoffs for agents within the class and thus that the classes are smaller. Lemma 8 shows that the top agents in a pool prefer revealing their class and matching to one another to remaining in the pool. Corollary 2 shows that an LPPE induces the finest partition of the type-space in terms of classes. Recall that $q(n-1)$ is the upper bound of a class n and define $q(LPPE, q(n-1))$ as the cutoff induced by $q(n-1)$ if everyone in the class accepts one another ex-ante and no one else, as in an LPPE and n_{LPPE} as the corresponding class starting at $q(n-1)$ with no cross-class pooling or within-class rejection and where balanced flow is satisfied. Also define $n_{LPPE,q(n)}$ as the class with upper bound $q(n-1)$, lower bound $q(n)$, limited pooling and balanced flow and define $q(LPPE, n)$ as the nth cutoff given an LPPE. Define M_n as the mass of agents in class n.

Lemma 7 *Given LRM, a class n starting at $q(n-1)$ where the probability of an agent in n accepting a date outside the class or rejecting a date inside the class is strictly positive must have a cutoff $q(n) < q(LPPE, q(n-1))$. Additionally, $q(n) < q(LPPE, n)$.*

Proof. See Appendix 1.6.2. ∎

Lemma 8 *Given LRM, consider an equilibrium where the probability of an agent in a class*

n accepting a date outside the class or rejecting a date inside the class is strictly positive. Further, define n as the first class where this is true. Then agents in class n above $q(LPPE, q(n-1))$ all strictly prefer an equilibrium with class n_{LPPE} to the one with class n, and if they could coordinate to reveal their true class and accept only others in their class, no agent outside their class would have an incentive to pool with them.

Proof. A class with a lower cutoff of $q(n)$ must have an expected match quality $q_E(n) = q(n)$ and one with $q(LPPE, q(n-1))$ must have $q_E(n_{LPPE}) = q(LPPE, q(n-1))$. By Lemma 7, $q(n) < q(LPPE, q(n-1))$, so agents must prefer the n_{LPPE}. Additionally, any agent outside the class that pools with them will be rejected ex-post, and thus has no incentive to pool. ■

Corollary 2 *In any LSSE with LRM, the nth class cutoff is maximal in the set of steady-state equilibria, and the number of classes is also maximal.*

Proof. By Lemma 7, $q(n) < q(LPPE, n)$ and if $q(k)$ exceeds the lower bound of the distribution, $q(LPPE, k)$ must as well. ■

1.4 Strategic Platforms

1.4.1 Per-Date Costs as Frictions

Monopolist Platform

Up until now, we've taken per-date costs as given, but a strategic platform such as a social planner or profit maximizing monopolist may be able to influence them, either by increasing them via per-date pricing, or decreasing them by providing easy ways for agents to communicate and verify type (e.g. video chat) or verifying certain aspects of an agent's report. We'll first consider the frictional case with a monopolist platform. Specifically, we'll consider a platform that charges a fixed fee for both sides of the market, can provide its service costlessly, faces an exogenous per-date friction $\bar{c}$, and can decrease the per-date cost to c at a cost $\tau(c)$, where $\tau(\bar{c}) = 0$ and τ is strictly decreasing in c. Thus the firm will charge a fixed fee p and every agent above a cutoff $q(p, c)$ will join the platform, yielding profit flow rate $\Pi_{fric}(p, c) \equiv p(F(\bar{q}) - F(q(p, c))) - \tau(c)$, and p will equal the expected payoff of the lowest joining agent. Formally, amend the game to include a first stage where the platform chooses p and c and consider the equilibrium where the maximal mass of agents join the platform. As before, agents have an outside option of zero[24]. Define $q_n(c)$ as the nth endogenous cutoff given c. Define $p^*(c)$ as the optimal price given a per-date cost c. Then it will generally be optimal for the platform to set a price such that the lowest type joining the platform is also the cutoff type for the last class joining the platform. Specifically,

Lemma 9 *Suppose LRM and $\bar{c}$ sufficiently high. Either:*

i) a positive c is optimal for the monopolist platform, or

ii) choosing a price that yields a cutoff $q(p, c)$ that does not coincide with the lowest joining class' endogenous cutoff is suboptimal for the platform if $q(p, c) < q_1(0)$.

Proof. Suppose the optimal contract yields a cutoff $q(p, c)$ such that $q(p, c) > q_1(c)$ and $q(p, c) < q_1(0)$. Then a positive c is optimal. Now suppose the platform induces an individual rationality (IR) cutoff $q(p, c) < q_1(c)$, and $q(p, c) \neq q_n(c)$ for any n. For sufficiently high $\bar{c}$, there must be a c' such that $\bar{c} \geq c' > c$ and $q(p, c) = q_1(c')$ since q_1 is continuous, decreasing in c, and goes to zero as c increases. Note that the cutoff type is equal to discounted expected utility. Thus, any class with the same lower cutoff yields the same expected match quality for agents in that class. Then if the firm chooses c' and a p' to induce the same cutoff, the quantity of agents on the platform is identical, but the cutoff agent is willing to pay $\phi(q(p, c))\psi(q(p, c))$, while under the original regime the cutoff agent is willing to pay $\phi(q_n(c))\psi(q(p, c))$. $q(p, c) > q_n(c)$ implies $p' > p$, and since cutoffs are decreasing in c, $\tau(c') < \tau(c)$. Thus the firm will increase profit by inducing $q(p', c') = q_1(c')$. ■

We can use this to result to show that the platform generally will not have an incentive to lower per-date costs to zero:

---

[24] This can be relaxed throughout Section 4. For example, the outside option can easily be re-specified as a time discounted random draw from the distribution of agents off the platform (large platform) or the overall distribution of agents in the market (small platform). Assuming that the search technology (draw rate) off platform is sufficiently slow, agents will optimally accept any draw off platform, rationalizing these specifications. All the results of Section 4 go through with these endogenous outside options, though the firm's optimal prices will change. If the search technology off platform isn't slow, high types may have better outside options than low types, limiting the extent to which a platform can profitably raise per-date prices or allow per-date frictions. Generally, the greater the efficiency advantage of the platform relative to the outside option, the more flexibility the platform will have to support high per-date costs and manipulate user behavior.

Proposition 8 *Suppose LRM. If $\overline{c}$ is sufficiently high and $q(p^*(0), 0) < q_1(0)$, the monopolist platform never has an incentive to lower per-date costs to zero.*

Proof. $q(p^*(0), 0) < q_1(0)$, so since $q(1, c) \to 0$ as c increases and q(1,c) is continuous in c, the intermediate value theorem (IVT) ensures that there will be a $c^* > 0$ such that $q(1, c^*) = q(p^*(0), 0)$, $p^*(c^*) \geq p^*(0)$, and $\tau(c^*) < \tau(0)$. Thus profit will be higher with a per-date cost c^* than with a zero per-date cost. ■

Thus, even when per-date costs are frictions, a monopolist platform can profitably withhold higher efficiency search technologies, even when the cost of implementing such technologies is minimal. Lemma 9 shows that inducing endogenous class cutoffs equal to the IR cutoff is typically optimal, since that is the highest class cutoff that indifferent agent can have and the higher a cutoff is, the higher the utility agents in the class receive due to the equality of the cutoff and the continuation value. Then, given that prices and per-date costs are chosen to induce this coincidence between the platform and last class cutoff, higher per date costs have a direct effect of decreasing total surplus via effort spent on dates, but at the same time transfer surplus from high types to low types. Generally, the net effect of these countervailing forces would be ambiguous, but because the cutoff is held constant they must cancel out exactly, again due to the equality between cutoffs and continuation values. Thus, in terms of revenue, the platform is indifferent between any per-date costs that induce the appropriate cutoff, and strictly prefers higher per-date costs in terms of its own cost τ.

Social Planner

We can now consider analogous social planner's problem. Here, it is much less likely that it will be efficient to have positive per-date frictions. However, as was mentioned before and will be elaborated in the next section, there are externalities that can make agents too picky, and this can lead to inefficiently small classes and large utility losses due to discounting–agents spend too much time searching due to other agents' pickiness, and end up worse off. Also, the last class is qualitatively different then the preceding J-1 classes–it may not end at the endogenous cutoff, but rather at the bottom of the support of the distribution, which may be above the endogenous cutoff. This can lead to tiny rump classes with very low payoffs since there are very few agents in the class and thus the average search time is extremely high. In addition to directly lowering payoffs through frictions, c can thus speed up matching and change the size of this last rump class, potentially avoiding inefficiently sized final classes. Note that the flow of total surplus can be expressed as

$$TS \equiv \sum_{n=1}^{J} q_{jE}(n) \int_{q(n)}^{q(n-1)} \psi(q) f(q) dq \tag{1.4.1}$$

That is, for each class n, the rate of total surplus generation is the expected match utility for each agent integrated over the distribution of inflow in class n, where expected match quality for the class $q_{jE}(n)$ is constant across agents and can be pulled out of the integral. Total surplus is the sum over these classes. For ease of exposition, suppose for now that $\psi(q) = 1$. As shown in Appendix 1.6.2, we can express $\frac{\partial TS}{\partial c}$ as the sum of the effects of the change in each class cutoff $\frac{\partial q(n)}{\partial c}$ on the surpluses generated in the classes above and below it. In particular, the net effect of a change in cutoff $q(n)$ is proportional to

$$F(q(n-1)) - F(q(n)) - f(q(n)) \cdot (q(n) - q(n+1)) \tag{1.4.2}$$

as shown in Figure 1.4.1. If $F(q(n-1)) - F(q(n)) > f(q(n)) \cdot (q(n) - q(n+1))$ for all n (with a caveat for the last class discussed in the Appendix), TS decreases in c, and if the inequality is reversed the opposite holds. While the condition itself is simple, the class cutoffs are determined by a highly non-homogeneous recurrence relation, so finding conditions for either case is quite difficult in general. Below, we treat the case when utility is highly supermodular, which ensures Equation 1.4.2 is positive because the trade-off is between utility for class n, $F(q(n-1)) - F(q(n))$, and utility for class $n+1$, $f(q(n)) \cdot (q(n) - q(n+1))$. Supermodularity ensures that higher classes generate more surplus since they are populated by higher type agents, so for sufficient supermodularity the effect on class n dominates the effect on class n-1.

Proposition 9 *Suppose LRM,*$\psi(q) = q^{\alpha}$*, and* $\tau(c) = 0$*. If* α *is sufficiently high, total surplus is decreasing in c.*

Proof. See Appendix 1.6.2. ■

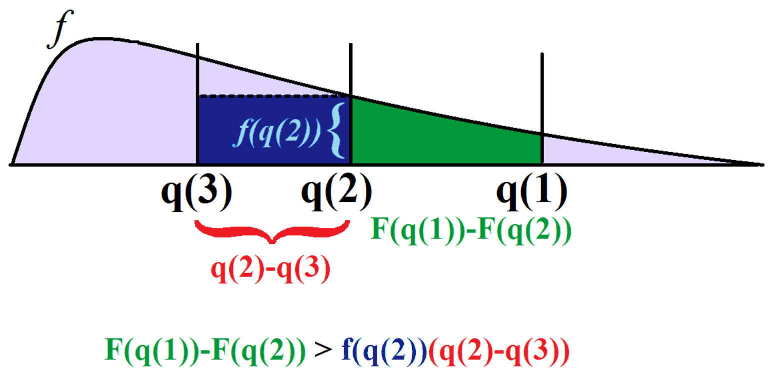

Figure 1.4.1: The trade-off induced by $q(2)$ shifting due to an increase in c.

Class 2 surplus decreases proportionally to its mass (in green) due to the class cutoff $q(2)$ shifting down and thus lowering expected match quality, which is equal to the class cutoff. Lowering $q(2)$ also bumps top class 3 agents up to class 2. The surplus increase induced by this is proportional to the change in payoff $q(2)-q(3)$ (red) multiplied by the density of agents at the cutoff $f(q(2))$ (blue). Geometrically, this is the blue rectangle in Class 3. If the former decrease (green area) exceeds the latter increase (blue area) for each class, increasing c will lower total surplus.

In the Appendix 1.6.2 we run numerical simulations for the modular utility case, and find that increasing c typically decreases total surplus, and must decrease it above a certain point (if c gets sufficiently high no one can get positive utility from joining the platform). The primary reason for this is that lower type classes tend to be smaller, since agents choose their reservation types based on a trade-off between quality and discounting, which causes a proportional decrease in match utility. Thus, agents who get high expected payoffs must be less selective, as waiting is more costly for them. Since $\frac{\partial TS}{\partial c}$ is negative when the surplus loss for higher classes of shifting a cutoff down outweighs the gain to lower classes, lower classes having less mass and thus generating less surplus makes $\frac{\partial TS}{\partial c} < 0$ likely. Equivalently, the direct effect of decreasing surplus due to frictions overwhelms any efficiency gains due to lower selectivity. However, at some values of c $\frac{\partial TS}{\partial c}$ can be increasing, due to the effect of c on time discounting and the rump class. This is especially common with decreasing distributions that yield more mass in lower classes. Decreasing distributions also put more weight on the rump class, creating periodicity in the surplus as the class cutoffs shift downward in c and the rump class goes from being a large, relatively efficient class to a small class whose size is limited by the support of the distribution, and then to a large class again as the last class cutoff passes the bottom of the support of the distribution and the next class becomes the last class. Thus, for friction reduction costs where small decreases in frictions are asymptotically costless like $\tau(c) = (\bar{c} - c)^2$, we can find cases where there is no incentive to decrease per-date frictions even a little bit, since total surplus is locally increasing in c.

1.4.2 Additively Separable Match Utility and Per-Date Prices

Monopolist Platform

We'll now consider the case of per-date costs as prices. We'll need to restrict our attention to the case when $\psi(q) = 1$ for tractability–platforms charging a per-date price proportional to agent type is inconsistent with the unobservable types assumption of this paper. Conversely, charging a fixed per-date price with supermodular match utility won't induce the class structure that makes this analysis tractable–higher type agents in any potential class will be less affected by the per-date cost than lower type agents in that class, and will thus have different cutoffs. We study numerical simulations for supermodular cases in Appendix 1.6.1. First consider the case of a monopolist platform. The platform charges a fixed fee p and a per-date price c. Consider the equilibrium where the maximal mass of agents join the platform. Since $\psi(q) = 1$, social surplus is modular–the total surplus is just the sum of the match payoffs $\phi(q)$ for each side, discounted by the expected time to match. Thus the structure of matching has no effect on payoffs, only the speed of assignment matters. Of course, individual agents benefit from matching to high types, but their benefit comes at a cost to other agents who don't get high type matches. Thus, externalities generally induce agents to be selective when matching, even though a social planner would assign every agent on the first draw. Thus, increasing per-date prices improves social surplus. It is also clear that, if a monopolist prefers positive per-date frictions, they will prefer positive per-date prices as well, which is the basis on which Corollary 3 extends the frictional results of the last subsection to the per-date price environment. Given profits under frictions of $\Pi_{fric}(p, c) \equiv p(F(\bar{q}) - F(q(p, c))) - \tau(c)$, profits under prices are $\Pi_{price}(p, c) \equiv (p + w(p, c)c)(F(\bar{q}) - F(q(p, c)))$, where w is the average of the expected discounting for each class weighted by the inflow rate of the class.

Corollary 3 *Suppose LRM. If $q(p^*(0), 0) < q_1(0)$, the monopolist platform never has an incentive to lower per-date prices to zero.*

Proof. Note that the proofs of Lemma 9 and Proposition 8 go through with the removal of $-\tau(c)$ in the profit function and the inclusion of $w(p, c)c(F(\bar{q}) - F(q(p, c)))$. ■

Social Planner

As discussed above, the social planner prefers agents to leave as quickly as possible to minimize time discounting, since assignment doesn't matter. We can prove this result very directly for constant returns to matching (CRM):

Proposition 10 *Suppose CRM. If c is a price, $\psi(q) = 1$, $f(x)$ is increasing or $xf(x)$ is increasing and c is sufficiently small, and there is more than one class when c is zero, a positive c maximizes social surplus.*

Proof. A sufficiently high c will ensure a single class, and we have multiple classes with zero per-date costs.$\psi(q) = 1$, so, given the inflow distribution F, social surplus is given by $SS = 2\int_{\underline{q}}^{\bar{q}} \phi(q)f(q)E[e^{-rt}|q]dq$. This is maximized when $E[e^{-rt}|q]$ is maximized for all q. A single class maximizes the exit rate, maximizing $E[e^{-rt}|q]$. Thus a single class maximizes social surplus and a positive c is necessary to induce a single class. Platforms can choose an appropriate fixed fee (possibly negative) to satisfy agent IR constraints. ■

In the LRM case, it's not necessarily true that the platform will want to maximize outflow for every agent type. Since the mass on the platform determines the rate of draws, having more agents on the platform can boost exit rate and thus discounted utility. If the mass of all agents is higher, this is a second order effect that must be overwhelmed by increased agent selectiveness in order to induce larger masses on the platform in the first place, but it is possible that the planner would benefit from inducing low types to reject other low types and only match to high types. High types give a higher payoff to their match, so time discounting is more costly for them. By contrast, a sufficiently low type agent contributes almost nothing no total surplus. Thus, inducing low types to reject other low types and increase their mass on the platform would allow high types to get quick matches and preserve their much more valuable contribution to match surplus. In this environment, the platform has no means to induce this partially negatively assortative matching[25] so inducing a single class will still be optimal given the instruments available, but we'll need to amend the proof to take into account the endogenous draw rate.

Proposition 11 *If c is a price, $\psi(q) = 1$, LRM holds, and there is more than one class when c is zero, a positive c maximizes social surplus.*

Proof. A sufficiently high c will ensure a single class, and we have multiple classes with zero per-date costs.$\psi(q) = 1$, so, given the inflow distribution F, mass on the platform will be 1 by (1.3.10) and every agent will leave upon a draw, which they get at rate 1. Suppose there is more than one class. Mass on the platform is $\sum_n \sqrt{(F(q(n-1)) - F(q(n)))}$ by (1.3.10), and the probability an agent in class n gets a draw they will accept is $\frac{\sqrt{(F(q(n-1))-F(q(n)))}}{\sum_n \sqrt{(F(q(n-1))-F(q(n)))}}$. Then the rate of accepted draws is $\sqrt{(F(q(n-1)) - F(q(n)))} < 1$ and all agents have a longer expected wait on the platform, getting lower utility in expectation. Platforms can choose an appropriate fixed fee (possibly negative) to satisfy agent IR constraints. ■

1.5 Conclusions

In this paper, we extended the NTU search literature to an environment with a cheap talk stage and costly type verification. We found that the partition or class based equilibria that have characterized this literature extend to this environment with informational frictions, with agents only matching to one another within their respective disjoint classes. When per-date costs are endogenously chosen by a strategic platform, positive per-date costs may be optimal, despite being distortionary, and, in the case of per-date frictions, having a negative direct effect on surplus. A social planner can take advantage of these per-date costs by using them to counter externalities that make agents too picky by allowing low type agents to pool with high type agents, preventing those high type agents from inefficiently rejecting them. A monopolist can use per-date costs to induce effective transfers from high to low types by forcing high types to match to low types, flattening the demand curve and allowing a monopolist that charges a fixed fee to extract more surplus from consumers.

[25] Negatively assortative matching occurs when higher types match to lower types rather than their own type.

Future avenues for study include analysis of more complex contracts in this environment to see how much profit the simple contracts commonly in use leave on the table, to see how more complex contracts interact with the externalities in these markets, and to more precisely capture the second degree price discrimination common to the menu of contracts in these markets. This is briefly studied in Appendix 1.6.1.

Studying an analogous model with transferable utility would be extremely useful, yielding results more applicable to job search, where cheap talk on both sides of the market can also be important. A two-type TU model is studied in Appendix 1.6.1, and we discuss how, under TU with linear returns to matching, agents are generally not picky enough since spending more time on the platform increases the mass of agents and thus increases the frequency of draws, a benefit agents do not internalize even with transfers known as the thick market externality. This is in contrast to the NTU case where agents are too picky, and thus yields opposite implications for per-date pricing, making negative per-date prices that incentivize an agent to stay on the platform longer optimal.

Including competition would be an obvious extension of this research program, though an even greater multiplicity of equilibria must be contended with due to the coordination issues with multiple platforms and network effects. Including exogenous exit and match dissolution would allow for more realistic modeling of matching behavior, especially on platforms that focus on short term matching like Tinder. This is not likely to substantively change the qualitative characteristics of the equilibrium, though.

1.6 Appendix

1.6.1 Extensions

More Complex Contracts

So far we've assumed a simple contract structure with a single fixed fee and per-date cost motivated by the observed simplicity of contracts in this market. However, in many cases this is not optimal. With frictions, it will generally be optimal for the monopolist to offer different contracts based on report. We will focus on a menu of fixed fees with a constant per-date cost. Optimal menus of per-date costs will be highly dependent on the distribution F, and aren't amenable to a simple analysis.

Given the IR cutoff induced by a single fixed fee contract, if it's possible to include multiple classes above that cutoff, the original IR cutoff can be maintained with additional higher price cutoffs for higher types, allowing more surplus extraction. Under modular utility it's optimal to set contract intervals coinciding with endogenous class cutoffs (everyone reporting in class n must pay p_n) for reasons analogous to those in Lemma 9. If utility is supermodular it could be possible to find cases where this is not true, due to higher utility for higher types, but we'll assume one and only one contract per class, noting that, for supermodular utility, this must be weakly worse than the optimal contract structure.

For the lowest class k, the IR binds as is standard, yielding $p_k = q(k)\psi(q(k))$ For class $k-1$, note that lower class agents cannot deviate to higher classes due to rejection, even if the contract is more favorable. Thus we only need to worry about deviations by higher type agents to lower reports. The incentive compatability (IC) constraint $IC_{k-1,k}$ requires $q(k-1)\psi(q(k-1)) - p_{k-1} \geq q(k)\psi(q(k-1)) - p_k$, based on the bottom agent $k-1$[26] so $(q(k-1) - q(k))\psi(q(k-1)) \geq p_{k-1} - p_k$ and so on, with each following price p_{n-1} rising based on the difference in class utilities $q(n-1) - q(n)$ scaled by the supermodular component $\psi(q(n-1))$ and with rents $q(n-1)(\psi(q) - \psi(q(n-1)))$ accruing to class $n-1$ agents based on ψ increasing in q over the class interval. This strategy maintains the same total surplus but allows the platform to extract more from users, and is a lower bound for maximal revenue with multiple contracts. Some rents are still left the to users if utility is supermodular.

In the modular case, however, the monopolist has full extraction and is essentially a social planner, and thus the earlier analysis of the social planner's problem in Section 1.4 applies: increasing per-date frictions is less costly in terms of revenue than the direct effect of frictions on utility would imply, just like the social planner case, and, generally, smaller τ's can rationalize high per-date costs than one woud expect based on the direct effect of frictions (in particular, a sufficient τ would be $\tau(c) = \overline{c} - c$), especially under certain distributional

[26] Higher type agents in the class benefit weakly more from higher quality matches, but pay the same price, so if the IC is satisfied for the lowest agent in must be satisfied for all others. This also makes showing local IC sufficient.

assumptions as discussed in the frictional social planner case. Generally, however, positive per-date frictions are undesirable. With per-date prices and modular utility, a single fixed fee is optimal since it's optimal for the social planner. With supermodular it may still be optimal, especially for low degrees of supermodularity, but it may not be.

Transferable Utility

While we don' treat a transferable utility model in the main body of this paper, it is of significant interest, since it may be more applicable to many job market applications, and some prefer the TU assumption in models of dating and marriage. In this environment, externalities resulting from socially inefficient acceptences and rejections are eliminated by transfers as in the Coase theorem, but externalities resulting from the effects acceptance and rejection have the on mass of agents on the platform and their distribution persist. Under LRM, there is no cost to having too many agents one does not want to match to on the platform (the congestion externality), but there is a cost to having fewer agents one does want to match to on the platform (the thick market externality). Thus, there is only one externality in play: staying on the platform longer benefits agents who would like to match to you and has no effect on agents who don't, so agents aren't picky enough because they don't interalize the benefits their presence has for others. This is the opposite of the net effect of externalities in the NTU case, and suggests that platforms ought to lower per-date costs as much as is feasible.

We'll illustrate this by studying a TU version of this paper's model. Unfortunately, transferable utility greatly complicates the analysis by making match payoffs contingent not just on agent types but also on the endogenous outside options of each agent. However, we can analyze a two-type analogue of the model, with high types h and low types l and symmetric distributions. We'll focus on the match surplus function $u(h,h) \equiv 1$, $u(h,l) \equiv \beta$, $u(l,l) \equiv \gamma$, $1 > \beta > \gamma$. We'll say u has (weakly) supermodular payoffs if $2\beta \geq 1 + \gamma$ and assume this for the remainder of the section. Suppose an inflow rate normalized to 1, with the inflow of h types f, and the proportion of h types on the platform g and a mass of agents on the platform N. Suppose per-date costs are zero. When high types only accept high types, $gN = \sqrt{f}$, $(1-g)N = \sqrt{1-f}$, $g = \frac{\sqrt{f}}{\sqrt{f}+\sqrt{1-f}}$. Expected discount is then $\frac{gN}{gN+r} = \frac{\sqrt{f}}{\sqrt{f}+r}$

for high types and $\frac{\sqrt{1-f}}{\sqrt{1-f}+r}$ for low types. Then, given inflow rates f and $1-f$, the rate of surplus generated by a match is $f\frac{\sqrt{f}}{\sqrt{f}+r}1$ for the high type and $(1-f)\frac{\sqrt{1-f}}{\sqrt{1-f}+r}\gamma$ for the low type. When all agents accept one another, $g=f$ and $N=1$. Expected discount is then $\frac{1}{1+r}$ for all types and the rate of surplus generated is $f\frac{1}{1+r}(f+(1-f)\beta)$ for high types and $(1-f)\frac{1}{1+r}(f\beta+(1-f)\gamma)$ for low types. Thus, separation is optimal if and only if $(1-f)\frac{1}{1+r}(f\beta+(1-f)\gamma)+f\frac{1}{1+r}(f+(1-f)\beta)\leq(1-f)\frac{\sqrt{1-f}}{\sqrt{1-f}+r}\gamma+f\frac{\sqrt{f}}{\sqrt{f}+r}$.

We can now study the TU equilibrium assuming Nash Bargaining. Then, given continuations values C_h and C_l, match payoffs after transfers are $u_h(h,l)=1/2(\beta+C_h-C_l)$, $u_l(h,l)=1/2(\beta+C_l-C_h)$, $u_l(l,l)=\gamma/2$, $u_h(h,h)=1/2$. Suppose high types reject all low types. Then a high type receives expected payoff and continuation value $\frac{\sqrt{f}}{2(\sqrt{f}+r)}$, and a low type receives $\frac{\gamma\sqrt{1-f}}{2(\sqrt{1-f}+r)}$ and deviation to accepting a low type yields $(\beta+\frac{\sqrt{f}}{2(\sqrt{f}+r)}-\frac{\gamma\sqrt{1-f}}{2(\sqrt{1-f}+r)})/2$. Then separation isn't supportable in equilibrium when

$$\frac{\sqrt{f}}{2(\sqrt{f}+r)}<(\beta+\frac{\sqrt{f}}{2(\sqrt{f}+r)}-\gamma\frac{\sqrt{1-f}}{2(\sqrt{1-f}+r)})/2 \tag{1.6.1}$$

However, the social planner cares about the changed utility of the low type agent who matches to high type. Thus, the surplus for the two agents that match when the high type deviates is $\frac{\sqrt{f}}{2(\sqrt{f}+r)}+\frac{\gamma\sqrt{1-f}}{2(\sqrt{1-f}+r)}$ without the deviation and β when deviating. Then surplus is increased by deviating if and only if $\frac{\sqrt{f}}{2(\sqrt{f}+r)}+\frac{\gamma\sqrt{1-f}}{2(\sqrt{1-f}+r)}<\beta$, equivalent to the high type's inequality (1.6.1) . If the distribution on the platform was exogenous, this would conclude the analysis and the TU equilibrium would maximize total surplus. However, the mass of agents on the platform shrinks as agents become less picky, so a high type accepting low types imposes costs on others, Computing g given that a proportion x of high types accept low types and taking the limit as $x\to 0$, we find that a small mass x of h types deviating to accepting all lowers expected discount by $x\frac{(\sqrt{f}r)}{2(\sqrt{1-f}+r)^2}$ for high types and $x\frac{(\sqrt{1-f}r)}{2(\sqrt{f}+r)^2}$ for low types by decreasing the rate of draws. Then a high type agent accepting low types cannot be socially efficient unless

$$\frac{\sqrt{f}}{2(\sqrt{f}+r)} < (\beta+\frac{\sqrt{f}}{2(\sqrt{f}+r)}-\gamma\frac{\sqrt{1-f}}{2(\sqrt{1-f}+r)})/2+\frac{r\left(-f^2\left(\sqrt{f}+r\right)-(f-1)^2g\left(\sqrt{1-f}+r\right)\right)}{4\left(\sqrt{1-f}+r\right)^2\left(\sqrt{f}+r\right)^2} \tag{1.6.2}$$

where $\frac{r\left(-f^2\left(\sqrt{f}+r\right)-(f-1)^2g\left(\sqrt{1-f}+r\right)\right)}{4\left(\sqrt{1-f}+r\right)^2\left(\sqrt{f}+r\right)^2} < 0$. Thus there is an interval where, under TU, high types will accept low types despite it being socially inefficient for them to do so–that is, agents are not picky enough.

Non-Multiplicatively Separable Utility

Multiplicative separability of the own-type component utility is a strong assumption in this paper. With modular utility and a constant per-date cost it is automatically satisfied since $\psi = 1$, but with supermodular utility it imposes a functional form restriction on match surplus and requires that per-date costs be a constant multiplied by ψ, meaning per-date costs must be higher for higher types and imposing a very strong relationship between match utility and per-date costs. This precludes constant per-date costs, making analysis of per-date prices with supermodular utility infeasible. Thus, we'd like to be able to say that this assumption, while necessary for tractability, is not driving our results. To assess this, we study a discretized analogue to our model, with five types ($q = .2, q = .4, q = .6$, $q = .8, , q = 1$) rather than a continuum. We need to limit the number of types because, without multiplicative separability, different agents in any candidate class will have different optimization problems and employ different cutoff strategies, precluding the discrete class structure that made analysis tractable. Without this, we'll instead find LSSE equilibria by brute force, testing every possible combination of cutoff strategies for each type for profitable deviations.[27] We'll also find optimal platform strategies as in Section 1.4 by testing every viable firm strategy (where price is the IR of the lowest joining type) and selecting the one that maximizes profit. We'll study the case where $\psi(q) = q^\alpha$ and per-date costs are constant. This will also allow us to look at per-date pricing when utility is supermodular. We'll use 3 different distributions, a decreasing distribution $(.35, .3, .2, .1, .05)$, the discrete uniform

[27] We ignore mixed strategies.

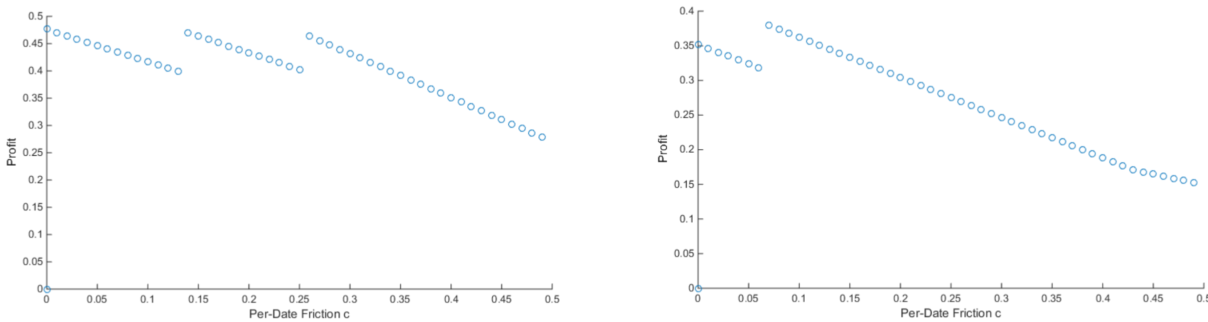

Figure 1.6.1: Cost per-date 1

From left to right: a) increasing distribution, $r = .05$, $\alpha = 0$, b) increasing distribution, $r = .1$, $\alpha = 1$.

$(.2, .2, .2, .2, .2)$, and an increasing distribution $(.05, .1, .2, .3, .35)$, and varying assumptions on r to study the equilibria under different conditions. For concision, we'll only report a few of the more salient examples here. Generally, the simulations using constant per-date costs are consistent with the results in Section 1.4 assuming per date costs of $c\psi(q)$.

In Figure 1.6.1 a), we see the discrete type analogue to the frictional modular utility case for a monopolist discussed in Proposition 8. As per-date frictions increase, the price that can be charged to the lowest joining type decreases, and there is a negative direct effect on profit. However, higher type agents become less selective, and when the per-date friction is high enough to induce a higher type to accept the lowest joining type there is a discontinuous increase in profit due to the effective transfer from the high type to the cutoff type which counterbalances the direct effect. Thus we see multiple levels of c that are consistent with maximizing revenue, as in the previous analysis. b) shows the case with $\psi(q) = q$, but unlike the formal analysis of Proposition 8, per-date frictions are c instead of cq, meaning that the class structure will not hold in equilibrium and the aforemention proposition does not apply. However ,we see qualitatively similar results, with a negative direct effect of c on profit and discrete jumps back to higher profit when higher types accept the lowest joining type.

In Figure 1.6.2 a), we see the discrete type analogue to the per-date price modular utility case for a monopolist discussed in Corollary 3. This case is very similar to the frictional case, but raising per-date prices doesn't decrease the amount of surplus that can be extracted from the lowest joining type, so the effective transfers from high to low types as per-date prices increase are the only salient effect, and profit increases with per-date price. b) shows the

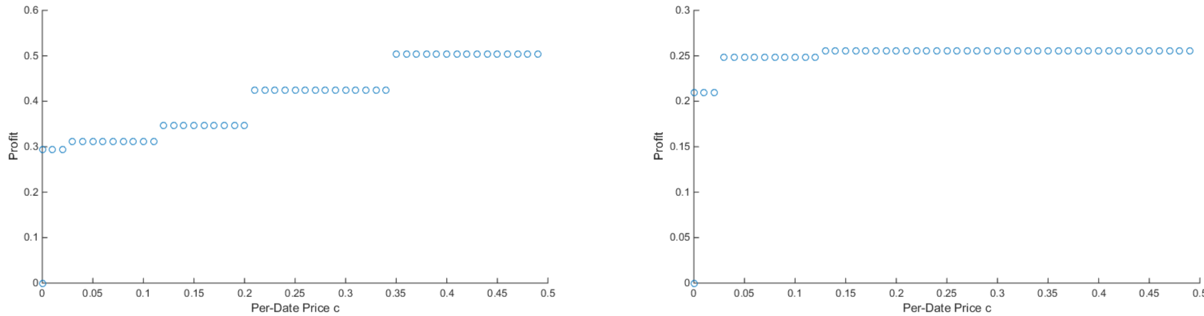

Figure 1.6.2: Cost Per Date 2

From left to right: a) increasing distribution, $r = .1$, $\alpha = 0$. b) uniform distribution, $r = .1$, $\alpha = 1$.

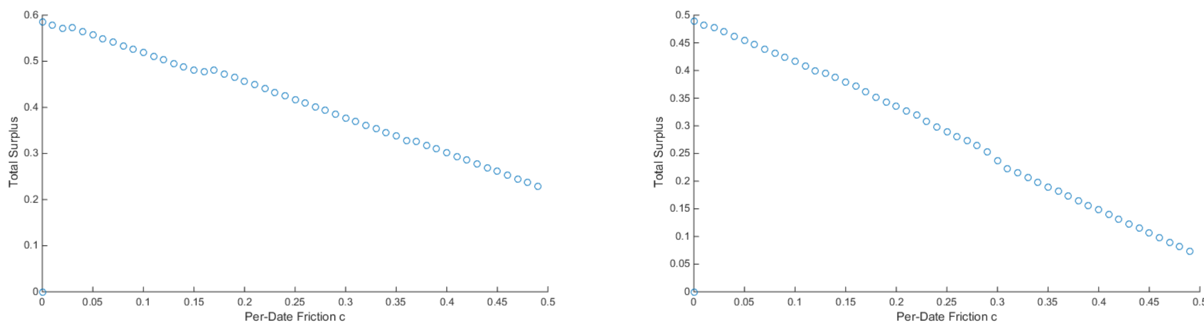

Figure 1.6.3: Cost Per Date 3

From left to right: a) increasing distribution, $r = .2$, $\alpha = 0$. b) increasing distribution, $r = .2$, $\alpha = 1$.

case with supermodular utility $\psi(q) = q$, and per-date price c, a case which could not be studied before due the lack of a class structure. In fact, however, we see the same situation, where higher per-date prices increase profit, and even though the class basis for the claim of Corollary 3 does not hold, the argument that the IR-cutoff agent's utility can be extracted through a combination of fixed fees and per-date prices, and higher per-date prices should make higher types less selective and thus force them to match to the IR-cutoff type, increasing their expected match utility should still hold.

In Figure 1.6.3 a), we see the discrete type analogue to the frictional modular utility case for a social planner discussed in Section 1.4.1. As discussed before, the direct negative effect of increasing per-date frictions dominates, and higher frictions generally lower surplus, although small local increases are possible due to the non-endogenous lower bound of the rump class. In b), we see the case with $\psi(q) = q$, and per-date friction c, and the effect of

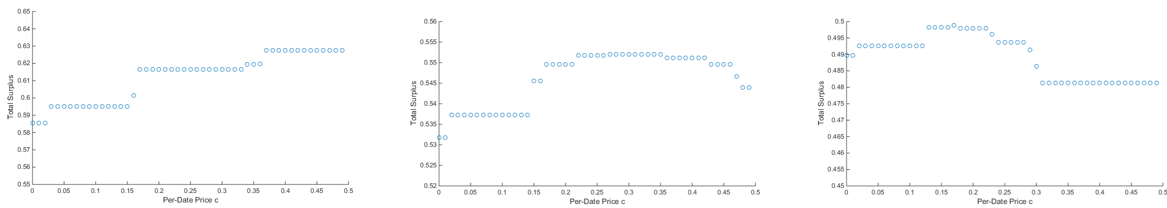

Figure 1.6.4: Cost Per Date 4

From left to right: a) increasing distribution, $r = .2$, $\alpha = 0$. b) increasing distribution, $r = .2$, $\alpha = .5$, c) increasing distribution, $r = .2$, $\alpha = 1$.

increasing c is qualitatively similar.

In Figure 1.6.4 a), we see the discrete type analogue to the per-date price modular utility case for a social planner discussed in Proposition 11. As discussed before, increasing per-date prices lowers the time costs of search, and because of modular utility sorting doesn't matter, so increasing per-date prices increases total surplus. We couldn't study optimal per-date prices with supermodular utility before due to tractability problems, but b) and c) we can examine numerical simulations of this case. As discussed before, with supermodular utility assortation increases surplus, so when per-date prices increase and agents become less picky, there will be a tradeoff between lowering time costs on the one hand and lowering sorting on the other. In fact, in b) with $\psi(q) = \sqrt{q}$ and moderate supermodularity we see exactly that, with total surplus initially increasing in per-date prices when decreasing time costs dominates and total surplus later decreasing when sorting effects dominate, yielding an optimal per-date price that is positive. With c), $\psi(q) = q$ and supermodularity is stronger. We see the same story here, but the optimal per-date price is significantly lower as the costs of lowering sorting are higher.

1.6.2 Proofs

Steady State–Linear Returns to Matching

We'll now provide a proof of Proposition 5 via two lemmas. This closely follows Burdett and Coles, but requires some adjustment to accommodate per-date costs. First, we'll transcribe a useful result from Burdett and Coles. Define $\Gamma(x_1, x_2) \equiv (F(x_1) - F(x_2))^2 - f(x_2) \int_{x_2}^{x_1} F(x_1) - F(x)dx$

Lemma 10 *An increasing hazard rate f(x)/(1-F(x)) implies* $\Gamma(.) \geq 0$

Define $\phi(q(n), q(n-1)) \equiv \int_{q(n)}^{q(n-1)} \frac{F(q(n-1))-F(x)}{F(q(n-1))-F(q(n))} dx - c$ for $q(n-1) > \underline{q}$. Since F is strictly increasing and twice differentiable, ϕ is well defined, continuous and twice differentiable almost everywhere, restricting ourselves to right differentiation at the lower bound. It can be shown that $\phi \to -c$ as class size goes to zero and $\phi \to \underline{q} - q(n) - c$ as $q(n-1) \to \underline{q}$. Lemma 10 implies ϕ is decreasing in $q(n)$:

$\frac{\partial}{\partial q(n)} \int_{q(n)}^{q(n-1)} \frac{F(q(n-1))-F(x)}{F(q(n-1))-F(q(n))} dx - c = \int_{q(n)}^{q(n-1)} \frac{(F(q(n-1))-F(x))f(q(n))}{(F(q(n-1))-F(q(n)))^2} dx - 1 \leq 0.$

We can also show that ϕ is strictly increasing in $q(n-1)$:

$\frac{\partial}{\partial q(n-1)} \int_{q(n)}^{q(n-1)} \frac{F(q(n-1))-F(x)}{F(q(n-1))-F(q(n))} dx - c = \int_{q(n)}^{q(n-1)} \frac{(F(x)-F(q(n)))f(q(n))}{(F(q(n-1))-F(q(n)))^2} dx > 0.$

Fix N and let $q_n(N)$, $\lambda_n(N)$, $J(N)$ satisfy

i) $q_0(N) = \overline{q}$

ii) if $q_{n-1}(N) > \underline{q}$, $q_n(N) = \phi(q_n(N), q_{n-1}(N))\frac{N\lambda_n(N)}{r}$, $q_n(N) = \phi(q_n(N), q_{n-1}(N))\frac{N\lambda_n(N)}{r}$, $\lambda_n(N) = \sqrt{F(q_{n-1}(N)) - F(q_n(N))}/N$

iii) if $q_{n-1}(N) <= \underline{q}$, $q_n(N) = \lambda_n(N) = 0$

The following lemma shows inductively that each cutoff is well behaved if the previous one is. The main challenge is to show uniqueness, especially in the presence of a per-date cost. In (1.3.9), the LHS is (obviously) increasing, so if we can show the RHS is decreasing, uniqueness is guaranteed. Thus the meat of the proof is establishing the properties of the RHS.

Lemma 11 *If* $q_{n-1}(N) > \underline{q}$ *and is continuous at N for some N>0, then there is a unique solution for* $q_n(N)$ *if participation in search in class n can be supported, where* $q_n(N)$ *is continuous at N,* $q_n(N) < q_{n-1}(N)$, $\lambda_n > 0$ *and is continuous at N.* q_n *and* λ_n *go to zero as* $q_{n-1} \to \underline{q}$. *Additionally,* $q_n(N)$ *is increasing in* $q_{n-1}(N)$,

Proof. $\frac{\partial}{\partial q(n)} \frac{1}{r}(\phi(q(n), q(n-1))\sqrt{(F(q(n-1)) - F(q(n)))}) = \frac{1}{r}(\phi_1\sqrt{(F(q(n-1)) - F(q(n)))} - \phi\frac{f(q(n))}{\sqrt{(F(q(n-1))-F(q(n)))}})$. Consider the minimal c such that $\phi_1\sqrt{(F(q(n-1)) - F(q(n)))} - \phi\frac{f(q(n))}{\sqrt{(F(q(n-1))-F(q(n)))}} \geq 0$.

Then $c = \frac{1}{r}(\int_{q(n)}^{q(n-1)} \frac{F(q(n-1))-F(x)}{F(q(n-1))-F(q(n))} dx - \frac{(F(q(n-1))-F(q(n)))}{f(q(n))}(\int_{q(n)}^{q(n-1)} \frac{(F(q(n-1))-F(x))f(q(n))}{(F(q(n-1))-F(q(n)))^2} dx - 1))$. Then the expected payoff is $\frac{1}{r}(\frac{(F(q(n-1))-F(q(n)))}{f(q(n))}\sqrt{(F(q(n-1)) - F(q(n)))})\phi_1$. ϕ_1 must be negative and the remainder of the expression is positive, so the RHS of (1.3.9) must be negative. Thus, either $\frac{1}{r}(\phi_1\sqrt{(F(q(n-1))-F(q(n)))}-\phi\frac{f(q(n))}{\sqrt{(F(q(n-1))-F(q(n)))}}) \leq 0$ and the RHS is decreasing while the LHS is increasing, ensuring a a unique solution, or c is high enough that any draw will be accepted ex-post, which also implies a unique cutoff. Direct inspection shows continuity given continuity of the constituent functions, and thus the continuity of $q_n(N)$ and $\lambda_n(N)$. The RHS is negative as $q_n \to q_{n-1}$ so $q_n < q_{n-1}$. Thus $\lambda_n(N) > 0$. $q_n \to 0$ goes to zero as $q_{n-1} \to \underline{q}$ since RHS is negative and $\lambda_n(N) \leq G(q_{n-1}) - G(\underline{q}) \to 0$ as $q_{n-1} \to \underline{q}$. Finally, $\frac{\partial}{\partial q(n-1)}\frac{1}{r}(\phi(q(n),q(n-1))\sqrt{(F(q(n-1))-F(q(n)))}) = \frac{1}{r}(q(n-1)-q(n)-c)\frac{f(q(n)}{2\sqrt{(F(q(n-1))-F(q(n)))}}$. Consider the minimal c such that $\frac{1}{r}(q(n-1)-q(n)-c)\frac{f(q(n)}{2\sqrt{(F(q(n-1))-F(q(n)))}} \leq 0$. Then $c = q(n-1)-q(n)$ and continuation value is $\frac{1}{G(q(n-1))-G(q(n))+r}\int_{q(n)}^{q(n-1)}(x-c)g(x)dx = \frac{1}{G(q(n-1))-G(q(n))+r}\int_{q(n)}^{q(n-1)}(x-q(n-1)+q(n))g(x)dx \leq \frac{1}{G(q(n-1))-G(q(n))+r}\int_{q(n)}^{q(n-1)}q(n)g(x)dx = \frac{G(q(n-1)_u)-G(q(n))}{G(q(n-1))-G(q(n))+r}q(n) < q(n)$. Thus $c \geq q(n-1)-q(n)$ ensures a corner solution for the cutoff, $\underline{q}$. Thus, q(n) is unchanging in q(n-1) if $c > q(n-1)-q(n)$, else the RHS is increasing in q(n-1), which, given that the LHS is unchanged and the LHS is increasing and the RHS decreasing in q(n), implies a rise in q(n-1) must induce a rise in q(n). ■

We'll now prove Lemma 7. Note that (1.3.3) can be rewritten as $q(n,\hat{q}) = \frac{(N\lambda_n(\hat{q})E[q'|n,\hat{q},match]-N\gamma_n(\hat{q})c)}{N\lambda_n(\hat{q})+r} = \frac{N\lambda_n(\hat{q})(E[q'|n,\hat{q},match]-c)}{N\lambda_n(\hat{q})+r} - \frac{N(\gamma_n(\hat{q})-\lambda_n(\hat{q}))c)}{N\lambda_n(\hat{q})+r}$ This is simply (1.3.4) minus c times a scalar on the RHS and with the values of $\gamma_n = Pr[date|n,\hat{q}]$, $\lambda_n = Pr[match|n,\hat{q}]$ and $E[q'|n,\hat{q},match]$ changed to reflect that agents may reject dates inside their class and accept agents outside their class, changing their probability of dating and matching over a given interval and changing the distribution of matches accepted. Every agent in a class must have the same expected match quality, so it will suffice to consider $q(n) = \frac{N\lambda_n(E[q'|n,match]-c)}{N\lambda_n+r} - \frac{N(\gamma_n-\lambda_n)c)}{N\lambda_n+r}$.

Lemma 12 *Given LRM and 1-F log-concave, a class n starting at q(n-1) where the probability of an agent in n rejecting a date inside the class is strictly positive must have* $\frac{N\lambda_n(E[q'|n,match]-c)}{N\lambda_n+r} \leq q_E(n_{LPPE,q(n)})$.

Proof. Forthcoming. ■

We can now prove Lemma 7:

Proof. Let k be the first class with a positive measure of agents deviating from the LPPE. Lemma 12 shows that $\frac{N\lambda_k(E[q'|k,match]-c)}{N\lambda_k+r} \leq q_E(k_{LPPE,q(k)})$, with strict inequality holding if the probability of rejecting a match within one's class is strictly positive, and $\frac{N(\gamma_k-\lambda_k)c)}{N\lambda_k+r} \geq 0$, with strict inequality holding if the probability of accepting a match outside one's class is strictly positive. Then, since $q_E(q(k)) = q(k)$, $q_E(k) < q_E(k_{LPPE,q(k)}) \leq q_E(LPPE, q(k-1)) \leq q_E(q(k))$. Contradiction. We can proceed inductively from here. Suppose that $q(n-1) \leq q(LPPE, n-1)$. Then $q(n) \leq q(LPPE, q(n-1))$ as before. Lemma 11 establishes that $q(LPPE, q(n-1)) \leq q(LPPE, n)$, so $q(n-1) \leq q(LPPE, n)$. ■

Steady State–Constant Returns to Matching

In addition to LRM, we can also study the analogous model with constant returns to matching. The CRM analysis largely follows Burdett and Coles (1997). While several proofs must be amended to account for per-date costs, some go through unchanged. Define the distribution of agents leaving the platform by $H(q)$ and the mass of agents leaving the platform by O.

Within a given class, we can get a simple characterization of outflow. Outflow in a class is given by the number of agents on the platform, N, times the proportion of agents in the class, λ_n, times the probability of an agent in the class drawing another agent in that class, λ_n. Then outflow from class n is $\lambda_n^2 N$. Then, in an LSSE,

$$\lambda_n = \sqrt{(F(q(n-1)) - F(q(n)))/N} \tag{1.6.3}$$

We also have that, for any $[z_1, z_2)$ in class n, $\lambda_n(G(z_2) - G(z_1))N = F(z_2) - F(z_1)$ and thus, with the differentiability of F,

$$g(q) = \frac{f(q)}{\lambda_n N} \tag{1.6.4}$$

Thus the density of agents on the platform in a given class is inflow density times a scalar. Combining equation 1.3.5 and balanced flow, we can get eliminate G terms, yielding class cutoffs solely in terms of inflows, N, and c.

$$q(n) = \frac{1}{r}(\int_{q(n)}^{q(n-1)} \frac{F(q(n-1)) - F(x)}{F(q(n-1)) - F(q(n))} dx - c)\sqrt{(F(q(n-1)) - F(q(n)))/N} \tag{1.6.5}$$

We're can now characterize the LSSE in this environment:

Proposition 12 *Given F, (G, N) defines a LSSE if and only if G satisfies (1.6.4) and $\{(\lambda_n, q(n))\}_{n=0}^{J}$ satisfies (1.6.3), (1.6.5), $q(0) = \overline{q}$, $q(J) <= \underline{q}$, and $\sum_n \lambda_n = 1$.*

Proof. $\sum_n \lambda_n = 1$, the boundary conditions, and (1.6.3)-(1.6.5) are necessary in an LSSE by construction. Conversely, the assumptions guarantee $G(\overline{q}) = 1$, $G(\underline{q}) = 0$ and G increasing, so G is a well defined steady state distribution and any G and N satisfying them form a valid LSSE. ■

To ensure existence of an LSSE, we'll need to make some distributional assumptions. The increasing hazard rate will ensure that, for each possible N, the class structure is unique. We'll now provide a proof of Proposition 13 via a lemma. This closely follows Burdett and Coles, but requires some adjustment to accommodate per-date costs.

Lemma 10 implies ϕ is decreasing in $q(n)$:
$\frac{\partial}{\partial q(n)} \int_{q(n)}^{q(n-1)} \frac{F(q(n-1)) - F(x)}{F(q(n-1)) - F(q(n))} dx - c = \int_{q(n)}^{q(n-1)} \frac{(F(q(n-1)) - F(x)) f(q(n))}{(F(q(n-1)) - F(q(n)))^2} dx - 1 \leq 0.$

We can also show thatϕ is strictly increasing in q(n-1):
$\frac{\partial}{\partial q(n-1)} \int_{q(n)}^{q(n-1)} \frac{F(q(n-1)) - F(x)}{F(q(n-1)) - F(q(n))} dx - c = \int_{q(n)}^{q(n-1)} \frac{(F(x) - F(q(n))) f(q(n))}{(F(q(n-1)) - F(q(n)))^2} dx > 0.$

Fix N and let $q_n(N)$, $\lambda_n(N)$, $J(N)$ satisfy

i) $q_0(N) = \overline{q}$

ii) if $q_{n-1}(N) > \underline{q}$, $q_n(N) = \phi(q_n(N), q_{n-1}(N)) \frac{\delta \lambda(N)}{r}$, $q_n(N) = \phi(q_n(N), q_{n-1}(N)) \frac{\lambda_n(N)}{r}$, $\lambda_n(N) = \sqrt{F(q_{n-1}(N)) - F(q_n(N))/N}$

iii) if $q_{n-1}(N) <= \underline{q}$, $q_n(N) = \lambda_n(N) = 0$

The following Lemma shows inductively that each cutoff is well behaved if the previous one is. The main challenge is to show uniqueness, especially in the presence of a per-date cost. In 2.8, the LHS is (obviously) increasing, so if we can show the RHS is decreasing, uniqueness is guaranteed. Thus the meat of the proof is establishing the properties of the RHS.

Lemma 13 *If $q_{n-1}(N) > \underline{q}$ and is continuous at N for some N>0, then there is a unique solution for $q_n(N)$, where $q_n(N)$ is continuous at N, $q_n(N) < q_{n-1}(N)$, and $\lambda_n > 0$ and is continuous at N. $q_n(N)$ and $\lambda_n(N)$ go to zero as $q_{n-1}(N) \to \underline{q}$. Additionally, $q_n(N)$ is increasing in $q_{n-1}(N)$,*

Proof. $\frac{\partial}{\partial q(n)}\frac{1}{r}(\phi(q(n),q(n-1))\sqrt{(F(q(n-1))-F(q(n)))/N}) = \frac{1}{r}(\phi_1\sqrt{(F(q(n-1))-F(q(n)))/N} - \phi\frac{f(q(n))}{\sqrt{N(F(q(n-1))-F(q(n)))}})$. Consider the minimal c such that $\phi_1\sqrt{(F(q(n-1))-F(q(n)))/N} - \phi\frac{f(q(n))}{\sqrt{N(F(q(n-1))-F(q(n)))}} \geq 0$.

Then $c = \int_{q(n)}^{q(n-1)} \frac{F(q(n-1))-F(x)}{F(q(n-1))-F(q(n))}dx - \frac{(F(q(n-1))-F(q(n)))}{f(q(n))}(\int_{q(n)}^{q(n-1)} \frac{(F(q(n-1))-F(x))f(q(n))}{(F(q(n-1))-F(q(n)))^2}dx-1)$. Then the RHS of (1.6.5) is $\frac{1}{r}(\frac{(F(q(n-1))-F(q(n)))}{f(q(n))}\sqrt{(F(q(n-1))-F(q(n))/N)}\phi_1$. By Lemma 10, ϕ_1 must be negative and the remainder of the expression is positive, so expected payoff must be negative. Thus, either $\frac{1}{r}(\phi_1\sqrt{(F(q(n-1))-F(q(n)))/N} - \phi\frac{f(q(n))}{\sqrt{N(F(q(n-1))-F(q(n)))}}) \leq 0$ and the RHS is decreasing while the LHS is increasing, ensuring a unique solution, or c is high enough that any draw will be accepted ex-post, which also implies a unique cutoff. Direct inspection shows continuity given continuity of the constituent functions, and thus the continuity of $q_n(N)$ and $\lambda_n(N)$. The RHS is negative as $q_n \to q_{n-1}$ so $q_n < q_{n-1}$. Thus $\lambda_n(N) > 0$. $q_n \to 0$ goes to zero as $q_{n-1} \to \underline{q}$ since RHS is negative and $\lambda_n(N) \leq G(q_{n-1}) - G(\underline{q}) \to 0$ as $q_{n-1} \to \underline{q}$. Finally, $\frac{\partial}{\partial q(n-1)}\frac{1}{r}(\phi(q(n),q(n-1))\sqrt{(F(q(n-1))-F(q(n)))/N})= \frac{1}{r}(q(n-1)-q(n)-c)\frac{f(q(n)}{2\sqrt{N(F(q(n-1))-F(q(n)))}}$. Consider the minimal c such that $\frac{1}{r}(q(n-1)-q(n)-c)\frac{f(q(n)}{2\sqrt{N(F(q(n-1))-F(q(n)))}} \leq 0$. Then $c = q(n-1)-q(n)$ and continuation value is $\frac{1}{G(q(n-1))-G(q(n))+r}\int_{q(n)}^{q(n-1)}(x-c)g(x)dx = \frac{1}{G(q(n-1))-G(q(n))+r}\int_{q(n)}^{q(n-1)}(x-q(n-1)+q(n))g(x)dx \leq \frac{1}{G(q(n-1))-G(q(n))+r}\int_{q(n)}^{q(n-1)} q(n)g(x)dx = \frac{G(q(n-1)_u)-G(q(n))}{G(q(n-1))-G(q(n))+r}q(n) < q(n)$. Thus $c \geq q(n-1)-q(n)$ ensures a corner solution for the cutoff, $\underline{q}$. Thus, q(n) is unchanging in q(n-1) if $c > q(n-1)-q(n)$, else the RHS is increasing in q(n-1), which, given that the LHS is unchanged and the LHS is increasing and the RHS decreasing in q(n), implies a rise in q(n-1) must induce a rise in q(n). ■

Proposition 13 *For all N>0, there exist unique, continuous solutions for $q_n(N)$ and $\lambda_n(N)$ satisfying (1.6.3)-(1.6.5), $q_0(N) = \overline{q}$ and $q_J(N) \leq \underline{q}$. such that $q_n(N) < q_{n-1}(N)$ and $\lambda_n(N) > 0$ if $q_{n-1}(N) > \underline{q}$.*

Proof. For the base case of $q_0 = \overline{q}$, q_0 is a constant function of N. Lemma 13 ensures that $q_{n-1}(N)$ continuous implies $q_n(N)$ continuous, and $q_n(N) < q_{n-1}(N)$, so induction follows. Continuity of $\lambda_n(N)$ follows from the continuity of $q_n(N)$, $q_{n-1}(N)$, and $\sqrt{(F(x)-F(y)/N}$. $\lambda_n(N) > 0$ follows from the fact that $q_n(N) < q_{n-1}(N)$ and $f(q) > 0$ for $q \in [\underline{q},\overline{q}]$. ■

However, N may not be consistent with $G(\overline{q}) = 1$, so we'll need an additional result. Log-concavity ensures the continuity of class sizes, and that, along with values of N inducing values $G(\overline{q}, N)$ above and below 1, ensures the existence of $\{\lambda_n\}_{n=1}^{J}$ such that $\sum_n \lambda_n = 1$. It's worth noting that the inclusion of a per-date cost makes uniqueness of the cutoffs harder to obtain than in the Burdett and Coles environment–without per date costs, the agent's optimization problem has convenient monotonicity properties that per-date costs militate against. However, the already necessary assumption of log-concavity of the survivor function also eliminates cases where per-date costs could induce a multiplicity of cutoffs.

Proposition 14 *An LSSE exists.*

Proof. Proposition 13 guarantees this result so long as $\sum_n \lambda_n(N) = 1$. $\sum_n \lambda_n(N) = \frac{1}{\sqrt{N}} \sum_n \sqrt{F(q_{n-1}(N)) - F(q_n(N))} > \frac{1}{\sqrt{N}} \sum_n F(q_{n-1}(N)) - F(q_n(N)) = \frac{1}{\sqrt{N}}(F(\overline{q}) - F(\underline{q})) = \frac{1}{\sqrt{N}}$, so $\underset{N \to 0}{lim} \sum_n \lambda_n(N) = \infty$. $\sum_n \lambda_n(N)$

$\Rightarrow q(1) = \frac{1}{r}(\int_{q_1(N)}^{\overline{q}} \frac{1-F(x)}{1-F(q_1(N))} dx - c)\sqrt{1 - F(q_1(N)))/N}$

$\frac{1-F(x)}{1-F(q_1(N))} < 1$ if $x \in (q_1(N), \overline{q}]$ so

$\int_{q_1(N)}^{\overline{q}} \frac{1-F(x)}{1-F(q_1(N))} dx < \overline{q} - q_1(N)$.

Then we have

$q_1(N) < \frac{1}{r}(\overline{q} - q_1(N) - c)\sqrt{(1 - F(q_1(N))/N}$

$q_1(N)\frac{1+r}{r} < \frac{1}{r}(\overline{q} - c)\sqrt{(1 - F(q_1(N)))/N}$

$q_1(N) < \frac{1}{1+r}(\overline{q} - c)\sqrt{(1 - F(q_1(N)))/N}$

so $q_1(N) \to 0$. For N sufficiently large, $q_1(N) < \underline{q}$, so $F(q_1(N)) = 0$. Then

$\sum_n \sqrt{F(q_{n-1}(N)) - F(q_n(N))} = \sqrt{F(q_0(N)) - F(q_1(N))} = 1$

$\sum_n \lambda_n(N) = \frac{1}{\sqrt{N}} \to 0$.

$\lambda_n(N)$ is continuous for all n, so $\sum_n \lambda_n(N)$ is continuous. Then, given $\sum_n \lambda_n(N) > 1$ for some N and $\sum_n \lambda_n(N) < 1$ for some N, the IVT ensures an N exists such that $\sum_n \lambda_n(N) = 1$.

■

Finally, we'd like to have uniqueness. This will require further distributional assumptions. Burdett and Coles only need that $xf(x)$ is increasing, but the inclusion of per-date costs again imposes stronger requirements for uniqueness. Unfortunately, in this case their assumptions are not strong enough to resolve the monotonicity issues with per-date costs.

For sufficiently small per-date costs, the increasing $xf(x)$ assumption is adequate, but to ensure uniqueness for any per-date cost we'll need the stronger assumption that $f(x)$ is increasing. This assumption is quite onerous, so we'll stick with the weaker assumption from Burdett and Coles and focus on sufficiently small per-date costs in the later analysis.

Lemma 14 *$q_n(N)$ is decreasing and differentiable in N.*

Proof. Differentiability follows from induction on (1.6.5). For the first class, we must have $q_1(N) = \frac{1}{r}(\int_{q_1(N)}^{\overline{q}} \frac{1-F(x)}{1-F(q_1(N))}dx - c)\sqrt{1-F(q_1(N)))/N}$. Denote the LHS L and the RHS R. Lemma 13 shows $R_{q_1} = \frac{\partial}{\partial q_1(N)}\frac{1}{r}(\int_{q_1(N)}^{\overline{q}} \frac{1-F(x)}{1-F(q_1(N))}dx - c)\sqrt{1-F(q_1(N)))/N}$ is negative while $L_1 = 1$ is positive. Suppose $q_1(N)$ is weakly increasing in N for some N. Then $L_N = q_1^{'}(N) > 0$ and $R_N = q_1^{'}(N)R_{q_1} - R/(2N)$ is negative. But L=R. Contradiction. We can now procede inductively. Suppose $q_{n-1}(N)$ is decreasing in N. (1.6.5) must hold, and $R_{q_n} = \frac{\partial}{\partial q_n(N)}\frac{1}{r}(\int_{q_n(N)}^{q_{n-1}(N)} \frac{F(q_{n-1}(N))-F(x)}{F(q_{n-1}(N))-F(q_n(N))}dx - c)\sqrt{(F(q_{n-1}(N)) - F(q_n(N)))/N}$ is negative and $R_{q_{n-1}} = \frac{\partial}{\partial q_{n-1}(N)}\frac{1}{r}(\int_{q_n(N)}^{q_{n-1}(N)} \frac{F(q_{n-1}(N))-F(x)}{F(q_{n-1}(N))-F(q_n(N))}dx - c)\sqrt{(F(q_{n-1}(N)) - F(q_n(N)))/N}$ is positive by Lemma 13. Then the $L_N = q_n^{'}(N) > 0$ and $R_N = q_n^{'}(N)R_{q_n} + q_{n-1}^{'}(N)R_{q_{n-1}} - R/(2N)$, which is negative since $q_{n-1}^{'}(N)$ is negative. ■

Lemma 15 *$\lambda_{n-1} \geq \lambda_n$ for any N>0 with $xf(x)$ for c sufficiently small or any c with $f(x)$ increasing.*

Proof. Differentiability follows from induction on (1.6.5). For the first class, we must have $q_1(N) = \frac{1}{r}(\int_{q_1(N)}^{\overline{q}} \frac{1-F(x)}{1-F(q_1(N))}dx - c)\sqrt{1-F(q_1(N)))/N}$. Denote the LHS L and the RHS R. Lemma 13 shows $R_{q_1} = \frac{\partial}{\partial q_1(N)}\frac{1}{r}(\int_{q_1(N)}^{\overline{q}} \frac{1-F(x)}{1-F(q_1(N))}dx - c)\sqrt{1-F(q_1(N)))/N}$ is negative while $L_1 = 1$ is positive. Suppose $q_1(N)$ is weakly increasing in N for some N. Then $L_N = q_1^{'}(N) > 0$ and $R_N = q_1^{'}(N)R_{q_1} - R/(2N)$ is negative. But L=R. Contradiction. We can now procede inductively. Suppose $q_{n-1}(N)$ is decreasing in N. (1.6.5) must hold, and $R_{q_n} = \frac{\partial}{\partial q_n(N)}\frac{1}{r}(\int_{q_n(N)}^{q_{n-1}(N)} \frac{F(q_{n-1}(N))-F(x)}{F(q_{n-1}(N))-F(q_n(N))}dx - c)\sqrt{(F(q_{n-1}(N)) - F(q_n(N)))/N}$ is negative and $R_{q_{n-1}} = \frac{\partial}{\partial q_{n-1}(N)}\frac{1}{r}(\int_{q_n(N)}^{q_{n-1}(N)} \frac{F(q_{n-1}(N))-F(x)}{F(q_{n-1}(N))-F(q_n(N))}dx - c)\sqrt{(F(q_{n-1}(N)) - F(q_n(N)))/N}$ is positive by Lemma 13. Then the $L_N = q_n^{'}(N) > 0$ and $R_N = q_n^{'}(N)R_{q_n} + q_{n-1}^{'}(N)R_{q_{n-1}} - R/(2N)$, which is negative since $q_{n-1}^{'}(N)$ is negative. ■

Lemma 16 *$xf(x)$ strictly increasing in x guarantees $\lambda_{n-1} \geq \lambda_n$ for c such that $c' \geq c > 0$ for some c'. $f(x)$ increasing guarantees $\lambda_{n-1} \geq \lambda_n$ for any c.*

Proof. Trivial for $n-1 \geq J(N)$. For $n-1 < J(N)$, we first want to show $\lambda_{n-1} \geq \lambda_n$ for all n. Define $\theta(q_l, q_h) = \int_{q_l}^{q_h} \frac{F(q_h)-F(x)}{F(q_h)-F(q_l)} dx - c)\sqrt{(F(q_h) - F(q_l))}$. Thus $\frac{1}{\sqrt{N}r} = q_l/\theta(q_l, q_h)$. $\lambda = \sqrt{(F(q_h) - F(q_l))/N}$ so it will suffice to show $\frac{\partial}{\partial q_h} F(q_h) - F(q_l(q_h))$ is increasing. The implicit function theorem yields $\frac{\partial}{\partial q_h} F(q_h) - F(q_l(q_h)) = f(q_h) - f(q_l) \frac{\frac{\delta}{\sqrt{N}(1-\delta)}\theta_2}{1-\frac{\delta}{\sqrt{N}(1-\delta)}\theta_1} = f(q_h) - f(q_l)\frac{q_l\theta_2/\theta}{1-q_l\theta_1/\theta}$. We can show that this is non-negative if and only if $\int_{q_l}^{q_h} x\, f(x) - q_l\, f(q_l) - c\, dx \geq 0$. If c is sufficiently small, this will be satisfied (clearly always satisfied for c=0 and $xf(x)$ increasing.) If $xf(x)$ strictly increasing, a strictly positive c can be supported. It can be shown that if $\frac{(q_h-q_l)f(q_l)}{F(q_h)-F(q_l)} \leq 1$, any c large enough to violate the above inequality will yield a corner solution for any agent's optimization problem, with agents accepting any match, a single class, and uniqueness thus ensured. if $f(x)$ is increasing, $\frac{(q_h-q_l)f(q_l)}{F(q_h)-F(q_l)} \leq 1$. ■

Proposition 15 *$xf(x)$ strictly increasing in x guarantees the existence of a unique LSSE for all c such that $c' \geq c > 0$ for some c'. $f(x)$ increasing guarantees uniqueness for any c.*

Proof. Total differentiation of $\lambda_{n-1}(N) = \sqrt{F(q_{n-1}(N)) - F(q_n(N))/N}$ yields $\frac{q'_{n-1}(N)f(q_{n-1}(N)) - q'_n(N)f(q_n(N))}{2\lambda_n N} - \frac{\lambda_n}{2N}$ for all but the last class and $\frac{q'_{n-1}(N)f(q_{n-1}(N))}{2\lambda_n N} - \frac{\lambda_n}{2N}$ for the last class. Summing over n, we have $-\frac{\lambda_J}{2N} + \sum_{n=1}^{J(N)-1} (\frac{1}{2N}(\frac{1}{\lambda_{n+1}} - \frac{1}{\lambda_n})q'_n(N)f(q_n(N)) - \frac{\lambda_n}{2N})$. λ_n are decreasing in n by Lemma 16, so $\frac{1}{\lambda_{n+1}} - \frac{1}{\lambda_n}$ increasing, and $q'_n(N)$ decreasing by Lemma 15. Thus the sum is negative, proving that $\sum_n \lambda_n$ is strictly decreasing in N. Thus, the N such that $\sum_n \lambda_n = 1$ must be unique, and so the LSSE. ■

Strategic Platforms

We'll now prove Proposition 9. Define $\lambda_F(n) \equiv F(q(n-1)) - F(q(n))$, the inflow mass in class n. Suppose $\psi(q) = q^a$ and define the mean ψ value in class n as $m_\psi(n) \equiv \int_{q(n)}^{q(n-1)} q^a \frac{f(q)}{\lambda_F(n)} dq$. Define the length of class n as $l(n) \equiv q(n-1) - q(n)$.

Proof. $\frac{\partial TS}{\partial c} = \sum_{n=1}^{J} q_{Ec}(n,c) m_\psi(n) \lambda_F(n) +$
$q_E(n,c)\left(q_c(n-1,c)q(n-1,c)^\alpha f(q(n-1,c)) - q_c(n,c)q(n,c)^\alpha f(q(n,c))\right)$

Using the fact that $q_E(n,c) = q(n,c)$ for all but the last class J, manipulating the summation, and suppressing c, we have

$\frac{\partial TS}{\partial c} = \sum_{n=1}^{J-2} q_c(n)\left(m_\psi(n) \cdot \lambda_F(n) - q(n)^\alpha \cdot f(q(n)) \cdot l(n+1)\right)$

$+q_c(J-1)\left(m_\psi(J-1)\cdot\lambda_F(J-1)-q(J-1)^\alpha\cdot f(q(J-1))(q(J-1)-q_E(J))\right)$

$+q_c(J)\left(m_\psi(J)\cdot\lambda_F(J)-q(J)^\alpha\cdot f(q(J))\cdot q_E(J)\right)$

If every term in this summation is negative, $\frac{\partial TS}{\partial c}$ is negative. By Proposition 7, $q_c(n)$ is decreasing, so it suffices to show $m_\psi(n)\cdot\lambda_F(n) > q(n)^\alpha\cdot f(q(n))\cdot l(n+1)$ for each n<J-1, and the corresponding inequalities for J-1 and J. By Jensen's inequality, $m_\psi(n) > E[q|q\in[q_n,q_{n-1})]^\alpha$. Given that F has full support, $E[q|q\in[q_n,q_{n-1})]^\alpha > q_n$. Thus, as $\alpha\to\infty$, $E[q|q\in[q_n,q_{n-1})]^\alpha/q^\alpha\to\infty$. Then $m_\psi(n)\cdot\lambda_F(n) > q(n)^\alpha\cdot f(q(n))\cdot l(n+1)$ for α sufficiently high, and since J is finite, an α exists ensuring $m_\psi(n)\cdot\lambda_F(n) > q(n)^\alpha\cdot f(q(n))\cdot l(n+1)$ for all $n < J-1$, as well as the inequalities for J-1 and J. ■

We'll now study a selection of simulations. We focus on the modular utility case, and find that increasing c typically decreases total surplus. This is not surprising given that a per-date friction of c decreases each agent's payoff by c. However, efficiency gains due to less selective agents and changes in the rump class due to c can outweigh the direct cost of c in some cases. In Figure 1.6.5, we assume a decreasing density f. This puts more weight on the rump class, creating more periodicity in the surplus as the class cutoffs shift downward in c and the

rump class goes from being a large, relatively efficient class to a small class whose size is limited by the support of the distribution, and then to a large class again as the last class cutoff passes the bottom of the support of the distribution and the next class becomes the last class. In both the case where r=.1 and r=.001, we clearly see the periodic component to total surplus, and over some intervals total surplus is actually increasing in c. Generally, the degree of discounting doesn't make a large difference unless agents are very impatient.

In Figure 1.6.7 a), we again assume a decreasing density f, but assume the a smaller support for f. This largely eliminates the periodic component to total surplus and leaves only the direct effect of c–total surplus is approximately linearly decreasing in c. Because the lowest type agents in the distribution are still half the quality of the highest types, there are far fewer endogenous classes, and efficiency losses due to excessive selectivity are lower. Thus, efficiency gains due to increasing c are much less relevant. In Figure 1.6.7 b) we consider an analogous case with a uniform distribution. The factors that could lead to increasing total surplus are weaker in this case, but we still see some periodic effect and a

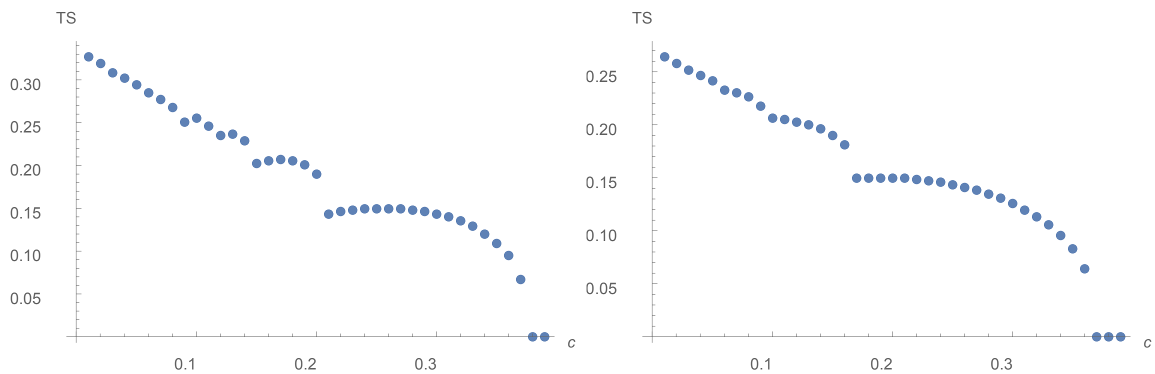

Figure 1.6.5: Total Surplus and Per-date Costs 1

From left to right: a) $F(x) = \sqrt{x-q}/\sqrt{1-q}$, $r = .001$, $q = .01$ b) $F(x) = \sqrt{x-q}/\sqrt{1-q}$, $r = .1$, $q = .01$

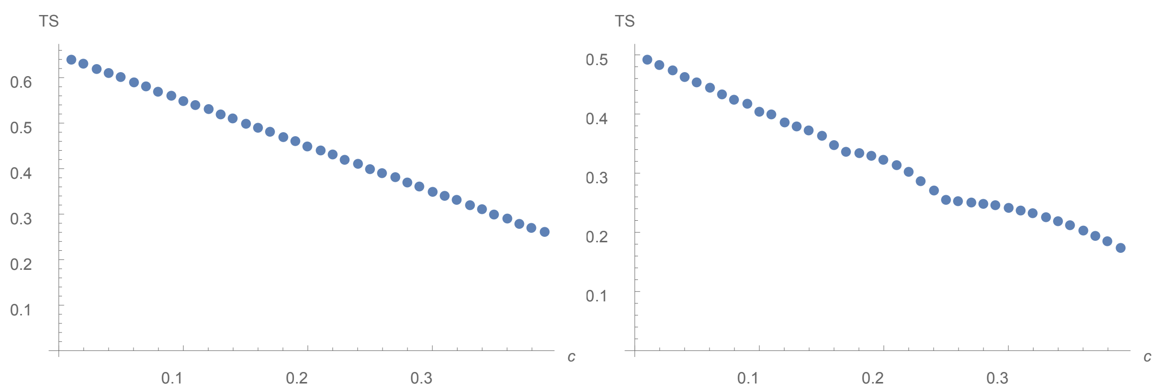

Figure 1.6.6: Total Surplus and Per-date Costs 2

From left to right: a) $F(x) = \sqrt{x-q}/\sqrt{1-q}$, $r = .001$, $q = .5$ b) $F(x) = (x-q)/(1-q)$, $r = .001$, $q = .01$

small region where total surplus is slightly increasing in c.

In Figure 1.6.7 a), we see the uniform distribution case corresponding to Figure 1.6.7 a). Again, the narrow range of qualities yields a monotonically decreasing total surplus. Finally, in Figure 1.6.7 b) we see total surplus when f is increasing. Here, even with a lower limit of support close to zero the second order effects are dominated by the linear friction costs.

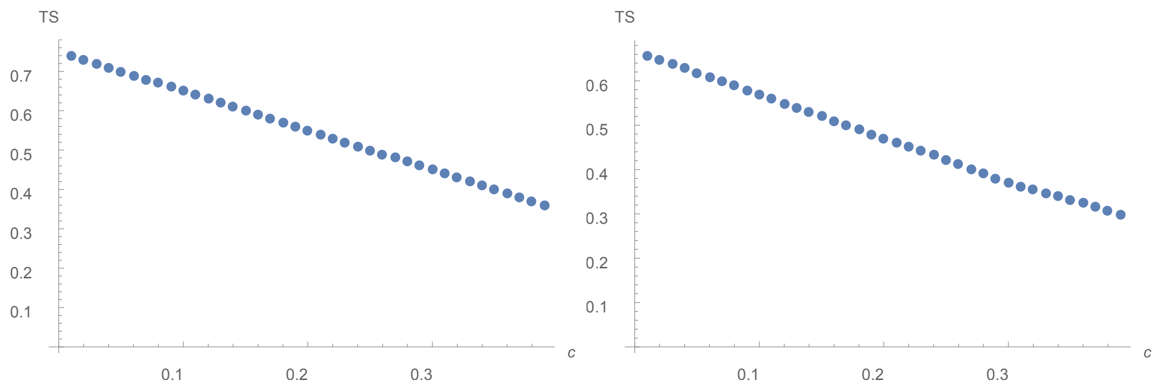

Figure 1.6.7: Total Surplus and Per-date Costs 2

From left to right: a) $F(x) = (x-q)/(1-q)$, $r = .001$, $q = .5$, $q = .001$ b) $F(x) = (x-q)^2/(1-q)^2$, $r = .001$, $q = .01$

CHAPTER 2

MATCHING WITH SINGLE PEAKED PREFERENCES

2.1 Introduction

In theoretical and empirical 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, preferences may be much more general, with different individuals having different preference orderings (violating the assumptions of vertical models) while not necessarily having a preference for their own type (violating the horizontal assumption).

In this paper we derive a simple, closed form matching function under single-peaked univariate preferences, where each agent is characterized by a univariate type and has an ideal type, preferring partners closer to that type, and where utility is nontransferable. Applications include marriage and dating, as well as job search when bargaining over wages is difficult or impossible, as in many public sector and entry level professional jobs. This result generalizes previous results like Becker (1973) and Clark (2007) to a much broader class of preferences. Specifically, we consider the case where there are continuous distributions of agents and, after removing perfect matches (agent a's ideal type is agent b's type and vice versa), the remainder distributions of unmatched agents can be separated by a monotonic curve[1], and where the masses on each side of this curve are not severely imbalanced in a way that will be described later. This assumption will often hold when the bivariate distributions

[1]More specifically, their graphs can be separated by this curve after one side of the market has transposed so that own type on one side corresponds to ideal type on the other. This transposition will be critical to our intuition throughout this paper.

of own and ideal types for each side have relatively similar dispersion and different mean vectors. For example, in marriage and dating, men may generally prefer women slightly shorter than them, with some variation, and women may generally prefer taller men, again with some variation. Similarly, some men may prefer women with similar incomes to them, while other men and all women may prefer higher incomes.

This paper makes contributions relevant to the empirical and theoretical literature on matching. First, it contributes to the wide empirical literature on *assortative* matching.[2] We show that matching over single peaked preferences exhibits several different forms of assortation, which form testable predictions. In particular, individuals who are perfect matches will match stably, and they will exhibit converse positive assortative matching (CPAM), where increases in an agent's ideal type correspond to increases in their match's type. Individuals who don't get perfect matches, however, exhibit other forms of assortation. If the supply for a given type of agent meets demand, in a sense that will be made explicit later, we find positive assortative matching (PAM), where higher types match to higher types irrespective of ideal type, or negative assortative matching (NAM), where higher types match to lower types, where the type of assortation depends on the relative orientation of the two distributions. We see PAM when agents generally prefer higher types than their stable matches (vertical preferences is one example of this). We see NAM when agents generally prefer lower types than their stable matches. Finally, when one side is locally in shortage, some will match as before, but others will be able to leverage their scarcity to match to agents of their ideal type (CPAM for the side in shortage) who do not find them ideal.

This paper also makes theoretical models involving embedded matching problems tractable in a much more general preference environment. For example, it provides a framework for studying theoretical models of marital sorting on income when different men and women have different preferences.

This paper follows a rich literature on stable matching problems, starting with the seminal paper by Gale and Shapley (1964). Becker (1973) found that PAM occurs when there

[2]In one dimension, assortation is a matching structure where the type of an agent's match is monotonic in the agent's own type.

is a continuum of types and the utility of a match is increasing in types and nontransferable. 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 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.

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.

The remainder of this paper is organized as follows: Section 2 demonstrates Clark's matching algorithm for the simplest case where preferences are homophilic–that is, where the peak preference is the agent's own type. Section 3 generalizes the model by allow agents to have arbitrary single peaked preferences, and matching algorithms are derived given some additional assumptions. Section 4 relates the single-peaked matching result to the horizontal and vertical preference literature. Section 5 provides interpretation for the results and empirical implications. Finally, the concluding section describes directions for further study.

2.2 Model with homophilic preferences

2.2.1 Baseline Model

[3] Consider a two sided one-to-one matching model with two continuous, integrable distributions A and B with full support on $[l, u] \times [l, u]$. We'll abuse notation by also letting A and B represent the set of agents on each side. We'll call the first dimension type θ, and the second preference p. Denote an agent i with type θ and preference p on side S as $s_{i\theta p}$, $s_{\theta p}$ if suppressing the index is appropriate, and s if suppressing both is possible. Let preferences be strictly single peaked. That is, for an agent $a_{\theta p}(b_{\theta p})$, if $\theta_2 > \theta_1 \geq \theta$ or $\theta \geq \theta_1 > \theta_2$, $b_{\theta_1 p_1} \underset{a_{\theta p}}{\succ} b_{\theta_2 p_2}$ ($a_{\theta_1 p_1} \underset{b_{\theta p}}{\succ} a_{\theta_2 p_2}$), with indifference over identical types. Before we address more general single peaked preferences, it is instructive to review the horizontal preference matching algorithm first derived in Clark (2007). Suppose for every agent $s_{\theta p}$ $p = \theta$ –that is, agents have homophilic preferences and we can suppress p. Suppose further that agents face no search costs or other limitations to matching, i.e. suppose agents optimize over the entire set of agents who are willing to match to them. Note that, while we can normalize either A or B to measure 1 without loss of generality, making both measure 1 is a simplifying assumption, requiring an equal mass of agents on each side. We will proceed for now using this assumption as it simplifies the problem, and relax it later.

- **Assumption 1 (MASS)** : Suppose an equal measure of agents on each side.

This scenario affords an extremely simple solution. First, we match and remove from consideration the area under both curves, if such an area exists.

Lemma 17 *For each type θ, a measure $\mu_\theta = min(\mu_{\theta A}, \mu_{\theta B})$ of θ agents on side A (B) matches to θ types on side B (A), where $\mu_{\theta S}$ is the mass of agents of type θ on side S.*

Proof. There are at least μ_θ agents of type θ on each side by definition. Since preferences are homophilic, $a_{i\theta}$ strictly prefers $b_{j\theta}$ to any agent b_{ks}, and symmetrically $b_{i\theta}$ prefers $a_{j\theta}$ to any a_{ks}. Then a mass μ_θ of a_θ's will strictly prefer to match with any of the measure μ_θ of

[3] I derived these horizontal results independently, being unaware of Clark's unpublished paper. However I believe the derivation is a good motivation for the later, novel result, so I have retained this section.

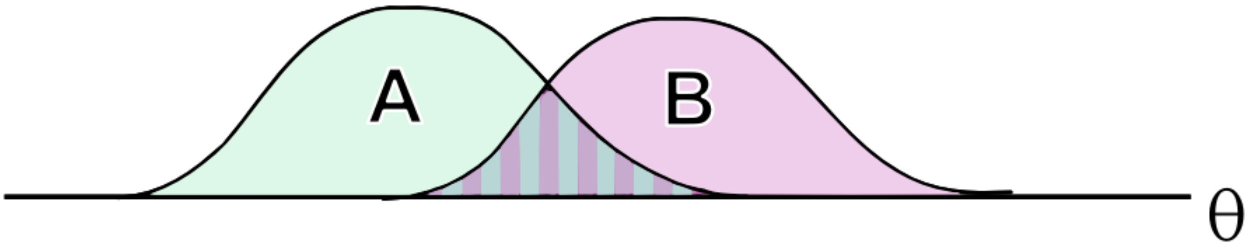

Figure 2.2.1: Initial distributions A and B of agents over type θ.

There is an area of overlap in the center where, for any type in the overlap, both distributions have at least as much mass as the overlapped region.

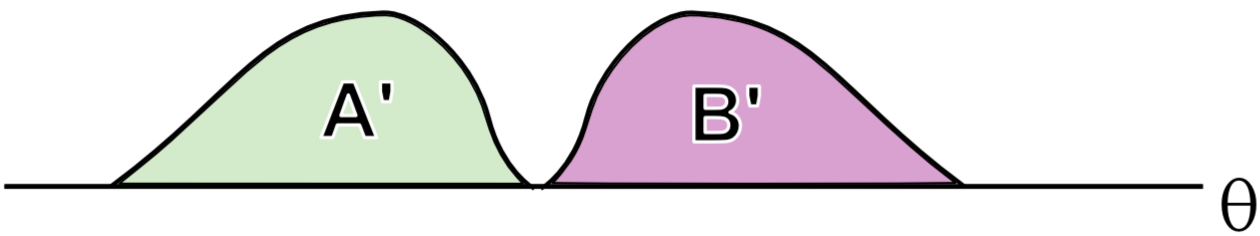

Figure 2.2.2: Remainder distributions A' and B'.

The remainder distributions A'.and B' that are left when the mass of overlapping agents is matched and removed from the market.

b_θ's, and the b_θ's will symmetrically strictly prefer a_θ's to any other agents, so they will form stable matches. ■

Now we can eliminate the stably matched overlap agents from consideration. The remainder distributions can be defined as

$$f_{A'}(\theta) = \max\{\frac{f_A(\theta) - f_B(\theta)}{\int_\Theta \max\{(f_A(\theta) - f_B(\theta), 0\} \, d\theta}, 0\}$$

and

$$f_{B'}(\theta) = \max\{\frac{f_B(\theta) - f_A(\theta)}{\int_\Theta \max\{(f_B(\theta) - f_A(\theta), 0\} \, d\theta}, 0\}$$

where the integral ensures a well-defined probability density function with mass 1. Note that the area under both A and B is the same for both distributions, so the scalar they must be multiplied by is also the same and we don't have any issues of miscounting the measures of agents on each side. Define A' and B' as the distributions with these respective densities.

We now inductively derive a very simple matching algorithm that yields the type s of the match b_{js} for $a_{i\theta}$ (a_{js} for $b_{i\theta}$) as an explicit function depending only on the remainder distributions and $a_{i\theta}$ ($b_{i\theta}$). The intuition here is that we start at the far right of the left remainder distribution (A in this example) and the far left of the right distribution (B in this example), or on other words the innermost points of each distribution, and then iteratively match outward, with the current (innermost remaining) matchers taking the already stably matched interior agents as unavailable. Because agents want the closest match possible, the current matchers on each side strictly prefer the current (innermost remaining) matchers on the other side to anyone else, so they match and the process continues. Note that, because matching is one-to-one, the measure of agents who have been matched on one side must equal the measure of agents who have been matched on the other. Before we complete this proof, we make two additional assumptions.

- **Assumption 2 (SEP)**: Suppose the probability density functions have the single crossing property i.e. the probability density functions intersect at only one point.

- **Assumption 3 (OUT)**: Suppose that agents prefer any match to no match.

Like MASS, SEP and OUT are not necessary for a tractable answer, but they allow for a very simple baseline result to be derived, against which deviations from these assumptions can later be compared.

Lemma 18 *Without loss of generality, assume A' is to the left of B'. Suppose that all agents in the interval (θ_A, θ_B) have been stably matched and are eliminated from consideration, while no other agents in A' or B' have matched. Then a mass of agents* $\min\{f_A(\theta_A), f_B(\theta_B)\}$ *of types θ_A and θ_B will match stably.*

Proof. We know $a_{i\theta_A}$ prefers $b_{j\theta_B}$ to any other $b_{k\theta}$ and vice versa, as they are mutually distance-minimal among the set of remaining potential matches, and all agents that have already been matched are closer to their match then they are to any remaining potential

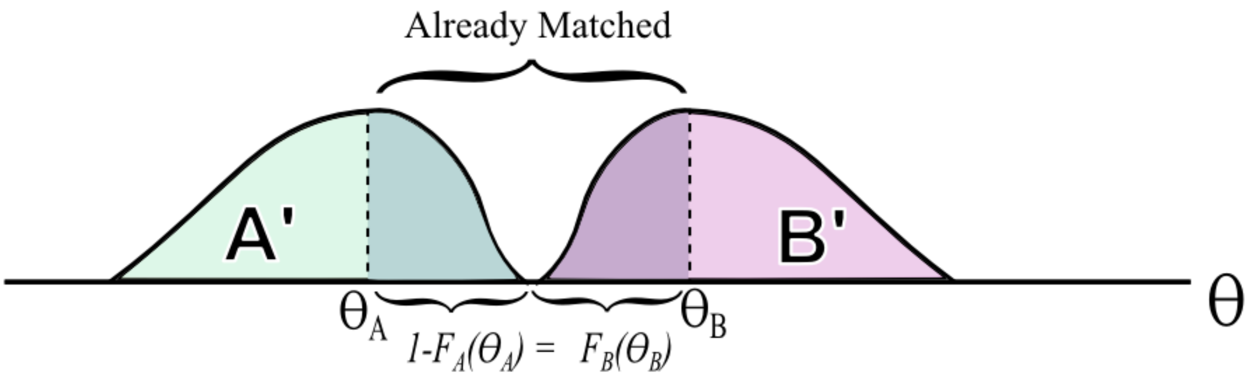

Figure 2.2.3: Inside-Out Matching.

The stage in the matching process when types θ_A and θ_B are the innermost unmatched types. Note that the darkened areas that have already been matched are of equal mass and θ_A and $\theta_B$are mutually closest to one another among the remaining agents.

match by construction, so they will not prefer to deviate to one of the current matchers. Thus the agents will match stably, as was to be shown. ■

We are now ready to present the algorithm and prove its validity.

Proposition 16 *(Homophilic Matching) Suppose MASS, SEP, and OUT. A measure of agents equal to the measure under both curves and with equal density over θ will match to their own type. For all remaining agents of all types θ, θ_A agents match to agents of type $F_B^{-1}(1 - F_A(\theta_A))$ and θ_B agents match to agents of type $F_A^{-1}(1 - F_B(\theta_B))$.*

Proof. The first portion of Proposition 16 is simply Lemma 17. The second is obtained by inductively applying Lemma 18 starting at the innermost points on the two remainder distributions and moving outwards, and by using the fact that the measures of agents matched on each side, $1 - F_A(\theta_A)$ for A and $F_B(\theta_B)$ for B, must be equal. ■

2.2.2 Extensions to the Baseline Model

To get the result above, we made three fairly restrictive assumptions. We will now relax them and find the matching outcome in the more general cases. The matching algorithm is remains quite simple, although relaxing SEP will require a new assumption.

MASS sets the measure of agents on each side of the matching market equal. This is

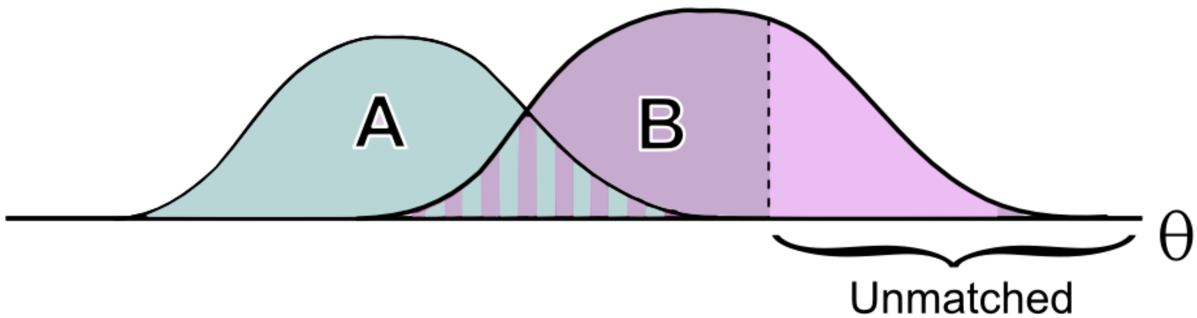

Figure 2.2.4: Matching with unbalanced distributions.

Here B has more mass than A, so the furthest (least attractive) B agents do not receive a match unless they are matched in stage one, the overlap matching phase.

a reasonable assumption in the broader heterosexual dating market, for example, where it is approximately true. However, if one wants to model, say, online dating platforms that attract men and women in disproportionate numbers, or the marriage market in countries that have a significant deficit of men due to war or women due to sex selective abortions, then $MASS$ must be abandoned. All this will do is leave the outermost agents of the larger side unmatched, as the supply of agents on the other side will have run out. Specifically, without loss of generality suppose $\mu_B > \mu_A$ and A is to the left of B. Then the rightmost $\mu_B - \mu_A$ B agents–that is, B agents to the right of $\theta_B = F_B^{-1}(\mu_B - \mu_A)$– will be unmatched, while the rest will match as before.

SEP ensures that all agents of one distribution are above or below all agents of the other. If this does not hold, we may have a situation like Figure 2.2.5 where some A agents are above all B agents and some A agents are below all B agents, and of course much more complicated situations of the same nature could occur. In order to find the matching here, we need to be able to relate preferences for types on the left of the agent to preferences for types to the right of the right of the agent. Assuming we have a utility function or some other means to compare potential matches to the left and right, we can find cutoff agents who are indifferent between their best available match on the right and left. In some cases we can proceed as before from the innermost points on each pair of adjacent "islands" in the two distributions, with indifferent agents determining cutoffs where agents switch from the available match on one side to the match on the other. However, we may run into situations

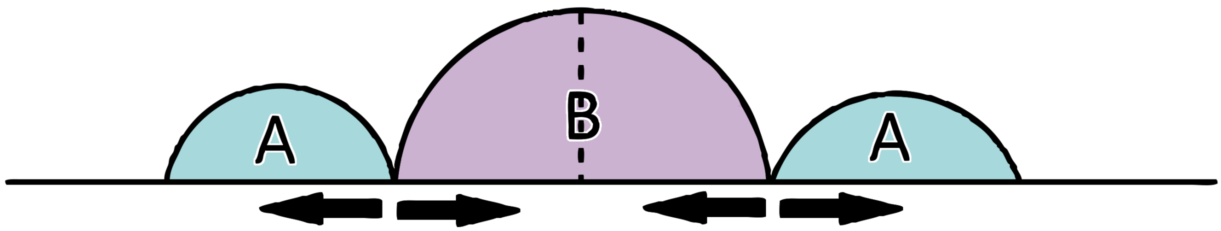

Figure 2.2.5: Matching without the single crossing property.

The remainder of A is separated into two "islands", with the remainder of B in the middle. B agents to the left of the dotted line will match to the left A island, while B agents to the right will match to the right island.

where agents in one island match to agents in a nonadjacent island. This is not the primary focus of this paper, so we will not explore this issue any further.

OUT requires that all agents accept whatever the best match available to them is. However the most obvious qualitative characteristic of the matching outcome in this model is that, for the agents of types that are over-represented relative to that type on the other side, the non-perfect matches quickly deteriorate in quality for fringe agents, as the best remaining match moves further away from them the further to the outside they are. The outermost agents will in fact get their worst possible match, so it seems reasonable that at a certain point agents will prefer no match to a terrible one. The result of dropping *OUT*, assuming that the reservation distance is the same for all types, is simply that matching will terminate once the distance between the innermost remaining agents is equal to the reservation distance, with the rest remaining unmatched.

2.3 Generalization to arbitrary single peaked preferences

2.3.1 Baseline Model

We now allow agent type and agent preference to vary independently, generalizing to arbitrary single peaked preferences. This allows agents to prefer types other than their own. For example, men may prefer women of a different level of femininity than their own, or may prefer someone of a complementary disposition to make up for their shortcomings. Also, two

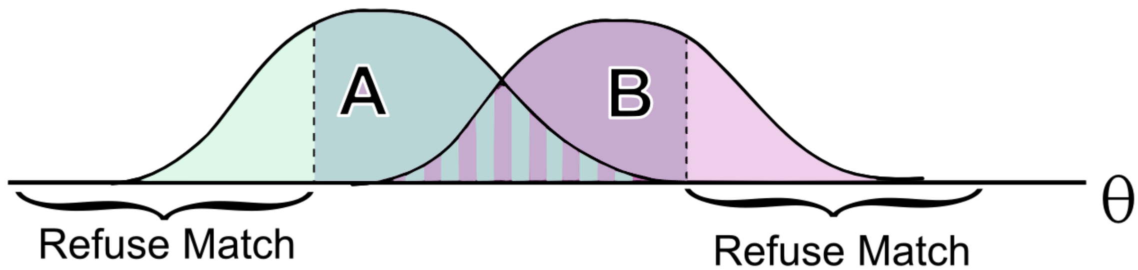

Figure 2.2.6: Matching when agents have the option to refuse.

Agents beyond the distance where the cutoff agents are indifferent between matching or staying single do not clear the matching market.

individuals with the same characteristics may have different preferences over their match's characteristics, rather than e.g. a man's height uniquely determining his height preference. Note that, since the "type" of an individual is now a pair of the form (characteristic, peak preference), so if we continue using the term "type" for an agent's characteristic, we introduce ambiguity. For that reason, an agent's characteristic (e.g. height, BMI, etc.) will now be referred to as their trait.

Now that agents have two characteristics, the set of agents on a given side is a bivariate distribution over own trait and peak preference. Denote an agent i with trait θ and peak preference p as $s_{i\theta p}$. Trait in this situation is the sole characteristic over which an agent's potential matches have preferences, while peak preference determines that agent's most preferred match. To facilitate easy visualization of the algorithm to be derived, we will overlay the distributions A and B, flipping the axes for B. This will put A agent traits and B agent peak preferences on the vertical axis, and B agent traits and A agent peak preferences on the horizontal axis.

The reason for representing the distributions like this is that A agents evaluate matches based on the distance between their preference and a B agent's trait, which is now the horizontal distance between $a_{i\theta p}$ and b_{jsn} on our graph, and B agents evaluate matches based on the distance between their preference and an A agent's trait, which is now the vertical distance between $a_{i\theta p}$ and b_{jsn}, so we can use the graph to easily compare agent preferences over potential matches.

In general, this is a more complicated problem, and no simple formula of great generality

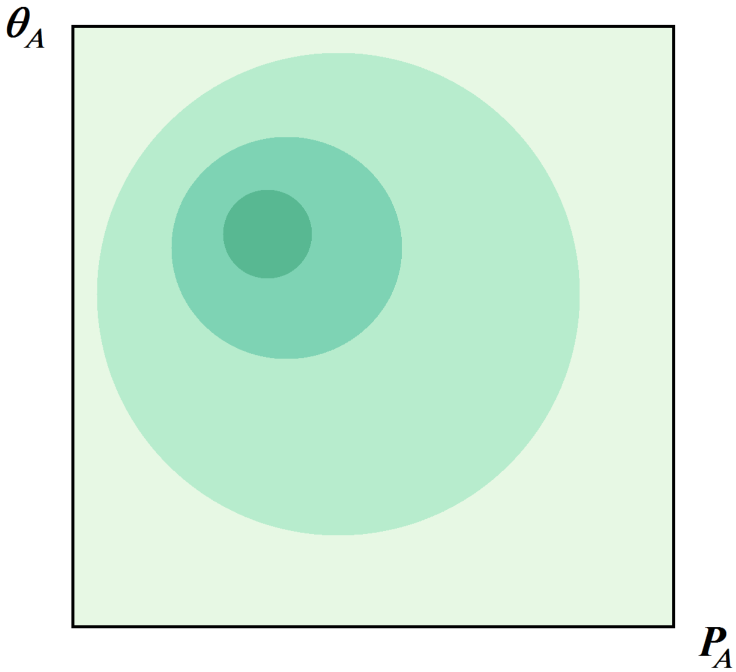

Figure 2.3.1: Contour plot of the trait/preference distribution A.

An A agent's trait runs along the vertical axis and its peak preference over the traits of potential matches runs along the horizontal axis. Darker colors indicate greater mass.

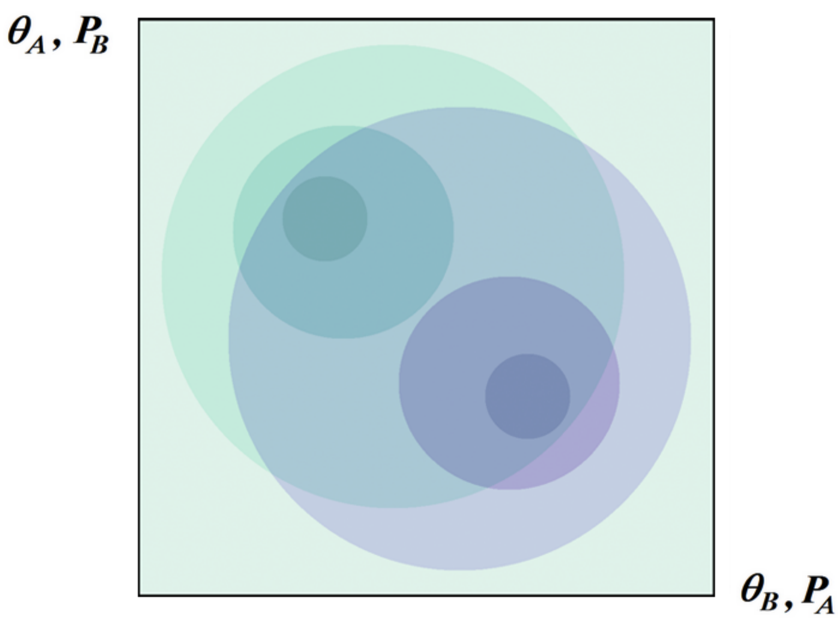

Figure 2.3.2: Overlaid contour plots of distributions A (teal) and B (purple).

Darker teal areas indicate more mass in A, and darker purple areas indicate more mass in B.

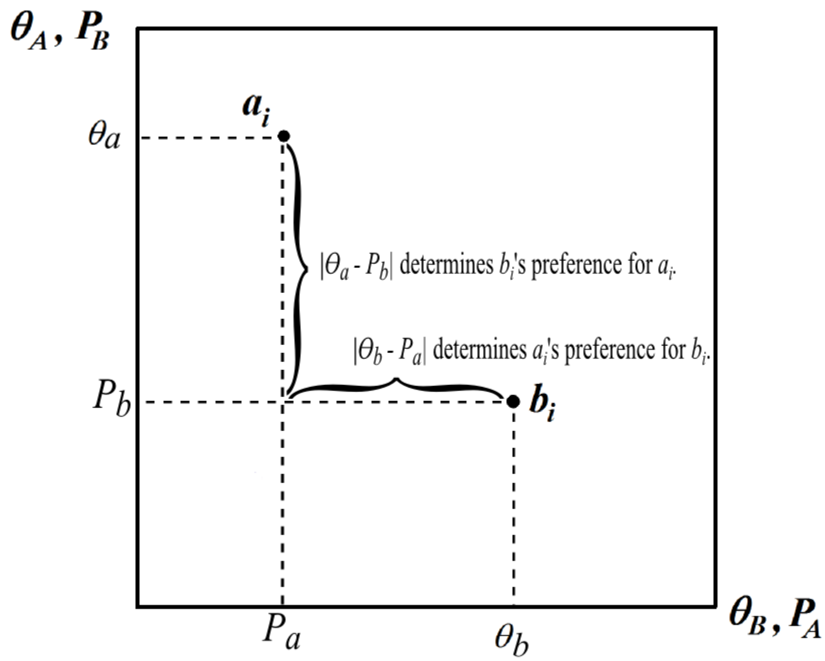

Figure 2.3.3: Comparing matches graphically using the overlay. As shown above, the vertical and horizontal distances between an A and a B agent on this representation tell us the agents' preferences for each other.

will be offered in this paper The agents under both distributions still match to their preferred traits, but the agents in the remainder distributions are more difficult to deal with. Whether a more general simple solution is possible is a topic for further study. However, under specific assumptions on the distributions, the problem is still very tractable. First, we'll make the same assumptions as in Section II, with a slight variation to account for the more general environment. Specifically, we keep $MASS$ and OUT as is and add the following amended assumption:

- **Assumption 2' (SEP')** : Define $h(x) = y$ all $\forall(x, y) \in s$. Suppose the remainder distributions A' and B' are separated by a single curve s. That is, $\forall x$ and $\forall y' > h(x)$, (x,y') has support only on A for all y' or only on B for all y', and $\forall y' < h(x)$, the only the opposite distribution has support at (x,y').

As before, this ensures no complications due to multi-modal distributions, varying tail weights, or flat areas between the remainder distributions. Further sufficient but not necessary conditions for a simple solution are as follows:

- **Assumption 4(CURVE1, CURVE2):**
  - h is monotonically increasing (decreasing) in x.
  - the marginal density of A at preference x equals the marginal density of B at preference y for $(x, y) \in s$.That is, $f_{AX}(x) = f_{BY}(y)$.

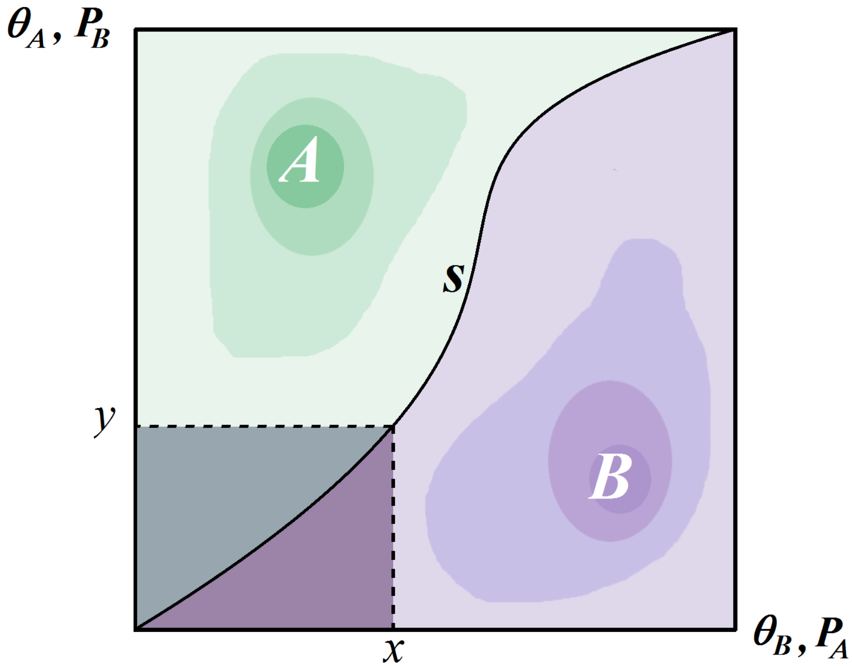

Figure 2.3.4: Matching from the southwest to northeast on the remainder distributions. B agents of trait x match to A agents of trait y, regardless of peak preference.

Under these conditions, the matching of the remainder distributions can be solved by matching in an "unzipping" fashion, where at any stage of the matching process the agents in the southwest (southeast) quadrant of the graph below and to the left (right) of some (x, y) on s have all matched (analogous to the interval which has already matched in the homophilic case), while no one else has, and the agents of interest are those on the edges of the quadrant (analogous to the innermost remaining types in the homophilic case). First, we eliminate the overlapping agents, as before.

Lemma 19 *For each point (x, y), a measure of agents on side $A(B)$ equal to $\mu(x, y) = min\{f_A(x, y), f_B(y, x)\}$, matches to the types on side $B(A)$ with transposed θ and p.*

Proof. Analogous to Lemma 17, there are at least $\mu_{\theta p}$ (θ, p) A agents and (p, θ) B agents by definition. Since preferences are homophilic, $a_{i\theta p}$ weakly prefers $b_{jp\theta}$ to any other B agent, and symmetrically $b_{ip\theta}$ prefers $a_{j\theta p}$ to any other A agent. Then a mass μ_θ of a_θ's will weakly prefer to match with any of the measure μ_θ of b_θ's, and the b_θ's will symmetrically strictly prefer a_θ's to any other agents, so they will form stable matches.

Note that in this case the matching outcome described here may not be the only one possible, since agents don't strictly prefer their mirror agent over agents with their ideal trait but a preference for someone other than them. This can be resolved by using lexicographic preferences where preference over θ are as before and, if two potential matches have the same θ, agents prefer matches whose preferences are closer to their own trait, with the rationale that, if someone likes you better, your relationship with them will generally be better. This gives an outcome with the same perfect matching in the first stage and where agents of a given trait match "inside out" in preference in the second stage, with agents closer to s matching to each other first, and agents further away from s matching later. However, this assumption of preferences presents measure theoretic complications for the solution in the more generalized case in section 3.2, so we will not make this assumption in the following sections. The proof of uniqueness of the stable matching in a finite version of this model is given in the appendix. ■

We now construct the remainder distributions. Define

$$f_{A'}(\theta, p) = \max\{\frac{f_A(\theta, p) - f_B(p, \theta)}{\int_p \int_\theta (\max\{f_A(\theta, p) - f_B(p, \theta), 0\})\, d\theta dp}, 0\}$$

and

$$f_{B'}(\theta, p) = \max\{\frac{f_B(\theta, p) - f_A(p, \theta)}{\int_p \int_\theta (\max\{f_B(\theta, p) - f_A(p, \theta), 0\})\, d\theta dp}, 0\}$$

We are now ready to prove the inductive lemma for this case:

Lemma 20 *Without loss of generality assume A' is to the northwest of B' (separated by a curve as per SEP'). Suppose that all agents with traits and preferences such that $(x, y) < (x^*, y^*)$ have been stably matched and are eliminated from consideration, while no other agents in A' or B' have matched. Then the set of A' agents $\{a_{i\theta p} : \theta = y^*\}$ will match stably and* *arbitrarily to the set of B' agents $\{b_{i\theta p} : \theta = x^*\}$.*

Proof. For each B agent b_{x^*n}, $n \leq y^*$, which is the trait of all A agents on the other edge of the quadrant defined by (x, y), so $a_{y^*p} \underset{b_{x*m}}{\succ} a_{\theta p}$, $\forall$ unmatched agents $a_{\theta p}$ where as $\theta > y^*$. A symmetric argument shows that all a_{y^*p} strictly prefer b_{x^*n} agents to any other unmatched B agents, so the edge agents of both distributions will match to one another. Since $f_{AX}(x) = f_{BY}(y)$, these sets of agents have equal measure, so they exactly and stably match to one another, leaving no remaining b_{x^*n} or a_{y^*p}. ■

Proposition 17 *(Single Peaked Two-sided Matching I) Suppose $MASS$, SEP', OUT, and $CURVE$. A measure of agents equal to the measure under both distributions and with equal density over (x, y) will match to their preferred type, which also finds them optimal. For agents in the remainder distributions A' and B' and for all $(x, y) \in s$, the set of A' agents $\{a_{\theta p} : \theta = y\}$ will match stably and arbitrarily to the set of B' agents $\{b_{\theta p} : \theta = x\}$.*

Proof. The first result is simply Lemma 17. Inductively proceeding with Lemma 2 northeast along s, we have that a_{yp} matches with any b_{xn} and vice versa $\forall (x, y) \in s$. Because the marginal densities are equal along this path, the measure of matched agents at any point in the inductive process is identical for both sides, so we don't violate the necessary condition of $1 - 1$ matching. ■

It's worth noting that this marginal density assumption is very important. If we did not have $CURVE2$ and tried to proceed as above, we'd have unequal measures of agents being matched at various points in the matching process, a clear contradiction. In fact, what would happen is that each "layer" of A' agents would not completely match out the corresponding layer of B' agents, and the remaining B' agents would match to the next layer of A' agents. We would then no longer be in the the extremely convenient situation where the current matchers are all of one trait and where every current matcher prefers the edge agents on the other side to any other available agent. Similarly, dropping CURVE1 would invalidate the procedure, with, for example, agents on one side matching to no one on the other side when the slope of s was negative.

We can also obtain an algorithm for the matching in a one-sided problem with single-peaked preferences from the two-sided algorithm by representing the one-sided problem as a two sided problem. Recall that a one-sided matching problem is one where there is a single set of agents who must be matched to one another in any way that satisfies stability, whereas the two sided problem imposes the additional constraint that agents can only match to individuals on the opposite side.

Corollary 4 *(Single Peaked One-sided Matching) For any distribution$f(p,\theta)$, define $f_A(p,\theta) = f_B(p,\theta) = f(p,\theta)$. Then if $f_A$and f_B* satisfy *$MASS$, SEP', OUT, and $CURVE$ and $s = (x,x)\forall x \in \mathbb{R}$, Proposition 17 holds. Equivalently, the one-sided matching problem with distribution f has the stable matching given by Proposition 17, where the match of a given agent $a(b)$ is inf rather than $f_B(f_A)$.*

Proof. The first claim follows directly from Proposition 17, as it is just a special case of the problem considered there. For the second claim, consider the one-sided matching problem with single peaked preferences and distribution f. An agent a with trait θ_a and preference p_a prefers matches b based on . Similarly, b prefers matches based on$|p_b - \theta_a|$. Then if we overlay $f_A = f$ with an axes-transposed copy of f, f_B, a's preferences over f are given by the vertical distance, and b's preferences over f are given by the horizontal distance. First, we remove agents with perfect matches, so a mass $min(f(p,\theta), f(\theta,p))$ is matched to its ideal match for each (θ,p), then we move on to the iterative stage. f, at every stage of the matching process, the set of unmatched agents on each side in the two-sided problem is equal to the set of unmatched agents in the one-sided problem, where all the matches thus far derived in the two-sided problem are stable in the one sided problem, then the current matches are optimal among the set of available matches for all agents on both sides, so they are optimal in the one-sided problem. There is one complication here–when agent a is matched to agent b in the two-sided problem, the a is removed from side A and b is removed from side B. However, in the one-sided problem both a and b are on the same side. If some matched agents are not removed from each side, the set of available matches will not correspond to the one-sided problem. However, because the distributions are identical, the remainder distributions are also identical by their definitions, and $x = y\ \forall (x,y) \in s$, the set of current matches is identical, and agents are indifferent between all possible matches in this set. Since there are infinitely many agents at every point with nonzero support, we can always have half the agents of a given type match to the other half, leaving no agents unmatched and all agents with their preference maximal match among the set of remaining agents.

Note that this yields a very simple matching outcome where as many agents get perfect matches as possibly can and the remainder match to their own trait (positive assortation). ■

2.3.2 Extensions to the Baseline Model

While the result in 3.1 is extremely simple, the assumptions, especially CURVE2, are unlikely to be even approximately satisfied in a real world application. Having the marginal densities equal at any particular point on s is unlikely, much less at every point. First, then, we will relax this assumption. This significantly complicates the problem, but does not render it insoluble. Without loss of generality, assume that A and B are separated by a monotonically decreasing h, with A above and to the right and B below and to the left. First, find all points $(x_i, y_i) \in s$ such that $1 - F_{B'Y}(y_i) = F_{A'X}(x_i)$ for $i \in \{1, ..., n\}$ (assume there are finitely many such points). Then, $\forall (x, y) \in s$ where $(x_i, y_i) \leqq (x, y) \leqq (x_{i+1}, y_{i+1})$, $(x_n, y_n) \leqq (x, y)$, or $(x, y) \leqq (x_1, y_1)$,$1 - F_{B'Y}(y) \leqq F_{A'X}(x)$ or $1 - F_{B'Y}(y) \geqq F_{A'X}(x)$. Without loss of generality, suppose $1 - F_{B'Y}(y) \leqq F_{A'X}(x)$. Then, with the following amended assumption, we can proceed to a matching solution.

- **Assumption 4' (CURVE2'):**
  - When $A'(B')$ has a larger mass matched out, the marginal density of $A'(B')$ at trait $x(y)$ is greater than or equal to the marginal density of $B'(A')$ at trait $y(x)$ for preferences greater than or equal to y(x), where y(x) is such that we have masses $m_a = m_{b1} + m_{b2}$ ($m_b = m_{a1} + m_{a2}$). That is, $f_{A'X}(x) \geqq \int_y^\infty f_{B'}(x, p)dp$ ($f_{B'Y}(y) \geqq \int_{-\infty}^x f_A(y, p)dp$), .

This assumption ensures that, for example, there is never more mass in β_1 in Figure 2.3.5 than in α, which would invalidate the matching algorithm since, as x moved outward as the matching progressed down and to the right, more mass would be matched in B than in A, even if y didn't decrease at all.

Finally, before we state the proposition, we'll need a definition and two equations.

Definition 6 *Define $M_A(x)$ as the set of B agents that an A agent of trait x can stably match to. That is. $M_A(x) = \{b_{j\theta p} : \theta = y\} \cap \{b_{j\theta p} : p = x \wedge \theta \geqq y\}$. Similarly, $M_B(y) = \{a_{j\theta p} : \theta = x\} \cap \{a_{j\theta p} : p = y \wedge \theta \leqq x\}$.*

We also define two equations guaranteeing equal masses of agents have been matched out at each step (this is equivalent to the equal masses condition in CURVE2').

$$\int_{-\infty}^{\infty} \int_{-\infty y}^{x} f_{A'}(\theta, p) d\theta dp = \int_{-\infty}^{x} \int_{y}^{\infty} f_{B'}(\theta, p) d\theta dp \tag{2.3.1}$$

$$\int_{-\infty}^{\infty} \int_{y}^{\infty} f_{B'}(\theta, p) d\theta dp = \int_{y}^{\infty} \int_{-\infty}^{x} f_{A'}(\theta, p) d\theta dp \tag{2.3.2}$$

Proposition 18 *(Single Peaked Two-sided Matching II) Suppose $MASS$, SEP', OUT, and $CURVE'$. A measure of agents equal to the measure under both distributions and with equal density over (x, y) will match to their converse type, who also finds them optimal. For agents in the remainder distributions $A^{'}$ and $B^{'}$ and for all $(x, y) \in s$, if $1 - F_{B'Y}(y) \leqq F_{A'X}(x)$ A' agents $\{a_{i\theta p} : \theta = x\}$ will match stably and arbitrarily to elements of the set of B' agents $M_A(x)$ and vice versa, where y satisfies Eqn. 3.1 If $1 - F_{B'Y}(y) > F_{A'X}(x)$ B' agents $\{b_{i\theta p} : \theta = y\}$ will match stably and arbitrarily to elements of the set of A' agents $M_B(y)$ and vice versa, where x satisfies Eqn. 3.2.*

Proof. Then A agents of trait x will match to b agents of trait y and $p \leqq x$ or b agents where$h^{-1}(x) \leqq p \leqq y$ and trait x. To check that this is stable, consider a_1, b_1 and a_2, b_2 matching this way, where $x_2 > x_1$and consequently $y_2 < y_1$. Without loss of generality, either (1) $1 - F_{B'Y}(y) \leqq F_{A'X}(x)$ for both 1 and 2, or (2)$1 - F_{B'Y}(y) \leqq F_{A'X}(x)$ for pair 1 and not for pair 2. If (1), $a_1 \succ_{b_1} a_2$ since $p_{b_1} \leqq \theta_{a_1} < \theta_{a_2}$. Then we need only consider the potential $a_1 - b_2$ blocking pair. $a_1 \succ_{b_2} a_2$ only if $p_{b_2} < \theta_{a_2}$, else a_2 is b_2's perfect match. But if so then the trait of b_2 is y_2. Then $b_1 \succ_{a_1} b_2$ as $y_2 < \theta_{b_1} < p_{a_1}$. If (2), $b_1 \succ_{a_1} b_2$ since $p_{a_1} > \theta_{b_1} > \theta_{b_2}$ and $a_1 \succ_{b_1} a_2$as $p_{b_1} \geqq \theta_{a_1} > \theta_{a_2}$, so there is no blocking pair. Finally, if both pairs come from the same stage in the matching algorithm, either the b's are indifferent between the a's or the a's are indifferent between the b's, so there is no blocking pair. ■

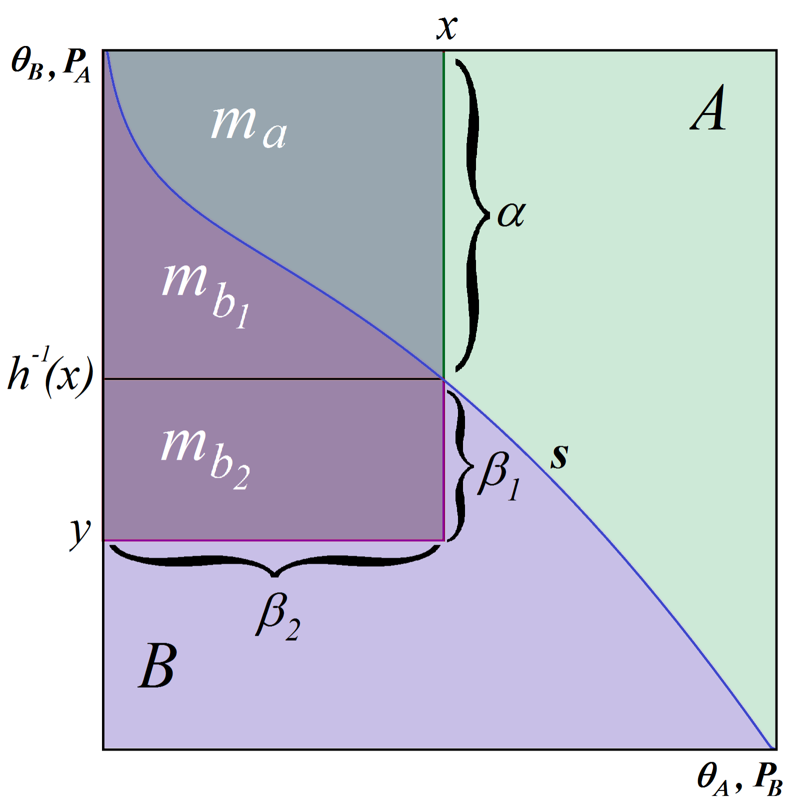

Figure 2.3.5: Matching with unequal marginal distributions.

A has more mass in m_a than B has in m_{b1}, so agents in m_{b2} have also been matched to equalize the mass on both sides.

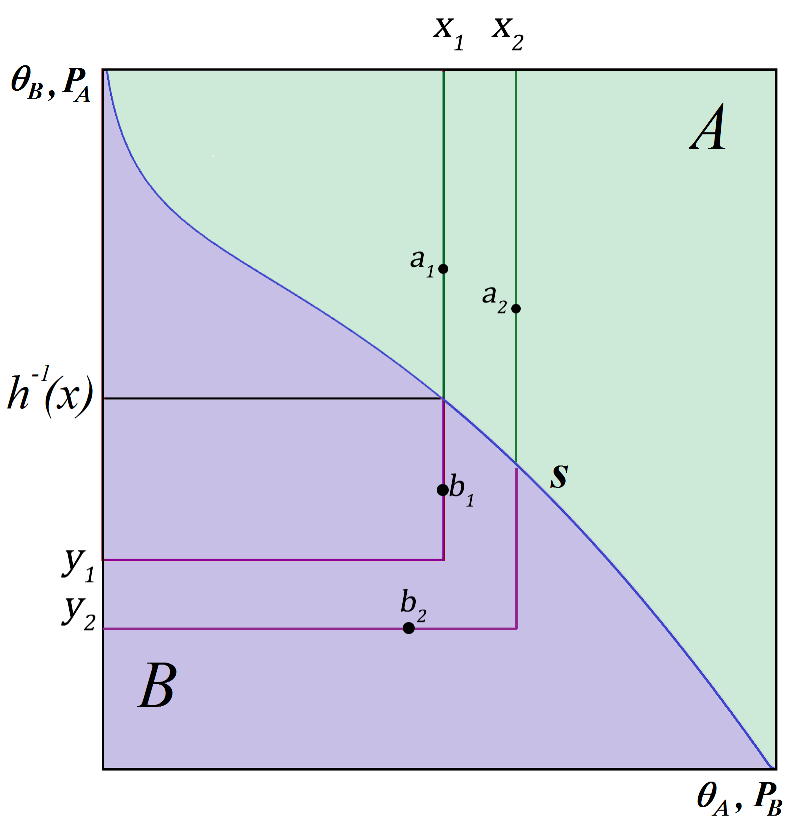

Figure 2.3.6: Proposition 18, Case 1 Example.

$1 - F_{B'Y}(y) \leqq F_{A'X}(x)$ for both(x_1, y_1) and (x_2, y_2).

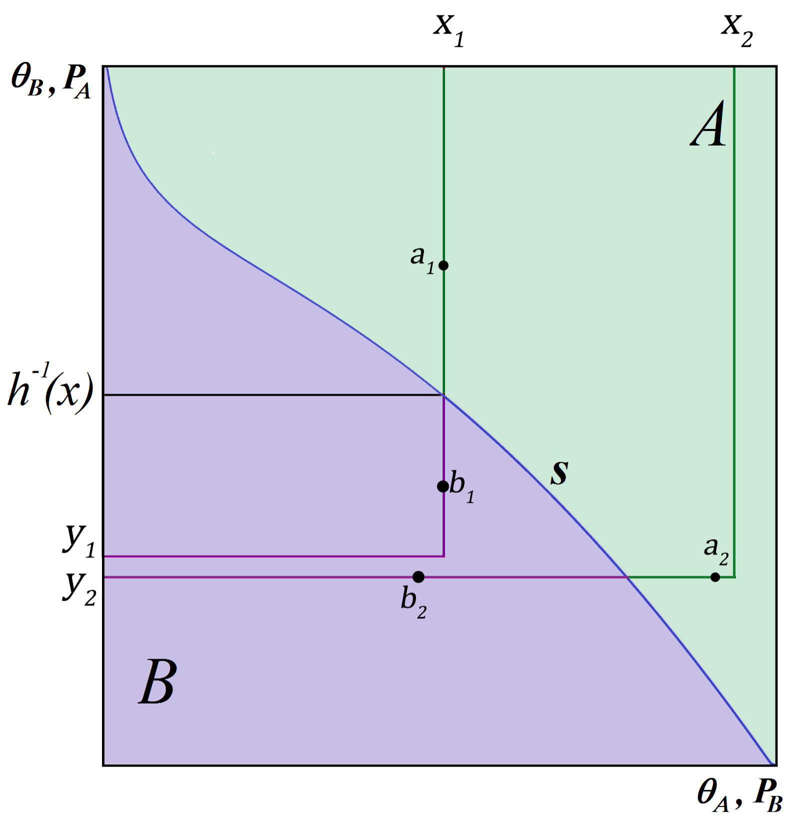

Figure 2.3.7: Proposition 18, Case 2 Example.

$1 - F_{B'Y}(y) \leqq F_{A'X}(x)$ for (x_1, y_1) only.

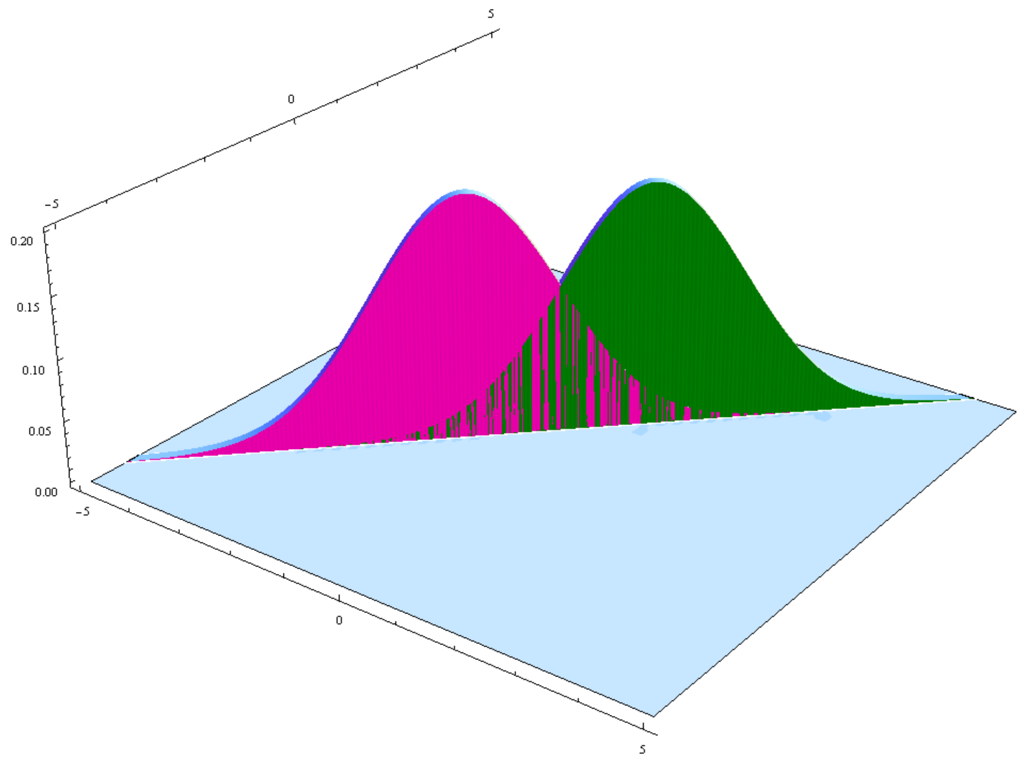

Figure 2.4.1: Horizontal preferences in the single peaked preference framework.

2.4 Relationship to the Literature

As mentioned previously, there are well known results for two-sided matching with vertical preferences, and Clark 2007(**?**) gives the results for horizontal preferences shown in section 2. It is obvious from the previous exposition that horizontal preferences are a special case of single-peaked preferences–specifically, they are the case where preference is set equal to trait. Then with these preferences, we should find that the single peaked algorithm reduces to the horizontal preference algorithm. In fact this is the case. In the horizontal preference case, the distributions have support only on the diagonal, where preference equals trait.

From here, we remove the overlap and can now easily draw a monotonically decreasing curve $s = (x, h(x))$ that separates the two sides, and by letting $h(x) = F_B^{-1}(1 - F_A(x))$, we have the appropriate matches and an equal mass of matched out agents at every step in the process, as desired.

Similarly, Becker's NTU model with vertical preferences is also a special case of single-peaked preferences, namely, when everyone's preference is for higher traits. If there is a

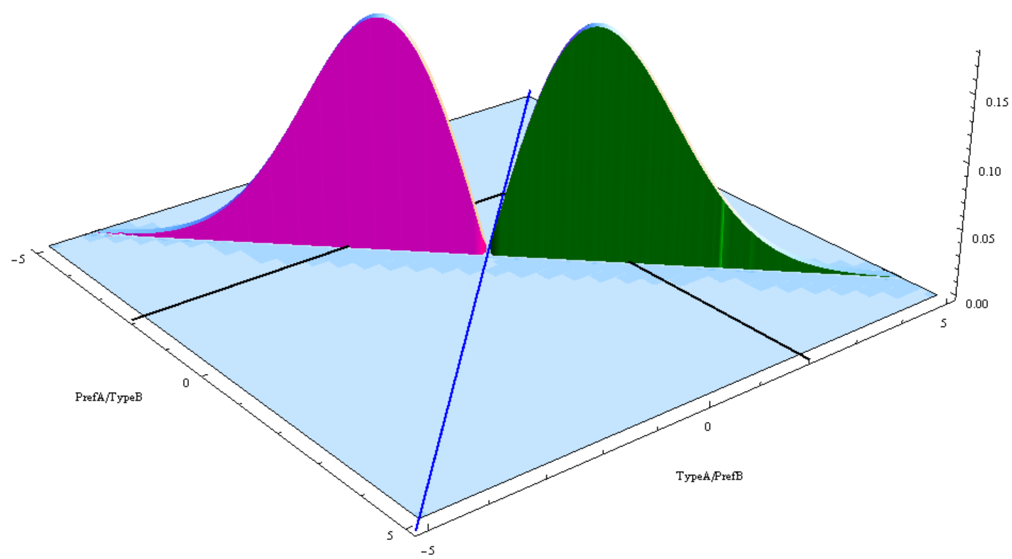

Figure 2.4.2: Horizontal matching in the single peaked preference framework

maximal trait $\bar{\theta}$, we can simply set preferences to $\bar{\theta}$. This affords a simple graphical representation, with the distributions varying along trait with support only at preference $\bar{\theta}$. Choosing$s = (x, h(x))$ and letting $h(x) =F_B^{-1}(F_A(x))$, we have the appropriate matches and an equal mass of matched out agents at every step in the process and the two distributions are separated by s, and we have the familiar positive assortative matching for vertical preferences with nontransferable utility.

2.5 Interpretation and Empirical Implications

2.5.1 Interpretation

These models are amenable to some interpretation. While in the vertical case we have PAM, and in the horizontal case we have PAM in the overlap and NAM in the remainders, in the equal-marginals case of single peaked preferences, we have two modes of matching that encompass these previous cases. First, we have CPAM over trait and preference in the overlap region, where we have positive assortation in A trait and B preference and in B trait and A preference–that is, increase in one parameter corresponds to increase in the other

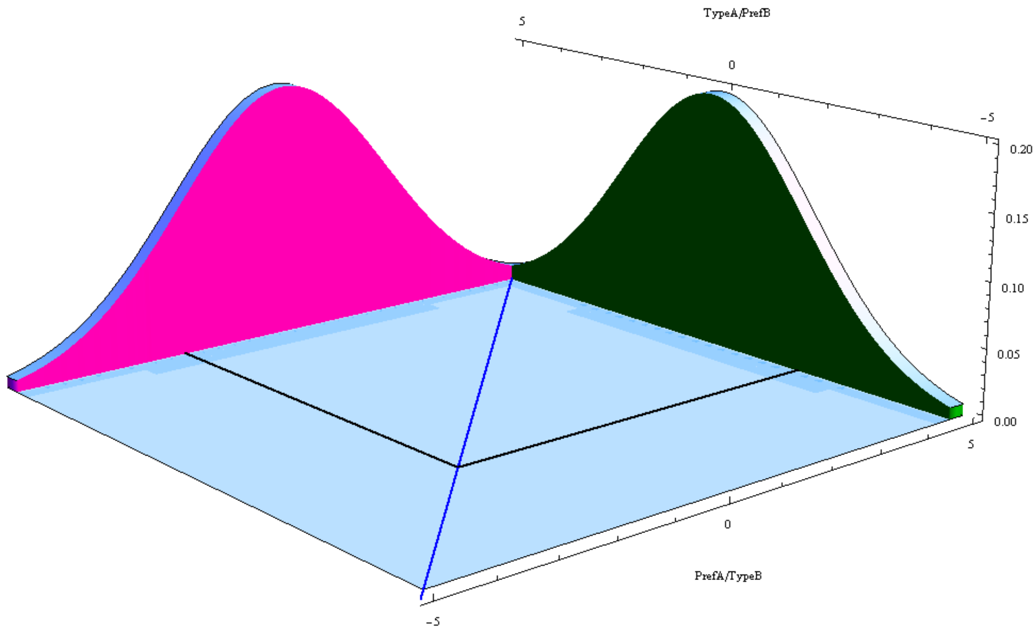

Figure 2.4.3: Vertical preferences in the single peaked preference framework.

parameter in one's match. Note that this is more than just PAM–the matches have exactly reversed trait and preference. We can also see that the standard PAM of the horizontal model is actually a special case of this CPAM, where only the fact that trait and preference are equal ensures that A trait equals B trait. For the remainders, we see PAM when the separating curve is increasing, which is to say that agents generally prefer higher types than their stable matches. We see NAM when the separating curve is decreasing, which is to say that agents generally prefer lower types than their stable matches. This again corresponds to the horizontal and vertical cases, with the line in the vertical case having a positive slope, while the horizontal case has a negative slope.

The more general model of proposition 18 is a bit more complicated, but also yields an intuitive interpretation. Without loss of generality, assume a strictly decreasing h and A at the top right, with higher mass on side A as in Figure 2.3.5. In this situation, B agents can be thought of as being in shortage at $(x, h^{-1}(x))$ in Figure 2.3.5, as there aren't enough of them to match to A agents as in the equal marginals outcome. As such, and given that matching utilities for each agent in a match need not correlate in any way, we'd expect that B agents will often be able to leverage this scarcity to get a matching outcome more favorable to their side, and this is in fact the case. Notice that the entire region β_1 in Figure

2.3.5 gets perfect matches (from their perspective), while β_2 agents have a similar matching outcome to B agents in the equal marginals case. A agents, on the other hand, match to B agents whose trait is further from their preferences than the agents with trait $h^{-1}(x)$ that they'd match to in the equal marginals case, so they are worse off. In terms of assortation, we see that, for the overlap region the result is the same as the equal-marginals case, while in the remainders we have some of the agents exhibiting PAM or NAM in trait as before (e.g. agents matching from α to β_2), while the shortage agents get CPAM.

2.5.2 Empirical Implications

The empirical implications of this model are also fairly straightforward, though they vary based on the orientation of the distributions. Again, without loss of generality assume a strictly decreasing h and A at the top right, with higher mass on side A as in Figure 2.3.5. As a's trait increases, the matching function that gives the distribution of possible matches exhibits distributions where the maximum preference of matches is equal to a's trait and thus increasing (β_2 in Figure 2.3.5 is the set of possible matches with preferences less than a's trait). Note that all potential matches for an A agent with trait x must have the same trait, y, unless their preference is equal to a's trait, and that this y is decreasing in x. The distribution of possible matches also includes B agents whose trait is greater than y but less than the minimum preference of a agents with trait x in A', β_1 in Figure 2.3.5. This region's upper and lower bounds in B agent trait decrease in x, while the B agent's preference is of course increasing in x as it is equal to x. Finally, we can expect a mass point of perfect matches where preference and trait are reversed from the first stage of the matching process. If we only observe trait, we would expect to see a mass point at the minimum trait a matches to, and as a's trait increases, we would expect that mass point to move downward. Analogous predictions can be recovered for other orientations and relative surpluses.

2.6 Conclusions

This paper derived an algorithm for finding matching outcomes in a generalization of several environments that have previously been explored in the matching literature, namely Becker's vertical model and Clark's horizontal model. By allowing for a wide variety of single peaked distributions, this algorithm can be used to explore matching behavior in a much richer environment.

There are several plausible extensions for this model which will not be explored fully in this paper but may be worth further consideration. For example further study could include determining how these models relate to analogous finite matching models or models with search.

2.7 Appendix

2.7.1 Uniqueness of the Stable Matching

As noted earlier, the stable matchings derived in this paper may not be unique. However, with the additional assumption of lexicographic preferences where agents first prefer match traits close to their peak preference, then secondarily prefer match preferences close to their trait, we can pursue a proof of uniqueness for the equal marginals case.

- **Assumption 5 (LEX)** : Without loss of generality, consider side A. For any A agent a_τ of type τ and B agents $b_{\theta i}$ and $b_{\theta j}$ where $\tau < i < j$ $(j < i < \tau)$,

  $b_{\theta i} \underset{a_\tau}{\succ} b_{\theta j}$ $(b_{\theta j} \underset{a_\tau}{\succ} b_{\theta i})$

There is an additional complication, however–the proof technique used here only works with finite sets of agents, so we must restrict ourselves to a finite version of the model analyzed above. It may well be possible to extend the proof to infinite case, but the measure theoretic complications haven't been resolved as of this writing. First, we need to prove the equal marginals proposition in this environment.

Lemma 21 *Suppose LEX. For each point (x, y), a measure of agents on side $A(B)$ equal to $\mu(x, y) = min\{f_A(x, y), f_B(y, x)\}$, matches to the types (y, x) on side $B(A)$.*

Proof. There are at least μ_{xy} (x, y) A agents and (y, x) B agents by assumption. $a_{i\theta p}$ strictly prefers $b_{jp\theta}$ to any other B agent, and symmetrically $b_{ip\theta}$ strictly prefers $a_{j\theta p}$ to any other A agent. Then a mass $\mu_{\theta p}$ of $a_{\theta p}$'s strictly weakly prefer to match with any of the mass $\mu_{\theta p}$ of

$b_{p\theta}$'s, and the $b_{p\theta}$'s will symmetrically strictly prefer $a_{\theta p}$'s to any other agents, so they will form stable matches. ■

We are now ready to prove the inductive lemma for this case:

Lemma 22 *Without loss of generality assume $A^{'}$ is to the northwest of $B^{'}$ (separated by a curve as per SEP'). Suppose that all agents with traits and preferences such that $(x, y) < (x^*, y^*)$ have been stably matched and are eliminated from consideration, while no other agents in $A^{'}$ or $B^{'}$ have matched. Then the set of A' agents $\{a_{i\theta p} : \theta = y^*\}$ will match stably to the set of B' agents $\{b_{i\theta p} : \theta = x^*\}$, where for matching pair (a,b), $p_b = f^{-1}_{B'T}(f_{A'X}(p_a))$.*

Proof. Note that the second order preferences don't change agent preferences for matches of different traits. They are only relevant when considering two agents of the same trait. Thus the types of agents in every match will remain the same as before, as proven in Lemma 20. For a given (x,h(x)), notice that this problem mirrors the horizontal matching problem of Section 2. Then by Lemma 18, we have that the innermost agents will match to each other at each step, and in order to ensure an equal mass matched out at every step, we must have $p_b = f^{-1}_{B'T}(f_{A'X}(p_a))$. ■

Proposition 19 *(Lexicographic Single Peaked Two-sided Matching I) Suppose $MASS$, SEP', OUT, and $CURVE$. A measure of agents equal to the measure under both distributions and with equal density over (x, y) will match to their preferred type, which also finds them optimal. For agents in the remainder distributions $A^{'}$ and $B^{'}$ and for all $(x, y) \in s$, the set of A' agents $\{a_{\theta p} : \theta = y\}$ will match stably to the set of B' agents $\{b_{\theta p} : \theta = x\}$, where for matching pair (a,b), $p_b = f^{-1}_{B'T}(f_{A'X}(p_a))$.*

Proof. The first result is simply Lemma 17. Inductively proceeding with Lemma 2 northeast along s, we have that a_{yp} matches with any b_{xn} and vice versa $\forall (x, y) \in s$. Because the marginal densities are equal along this path, the measure of matched agents at any point in the inductive process is identical for both sides, so we don't violate the necessary condition of $1 - 1$ matching. ■

Define a member of the family of matchings given by Proposition 19 as L. Define the type of a match for agent $a_i(b_i)$ under assignment M as as $m_M(a_i)$ $(m_M(b_i))$

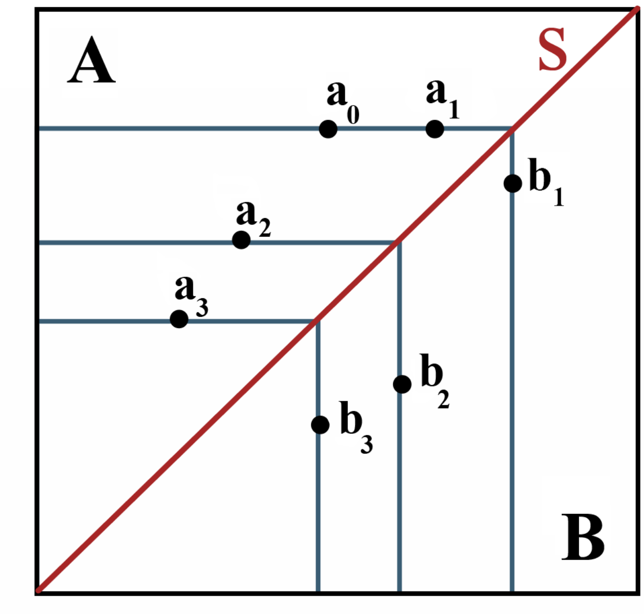

Figure 2.7.1: A sequence of matching pairs under L and a_0, b_1's assignment under M.

Proposition 20 *(Single Peaked Two-sided Matching I Uniqueness) Suppose the conditions of Proposition 17 are satisfied and A and B are finite sets. Then L is the unique stable matching up to agents with identical type.*

Proof. Define b_i as $m_L(a_i)$. We'll suppose a stable matching M exists such that $m_L(a_i)$ has a different type vector than b_i for some i. First consider stage one. Every agent gets their unique (in terms of type vector) ideal match, which also considers them ideal. Then no agents in either stage 1 or stage 2 could form a blocking pair with a stage 1 agent, since the stage 1 agent strictly prefers their current match. Then any stable matching M cannot differ from L for these agents). Now consider stage 2. If the matching M differs from L in this stage, at least one agent has a different match, and their new match in M also has a new match relative to L. Suppose there exists $b_1 \in B'$ such that $m_M(b_1) = a_0$ and $\theta_{a_0} > \theta_{a_1}$ or $\theta_{a_0} = \theta_{a_1}$ and $P_{a_0} < P_{a_1}$. If not, a_0 is such that $\theta_{a_0} < \theta_{a_1}$ or $\theta_{a_0} = \theta_{a_1}$ and $P_{a_0} > P_{a_1}$. Thus b_1 strictly prefers a_0 to a_1. Then a_0 must strictly prefer their match under L, $b_0 = m_L(a_0) \neq b_1$, to b_1 or (b_1, a_0) from a blocking pair in L, contradicting its stability. In this case we could proceed with the following argument with a_0 rather than b_1. Then without loss of generality, assume $\exists b_1 \in B'$ as before. Since b_1 matches to a_0, despite preferring a_1, it must be that some a_1 type agent matches to some b_2 such that $\theta_{b_1} > \theta_{b_2}$ or $\theta_{b_1} = \theta_{b_2}$and $P_{b_1} < P_{b_2}$, else there would be a blocking pair (a_1, b_1) in M since the only agents a_1 would be indifferent to are other $b_1$agents, and one of them is matching to a_0 so there are not enough remaining b_1 agents to match to all a_1 agents. We see that b_2 is matching to a_1, despite preferring a_2. From here, we can proceed by induction–for any t, if b_t is such that $m_M(b_t) = a_{t-1}$ and $\theta_{a_{t-1}} > \theta_{a_t}$ or $\theta_{a_{t-1}} = \theta_{a_t}$ and $P_{a_{t-1}} < P_{a_t}$, a_1, it must be that some a_t type agent matches to some b_{t+1} such that $\theta_{b_t} > \theta_{b_{t+1}}$ or $\theta_{b_t} = \theta_{b_{t+1}}$ and $P_{b_t} < P_{b_{t+1}}$. But then $m_M(b_{t+1}) = a_t$ and $\theta_{a_t} > \theta_{a_{t+1}}$ or $\theta_{a_t} = \theta_{a_{t+1}}$ and $P_{a_t} < P_{a_{t+1}}$. Since there are finitely many agents in A and B, we will find a t such that, e.g. a_t must match to some b_{t+1} such that $\theta_{b_t} > \theta_{b_{t+1}}$ or $\theta_{b_t} = \theta_{b_{t+1}}$and $P_{b_t} < P_{b_{t+1}}$, but there are no remaining agents in B that satisfy this condition, so we have a blocking pair. ■

CHAPTER 3

MATCHING MARKETS WITH N-DIMENSIONAL PREFERENCES

3.1 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 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 preferences[1], and even more general single peaked preferences[2], 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.[3] While research on assortation 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 assortation among all traits simultaneously. Under a special case of symmetry, it predicts *positively assortative matching*[4] (PAM) along all but one trait, and *negatively assortative matching*[5] (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 assortation 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

[1]For example, there are several ethnic categories and agents prefer their own category.

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

[3]In one dimension, assortation is a matching structure where the type of an agent's match is monotonic in the agent's own type.

[4]PAM means an agent's match's type is monotonically increasing in the agent's own type.

[5]NAM means agent's match's type is monotonically decreasing in the agent's own type.

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 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 environments[6]. 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

[6]Environments that satisfy the assumptions of this paper and those of e.g. Lauermann and Nï¿œldeke (2014).

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.

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, Hortaï¿œsu, 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.

3.2 Theoretical Preliminaries

3.2.1 The Model

The environment we'll be considering is a matching market with two sides, or sets of agents, A and B[7]. 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 $b \underset{a}{\succ} b'$ we'll say a strictly prefers b to b', and if $b \underset{a}{\succsim} b'$ 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

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

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 $u(a, b)/2$. 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 $b \underset{a}{\succ} \mu(a)$ and $a \underset{b}{\succ} \mu(b)$. Such an (a, b) is a 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 $v : A \cup B \to \mathbb{R}$ 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 ith 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.

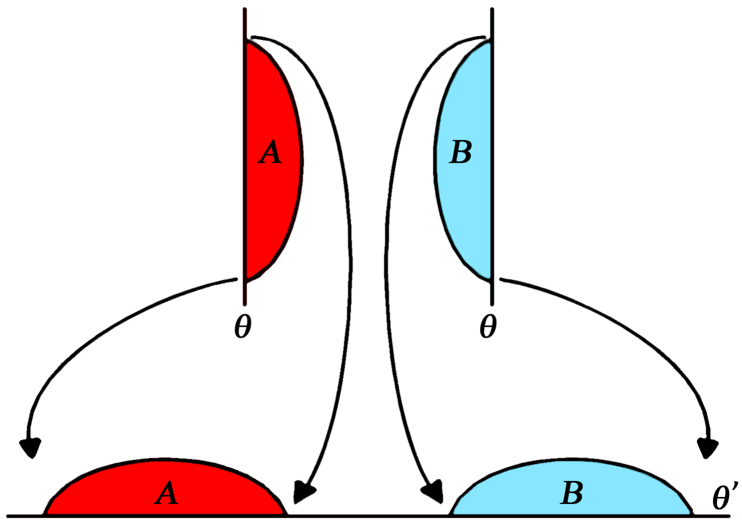

Figure 3.2.1: Mapping vertical preferences to a horizontal model.

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.

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

3.2.2 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.[8] Luckily, while the horizontal preference assumption

[8] A 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.

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.

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 3.2.3. 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. [9]

3.2.3 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

[9] Note 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.

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 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 parameters[10] 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 salient[11]. Thus, we'll now consider the problem of explicitly solving multivariate matching problems.

[10]e.g. a price of entry into the matching platform.

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

3.3 Theoretical Results for Symmetric Distributions

3.3.1 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, 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 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'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, Hortaï¿œsu, 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 3.2. 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 the 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 3.4.1. 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 21 *(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 21 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 $d(a_1, b_1) = 2d(a_1, h) = d_\eta(a_1, h)$, $d(a_2, b_2) = 2d(b_2, h) = d_\eta(b_2, h)$. 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

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

$$= (d(a_1, b_1) + d(a_2, b_2))/2$$

$$\geq min\{d(a_1, b_1), d(a_2, b_2)\}$$

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 assortation for multiple dimensions. Define a_i as the value of the ith trait of agent a, and $a_{\neg i}$ as the vector of a's traits excluding i.

Definition 7 *(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 assortation 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 $a_{\neg i}$ vector ($a_{\neg i} = a'_{\neg i}$). Then we could have, for example, PAM in income for low education individuals and NAM in income for high education individuals. Definition 7, by contrast, requires a much stronger notion of assortation.

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

Corollary 5 *(n-Dimensional Assortation)* for each i, μ_i either satisfies unconditional PAM, satisfies unconditional NAM, or, for each $a_{\neg i}$, all $(\cdot, a_{\neg i})$ type agents match to B agents of type $(b_i, \cdot)$ for some b_i.

While the direction of assortation 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 21:

Corollary 6 *(Normal n-Dimensional Assortation) Where defined,* $\dfrac{\partial \mu(a)}{\partial a_1} = -e_1$, $\dfrac{\partial \mu(a)}{\partial a_i} = e_i$ *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 $i = 1$. ■

Corollary 6 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 6 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 assortation 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 6 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 6 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.

3.3.2 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.[12]

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.

[12]Becker 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.

Proposition 22 *(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 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 3.7.1.2 ■

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.

3.3.3 Matching with Asymmetric Distributions

Propositions 21 and 22 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 assortation in the symmetric case are still at work in an asymmetric market. In an asymmetric case like Figure 3.3.1, 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

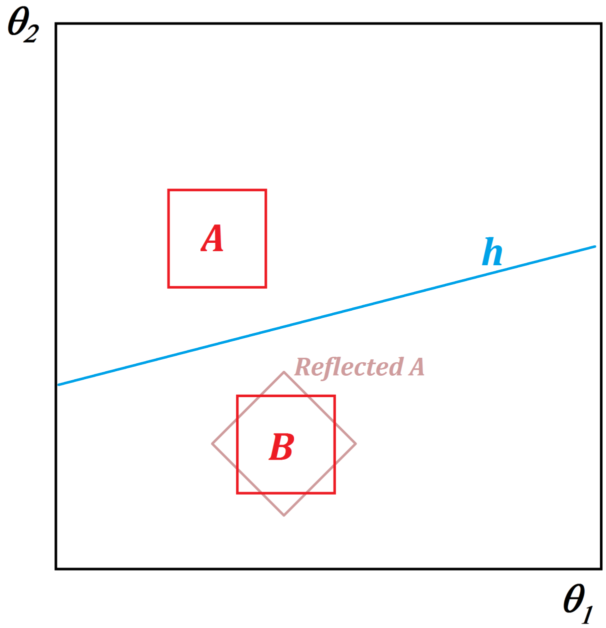

Figure 3.3.1: Failure of the symmetry assumption.

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.

there is not 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 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 3.2.2. 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.

3.4 Simulation Framework

3.4.1 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 21, 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 3.3.1. 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 3.4.1. In this case A and B are symmetric about h, h is horizontal, and the assortation 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 3.3.1. 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 3.4.1.

3.4.2 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

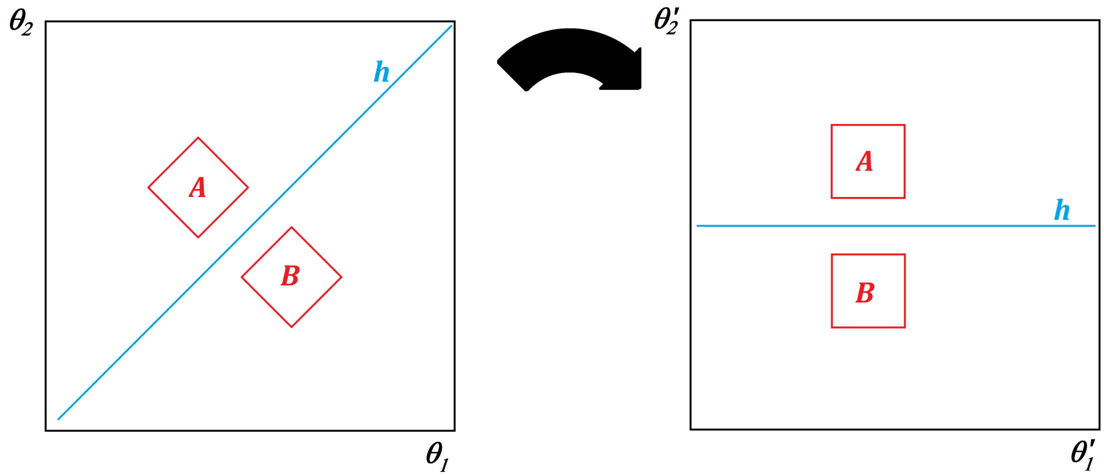

Figure 3.4.1: Rotation example

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 assortation, we must rotate the typespace so that h is normal to one of the newly created synthetic traits.

Klumpp's inside-out algorithm in the NTU case to find the stable matching outcome.[13] 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)).[14] 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 3.7.2 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

[13]See Online Appendix 1.4.3 for a summary of the algorithm.

[14]See Online Appendix 1.4.2 for a summary of the algorithm.

Predicting $\mu_1(a)$ (vertical characteristic) by a

	a_1	a_2	R^2
Symmetric, n=100			
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

Table 3.1: MC Results for the Vertical trait: NTU, 200 iterations, baseline specification.

convexity with TU, and NTU and TU simulations for various market sizes and numbers of preference traits.

3.5 Results

First, we'll look at the two dimensional case with NTU and a very coarse market of 100 agents on each side (Table 3.1). 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 (Table 3.2). 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

Predicting μ_2(a) (horizontal characteristic) by a

	a_1	a_2	R^2
Symmetric, n=100			
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

Table 3.2: MC Results for the Horizontal trait: NTU, 200 iterations, baseline specification.

Predicting μ_1(a) (vertical characteristic) by a

	a_1	a_2	R^2
Symmetric, n=100			
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

Table 3.3: MC Results for the Vertical trait: TU, 60 iterations, baseline specification.

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. 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 (Table 3.3). 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 $\mu_2(a)$ (horizontal characteristic) by 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

Table 3.4: MC 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 3.4). 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 assortation 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 3.7.2.2.

So far, we've considered very small, coarse matching markets. How do the results we've seen change when 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 models with just two traits over which agents have preferences. What happens if there are

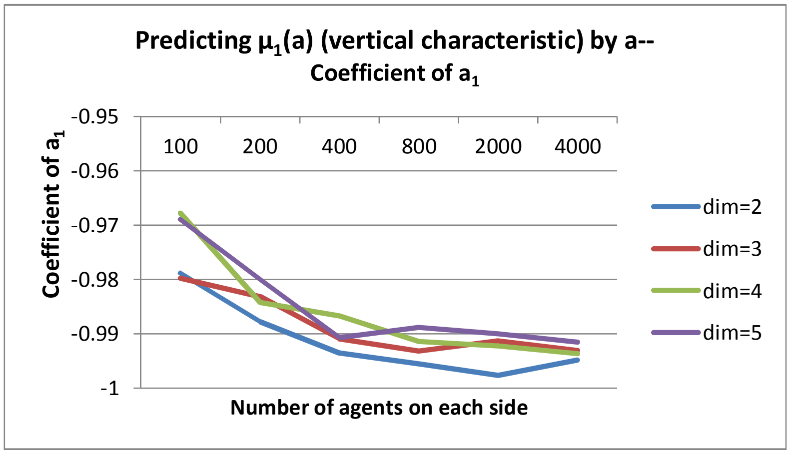

Figure 3.5.1: NTU with baseline specifications 1.

Predicting match's vertical trait from own vertical trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

more traits? We now address these questions.[15]

In Figure 3.5.1 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.

[15]Before 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.

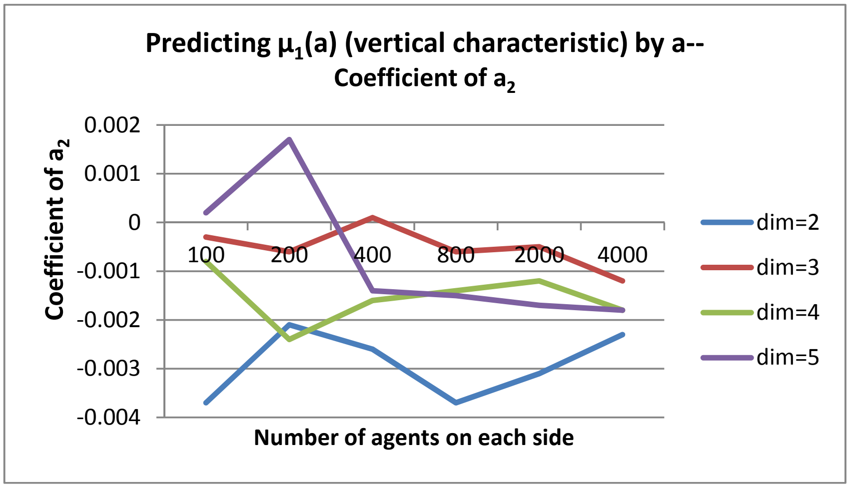

Figure 3.5.2: NTU with baseline specifications 2.

Predicting match's vertical trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

In Figure 3.5.2 we see that, for an agent's match's vertical trait, the 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 3.5.3 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

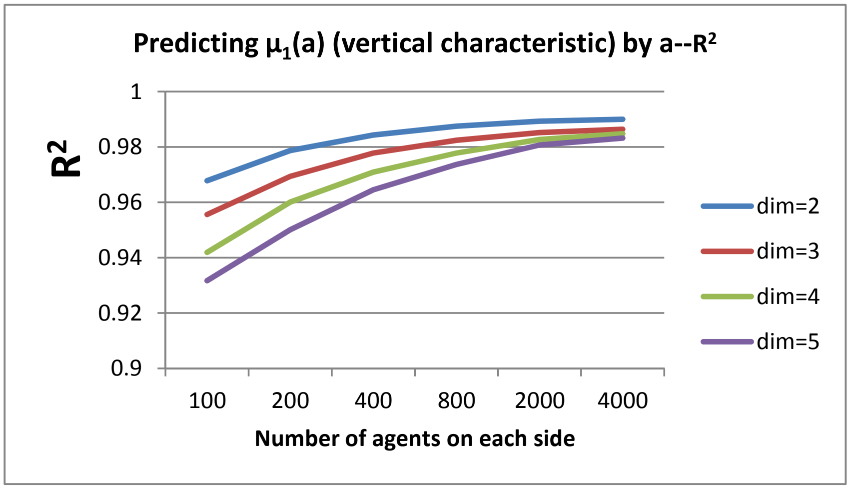

Figure 3.5.3: NTU with baseline specifications 3.

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.

one's match trait solely attributable to the predicted same own trait effect.

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 assortation in the vertical trait. As seen in Appendix 3.7.2.3, we also come fairly close to the predicted assortation 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

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

Table 3.5: MC Results for Five Dimensions.

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.

here 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 Table 3.5, 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 3.7.2.4.

3.6 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 Lauermann and Nï¿œldeke (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 environments[16] where finite[17] 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

[16]e.g. those satisfying the assumptions of one of the above papers and this paper.

[17]The NTU uniqueness result in this paper is only proved for finite markets.

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.

3.7 Appendix

3.7.1 Proofs

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 23 *(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. 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, $\sqrt{\sum \begin{matrix} n-1 \\ i=1 \end{matrix} d_{h_i}(a, b)^2 + d_\eta(a, b)^2}$. 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)). ■

Continuous Symmetric TU Matching.

Proof. First, we'll show stability. For a contradiction, consider the matching outcome of Proposition 22 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

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

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

$$= d_\eta(a_1, b_2)$$

$$= d_\eta(a_1, h) + d_\eta(b_2, h)$$

$$= (d(a_1, b_1) + d(a_2, b_2))/2$$

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.

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.[18] 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 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. [19] ■

[18] If 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.

[19] This 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.

3.7.2 Monte Carlo Simulations

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 3.5. 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 3.2.5 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, and 3.2.5 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.[20] 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)).[21] 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

[20] See Online Appendix 1.4.3 for a summary of the algorithm.

[21] See Online Appendix 1.4.2 for a summary of the algorithm.

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 th e synthetic traits.

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 assortation 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 3.7.1 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 assortation prediction along the horizontal dimension, and approximate the vertical assortation prediction relatively less well? First,

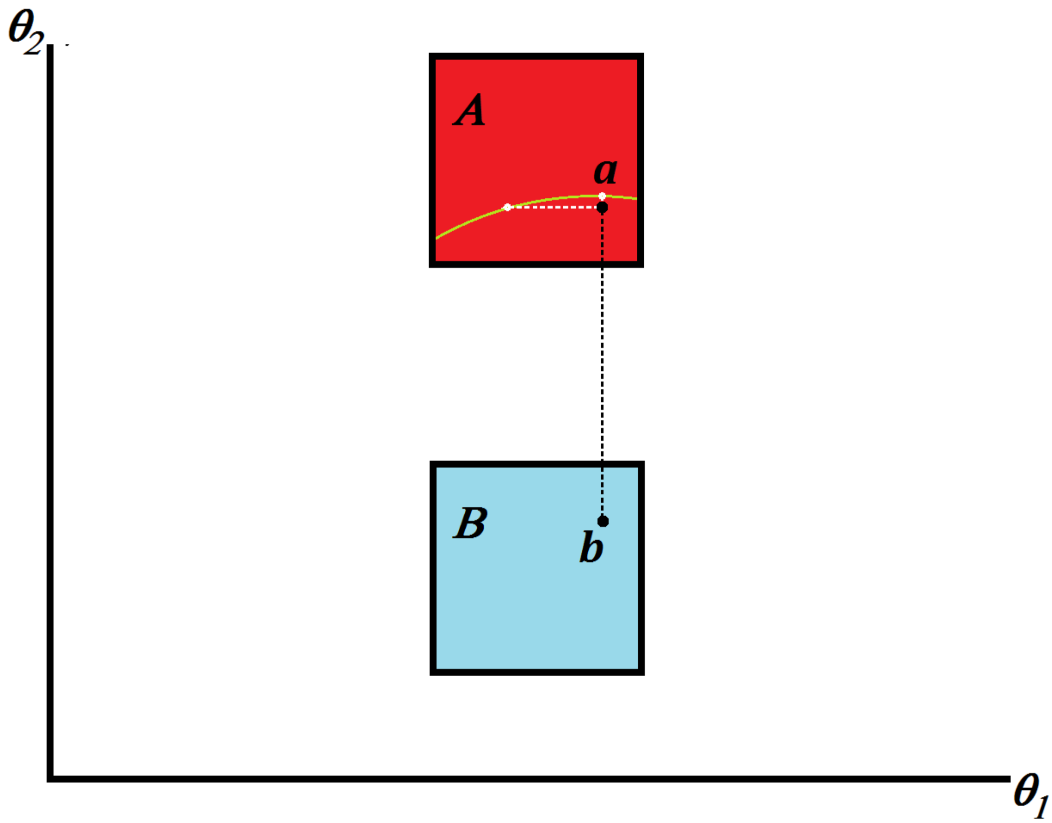

Figure 3.7.1: Horizontal deviations have small utility effects relative to vertical deviations.

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.

let's consider the simplified example in Figure 3.7.2. 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 types differ. In the NTU case we have the predicted vertical assortation, but not the horizontal assortation, and in the TU case we have the predicted horizontal assortation, but not the vertical assortation. Note also that we cannot have both at the same time–the predicted matches do not actually exist, and the two configurations shown in 3.7.2 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 3.7.2 distills an important quality of n-dimensional matching problems of the sort we've been studying: there will generally be a tradeoff between assortation along one dimension and assortation along another. In the two dimensional case, we can come up with matchings (not necessarily stable) that get closer to the horizontal assortation prediction, but this will often result in less fidelity to the vertical assortation prediction, and vice versa. We will see below that TU stable matching puts a greater premium on horizontal assortation than vertical assortation relative to NTU, so the stable matches in the TU framework will exhibit better horizontal assortation and commensurately worse vertical assortation.

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 3.7.2. 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 assortation affects individual utility: everyone agrees on the rank ordering of potential matches, so if one A agent matchesf 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 assortation 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 assortation 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 assortation at the expense of vertical assortation.

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 ith trait, or equivalently we were to increase the weight on the ith trait in the distance metric, the relative fidelity of the simulations to the horizontal and vertical assortation predictions would change. Specifically, putting more weight on a trait will generally improve the assortation along that trait, while worsening it along all others. An illustration of this can be seen in Table 3.5 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 assortation

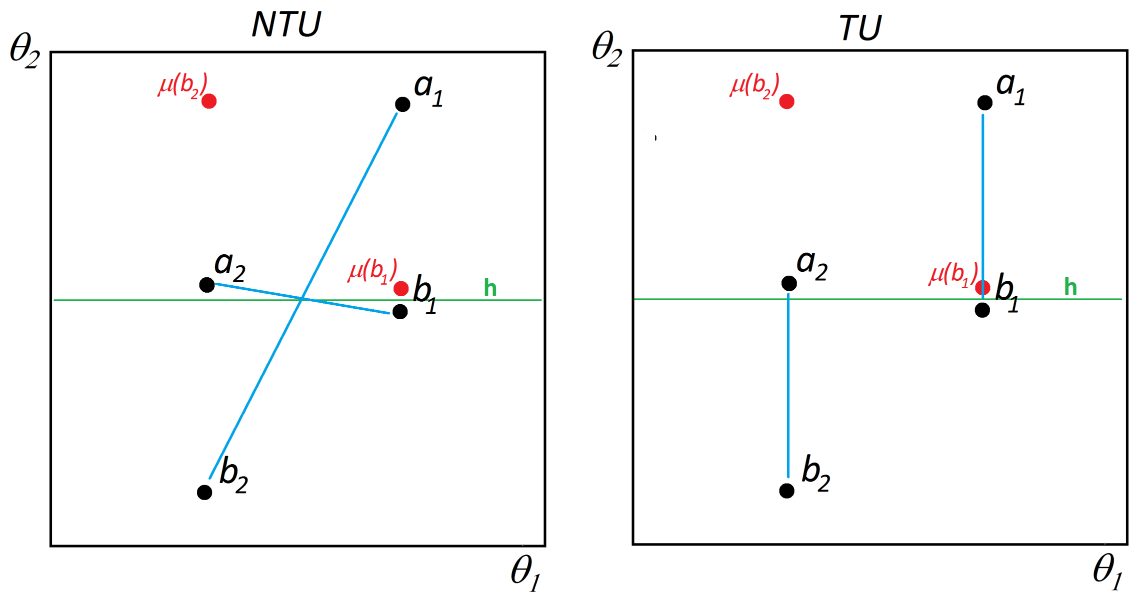

Figure 3.7.2: TU and NTU stable assignments differ when the symmetry assumption fails. Here we consider the simplest nontrivial matching problem–two agents on each side–and assume mild convex preferences for distance for the TU case. Assume $u_{TU}(a) = -\sqrt{d(a, \mu(a))}$. The predicted but nonexistent symmetric matches for B agents are shown as $\mu(b_i)$.

results worsen. Intuitively, we can think about this in much the same way we looked at the better vertical assortation 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 assortation. 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 assortation result will be stronger along the higher-weight axis and weaker along all others.

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 the horizontal same trait coefficients and the corresponding R^2.

In Figure 3.7.3 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.

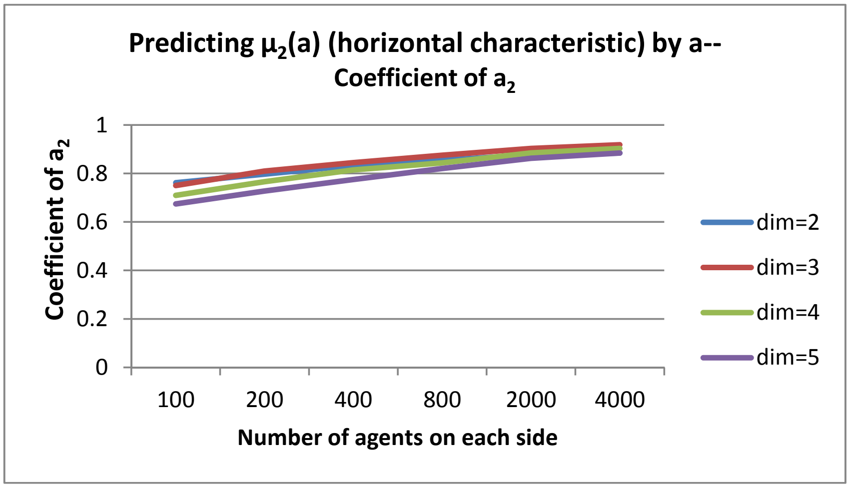

Figure 3.7.3: NTU with baseline specifications 4.

Predicting match's horizontal trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 200 iterations per specification.

Recall that this this 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 3.7.4 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 this 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

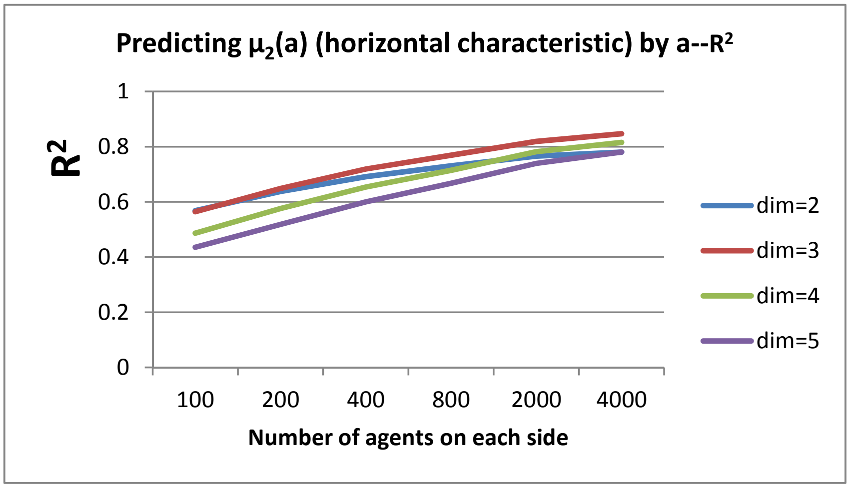

Figure 3.7.4: NTU with baseline specifications 5.

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.

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.

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

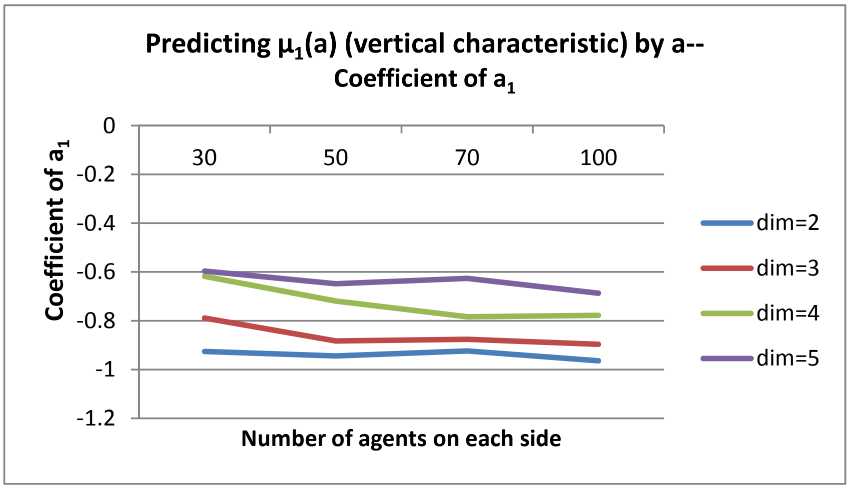

Figure 3.7.5: TU with baseline specifications 1.

Predicting match's vertical trait from own vertical trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

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 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.

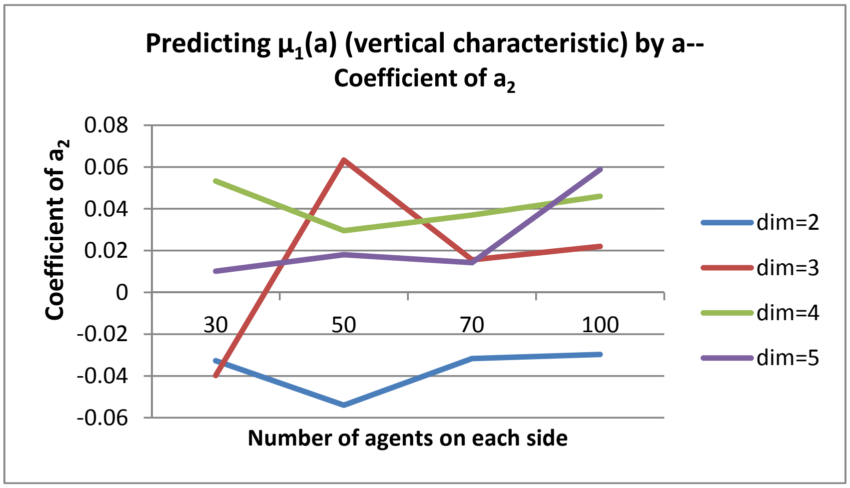

Figure 3.7.6: TU with baseline specifications 2.

Predicting match's vertical trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

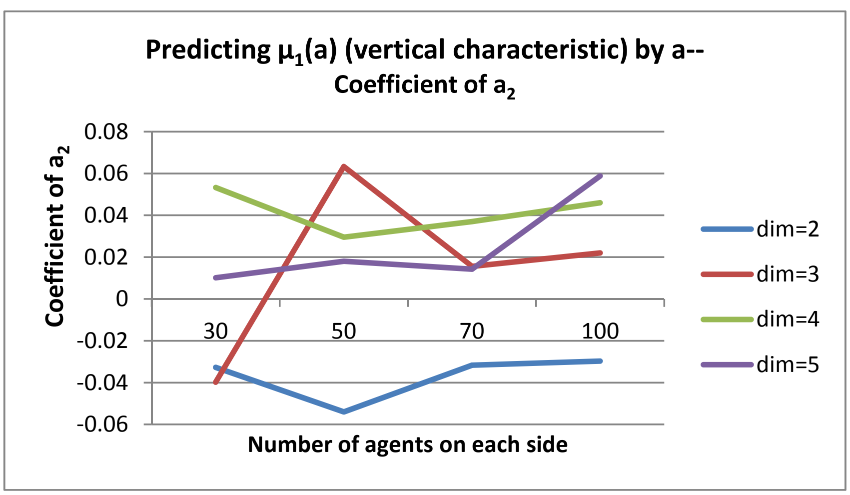

Figure 3.7.7: TU with baseline specifications 3.

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.

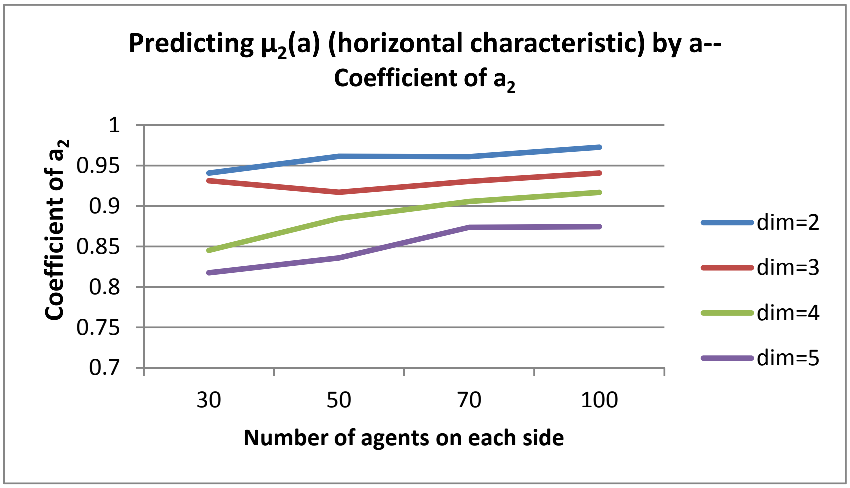

Figure 3.7.8: TU with baseline specifications 4.

Predicting match's horizontal trait from own horizontal trait for different number of agents and different numbers of preference dimensions. 20 iterations per specification.

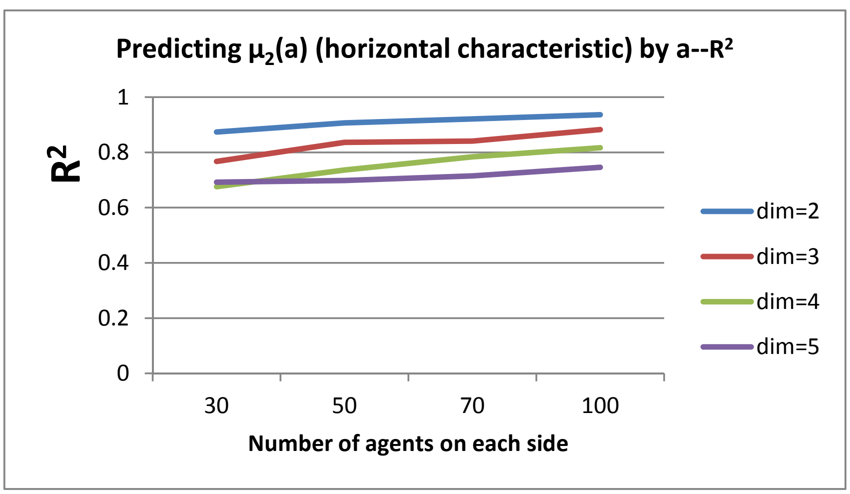

Figure 3.7.9: TU with baseline specifications 5.

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.

BIBLIOGRAPHY

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.

Armstrong, M. (2006). Competition in Two-Sided Markets. *RAND Journal of Economics 37(3)*, 668-691.

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.

Bedre-Defolie, Özlem, and Emilio Calvano. "Pricing payment cards." American Economic Journal: Microeconomics 5, no. 3 (2013): 206-231.

Bilancini, Ennio, and Leonardo Boncinelli. "Disclosure of information in matching markets with non-transferable utility." Games and Economic Behavior 82 (2013): 143-156.

Bloch, Francis, and Harl Ryder. "Two-Sided Search, Marriages, and Matchmakers." International Economic Review 41, no. 1 (2000): 93-116.

Browning, Martin, Pierre-André Chiappori, and Yoram Weiss. Economics of the Family. Cambridge University Press, 2014.

Burdett, Ken, and Melvyn G. Coles. "Marriage and class." The Quarterly Journal of Economics 112, no. 1 (1997): 141-167.

Burdett, Ken, and Melvyn G. Coles. "Marriage and class." The Quarterly Journal of Economics 112, no. 1 (1997): 141-167.

Chade, Hector, Jan Eeckhout, and Lones Smith. Search and matching models in macroeconomics. working paper, 2015.

Chade, Hector. "Two-sided search and perfect segregation with fixed search costs." Mathematical Social Sciences 42, no. 1 (2001): 31-51.

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.

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.

Clark, S. (2003), “Matching and Sorting with Horizontal Heterogeneity”, ESE Discussion Paper 98,

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

Clark, S. (2007). Matching and sorting when like attracts like. Working paper.

Crawford, Vincent P., and Joel Sobel. "Strategic information transmission." Econometrica: Journal of the Econometric Society (1982): 1431-1451.

Damiano, Ettore, and Hao Li. "Competing matchmaking." Journal of the European Economic Association (2008): 789-818.

Damiano, Ettore, and Hao Li. "Price discrimination and efficient matching." Economic Theory 30, no. 2 (2007): 243-263.]

Flanders, S. (2014). Characteristic Aggregation in Models of Matching Markets. Working paper.

Gale, David, and Lloyd S. Shapley. "College admissions and the stability of marriage." The American Mathematical Monthly 69, no. 1 (1962): 9-15.

Hagenbach, Jeanne, Frédéric Koessler, and Thomas Trégouët. "Truthtelling in Matching Markets." (2015).

Hitsch, Günter J., Ali Hortaçsu, and Dan Ariely. "Matching and sorting in online dating." The American Economic Review (2010): 130-163

Hopkins, Ed. "Job market signaling of relative position, or Becker married to Spence." Journal of the European Economic Association 10, no. 2 (2012): 290-322.

Hoppe, Heidrun C., Benny Moldovanu, and Aner Sela. "The theory of assortative matching based on costly signals." The Review of Economic Studies 76, no. 1 (2009): 253-281.

J. Eeckhout, Bilateral search and vertical heterogeneity, International Economic Review 40 (1999) 869–887.

Jan Eeckhout, 2006. "Local Supermodularity and Unique Assortative Matching," 2006 Meeting Papers 127, Society for Economic Dynamics.

Klumpp, Tilman. "Two-sided matching with spatially differentiated agents."Journal of Mathematical Economics 45, no. 5 (2009): 376-390.

Ko, Chiu Yu, and Hideo Konishi. "Profit-maximizing matchmaker." Games and Economic Behavior 75, no. 1 (2012): 217-232.

Lauermann, Stephan, and Georg Nöldeke. "Stable marriages and search frictions." Journal of Economic Theory 151 (2014): 163-195.

Lee, Robin S. "Vertical integration and exclusivity in platform and two-sided markets." The American Economic Review 103, no. 7 (2013): 2960-3000.

Legros, Patrick, and Andrew F. Newman. "Beauty is a beast, frog is a prince: assortative matching with nontransferabilities." Econometrica 75, no. 4 (2007): 1073-1102.

Lindenlaub, I. (2013). ”Sorting Multidimensional Types: Theory and Application ”, Job Market Paper.

McNamara, John M., and E. J. Collins. "The job search problem as an employer-candidate game." Journal of Applied Probability (1990): 815-827.

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

Rochet, Jean-Charles, and Jean Tirole. "Platform competition in two-sided markets." Journal of the European Economic Association (2003): 990-1029.

Rochet, Jean-Charles, and Jean Tirole. "Two-sided markets: a progress report." The RAND Journal of Economics 37, no. 3 (2006): 645-667.

Shapley, Lloyd and Martin Shubik, ’The Assignment Game 1: The Core’, International Journal of Game Theory, 1 (1972),

Shimer, Robert, and Lones Smith. "Assortative matching and search." Econometrica 68, no. 2 (2000): 343-369.

Smith, Lones. "The marriage model with search frictions." Journal of Political Economy 114, no. 6 (2006): 1124-1144.

Weyl, E. Glen. "A price theory of multi-sided platforms." The American Economic Review (2010): 1642-1672.