August 2008

Preference Structure and Random Paths to Stability in Matching Markets

James W. Boudreau
University of Connecticut

University of Connecticut

Department of Economics Working Paper Series

Preference Structure and Random Paths to Stability in Matching Markets

James W. Boudreau
University of Connecticut

Working Paper 2008-29

August 2008

---

341 Mansfield Road, Unit 1063
Storrs, CT 06269–1063
Phone: (860) 486–3022
Fax: (860) 486–4463
http://www.econ.uconn.edu/

Abstract

This paper examines how preference correlation and intercorrelation combine to influence the length of a decentralized matching market's path to stability. In simulated experiments, marriage markets with various preference specifications begin at an arbitrary matching of couples and proceed toward stability via the random mechanism proposed by Roth and Vande Vate (1990). The results of these experiments reveal that fundamental preference characteristics are critical in predicting how long the market will take to reach a stable matching. In particular, intercorrelation and correlation are shown to have an exponential impact on the number of blocking pairs that must be randomly satisfied before stability is attained. The magnitude of the impact is dramatically different, however, depending on whether preferences are positively or negatively intercorrelated.

Journal of Economic Literature Classification: C78, C15, P41

Keywords: Marriage matching, stability, random paths.

1 Introduction.

Despite the success of centralized matching procedures in the markets matching new doctors to residencies (Roth and Peranson, 1990) and city children to public schools (Abdulkadíroğlu, Pathak and Roth, 2005; Abdulkadíroğlu, Pathak, Roth, and Sönmez, 2005) most matching markets remain decentralized. While there are many reasons why these markets are not centralized, it is an open question whether or not decentralized matching markets actually manage to attain *stable* outcomes. A *stable* matching is one in which no group of matched agents would prefer to be with each other rather than those they are currently matched with. In one-to-one two-sided matching problems (also known as marriage matching problems), such pairs of agents are known as *blocking pairs* because they are able to block the match by breaking relationships with their current partners and instead matching with each other. Stability is a critical characteristic for any matching market for obvious reasons, and thus a vast literature has been dedicated to the study of deterministic mechanisms that are guaranteed to reach such outcomes (see Roth and Sotomayor (1990) or Roth (2008) for surveys). The problem of course is that this literature is for the most part inapplicable to markets that are not centralized.

Yet there is good reason to believe that decentralized matching markets can indeed achieve stability. In particular, the work of Roth and Vande Vate (1990) outlines a random process by which any one-to-one two-sided matching market can converge to a stable outcome. Beginning from any initial matching of agents, the process goes step by step, randomly selecting one blocking pair from the set of all existing blocking pairs at any step and allowing them to pair with each other while leaving their former partners single. Roth and Vande Vate (1990) show that this process converges to a stable matching with probability one in a finite number of steps, and their results are generalized (with some minor restrictions on preferences) by Klaus and Klijn (2007) to the case of many-to-one matching with couples, and by Kojima and Ünver (2008) to the case of many-to-many matching problems. Furthermore, Eriksson and Häggström (2008) show that even in cases of incomplete information, stable matchings can still be attained by decentralized processes given certain restrictions on preferences and some degree of search effort.

Thus, it is at least possible for stability to be obtained by a decentralized process of agents simply making and breaking partnerships over time. But in many situations, forming new partnerships over and over again is not a trivial exercise. Though individuals may be willing to quit their jobs to obtain more preferable employment and employers may be willing to fire employees to hire better help, it is generally not a process they wish to engage in frequently since forming new partnerships often requires significant sunk costs from both involved parties. At the least such a process necessarily involves some expenditure of the ultimate scarce resource: time.

The aim of this paper is to identify and quantify relationships between agents' preferences and the length of time necessary to achieve stability in matching markets as measured by the number of steps taken in the decentralized random process defined by Roth and Vande Vate (1990). Obviously, in this sense larger markets with more participants will have longer paths. Large markets require more steps because they have more blocking pairs to satisfy. Here, however, the number of participants is held fixed while the preference structure of participants varies. Preferences are fundamental characteristics of any matching market since they directly determine which outcomes are stable, and it therefore seems prudent to evaluate their impact on the length of decentralized paths to stability.

The relationship between preference structure and matching market outcomes has only recently gained much attention, probably because the vast amount of variety and complexity that preferences can entail tends to limit theoretical results to rather extreme cases. Examples of such work include Wilson (1972) and Knoblauch (2008), who find limiting results for upper and lower bounds respectively on men's partner satisfaction in marriage problems when men propose to women in a deferred-acceptance algorithm (see Gale and Shapley, 1962) and men's preferences are random.

Taking advantage of simulation techniques, however, recent work has been able to go beyond the limitations of theory to analyze match outcomes with more flexible preference characteristics. Caldarelli and Cappoci (2001) and Celik and Knoblauch (2007) use simulation methods to analyze the relationship between preference correlation and men's and women's partner satisfaction in marriage problems, while Boudreau and Knoblauch (2007) examine the relationship between intercorrelation of preferences and partner satisfaction in marriage problems. Teo, Sethuraman and Tan (2001) use simulation to investigate the potential for strategic misrepresentation in marriage markets. All four papers restrict attention to the case of centralized matching, using the men-propose deferred-acceptance algorithm to find stable matchings. In a different but related framework, Alpern and Reyniers use simulations to examine how perfectly intercorrelated (1999) and perfectly correlated (2005) preferences alter bilateral search/accpetance strategies in populations seeking mates.

This paper uses simulation methods akin to those in Boudreau and Knoblauch (2007) to examine just how long a random path to stability is likely to be for a (one-to-one and two-sided) marriage matching market, given the market's preference structure. Simulation experiments begin by randomly generating agents' preferences with specific levels of correlation and intercorrelation, and then assigning agents to an arbitrary initial matching. The path to stability then proceeds by identifying all possible blocking pairs and satisfying one of them at random, with each pair having an equal chance of being chosen. The process then repeats until no blocking pairs exist.

Regression analysis of the simulated data shows that both correlation and intercorrelation have an exponential impact on the length of the market's random path to stability. When all agents on one side of the market agree on which partners on the other side are most attractive (preferences are correlated), the path to stability tends to be relatively short. If agents on one side of the market rank highly those on the other side who reciprocate their preference (preferences are positively intercorrelated), the path to stability is even shorter. If, however, agents prefer those on other side of the market who do not prefer them (preferences are negatively intercorrelated), then the path to stability can be quite long. These results thus help to characterize which types of matching markets could benefit most from centralized mechanisms that facilitate their progress toward stability.

2 Preliminaries.

The model considered here is the simple marriage matching problem, popularized by Gale and Shapley (1962). There are two finite and disjoint sets, $M = \{m_1, m_2, \ldots, m_n\}$ and $W = \{w_1, w_2, \ldots, w_n\}$, of agents known respectively as men and women that are seeking to match with each other[1]. Each agent has a complete, strict, and transitive preference ordering over the n agents on the other side of the market and the prospect of remaining single. Each man i's preferences are represented by a one-to-one and onto ranking function $r_{m_i} : W \cup \{m_i\} \to \{1, 2, \ldots, n+1\}$, where w_j is preferred to w_k by m_i if $r_{m_i}(w_j) < r_{m_i}(w_k)$. Women's preferences are similarly represented by r_w. Here, as in the basic marriage matching model, the option of remaining single is always the least preferred option, so $r_{m_i}(m_i) = n + 1$ for all men and $r_{w_j}(w_j) = n+1$ for all women. Letting $R \equiv \{r_{m_1}, \ldots, r_{m_n}, r_{w_1}, \ldots, r_{w_n}\}$ be the set of all agents' ranking functions, the model is therefore specified by $\{M, W, R\}$.

The outcome of a marriage problem is a matching of men to women, represented by a one-to-one and onto function $\mu : M \cup W \to M \cup W$. A matching is said to be *stable* if there does not exist a *blocking pair* $\{m_i, w_j\}$ such that $\mu(m_i) \neq w_j$ but $r_{m_i}(w_j) < r_{m_i}(\mu(m_i))$ and $r_{w_j}(m_i) < r_{w_j}(\mu(w_j))$. As proved by Gale and Shapley (1962), at least one such matching will always exist in the case of marriage markets, and the set of all such matchings is known as a particular market's *stable set*. Furthermore, as proved by Roth and Van Vate (1990), stability can be reached from any initial (unstable) matching in a finite number of steps via the following process:

Step 1. Find all possible blocking pairs existing in the current match, μ_t.

Step 2. Randomly select one of the blocking pairs, $\{m_i, w_j\}$, and satisfy them by matching them together so that $\mu_{t+1}(m_i) = w_j$. If $\{m_i, w_j\}$ were not single in μ_t, this leaves their former partners, $\mu_t(m_i)$ and $\mu_t(w_j)$ single in μ_{t+1}. Note that in the current model, this pair is necessarily then a blocking pair for the next round.

[1] Alternative interpretations categorize the sets as firms and workers, or workers and machines.

⋮

Step k**.** Repeat steps 1. and 2. until no more blocking pairs exist.

If more than one stable matching is possible for $\{M, W, R\}$ (as is often the case), then the outcome that realizes from this decentralized process depends on the initial starting point and which particular blocking pairs end up being satisfied along the way.

3 Experimental Settings.

3.1 Preferences.

Preferences for individuals can be as complicated and varied as individuals themselves. Ultimately in matching markets, however, all that matters is that agents can indeed rank their prospective partners. Markets can then be categorized according to how the agents' individual ranking lists compare with each other. If men (women) all tend to agree on which women (men) are most and/or least preferable, then their preferences are positively *correlated.* Perfunctory examples include the possibility of all men agreeing on which woman is the most beautiful, or all women agreeing on which man is the best provider. If agents across sets have preferences related to each other in some way, then their preferences are *intercorrelated.* An example of positive preference intercorrelation is the tendency of both men and women to prefer those with similar educational backgrounds. A possible example of negative intercorrelation, on the other hand, could be poor men strongly preferring wealthy women while wealthy women prefer only wealthy men. Of course, as even these simple examples illustrate, the two concepts are often intertwined.

Because of the many varied types and degrees of correlation and intercorrelation, as well as the many degrees of combination between the two, theoretical results on the effects of these characteristics are difficult to obtain beyond extremely rigid cases and assumptions. To perform a more general examination of the effects of preference correlation and intercorrelation on a marriage market's path to stability, the approach here follows that of Caldarelli and Capocci (2001) and Boudreau and Knoblauch (2007), simulating the matching process of many experimental markets and analyzing the results. The critical difference between this paper and those previously mentioned, of course, is the use of a decentralized (random) mechanism rather than the centralized deferred-acceptance algorithm.

As in Boudreau and Knoblauch (2007), individual agents in each experimental market are endowed with preferences by the following equation (1), which shows the score given by m_i to w_j.

$$X^m_{i,j} = \eta^m_{i,j} + U^m I^m_j + V^m|i-j|_n \tag{1}$$

For any $w_j, w_k \in W$, $r_{m_i}(w_j) < r_{m_i}(w_k)$ if $X^m_{i,j} < X^m_{i,k}$, with women forming preferences according to a symmetrically specified score $X^w_{i,j}$.

The first term on the right-hand-side of equation (1), $\eta^m_{i,j}$, represents the unique preference of m_i for w_j, and is randomly drawn from a uniform distribution on $[0,1]$. The next two terms introduce a degree of correlation to the preferences of agents from the same set. I^m_j is another uniform random draw from $[0,1]$, but one that is common to all men. Together with U^m, which is an adjustable weighting parameter, the second term represents the common degree of preference for w_j that is shared by all men. A higher value for U^m means a higher degree of correlation for men's preferences. The final term in equation (1) introduces intercorrelation across the groups of men and women. The term $|i-j|_n = \min\{|i-j|, n-|i-j|\}/\frac{n}{2}$ is the normalized distance of m_i from w_j when they are placed on an n-hour clock at hours i and j respectively. If V^m is negative then men like women closest to them, while if V^m is positive men prefer women further away from them. Intercorrelation between the two sets depends on the signs of both V^m and V^w; when both are of the same (opposite) sign, preference lists will be positively (negatively) intercorrelated.

3.2 Measurement.

Simulation data is gathered by generating preferences for a market of fixed size according to (1), randomly assigning the n men and women for an initial matching, and then proceeding according to the decentralized process outlined in section 2. Blocking pairs are selected at random, with each possible pair having an equal likelihood of being chosen, and the length of the path to stability is measured by the number of blocking pairs satisfied. Because the path can be quite long, a relatively small market size of $n = 10$ men and women is used for most of the analysis, but eventually is allowed to vary.

Preference correlation and intercorrelation are controlled in simulations by varying the parameters $\{U^m, U^w\} \in \{0, 1, 2, 3, 4\}$ and $\{V^m, V^w\} \in \{-2, -1.5, -1, -0.5, 0.5, 1, 1.5, 2\}$. 100 independent trials are conducted for each ordered 4-tuple $\{U^m, U^w, V^m, V^w\}$, with results presented as the average over trials. In addition to altering the preference parameters, however, it is also necessary to measure the actual levels of correlation and intercorrelation in agents' realized preference lists. Since the realized lists are in part stochastically determined, although the specified parameters are heavily influential they are not complete descriptors.

Fortunately, more general measures of preference correlation and intercorrelation have already been devised and are available for use. To measure correlation, the present analysis therefore employs the definition of Celik and Knoblauch (2007), denoting correlation in men's and women's lists by $\rho^m, \rho^w \in [0, 1]$ with a larger number indicating a larger degree of correlation. The measure of intercorrelation across men's and women's lists is as defined by Boudreau and Knoblauch (2007), and is denoted by $\Phi \in [0, 1]$. A value of $\Phi = 0$ indicates perfectly negative intercorrelation, a value of $\Phi = 1$ indicates perfectly positive intercorrelation, and a value of $\Phi = 0.5$ indicates neutral intercorrelation. Complete descriptions of these measures are available in the appendix of this paper, or in the original citations.

4 Results.

The analysis proceeds in four steps, first considering the effects of intercorrelation and correlation separately, then considering the two together, and finally taking market size into account. To examine the relationship between preference intercorrelation and the length of the market's path to stability, figure 1 illustrates the simulated data as arranged by the logged number of steps on the path. From this data, two characteristics are immediately evident.

Figure 1: Simulation Results on Intercorrelation

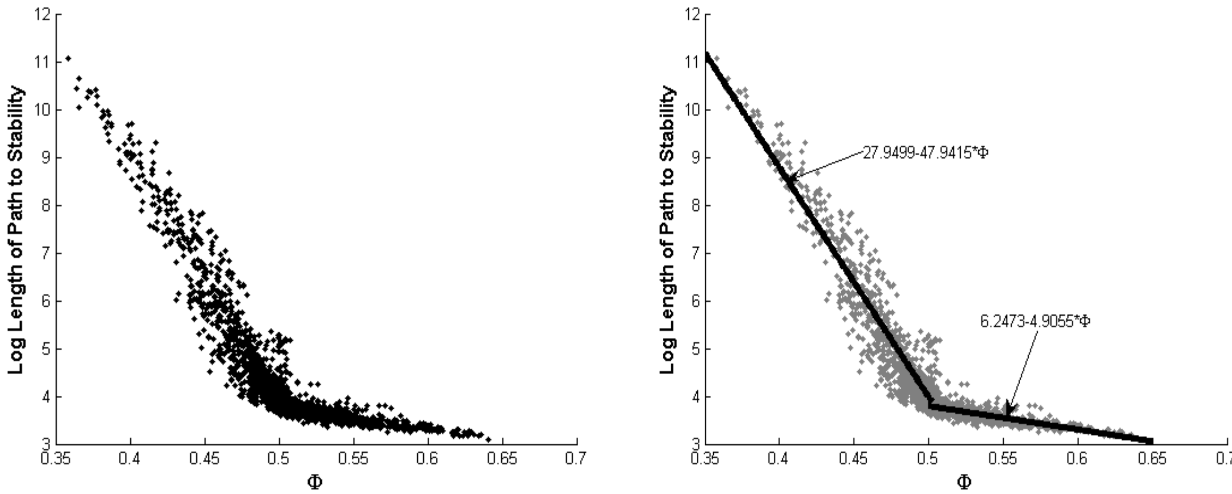

1. The level of intercorrelation has an exponential impact on the length of a market's random path to stability, suggesting a relationship of $log(T) = \alpha_0 + \alpha_1\Phi$. Estimates of α_0 and α_1 are presented on the right side of figure 1, along with the linear trends they imply.
2. The effect of positive intercorrelation ($\Phi > 0.5$) is drastically different from the effect of negative intercorrelation ($\Phi < 0.5$). Such asymmetric effects are consistent with the results of Boudreau and Knoblauch (2007), who find a similar relationship for preference intercorrelation and the size of marriage market's stable set.

Together and simply put, these observations suggest that when more stable sets are possible to reach the time taken to actually reach one in a decentralized fashion increases dramatically. An increased number of stable outcomes means that the random selection of blocking pairs can entail switching back and forth in terms of direction toward any outcome in particular. There are thus two avenues for preferences to affect the length of a market's path to stability: by increasing the total number of blocking pairs and by complicating the path with twists and turns. The following two examples serve to illustrate these effects.

Example 1. Consider a marriage market of $n = 3$ in which men most strongly prefer women that have the same indices as themselves, so $r_{m_i}(w_i) = 1$ for all $i = 1, 2, 3$. Preferences are negatively intercorrelated, however, because women most strongly prefer men with indices different than their own. In particular, $r_{w_1}(m_3) = r_{w_2}(m_1) = r_{w_3}(m_2) = 1$. With the remaining rankings are arbitrary, this market has at least two possible stable matchings: $\{(m_1, w_1); (m_2, w_2); (m_3, w_3)\}$ and $\{(m_1, w_2); (m_2, w_3); (m_3, w_1)\}$.

Suppose for simplicity that the market begins with all agents unmatched and proceeds toward stability by randomly satisfying blocking pairs. Since remaining unmatched is all agents' least preferred option, any two agents constitute a blocking pair in the market's initial state. Suppose then that the first pair to be satisfied is (m_1, w_1). Though this pair is a member of one of the stable sets, it can be broken by the blocking pair (m_3, w_1), a pair that is also included in one of the market's stable sets, because w_1 prefers m_3 to all others and m_3 prefers w_1 to being single. This matching can similarly be broken up, however, by the blocking pair (m_3, w_3). That pair can then be broken by the pair (m_1, w_3), which can in turn be broken by the pair (m_1, w_1) to complete a full circle.

Because the two sides' preferences are negatively intercorrelated, a cyclical pattern is possible in this case despite the fact that each pairing involved belongs to one of the market's stable sets. No matter which pair is formed, one member always prefers to be matched with an agent that is not their current partner. Thus, if that preferred agent happens to be single at the time, there is always a chance that the pair will be broken up.

Example 2. Now consider a very similar market with the same men's preferences, but with women's preferences positively intercorrelated so that $r_{w_i}(m_i) = 1$ for all $i = 1, 2, 3$. In this case there is only one possible stable outcome: $\{(m_1, w_1); (m_2, w_2); (m_3, w_3)\}$. Cycling is still possible on this market's path to stability depending on the remaining preference orderings, but once any pair belonging to the stable set is formed it can not be broken. The path to stability is therefore dramatically shortened since those two agents can no longer be involved in any blocking pairs.

Figure 2: Simulation Results on Correlation

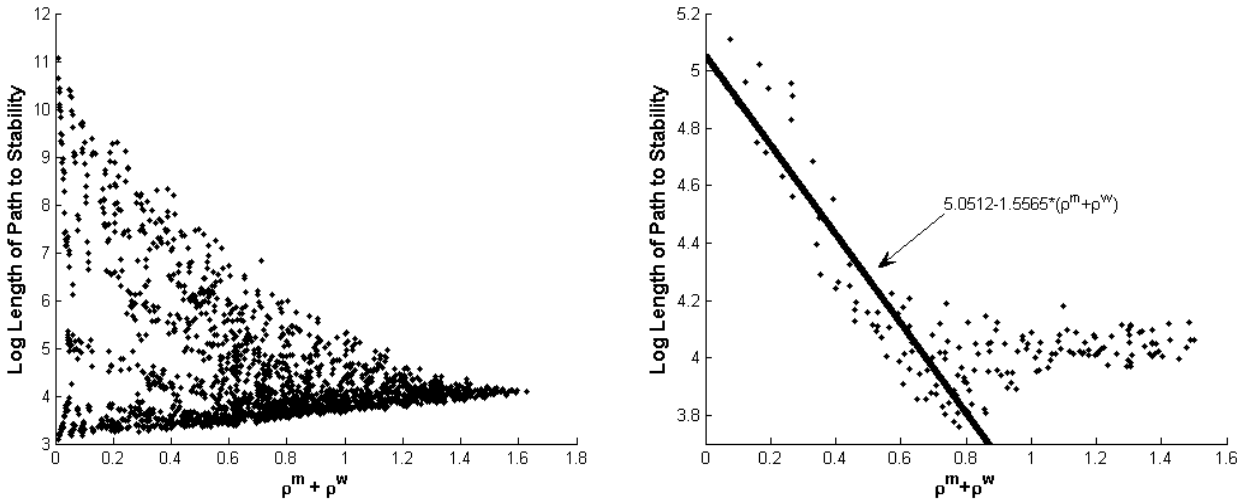

Next, the relationship between the correlation of men's and women's preferences and the length of the market's path to stability is illustrated in Figure 2. Because the effects of the two genders' preferences are symmetric in the case of the decentralized process studied here (which is certainly not the case for most centralized mechanisms, especially the oft-studied deferred-acceptance mechanism), correlation in the market as a whole is simply categorized by adding the correlation scores of men and women.

Due to the intermingling effects of intercorrelation, the impact of correlation is not as readily apparent on its own when considering all 2025 data points that result from all permutations of $\{U^m, U^w\} \in \{0, 1, 2, 3, 4\}$ and $\{V^m, V^w\} \in \{-2, -1.5, -1, -0.5, 0.5, 1, 1.5, 2\}$. Thus, the right side of Figure 2 depicts simulation results for $V^m = V^w = 0$, with $\{U^m, U^w\} \in [0, 3]$ in increments of 0.25. Holding intercorrelation more or less constant in this manner, the effects of correlation become more evident. Like intercorrelation, correlation in preferences also exponentially reduces the number of blocking pairs satisfied before reaching a stable match. This effect is not always present, however, as it levels off once the aggregate correlation in the market becomes high enough. With sufficient correlation, not many blocking pairs are likely to exist, so the path to stability is a relatively short trip from any initial matching. The estimated coefficients presented on the right side of figure 2 are therefore based only on data with $\rho^m + \rho^w = 0.8$.

Based on the two previous treatments of preference intercorrelation and correlation separately, the following is speculated as the joint relationship between the length of a market's path to stability and its fundamental preference characteristics:

$$log(T) = \beta_0 + \beta_1 \Phi + \beta_2(\rho^m + \rho^w). \tag{2}$$

Given the asymmetric effects of intercorrelation, however, estimation of equation (2) is conducted separately for $\Phi < 0.5$ and $\Phi > 0.5$. These results are presented in table 1.

For the most part the estimated coefficients have signs and magnitudes that justify with earlier results. Intercorrelation is the more dominant influence whether it is positive or negative. In fact, the coefficient on preference correlation is generally not significant in the presence of positive intercorrelation. The explanation for this absence of an effect lies in the fact that the path to stability is already quite short when intercorrelation is neutral ($\Phi = 0.5$). In accordance with figure 2, correlation's effects tend to taper off quickly. This also serves as an explanation for the lopsided impact of intercorrelation; the path to stability can only be so short.

To complete the analysis, market size must also be taken account. A larger market will obviously entail a longer path to stability, but how much longer? And how are the effects of preference structure altered when more agents' preferences are involved? Simulations of larger markets are difficult because of the tremendous complexity their paths can entail. Table 1 does present estimates with relatively small variations in market size, however, indicating that a larger number of participants can mean a pronounced increase in the effects of preference structure.

Again, the case of preferences being negatively intercorrelated is quite different from the case of positive intercorrelation. As the size of the market increases, the increase in the impact of intercorrelation is drastic when intercorrelation is negative, but only moderate when intercorrelation is positive. The effect of correlation also differs from case to case. Though its influence is small it does grow consistently with market size in the case of negative intercorrelation, but

Table 1: Estimated Coefficients from Equation (2)

$\Phi < 0.5$	β_0	β_1	β_2	$\Phi > 0.5$	β_0	β_1	β_2
$n = 4$	3.4931*	-3.4361*	0.0166	$n = 4$	2.6566*	-1.9455*	0.1110*
	(0.0542)	(0.1314)	(0.0143)		(0.0933)	(0.1488)	(0.0168)
$n = 6$	8.5785*	-11.4942*	-0.0774*	$n = 6$	4.6613*	-3.9010*	0.0100
	(0.1214)	(0.2790)	(0.0298)		(0.2433)	(0.3917)	(0.0456)
$n = 8$	15.8055*	-24.3367*	-0.2097*	$n = 8$	5.8852*	-4.9997*	-0.0548
	(0.2433)	(0.5611)	(0.0629)		(0.4522)	(0.7399)	(0.0804)
$n = 10$	26.4963*	-44.2547*	-0.4015*	$n = 10$	6.7929*	-5.8009*	-0.1060
	(0.3956)	(0.9297)	(0.0993)		(0.3802)	(0.6240)	(0.0708)
$n = 12$	37.8565*	-65.5858*	-0.5336*	$n = 12$	7.3628*	-6.2975*	-0.0505
	(0.9174)	(1.9134)	(0.0993)		(0.8093)	(1.3549)	(0.1141)
$n = 14$	51.5427*	-91.5156*	-0.8097*	$n = 14$	7.6192*	-6.2143*	0.1167
	(1.4629)	(3.0022)	(0.1400)		(0.8515)	(1.4231)	(0.1191)

Standard errors in parentheses, * indicates 99% significance level.

remains irrelevant regardless of market size when preferences are positively intercorrelated.

5 Discussion.

The appeal of the decentralized process considered here lies in its similarity with real-life market functioning. Partnerships are often made hastily and then dissolve if and when better opportunities present themselves. Thus, the question addressed by this paper is the following: how long is such a process likely to take before stability is attained? If the answer is exceptionally long, and if making and breaking relationships is costly or if time is considered a scarce and valuable resource, then that provides substantial justification for an increased use of centralized mechanisms in matching markets.

It turns out that the length of a market's decentralized path to stability hinges critically on the market's general preference characteristics. When the two sides of a matching market are fundamentally opposed, with agents on one side preferring those agents from the other side that do not prefer them, the market's path to stability is likely to be quite long. Even a small degree of common ground is therefore vital if stability is to be attained in both a decentralized and timely manner. Just a bit of positive intercorrelation across the preferences of a market's two sides can mean a relatively short path to stability.

Moreover, those markets with fewer stable outcomes tend to converge more rapidly to stability. A high degree of correlation in the preferences of one or both sides of the market is therefore also instrumental for a shorter path. If both genders' preferences are perfectly correlated, for example, then there is obviously only one possible stable outcome and the market is likely to reach it very quickly without any external coordination. A note of irony, however, is that in this case stability presents the least gain for the market as a whole. When both genders' preferences are perfectly correlated, no matter what matching is attained one member will always receive their first choice, one will always receive their second choice, and so on. Thus, in the case when stability is easiest to reach in a decentralized fashion, it matters least to the market's overall welfare.

Finally, in addition to characterizing which types of matching markets are most in need of centralized mechanisms to facilitate their progress toward stability, this paper also reveals an interesting avenue for future theoretical research.

In the simulation experiments presented here, agents behave according to their own true preferences, always making an improvement if the opportunity presents itself. The potential length of the market's path to stability, however, may mean that some agents are better off quitting while they are ahead, leaving the market once they obtain an adequate mate. There may be some strategic element to agents' decisions regarding which blocking pairs they are willing to be a part of. Though previous work has studied similar questions of acceptance rules for the extreme cases of perfectly correlated (Eeckhout, 1999; Alpern and Reyniers, 2005) and perfectly intercorrelated (Alpern and Reyniers, 1999) preferences, the question remains open for the general case of preferences and in the more general decentralized matching environment presented here. The potential for such manipulation is far from trivial, since an outcome that is stable to the market's true preferences may then be prevented from occurring. Thus, the question of how a decentralized matching market may evolve in the presence of more sophisticated behavior is not addressed in this paper, but presents a research agenda for the future.

References

[1] Abdulkadíroğlu, A, P. Pathak & Roth, A.E. (2005) The New York City High School Match. *American Economic Review*, *95*(2), 364–367.

[2] Abdulkadíroğlu, A, P. Pathak, Roth, A.E. & T. Sönmez (2005) The Boston Public Schools Match. *American Economic Review*, *95*(2), 368–371.

[3] Alpern, S., D. Reyniers (1999) Strategic Mating with Homotypic Preferences. *Journal of Theoretical Biology*, *198*, 71–88.

[4] Alpern, S., D. Reyniers (2005) Strategic Mating with Common Preferences. *Journal of Theoretical Biology*, *237*, 337–354.

[5] Boudreau, J.W., V. Knoblauch (2007) Marriage Matching and Intercorrelation of Preferences. *Working Paper*, University of Connecticut.

[6] Caldarelli, G., A. Cappoci (2001) Beauty and Distance in the Stable Marriage Problem. *Physica A*, *300*, 325-331.

[7] Celik, O. B., V. Knoblauch (2007) Marriage Matching with Correlated Preferences. *Working Paper*, University of Connecticut.

[8] Eeckhout, J. (1999) Bilateral Search and Vertical Heterogeneity. *International Economic Review*, *40* (4), 869–887.

[9] Eriksson, K., O. Häggström (2008) Instability of Matchings in Decentralized Markets with Various Preference Structures. *International Journal of Game Theory*, *36*, 409–420.

[10] Gale, D., L.S. Shapley (1962) College Admissions and the Stability of Marriage. *Amer. Math. Monthly*, *69*(1), 9–15.

[11] Klaus, B., F. Klijn (2007) Paths to Stability for Matching Markets with Couples. *International Journal of Game Theory*, *36*, 473–488.

[12] Knoblauch, V. (accepted 2008) Marriage matching and gender satisfaction. *Social Choice and Welfare*, *forthcoming* DOI 10.1007/s00355-008-0303-2.

[13] Kojima, F., M. Utku Ünver (2008) Random Paths to Pairwise Stability in Many-to-many Matching Problems: A Study on Market Equilibration. *International Journal of Game Theory*, *36*, 473–488.

[14] Roth, A.E. (2008) Deferred Acceptance Algorithms: History, Thought, Practice and Open Questions. *International Journal of Game Theory*, *36*, 537-569.

[15] Roth, A.E., E. Peranson (1999) The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. *American Economic Review*, *89*(4), 748–780.

[16] Roth, A.E., M. Sotomayor (1990) *Two-sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press, Cambridge, MA.

[17] Roth, A.E., J.T. Vande Vate (1990) Random Paths to Stability in Two-Sided Matching. *Econometrica*, *58* (6), 1475–1480.

[18] Teo, C.P., J. Sethuraman, & Tan, W.P. (2001) Gale-Shapley Stable Marriage Problem Revisited: Strategic Issues and Applications. *Management Science*, *47*(9), 1252–1267.

[19] Wilson, L.B. (1972) An Analysis of the Marriage Matching Assignment Algorithm. *BIT*, *12*, 569-575.

6 Appendix.

The measure of preference correlation used in this paper comes directly from Celik and Knoblauch (2007). As in their paper, $Avei = \frac{1}{n}\sum_{j=1}^{n} r_{w_j}(m_i)$ is the average ranking of m_i by all n women, so the total measure of correlation in women's preferences is

$$\rho_w = \frac{\sum_{i=1}^{n}(Avei)^k - n(\frac{n+1}{2})^k}{\sum_{i=1}^{n} i^k - n(\frac{n+1}{2})^k},$$

where $k \geq 2$. All experiments here use $k = 9$, as recommended by Celik and Knoblauch (2007). The measure of correlation in men's preferences, ρ_m, is defined symmetrically.

The measure of preference intercorrelation used here comes directly from Boudreau and Knoblauch (2007). First, for each man i, square the difference between the rank m_i gives w_j and the rank w_j gives m_i and add over all women:

$$\phi_{m_i} = \sum_{j=1}^{n}(r_{m_i}(w_j) - r_{w_j}(m_i))^2.$$

Then sum across the men and divide by n to get

$$\phi_{ave} = \frac{\sum_{i=1}^{n} \phi_{m_i}}{n}.$$

Finally, the score is normalized by using the maximum possible ϕ_{ave} score that is obtained when no two men agree on the rank of any woman and each man is ranked last by his first-ranked woman, second last by his second ranked woman, third last by his third ranked woman, etc.

$$\Phi = \frac{\sum_{k=1}^{n}(n+1-2k)^2 - \phi_{ave}}{\sum_{k=1}^{n}(n+1-2k)^2}$$

Perfect positive intercorrelation therefore yields $\Phi = 1$ and perfect negative intercorrelation yields $\Phi = 0$.