NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

DISSERTATION

TWO-SIDED MATCHING IN HIERARCHICAL ORGANIZATIONS—AN APPLICATION FOR THE ASSIGNMENT OF MILITARY PERSONNEL

by

Paul Anthony Robards

June 2011

Dissertation Supervisor: William Gates

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.			
1. AGENCY USE ONLY ***(Leave blank)***	**2. REPORT DATE** June 2011	**3. REPORT TYPE AND DATES COVERED** Dissertation	
4. TITLE AND SUBTITLE: Two-Sided Matching in Hierarchical Organizations—An Application for the Assignment of Military Personnel		**5. FUNDING NUMBERS**	
6. AUTHOR(S): Paul Anthony Robards			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000		**8. PERFORMING ORGANIZATION REPORT NUMBER**	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A		**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol Number NPS.2005.0033.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited		**12b. DISTRIBUTION CODE** A	

13. ABSTRACT (maximum 200 words)

Many large organizations rely on manual assignment processes, despite the theory of bounded rationality indicating that time and cognitive constraints would limit the quality of assignments. This research used participant experiments to explore the effect of information load on assignment quality: participants, motivated by induced value theory, performed the role of decision makers; and information load was identified by the number of personnel requiring assignment and the number of attributes to be considered. Results varied considerably between participants, despite a relatively homogenous group of participants and low information loads compared to what would be experienced in actual military assignment processes. Having analyzed the shortcomings of manual assignment processes, this research examined two-sided matching as the basis for a decision support system. It was demonstrated that two-sided matching could be used to assign personnel to positions in hierarchies. Multi-attribute utility functions were used to generate position preferences based on a variety of attributes, some relevant to the organization and others to its subordinate units. Computational experiments showed that assignments are responsive to the utility function weights, allowing decision makers to quickly examine various assignment sets under different conditions. The effects of preference list indifference on two-sided matching were also examined.

14. SUBJECT TERMS Two-sided matching; military assignments; decision making; information load; multi-attribute utility functions; preference list indifference.			**15. NUMBER OF PAGES** 287
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UU

NSN 7540-01-280-5500 Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution unlimited

TWO-SIDED MATCHING IN HIERARCHICAL ORGANIZATIONS—AN APPLICATION FOR THE ASSIGNMENT OF MILITARY PERSONNEL

Paul A. Robards
Lieutenant Colonel, Australian Regular Army
B.ENG., University of New South Wales, 1993
M.S.M., Naval Postgraduate School, 2001

Submitted in partial fulfillment of the
requirements for the degree of

DOCTOR OF PHILOSOPHY IN INFORMATION SCIENCES

from the

NAVAL POSTGRADUATE SCHOOL
June 2011

Author: _____
Paul A. Robards

Approved by:

_____	_____
William Gates Professor of Business & Public Policy Dissertation Committee Chair & Supervisor	Mark Nissen Professor of Information Sciences
_____	_____
Dan Boger Professor of Information Sciences	Dan Dolk Professor of Information Sciences
_____	
Peter Coughlan Professor of Business & Public Policy	

Approved by: _____
Dan Boger, Chair, Department of Information Sciences

Approved by: _____
Douglas Moses, Vice Provost for Academic Affairs

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Many large organizations rely on manual assignment processes, despite the theory of bounded rationality indicating that time and cognitive constraints would limit the quality of assignments. This research used participant experiments to explore the effect of information load on assignment quality: participants, motivated by induced value theory, performed the role of decision makers; and information load was identified by the number of personnel requiring assignment and the number of attributes to be considered. Results varied considerably between participants, despite a relatively homogenous group of participants and low information loads compared to what would be experienced in actual military assignment processes. Having analyzed the shortcomings of manual assignment processes, this research examined two-sided matching as the basis for a decision support system. It was demonstrated that two-sided matching could be used to assign personnel to positions in hierarchies. Multi-attribute utility functions were used to generate position preferences based on a variety of attributes, some relevant to the organization and others to its subordinate units. Computational experiments showed that assignments are responsive to the utility function weights, allowing decision makers to quickly examine various assignment sets under different conditions. The effects of preference list indifference on two-sided matching were also examined.

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

The effective assignment of personnel to positions (jobs) has been described as a crucial function for organizations. However, many large organizations rely on manual processes to generate assignments, despite the theory of bounded rationality indicating that time and cognitive constraints would limit the quality of assignments produced. While the effect of information overload has been examined in fields such as consumer behavior, it has not previously been examined in relation to assignment decision making, which is arguably more complex because each individual assignment affects the remaining options available. Experiments were used in this research to explore the effect of information overload in assignment processes conducted without decision support, with participants performing the role of assignment decision makers. These experiments revealed that: decision makers tend to focus on one group (positions) to the detriment of the other (personnel); there is considerable variability in decision quality, despite the use of a relatively homogenous group of decision makers in the experiments; a sequential heuristic was used by some decision makers, resulting in poorer outcomes for the last personnel to be assigned; and decision makers are likely to over-estimate their abilities.

Having identified the problem with existing assignment processes that lack decision support, a two-sided matching algorithm was explored as the basis for a decision support system. Two-sided matching was selected because it produces stable results, and under certain conditions it is a dominant strategy for participants to submit their true preferences.

Two-sided matching has been successfully applied in a number of applications, such as the National Resident Matching Program (NRMP). Most existing applications are characterized as markets; that is, they involve the assignment of autonomous participants (e.g., medical students and residency programs). Some more recent applications, such as the New York City high school match, demonstrate some features of hierarchical organizations; that is, subordination exists, and the organization and its subordinate elements (e.g., Department of Education and schools) each have an interest in the assignment of students. Using the example of military personnel assignments, this

research demonstrates that multi-attribute utility functions may be used to combine attributes that are relevant to the organization and its subordinate units. Computational experiments were used to demonstrate that two-sided matching assignments are responsive to the utility function weights, and therefore it is possible to produce assignments that will tend to favor some attributes over others. However, the responsiveness of the assignments to the utility function weights depends on the correlation of preferences with respect to each attribute.

Two-sided matching requires rank-ordered preference lists from all participants (e.g., the personnel and positions in a military assignment process). When the positions' preferences are determined through multi-attribute utility functions, many preference ties are likely, due to the limited number of attributes available and the typically discrete nature of these attributes. When preference list indifference exists and ties are broken randomly, the two-sided matching assignments remain weakly stable. Further, different results are possible from different tie breakings. While existing applications typically require participants to submit strictly ordered preference lists, this research examined how assignments could be improved by exploiting preference list indifference. This is particularly relevant to hierarchical organizations where subordination exists. Computational experiments were used to demonstrate the range of results that are possible when tied preferences are broken in different ways. Each experiment involved multiple trials, with tied preferences randomly broken at each trial. These experiments demonstrated how assignments are influenced by the number of agents seeking assignment, their preference list length, the degree of indifference in their preferences, and the extent to which their preferences are correlated.

It is not suggested that a decision support system should replace career managers or that such systems should produce assignments autonomously. However, a decision support system could quickly provide career managers with a range of options. Such a system would ensure a consistent consideration of attributes and with less variability than occurs in the absence of decision support. Further, if the decision support system is based on a two-sided matching algorithm, the assignments will be stable and it will be a dominant strategy for personnel to submit truthful preferences.

TABLE OF CONTENTS

I. INTRODUCTION 1
- **A. MARKET DESIGN 1**
- **B. THE ASSIGNMENT OF PERSONNEL 4**
- **C. THE COMPLEXITY OF ASSIGNMENT DECISIONS 5**
- **D. ALTERNATIVE ASSIGNMENT MECHANISMS 7**
- **E. APPLYING TWO-SIDED MATCHING 8**
- **F. RESEARCH CONTRIBUTIONS 10**
- **G. DISSERTATION OUTLINE 11**

II. LITERATURE REVIEW 13
- **A. HUMAN DECISION MAKING 13**
  - **1. Rational Behavior and Bounded Rationality 15**
  - **2. Prior Investigation of Information Overload 16**
- **B. TWO-SIDED MATCHING 21**
  - **1. The Stability of Matches 24**
  - **2. The Nature of Matches Generated by the Deferred-Acceptance Algorithm 26**
    - ***a. One-to-One Matching 26***
    - ***b. Many-to-One Matching 29***
  - **3. Determining an "Optimal" Stable Match 31**
  - **4. Market Complexities 32**
  - **5. Indifference / Preference List Ties 33**
  - **6. Applications 37**
- **C. MULTI-ATTRIBUTE UTILITY THEORY 39**
  - **1. Decision-Making Techniques 41**
    - ***a. Additive Value Function 41***
    - ***b. Weighted Sum Model and Weighted Product Model 42***
    - ***c. Analytical Hierarchy Process 42***
    - ***d. SMART, SMARTS and SMARTER 43***
    - ***e. ELECTRE 45***
  - **2. Determination of Weights 45**
    - ***a. Direct Allocation 45***
    - ***b. Swing Weights 46***
    - ***c. Rank Weights 46***
    - ***d. Tradeoff Weights 48***
  - **3. Multi-Attribute Utility Theory Summary 48**

III. RESEARCH DESIGN 51
- **A. HUMAN DECISION-MAKING QUALITY IN ASSIGNMENT PROCESSES 51**
  - **1. Information Load 52**
  - **2. Design 55**
  - **3. Experiment Setting 62**

B. PREFERENCE LIST INDIFFERENCE IN TWO-SIDED MATCHING....64
1. The Likelihood of Preference List Indifference in Large Hierarchical Organizations....65
2. Simulation Model....67
3. Model Parameters....70
4. Model Development....75
5. Exploring Outcomes....78
C. DEVELOPING TWO-SIDED MATCHING IN A HIERARCHICAL ORGANIZATION....82
1. Agent Specification....84
2. Generating Preferences for Positions....88
a. Rank Utility....89
b. Branch Utility....90
c. Skill Utility....91
d. Locality Utility....91
e. Career Development Utility....91
f. Performance Utility....91
3. Simulating the Preferences for Personnel....92
a. Rank Utility....93
b. Employment Utility....94
c. Locality Utility....94
d. Career Development Utility....95
e. Performance Utility....95
f. Determining Attribute Weights....95
4. Exploring Outcomes....96
IV. DECISION-MAKING ABILITY IN ASSIGNMENT PROCESSES....97
A. MEASURING COMPLEXITY....98
B. RESULTS....98
1. Subject Demographics....98
2. Excluded Data....99
3. Overview of Results....100
4. Comparing Results Between Classes....103
a. Scenario A....104
b. Scenario B....105
c. Scenario C....106
d. Scenario D....107
e. Scenario E....108
f. Scenario F....109
g. Summary of Class Results....110
5. Trading Off Between Personnel and Positions....115
6. The Effect of More Attributes to Consider....120
7. The Effect of More Personnel to Assign....123
8. Subject Responses....126
a. Self-Assessment of Performance....126

b. Decision-Making Strategies *128*
c. Feedback on Complexity *131*
C. SUMMARY 134

V. THE EFFECT OF PREFERENCE LIST INDIFFERENCE ON TWO-SIDED MATCHING OUTCOMES 137
A. THE EFFECTS OF PREFERENCE LIST INDIFFERENCE 141
1. Number of Agents Assigned 142
2. Average Preference Rank of Assigned Agents 144
3. Determining the Required Number of Trials 146
B. THE EFFECT OF PREFERENCE LIST INDIFFERENCE ON DIFFERENT TYPES OF SCENARIOS 150
1. Preference List Indifference and Number of Agents 150
2. Preference List Indifference and Preference List Length 155
3. Preference List Indifference and Preference Correlation 161
C. SUMMARY 169

VI. APPLYING TWO SIDED MATCHING IN A HIERARCHICAL ORGANIZATION 173
A. PROBLEM SPECIFICATION 173
B. RESULTS 176
1. Scenario 1 176
2. Scenario 2 181
3. Analysis of Scenarios 1 to 8 184
C. OTHER IMPLEMENTATION CONSIDERATIONS 193
1. Two Assignment Rounds 193
2. Limiting Position Preference Lists Based on Utility 196
3. Other Considerations 198
D. SUMMARY 199

VII. SUMMARY, CONTRIBUTIONS, AND FURTHER RESEARCH 201
A. CONTRIBUTIONS 201
1. Human Decision Making 201
2. The Effects of Indifference on Two Sided Matching Assignments 203
3. Using Two Sided Matching With Preferences Determined from a Multi-Attribute Utility Function 206
B. FURTHER RESEARCH 209
C. SUMMARY AND CONCLUSIONS 212

APPENDIX A 217
A. ASSIGNMENT EXPERIMENT—INSTRUCTIONS 217

APPENDIX B. 223
A. DISTRIBUTION OF PERSONNEL AND POSITIONS FOR SIMULATION 223

APPENDIX C. 231

A. DECISION-MAKING RESULTS—POINTS FOR EACH SUBJECT AND SCENARIO 231

APPENDIX D. 237

A. ASSIGNMENT EXPERIMENT SCENARIOS 237

LIST OF REFERENCES 251

INITIAL DISTRIBUTION LIST 263

LIST OF FIGURES

Figure 1. Representation of Consumer Decision....53
Figure 2. Representation of Assignment Decision....54
Figure 3. Example of Agents and Homophily....69
Figure 4. Demonstration of Preference Correlation Parameter....72
Figure 5. Low Correlation....73
Figure 6. High Correlation....74
Figure 7. Screenshot of Form for Creating a Scenario to Investigate Indifference....76
Figure 8. Example of Randomly Generated Position Quality Values....76
Figure 9. Example of Position Preference List....77
Figure 10. Example Results from Indifference Simulation....78
Figure 11. Summary of Results by Subject....101
Figure 12. Results for Scenario A: 5 Personnel, 3 Attributes....105
Figure 13. Results for Scenario B: 5 Personnel, 5 Attributes....106
Figure 14. Results for Scenario C: 10 Personnel, 3 Attributes....107
Figure 15. Results for Scenario D: 10 personnel, 5 Attributes....108
Figure 16. Results for Scenario E: 15 personnel, 3 Attributes....109
Figure 17. Results for Scenario F: 15 Personnel, 5 Attributes....110
Figure 18. Number of times Personnel Slammed in Scenario A: 5 Personnel, 3 Attributes....112
Figure 19. Number of times Personnel Slammed in Scenario B: 5 Personnel, 5 Attributes....112
Figure 20. Number of times Personnel Slammed in Scenario C: 10 Personnel, 3 Attributes....113
Figure 21. Number of times Personnel Slammed in Scenario D: 10 Personnel, 5 Attributes....113
Figure 22. Number of times Personnel Slammed in Scenario E: 15 Personnel, 3 Attributes....114
Figure 23. Number of times Personnel Slammed in Scenario F: 15 Personnel, 5 Attributes....114
Figure 24. Personnel and Position Points – Scenario A....116
Figure 25. Personnel and Position Points – Scenario B....117
Figure 26. Personnel and Position Points – Scenario C....117
Figure 27. Personnel and Position Points – Scenario D....118
Figure 28. Personnel and Position Points – Scenario E....118
Figure 29. Personnel and Position Points – Scenario F....119
Figure 30. Self Assessment of Performance (Average Quintile) Compared to Optimal Outcome for Each Scenario....134
Figure 31. Variability within Identically Defined Scenarios – Number of Agents Assigned....142
Figure 32. Preference Ranking of Assigned Agents – 1,000 Trials of Scenario 3....145
Figure 33. Best Average Person and Position Preference Ranking for Each Scenario....146
Figure 34. Frequency of Trials Showing Number of Agents Assigned....148

Figure 35. Percentage of Trials Showing Number of Agents Assigned 148
Figure 36. Best Preference Rankings of Assigned Agents – 1,000 Trials v 25,000 Trials 150
Figure 37. Frequency of Number of Assignments – n = 2,000 152
Figure 38. Frequency of Number of Assignments – n = 1,000 153
Figure 39. Frequency of Number of Assignments – n = 500 153
Figure 40. Frequency of Number of Assignments – n = 100 154
Figure 41. Effect of DOI on Range of Number Agents Assigned 155
Figure 42. Effect of DOI on Range of Agents Assigned as Percentage of Number of Agents 155
Figure 43. Effect of Preference List Length on Number of Agents Assigned 157
Figure 44. Effect of DOI and L on Range of Outcomes 158
Figure 45. Average Preference Rank of Assigned Agents - 1,000 Trials for L = 5, DOI = 0.1 159
Figure 46. Average Preference Rank of Assigned Agents for L = 5.....159
Figure 47. Average Preference Rank of Assigned Agents for L = 10.....160
Figure 48. Average Preference Rank of Assigned Agents for L = 20.....160
Figure 49. Effect of Preference Correlation on Number of Agents Assigned 163
Figure 50. Preference List Length Frequency for C = 0.01 164
Figure 51. Effect of DOI and C on Range of Outcomes 165
Figure 52. Average Preference Rank of Assigned Agents for C = 0.0025 166
Figure 53. Average Preference Rank of Assigned Agents for C = 0.005 166
Figure 54. Average Preference Rank of Assigned Agents for C = 0.01 167
Figure 55. Percentage of Position Agents Assigned by Preference List Length.....168
Figure 56. Number of Position Agents by Preference List Length and Assignment Outcome.....168
Figure 57. Frequency of Assignments – 1,000 Trials for Scenario 1.....178
Figure 58. Average Personnel Preference Rank – 1,000 Trials for Scenario 1.....178
Figure 59. Preference Rank of Assigned Personnel Agents - Scenario 1, Trial 903.....179
Figure 60. Utility of Assigned Position Agents – Scenario 1, Trial 903.....180
Figure 61. Percentage of Assignments that Satisfy Position Attributes – Scenario 1, Trial 903.....181
Figure 62. Number of Assignments by Personnel Preference Rank – Scenarios 1 to 8 .185
Figure 63. Personnel Preference Rank by Percentage of Assignments – Scenarios 1 to 8.....186
Figure 64. Percentage of Assignments that Satisfy each Positions' Attribute – Scenarios 1,3, 5 and 7 190
Figure 65. Percentage of Assignments that Satisfy each Positions' Attribute – Scenarios 2, 4, 6 and 8 190
Figure 66. Percentage of Assignments that Satisfy each Positions' Attributes - Comparison of Scenario 2 and the Two-round Process.....195

LIST OF TABLES

Table 1 Human Decision-Making Literature 14
Table 2 Two-Sided Matching Literature 22
Table 3 Multi-Attribute Utility Theory Literature 40
Table 4 Weights According to Rank Based Weighting Techniques 47
Table 5 Assignment Instances to be Examined. 57
Table 6 Position Attributes 60
Table 7 Information Provided for Market Instance A 61
Table 8 Example Participant Preference List 81
Table 9 Sample Personnel Data 87
Table 10 Sample Position Data 88
Table 11 Rank Utility Table 89
Table 12 Branch Utility Table 90
Table 13 Performance Utility Table 91
Table 14 Locality Preferences 94
Table 15 Rank Order Centroid Weights for Five Attributes 96
Table 16 Assignment Scenarios 100
Table 17 Scenario A – Presentation Order by Class 104
Table 18 Scenario B – Presentation Order by Class 105
Table 19 Scenario C – Presentation Order by Class 106
Table 20 Scenario D – Presentation Order by Class 107
Table 21 Scenario E – Presentation Order by Class 108
Table 22 Scenario F – Presentation Order by Class 109
Table 23 Slope of Line of Best Fit for Scenarios A to F 119
Table 24 The Effect of More attributes on Scenario Outcomes 121
Table 25 Class Median Points as Percent of Optimal Result 125
Table 26 Comparison of Self-Assessed Performance v Actual Performance – All Classes 127
Table 27 Comparison of Self-Assessed Performance v Actual Performance – Class 1 127
Table 28 Comparison of Self-Assessed Performance v Actual Performance – Class 2 127
Table 29 Comparison of Self-Assessed Performance v Actual Performance – Class 3 128
Table 30 Decision Making Strategies 129
Table 31 Aggregate Points Across all Scenarios – Strategy Group 1 (n = 18) 130
Table 32 Aggregate Points Across all Scenarios – Strategy Group 2 (n = 10) 130
Table 33 Aggregate Points Across all Scenarios – Strategy Group 3 (n = 52) 130
Table 34 Subjects Responses to Question Regarding Scenario Complexity 131
Table 35 Parameters and Variables of Interest 140
Table 36 Results from 25,000 Trials – Frequency of Number of Agents Assigned 149
Table 37 Summary of Key Findings from Preference List Indifference Experiments 171
Table 38 Scenario Specifications 176

Table 39 Personnel Assignments by Preference Rank – Scenario 1 and 2 182
Table 40 Position Assignments by Utility – Scenario 1 and 2 183
Table 41 Achievement of Position Attributes – Scenario 1 and 2 184
Table 42 Number of Assignments – Scenarios 1 to 8 184
Table 43 Effect of Increasing L_{pers} on Personnel Assigned at Scenario 1 187
Table 44 Preference Rank of Personnel Assigned at Scenario 2 but not at Scenario 1 187
Table 45 Number of Assignments Satisfying Unit Attributes – $L_{pers} = 10$ 188
Table 46 Number of Assignments Satisfying Unit Attributes – $L_{pers} = 20$ 188
Table 47 Number of Assignments Satisfying Service Attributes – $L_{pers} = 10$ 188
Table 48 Number of Assignments Satisfying Service Attributes – $L_{pers} = 20$ 189
Table 49 Number of Assignments that Satisfy Specified Number of Position Attributes 192
Table 50 Number of Assignments Satisfying Indicated Number of Attributes – Scenario 2 197
Table 51 Number of Assignments Satisfying Indicated Number of Attributes – Modified Scenario 2 197

LIST OF ACRONYMS AND ABBREVIATIONS

AHP	Analytical Hierarchy Process
DOI	Degree of Indifference
DSS	Decision Support System
EPAR	Employment Preferences And Restrictions
EMPLID	Employee ID
GPA	Grade Point Average
M-Optimal	Male Optimal Two-Sided Match
NIMP	National Intern Matching Program
NRMP	National Resident Matching Program
ROC	Rank Order Centroid
ROL	Rank Order List
RR	Rank Reciprocal
RS	Rank Sum
SMART	Simple Multi Attribute Rating Technique
SMART	Simple Multi Attribute Rating Technique Using Swing Weights
SMARTER	Simple Multi Attribute Rating Technique Exploiting Ranks
W-Optimal	Woman Optimal Two-Sided Match
WPM	Weighted Product Model
WSM	Weighted Sum Model

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGEMENT

To my wonderful wife Kylie, and children Amelia, Victoria, William and James; this has been a long journey, and it would not have been possible without your support, understanding and patience. I cannot thank you enough.

A special thanks to my supervisor, Professor Bill Gates; you have provided fantastic support, advice and encouragement. I am incredibly fortunate to have had your guidance through this process. Thank you to each of my dissertation committee members; I appreciate all of the assistance that you have provided.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

The "assignment of individuals to positions in a manner that makes use of their skills while providing adequate job satisfaction is crucial to the smooth and economical function of the organizations." (Klingman & Phillips, 1984, p. 1362)

A. MARKET DESIGN

Markets exist in a variety of forms to facilitate the exchange of information, goods, services and payments (Bichler, Field, & Werthner, 2001). While the existence of markets can be traced back over the centuries to barter economies, the design of markets is regarded as a new discipline that has been emerging since the 1990s (Roth, 2002). Large markets that have undergone successful design during this period include the (re)design of the National Resident Matching Program (NRMP) (Roth & Peranson, 1999), and the design of auctions for the sale of rights to the radio spectrum by the U.S. Federal Communications Commission (Banks, Olson, Porter, Rasssenti & Smith, 2003). However, market design is not necessarily a simple matter; poor design, testing and implementation can lead to market failure as indicated by the failure of some regional British residency markets (Roth, 1991), and the more spectacular failure of the Californian electric market (Crow, 2002), which had significant consequences on customers and almost led to bankruptcy of the utility companies.

This research focuses on the design of a labor clearinghouse—a market in which personnel are assigned to positions, where positions are the jobs or appointments within an organization. Labor clearinghouse assignments have been used in a variety of contexts, including the allocation of medical residents to hospital residency positions (Roth, 1991), the allocation of law students to U.S. federal judicial clerkships (Avery, Jolls, Posner & Roth, 2001), the allocation of students to New York City high schools (Abdulkadiroğlu, Pathak & Roth, 2005), and the allocation of personnel to tasks in manufacturing (Zulch, Rottinger & Vollstedt, 2004).

A well-designed labor clearinghouse can benefit both participants and the administering system: for example, after the redesign of the New York City high school

assignment process involving over 90,000 students annually, the number of students receiving their first preference increased from approximately 24,000 to 31,000, and the number of administrative assignments (cases where students were assigned to schools that they did not list) managed by the Department of Education decreased from 30,000 to 3,000 (Abdulkadiroğlu et al., 2005). However, Roth (1991) provides examples of labor clearinghouses that have failed, and indicates that to avoid such failures, it is necessary to ensure that the design of the labor clearinghouse meets the needs and nature of the market.

In the broad use of the term, a market is "any place or mechanism in which goods and services are traded" (Thompson, 2001, p. 3). However, it is useful to distinguish between the different types of institutions in which labor clearinghouses may exist. Williamson (1975) refers to two institutional forms that exist for exchange; markets and hierarchies. He defines a market more narrowly as an exchange between autonomous economic entities, which is in contrast to hierarchical transactions where a "single administrative entity spans both sides of the transaction" and subordination exists in some form (Williamson, 1975, p. xi). This delineation between markets and hierarchies closely resembles the distinction between external and internal labor markets respectively. Doeringer and Piore (1985) define an external labor market as an exchange in which the allocation of labor is determined from "conventional economic theory where pricing, allocating and training decisions are controlled directly by economic variables," while an internal labor market occurs where "the pricing and allocation of labor is governed by a set of administrative rules and procedures" (p. 2). When Doeringer and Piore (1985) and others refer to internal labor markets, it is clear that these are not markets as defined by Williamson (1975), but rather the exchanges that occur within a hierarchy. This research is predominantly focused on internal labor market clearinghouses, that is, those that exist in hierarchies.

The nature of the institution, whether it is a market or hierarchy, determines the signals that are produced in the event of assignment inefficiency or instability. In external labor assignment markets, inefficiencies or instability can lead to failure that is often evident and manifested by declining participation or market unraveling; situations that

have been explored by Roth (1991) and Roth and Peranson (1999). However, in a hierarchical organization where participation is compulsory, such direct manifestations cannot occur to reveal the extent of market inefficiencies, and hence market inefficiencies may only appear indirectly; for example, lower productivity or higher attrition due to dissatisfaction. Therefore, managers of hierarchical organizations may not be aware of inefficiencies in their assignment processes, and may incorrectly attribute indicators such as higher attrition to other phenomena.

While the nature of hierarchical institutions may mask the signals that reveal inefficiencies in labor assignments, the problem is exacerbated if these hierarchies use manual assignment processes, since bounded rationality (Simon, 1955) limits the ability of decision makers, and therefore assignments are likely to be inefficient. Here the efficiency is used in the economic sense and refers to the value that is available to the participants, both personnel and positions. An inefficient outcome is one where it is possible for greater value or utility to be provided to the participants; for example, personnel being assigned to preferable positions, and / or positions being assigned personnel that better meet requirements.

Militaries are exemplars of hierarchical institutions where, for uniformed members, participation in the labor assignment process is compulsory, and assignments are typically conducted manually or with limited decision support: Blanco and Hillery (1994) state that, until the 1980s, for the U.S. Navy, “the process of matching people to jobs and making assignment decisions was exclusively a manual process” (p. 814). While some decision support has since been provided, Short (2000) reveals that “there is no single tool to help the detailers ‘mentally juggle’ diverse policies, procedures, and information” (p. 22), and so the process remains bounded by the cognitive ability of the decision makers.

The preceding discussion highlights two problems that may exist in a hierarchical organization that uses manual labor assignment processes; firstly, the assignments are likely to be inefficient due to bounded rationality, and secondly, the managers of such organizations may not be aware of the full extent of the problem. This research addresses these two issues. Experimental work examines the nature of manual labor assignments to

determine the conditions where human decision makers may perform well, and conditions where bounded rationality may cause inefficiency. Then, for situations where human decision makers perform poorly in the absence of decision support, computational experimentation is used to examine how a decision support system utilizing two-sided matching processes may be used to improve the efficiency of the assignments.

B. THE ASSIGNMENT OF PERSONNEL

It is frequently acknowledged that people are an organization's most important asset; Sikula (2001) states, "It is unquestionably true that human resources are every organization's most important and most valuable assets. This has been recognized for decades now, but it still is lacking in practice in business" (p. 421). In order to achieve best practice with human resources, organizations must address the various functions that are involved, and depending on organization size, these may include recruiting, training, mentoring, career management, personnel assignments (or allocations), payroll and compensation. This research addresses just one of these functions; the assignment of personnel to positions. The assignment process seeks to ensure that an organization's personnel are employed in a manner that suits the needs of the organization while satisfying the individual. While an organization's personnel may have the correct skill mix, unless the assignment process is effective, the personnel may be employed in the wrong positions, thus leading to sub-optimal organizational outcomes. To meet the organization's requirements, it is necessary to have a well-designed labor assignment process.

When hiring new personnel, many organizations undertake a process of classification and selection to ensure that the personnel they employ have the requisite knowledge, skills and abilities to fill their initial appointment. While this initial classification and selection is often well planned and structured, in many organizations the subsequent reassignment of personnel is ad hoc, primarily conducted as necessitated by staff retirements, resignations, promotions etc. It is perhaps for this reason that the literature on classification and selection is quite extensive (for example, Monahan & Muchinsky (1983) reviewed 394 personnel selection studies published solely in the

Personnel Psychology journal) but limited when it comes to examining internal labor assignment processes. Despite this lack of scholarly attention, there are some hierarchies, most notably military organizations, where the internal assignment of labor is a significant and regular process, with large batches of personnel and positions considered for assignment simultaneously. In large military forces, the number of annual assignments undertaken is considerable. At a time when the U.S. Navy had approximately 500,000 enlisted personnel, Liang and Buclatin (1988) claim that the Navy assigned approximately 200,000 people to jobs each year. An annual assignment of 40 per cent of a force each year is a reasonable approximation for other services and nations where tenure in a position generally varies between one and three years.

There is evidence that military personnel feel that more could be done to incorporate their preferences in the assignment decision-making process. The Australian Defence Organisation administers the annual Defence Attitude Survey to 30% of uniformed and civilian personnel to gauge attitudes towards a variety of issues, one of which is career management. The responses to two questions reveal the satisfaction with extant assignment decision processes. From the 2009 Defence Attitude Survey, 70% of Army respondents agreed with the statement, "Individual posting preferences need to have more influence," while 50% agreed with the statement, "My views are considered when postings are planned."[1] Such responses provide the motivation for this work.

C. THE COMPLEXITY OF ASSIGNMENT DECISIONS

Not only are military assignments large-scale problems, they are also complex, where complexity is a function of the number of criteria to be considered in making the assignments. In examining the U.S. Marine Corps assignment process, Klingman and Phillips (1984) state that each assignment may consider from 3 to 15 criteria. Regarding the U.S. Navy, Liang and Buclatin (1988) state, "For each person, a maximum of five trained skills may be recorded. A job may have a requirement for up to nine skills, some of them alternative" (p. 183). Further to the consideration of skills criteria, Liang and Buclatin investigate a significant factor in the U.S. Navy assignment process; making the

[1] Australian Defence Organisation, 2009 Defence Attitude Survey.

best use of training spaces when the skills of the distributable personnel do not meet the requirements of the positions. Additional complexity is provided by the need for decision makers to consider relevant policies and the preferences of the personnel who are being assigned.

With such a large number of assignments to be made and multiple criteria to be considered for each assignment, the process is far from simple. Despite this, many military services continue to assign personnel with manual processes, or processes that involve limited system support. While such processes may yield satisfactory results, it is worth considering whether more effective results could be achieved more efficiently through a well-designed market. Eppler and Mengis (2004) review decision-making research in disciplines such as accounting, organization science and consumer behavior, and show that there is considerable evidence that humans make poor quality decisions under excessive information loads. While the effects of information load on decision quality have been examined in a number of disciplines, such behavior has not been examined in relation to assignment decision making, which it could be argued is more problematic due to preferences of participants on both sides of the market, and the need for multiple decisions and tradeoffs. Therefore, to examine the effects of information load on manual assignment decisions, the following research question is posed:

Research Question One: How does human decision quality vary in relation to the size and complexity of an assignment problem?

The results of this research will demonstrate situations where manual decision processes are appropriate, and situations where decision support systems (DSS) should be used. This process of understanding the problem helps define the crisis with the status quo, which according to Kotter (1996), is the first step to managing change.

While a DSS may be designed to improve the quality of assignments in labor markets, different assignment mechanisms are possible for such a DSS. The next section explores the relative merits of three different assignment mechanisms.

D. ALTERNATIVE ASSIGNMENT MECHANISMS

The first approach considered for assigning personnel to positions is pure binary integer programming (Render & Stair, 2000). This approach is based on a linear programming model that maximizes (or minimizes) an objective function, but is termed pure because only integer outcomes are acceptable, and binary because only 0-1 outcomes are acceptable (each person may be assigned to only one position). An integer programming approach is suitable if assignments are primarily expected to maximize an objective function, but there are disadvantages to such an approach. One disadvantage is that it is not necessarily a dominant strategy for personnel to state their true preferences under an integer programming approach, which therefore enables personnel to strategically manipulate the process to their advantage. Also, optimization may return unstable assignments; that is, assignments where a person and position are not matched, but prefer each other to their assigned partners.

Genetic algorithms provide another potential assignment approach. Genetic algorithms use biological evolutionary concepts to develop an outcome, and are best suited to complex situations. Goldberg (1989) explores how genetic algorithms differ from traditional optimization techniques, while Biethahn and Nissen (1995) provide a table outlining the advantages and disadvantages of genetic algorithms. Herrera, Lopez, Mendana and Rodriguez (1999) state, “GAs are powerful in difficult environments where the space is usually large, discontinuous, complex and poorly understood. They are not guaranteed to find the global optimum solution to a problem, but they are generally good at finding acceptably good solutions to problems quickly” (p. 311). That is, genetic algorithms are particularly suited to problems where the decision space is discontinuous and poorly defined. However, as will be described in this research, the situation within a labor assignment process may be relatively well defined, and therefore genetic algorithms do not offer significant outcome advantages over optimization techniques, although there may be efficiency gained in the process of achieving the outcome. Furthermore, like integer programming, genetic algorithms do not necessarily make truthful revelation of preferences a dominant behavior.

The third approach to the assignment problem is through algorithms that provide stable outcomes, such as deferred-acceptance two-sided matching algorithms, which are reviewed in detail by Roth and Sotomayor (1990). There are a number of advantages to using two-sided matching algorithms. First, the algorithm necessarily considers the preferences of both sides of the market. Second, under certain conditions it is a dominant behavior for participants on at least one side of the market to truthfully reveal their preferences. Third, alternative results appear when preference list ties are broken in different ways thereby enabling alternative outcomes to be explored by the decision maker. Fourth, while the algorithm does not produce an optimized outcome, it does produce a stable outcome. While this will be described in more detail later, it can be considered a situation where there is no person and position that both prefer one another to their assigned partners. A disadvantage of the two-sided matching algorithm is that unless complete preference lists are used (i.e., every person provides a rank-ordered preference list of every position and vice versa), the algorithm does not guarantee that all participants will be matched. Additionally, there is no inherent mechanism to ensure high priority positions are filled. However, different mitigation strategies are available; for example, longer preference lists and priority or secondary matching rounds.

Given the nature of the assignment problem, there are considerable benefits available to personnel who would seek to manipulate the system by misrepresenting their true preferences. Such misrepresentation under some matching mechanisms could lead to the attainment of a preferable assignment. Therefore, a key consideration is to select an assignment mechanism that makes it a dominant strategy for participants to truthfully express their preferences. Two-sided matching is an assignment mechanism that uniquely satisfies this consideration while having other desirable properties, and therefore this research will focus on the use of two-sided matching algorithms.

E. APPLYING TWO-SIDED MATCHING

Based on a comparison of the relative merits of each assignment approach, this research is focused on the use of two-sided matching algorithms as the assignment mechanism at the heart of the decision support system. These algorithms have been used

in a number of different applications as outlined by Roth (1991), Roth and Peranson (1999), Irving (2000) and Abdulkadiroğlu et al. (2005), however they have not been applied to strictly hierarchical organizations, and therefore the application is not straightforward. As indicated by the name, two-sided matching algorithms are based upon the preferences submitted by participants on two sides of a market (or hierarchy), e.g., students and schools. In a hierarchy, the workers and jobs are analogous; however, there may also be system level (organization) factors to be considered. For example, in a military organization, personnel have preferences for positions, commanders have preferences for personnel to fill positions within their command (or at least can express preferences over different criteria), and at the system level, the military service has preferences regarding the characteristics of assignments. For a large hierarchical organization with system level criteria to be considered, this paper proposes a multi-attribute utility approach to facilitate the determination of rank-ordered preference lists of personnel for each position. Such an approach ensures a consistent consideration of relevant criteria for all assignments.

Given that this paper explores the design of a new type of two-sided matching market, there are new issues to be considered as part of the design process. In a large hierarchical organization such as a military internal labor market, there are clusters of homogeneity (groups of personnel or positions that have relatively similar characteristics, at least in terms of measurable and available data), which implies that there will be clusters of indifference in preference lists. While existing applications of two-sided matching require participants to submit strictly ordered preference lists, this may not be reasonable in a hierarchical organization, and therefore it is necessary to examine the issue of preference list indifference and its impact on two-sided matching outcomes. This leads to the second research question:

Research Question Two: What effect does preference list indifference have on a two-sided matching outcome, and how can a decision maker use this to produce better outcomes?

Another issue to be examined as part of the market design is how position preferences will be determined. In external labor markets such as the National Resident

Matching Program, this is the responsibility of individual hospitals; however, such an approach is not feasible in an internal labor market where position and system level criteria must each be considered. It is proposed that a multi-attribute utility approach be used to determine the rank-ordered list of preferred personnel for each position. This approach would allow each position's preference list to include a consideration of both position and system level criteria, with weights indicating the relative importance of different criteria. It is necessary to examine the effects that such a multi-attribute approach would have on the stability of assignments, the incentive behavior of participants, and the sensitivity of assignments to differing weights. This leads to the third research question:

Research Question Three: How can a large hierarchical organization use two-sided matching when position and organization-level factors require consideration?

The second and third research questions will be addressed through computational experimentation similar to that demonstrated by Nissen and Buettner (2004). Such an approach is appropriate for the following reasons: firstly, field observation is not possible because markets of the type under examination do not currently exist; and secondly, when compared to laboratory experimentation, the computational experimentation allows a far greater number of scenarios to be examined in a relatively short time.

F. RESEARCH CONTRIBUTIONS

The first and broadest contribution of this research relates to the concept of market design; the research examines the labor assignment process within hierarchical organizations and builds the case for redesign of such processes. The research then proposes a design that may be used for military labor assignments, and may well be applicable to other hierarchical organizations.

In establishing the requirement for redesigning the labor assignment process, the research examines the effect of information load on assignment decision making in a complex situation. While the effects of information load on decision quality have been explored in other contexts, this paper makes a new contribution by specifically examining how assignment decision making is affected by complexity. Through understanding the

deleterious effects of complexity on manual assignment decision quality, it will be possible to demonstrate to managers the advantages of assignment decision support systems and how such systems may be implemented.

For managers seeking to improve the efficiency and effectiveness of assignment decision making, under certain conditions two-sided matching is a useful mechanism. However, determining position preferences is not straightforward because, in a hierarchical organization, both position level and system level criteria may be relevant. This research investigates how a multi-attribute utility approach that considers both position and system level criteria may be used to construct the preferences for the two-sided matching process, and provides indicative results and analysis. This demonstrates that two-sided matching algorithms may be applied to hierarchical institutions where system-level criteria are relevant.

While other research makes it clear that preference list indifference will have an effect on assignment outcomes, the dimensions of this effect have not been explored, nor has there been analysis of whether the outcome may be improved by allowing preference list indifference. This research demonstrates that when two-sided matching is applied to hierarchical institutions, it may be to the benefit of the decision maker to allow participants to express indifference in their preference lists.

Finally, while existing two-sided matching applications are entry-level markets, this research demonstrates the application of two-sided matching to a hierarchical organization. A military personnel assignment process is used for demonstration purposes.

G. DISSERTATION OUTLINE

The remainder of the dissertation is organized as follows. Chapter II presents a review of the relevant literature from the fields of human decision making, two-sided matching and multi-attribute utility theory. Chapter III outlines the research design that will be used to address the research questions already posed. Chapter IV explores the performance of human decision making in complex assignment problems. Chapters V and VI each examine two-sided matching issues; the impact of preference list

indifference on two-sided matching assignments, and the effect of determining preferences through a multi-attribute approach. Chapter VII provides the conclusion and recommendations for further work.

II. LITERATURE REVIEW

This chapter reviews three key fields pertinent to the research; human decision making, two-sided matching and multi-attribute utility theory.

A. HUMAN DECISION MAKING

The literature to be covered in this section of the paper is summarized in Table 1.

Reference	Relevant Concepts
Cosmides and Tooby (1994)	Identified that the human mind performs particularly well for some cognitive activities (e.g., speech recognition, grammar acquisition).
Simon (1955)	Established that human decision-making is limited by bounded-rationality.
Miller (1956)	Established that humans have a limit to the amount of information that can be stored and processed.
Arthur (1994)	Identified that there is a limit to which human logical capacity is able to cope.
Gates and Nissen (2002)	Identified that some organizations rely upon the cognitive skills of decision makers to assign personnel to jobs in a complex situation.
Eppler and Mengis (2004)	Reviewed information overload in a variety of disciplines.
Jacoby, Speller and Kohn (1974)	Found that information overload, as measured by increasing complexity, leads to poorer quality decisions in a consumer domain.
Jacoby, Speller and Berning (1974)	Confirmed the findings of Jacoby, Speller and Kohn (1974) in a replication study.
Summers (1974)	Criticized the methodology of Jacoby et al.
Wilkie (1974)	Criticized the methodology of Jacoby et al.
Jacoby (1977)	Responded to earlier criticisms of the work of Jacoby et al.

Malhotra et al. (1982)	Criticized the methodology of Jacoby et al.
Malhotra (1982)	Confirmed the effect of information overload on consumer decision making.
Keller and Staelin (1987)	Confirmed the effect of information overload when information defined by dimensions of quality and quantity.
Henry (1980)	Examined the effect of individual information processing ability on decision-making accuracy.
Houser et al. (2004)	Classified subjects by decision-making strategies in a complex decision-making scenario.
Hahn et al. (1992)	Confirmed the effect of information overload when subjects were placed under time-pressure.
Casey (1980)	Confirmed the effects of information overload on bank loan officers.
Larichev et al. (1988)	Confirmed the effects of information overload when using a classification schema to examine decision-making accuracy.

Table 1 Human Decision-Making Literature

Humans make countless decisions on a daily basis. For most people, many simple decisions are conducted frequently and require little thought (e.g., what to wear to work), while more complex decisions tend to be less frequent and require more significant consideration (e.g., the purchase of a car). But do we necessarily make the best decisions in all cases?

For certain cognitive processes, humans perform particularly well. Cosmides and Tooby (1994) state, “On evolutionary recurrent computational tasks, such as object recognition, grammar acquisition, or speech comprehension, the human mind greatly outperforms the best artificial problem-solving systems that decades of research have produced” (p. 329). However, as will be reviewed, humans do not perform as well in all cognitive aspects. Of interest to this research are the conditions where assignment decisions are likely to be degraded if conducted by humans without decision support.

1. Rational Behavior and Bounded Rationality

Simon (1955), states the traditional economic definition of a rational person as one who "is assumed to have knowledge of the relevant aspects of his environment, … and a skill in computation that enables him to calculate, for the alternative courses of action that are available to him, which of these will permit him to reach the highest attainable point on his preference scale" (p. 99). However, cognitive limits prevent humans from always achieving such a utility maximizing outcome. According to Miller (1956) in his paper on the magical number seven, "the span of absolute judgment and the span of immediate memory impose severe limitations on the amount of information that we are able to receive, process, and remember" (p. 95). Due to cognitive constraints, Simon introduced the idea that decision making is not always rational, but may be bounded, with bounds being imposed due to information or time constraints. Consequently, Simon (1955) states, "actual human rationality-striving can at best be an extremely crude and simplified approximation" (p. 101). Further, Arthur (1994) states, "There are two reasons for perfect or deductive rationality to break down under complication. The obvious one is that beyond a certain level of complexity human logical capacity ceases to cope—human rationality is bounded. The other is that in interactive situations of complication, agents cannot rely upon the other agents they are dealing with to behave under perfect rationality, and so they are forced to guess their behavior" (p. 406). While assignment decision making in the context explored here is not interactive, the first of Arthur's two reasons is relevant and will be explored.

There are a wide variety of tasks where humans are called upon to use their cognitive abilities to make assignments or allocations. Some diverse examples include the assignment of projects to students (Harper, de Senna, Vieira & Shahani, 2005), referees to games (Scarelli & Narula, 2002), court cases to judges (Yang & Dean, 1993), and personnel to positions (Suh, Byun & An, 1993). In some of these situations, only agents on one side of the relationship have preferences over potential assignments (e.g., projects to students), while in other situations all agents have preferences over potential assignments (e.g., personnel to positions). Assignment processes may also be characterized by size and complexity; size referring to the number of objects requiring

assignment, and complexity referring to the number of criteria to be considered in making assignments. Despite the existence of large and complex assignment mechanisms, Gates and Nissen (2002) find that some of these mechanisms (e.g., military forces and government organizations) rely upon administrative procedures to match people with jobs. Specifically, they state, "Hierarchical job assignments rely upon the cognitive process of centralized, administrative professionals to match individual capabilities and job requirements and to reflect both the job's relative priority and the individual's job preferences" (p. 5). While bounded rationality indicates that in such large and complex situations human decision makers will achieve less than optimal outcomes, bounded rationality does not indicate where the limits of human rational behavior lie; for example, the number of criteria or the number of objects that a decision maker can consider, and how these multiple dimensions interact. Using experimental techniques, it is possible to test the limits of human rationality under varying assignment conditions, such as decision-space size and complexity. The next section reviews studies that have examined the effects of decision making under complexity in other fields.

2. Prior Investigation of Information Overload

As a decision situation becomes more complex, there is a greater information load on the decision maker. In a comprehensive review of information overload in various disciplines, Eppler and Mengis (2004) state that the performance of a decision maker "correlates positively with the amount of information he or she receives—up to a certain point. If further information is provided beyond this point, the performance of the individual will rapidly decline" (p. 326). The effect of information overload has been studied extensively in the field of consumer research and behavior, and this field provides useful parallels to assignment decision making.

Starting with the works of Jacoby, Speller and Kohn (1974), the effects of complexity and information overload with respect to consumers has been examined by a number of researchers. In these studies, complexity is typically defined by the number of brand choices and the number of attributes describing each brand. While it is now generally acknowledged that increasing complexity leads to consumer information

overload, in turn leading to poorer quality decisions, such beliefs did not always prevail. Jacoby, Speller and Kohn (1974) defied the notion that more product information is necessarily better for consumers. They provided subjects with a variety of fictional laundry detergent brands with varying amounts of information provided about the brands. The total information was classified as the number of brands multiplied by the number of items of information per brand. Decision quality was measured as the ability of each subject to select a brand that best met their individual needs, where the individual needs were ascertained from a questionnaire conducted at the start of the experiment. A curvilinear relationship (an inverted U-shape) was found such that the ability of subjects to select a detergent that best satisfied their needs was poorest at the low and high levels of total information, while subjects provided with intermediate levels of information made the best quality decisions. Jacoby, Speller and Berning (1974) confirmed these findings in a replication study where different product categories were chosen (rice and prepared dinners), and the subjects were housewives rather than students.

The works of Jacoby, Speller and Berning (1974) and Jacoby, Speller and Kohn (1974) were criticized by Summers (1974) and Wilkie (1974), primarily on methodological grounds. Jacoby (1977) responded, and summarized the two primary criticisms as; “(1) how does one define information load in the consumer context, and (2) how does one assess decision quality?” (p. 570). In the conclusion of his rejoinder, Jacoby (1977) states, “though our critics may disagree with our methods, measures, and analyses, they do not contest the existence of information overload in the consumer context” (p. 572).

Malhotra, Jain and Lagakos (1982) state that they have three major reservations regarding the work of Jacoby et al., and that these reservations are similar to those raised previously by Summers (1974) and Wilkie (1974). First, Malhotra et al. (1982) argue that “the number of brands and the number of information items (attributes) per brand are conceptually as well as psychologically different dimensions” (p. 29), and therefore Jacoby et al., should not have defined the total amount of information as the product of the dimensions. Second, Malhotra et al. (1982) state that “the probability of making a correct choice by chance alone is inversely proportional to the number of brands” (p. 29),

and so the effects of chance factors should have been included. Third, Malhotra et al. (1982) question the approach "of using weighted Euclidean distances of the brands in the choice set from the ideal brand to measure correct choice" (p. 29), because of the problems with ideal point measurement. Despite the methodological criticism of Jacoby et al., in another article Malhotra (1982) does claim to find strong evidence for information overload effects.

Given the preceding controversy regarding information overload and techniques for examining it, Keller and Staelin (1987) took a new approach of defining information in the consumer context by the dimensions of quality and quantity, with the perspective that information overload occurs "when the incremental decreases in decision effectiveness due to additional information quantity are greater than the incremental increases in decision effectiveness due to additional information quality" (p. 202). They too found that too much information causes a decrease in decision effectiveness. Therefore, while there have been a variety of criticisms regarding methods used to test for information overload, there is at least general agreement that the information overload effect does exist.

Henry (1980) also examined information processing accuracy in the consumer behavior domain, however his approach was different for two primary reasons. Firstly, he examined the effect of an individual's information processing ability on processing accuracy. Secondly, he used a more detailed concept of complexity; defined as an interplay between the number of brands, number of attributes per brand and the scaling employed by each attribute. Based on his results, Henry produced a linear regression equation of the relationship between the dependent variable, information processing accuracy, and the independent variables; complexity (measured as the number of dimensions presented) and individual processing ability (measured by a score in a paragraph completion test). The coefficient for complexity was negative and statistically significant, confirming the information overload effect, while the coefficient for individual processing ability was positive and statistically significant, indicating the higher accuracy of people of higher processing ability. Henry also claims that when individual cognitive limits are approached, consumers attempt to limit stress by

simplifying the problem, e.g., restricting either the number of alternatives considered or the number of attributes considered, or simplifying the attribute scaling.

While Henry (1980) examined the effect of the decision makers' ability on decision outcome, Houser, Keane and McCabe (2004) sought to classify human subjects by the behavioral strategies that they employed in a complex situation. From the results of controlled laboratory experiments, subjects were classified into three types according to their behavior; those that are "near-rational" and achieve outcomes close to optimal, those that are "fatalist" and give up on achieving a near-optimal strategy if they make some poor early choices, and those that are "confused" and irrational in their decision making. They confirmed that people use different decision strategies and that outcomes vary between decision makers.

Like Jacoby et al., Hahn, Lawson and Lee (1992) also found an inverted U-shape function relating decision quality to information load, however Hahn et al., claim that this inverted U-shape only appears when the subjects are under time pressure. In the absence of time pressure, they claim that subjects continued to show improvements in decision quality as the number of attributes increased to a maximum of 20 (for 10 alternatives). While Hahn et al., have an interesting hypothesis regarding the effects of time pressure, there are concerns regarding the methodology, including: the selection of subjects (high school students), an apparent lack of interest from the subjects (the authors state, "72 subjects [17%] were excluded due to failure to follow the instructions seriously" (p. 371)), a confounding of one of the focus variables with type of student (groups comprised either students from vocational high schools or from college preparatory schools), and poor / unclear experimental rewards (rewards appeared to have been unspecified at the start of the experiment other than "the persons who made the best choice would receive awards"(p. 370)). Therefore, it is questionable whether the experiments by Hahn et al., proves that information overload only appears when subjects are under time pressure.

In a financial study, Casey (1980) examined the effects of information load on bank loan officers. This study examined the ability of bank loan officers to predict bankruptcy and found that subjects with the highest information load were slower and had no greater predictive ability than those with lesser information loads. This provides

additional support for the idea that humans become overloaded by information, and although the additional information may be important and relevant, humans are unable to effectively process large volumes of information.

Larichev, Moshkovich and Rebrik (1988) used a different type of experiment to examine the effects of complexity on human decision making. They used multi-attribute classification problems whereby subjects were required to separate objects into decision classes of various sizes. The ability of subjects to perform these tasks consistently, i.e., without making errors involving the assignment of a dominating object to an inferior decision class (and vice versa), was the conceptual measure of a subject's accuracy. Complexity was defined by three parameters; the number of attributes that characterize the objects, the number of estimates of each attribute on ordinal scales, and the number of decision classes to which the objects were to be assigned. Their work confirms the effects of information overload found elsewhere, and an understanding of this limitation of human ability with regard to the classification process is useful in that it may actually be one of the intermediate cognitive steps used by decision makers to generate effective assignments. For example, to determine the suitability of personnel for a position, the decision maker may classify the personnel into classes such as 1) ideal for the position considered, 2) suitable for the position considered, and 3) unsuitable for the position considered.

While the results of experiments regarding information overload in other disciplines are informative, also of value is the definition of complexity, primarily established in the consumer research domain. Complexity and information load are typically defined by the number of alternatives available, the number of attributes used to measure the alternatives and the types of scales used to measure the attributes. In the assignment context it is possible to measure complexity with similar dimensions, however, there are additional dimensions of complexity not found in consumer behavior. Firstly, the decision maker is not seeking the best person for a single position, but rather seeks to evaluate the best person for multiple positions, which likely necessitates tradeoffs since each assignment made restricts the set of personnel and positions available

for the remaining assignments. Furthermore, in addition to the positions having preferences over the characteristics of personnel, the personnel also have preferences over the positions to which they wish to be assigned.

As outlined, a variety of experiments have examined the effects of information load on decision quality. Such experiments have been primarily in the consumer behavior domain, although examples in other areas do exist. However, none have dealt with the more complex environment of an assignment process.

B. TWO-SIDED MATCHING

The literature to be covered in this section of the paper is summarized in Table 2.

Reference	Relevant Concepts
Gale and Shapley (1962)	Provided first examination of two-sided matching markets
Roth and Sotomayor (1990)	Reviewed two-sided matching theories
Roth (1991)	Reviewed matching markets in the UK
Erdil and Ergin (2007)	Outline a 'stable improvements cycle' for dealing with preference list ties
Dubins and Freedman (1981)	Examined dominant strategies for participants in two-sided matching market
Gale and Sotomayor (1985)	Examined the possibility of strategic manipulation of preferences to obtain preferable outcome
Roth and Rothblum (1999)	Examined practical limits on ability of participants to strategically manipulate outcomes
Crawford (1991)	Examined the effect of introducing new participants into a two-sided matching market
Irving and Leather (1986)	Examined the number of stable matches in a market with complete preference lists
Roth and Peranson (1999)	Reviewed the modification of the NRMP algorithm and its effects on assignment outcomes
Roth (1985)	Examined the differences between one-to-one and many-

	to-one matching markets
Irving et al. (1987)	Examined techniques to obtain all possible stable assignments in a two-sided matching market
Roth and Vande Vate (1990)	Developed instability chain algorithm to handle two-sided matching market complexities
Irving (1994)	Introduced new definitions of stability in markets that have preference list indifference
Manlove (1999) and Iwama et al. (1999)	Examined computational complexity of finding strong and super stable matches in markets with preference list indifference
Irving et al. (2000)	Examined the probability of finding super-stable solutions in markets with preference list indifference
Gent and Prosser (2002)	Developed process to determine the smallest and largest number of matches in a market with preference list indifference
Abdulkadiroğlu et al. (2007)	Examined the tradeoff between incentives and efficiency in the NYC schools match
Mongell and Roth (1991)	Examined application of two-sided matching to sorority rush
Barron and Várdy (2004)	Recommended use of two-sided matching for International Monetary Fund Economist Program
Teo et al. (2001)	Recommended use of two-sided matching for allocation of high school students in Singapore
Abdulkadiroğlu and Sönmez (2003)	Recommended use of two-sided matching for certain high schools in U.S.A.
Abdulkadiroğlu et al. (2005)	Reviewed application of two-sided matching to New York City Department of Education high schools
Abdulkadiroğlu et al. (2006)	Outlined two potential assignment mechanisms for Boston Public Schools

Table 2 Two-Sided Matching Literature

If assignment decisions are adversely affected by human cognitive limitations, then the efficiency and effectiveness of assignment processes may be improved by implementing a decision support system. For reasons outlined in the introduction, this research focuses on two-sided matching processes as the mechanism by which assignments are generated.

Two-sided matching is a systematic process for generating assignments in situations where members belong to one of two distinct groups, and members of one group seek a partner from the other group. There are two broad classes of two-sided matching; one-to-one matching and many-to-one matching. One-to-one matching is generally visualized as a marriage market where men and women seek partners in marriage, with participants only able to be paired in a one-to-one monogamous relationship with a partner from the opposite sex. Many-to-one matching is often characterized as a matching between workers and firms, where each firm can employ many workers, but each worker can work for only one firm.

The theoretical study of two-sided matching began with the paper "College Admissions and the Stability of Marriage" by Gale and Shapley (1962). However, systematic two-sided matching existed in practice from 1951 in the National Intern Matching Program (NIMP), known today as the National Resident Matching Program (NRMP), a many-to-one matching market that assigns medical students to hospital residency programs. A number of concerns relating to the matching process existed in the medical matching market prior to implementation of the NIMP. For example, Roth and Sotomayor (1990) state, "Students who were pressured into accepting offers before their alternate status was resolved were unhappy if they were ultimately offered a preferable position, and hospitals whose candidates waited until the last minute to reject them were unhappy if their preferred alternate candidates had in the meantime already accepted positions" (p. 3). As such, during the 1940s, the U.S. medical residency market went through a period of ever decreasing time intervals during which students had to accept or decline offers. The interval was reduced from ten days in 1945 to 12 hours in 1949, although the American Hospital Association determined that even the 12-hour waiting period was too long. To avoid the uncertainty of whether waiting list offers would

become confirmed, and the associated last minute rejection of offers, the NIMP was introduced for the 1951–52 market. Since inception, the NIMP (and now the NRMP) has been a voluntary process, and the high participation rate is testimony to the ability of two-sided matching to solve at least some of the problems that were faced by medical students and residency programs. In 2003, 16,538 medical students matriculated from U.S. medical schools, and 15,101 (91.3%) of these participated in the NRMP together with 15,903 independent applicants (former graduates of U.S. medical schools or graduates of foreign medical schools).

Gale and Shapley (1962) describe a problem in the college market similar to that described in the medical market; the problem for colleges is that they do not know "whether a given applicant has applied elsewhere; …how he ranks the colleges to which he has applied; … which of the other colleges will offer to admit him" (p. 9). Applicants face similar problems in that they do not know how competitive they will be for certain colleges, and they do not know which colleges are likely to provide offers. These problems impact decision making when applicants are placed on college waiting lists pending vacancies, for the applicants must decide whether to accept a confirmed offer that may have been provided by a college, or wait and hope that a vacancy occurs at another preferred college where they are on the waiting list. If the applicant accepts a college's confirmed offer, but later rejects it when a preferred college's waiting list offer becomes a confirmed offer, this either leaves the first college with a vacancy, or causes a ripple effect as the first college sends out new offers to applicants that were on its waiting list. Therefore, one of the problems that two-sided matching resolves is the strategic manipulation that participants encounter when deciding what preferences to submit, and what offers to accept.

1. The Stability of Matches

One of the desirable qualities exhibited by two-sided matching assignments is stability. A matching outcome is stable if no agent is forced into an unacceptable assignment and no individual assignments are unstable. Using the example of the NRMP, an unstable assignment would be one in which a medical student and residency program

are not matched to each other, but prefer each other to their assigned partners. Such a pair is referred to as a blocking pair, and in a voluntary procedure, blocking pairs could withdraw from the process and organize their own match externally. However, two-sided matching guarantees a stable outcome; one that has no blocking pairs.

Gale and Shapley (1962) describe a "deferred-acceptance" algorithm for both a one-to-one and many-to-one market and prove that the algorithm produces stable matches. The algorithm for a many-to-one market is given as follows;

> First, all students apply to the college of their first choice. A college with a quota of q then places on its waiting list the q applicants who rank highest, or all applicants if there are fewer than q, and rejects the rest. Rejected applicants then apply to their second choice and again each college selects the top q from among the new applicants and those on its waiting list, puts these on its new waiting list and rejects the rest. The procedure terminates when every applicant is either on a waiting list or has been rejected by every college to which he is willing and permitted to apply. At this point each college admits everyone on its waiting list. (p. 13)

To see that the process must generate a stable outcome, consider a student S and college C who are not matched to each other, but S would prefer C to the college to which he / she has been matched. Given that S prefers C to his / her assigned college, S must have proposed to C during the process. However, C must have rejected S in favor of another student(s) that C prefers, because C will only keep on its waiting list its q most preferred applicants. Given that C prefers its assigned applicants to S, C and S cannot be a blocking pair and hence the matching is stable. The deferred-acceptance procedure for the one-to-one market operates in the same way as described for the many-to-one market if it is assumed that each college has a quota of one, and so the proof of stability is unchanged.

To examine whether the success of the NRMP is due to its stable outcomes or some other factors, Roth (1991) examined the experiences from several centralized matching markets for new physicians in the United Kingdom. Rather than a national market for medical students as exists in the U.S., the UK market is divided into a number of smaller regional markets. Of the seven different UK markets that have used computerized matching techniques, Roth found that the two that use stable matching

techniques have survived, while three of the five unstable techniques have failed. While not a statistically large sample, he finds that stability is in fact an important concept that relates to the success of the market.

Two further concepts appear to have been related to the markets that failed; firstly, there was no dominant strategy for any agents to submit their true preferences, and secondly, the students and hospitals in these failed markets increasingly negotiated private arrangements in advance of the centralized match. In the two nonstable markets that have survived, Roth hypothesizes that they continue to operate due to the smaller sizes of these markets that puts social pressure on students to accept appointments, regardless of whether they are stable or not.

When different matching mechanisms are possible, the stability provided by two-sided matching is a desirable property. Erdil and Ergin (2007) outline that legal and political issues favor stable mechanisms in certain applications, such as school choice.

2. The Nature of Matches Generated by the Deferred-Acceptance Algorithm

Some of the features of the assignment outcomes differ according to whether the relationship is a one-to-one matching or many-to-one matching, and hence these two situations will be examined separately. The findings discussed in this section, unless otherwise stated, are based upon participants having strict preferences. The issue and implications of preference list ties will be addressed later.

a. One-to-One Matching

Gale and Shapley (1962) showed that when preferences are strict, a stable matching exists for every marriage market. It is possible that there is more than one stable matching for a market. Roth and Sotomayor (1990) show that if multiple stable matches exist for a single market, the set of people who are single is the same for all stable matches.

Roth and Sotomayor (1990) describe an M-optimal stable matching as one that every man likes at least as well as any other stable matching. Similarly, the W-

optimal stable matching is the stable matching that all women agree is best. Furthermore, "When all agents have strict preferences, the M-optimal stable matching is the worst stable matching for the women; that is, it matches each woman with her least preferred achievable mate. Similarly, the W-optimal stable matching matches each man with his least preferred achievable mate" (p. 33). The matching produced by the deferred acceptance algorithm is the M-optimal stable matching when the men propose, and W-optimal when the women propose.

While the assignment outcome may be stable with respect to the stated preferences, it is necessary to consider whether the process encourages participants to state their true preferences. Roth and Sotomayor (1990) identify that there is no stable matching mechanism where stating true preferences is a dominant strategy for every agent. However, Dubins and Freedman (1981) show that, when the male proposing algorithm is used, it is a dominant strategy for the men to state their true preferences, and similarly the female proposing algorithm makes it a dominant strategy for women to state their true preferences. So, under the male proposing algorithm, for each man operating independently, there is nothing he can do better than to state his true preferences.

While the male proposing algorithm makes it a dominant strategy for men to tell the truth, it is not a dominant strategy for women to tell the truth under the male proposing algorithm. In fact, based on the male-proposing algorithm, Gale and Sotomayor (1985) show that "If there is more than one stable matching, then there is at least one woman who will be better off by falsifying, assuming the others tell the truth" (p. 264). If the women truncate their preferences by maintaining the true order, but deleting all men ranked below their W-optimal partners, it is possible for the women to obtain the W-optimal outcome by this misrepresentation even though the men propose (the proof of this follows from the fact that the same participants are single at every stable match). However, what makes this difficult for the women is that they must know their W-optimal partners.

In addition to truncation strategies, women can also benefit by submitting preferences that are not of the true order. By altering the true order of preferences, a woman is indicating which of two men to reject when she has an offer from both. The

one that is rejected can lead to a chain of further rejections, which in turn may lead to a better offer for the woman. Therefore it is necessary for the woman to know what further rejections and offers will result from her altering of her preferences. Roth and Rothblum (1999) show that such a strategy of misstatement requires detailed information about the preferences of the other men and women, and therefore conclude that submitting truncated preferences is a much easier form of strategic manipulation, although participants who truncate their preferences excessively run the risk of being unmatched at the conclusion of the matching process.

Roth and Sotomayor (1990) show that the male proposing algorithm produces a result that is weakly Pareto optimal for men, that is, there is no other matching, stable or not, that all men prefer. However, it is not possible to guarantee that the male proposing outcome is strongly Pareto optimal; an outcome that some men would prefer and others would be indifferent to. As shown by Dubins and Freedman (1981), it is possible to find unstable matches that leave some males no worse off, while improving the situation of the rest. Therefore, depending on the size of the market and the degree of knowledge of each participants' preferences, it is possible that a group of men may, in collusion, manipulate their preferences to obtain an outcome that will benefit some members of the group while leaving other members of the group indifferent to the outcome. While theory demonstrates the ability of participants to submit false preferences to obtain a strongly Pareto optimal outcome, experimentation is warranted to determine the ease with which participants could engage in such behavior.

A dominant strategy is defined as one that is a best response to whatever other participants may do. Under the male proposing algorithm, it is a dominant strategy for men to state their true preferences. While there is no such dominant strategy for women under the male proposing algorithm, there are equilibrium strategies as discussed by Roth and Sotomayor (1990), where each participant's strategy is in equilibrium if no player can achieve a better outcome by changing their strategy if all other participants leave their strategies unchanged. Such equilibrium is known as Nash equilibrium. Roth and Rothblum (1999) examine equilibrium strategies and consider the information that participants require in order to pursue equilibrium strategies. For example, a strong

equilibrium strategy is for the women to submit preferences truncated such that the last male appearing on each woman's preference list is the woman's W-optimal partner. However, as mentioned already, such a strategy requires the women to have detailed knowledge in order to know their W-optimal partners.

Roth and Sotomayor (1990) show that if a man in the market extends his preference list to include more women, this cannot help the other men and cannot harm the women. Furthermore, applying Gale and Sotomayor's Lemma 1 (1985), any men who prefer the matching derived from the men's original shorter preference list will be matched with women who prefer the matching derived from the men's longer preference list. That is, for two stable matches μ and μ' from a single market, if a man *m* prefers his partner at match μ to his partner at μ', then both of his female partners at μ and μ' will prefer the match μ' to the match μ.

If a new woman enters the market, it weakens the competitive position of the other women, while strengthening the competitive position of the men. As shown by Roth and Sotomayor, (1990) and Crawford (1991), at the M-optimal matching, no man can do worse by the introduction of a new woman into the market, and assuming that the new woman becomes matched to a man who was previously matched, there are some women who will have a less desirable partner due to the introduction of the new woman.

From Irving and Leather (1986), in a marriage market consisting of n men and n women who each submit complete preference lists, the number of stable matches grows exponentially as n increases. However, Roth and Peranson (1999) examine the number of stable matches in the NRMP, where preference lists are incomplete, and conclude, "When preferences are highly correlated (i.e., when similar programs tend to agree which are the most desirable applicants, and applicants tend to agree which are the most desirable programs), the set of stable matchings is small" (p. 767).

b. Many-to-One Matching

Many-to-one assignment situations are described here in terms of workers and firms, where each firm can accept multiple workers, but each worker can accept a job at only one firm. Many, but not all, of the findings relating to one-to-one matching

generalize to many-to-one matching. In some literature, e.g., Dubins and Freedman (1981) and Roth and Sotomayor (1990), the many-to-one situation is conceptually broken down to resemble a one-to-one situation, by considering each firm to be divided into multiple pieces of itself, with each piece having a quota of one, and the preferences of each piece of a firm being identical to the preferences of the firm. By considering a many-to-one situation in such a way, it is apparent that a stable outcome will always exist. Furthermore, it follows that for a given situation the same workers will be unmatched at every stable outcome and firms will have the same number of jobs vacant at every stable outcome.

Stability in a many-to-one situation can be determined by considering individual stability and pairwise stability. Individual stability is based on the notion that the process is voluntary and any worker or firm can remain unmatched rather than being forced to accept an undesirable partner, and hence a matching is individually stable if no participant is forced into an unacceptable match. A matching is pairwise stable if there are no blocking pairs consisting of a worker and firm, where each would prefer to be matched to the other than to their assigned partner(s). The implication of this is that stability can be identified by knowing the preferences of the workers and firms, and does not require knowledge of a firm's preferences over groups of workers.

Just as there exist M-optimal and W-optimal outcomes for one-to-one matching situations, there also exist worker and firm-optimal outcomes for many-to-one matching situations. Roth and Sotomayor (1990) describe these outcomes in terms of the residency matching process: "When preferences are strict, there exists a hospital-optimal stable matching that every hospital likes as well as any other stable matching, and a student-optimal matching that every student likes as well as any other stable matching" (p. 139).

While the results above mirror those from the one-to-one matching situation, the opportunities for strategic behavior are somewhat different. Roth and Sotomayor (1990) state, "When preferences over individuals are strict, the student-optimal stable matching is weakly Pareto optimal for the students, but the hospital-optimal stable matching need not even be weakly Pareto optimal for the hospitals" (p.

139). This is realized by first considering the one-to-one situation where the M-optimal solution is weakly Pareto optimal for the men, indicating that there is no other matching, stable or not, that all men prefer. However, in the many-to-one situation, while it is not possible for every job in a firm to prefer a matching to that provided by the firm optimal solution, an alternative solution may result in some of a firm's jobs being indifferent to the alternative solution while the other jobs at the firm prefer the alternative solution. Therefore, it is possible that some firms would prefer an alternative matching to the firm-optimal outcome. Or as stated by Roth (1985), "There may exist outcomes that all colleges strictly prefer to the C-optimal [college-optimal] stable outcome" (p. 283). Therefore, under either the firm or worker optimal algorithms, firms have an incentive to misrepresent their preferences to obtain a better outcome. For workers it remains a dominant strategy to state their true preferences when the worker-optimal algorithm is used, but not under the firm-optimal algorithm.

One of the problems with the early version of the NIMP algorithm was that it was a hospital-optimal algorithm; therefore, there was no dominant strategy for either hospitals or students to state their true preferences. Roth and Peranson (1997, 1999) examine the change of the NRMP algorithm to a student-optimal solution, an algorithm that does at least give students the incentive to state their true preferences. Interestingly, by using actual rank-order lists submitted to the NRMP in 1987 and from 1993–1996, and comparing the matches generated by both the hospital and student-optimal algorithms, it was found that less than 0.1% of applicants would have been affected by a change to the applicant-proposing algorithm (Roth & Peranson, 1997). This is primarily due to the correlation of preferences, which leads to minimal difference between the hospital and student optimal outcomes.

3. Determining an "Optimal" Stable Match

While the deferred-acceptance algorithm leads to a solution that either the men or the women regard as providing their best achievable partner, Irving, Leather and Gusfield (1987) consider a procedure to improve the average satisfaction of men and women, where satisfaction is measured by each partner's rank. The procedure operates as follows;

using preference lists that include only acceptable partners, the male proposing algorithm starts, and when a woman *w* receives a proposal from a man *m*, woman *w* deletes all men ranked lower than *m*. The men deleted from *w*'s list also delete *w* from their preference lists. This produces what Irving et al., refer to as the "male oriented shortlists." From the shortlists, rotations are identified, where "a rotation ρ is a sequence

$\rho = (m_0,w_0), \ldots, (m_{r-1},w_{r-1})$

of man / woman pairs such that, for each i ($0 \leq i \leq r$-1),

(i) w_i is first in m_i's shortlist;

(ii) w_{i+1} is second in m_i's shortlist (i + 1 taken modulo r)" (Irving et al., 1987, p. 535)

Within a rotation defined as above, if each man m_i exchanges his partner w_i for w_{i+1} then the resulting match is also stable. Irving et al., continue to show that by eliminating rotations it is possible to reveal every stable matching. Of the stable matches for a given instance, the stable match that provides the best average satisfaction for the men and women is selected.

One of the key facets focused on by Irving et al. (1987), is the computation time required to determine the "optimal" stable matching. Given that the maximum number of stable matchings for a situation of size n by n grows exponentially, the procedure produced by Irving et al. (1987) reduces the computation time from an exponential worst case to $O(n^4)$-time.

4. Market Complexities

Roth and Peranson (1999) describe four types of complexities that exist in the NRMP; couples seeking positions close to one another; applicants with supplemental lists for second-year positions; positions in residency programs that revert to different programs if unfilled; programs who seek to have an even number of positions filled. The deferred acceptance algorithm, as described previously, cannot handle such complexities, so Roth and Peranson (1997, 1999) use the instability-chaining algorithm, first described

in Roth and Vande Vate (1990), to handle these complexities in the redesign of the NRMP. For the NRMP, the instability-chaining algorithm considers the entire set of residency positions, introduces the medical students one at a time, and resolves any instability due to market complexities as they arise, before introducing new students into the algorithm.

5. Indifference / Preference List Ties

When preferences are not strict, different outcomes are achievable for different orderings of tied preferences. It is possible for some participants to be matched at one preference ordering yet unmatched at a different ordering of the tied preferences. Consider the example of Roth and Sotomayor (1990) that comprises "a marriage market with two men and only one woman. If the woman prefers either man to being single, but is indifferent between them, and if both men prefer the woman to being single, then there are exactly two stable matchings, and each of the men is single at one of them and married at the other" (p. 48). Since each man would prefer to be matched to the woman rather than be left unmatched (because only acceptable partners are listed), the men do not agree on which of the two outcomes is the better match. Therefore, when preferences are not strict, there need be no M-optimal (or W-optimal) stable matching since an M-optimal matching is defined as one that every man likes at least as well as any other stable match.

In the case that some man or woman is indifferent between two or more possible partners, Roth and Sotomayor (1990) propose a fixed tie-breaking rule, such as to alphabetically order preferences where ties occur. While this may ensure a standard and reproducible result, it may not necessarily ensure the "best" outcome. While under strict preferences the men agree that the M-optimal outcome is best and women agree that the W-optimal outcome is best, there is no M or W-optimal outcome under indifference, and therefore there is no definition of what constitutes the "best" outcome.

Irving (1994) suggests, "In practical instances, it is perhaps unrealistic to expect each participant to provide a strict preference ordering of all of his/her prospective partners. If indifference is permitted, so that each person's preference list becomes a

partial order, the character of the problems changes somewhat" (p. 262). Because participants may be matched to partners that they either strictly prefer or are indifferent towards, Irving introduces new definitions of stability for use when participants may express indifference. The notion of who constitutes a blocking pair is different for each definition:

- A weakly stable match has no blocking pairs, where a blocking pair is defined as any couple, each of whom strictly prefers the other to their assigned partner. When ties are broken arbitrarily and the Gale-Shapley Deferred Acceptance algorithm is applied, the resulting match is weakly stable.
- A strongly stable match has no blocking pairs, where a blocking pair is defined by any couple (*m*, *w*), where *m* strictly prefers *w* to his assigned partner, and *w* either strictly prefers *m* or is indifferent between her assigned partner and *m* (and vice versa for *m* and *w*).
- A super stable match has no blocking pairs, where a blocking pair is defined as any couple that either strictly prefers each other to their assigned partners, or are indifferent between each other and their assigned partners.

Irving (1994) shows that, while the deferred acceptance algorithm guarantees a stable match in a situation without indifference, the deferred acceptance algorithm guarantees a weakly stable match when indifference is expressed. Irving also provides algorithms to produce both strongly and super stable matches, however there is no guarantee that such solutions will exist.

Irving (1994) sought to prove when participants have ties in their preference lists, that super and strongly stable solutions can be found in polynomial time if they exist; a super stable solution can be found in $O(n^2)$ time and a strongly stable solution can be found in $O(n^4)$ time. Further work by Manlove (1999) and Iwama, Manlove,

Miyazaki and Morita (1999) examine the computational complexity of finding strong or super stable matches in marriage problems when ties are permitted in preference lists.

The work by Irving (1994), Manlove (1999) and Iwama et al. (1999) all focus primarily on the computational complexity of finding strong or super stable solutions, however none of them discuss the likelihood of finding such solutions. Irving, Manlove and Scott (2000) examine the likelihood of obtaining super stable solutions under various parameters, including the number of agents involved, lengths of preference lists and the number, position and sizes of the ties. They state, "the probability of a super-stable matching existing decreases as the size of the instance increases, and also decreases as the number and length of the ties increase. However, it was found that the probability of a super-stable matching existing is likely to be much higher if the ties occur on one side only; for example in the hospitals' lists and not in the residents' lists" (p. 270).

Irving, who has been an author of a number of previously mentioned papers that examine matching with preference list ties, was a design consultant for the Scottish PRHO Allocations Scheme (a medical student matching scheme), which first operated with a two-sided matching algorithm in 1999–2000. In this application, hospitals are permitted to express ties in their preferences lists, although medical students may not express ties. Interestingly, despite Irving's work on strong and stable matches, the algorithm used for the Scottish scheme did not involve a strong or super stable match. Rather, the approach taken was to randomly break the hospitals' ties and then use the student-optimal Gale-Shapley algorithm, modified for circumstances that require students to undertake both a surgical and a medical residency period. Unfortunately, it is not clear if the choice of approach was due to a low probability of obtaining a strong or super stable match or for other reasons. However, Irving does note that a subsequent analysis of the medical match revealed that an additional three students could have been matched (484 rather than 481) if a different tie breaking was used.

Gent and Prosser (2002) show that it is possible to use constraint programming to determine the smallest and largest number of matches in a stable matching with incomplete lists and indifference. However, unless the matching market is relatively small, the search cost (computing time) to determine the smallest and largest number of

matches increases significantly. Gent and Prosser did not examine any markets of dimension greater than 60, and primarily focused on markets of size 10. As market size increased from 10 to 60, and the number of preference list ties increased, it became too computationally time consuming to determine the largest and smallest number of stable matches.

Abdulkadiroğlu et al. (2007) examine the tradeoff between incentives and efficiency in the context of the New York City high schools match, and show that random tie-breaking preserves the strategy proof nature of the mechanism with regards to the students in the student-proposing mechanism, but that the tie breaking introduces artificial constraints that harm the student welfare. Therefore, attaining a strategy proof outcome has a cost in terms of efficiency, where efficiency is measured by the number of students who can be matched to a school that they rank higher. By relaxing the requirement for a strategy-proof student optimal approach, Abdulkadiroğlu, Pathak and Roth (2009) find that 6,854 students of the 63,795 students matched in the 2003-04 NYC high school match could have been matched to schools that they rank higher in their preferences. However, using data from the 2005–06 Boston Public School match, it was found that no Pareto improvements at all could be made for students. Abdulkadiroğlu, Pathak, Roth and Sonmez (2006) provide two reasons why the Boston school match is less likely to yield Pareto improvements; first, there were a significantly higher proportion of students receiving their first preferences in Boston as compared to NYC, and second, in Boston the magnitude of indifference is less than in NYC where some schools are indifferent between all applicants.

Erdil and Ergin (2007) examine the inefficiency that is created when preference list ties are broken with a lottery. They propose a procedure called the stable improvement cycle, which in situations involving preference list ties, can match some students to schools ranked higher on their preference lists without harm to others in the process. They also show that the stable improvement cycle is not a computationally hard process, unlike the processes used by Irving (1994), Manlove (1999) and Iwama et al. (1999).

6. Applications

Having examined the different types of two-sided matching processes and some of the issues relating to these processes, it is worth examining some of the applications of two-sided matching.

The National Resident Matching Program (NRMP) is perhaps the most widely known example of a two-sided matching market. The NRMP is in fact an organization that conducts many matches; the largest being the match for first and second year medical residency positions, however 35 medical specialty fellowship matches are also operated. Two-sided matching is also used in a number of other medical matching markets, primarily within the United States of America, Canada and the United Kingdom. Two-sided matching has also been used for the placement of personnel into articling positions in Canadian law firms and sororities (Mongell & Roth, 1991).

There have also been a number of cases where two-side matching has been, or is in the process of being examined for use in a new situation. Barron and Várdy (2005) recommend the use of two-sided matching as the mechanism for matching participants in the International Monetary Fund's Economist Program. Teo, Sethuraman and Tan (2001) propose using a stable matching process for allocating students to high schools in Singapore, a centralized matching process that did not produce stable outcomes. Abdulkadiroğlu and Sönmez (2003) examine two-sided matching as a mechanism to match students to high schools in areas such as Boston, Seattle, Columbus and Minneapolis, where school choice (intra and inter-district choice) has been implemented. Hansen (2003) describes a proposal to use two-sided matching for adopting children in the U.S.A.

Abdulkadiroğlu et al. (2005) provide a preliminary review of a relatively new, and the largest known, application of two-sided matching. Operating for the first time in 2004, over 90,000 students applied to New York City public high schools using a centralized process based on two-sided matching. Students submitted rank-order preference lists of up to 12 schools, and schools independently determined their preferences for students. The authors compare the processes before and after the process

redesign and find that the redesign produced favorable results; more students received schools on their preference lists, and less students were administratively assigned. The similarity between the high school process and the NRMP is clear; in each, students have preferences for schools / hospitals, and schools / hospitals have preferences for students. However, there is also a difference between these two processes; while the NRMP is an impartial matchmaker that has no direct interest in the assignment outcome, the NYC Department of Education does have an interest in how students are matched. One manner in which the Department of Education imposes its interests is through constraints on schools. In the NYC match, schools classified in the Educational Option program may individually evaluate students for half of the school's seats, but this is subject to 16% of seats being allocated to top performers in a standardized English Language Arts exam, 68% of seats to middle performers and 16% of seats to bottom performers. The other half of the school's seats are allocated by lottery, but again distributed in accordance with the standardized test scores.

Another school system that redesigned its assignment mechanism is the Boston Public Schools system. Abdulkadiroğlu et al. (2006) highlight that the strategic manipulation of preferences was a significant concern with the matching process that was in use prior to 2006. The authors provided two potential assignment mechanisms to the Boston Public Schools; the two-sided matching deferred acceptance algorithm and the top trading cycles mechanism. In a process where schools have priorities for students, the top trading cycles mechanism allows students to trade their priorities; for example, student A might be a priority for school Y but prefers school Z, and under this mechanism student A could swap the priority with another student who has a priority at school Z. The deferred acceptance algorithm was chosen because of concerns that the top trading cycles mechanism would be less transparent, more difficult to explain to parents and could perpetuate the perceived need to game the system. The Boston Public Schools implemented the deferred acceptance process for the 2006–07 school year.

While markets, such as the NYC Department of Education, have the potential for decisions to be aligned with the decision maker's preferences, this concept has not been

fully developed. This research takes a new approach; it develops a two-sided matching process where the decision maker has preferences over outcomes, and examines the impact on issues such as stability and participant behavior.

C. MULTI-ATTRIBUTE UTILITY THEORY

The literature to be covered in this section of the paper is summarized in Table 3.

Reference	Relevant Concepts
Keeney and Raiffa (1976)	Examined decision making for problems with multiple objectives
Triantaphyllou (2000)	Reviewed a variety of multi-criteria decision-making techniques
Saaty (1980)	Developed the Analytical Hierarchy Process (AHP)
Belton and Gear (1983)	Raised problem of rank-reversal in the original AHP and proposed a modified version
Edwards (1977)	Developed the Simple Multi-Attribute Rating Technique
Edwards and Barron (1994)	Proposed modifications to SMART to account for range sensitivity; SMART using Swings (SMARTS) and SMART Exploiting Ranks (SMARTER)
Benayoun et al. (1966)	Examined the ELECTRE decision-making technique
León (1997)	Provided support for additive functions
Zahedi (1986)	Reviewed applications of AHP
Belton (1986)	Examined limitations of the AHP
Taylor et al. (1998)	Applied the AHP to a university dean selection
Stewart (1992)	Discussed problem of range sensitivity
Nitzsch and Weber (1993)	Discussed problem of range sensitivity
Barron and Barrett (1996)	Developed case for the use of Rank Order Centroid weights
Srivastava et al. (1995), Stillwell et al. (1981), Fischer (1995),	Compared different techniques available for determining criteria weights

Borcherding et al. (1991), Weber and Borcherding (1993), Roberts and Goodwin (2002), Pöyhönen and Hämäläinen (2001)	

Table 3 Multi-Attribute Utility Theory Literature

As outlined, two-sided matching is a process that relies on rank-ordered preferences to generate assignments. Typically these preferences are determined cognitively by individuals involved in the process, either as participants (e.g., students in the NRMP) or as representatives of organizations (e.g., hospital staff determining residency preferences on behalf of a hospital in the NRMP). However, in situations that involve very large numbers of personnel who must be evaluated according to multiple attributes, a structured process enables rapid and consistent preference rank ordering. In such a structured process, participants may be ranked according to the degree to which they satisfy stated attributes. Multi-Attribute Utility Theory refers to a class of decision-making techniques that can support the evaluation and ranking of personnel as required for the two-sided matching algorithm.

There are numerous multi-attribute utility theory techniques to consider, including the additive model proposed by Keeney and Raiffa (1976), the Weighted Sum Model (WSM) and Weighted Product Model (WPM) outlined by Triantaphyllou (2000), the Analytical Hierarchy Process (AHP) proposed by Saaty (1980), the revised version of the AHP proposed by Belton and Gear (1983), the Simple Multi-Attribute Rating Technique (SMART) proposed by Edwards (1977), a modification of SMART titled SMART using Swings (SMARTS) proposed by Edwards and Barron (1994), the SMART Exploiting Ranks (SMARTER) proposed by Edwards and Barron (1994), and the ELECTRE method proposed by Benayoun, Roy and Sussman (1966). With such a variety of techniques, not surprisingly Triantaphyllou (2000) states the paradox; “What decision making method should be used to choose the best decision-making method?” (p. 5).

1. Decision-Making Techniques

A brief review of different decision-making techniques follows, but first it is important to clarify some terminology.

Alternatives—These are the options from which the decision maker is selecting. For example, when purchasing a car, the alternatives are the different models of cars that are being considered. In school matching situations, which have been described previously, the individual students are the alternatives that the schools must assess and rank order.

Attributes—An attribute is the means by which an alternative is assessed. Some literature uses the term criteria in place of attributes. In assessing the best car to purchase, attributes may include fuel efficiency and safety. In assessing students, attributes may include GPA and attendance rates.

a. Additive Value Function

From Keeney and Raiffa (1976), a value function may be expressed in an additive form if and only if the attributes are pairwise preferentially independent. They define that a "pair of attributes X and Y is preferentially independent of Z, if the conditional preferences in the (x,y) space given z´ do not depend on z´" (p. 101). They continue to show that, in a case that evaluates alternatives over multiple attributes, the additive value function will have the form

$$\upsilon(x_1, x_2, ..., x_n) = \sum_{i=1}^{n} \lambda_i \upsilon_i(x_i)$$

Where the weights υ are scaled from zero to one and the sum of the weights equals one. It is important to note that the component value functions are "normalized" so that the worst outcome on the i^{th} attribute has a score of 0 and the best outcome on the i^{th} attribute has a score of 1.

In support of the additive functional form, León (1997) claims, "additive structures produce excellent approximations to nonadditive ones and that most weight elicitation procedures for nonadditive MAU models are difficult and technical" (p. 250).

However, Edwards and Barron (1994) claim there may be times when additive models "may be lousy even as approximations. Fortunately, an easy-to-use test will weed out almost all instances in which an additive model would be really bad. It consists of looking for instances in which, at one level of value attribute x, more of y is better than less, while at another level of x, less of y is better than more" (p. 315).

b. Weighted Sum Model and Weighted Product Model

Similar to the above process is the Weighted Sum Model (WSM), described by Triantaphyllou (2000) as "probably the most commonly used approach, especially in single dimensional problems" (p. 6). The only difference between the WSM and the additive value function is that the WSM does not use a normalization process to convert values on each attribute to 0-1 utility values, and therefore it can only be used when the dimensions of each attribute are the same (e.g., dollars). Therefore, the WSM may be regarded as a limited application of the additive value function. When attributes are not all the same dimension, Triantaphyllou (2000) proposes the Weighted Product Model (WPM), which uses ratios to avoid the problem of different dimensions. Using the WPM, pairs of alternatives are compared with a ratio for each attribute compared. For example, in comparing car A to car B (two of potentially many alternatives), the fuel efficiency of car A is compared to car B (say 10mpg : 20mpg = 0.5) and the safety of car A is compared to car B (say 4 star : 5 star = 0.8). Attribute weights are then applied to each attribute ratio.

c. Analytical Hierarchy Process

The Analytical Hierarchy Process (AHP) was developed by Saaty (1980) and has been used in a large number of applications as reviewed by Zahedi (1986). Calhoun (1989) states, "Although the AHP process is mathematically complicated, it yields useful results and is not difficult to apply" (p. 21). The AHP is a process that relies on the decision maker providing subjective pairwise comparisons of the alternatives against each attribute, with the comparisons being given numeric values on a scale of 1 to 9 to indicate relative importance. There have been criticisms of the scales used to indicate relative importance, including Belton's (1986) comment, "The limitation of the scale to

1–9 imposes unnatural restrictions on judgments" (p. 11). A Revised AHP was developed by Belton and Gear (1983) to deal with potential rank inconsistencies that could occur with the AHP.

The number of comparisons required to undertake the process grows rapidly as the number of alternatives increases. For this reason, the AHP and Revised AHP are not appropriate for this research. Taylor, Ketcham and Hoffman (1998) show that the number of pairwise comparisons required per attribute is given by

$$T = \sum_{j=1}^{n-1} j$$

where T is the number of comparisons and n is the number of alternatives being evaluated. Therefore the number of comparisons grows rapidly even for modest sized problems, as demonstrated by Taylor et al. (1998) who applied the AHP to the selection of a college dean at Texas A&M University and found that the size of the problem, involving 33 candidates and four criteria, was excessive. With n = 33 (alternatives), the number of comparisons to be made for each attribute is 528. This number of comparisons must be made for each of the four attributes, resulting in a total of 2,112 comparisons.

d. SMART, SMARTS and SMARTER

Edwards (1977) developed the technique known as SMART (Simple Multi-Attribute Rating Technique) to provide an easy to use decision-making tool. There are three basic steps to determining attribute weights in SMART: first, attributes are ordered according to their importance; second, each attribute is assigned a relative weight according to its importance, usually with the least important attribute assigned a value of 10 and the most important a value of 100; and third, the importance weights are normalized by dividing each weight by the sum of the weights.

While this provides an easy process for weight determination, the concept of range sensitivity has been discussed by a number of authors including Stewart (1992) and Nitzsch and Weber (1993); they indicate that weights should not be based on importance alone, but should also be related to the range of outcomes for the criterion.

For example, in selecting a machine, the price attribute may be important to the decision maker, however if the price of available machines only varies from $20,000 to $20,100, there is little benefit in giving a high weight to price.

Recognizing that SMART failed to take range sensitivity into consideration, Edwards and Barron (1994) proposed two new techniques known as SMARTS (SMART using Swing weights) and SMARTER (SMART Exploiting Ranks). Each of these new techniques overcomes the problem of range sensitivity that existed in SMART. The technique of SMARTS is perhaps the most rigorous and accurate in terms of eliciting attribute weights, however the weight elicitation process is time consuming. The weight elicitation involves questioning the decision maker to determine what tradeoffs he / she would be willing to make in one attribute for variations in another of the attributes. Consequently, Edwards and Barron (1994) recommend SMARTER, which is based on Rank Order Centroid (ROC) weights (to be described below), the case for which is developed by Barron and Barrett (1996). Using ROC weights, the decision maker need only determine an ordinal ranking of attributes, from which the weight of the i^{th} most important attribute is determined as:

$$w_i = \frac{1}{n}\sum_{j=1}^{n}\frac{1}{j} \ , \quad \text{i} = 1,\ldots, \text{n}$$

This approach appeals to situations where decision makers may be able to relatively easily define an ordinal ranking of attributes but have difficulty with determining weights, a situation of the type discussed by Steuer and Schuler (1978). Consequently, Edwards and Barron (1994) recommend SMARTER due to the greater ease of elicitation. Furthermore, in terms of relative accuracy of SMARTS and SMARTER, they recommend SMARTER because it eliminate the difficult judgmental steps involved in SMARTS, and claim that ROC weights provide around 98 to 99% of the utility obtainable by full elicitation of the weights.

e. ELECTRE

ELECTRE is an outranking method developed by Roy (1968) that orders alternatives by determining which alternatives "outrank" others. Roy (1991) reviews different versions of ELECTRE; ELECTRE I, II, III, IV, TRI and IS.

As stated by Triantaphyllou (2000), "The basic concept of the ELECTRE method is to deal with "outranking relations" by using pairwise comparisons among alternatives under each one of the criteria separately" (p. 13). The notation used for ELECTRE, aSb, indicates that alternative a is "at least as good as" b. Concordance and discordance indices are then used to measure, for each pair of alternatives, whether aSb. One of the desirable features of ELECTRE is that it is noncompensatory, meaning that a bad score on one criterion cannot compensate for a good score on another criterion. However, for situations with a large number of alternatives, the pairwise comparisons required by ELECTRE are a limitation in the same way that such comparisons are a limitation of the AHP.

2. Determination of Weights

For techniques, such as the AHP, SMART, SMARTS and SMARTER, determining weights is embedded in the process. However, for additive value functions, WSM and WPM, a functional form is provided but requires the decision maker to determine weights through other means. A variety of techniques have been reviewed and compared by different authors; Srivastava, Connolly and Beach (1995), Stillwell, Seaver and Edwards (1981), Fischer (1995), Borcherding, Eppel and von Winterfeldt (1991), Weber and Borcherding (1993) and Roberts and Goodwin (2002). An overview of the main weighting techniques follows.

a. Direct Allocation

Direct allocation is probably the simplest of the weighting schemes. As discussed by Pöyhönen and Hämäläinen (2001), in this process weights are allocated to

the attributes directly; for example, by dividing 100 points between the attributes. This approach however suffers from the range sensitivity principle discussed previously as part of the SMART review.

b. Swing Weights

Edwards and Barron (1994) provide a description of Swing weights, as it is one of the intermediate steps of the SMARTS technique. In Swing weighting, all criteria are shown along with the best and worst outcome for each criterion. The decision maker must first assume they are given an alternative with the worst outcome for each criterion. They then must specify the criterion for which a "swing" from worst to best outcome would produce the greatest improvement in overall satisfaction. The decision maker is again presented with an alternative that has all the worst outcomes, and must identify which criterion, other than the one already identified, they would most like to change from worst to best outcome. This process is repeated until all criteria have been ranked. Once this phase is complete, a value of 100 is assigned to the first criterion (the one judged to have the largest swing importance), and all other criteria are expressed as a percentage of the first criterion. To do so, the decision maker is asked to indicate what weight, relative to the 100 assigned to the first criterion, they would provide to a swing from the worst to best outcome of the second most important criterion. This is repeated until all criteria are weighted on the scale 0–100. Once complete, the points are normalized to sum to one to provide criteria weights.

By asking decision makers to compare the best and worst outcomes for each criterion, this procedure considers the range sensitivity principle. Pöyhönen and Hämäläinen (2001) identify that the process of eliciting Swing weights explicitly incorporates the attribute ranges, while Fischer (1995) found that the Swing method performs better than direct weighting in reflecting the ranges of alternatives.

c. Rank Weights

A variety of rank-based weighting schemes are available and discussed by Barron and Barrett (1996). These methods include Rank Order Centroid (ROC), Rank Sum (RS) Weights and Rank Reciprocal (RR) Weights. For these weighting techniques,

the decision maker specifies the order in which the attributes are ranked, and the procedures derive weights based upon the ordinal ranking of the attributes. After determining the number and ordering of attributes in a problem, the weights are applied according to the following equations:

$$w_i(ROC) = \frac{1}{n}\sum_{j=i}^{n}\frac{1}{j}$$

$$w_i(RS) = \frac{n+1-i}{\sum_{j=1}^{n} j} = \frac{2(n+1-i)}{n(n+1)}$$

$$w_i(RR) = \frac{1/i}{\sum_{j=1}^{n} 1/j}$$

where i is the i^{th} of n attributes.

For example, a problem that is defined by five attributes would have weights for each method, as shown in Table 4.

Rank	ROC	RS	RR
1	0.4567	0.3333	0.4379
2	0.2567	0.2667	0.2190
3	0.1567	0.2000	0.1460
4	0.0900	0.1333	0.1095
5	0.0400	0.0667	0.0876

Table 4 Weights According to Rank Based Weighting Techniques

Using probability density functions, Barron and Barrett (1996) explain that ROC provides the best estimation of weights when all that is known is the rank ordering of attributes. Additionally, in comparing the rank order methods, Srivastava et al. (1995) state a preference for ROC weights because they “have an appealing theoretical rationale,

are simple to calculate, and appear to perform better than earlier rank-based schemes in reproducing preference generated by full-scale multiattribute models" (p. 113).

These procedures are quite simple in that the decision maker need only determine the ranking order of the attributes. As discussed by Barron and Barrett (1996), this can be valuable if the decision maker is "unavailable, unable, or unwilling to specify sufficiently precise weights; or … the decision-making group may be able to reach agreement only on a ranking of attribute weights" (p. 1516).

d. Tradeoff Weights

The tradeoff method proposed by Keeney and Raiffa (1976) requires the decision maker to consider two attributes at a time and determine how much of one attribute they would be willing to tradeoff for a given change in another attribute, all other attributes held constant. While the tradeoff method has certain technical appeal, Borcherding et al. (1991) identify that it is difficult and time consuming. Furthermore, Edwards and Barron (1994) believe that due to the complexity of this method, it is likely to produce elicitation errors. Furthermore, as stated by Pöyhönen and Hämäläinen (2001), this method "is problematic in many decision making problems because the method assumes that the attributes are measured on a continuous scale" (p. 572).

3. Multi-Attribute Utility Theory Summary

A variety of decision-making techniques are available, each having their own strengths and weaknesses, and selection of a technique depends on the decision-making situation. Most decision-making techniques are used to select a single outcome from a variety of alternatives, however the approach used in this research is to use the decision-making technique to provide an ordered set of alternatives, potentially including dominated alternatives, which will then be used as a preference list for a two-sided matching algorithm.

In order to derive the ordered set of alternatives in this research, an additive value function will be used. Techniques, such as AHP and ELECTRE, which require pairwise comparisons, were eliminated due to the large number of comparisons required. Of the

remaining techniques, the additive value function is selected for its simplicity and ease of explanation and understanding. Weights for the additive value function will be determined from the Rank Order Centroid technique due to the simplicity and ease of elicitation.

THIS PAGE INTENTIONALLY LEFT BLANK

III. RESEARCH DESIGN

The issue and importance of market design was discussed at the outset. This research examines the design of a personnel allocation process within a large hierarchical organization. While the inefficiencies that may exist in open personnel markets are likely to be readily apparent (e.g., market unraveling and declining participation rates), inefficiencies in large hierarchical organizations are not likely to be as apparent. Through experimentation with human subjects, the first research question examines the inefficiencies that may exist in an organization where assignments are conducted in the absence of decision support, situations typical of large hierarchical organizations. An examination of these inefficiencies will establish the requirement for redesign of the personnel assignment process in such organizations.

As was examined in the introduction, two-sided matching algorithms have features that are desirable for a variety of applications, and have been successfully applied in various markets. However, two-sided matching algorithms have not been applied to applications of the type explored in this research; large hierarchical organizations where considerable indifference may occur and system level attributes can be incorporated. The effect of these issues on two-sided matching outcomes is examined through computational experimentation in the second and third research questions. This chapter addresses the research design for each of the research questions in turn.

A. HUMAN DECISION-MAKING QUALITY IN ASSIGNMENT PROCESSES

The inefficiencies in a hierarchical organization's personnel assignment process may not be directly apparent. Short (2000) used survey responses from U.S. Navy detailers (those who conduct the assignments) and U.S. enlisted Sailors to examine the efficiency and effectiveness of the extant assignment process. Concerns were raised from both the Sailor and detailer perspective. Sailors expect to be treated fairly, but expressed distrust of detailers and the process. Meanwhile, detailers responded that the process was inefficient, complicated, burdensome, and were frustrated by Sailors' unrealistic

expectations. While such concerns were identified, survey work has limitations (Kerlinger & Lee, 2000). Results may be limited by sampling, how respondents feel their responses will be used and the degree of anonymity perceived. Furthermore, survey work is not able to indicate the quality of detailers' assignment decisions.

This research uses human subjects in experiments that incorporate induced value theory (Smith, 1976) to examine how human assignment decision making is affected by changes in the size and complexity of the assignment problem. Experimental methods are the accepted methodology used in other studies of information overload effects (see Jacoby et al., 1974; Henry, 1980; Casey, 1980; Keller & Staelin, 1987; Larichev et al., 1988), and the use of induced value theory enhances internal validity (Campbell & Stanley, 1963). A full factorial experimental design (Box, Hunter, & Hunter, 1978) is used.

1. Information Load

As outlined previously, consumer behavior research typically defines information load by the number of objects to be considered, and the number of attributes that identify each object. The subject then evaluates the attributes against his / her own preferences for each attribute, and selects a single object that maximizes his / her utility. This may be represented in Figure 1, where three objects X, Y and Z (e.g., different brands of laundry detergent as used by Jacoby, Speller & Kohn, 1974) are represented by three characteristics A, B and C (e.g., price, quality and phosphorous content). The decision maker, who is the experimental subject, has preferences for the three characteristics and evaluates the objects according to these characteristics.

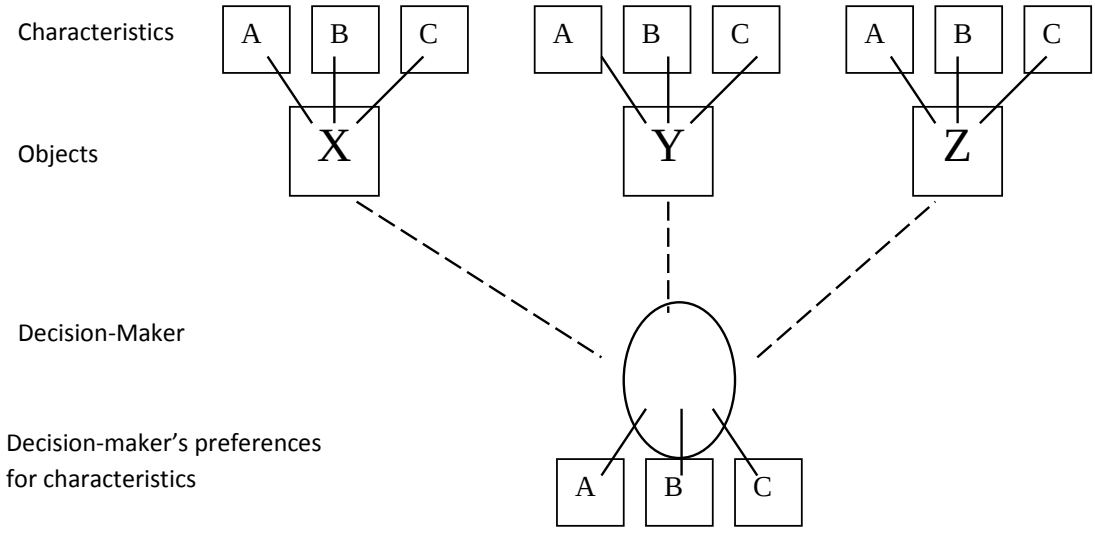

Figure 1. Representation of Consumer Decision

As shown in Figure 2, an assignment application has some similarities to a consumer decision; for example, personnel have various characteristics (e.g., skills attained, years of experience and specialization). However, in the assignment context, the decision maker is not making choices based upon his / her own preferences for the characteristics, but rather upon each position's preferences for the characteristics. The decision maker must consider each position's preferences for the characteristics, and determine the most suitable person available for each position. However, multiple assignments are required and tradeoffs may be necessary, because a person may be regarded as most suitable by more than one position. An additional consideration not shown is the preferences that the personnel have for positions. Therefore, while an assignment context has similarities to a consumer decision, assignment decisions are considerably more complex.

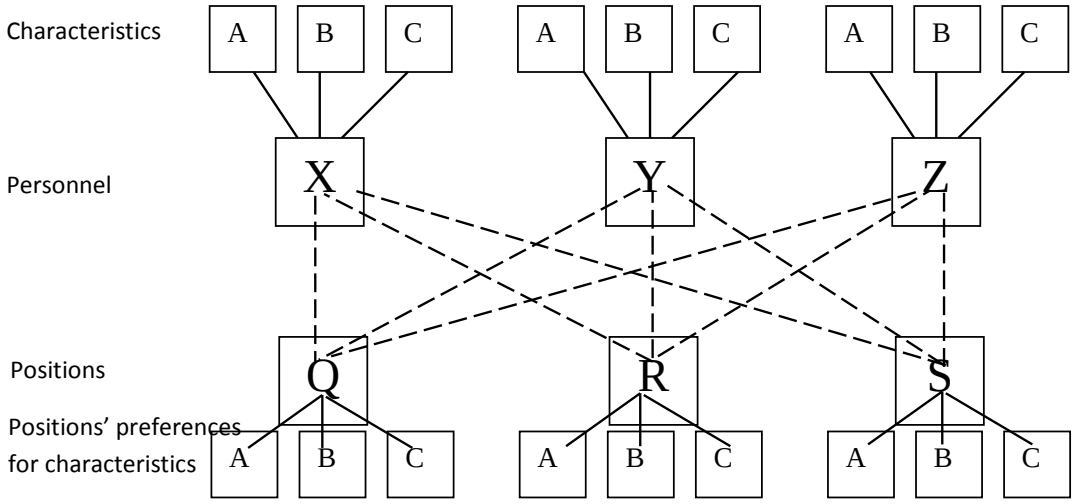

Figure 2. Representation of Assignment Decision

In the assignment experiments to be conducted in this research, subjects will be provided with information that is similar to what would be available in the assignment decision-making context for a large organization; a list of personnel and positions to be assigned, the characteristics of the personnel, the position's preferences for the various characteristics, and each person's preferences for positions.

As previously outlined, various processes are available for allocating weights to the attributes (characteristics) in a multi-attribute decision-making process. The results of such processes are that some attributes are likely to have a higher weighting than others, reflecting greater importance. To test whether subjects appropriately consider attributes of different importance, this research conceptually divides the attributes between primary and secondary attributes; primary attributes are the most important and are given greater weights, while secondary attributes are given less weight. Therefore, the secondary attributes are relevant and are to be taken into account where possible, but not to excessive detriment of the primary attributes. This leads to a situation of complexity similar to that used by Henry (1980) who defined complexity by the number of brands, number of attributes per brand and the scaling employed by each attribute (here the weighting of the attributes is analogous to Henry's scaling).

From related research on consumer behavior, it is expected that the subjects will suffer from information overload as the size of the market and the number of attributes

increase; however, it is unclear how this will manifest itself in the assignment decision makers' behavior. It is possible that human decision makers may ignore the less important secondary attributes (to ease the information load), or may attempt to assimilate the secondary attributes but in doing so increase the information load and consequently make poorer quality assignments. It is of course possible that different people employ different strategies.

The aim of this research is to confirm the existence of information overload in an assignment context, and examine how different subjects respond to information load. Assuming that information overload exists, the research will explore strategies that decision makers use to overcome it. The effects of assignment problem size and complexity will be examined.

2. Design

Of interest in this experiment is the performance of subjects as the complexity and size of the assignment problem increases. These two factors of interest, complexity and size, are the two factors to be controlled in the experiment, which will make use of a full-factorial design (Box et al., 1978). In such a design, each combination of factor levels is tested. A full-factorial design is appropriate where a limited number of factors of interest are used with limited levels for each factor. The advantage of a full-factorial design is its ability to uncover interaction effects (Campbell & Stanley, 1963).

Trochim (1999) states that the randomized block design is "research design's equivalent to stratified random sampling" (p. 207). Blocking is used where the sample is divided into subgroups that have similar characteristics. This ensures that the variability within blocks is less than the overall sample. A classic example of blocking is the experiment on durability of different materials in boy's shoes (Box et al., 1978). In some experimental work, such as the analysis of crops (Murphy, Samaranayake, Honig, Lawson & Murphy, 2005), blocking leads to differences during the implementation of an experiment. However, as indicated by Trochim (1999), blocking need not affect anything done with research participants; rather, blocking can be "a strategy for grouping people in your data analysis in order to reduce noise – it is an analysis strategy" (p. 207). In the

experiments to be conducted here, each person will complete all the factorial combinations of assignments. Given the differences in individual cognitive skills, blocking at the individual level will be used during analysis to "reduce the noise" of individual cognitive abilities. This will allow a clearer examination of the relative effects of increasing market size and complexity on decision quality.

The number of position attributes to be considered in an assignment round is a measure of complexity. Klingman and Phillips (1984) and Liang and Buclatin (1988) establish that, in the U.S. Marine Corps and U.S. Navy, the number of attributes to be considered varies between three and an upper limit of nine to 15 attributes. The lower limit of three attributes will be used in these experiments, however nine to 15 attributes is considered too many for subjects to assimilate in the time period available for experimental work. With such considerations in mind, the number of attributes for subjects to consider will be established at two control levels; three and five attributes. If, as is anticipated, the subjects have difficulty producing better assignments as more relevant data is provided (i.e., an increase in the attributes from three to five), it can be induced that subject performance would not be any better if the number of attributes increased further, say to the maximum number of attributes described by Klingman and Phillips or Liang and Buclatin.

Koh's (2002) study of the U.S. Navy enlisted detailing (the personnel assignment processes) reveals that, in each two-week requisition cycle, detailers (those who do the assigning) must consider approximately 45 Sailors. This volume of personnel and positions is too excessive for subjects to consider in the relatively short time frame available for an experiment, so a smaller number of assignments will be included in the experiments. The number of assignments will be set at levels of 5, 10 and 15; levels that will allow comparison of the effects of increasing numbers of assignments, while still allowing subjects to complete the experiment in a reasonable time period.

As indicated by Koh (2002), detailers typically have more positions available for consideration than there are personnel to be assigned; in the U.S. Navy enlisted detailing process, a ratio of 45 Sailors to 60 positions is typical. The experiments in this research use a similar ratio: for the case of five personnel to be assigned, an additional position

(i.e., six positions) will be provided for assignment. For consistency, the same ratio of six positions for every five personnel to be assigned will be used in the larger instances. Therefore, the three assignment levels analyzed in the experiments will be 5, 10 and 15 personnel, and 6, 12 and 18 positions. As indicated in Table 5, the assignment problems have three levels for size and two levels for attributes, providing a 3 x 2 combination, or six instances to be used in the experiment.

		attributes	
		3	5
Assignment Size, Number of Personnel to be Assigned	5	A	B
	10	C	D
	15	E	F

Table 5 Assignment Instances to be Examined.

In prior information overload studies, particularly in the consumer behavior domain, a particular brand is assumed to be the best outcome regardless of the number of alternatives and attributes presented. However, one of the difficulties with examining an assignment situation is comparing results between experiments. While position utility can be measured on a 0-1 scale, which facilitates comparison between subjects, this does not allow for comparing different instances since the average utility value differs between each instance as the problem specification changes. To maintain a common reference point for instances of each size, the characteristics of the three primary attributes will remain unchanged as the number of position attributes increases from three to five.

For example, in instance C of Table 5, the positions have three equally weighted primary attributes to be considered. Instance D comprises the same 10 personnel and 12 positions (same in terms of each person's characteristics and preferences for positions, and each position's primary attributes), but each position has two additional secondary attributes. The two secondary attributes will have significantly smaller weights such that, if the decision maker generates assignments with the position attributes foremost (i.e., with each person's preferences being subordinate), the outcome of instance D with respect to the primary attributes should be no worse than the outcome from instance C. If

the additional attributes lead the decision maker to return an outcome that is significantly worse in terms of the three primary attributes, then it may be concluded that the additional information has caused information overload for the decision maker. To prevent the subjects from recognizing that the primary attributes are identical for a given instance size, the ordering in which subjects are presented scenarios will ensure that a given instance size is not immediately repeated (e.g., not Scenario A followed by B).

The order in which subjects are presented scenarios will also be varied to examine for learning effects. That is, different groups will receive the scenarios in different orders. Three groups of subjects will allow for controlling for learning effects; the first group of subjects will be presented the scenarios in the order A, C, E, B, D, F, the second group will be provided the scenarios in the order E, C, A, F, D, B (i.e., starting with same number of attributes but larger scenario size first), and the third group will be provided the scenarios in the order B, D, F, A, C, E (i.e., starting with the same size scenario as the first group, but with more attributes first). These orderings are designed to account for learning effects, and also are designed such that the subjects should have forgotten the original specification and the results that they gave before they repeat a market size. Subjects will be randomly assigned to each of the groups.

Of interest in the experiments is the difference between quality of assignments when three attributes are presented versus five attributes, and whether such an effect is influenced by the number of assignments to be made. Henry (1980) demonstrated that individuals have different abilities with regard to decision making. Given the difference in abilities, blocking will be used to compare the results of individuals under each of the assignment instances.

Friedman and Sunder (1994) state the principle of induced value theory; “that proper use of a reward medium allows an experimenter to induce prespecified characteristics in experimental subjects, and the subjects’ innate characteristics become largely irrelevant” (p. 12). To induce behavior concordant with the experiment design, payments are made to the subjects according to their behavior. However, there is the potential that, if professionals are presented with familiar information, they will behave according to the business norms rather than according to the induced value. For example,

Friedman and Sunder (1994) provide the example of experiments with traders from the Minneapolis Grain Exchange who were provided extensive instructions on how to behave in the experiments, but then proceeded to behave by the Minneapolis Grain Exchange rules. Therefore, it is important to consider how information should be presented to subjects; familiar conditions and terms provide "cognitive handles" for subjects to grasp and therefore eliminate some of the information load that may be imposed by abstract or unfamiliar conditions and terms; however, this familiarity may cause problems with induced value theory, because the subjects may behave as they would expect in the real world rather than as directed in the experiment.

In this series of experiments, there is a conflict between internal and external validity, and it is necessary to develop a balance between these. Unfamiliar terms would provide the greatest internal validity, since we could conclude that subjects are not participating according to business norms, but rather according to the experiment design. However, completely unfamiliar terms would provide low external validity because decision makers in the real world do not typically make decisions based upon unfamiliar attributes. Therefore, unfamiliar terms would be likely to introduce an additional cognitive load that is not normally experienced by decision makers.

The subjects in the experiments are military officers who are undertaking master's-level education. In an attempt to balance internal and external validity, attribute descriptions have been selected to provide familiarity for the subject participants, while the values used to describe the attributes are more generic to avoid familiarity that may undermine the internal validity of the experiment. Table 6 outlines the attributes and values used to specify position requirements.

	Attributes	Possible Values
Primary Attributes	Rank	O-2, O-3, O-4, O-5
	Branch	Infantry, Artillery, Armor, Engineer
	Primary Skill	A, B, C, D
Secondary Attributes	Secondary Skill	Q, R, S, T
	Location	North, South, East, West

Table 6 Position Attributes

While the five attribute descriptions provided in Table 6 are not the only attributes that may be considered when determining assignments, the attributes were chosen because they are significant and meaningful to experimental subjects (external validity). Four possible values of each attribute are provided, as this allows the positions to be sufficiently distinct, yet allows overlap so that no individual assignment stands out as being clearly the best.

Subjects are provided with a description of each person and position in terms of the attributes outlined in Table 6, and the subject must match position requirements with personnel characteristics as best as possible, while also considering each person's preferences. Attribute utility values indicate which of the possible attribute characteristics is the most desirable for a position; a utility value of 1.0 indicates the most desirable attribute value and a value of 0.0 indicates the least desirable attribute value. Weights indicate the relative importance of each attribute, with the sum of the weights totaling to 1.0. The overall suitability of a person for a position may therefore be determined from the sum of the attribute utility values multiplied by the attribute weights. When only the three primary attributes are presented to the subjects, each attribute is weighted equally (i.e., 0.33 for each criterion). When all five attributes are presented to the subjects, the attribute weights become 0.26 for each of the three primary attributes and 0.11 for each of the secondary attributes. These lower weights to the secondary attributes ensure that a perfect match between a position and person on both the secondary attributes does not "outscore" a match on any of the primary attributes. Consequently, a subject should not

reduce the primary attribute utility to improve the secondary attribute utility, unless they consider it warranted due to a corresponding improvement in the personnel ranking.

Table 7 illustrates the information provided to subjects for market instance A.

	Weight	Criteria	Positions 1	2	3	4	5	6
Primary Criteria	0.333	**Rank**						
		O-2	0.5	0.5		0.5	0.5	
		O-3	1.0	1.0	0.5	1.0	1.0	0.5
		O-4			1.0			1.0
		O-5						
	0.333	**Branch**						
		Infantry				1.0		1.0
		Artillery	1.0		1.0			
		Armor		1.0				
		Engineers					1.0	
	0.333	**Primary Skill**						
		A		1.0		1.0		
		B			1.0			
		C					1.0	
		D	1.0					1.0

	Person				
Characteristics	**101**	**102**	**103**	**104**	**105**
Rank	O-3	O-3	O-4	O-3	O-2
Branch	Engineers	Armor	Infantry	Artillery	Artillery
Primary Skill	A	D	B	A	C

	Person				
Preferences	**101**	**102**	**103**	**104**	**105**
1st	2	2	3	2	5
2nd	4	6	6	4	1
3rd	5	1		1	3

Table 7 Information Provided for Market Instance A

As an example, Position 3 evaluates person 104 with a utility value of 0.50, derived from the Rank score (0.33 x 0.5) + Branch score (0.33 x 1.0) + Primary Skill score (0.33 x 0.0). It is also noted that person 104 does not rank Position 3 as a preference.

As will be noted at the top of Table 7, the position attributes are ordered Rank, Branch and Primary Skill. Subjects will randomly be provided the attributes in different orders. Therefore, if subjects weight the attributes in a manner other than specified for the experiment, this will help determine whether it reflects the ordering or the attribute names that lead to the different weights.

3. Experiment Setting

In a real-world assignment setting, the decision maker has a finite amount of time to make a given number of assignments, and the decision maker chooses whether to consider all available attributes in the time available, or only the most important attributes. To imitate this real-world situation, regardless of the number of attributes to be considered, subjects are provided with a fixed time period per assignment decision; 45 seconds per person to be assigned. The time of 45 seconds per assignment was selected from pilot experiments as a time that should allow subjects to complete the assignments, yet provide some time pressure. Therefore, for the market instances involving 5, 10 and 15 personnel, the subjects have 3:45, 7:30 and 11:15 minutes respectively to make the assignments. Subjects cannot move onto the next part of the experiment until the designated time has elapsed.

Subjects accrue points during the experiments, with the number of points accrued contingent on how well they maximize the average utility of each position and the personnel preferences. For each assignment instance, points are awarded in two parts. The first part is based on the extent to which each position is assigned the type of person sought, and is the average utility of the positions multiplied by 200 points. The second part of the payment is based on the extent to which each person achieves one of his or her preferred positions, and is the inverse of the average person's preference rank multiplied by 100 points. The greater number of points is accorded to the position requirements for

two reasons; firstly, most military assignment markets operate in a manner such that the service is given priority, and secondly because it requires greater cognitive effort to maximize the position utility as compared to the personnel preferences that are already provided as rank-ordered lists.

The experiments are to be conducted in a classroom setting, with details being provided to subjects in an envelope package and assignments to be completed on paper.

The experiments involve approximately 90 military, master's degree students participating as subjects in the experiments. The subjects are likely to have between five to twenty years of military service. Such personnel have had experience in the assignment process, most commonly indirectly as a person being assigned; however, the subjects may include personnel who have performed the role of detailer. Questionnaires will be used to determine existing differences in experience, with this information potentially being used to identify different blocking groups during the analysis phase.

Using military personnel as subjects provides external validity with regards to a military assignment process, because it is similar military personnel that make assignment decisions. However, it is not clear the extent to which the results may generalize to other hierarchical organizations where personnel of different characteristics generate the assignments. Furthermore, the military personnel participating in these experiments are relatively homogenous; they are required to achieve certain levels on standardized military entry tests, they have undertaken similar military training, are likely to have had similar types of experience, and are from a common culture. Such a homogeneous background could lead to less variability of responses than may actually occur in other hierarchical organizations where the assignment decision makers are less homogeneous. A follow up comparative study in a different organizational setting would be beneficial.

A copy of the experimental instructions are provided in Appendix A.

B. PREFERENCE LIST INDIFFERENCE IN TWO-SIDED MATCHING

The benefits of two-sided matching have been demonstrated in a number of external labor assignment processes (e.g., the NRMP), however, such processes have not been applied to the internal labor assignment practices of hierarchical organizations. Consequently, there are issues not yet explored that are relevant to implementing two-sided matching within hierarchical organizations.

In an external labor assignment situation that uses two-sided matching, it is important that the matching authority act without bias, and merely facilitate the assignment process based upon submitted preferences. Any bias on the part of the matching authority against participants from either side of the process may result in dissatisfaction and potential process boycotts by some participants. For example, Roth and Peranson (1999) state that in the 1990s, the NRMP "began to suffer a crisis of confidence concerning whether the matching algorithm was unreasonably favorable to employers at the expense of applicants" (p. 748). This concern centered over the use of the program versus the applicant-proposing algorithm, and reflected the stakes that are involved for the medical applicants; for the medical applicants, entry to a highly respected residency program could lead to greater opportunities in the early career stages. Such concerns motivated a NRMP redesign, and help explain why existing two-sided matching applications require participants to submit strict preferences.

Strict preferences and a single algorithmic approach such as the applicant-proposing algorithm ensure that a single assignment outcome is returned (Roth & Sotomayor, 1990), and so the matching authority simply implements the algorithm without making any decisions. However, if preference list ties are permitted, multiple alternative outcomes are possible, with each outcome being more highly regarded by some participants, and less regarded by others. Therefore, if participants were permitted to express ties in their preference lists, the matching authority would be required to choose one of the alternative outcomes. If participants had knowledge of the different possible outcomes, those who were aggrieved by the manner in which the ties were broken would have reason to complain to the matching authority and dispute the outcome.

While processes such as the NRMP require strict preferences to avoid disputes, the reasons for this requirement do not necessarily extend to hierarchical organizations. Subordination exists in hierarchies (Williamson, 1975), and so a matching authority within a hierarchical organization may permit preference list indifference without concern that disputes regarding selection of one alternative over another will undermine the assignment process. It is perhaps for this reason that two-sided matching with indifference can work in a school district scenario; the schools and students are subordinate to the school board / department of education, and so it is reasonable for the board / department to choose which outcome offers the greatest overall welfare.

The effect of preference list ties on efficiency has been explored by Abdulkadiroğlu et al. (2009) and Erdil and Ergin (2007), but only with respect to preference rank efficiency. That is, beginning with a stable match determined from random tie breaking, they use the stable improvements cycle to find the student optimal stable match. However, this is a student optimal stable match with respect to the initial tie breaking, and it should be considered that other tie breaking processes could produce different outcomes with a different number of participants matched. This research examines the different results that can occur when different tie breaking mechanisms are used. In particular, whether different tie breaking approaches could yield greater efficiency in terms of the number of agents matched, and whether different approaches to tie breaking may improve efficiency in terms of both number of participants matched and the preferences to which they are matched.

1. The Likelihood of Preference List Indifference in Large Hierarchical Organizations

As new applications for two-sided matching continue to appear, and as theory and practice continue to inform each other, the effect of preference list indifference is attracting greater interest. Recent research dealing with the effects of preference list indifference from a primarily theoretical viewpoint include Erdil (2006), Erdil and Ergin (2006), and Erdil and Ergin (2007). From a practical perspective, recent research on the effects of preference list indifference include Abdulkadiroğlu et al. (2005),

Abdulkadiroğlu et al. (2006) and Abdulkadiroğlu et al. (2009). This literature has focused primarily on many-to-one matching markets, with the increase in interest related to the growing use of two-sided matching to allocate spaces in schools.

Roth and Sotomayor (1990) view preference list indifference as indicating a lack of information available to a participant, and that if more information, possibly including new dimensions, was available, then the participants could differentiate between alternatives. However, there may be situations where such additional information is not forthcoming or it is unreasonable to seek new dimensions to differentiate.

In a large two-sided matching process, such as the New York City Department of Education where schools are required to rank-order students, the schools only have limited information dimensions to differentiate students. Furthermore, unlike the NRMP, it is not feasible for each school to interview each student to determine a preference ranking. In the Boston Public School assignment process, Erdil and Ergin (2007) outline that the school preferences are determined centrally based on four indifference categories for each school. In order of priority, schools rank students according to the following categories: 1) students who have siblings at the school and are in the school's walk zone, 2) students who have siblings at the school, 3) students within the school's walk zone, and 4) all other students. It is clear that in such a case there will be considerable indifference in the schools' preferences. While the school rankings in the New York City high school match are determined differently, considerable indifference also exists; in fact, Abdulkadiroğlu (2009) outlines that some schools have preference lists that are completely indifferent between all student applicants.

Similarly, military organizations are sufficiently large that widespread interviews are not feasible, and some positions tend to be relatively generic; for example, numerous people may be equally qualified (in terms of measurable and recordable attributes) for a staff position within a headquarters, and such positions would therefore have high levels of indifference in their preference lists. On the other side of the application, there are potentially a number of similarly defined positions to which a person would be equally willing to be assigned, and therefore would like to express this as indifference in their preference lists.

Due to the likely existence of preference list indifference in certain organizations, and because of its relevance to hierarchical organizations, the next part of this research examines the effect of preference list indifference on two-sided matching outcomes. While two-sided matching literature (Roth & Sotomayor, 1990) shows that different results are obtainable with preference list indifference, the extent and frequency of such variation has not been empirically examined. Therefore, this research examines how the outcome properties vary under different indifference conditions. The outcome properties of interest are the number of participants assigned and the "quality" of the assignments.

Computational experimentation is used to examine the effects of preference list indifference because it allows examination of the effects of different market characteristics and their interaction effects, and it enhances internal validity.

2. Simulation Model

An agent-based model will be used to demonstrate the effects of preference list indifference in a two-sided matching process. The model is developed in a Microsoft Access database where the tables store data on the agents' characteristics and preferences, and a Visual Basic program is developed to determine assignments based upon agent preferences. Further detail on the model is provided later. The assignment algorithm is modeled on the instability-chaining algorithm presented by Roth and Peranson (1999), which allows for complexities such as married couples. While such complexities are not involved in this part of the research, they are relevant for the final section of research that models assignments within a hierarchical organization, and so for ease of simulation development, a single program based upon the instability-chaining algorithm is used. In the case of an application without complexities, this algorithm produces the same results as the Gale and Shapley (1962) deferred-acceptance algorithm.

Two classes of agents are used in the model, personnel and positions, with each agent behaving independently of other agents in its own class. The conceptual model for the behavior of agents in this research is homophily, described by Hinds, Caley, Krackhardt and Wholey (2000) as "the tendency for people to be attracted to others who have similar attitudes, beliefs and personal characteristics" (p. 228). The concept is used

in this model to determine whether an agent of one class will be attracted to an agent of the other class. In friendship determination, homophily arises from multidimensional factors; for example, age, gender, education, social class and beliefs (McPherson and Smith-Lovin, 1987). Similarly, in a situation such as the NRMP, the extent to which a medical student and residency program will be "attracted" to each other may be a reflection of multidimensional factors; for example, a high reputation program seeks high performing students and vice versa, and a rural based program seeks students willing to work in rural areas and vice versa. Homophily based on geographic location is used by Erdil and Ergin (2007) to generate school and student preferences, on the basis that, "If a student has a high priority at a school because she is in the walk zone of that school, then she would in turn be more likely to favor that school because of its locational proximity" (p. 12).

While the attraction of agents may be multidimensional, certain dimensions may dominate. Roth and Peranson (1999) find from historical NRMP data that preferences tend to be highly correlated, indicating, "similar programs tend to agree which are the most desirable applicants, and applicants tend to agree which are the most desirable programs." (p. 767). While their analysis does not indicate specifically what dimension the programs and applicants agree upon (quality, measured by hospital reputation and student grades, may be a reasonable assumption), it does indicate that there exists a dominant dimension upon which preferences are determined. The model used for this research uses homophilous attraction based upon a single dimension. While additional dimensions could be used, the intent is to determine the limits of expected outcomes, not specific outcomes for specific organizations, and therefore a single dimension is considered sufficient.

Each agent in the model, both personnel and positions, is assigned a "quality" score (the name of the dimension is unimportant) on the interval zero to one; those with a score closer to zero are considered lower quality, while those with a score closer to one are considered higher quality. Based upon homophilous attraction, agents from each class seek agents of the other class that are alike in terms of the quality score. In this model, position agents have perfect knowledge of the quality scores of personnel agents.

Consequently, each position is able to determine a list of personnel, ordered from those who are most alike the position (and so most able to satisfy the position's requirements) to those least alike the position (and so least able to satisfy the position's requirements). This resembles the situation in a hierarchical assignment situation where the organization determines attributes against which the personnel are assessed, rates the personnel against those attributes, and is able to determine an ordered list of the personnel who are most suited to each position. The personnel agents in this model have knowledge of the positions that identify them as suitable, and the personnel are able to determine which positions they would prefer to be assigned to.

Figure 3 demonstrates some sample agents and indicates the manner by which an agent's preferences are determined. For example, person Z ranks Positions B, C and D in his / her preference ranking because those positions are sufficiently close to person Z in terms of the quality score. Further details on the determination of agents' preferences are provided in the next section.

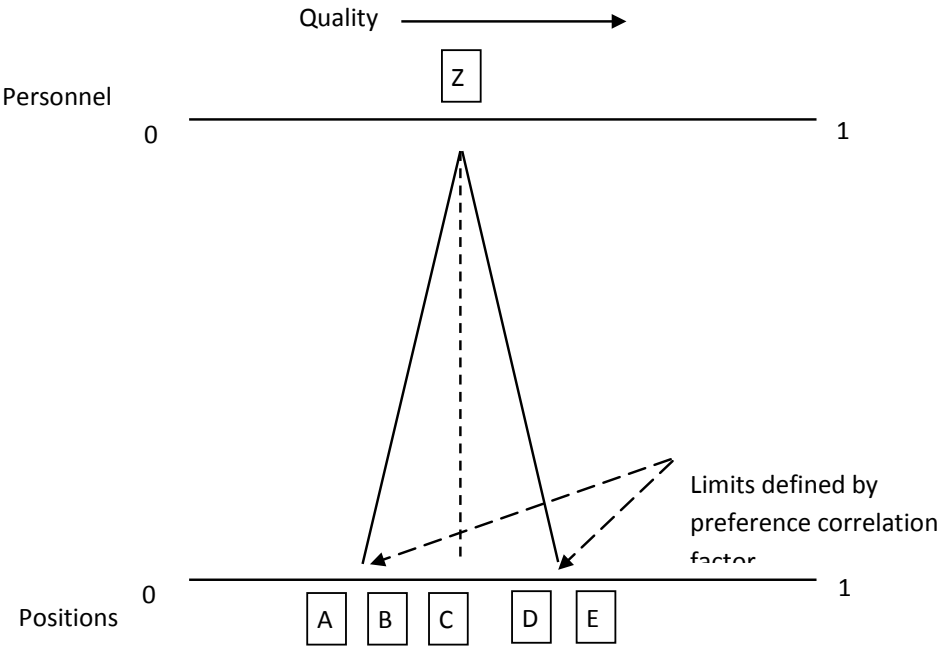

Figure 3. Example of Agents and Homophily

The focus of the simulation is determining the upper bounds of how preference list indifference affects two-sided matching assignments, without addressing the detailed

characteristics of any single application. This approach favors internal over external validity; it establishes the limits on outcomes likely to be achieved, without specifying the exact outcomes likely to be achieved in a specific application.

3. Model Parameters

To examine the effects of preference list indifference on matching applications of varying specifications, a computational experimentation approach is used. Nissen and Buettner (2004) discuss the advantages of computational experimentation; namely that it bridges the gap between laboratory and field research by mitigating the weaknesses of each. However, they warn that computational experimentation has its own limitations, and is best used as a complement rather than a replacement for laboratory and field research.

For the purposes of this experimentation, the assignment environment is designed according to four key parameters; the number of agents, the preference list length that agents submit, the extent to which agents' preferences are correlated and the extent to which indifference exists in each agent's preferences. The reason for selecting these factors and an explanation of the factors follows.

Number of Agents (n). Roth and Peranson (1999) find that the number of participants in the process is one of the important determinants of the number of stable matches. They state that while the number of participants increases, there remains a limit to the number of positions to which a person can interview and hence submit as preferences. Consequently, as assignment problem size grows, the number of possible stable matches remains small. Given the importance of assignment problem size in establishing limits on the number of stable matches when preferences are strict, it is an important variable to consider when preferences may include indifference. The number of agents to be generated within the process is specified by the parameter n, with a symmetrical process (i.e., n workers and n positions) being used. Such a specification is valid given that the majority of two-sided matching processes are relatively symmetrical.

Each of the *n* person and position agents is allocated a random score that signals the quality of the agent to agents of the other type. The quality scores are measured on the interval $0 \leq \text{quality} < 1$, with the values having no significance other than to identify which agents are most alike and those which are least alike.

Preference Correlation (*C*). As discussed previously, preferences tend to be correlated; agents seek assignment with agents of similar quality. This parameter allows investigation of whether the strength of correlation has an effect on assignment outcomes. Based upon the quality score that each agent is allocated, the preference correlation parameter indicates how similar two agents (a person and position) must be in order to consider ranking each other. If agents only consider other agents who are very similar in terms of quality scores, the preferences are regarded as highly correlated.

The correlation parameter is a single value for the entire market (i.e., not unique for each agent). It is a value measured on the interval $0 \leq C \leq 1$ to reflect the difference in quality scores that each agent is willing to consider. A correlation factor close to zero indicates that agents would only consider other agents that are very similar in terms of their quality score, while a correlation factor close to one indicates that agents are less discriminating when considering other agents. A correlation factor of one would indicate completely random preferences. Figure 4 demonstrates two different correlation parameters that a market may have. When the personnel and positions are ordered according to their quality scores, a Low *C* value indicates that a person will only consider positions that are similar in terms of quality scores. Similarly, a High *C* value indicates that a person is more willing to consider positions less like him / herself.

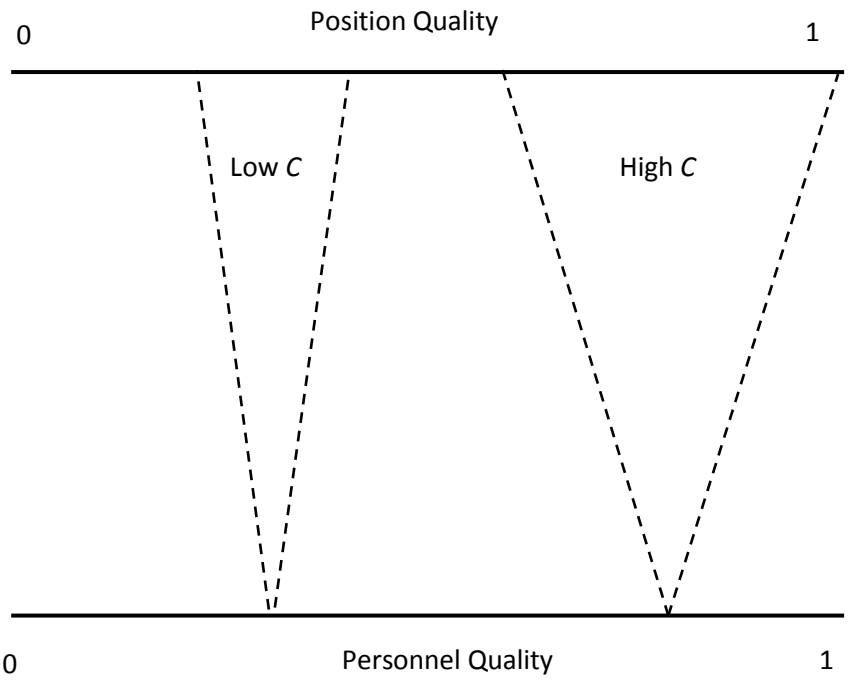

Figure 4. Demonstration of Preference Correlation Parameter

Positions strictly order personnel who are within the region of consideration (defined by *C*) based on the difference in quality scores; a position prefers personnel who are closer in quality scores. Personnel randomly order position agents that are within each person's region of consideration (defined by *C*). An additional consideration for the personnel ranking is that a person will only rank a position if the position has ranked the person. For positions, the use of strict ordering based on quality scores is not unrealistic if preference rankings are determined in a hierarchical situation by an algorithmic approach. However, even if personnel are generally attracted towards those positions which are alike, there is likely to be a degree of randomization that reflects a variety of factors that are different for each person.

The use of homophily that is bounded by a correlation demonstrates how outcomes vary as preferences vary from being highly correlated at one extreme (low *C* value), to totally random at the other extreme (*C*=1). Within the computational experimentation, varying the correlation parameter indicates the effect of stronger or weaker homophily.

The effects of different values of the correlation parameter are demonstrated in Figure 5 and Figure 6. In these figures, each point represents an assignment, with the x and y-coordinates given by the position's and person's quality scores respectively. Figure 5 is a plot of potential agent assignments in a completely random market (C = 1), while Figure 6 plots potential assignments in a market with high correlation (say C = 0.1).

Figure 5. Low Correlation

Figure 6. High Correlation

Maximum Preferences (*L*). It is known that a relationship exists between the length of preference lists submitted and the number of participants who are matched; reviewing the 1996–2006 results, the NRMP reports that "matched applicants and filled programs consistently have longer ROLs (Rank Order Lists) than unmatched applicants and unfilled programs." In existing two-sided matching markets, upper limits are set on the number of preferences that participants may submit, however not all participants submit preference lists up to the maximum length allowed. Of interest to this research question is the impact that preference list length has on assignment outcomes when preference list indifference is allowed. How preferences are determined for each agent has been described above, and is based upon agents' quality scores, preference correlation and a random factor.

In the simulation, positions determine their preferences first, and rank list a maximum of *L* personnel in their preference list, provided that *L* or more personnel exist within the quality region bounded by the position's preference correlation factor. If fewer than *L* personnel exist in the position's quality region, all personnel within the quality region will be rank ordered. After establishing all position preferences, personnel

preferences are determined, such that personnel only rank positions if they have already been ranked by the position, and the positions are within the person's preference correlation zone.

Degree of Indifference (*DOI*). This parameter is used to examine different market characteristics. If a spectrum is considered such that preferences are strictly ordered at one end and preferences are completely indifferent at the other end, this parameter indicates where the simulated market falls along that spectrum. When generating an agent's preferences, the *DOI* parameter determines whether a preference (with the exception of first ranked preferences) is considered indifferently to the previous preference. The *DOI* is defined on the interval $0 \leq DOI < 1$. An initial strict preference ordering is randomly generated; for each worker and position preference (except the first), a random number p is generated (defined $0 \leq p < 1$). If $p < DOI$ then the preference is ranked indifferently to the previous listed preference, otherwise it is considered strictly less desirable than the previous preference. If $DOI = 0$ then all preferences are strictly ordered (no indifference); if $DOI = 1$ then all preferences are completely indifferent. A similar process for generating preferences was used by Gent and Prosser (2002).

4. Model Development

The model is built within a Microsoft Access database, where VBA generates the agents and their preferences according to the parameters already described. Information about the agents is stored in the database tables and VBA programming subsequently generates assignments based upon the stored agent preferences.

The form shown in Figure 7 is created in the Access database to input the parameters which specify a scenario.

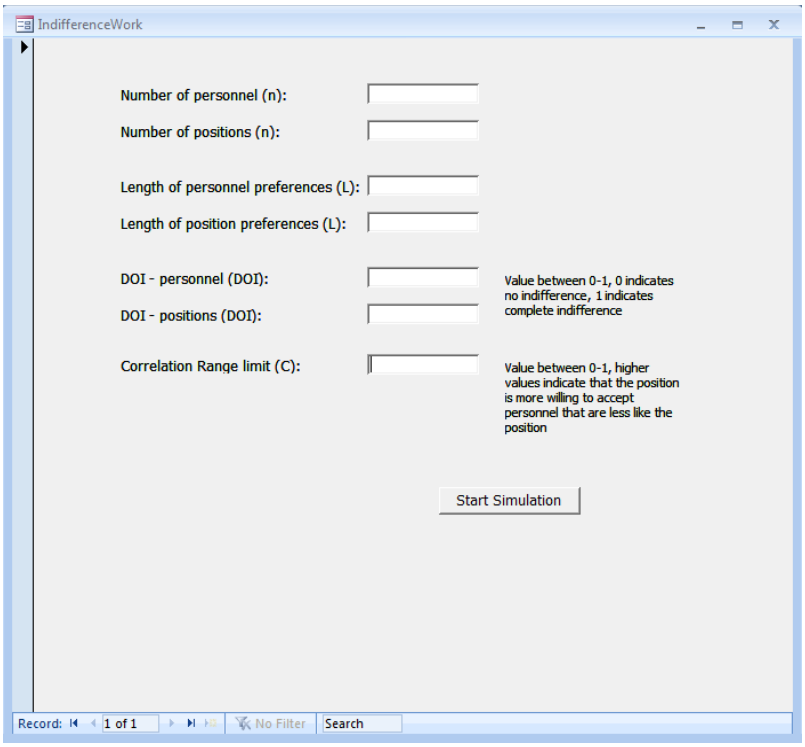

Figure 7. Screenshot of Form for Creating a Scenario to Investigate Indifference

Once the simulation is started from Figure 7, the Visual Basic code creates and populates a table with the quality values for each person and position up to the number *n* indicated. These quality values are generated through a random number generator. Figure 8 shows an example of the table containing position quality values.

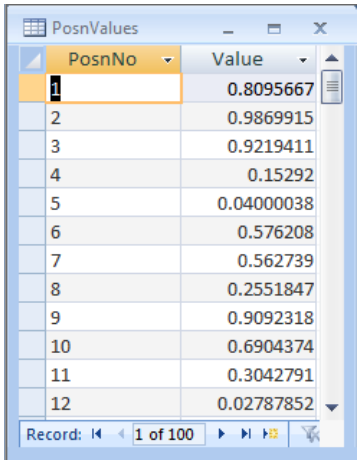

Figure 8. Example of Randomly Generated Position Quality Values

Once the position and personnel quality values have been stored in the appropriate tables, the Visual Basic code then constructs the position preference lists. All personnel whose quality values are within the range ± (1/2 x *C*) from each position's quality value are placed into an interim list which is ordered based on differences between each position's and person's quality scores. If this interim list is larger than the designated maximum position preference list length (*L*), then only the first *L* personnel from the interim list are selected and placed into the final position preference list, otherwise all personnel from the interim list are placed into the final position preference list. The Visual Basic code then proceeds down each position's strictly ordered preference list, and for each preference (except a position's first preference) a random number *p* is generated (defined $0 \leq p < 1$). If $p < DOI$ then the preference is ranked indifferently to the previous listed preference, otherwise it is considered strictly less desirable than the previous preference. This produces a preference list for the positions, a section of which is shown in Figure 9, where Rank is the strictly ordered preference list and Rank2 is the preference list with indifference.

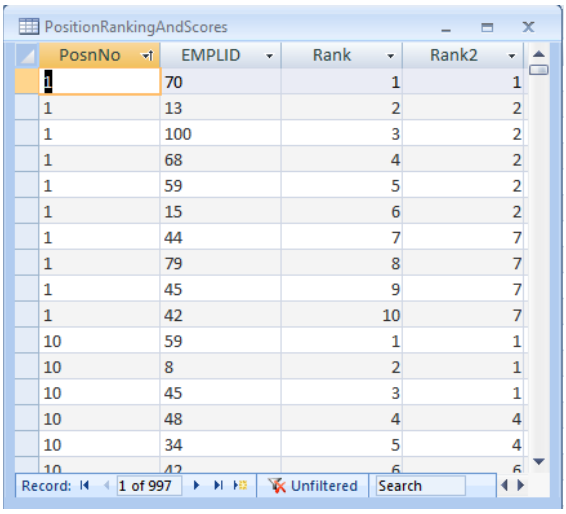
PositionRankingAndScores

PosnNo	EMPLID	Rank	Rank2
1	70	1	1
1	13	2	2
1	100	3	2
1	68	4	2
1	59	5	2
1	15	6	2
1	44	7	7
1	79	8	7
1	45	9	7
1	42	10	7
10	59	1	1
10	8	2	1
10	45	3	1
10	48	4	4
10	34	5	4
10	42	6	6

Record: 1 of 997 Unfiltered Search

Figure 9. Example of Position Preference List

Once the position preferences have been determined as outlined above, the personnel preferences are determined in a similar manner but with two key differences. The first difference is that a position is only placed on a person's interim list if the

position has already ranked the person. The second difference is that the preference list is not strictly ordered according to the difference in quality scores.

Visual basic code was written to implement the two-sided matching algorithm based on preferences determined as outlined above. Assignments are calculated from the two-sided matching algorithm for a specified number of trials, with the indifferent preferences randomly reordered for each trial. The assigned personnel and positions are recorded for each trial to analyze the results at a later stage. An example of the table that records these results is shown in Figure 10. Figure 10 illustrates the results from ten trials for two personnel (EMPLID 1 and 10), and shows the positions (identified by POSNNO) that were assigned to these personnel in each trial, along with how the personnel and positions ranked their partners. In this example, EMPLID 1 was assigned to the same position at every one of the ten trials. However EMPLID 10 was assigned to three different positions across the ten trials, depending on how tied preferences are ordered

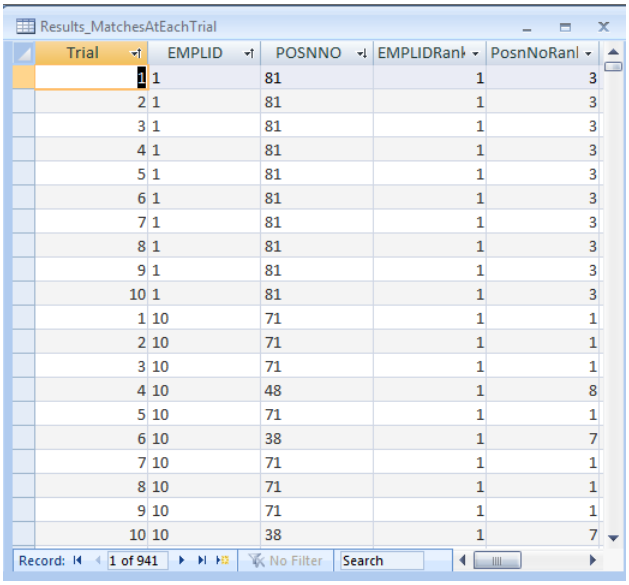
Results_MatchesAtEachTrial

Trial	EMPLID	POSNNO	EMPLIDRanl	PosnNoRanl
1	1	81	1	3
2	1	81	1	3
3	1	81	1	3
4	1	81	1	3
5	1	81	1	3
6	1	81	1	3
7	1	81	1	3
8	1	81	1	3
9	1	81	1	3
10	1	81	1	3
1	10	71	1	1
2	10	71	1	1
3	10	71	1	1
4	10	48	1	8
5	10	71	1	1
6	10	38	1	7
7	10	71	1	1
8	10	71	1	1
9	10	71	1	1
10	10	38	1	7

Record: 1 of 941 No Filter Search

Figure 10. Example Results from Indifference Simulation

5. Exploring Outcomes

Computational experimentation is used to examine the effects of the four parameters on the number of agents matched and the quality of assignments. To examine

the individual and interaction effects of the four parameters described in the last section, a factorial design is used. The levels to be used for each of the variables in the experimental design are discussed below.

It is not intended to examine how preference list indifference affects a specific assignment situation, such as a medical resident market, high school or military assignments. The intent is to examine how preference list indifference affects any generalized assignment situation. The four parameters incorporate small to large assignment situations (by varying *n*), situations where preference lists are short or long (by varying *L*), situations where agents may have very little or extensive indifference in their preferences (by varying *DOI*), and situations where agents may be more or less willing to be assigned to partners similar to themselves (by varying *C*).

Number of Agents (*n*). There is a broad range in the size of two-sided matching processes; ranging from medical specialty processes such as thoracic surgery with approximately 200 applicants, to the New York City high schools with approximately 94,000 students. The time demands for computational experimentation increase significantly as the number of agents increase beyond small to moderate sized situations. Therefore, scenarios will be used that vary the number of agents between 100 to 2,000 agents, which represents realistic numbers of agents for small to moderately sized processes. While this parameter is at the lower range of values expected in actual applications, by varying the parameter between different levels it will be possible to examine the effect of changing market size.

Maximum Preferences (*L*). In existing markets, and regardless of market size, applicants are typically restricted to between 5 and 20 ranked preferences, although many choose to submit shorter preferences. The preference list length that agents may submit in this experimental work is set between 5 and 20, as this represents a reasonable number of preferences based upon actual two-sided matching applications.

Preference Correlation (*C*). To determine the correlation of preferences that occurs in actual two-sided matching markets would be difficult and largely unproductive for this research, given that the intent is to determine whether, and if so how, the

correlation affects market outcomes. For two scenarios that have differing numbers of agents, the number of agents located within the bounds of the correlation factor is larger for the scenario that has more agents. For example, given that agents have quality values randomly distributed between zero and one, a scenario with 100 agents would be likely to have 10 agents within a correlation bound of 0.1, whereas a scenario with 1,000 agents would be likely to have 100 agents within the same correlation bound. Therefore, preference correlation values will be set at levels according to the number of agents in the scenario. The correlation values will be set such that the number of agents within the correlation bound will be similar to the preference list lengths identified above (5 to 20). This will reveal how outcomes may vary as an assignment market moves from highly correlated preferences to increasingly random preferences.

Degree of Indifference (*DOI*). There is no market data available to indicate the degree of indifference that participants may desire to express in existing applications. Therefore, values of 0.1, 0.3 and 0.5 will be used to measure the effects of preferences that progress from being relatively strict towards increasing indifference.

Within the computational setting, a variety of simulated markets will be generated according to the described parameters. For each scenario defined, multiple outcomes are possible because, depending upon how the tied preferences are broken, each different ordering of ties can lead to a different outcome. While the ideal situation would be to algorithmically determine the different outcomes that are possible for a given scenario, no known algorithm provides this information, and Irving (2000) claims that it is unlikely that any such algorithm can be found. Furthermore, a brute force method of examining every possible alternative is not feasible since the number of possible combinations of preference orderings can be very large for even moderately sized markets. For example, consider a single participant whose preferences for positions are as given in Table 8.

Position	Initial Strict Preference Rank	Indifferent Preference Rank
p1	1	1
p2	2	1
p3	3	3
p4	4	3
p5	5	3
p6	6	6
p7	7	6

Table 8 Example Participant Preference List

In Table 8, the total number of possible reorderings of that one participant's preferences is 24 (2! x 3! x 2!), each of which may yield a different outcome depending on how other participants reorder their indifferent preferences. Even in a small market with indifference, the number of combinations can be very large due to the factorial nature of the problem. Therefore, in a larger market with extensive preference lists and indifference, it is infeasible to run the assignment algorithm for every possible way that ties can be broken.

Consequently, a statistical approach is used. For a given scenario, multiple trials will be conducted with tied preferences to be broken randomly in each trial. Once tied preferences are randomly broken in a trial, the preferences are treated as being strictly ordered, and a two-sided matching algorithm applied. This determines a weakly-stable match for the random preference ordering. The outcome properties (number and quality of assignments) of the trial are recorded before the process is repeated with a new random breaking of the originally tied preferences. While such an approach cannot guarantee that all outcomes are explored, statistical techniques will enable limits on the characteristics to be determined with a degree of probability.

Abdulkadiroğlu et al. (2009) discuss the impact of using multiple random tie-breaking where lottery numbers are assigned to students at each school, versus single random tie breaking where each student is given a single lottery number to break ties at

every school. They indicate that the single tie-breaking process has superior welfare properties. The difference between the two processes will be examined in this experimentation.

The algorithm to be used in this work is the worker-optimal instability chain algorithm. This algorithm was described by Roth and Peranson (1999), and has been used in the NRMP since 1998. This algorithm has been used because it generates a stable outcome when couples participate, and although couples are not considered in this section, the outcome of markets that include couples is of interest in the next research question. Therefore, for consistency, the instability chain algorithm is used throughout this research.

C. DEVELOPING TWO-SIDED MATCHING IN A HIERARCHICAL ORGANIZATION

This research has been divided into three parts. The first part of the research demonstrates the need for redesign of assignment processes in hierarchical organizations. The second part of the research lays some of the foundations that are necessary, by considering the effects of indifference and examining how to deal with indifference in two-sided matching algorithms. This third part of the research demonstrates how such assignment processes in hierarchical organizations may be redesigned.

The New York City high school match has features that resemble the hierarchical organizations that will be outlined in this part of the research. In that high school match, the students are participants who are subordinate in the sense that they must abide by the organization's rules and decisions, and they independently submit their preferences for the schools. The hierarchical organization consists of the schools and the Department of Education, both of which have interests in the final assignment of students to the schools. As outlined by Abdulkadiroğlu et al. (2005), some schools have preferences for students based on scores, while other schools have preferences for students with good attendance records. The Department of Education also has preferences over the assignment of students, as evidenced by the constraints placed on schools; the Department requires

schools to have a distribution of students according to scores in standardized English Language Arts exams. Therefore, the preferences submitted by schools are not solely the preferences of the schools, but also include factors of interest to the Department of Education.

In other hierarchical organizations, such as the military, the assignment process is similar to that outlined above. In the military context, the individual personnel have independent preferences for the positions to which they may be assigned. The hierarchical organization consists of the units (or commands) that are analogous to the schools, and the service headquarters which is analogous to the Department of Education, and each of these have preferences over the assignment of the personnel to the units. The units, who receive the personnel being assigned, have preferences for personnel of certain skills, rank, ability etc., while the service headquarters will have preferences over factors such as minimization of relocation costs. Consequently, it is reasonable that the preference lists of units be constructed in a manner that includes factors of relevance to the service, in the same way that New York City school preferences include constraints of relevance to the Department of Education.

This part of the research will demonstrate that it is possible to construct preferences for the hierarchical organization with a multi-attribute utility function. A utility function will be developed to rank personnel according to the attributes of interest for the units and service. Through changes to the weights of the attributes, computational experimentation will examine the different results produced by the deferred acceptance algorithm as the weighting shifts from unit to service preferences, and various combinations in between. For reasons previously outlined, dealing with indifference will undoubtedly be important.

Multi-attribute utility functions lend themselves to determining the preference rank ordering of personnel for hierarchical organizations, such as military units and schools. Features of such organizations that make the multi-attribute utility functions an appealing option include: the large size of the organizations, which makes it infeasible for military units or schools to know, interview or evaluate each person individually; the requirement to combine system level attributes with the attributes relevant to the units or

schools; and the central decision maker's requirement to apply attributes consistently to all participants without bias. While using multi-attribute utility functions to determine preference lists has appeal in situations that have the above features, such an approach has not previously been used in two-sided matching applications. Therefore, it is desirable to learn the effects on assignment outcomes if preferences are determined from multi-attribute utility functions.

Computational experimentation is used to examine how preferences derived from utility functions will affect two-sided matching assignment outcomes. The advantage of this approach is the ability to change the variables and the weightings applied, and examine the different outcomes returned by the assignment process.

The process developed for this part of the research is similar to that used for the decision-making experimental work described previously; military personnel are to be assigned to individual positions. The personnel in these experiments have preferences over the positions to which they wish to be assigned, and for the positions, utility functions are used to rank order the available personnel based on attributes of interest to the units and the service. Each position's rank ordered list becomes the preference list for the purpose of the two-sided matching.

Unique preferences lists are to be determined for each individual position within the organization, and so this is considered a one-to-one two-sided matching process. That is, the personnel and positions are the two groups of agents within the process, and each person can be matched to one and only one position, and each position can be matched to one and only one person. In this respect, this is different to the previously outlined schools matching which are constructed as many-to-one processes. In the schools matching process, the seats within the schools are identical and it does not matter which student is assigned to which seat at a school.

1. Agent Specification

The personnel in this process are described by six characteristics, of which three are of interest to the units and three are of interest to the Service. The three characteristics of interest to the units are Rank, Branch and Skills. Units are interested in such

characteristics because they define some of the key attributes necessary for satisfactory work performance in the unit. Although the terminology may be different, such characteristics are not dissimilar to the way positions are advertised in other contexts; for example, an academic position may be advertised according to the desired level (assistant professor, associate professor, etc., being analogous to rank), the desired academic stream (mathematics, physics, etc., being analogous to branch), and the desired specialization (statistics, graph theory, etc., mathematics specializations being analogous to military skills).

The three characteristics of interest to the Service in this process are Career Development, Location and Performance. The Service is interested in career development because the effective development of personnel through different types of assignments is what prepares a person for higher appointments within an organization. Militaries are not the only organizations using varying assignments for personnel career development; Kohonen (2005) details the career development aspects of international assignments in multi-national organizations. Organizations such as militaries are also interested in location considerations when making assignments, because each person who is required to geographically relocate for a new assignment is an additional organizational cost, and given budget constraints, *ceteris paribus*, the organization would prefer to minimize relocation costs. While each unit would ideally like to receive the highest performing personnel, from a Service perspective it is preferable that the highest performing personnel are assigned to positions that require personnel of such talent; for example, positions that are representational in nature or involve training functions may require higher performing personnel as compared to some other types of positions.

These experiments utilize actual data for Australian Army personnel and positions to the maximum extent possible, thereby increasing the external validity of the research. Data is not available for all attributes of personnel and positions; consequently, in some cases it has been necessary to randomly assign characteristics to the personnel and positions based on distributions (this will be explicitly identified where it occurs).

The purpose of these experiments is to demonstrate how assignment outcomes vary when a hierarchical organization's preferences are constructed using multi-attribute

utility functions. Using the entire population of Australian Army personnel and positions would be time consuming and unnecessary for the purpose of these experiments. Therefore, a segment of the population has been selected for use; the experiments are based on personnel and positions of rank E06 (SGT), E08 (WO2) and E09 (WO1) (the E07 (SSGT) rank is no longer used in the Australian Army). The total number of personnel and positions in these ranks is 5,370 and 5,991. The majority of assignments for the Australian Army occur in January of each year, and most personnel complete three year assignments. Therefore, one third of all personnel and positions at these ranks are assumed to be available for assignment. Data was randomly removed so that one third of personnel and positions remain for these experiments. A total of 1,790 personnel and 1,997 positions remain for assigning in these experiments.

The distribution of the personnel and positions by rank, branch and skill is shown in Appendix B; these are actual distributions of personnel and positions, however the names of branches and skills have been de-identified for security reasons. As can be observed in Appendix B, multiple skills can belong to a branch, but no skill can belong to more than one branch. Every person is identifiable by a specific rank, branch and skill. While every position is identifiable by a specific rank, the branch and skill specification of positions is different: some positions require a person of a particular branch and skill; some positions are not skill specific but require a person from a specific branch; some positions are not skill specific but require a person from a select group of branches (identified by position branches 161, 162 and 163 in Appendix B); some positions are not specific to any skill or branch (identified by position branch 151 in Appendix B).

Actual data on the location of all personnel and positions was obtained, and to simplify the experiments, the many locations were grouped into localities (regions). The primary purpose of this simplification was to facilitate determining each person agent's preferences; as will be described in further detail later, personnel preferences were not available, and so these are constructed using other information and some assumptions. A summary of the number of personnel and positions by rank and locality is provided in Appendix B.

Career development is one of the Service attributes used to construct positions' utility values. The rationale is that each position provides a different experience, and each person requires particular types of experiences to be developed through their career. While career development is an explicitly stated consideration for assignments in the Australian Army, the career development which each position provides is tacit knowledge. Therefore, as a simplification, it is assumed that a position provides and a person requires one of four types of career development: experience in a regimental environment, experience in a headquarters, experience with the Reserves, and experience in a training role. All units are classified according to one of these four types of career development, and it is assumed that all positions within each unit provide that type of career development. Personnel agents are randomly assigned career development requirements based on the position career development distribution. A summary of the number of personnel and positions by rank and career development is provided at Appendix B.

Performance reporting data for personnel is not available due to privacy issues. Therefore, using a uniform distribution, each person agent is randomly assigned a value from one to four, to indicate a performance quartile to which each person belongs. Similarly, each position is randomly allocated to a performance quartile to indicate whether the position requires a high performing person, or whether a lower performing person is acceptable.

A table is used to store the six attributes (rank, branch, skill, location, career development and performance) for all 1,790 personnel agents, and a separate table is used to store the same six attributes for the 1,990 position agents. Extracts from these tables are shown in Table 9 and Table 10.

PERS ID	Rank	Branch	Skill	Locality	Career Development Required	Performance
200001	E08	103	218	QLD NORTH	REGT	3
200070	E06	102	211	QLD NORTH	REGT	2
200072	E06	115	227	VIC MELB	REGT	0
200100	E08	103	228	VIC MELB	HQ	1
200104	E09	106	222	QLD SOUTH	REGT	0
200106	E08	105	243	NSW SYD	REGT	3

Table 9 Sample Personnel Data

POSN ID	Rank	Branch	Skill	Locality	Career Development Provided	Performance Requirement
100001	E09	106	208	NSW OTHER	HQ	0
100007	E06	104	202	ACT	HQ	3
100008	E09	151	Any	OVERSEAS	HQ	3
100012	E06	106	208	OVERSEAS	HQ	1
100020	E08	103	228	NT DARWIN	REGT	1
100045	E08	102	211	QLD SOUTH	REGT	3

Table 10 Sample Position Data

2. Generating Preferences for Positions

This section outlines how the positions' preference lists will be determined from a utility function based on the six attributes. This is a necessary step for the purpose of these experiments. It also demonstrates how hierarchical organizations may use multi-attribute utility functions to construct preference lists based on attributes of interest to the organization as a whole, as well as elements of the organization. In these experiments, the Service (the Australian Army) represents the organization, while units represent the organization elements; this is similar to the Department of Education and the schools in a school assignment context.

Based on the review of available multi-attribute utility theory techniques, an additive utility function is used to generate the position preferences. An additive functional form is appropriate because the attributes to be used in the utility function are preferentially independent. Furthermore, as outlined in the literature review, the additive functional form is simpler to use and easier to explain than many other more complicated forms, yet provides quite good results.

The positions' preferences are generated according to the combined utility of the unit and Service attributes. When implemented, a decision maker will need to set such weights to provide the appropriate balance between unit and Service requirements. In these experiments the weights will be varied to determine how different weights affect the assignment outcomes.

The utility function for determining the position preferences is:

$$U_t = (W_r \times U_r) + (W_b \times U_b) + (W_s \times U_s) + (W_l \times U_l) + (W_{cd} \times U_{cd}) + (W_p \times U_p)$$

Where, U is the utility, W is the weight, and the subscripts t, r, b, s, l, cd and p indicate total, rank, branch, skill, location, career development and performance respectively. For each position, the person with the highest U_t value is the position's first ranked preference, and subsequent preferences are in order of U_t. Details for calculating each of the attribute utilities follows.

a. Rank Utility

The rank utility table shown in Table 11 was constructed to identify the rank utility that a position assigns to a person. This table is not based on any existing processes, but was developed to demonstrate how such a process could operate. The table identifies that the rank utility value will be one (ideal situation) if a person and position are of the same rank. Further, personnel whose rank is one below that of a position (rank utility of 0.6) are preferred to personnel whose rank is one above that of a position (utility of 0.3). The rank utility value is zero for personnel whose rank is two above or below that of a position. No veto rules have been included because the scenarios only involve three consecutive ranks and it is considered that a person from any of these ranks is at least acceptable for positions of these ranks. However, if the process involved assignment of positions and personnel whose ranks were more varied, a veto rule could be included to prevent personnel being assigned to positions if the rank difference between the two is considered excessive.

		Pers Rank		
		E06	E08	E09
Posn Rank	E06	1	0.3	0
	E08	0.6	1	0.3
	E09	0	0.6	1

Table 11 Rank Utility Table

b. Branch Utility

The branch utility table shown in Table 12 was constructed to identify the branch utility that a position assigns to a person. Branches 101 to 118 are those that personnel can belong to (e.g., Infantry, Artillery and Engineers). Branches 151 and above are "position only" branches. Positions that belong to branch 151 can accept personnel of any branch. Positions that belong to branches 161, 162 or 163 can accept personnel from a limited number of branches; for example, some positions belong to an "Any Health" branch, which can accept any personnel who belong to the branches Medical, Dental or Nursing.

		Pers Branch																	
		101	102	103	104	105	106	107	108	109	110	111	112	113	114	115	116	117	118
Position Branch	101	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	102	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	103	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	104	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	105	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	106	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	107	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	108	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	109	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	110	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1
	111	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1
	112	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1
	113	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1
	114	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1
	115	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1
	116	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1
	117	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1
	118	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1
	151	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
	161	1	-1	-1	-1	-1	-1	1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	162	-1	1	1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
	163	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	1	-1	-1

Table 12 Branch Utility Table

A veto rule has been included such that if a position records a negative branch utility for a person, then the person will not be included as one of the position's preferences.

c. Skill Utility

A skill utility value of one is recorded where the skill of a person and position match; otherwise a skill utility value of zero is recorded. Although possible, no veto rule has been included because a veto already exists at the branch level.

d. Locality Utility

The locality utility is used to minimize the number of geographic moves required. A locality utility value of one is recorded where the locality of a person and position match; otherwise a locality value of zero is recorded. Other options such as calculating the cost of moving between different locations would be possible but have not been included for these experiments.

e. Career Development Utility

The career development utility is used to maximize the number of assignments that provide personnel with their ideal career development. A career development utility value of one is recorded where the career development that a position provides matches what a person requires; otherwise a career development value of zero is recorded.

f. Performance Utility

The performance utility is used to provide the best possible alignment between the performance of personnel and the performance requirements of positions. Table 13 shows the performance utility values developed for these experiments. Decision makers may choose to use performance quality gradations other than the quartiles shown here and similarly may choose to use utility values different to those shown in Table 13.

		Personnel Quality			
		0	1	2	3
Position Quality	0	1	0.5	0.25	0
	1	0.5	1	0.5	0.25
	2	0.25	0.5	1	0.5
	3	0	0.25	0.5	1

Table 13 Performance Utility Table

3. Simulating the Preferences for Personnel

In an actual assignment situation, personnel would submit their list of preferred positions to a central department within the hierarchical organization, just as medical students submit their preferences to the NRMP and students submit their preferences to the agency responsible for the schools match. In this experimentation, it is necessary to simulate personal preferences.

The first consideration with respect to the personnel preferences is that personnel will only rank positions for which they have been included on the positions' preference lists. In existing markets, such as the NRMP, participating students are unaware of how or if they are ranked by the residency programs. However, in a hierarchical organization it is feasible that personnel within the organization could be informed of whether they are suitable and ranked by the units within the organization. In fact, in very large organizations such as militaries, it may be preferable that the personnel are informed of how highly they are ranked by the positions. Such information need not necessarily be overly specific (to the extent of specific rankings), but could provide indicative feedback as to whether a person is highly, moderately or poorly ranked by a position. Such feedback would assist with managing the expectations of personnel within the organization. Furthermore, in situations where personnel do not meet essential position requirements, the feedback would prevent personnel wasting their preferences on positions to which they cannot be assigned.

This simulation attempts to generate each person's preferences in a manner that may reflect how people behave. While Butler and Molina (2002) indicate factors that motivate personnel within the U.S. Navy Aviation Support Equipment Technician community, broader data is limited regarding the factors that influence personnel selecting their preferred positions. In the model, there are six known characteristics of the personnel and positions, and each person's inclination towards each of the six position characteristics will be used to generate the personnel preferences. In a manner similar to the utility function that generates position preferences, personnel preferences will also be generated from a utility function.

Different factors motivate different people, and the strength or intensity of the motivation may differ between people. Some people may be strongly motivated by the desire to work in a particular location, and for them other considerations are minor. Other people may be motivated by a strong desire to succeed in their employment, and so will seek work that enhances their career (career development assignments). Others may seek work in a domain where they feel comfortable and familiar, for example within their branch and particular skills area. Therefore, the model should simulate that different personnel are motivated by different factors.

A utility function is created for each person to determine the order in which each person will rank the available positions. The utility function will be based on the six personnel and position characteristics, with the concept of homophily continued here. The model assumes that personnel seek to be assigned to positions that are most alike themselves with regards to five of the six simulation characteristics, these being rank, branch, skills, career development and performance. With respect to location, each person has preferences for the different locations, and a person does not necessarily seek to remain in their current location (even though from a Service perspective it is desired that personnel remain in their current location).

The utility function used to generate personnel preferences is given by:

$$U_t = (W_r \times U_r) + (W_e \times (U_b + U_s)) + (W_l \times U_l) + (W_{cd} \times U_{cd}) + (W_p \times U_p)$$

where, U is the utility, W is the weight, and the subscripts t, r, e, b, s, l, cd and p indicate total, rank, employment, branch, skill, location, career development and performance respectively. The component utility values are outlined further below.

a. Rank Utility

The personnel utility values are based on the same utility table used for position rank utility. This reflects that personnel will mostly seek positions of the same rank, but in some cases may identify positions of different ranks in their preference lists.

b. Employment Utility

The employment utility is determined from consideration of each position's branch and skill. The approach of combining branch and skill is used because some positions are not specific to any particular skill set, therefore it is less appropriate to use a skill utility and apply a separate weight to the skill. The personnel branch utility is based on the position branch utility table previously provided. Personnel record a skill utility value of one for positions whose skill matches the person's skill, or for positions that are not specific to any skill set; otherwise, a skill utility value of zero is recorded.

c. Locality Utility

Whereas the organization's preferences are structured in a manner that seeks to minimize the number of geographic moves, the personnel do not necessarily seek to remain in their current location. In the Australian Army, personnel can choose to submit an Electronic Preferences and Restrictions (EPAR) form to identify to their career manager any preferences or restrictions that should be taken into account when assignments are determined. Personnel can identify up to three desired locations in the EPAR. While the EPAR information is considered confidential, nonidentifiable data was released for the purpose of this research. Table 14 shows the percentage of personnel who seek each locality as a first, second or third preference, and is based on 11,354 EPAR forms submitted from 1 January 2005 to 1 October 2010.

Locality	Preference 1	Preference 2	Preference 3
ACT	6.0%	5.5%	6.0%
NSW OTHER	7.1%	7.0%	10.6%
NSW SYD	14.9%	15.0%	10.9%
NT DARWIN	7.9%	6.4%	7.4%
OVERSEAS	3.8%	3.9%	6.2%
QLD NORTH	16.0%	13.5%	14.5%
QLD SOUTH	25.3%	28.0%	19.7%
SA ADELAIDE	3.8%	4.6%	4.6%
TAS	0.9%	1.2%	1.8%
VIC MELB	5.8%	6.7%	7.3%
VIC OTHER	2.3%	2.7%	4.0%
WA OTHER	0.5%	0.5%	0.9%
WA PERTH	5.9%	5.0%	6.0%

Table 14 Locality Preferences

Personnel agents are randomly allocated three locality preferences according to the preference distribution in Table 14. The same locality is not identified more than once for any person agent.

A locality utility value of 1.0 is recorded if a position's locality is the same as a person's first preference, a locality utility value of 0.5 is recorded if a position's locality is the same as a person's second preference, and a locality utility value of 0.25 is recorded if a position's locality is the same as a person's third preference; otherwise a locality utility value of zero is recorded.

d. Career Development Utility

A career development utility value of one is recorded where the career development that a position provides matches what a person requires; otherwise a career development value of zero is recorded.

e. Performance Utility

Continuing the model of homophilous attraction, personnel self-select to those positions that are most alike in performance requirements. The personnel performance utility is determined using the same performance utility table used for calculation of position performance utility.

f. Determining Attribute Weights

Weights are used to indicate the extent to which personnel are motivated by each of the attributes (rank, employment, locality, career development and performance); however, not all personnel will be motivated to the same extent by each of the attributes. Personnel primarily seek assignment to positions of the same rank; therefore, the rank is the primary attribute for every person. The remaining four attributes are then randomly ordered for each person to indicate attribute priorities. Based on each person's ordering of attributes, weights are applied to the attributes using Rank Order Centroid weights (outlined earlier). Table 15 shows the Rank Order Centroid weights for five attributes.

Rank	First	Second	Third	Fourth	Fifth
Weight	0.4566	0.2567	0.1567	0.09	0.04

Table 15 Rank Order Centroid Weights for Five Attributes

4. Exploring Outcomes

A series of computational experiments is to be conducted. The data for these experiments primarily utilizes actual Australian Army data, supplemented where necessary with data based on assumptions. The preferences required for two-sided matching are generated from utility values, with weights able to be varied to apply different importance to the various attributes. The aim is to investigate how assignment outcomes are affected by varying the attribute weights; specifically, how the number of participants matched and the utility of matched personnel and positions changes. Further, for the matched positions, the experiments will examine how well the service and unit attributes are satisfied as the weights are adjusted.

Data for the agents and their preferences are stored in Microsoft Access tables. Visual Basic code is used to implement the two-sided matching algorithm and generate the assignments.

IV. DECISION-MAKING ABILITY IN ASSIGNMENT PROCESSES

In Chapter II, a number of studies on the effects of information overload were reviewed. While the existence of information overload is generally no longer debated, mixed results have been encountered in previous studies. Some authors (Keller & Staelin, 1987) state that too much information will lead to information overload. Others (Hahn, Lawson & Lee, 1992) suggest that information overload is only encountered under time pressure.

While information overload has been explored in a variety of contexts including consumer behavior, bankruptcy prediction, finance and management, no previous research has explored decision maker outcomes and the effects of information overload in assignment scenarios. That is, a situation where a decision maker must match elements of one group with elements from another group. Such assignment scenarios are more complex because multiple decisions are required. For example, it is not a matter of simply selecting the best product from a list as in the case of consumer decision making (e.g., Jacoby, Speller & Berning, 1974). Instead, preferences or attributes from each group require consideration, and each individual assignment will limit the subsequent options by removing the matched pair from further consideration. While matched pairs could potentially be reconsidered if the decision maker finds a better alternative as they progress through the process, it is unclear whether the decision maker would systematically review decisions to see if different alternatives are available.

Chapter III, Section A outlined the design of experiments to examine human decision making in assignment processes. Scenarios involving assigning personnel to positions were created. Each person and position was given unique characteristics, with homophily forming the basis for optimal assignments. That is, the subjects were required to match the personnel and positions so that, to the maximum extent possible, the characteristics of each matched pair are aligned. The personnel and position options were designed so that no matches clearly dominated all others, and tradeoffs would be necessary.

This chapter will examine the results of the experiments, and explore the following issues: how "quality" varies as assignment scenario complexity increases, the different tradeoffs that participants make, the learning effect that participants exhibit, how well subjects self-assess their performance, and a comparison of results to optimal outcomes.

A. MEASURING COMPLEXITY

In these assignment decision-making experiments, the complexity of each scenario is related to both the number of personnel to be assigned and the number of attributes to be considered. This is similar to the study by Henry (1980) who examined information processing accuracy in consumer behavior, where complexity was defined as the interplay between the number of brands, number of attributes per brand and the scaling employed by each attribute.

Assignment quality is scored in two parts: how well within an assigned pair the person's preferences are met; and how well the position's attributes are satisfied (with the position seeking a person that exhibits characteristics closest to its requirements - homophilous attraction). Further detail is provided in Chapter III, Section A3.

B. RESULTS

1. Subject Demographics

A total of 81 subjects participated in the experiments. These subjects were from three master's-level classes of approximately equal size, with 28, 27 and 26 subjects in the classes. Participation in the experiments was voluntary for all subjects. There was no apparent systematic distribution of subjects between classes, and so the distribution of subjects was considered random. All subjects within a class completed the scenarios in the same order, but the order in which scenarios were presented differed between classes.

Participants in the experiments were all military officers between rank of O-2 to O-5, and were between 24 and 44 years of age, with an average age of 33 years. The participants had between 2 and 24 years of commissioned service, with a median length of service of 12 years. 69 of the subjects were male with the remaining 12 being female.

2. Excluded Data

There were a total of 25 subjects with some form of anomaly in their responses. These subjects with anomalies were spread across the three classes. For 18 of the 25 subjects with anomalies, the anomalies were minor and only affected one of their assignment scenarios. These anomalies occurred where two personnel were assigned to the same position (e.g., person 103 and 107 both being matched to position 9). These may be considered oversight errors, and would likely be recognized and rectified if time allowed; Hahn et al. (1992) claim that information overload is created by time pressure. Six of the cases were in assignment contexts involving 10 personnel, and the remaining eleven were in contexts involving 15 personnel. None of the assignment contexts involving 5 personnel were affected. This in itself provides an early confirmation that the subjects experienced greater problems with scenarios that assigned more personnel. That is, as complexity increases, so too does the error rate.

The common mistake performed by the 18 subjects described above was that two personnel were assigned to the same position, a situation that was not allowed. To resolve these anomalies and maximize the number of usable results, a standard rule was used to resolve the anomalies. Of the two personnel who were incorrectly assigned, the highest numbered person was rematched to their next available preference, and if none of their preferences were available, they were matched to the most suitable available position (suitable from the position's perspective).

Seven subjects had errors that were more complex and could not be resolved in a standard manner as described above. These errors involved multiple personnel and positions in an assignment scenario, or in some cases, there were numerous personnel left unmatched in an assignment scenario such that a complete response could not be determined. Four of these subjects had complex errors in only one assignment scenario. The other three subjects had complex errors across almost all of the assignment instances. The errors associated with these subjects indicate either a complete misunderstanding of the assignment process or an inability to complete the tasks in the allocated times. All results for these seven subjects were excluded from analysis.

After the seven subjects were excluded, there were 74 subjects with usable data remaining. The number of subjects by group were 26 (two excluded), 24 (three excluded) and 24 (two excluded).

3. Overview of Results

As previously described, subjects were presented six scenarios, with each scenario representing a different number of personnel to assign and attributes to consider. Table 16 shows the scenario naming convention, which will be used in this analysis. To examine issues such as learning effects, which will be examined later, the classes were provided the scenarios in different orders; Class 1 in the order A, C, E, B, D, F, Class 2 in the order E, C, A, F, D, B and Class 3 in the order B, D, F, A, C, E.

		Attributes	
		3	5
Assignment Size, Number of Personnel to be Assigned	5	A	B
	10	C	D
	15	E	F

Table 16 Assignment Scenarios

As previously outlined, points were awarded to subjects according to how well they satisfied the positions' requirements and the personnel preferences in each scenario. For each scenario, the average preference ranking of the assigned personnel is calculated, as is the average utility of assigned positions. These are converted to points: the personnel points are calculated by taking the inverse of the average preference rank (inverse because a lower rank is preferred), and multiplying by 100; the position points are calculated by multiplying the average position utility by 200. The different multiplicative factors were used so that the position requirements should factor foremost in the assignment considerations, in the same way that military assignments put the service needs ahead of the individual (Arnhart, 2007).

A summary of each subject's points in each scenario is provided at Appendix C.

The first issue to examine is whether subjects earned points in accordance with the experimental design. That is, whether subjects earned more points from satisfying the position attributes as opposed to the personnel preferences. The position points and personnel points that each subject earned from each scenario are summed across the six scenarios to provide total position and total personnel points for each subject. The results for each subject are displayed in Figure 11, ordered within each class group from best to worst overall outcome. Also shown are the total points that each subject achieved, that is, the sum of all position and personnel points from all six scenarios. The dotted lines indicate the optimal solution. As can be observed from the diagram, Class 1 covered subjects 1 to 26, Class 2 included subjects 27 to 50, and Class 3 involved subjects 51 to 74. Across all subjects, the average personnel score was 292.4 points, and the average position score was 740.4. Therefore, in accordance with the experiment design, positions points were approximately twice the level of the personnel points, reflecting the design intent that the subjects should give position requirements priority.

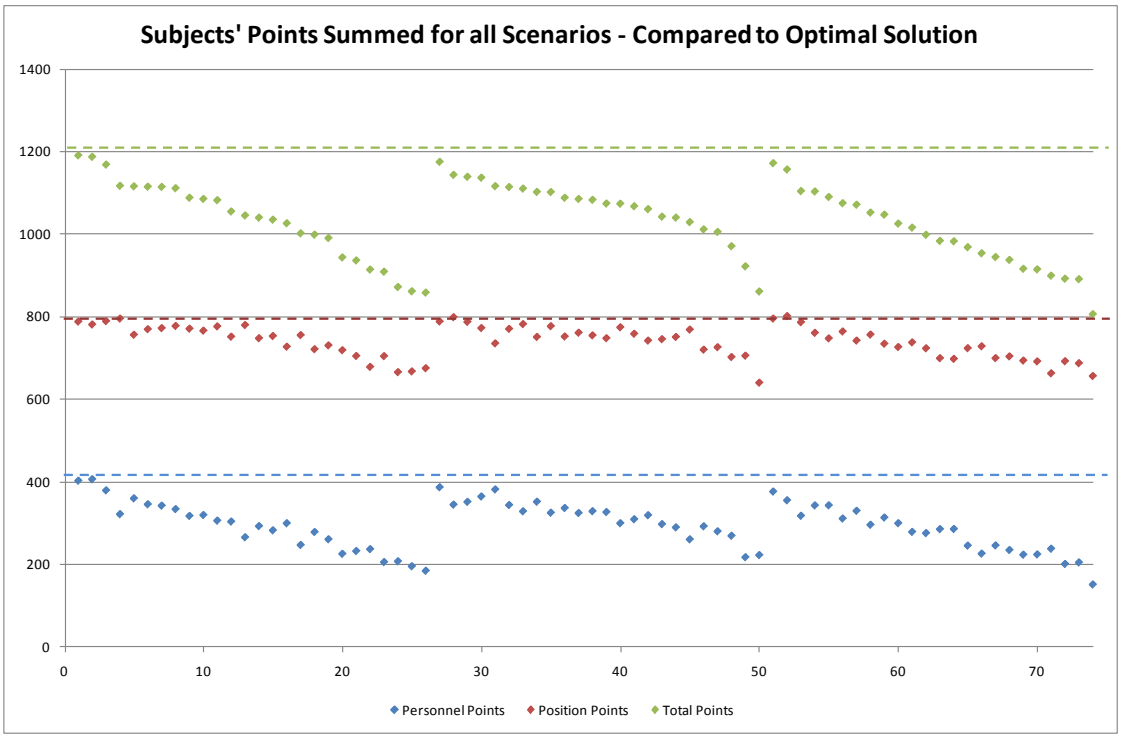

Figure 11. Summary of Results by Subject

Overall, Class 2 performed better than the other two classes in terms of total points achieved across all scenarios. Class 2 had the highest average and median total points of all three classes, with Class 1 ranked second overall and Class 3 third in terms of these metrics. Further, only three subjects from Class 2 achieved less than 1,000 overall points, whereas there were nine from Class 1 in this situation and thirteen from Class 3. The reason for this variation between classes is not clear, however two possibilities are suggested: first, that this relates to the order in which results were presented to the classes (this will be examined further later); second, that the classes were not randomly distributed in terms of abilities. The results will primarily be analyzed within the classes to minimize the effect of any differences between classes.

Figure 11 provides some interesting initial insights before further analysis of the results is undertaken. The variability in the personnel points is greater than in the position points, with personnel points ranging between 150 to 405 (a range of 255, or 61% of the personnel points at the optimal solution), and position points ranging between 639 to 802 (a range of 173, or 22% of the position points at the optimal solution). In part this may reflect the experimental design, where position points were afforded more weight than personnel points. Consequently, it is reasonable to expect that the subjects would focus on position requirements and therefore position points would generally be closer to an "optimal" outcome than personnel points.

Across all classes, the average total position points attained by the subjects was 740, representing 93% of the optimal solution, while the average total personnel points was 292, representing 70% of the optimal solution. The highest performing subjects achieved results that were close to the optimal outcomes for both personnel and positions. However, for the lowest performing subjects the reduction in personnel points was relatively greater than the reduction in position points; for example, the 10 lowest performing subjects had an average total position points of 672 (85% of the optimal) and an average total personnel points of 203 (49% of the optimal). This suggests that, even though the subjects were aware that it was important to consider both the position and personnel outcomes, the subjects may have been overly focused on one aspect of the situation (meeting the position requirements) to the detriment of the other (meeting the

personnel requirements). Such a situation is likely to occur in real life decision-making scenarios; that is, the decision maker may become too focused on one aspect of the decision-making process to the detriment of other elements. This highlights the need for decision support in complex assignment processes, because without it, decision makers may not appropriately consider each of the attributes.

This section has provided a broad overview of the results. The next section will examine results achieved by each class in each scenario. This includes a review of learning effects.

4. Comparing Results Between Classes

The results of experiments of the type undertaken here can be expected to demonstrate signs of learning effects (Kagel & Roth, 1995). That is, as subjects undertake more assignment decision-making scenarios, they gain experience and enhance their decision-making techniques. These learning effects provide a number of problems for analysis, with Friedman and Sunder (1994) stating that "no valid statistical correction presently is available" (p. 98). Therefore, in the analysis undertaken here it is necessary to compare results for classes who are at similar stages of the learning process; that is, classes that have undertaken a similar number of scenarios.

There are various ways to compare the scenario results between classes, including a comparison of median points, maximum points, points range and inter-quartile range to name some. To examine the learning effects, it is useful to compare how each class performed on a particular scenario, noting that the classes undertook the scenarios in different orders. Box and Whisker plots are used to compare these key statistics from each scenario. Figure 12 to Figure 17 show the results for each of the six scenarios examined with each plot showing the results for all three subject pools. For comparative purposes, the optimal outcome for each scenario is shown by the dotted red line on each plot.

The points plotted on the y-axis are each subjects combined personnel and position points from each scenario. Using these combined scores, the box and whisker plots present, by class, the range and quartiles for the subjects within each class.

While some subjects performed well on the first assignment scenario they were presented, it is clear the majority of subjects were learning the process and developing their strategies for at least the first scenario they were presented. This is seen by comparing the results of classes who were undertaking their first scenario and comparing to classes who had already completed some scenarios. A description of the results from each scenario is outlined below.

a. Scenario A

Class	1	2	3
Presentation Order	First	Third	Fourth

Table 17 Scenario A – Presentation Order by Class

Classes 2 and 3 show similar results, and their results are distinctly better when compared to those of Class 1. For Classes 2 and 3, there are identical outcomes in terms of lower quartile Q1, median, upper quartile Q3 and maximum outcomes. Only the minimum points for Classes 2 and 3 differ. By comparison, with the exception of the maximum points, all these key statistics for Class 1 are lower than for Classes 2 and 3. Further, the range and inter-quartile range for Class 1 is higher than for Classes 2 and 3. The similarity of results between Classes 2 and 3, and their difference from Class 1 supports the prediction of learning effects; that is, subjects will perform better after they have had experience with other scenarios.

This scenario is expected to be the simplest of the scenarios presented, given that it involves the least number of attributes to consider and least number of personnel to assign. Yet despite the apparent simplicity, there are subjects in Classes 2 and 3 (who undertook this scenario as their third and fourth scenario respectively) who achieved lower outcomes than some people from Class 1 (who undertook this as their first scenario). This is likely to be a result of different cognitive abilities; Henry (1980) demonstrated that information processing accuracy is related to cognitive abilities. Therefore, with a distribution of cognitive abilities between classes, it is hypothesized that the subjects from Class 1 who had higher cognitive abilities were able to achieve

superior results than subjects from Classes 2 and 3 who had lower cognitive abilities, despite the learning effect advantage for those in Classes 2 and 3. Ideally, information about the cognitive abilities of the subjects would be used to confirm this hypothesis, however permission to use such information was not sought at the time, nor were tests undertaken to assess cognitive abilities (as was done by Henry (1980)).

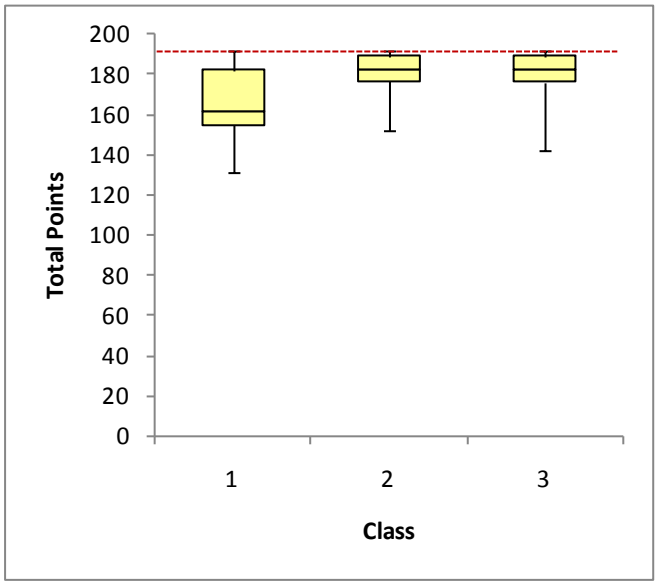

Figure 12. Results for Scenario A: 5 Personnel, 3 Attributes

b. Scenario B

Class	1	2	3
Presentation Order	Fourth	Sixth	First

Table 18 Scenario B – Presentation Order by Class

In Scenario B, the performance from the three classes is generally as expected based on the order in which they completed the scenarios. Class 2, who was completing its sixth scenario and had the benefit of most learning opportunity, showed the best performance in terms of key statistics. This was followed by Class 1 who was completing its fourth scenario. The differences in outcomes between Classes 1 and 3 are perhaps not as great as would be expected based on learning opportunities. While the top 50% of subjects in Class 1 clearly achieved better results than the top 50% of Class 3, the

bottom 50% of subjects from each class show similar statistics. Further, the inter-quartile range for Class 1 is larger than for Class 3, although it is not clear why this would be the case. Generally it would be expected that the range and inter-quartile range would decrease as experience increases. With greater experience it is expected that the results would improve and converge towards an optimal outcome, while those completing the scenarios with less experience would be expected to show a larger "tail" of results at the lower end.

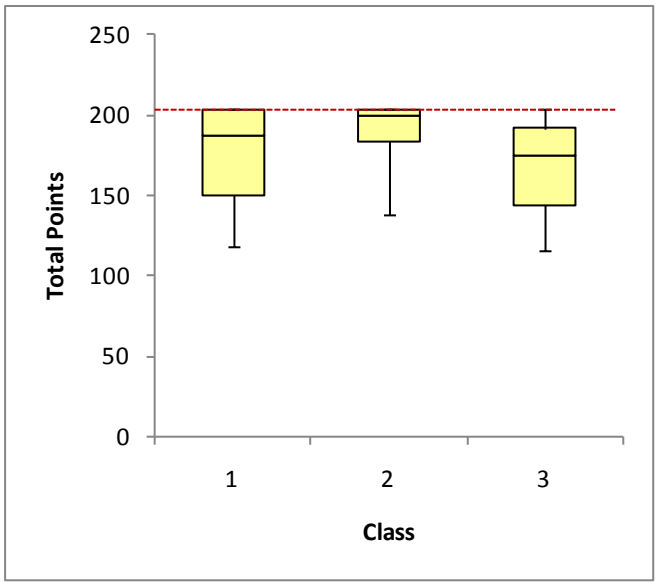

Figure 13. Results for Scenario B: 5 Personnel, 5 Attributes

c. Scenario C

Class	1	2	3
Presentation Order	Second	Second	Fifth

Table 19 Scenario C – Presentation Order by Class

As hypothesized, Class 3 (who are undertaking its fifth scenario) outperformed Class 1 (completing its second scenario) on every statistic shown except the maximum points. Also as expected, the inter-quartile range for Class 3 is lower than Class 1. However, the results for Class 2 (completing its second scenario) are not as expected when compared to Class 3. Class 2 outperforms Class 3 in most respects; higher

maximum, median and Q1 quartile, and a smaller inter-quartile range. The only key statistic where Class 2 has a lower performance than Class 3 is the minimum points. The reason for this is not clear; however, across the various scenarios, Class 2 consistently performs well. This is seen in Scenario C by the performance of Class 2 compared to Class 1, with both of these classes completing their second scenario. The relatively strong performance of Class 2 compared to Classes 1 and 3 may indicate an uneven spread of cognitive abilities between classes.

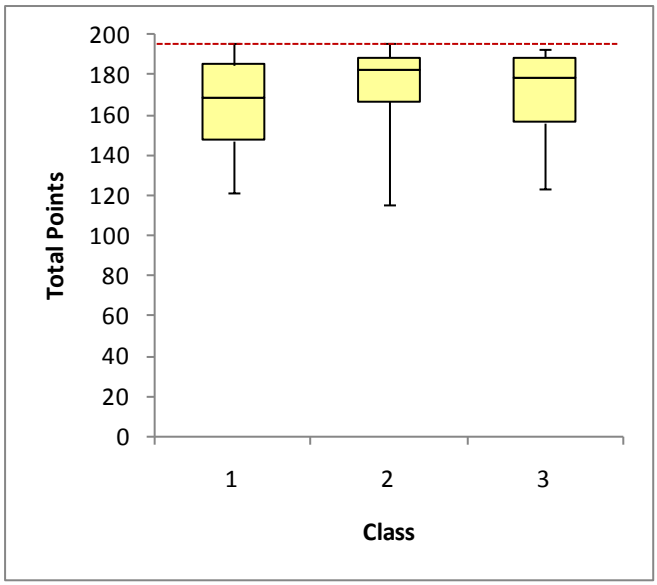

Figure 14. Results for Scenario C: 10 Personnel, 3 Attributes

d. Scenario D

Class	1	2	3
Presentation Order	Fifth	Fifth	Second

Table 20 Scenario D – Presentation Order by Class

Classes 1 and 2 each completed Scenario D as their fifth scenario. Following the comparison of these two classes in Scenario C, once again Class 2 outperforms Class 1 in every respect in Scenario D, apart from the two classes having the same minimum points. This again lends support to the likelihood that Class 2 may have an above average spread of cognitive abilities when compared to the other classes.

As hypothesized, with the exception of the maximum points the key statistics for Class 3 (undertaking its second scenario) are lower than Classes 1 and 2. The inter-quartile range for Class 3 is greater than for Classes 1 or 2.

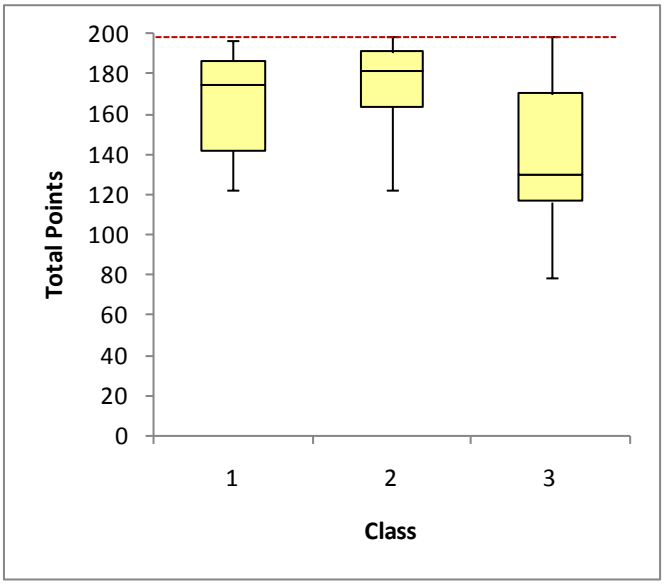

Figure 15. Results for Scenario D: 10 personnel, 5 Attributes

e. Scenario E

Class	1	2	3
Presentation Order	Third	First	Sixth

Table 21 Scenario E – Presentation Order by Class

The results for Scenario E are entirely as expected based on the order in which the classes performed this scenario. That is, the results for Class 2 (undertaking its first scenario) are dominated by Class 1 (undertaking its third scenario), and the results for Class 1 in turn are dominated by Class 3 (undertaking its sixth scenario). As expected, the inter-quartile range also decreases with experience.

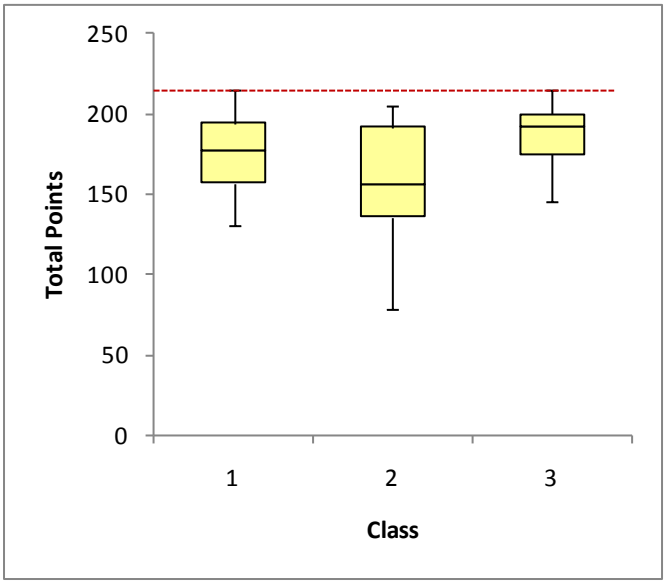

Figure 16. Results for Scenario E: 15 personnel, 3 Attributes

f. Scenario F

Class	1	2	3
Presentation Order	Sixth	Fourth	Third

Table 22 Scenario F – Presentation Order by Class

As expected, Class 1 (undertaking its sixth scenario) achieves higher results for each of the key statistics when compared to Class 3 (undertaking its third scenario). Unexpectedly the inter-quartile range for Class 3 is marginally smaller than for Class 1, although the inter-quartile range is at a much lower level for Class 3 when compared to Class 1.

Contrary to the hypothesis, Class 2 (undertaking their fourth scenario) achieves a higher Q1, median and Q3, and has a lower inter-quartile range when compared to Class 1. Once again, this suggests that the cognitive abilities in Class 2 were above those of the other classes.

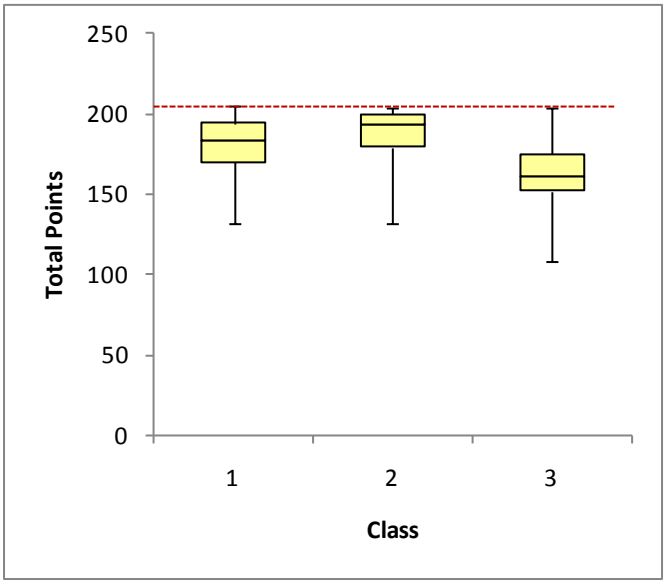

Figure 17. Results for Scenario F: 15 Personnel, 5 Attributes

g. Summary of Class Results

The results indicate that there may not have been an even spread of abilities between classes. Class 2 achieved the best results in terms of highest median points, quartile Q1 and Q3 points, and lowest inter-quartile range for every scenario except for Scenario E, which was the first scenario undertaken by Class 2. If repeating the experiments, it would be advantageous to gather additional information regarding individual cognitive abilities. This may include seeking information such as GPA or grades for particular subjects.

It was hypothesized that the range and inter-quartile range would decrease with experience, holding complexity as a constant. It was also hypothesized that the range and inter-quartile range would increase with more complex scenarios, holding experience levels constant. The results do not provide conclusive support for or against these hypotheses. Further experiments would be required to examine these hypotheses. In conducting further experiments it would be advantageous to randomize the order in which the personnel and positions appear in each scenario. In the experiments conducted here, although the classes received the scenarios in different orders, the layout of personnel and positions within each scenario was identical. That is, person 101 appeared

on the left of the page, followed by person 102, then 103 and so on. It appears that the layout of the information may have led the subjects to a certain range of outcomes and reduced the variability between outcomes.

Based on an analysis of the results, a number of subjects developed a simple heuristic whereby they worked left to right through the information presented on the experiment paperwork. Every person in the experiments needed to be assigned to a position, although not every position needed to be assigned to a person given that there were more positions than personnel. On this basis, an analysis of each person's assignment against their listed preferences was undertaken. Across each of the 74 subjects, a count was taken to determine the number of times each person was assigned to a position that was not one of the person's preferences. In the U.S. military, a person is considered to be "slammed" if they are assigned to a position that they do not wish to be assigned to. This terminology is used here.

Figure 18 to Figure 23 show the number of subjects who slammed each person in the assignment scenarios. A feature of these figures is that personnel who were listed first (reading left to right) were less likely to be slammed, while personnel who were listed last were more likely to be slammed. It is only in Scenario B where this does not appear so clearly, although the number of subjects who slammed personnel in the 5 personnel scenarios was less than for the 10 or 15 personnel scenarios (vertical scales are held constant for ease of reference). No personnel are slammed at the optimal solutions.

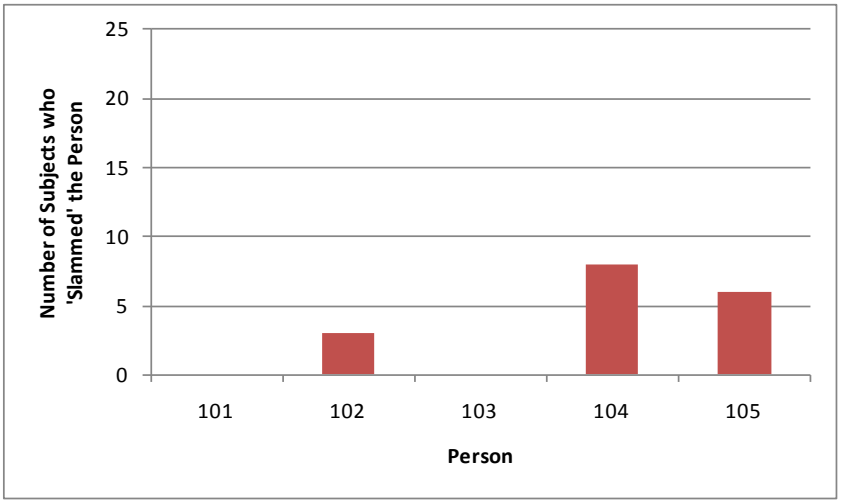

Figure 18. Number of times Personnel Slammed in Scenario A: 5 Personnel, 3 Attributes

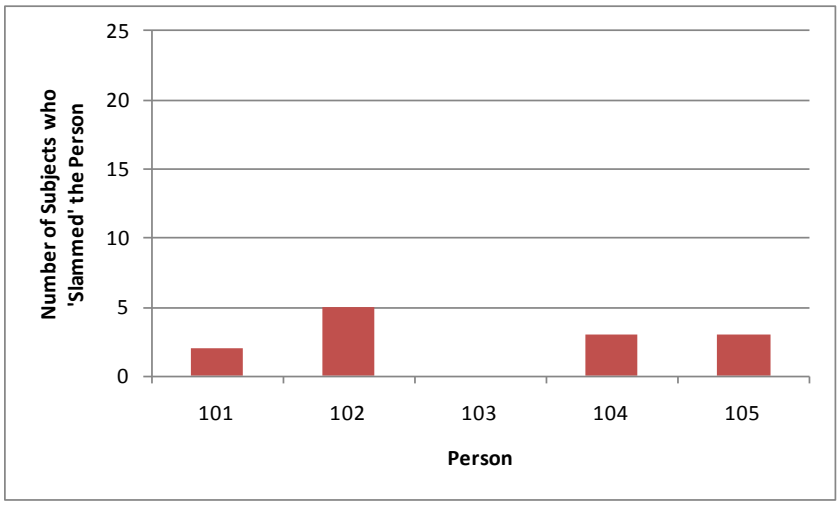

Figure 19. Number of times Personnel Slammed in Scenario B: 5 Personnel, 5 Attributes

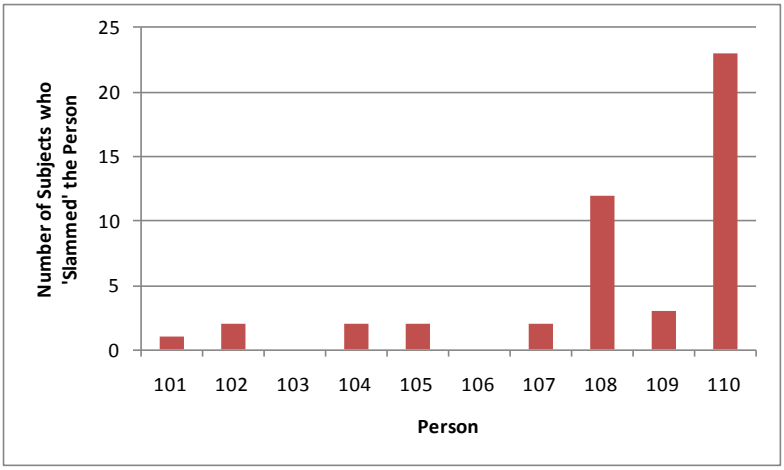

Figure 20. Number of times Personnel Slammed in Scenario C: 10 Personnel, 3 Attributes

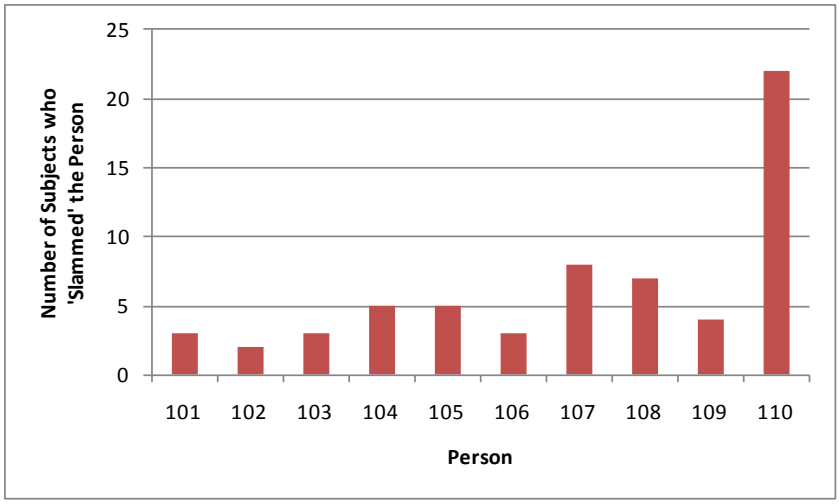

Figure 21. Number of times Personnel Slammed in Scenario D: 10 Personnel, 5 Attributes

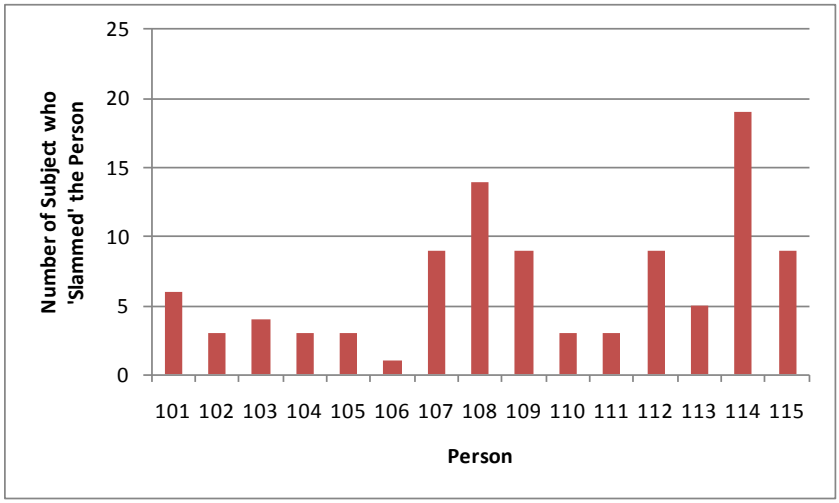

Figure 22. Number of times Personnel Slammed in Scenario E: 15 Personnel, 3 Attributes

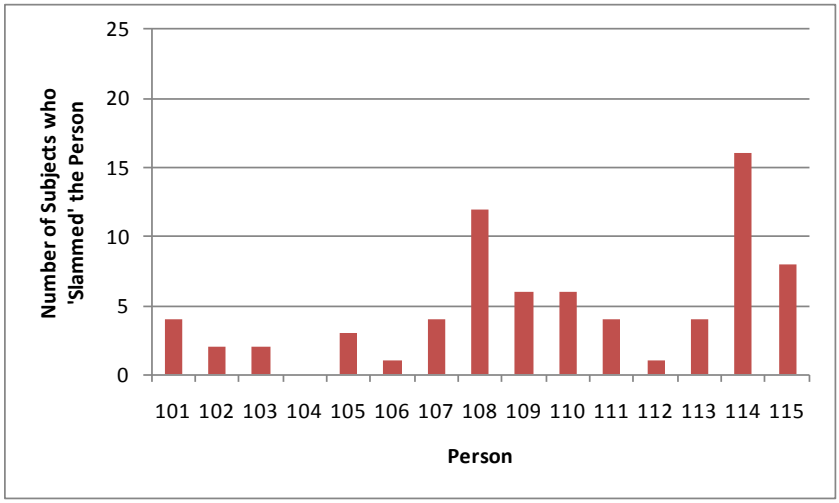

Figure 23. Number of times Personnel Slammed in Scenario F: 15 Personnel, 5 Attributes

The observation that some subjects developed a heuristic whereby they assigned personnel according to the order in which the personnel were presented has important ramifications. In real life situations, the decision outcome could be influenced by the order in which a decision maker is presented the information, even if the decision maker is provided with all of the information simultaneously (as opposed to the information being trickle-fed). Further experiments would be necessary to eliminate the possibility that these observations are related to the design of the scenarios. The same

experimental scenarios could be used, but with the information within scenarios being presented in different orders for different groups of subjects.

5. Trading Off Between Personnel and Positions

Some subjects performed well for both personnel and positions, however other subjects achieved a poor tradeoff. That is, they could have achieved more personnel points with minimal impact on the position points, and vice versa.

Figure 24 to Figure 29 show, for each scenario, the personnel and positions points that each subject achieved. The results for all subjects from all classes are shown on each figure. Each subject is represented on the figures by a single marker. In some scenarios, particularly Scenarios A and B, several subjects achieved the same outcomes, and therefore the number of markers on the figures is less than the number of subjects.

In each of the scenarios, it is possible to see situations where there are two or more points aligned vertically. In these cases, the subjects who are lower on the y-axis (have lower position points) could have achieved a better average outcome for the positions without the average personnel outcome being worse; this does not guarantee that some individual personnel or positions in the scenario would not necessarily be worse off, but the average result would be better, and that reflects the incentives included in the experimental design. Similarly, it is possible to see situations where there are two or more points aligned horizontally in each of the scenarios. In these cases, the subjects who have the lower score on the x-axis (have lower personnel points) could have increased the personnel points without sacrificing the positions points; again, this is measured on average and does not guarantee that some individual personnel or positions would not have worse outcomes. Finally, the figures demonstrate that there are many dominated outcomes; that is, subjects who achieved lower position and personnel points when compared to other subjects' outcomes. Subjects who have dominated outcomes could have achieved better average results for both the positions and people.

If a linear trend line was drawn through each of the points in Figure 24 to Figure 29, the slope in every case would be positive. A negative trend line would indicate that subjects were systematically trading off personnel points against position points; that is,

the increase in one of these would be at the expense of the other. However, the positive trend line demonstrates that the points achieved are more related to the subjects' cognitive ability; those towards the lower left on the charts have achieved inferior results, while those who are towards the upper right have achieved superior results. The subjects represented in Figure 24 to Figure 29 are from classes who have different levels of experience with the experiments, and so some dominated outcomes may potentially represent subjects who have completed fewer scenarios and do not have as much experience.

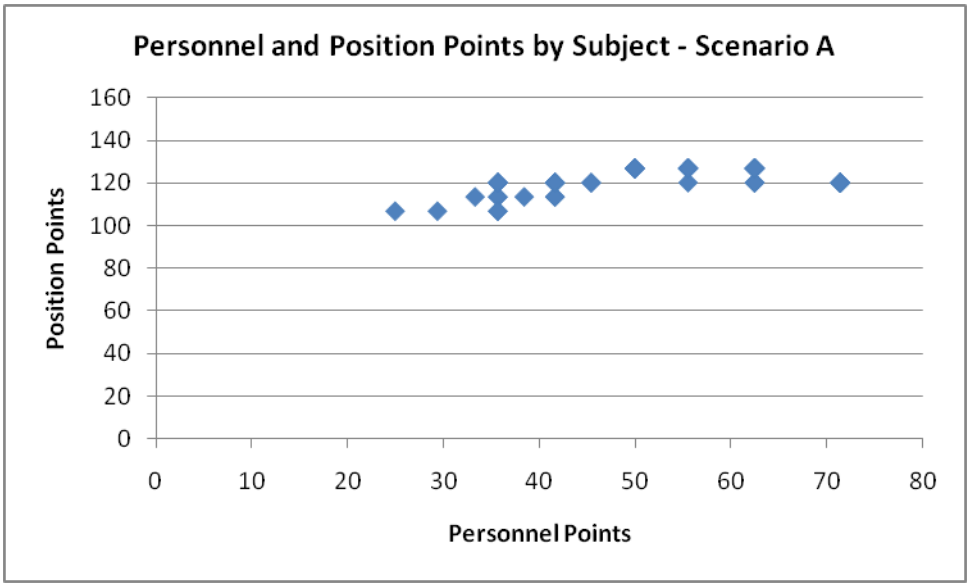

Figure 24. Personnel and Position Points – Scenario A

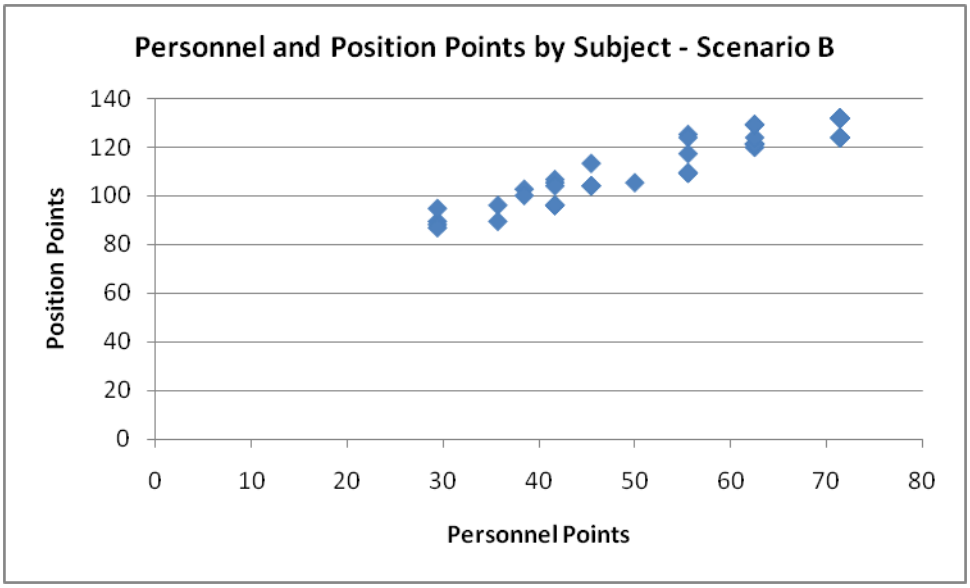

Figure 25. Personnel and Position Points – Scenario B

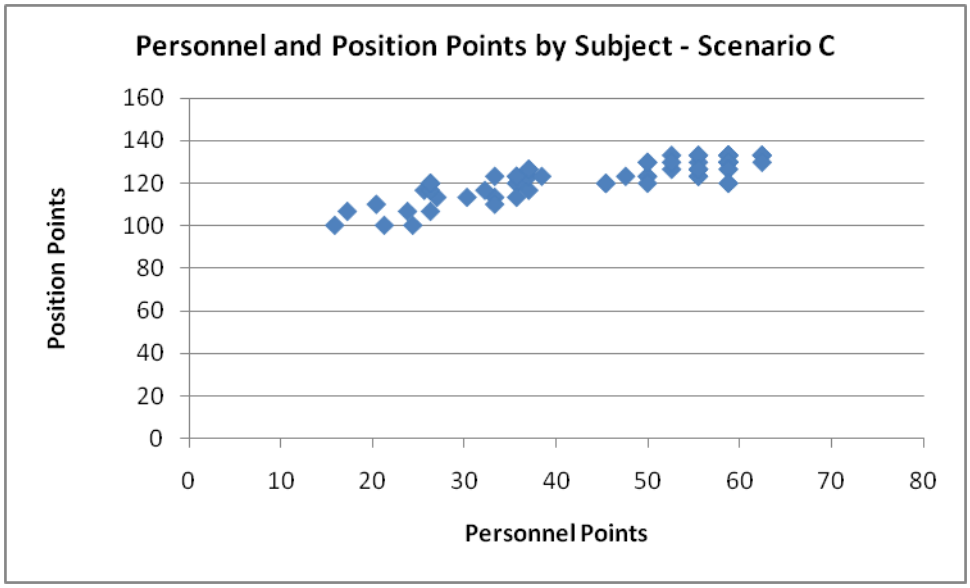

Figure 26. Personnel and Position Points – Scenario C

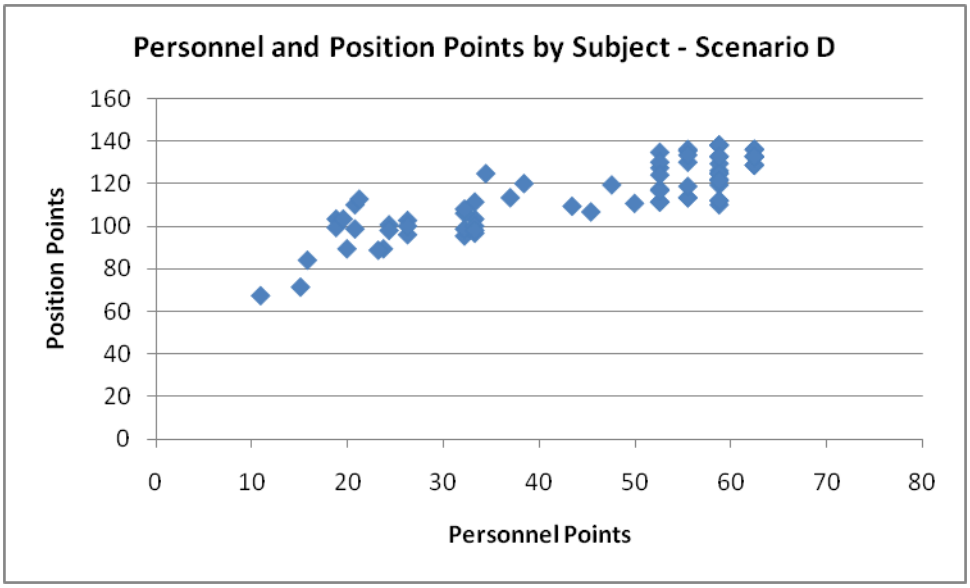

Figure 27. Personnel and Position Points – Scenario D

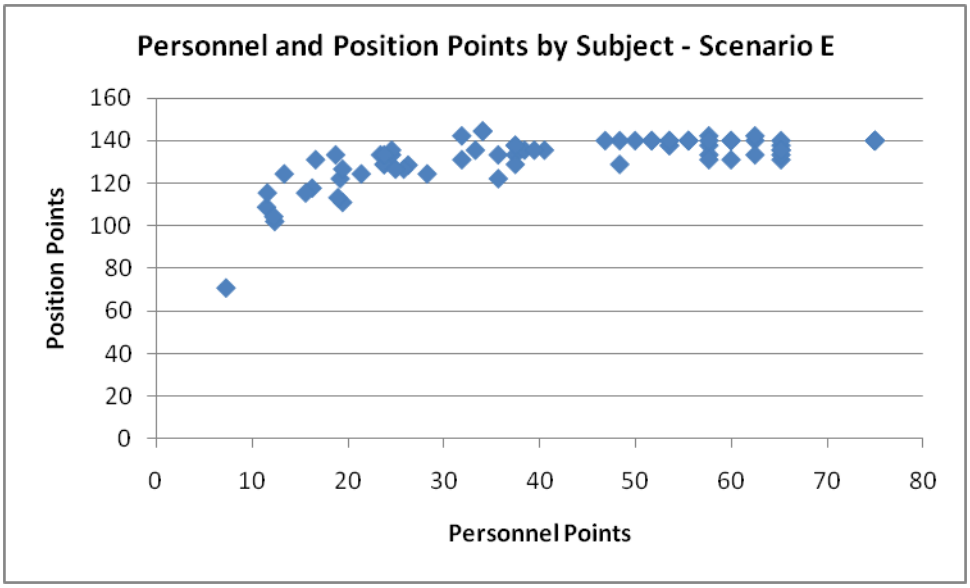

Figure 28. Personnel and Position Points – Scenario E

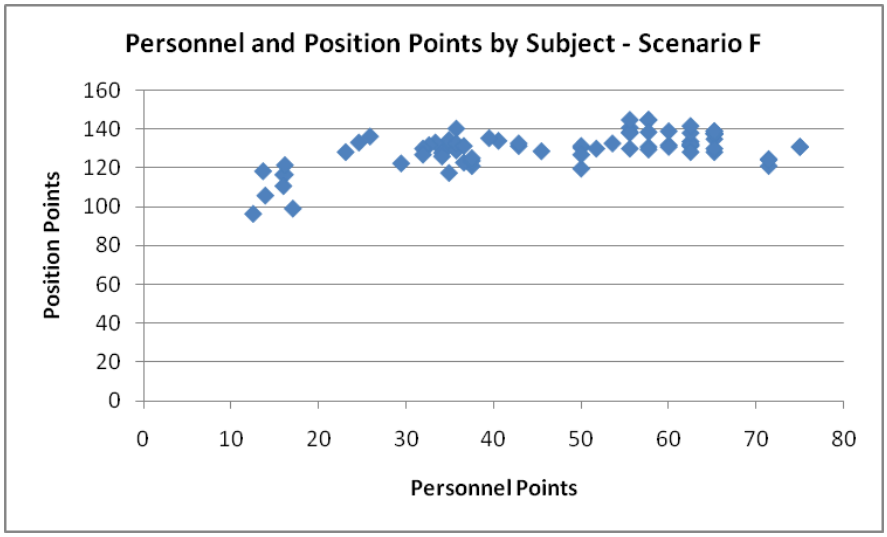

Figure 29. Personnel and Position Points – Scenario F

The slope of a linear line of best fit was determined for each scenario and is presented in Table 23. In all cases, the slope of the lines is less than 1.0. This indicates that the potential improvement in personnel outcomes is greater than the potential improvement in position outcomes, on average across all subjects. This also confirms that the subjects have focused on the position points, but potentially excessively so. If the consideration of personnel and position points was balanced, the slope of the trend line would be 1.0.

Scenario	Slope
A	0.24
B	0.99
C	0.54
D	0.87
E	0.41
F	0.31

Table 23 Slope of Line of Best Fit for Scenarios A to F

6. The Effect of More Attributes to Consider

The experiments were designed to examine how assignment results would be affected by a modest increase in the number of attributes that decision makers must consider. Subjects made assignment decisions in scenarios that involved either three or five attributes, a number that is quite within the bounds of what could be expected in a military assignment situation; for example, Klingman and Phillips (1984) state that the U.S. Marine Corps assignments may consider from 3 to 15 attributes, while Liang and Buclatin (1988) state that U.S. Navy assignments may involve up to nine attributes. Therefore, the number of attributes in these experiments remains at a lower level of complexity compared to what might be experienced in a military assignment situation.

Prior to undertaking the experiments, two issues were identified as being likely to cause difficulty in analyzing the effect of more attributes. The first issue relates to controlling for the learning effects. The subjects can be expected to improve their decision-making technique as they gain experience. For example, if a subject performs better for the 5-3 scenario (5 personnel – 3 attributes) after having earlier completed the 5-5 scenario, it is necessary to determine whether the improvement is related to a change in complexity or the benefit of experience.

The second issue to resolve is the ability to compare results between scenarios involving the same number of personnel to be assigned. That is, the results of the 5-3 scenario must be compared to the 5-5 scenario. If the specification of the 5-5 scenario was completely unrelated to the 5-3 scenario, then there would be no way to directly compare results between the two scenarios. As a result, the scenarios with 5 attributes were designed so that the personnel preferences were identical, and the three primary attributes for the positions were the same. The only difference is that two secondary attributes were added to the five attribute scenarios, and these two secondary attributes were given weightings, the sum of which were less than any of the individual primary attributes. The assignment scenario designs are outlined in Appendix D. Effectively, removing the secondary attributes from the scenarios with five attributes makes the scenarios identical to the three attribute scenario (for scenarios involving the same number of personnel). This makes the situation akin to a person simplifying the five

attribute scenario by ignoring or removing the secondary attributes. This design makes it possible to determine what outcome each subject could have achieved for the five attribute scenarios if they had used the identical assignments from the three attribute scenarios.

To determine whether the subjects' performance improved or worsened with the additional attributes, each subject's three attribute scenario assignments were applied to the five attribute scenarios. If the results that a subject achieved from a five attribute scenario are better than the results determined from applying the related three attribute assignments to the five attribute scenario, then the subject has performed better under the additional complexity of the five attribute scenario. Otherwise, the subject would have been better off ignoring the two secondary attributes for the five attribute scenario and treating the scenario the same as they did for the three attribute scenario. To account for learning effects, results are grouped by class in Table 24. The number of subjects in each situation is indicated in the table.

	Outcome	5 Personnel Scenario	10 Personnel Scenario	15 Personnel Scenario
Class 1	Better outcome with -5 scenario	20	18	18
	Worse outcome with -5 scenario	6	8	8
Class 2	Better outcome with -5 scenario	17	17	21
	Worse outcome with -5 scenario	7	7	3
Class 3	Better outcome with -5 scenario	11	10	8
	Worse outcome with -5 scenario	13	14	16

Table 24 The Effect of More attributes on Scenario Outcomes

To explain how the results of Table 24 were determined, consider the following example. A subject achieved 165 points (positions and personnel points combined) for the 5-3 scenario. The same subject achieved 137 points for the 5-5 scenario, however the points for the 5-5 scenario are not directly comparable to the 5-3 scenario because the 5-5 scenario involves additional attributes and different weighting. Therefore, it is not possible to determine whether or not the subject has achieved a better outcome in the 5-5

scenario. So, the subject's assignment decisions from the 5-3 scenario are applied to the 5-5 scenario. This produces an adjusted 5-5 scenario result of 145 points. That is, if the subject had applied the assignments from the 5-3 scenario to the 5-5 scenario, the outcome would have been better than they actually achieved in the 5-5 scenario (145 points compared to 137 points). This person is recorded as one of the subjects having a worse outcome with the 5-5 scenario in Table 24.

Table 24 shows that the majority of subjects from Classes 1 and 2 achieved better outcomes with the scenarios involving five attributes than they would have if their assignments from the three attribute scenarios were used. Both these classes completed the scenarios with three attributes before they undertook the scenarios involving five attributes. Therefore, they had the advantage of developing their decision-making strategies initially and could then apply these to the scenarios with five attributes. Consequently, the improved outcomes for the scenarios with five attributes could potentially be related to the learning effects. Despite the benefit of the learning effects, there was a small group of subjects whose results were worse with the five attribute scenarios. That is, they would have achieved better outcomes for the five attribute scenarios if they had simply used the assignments they made for the three attribute scenarios. For this group of people, the increased complexity apparently negated any benefit that they may have attained from practice and developing a decision-making strategy during the three attribute scenarios.

The situation was different for Class 3 where a smaller proportion of subjects achieved better outcomes for the five attribute scenarios. This likely reflects that Class 3 was presented with the five attribute scenarios first; these subjects were still learning their decision-making strategies while they were completing the more difficult five attribute scenarios. Despite this, there were still a number of subjects who managed to achieve superior results for the five attribute scenarios, even though they were potentially still developing their decision-making strategies.

The learning effects explains why a greater proportion of Class 3 achieved worse outcomes for the five attribute scenarios as compared to Classes 1 and 2, but it does not explain why some people in Class 3 would have achieved better results for the five

attribute scenarios. It is likely that this is because the five attribute scenarios resolved some indifference / ties that were harder to resolve in the three attribute scenarios. Therefore, when making the assignments for the three attribute scenarios, the subjects faced decisions that could have gone one way or the other, but when those results were applied to the five attribute scenarios, it made more sense to make the decisions differently.

The data from Table 24 was also examined at the individual level. For Classes 1 and 2, there was no consistency in terms of subjects who performed worse with the five attribute scenarios. That is, one subject may have achieved a worse outcome for the 5-5 scenario, but may have achieved a better outcome for the 10-5 scenario (when compared to the situation if the relevant three attribute scenario assignments were applied to the five attribute scenarios). However, for Class 3, six of the subjects returned worse outcomes for all of their five attribute scenarios.

The analysis shown here demonstrates that some subjects managed to handle the additional complexity of five attributes, however not all subjects accommodated complexity. The number of attribute examined in these experiments is at the lower end of the spectrum of what might be considered in military assignment processes. Therefore, unless decision makers are screened to ensure that they have the skills required to produce quality assignment decisions, there is a high risk that their assignment decisions may be sub-optimal absent decision support tools. This risk is likely to increase as the number of attributes (complexity) increases.

7. The Effect of More Personnel to Assign

When examining the effect of more attributes the scenarios were designed so that the five attribute scenarios were an extension of the three attribute scenarios; that is, the preferences and primary attributes were identical, but then secondary attributes were added to the five attribute scenarios. This enabled a comparison between scenarios that involved the same number of assigned personnel, e.g., a comparison between the 10-3 and 10-5 scenario. However, no similar method was found to enable a direct comparison

of scenarios involving different numbers of personnel to assign. Therefore a different method is required to examine how the assignment outcomes are affected by increasing the number of assigned personnel.

In some decision-making studies, decisions can be assessed as correct / incorrect. These studies then relate the probability of making correct decisions to the quantity of information presented, where quantity may be related to the number of items to consider or the amount of information that is available about each item. For example, Casey (1980) examined the effects of information load on bank loan officers assessing the risk of customers going bankrupt. In this study, information was available to identify whether the customers actually did / did not go bankrupt, although these outcomes were clearly not made known to the subjects during the experiment. By presenting more cases or more information per case, it is possible to relate the probability of correct assessments to the quantity of information presented.

In experiments of the type conducted in this research, each scenario is independent and cannot be compared directly to another scenario. That is, the points obtained in a five-person assignment scenario cannot be compared to the points obtained in 10-person assignment scenario. The experiment design is such that it may only be possible to achieve 180 points in a scenario with five personnel, but it may be possible to achieve 200 points in a scenario with 10 personnel. This does not necessarily mean that a subject has achieved a better outcome with the 10 personnel scenario, but rather it may be possible to better satisfy the position requirements and the personnel preferences. As a result, there is no simple comparison of points.

One approach to examining the relationship between decision-making quality and the number of personnel assigned is to compare the results achieved by the subjects to the optimal results. It is hypothesized that subjects would achieve outcomes relatively closer to the optimal at lower numbers of personnel to be considered, but the results would diverge as the scenario assignment size increases. Once again, learning effects need to be considered. Friedman and Sunder (1994) suggest that an extreme manner of controlling for learning effects is to discount all trials except for the last or next-to-last, but they also refer to other research that suggests that learning effects will take the form of an

exponential decay towards a behavioral equilibrium. As a result, in this analysis learning effects will be controlled by ignoring any results from a class that was performing one of its first three scenarios. It will be important to acknowledge that some learning effects are likely to continue beyond this point, but this process should exclude the most significant effects. Additionally, only scenarios involving the same number of attributes will be compared. That is, the 5-3, 10-3 and 15-3 scenarios will be compared as a group, and the 5-5, 10-5 and 15-5 scenarios will be compared as a group.

Once the results have been limited as outlined above, the results from the three attribute scenarios are limited to Class 3 (Class 3 performed all three attribute scenarios as their fourth to sixth scenarios), and the results from the five attribute scenarios are limited to Classes 1 and 2 (Classes 1 and 2 performed all five attribute scenarios as their fourth to sixth scenarios). Table 25 shows the median points from each class expressed as a percentage of the optimal points from the same scenarios. Results are read down the columns to examine the effect of increasing numbers of personnel.

	3 Attributes (Class 3)	5 Attributes (Class 1)	5 Attributes (Class 2)
5 Personnel	95%	92%	98%
10 Personnel	91%	88%	91%
15 Personnel	89%	89%	94%

Table 25 Class Median Points as Percent of Optimal Result

The results from Class 3 performing the three attribute scenarios are as hypothesized. That is, relative to the optimal results, Class 3 performed best with the least number of personnel to be assigned, and progressively achieved relatively worse outcomes as the number personnel to assign increased. However, the results from Classes 1 and 2 were not entirely as expected; both achieved the best outcomes for the scenarios with 5 personnel, but for both classes the results from the scenarios with 15 personnel were better than the scenarios with 10 personnel.

A limitation to the approach of comparing subject results to the optimal outcomes is that in the scenarios that involve fewer personnel there may be less

opportunity to deviate from the optimal outcome. As the number of personnel involved increases, the number of potentially different outcomes will increase, and therefore so will the potential for greater variability.

The results here are not sufficient to support or reject the hypothesis. Further work involving a larger number of subjects would be required to further test the hypothesis. Also, the experiments would need to be designed so that the information presented to the subjects appears in a random manner to ensure that the presentation order does not affect the results.

8. Subject Responses

a. Self-Assessment of Performance

Immediately after completing the experiments the subjects answered a series of questions about their decision-making strategies and provided a self-assessment of their performance. These questions were asked before the subjects received any feedback on their own performance or the performance of others in the class. One question asked was, “Based on overall points that you earned, please indicate how well you believe you performed in relation to the other members of your class.” Subjects were then asked to place themselves into quartile ranges. Of the 74 subjects with usable results, 69 responded to this question.

If subjects had perfect information about their own outcomes and that of the other participants, there would be an even distribution between the quartile ranges; that is, 25% of subjects would provide a self-assessment that they were in the top quartile, 25% would provide a self-assessment that they were in the next quartile, and so on. The aggregated results across all three classes are shown in Table 26 (quartile 1 represents the top quartile). As it turns out, the subjects had inflated opinions of their own assignment decision-making abilities; only 13% (9 of 69) of subjects felt that they would be in the bottom 50% of participants, and only 3% (2 of 69) felt that they would be in the bottom quartile. While 27 subjects (39%) correctly assessed their actual quartile, 34 (49%) believed that they would be in a higher quartile than their actual performance, and only 8 (12%) believed that they would be in a lower quartile than they actually were.

Table 27 to Table 29 show the results by individual class. The numbers in these tables represent the number of subjects fitting within each quartile. The diagonal boxes from top left to bottom right are subjects who correctly assessed their performance. Subjects falling in the boxes to the right of the diagonal are the subjects who assessed their performance as worse than actual. Correspondingly, subjects to the left of the diagonal are the subjects who assessed their performance as better than actual.

		Self assessed quartile			
		1	2	3	4
Actual quartile	1	12	5	0	0
	2	5	10	2	0
	3	5	8	4	1
	4	2	13	1	1

Table 26 Comparison of Self-Assessed Performance v Actual Performance – All Classes

		Self assessed quartile			
		1	2	3	4
Actual quartile	1	5	1	0	0
	2	2	3	1	0
	3	2	4	1	0
	4	1	4	1	1

Table 27 Comparison of Self-Assessed Performance v Actual Performance – Class 1

		Self assessed quartile			
		1	2	3	4
Actual quartile	1	5	0	0	0
	2	3	3	0	0
	3	3	2	0	1
	4	1	4	0	0

Table 28 Comparison of Self-Assessed Performance v Actual Performance – Class 2

		Self assessed quartile			
		1	2	3	4
Actual quartile	1	2	4	0	0
	2	0	4	1	0
	3	0	2	3	0
	4	0	5	0	0

Table 29 Comparison of Self-Assessed Performance v Actual Performance – Class 3

The results for individual classes were not substantially different to the aggregate, although the top actual quartile from Class 3 were the group most likely to believe their performance was worse than actual.

Analyzing the responses reveals that decision makers are likely to over-estimate their performance. Of the 35 subjects who were in the bottom 50% of actual outcomes, 28 (80%) self-assessed that they would be in the top 50% of outcomes. The results of these experiments confirm the findings of Kagel and Roth (1995) who identified that subjects tend to be overconfident.

Without feedback regarding the quality of assignments, decision makers are unlikely to recognize the quality of their decisions. The overconfidence that decision makers generally feel may contribute to their resistance towards introducing decision support systems. Therefore, the choice regarding whether or not a decision support system should be introduced should not be left entirely to the existing decision makers; in the absence of feedback, they are likely to believe that their current processes are producing outcomes that are better than actual.

b. Decision-Making Strategies

The subjects were also asked about their decision-making strategies and their perception of which scenarios were more complex. Regarding decision-making strategies, the subjects were asked whether they focused primarily on position requirements, personnel preferences, or attempted to balance the two; for ease of

reference, these will be referred to as Strategy Groups 1, 2 and 3. The percentage of subjects identifying themselves in each of the strategy groups is shown in Table 30.

Response	Response Rate
Strategy Group 1 - I focused primarily on position requirements	22% (18 / 81)
Strategy Group 2 - I focused primarily on personnel preferences	12% (10 / 81)
Strategy Group 3 - I tried to balance the position requirements and personnel preferences	64% (52 / 81)
Other. Please state.	1% (1 / 81)

Table 30 Decision Making Strategies

The majority of subjects were in Strategy Group 3, and this strategy is theoretically what would most likely produce the best possible outcome. However, 34% of subjects focused primarily on either the position requirements or the personnel preferences (Strategy Groups 1 and 2). It is not possible to verify why the subjects would select a strategy that does not balance all the requirements and would be less likely to produce the optimal outcome. However, it is likely that these subjects did so to simplify the process. The works of Simon (1955) and Arthur (1994) show that rational behavior is likely to breakdown in certain circumstances, and it is possible that some, but not all subjects sought strategies that would reduce the information load. Henry (1980) claims that as individuals approach their cognitive limits they will attempt to limit stress by simplifying the problem, and this includes reducing the number of attributes being considered.

The subjects' performance from each strategy group is examined in Table 31 to Table 33. The results from the three groups have been separated to determine if there is any correlation between the decision-making strategy and the overall outcomes.

	Minimum	Median	Maximum
Personnel Points	150.4	244.5	350.4
Position Points	656.6	726.9	796.4
Total Points	806.4	968.5	1138.9

Table 31 Aggregate Points Across all Scenarios – Strategy Group 1 (n = 18)

	Minimum	Median	Maximum
Personnel Points	277.5	346.3	385.8
Position Points	698.0	748.8	802.4
Total Points	982.8	1102.9	1175.1

Table 32 Aggregate Points Across all Scenarios – Strategy Group 2 (n = 10)

	Minimum	Median	Maximum
Personnel Points	183.7	303.9	405.3
Position Points	639.8	753.2	799.8
Total Points	858.8	1057.9	1190.8

Table 33 Aggregate Points Across all Scenarios – Strategy Group 3 (n = 52)

The results from Table 31 to Table 33 show that the subjects in Strategy Group 2, those who focused primarily on personnel preferences, achieved the best median personnel points, although the highest maximum personnel points was achieved within Strategy Group 3. Counter-intuitively, Strategy Group 1, which focused primarily on position requirements, achieved the worst outcomes in terms of position points achieved; in fact, this group had the worst outcomes for personnel, position and total points. It is not clear why the subjects who focused primarily on position requirements did not achieve superior results in terms of position points. The best overall outcome, measured in terms of the maximum total points, was achieved from Strategy Group 3, although the median total points of this group was lower than Strategy Group 2.

Therefore, while it is possible to conclude that Strategy Group 1 had the poorest overall performance, it is not possible to make any conclusive assertions regarding the superiority of outcomes from Strategy Groups 2 or 3.

A limitation of this review is that some subjects may actually have used different decision-making strategies for different scenarios. For example, they may have attempted to balance position requirements and personal preferences in some scenarios (potentially the simpler ones), but used simplifying strategies that focused on only position requirements or personal preferences in other scenarios (potentially the more complex scenarios).

c. Feedback on Complexity

Subjects were asked for their opinion on whether complexity was increased by increasing the personnel to assign, increasing the attributes, or a combination of the two. Specifically, subjects were asked, "How would you describe the difficulty in making the assignments?" Subjects were required to select one of the responses shown in Table 34.

Response	Response Rate
All scenarios were of equal difficulty.	7% (6 / 81)
The scenarios with more personnel were more difficult.	15% (12 / 81)
The scenarios with more attributes were more difficult.	28% (23 / 81)
The combination of more attributes and more personnel made the problem more difficult.	46% (37 / 81)
Other.	4% (3 / 81)

Table 34 Subjects Responses to Question Regarding Scenario Complexity

A literature review of information overload in various disciplines by Eppler and Mengis (2004) showed that, up to a certain point, decision-making performance improved with increases in the amount of information that the decision maker has available, but that beyond that point, the performance decreases. Keller and

Staelin (1987) defined information load in the consumer context as being a factor of the dimensions of quality and quantity of information. In this context, quality may be considered to be similar to the number of attributes, because more attributes provides more information and hence the opportunity for better quality assignments. Quantity of information in this context is measured by the number of personnel to assign. Therefore, using Keller and Staelin's definition of information load as being the combination of quality and quantity of information, 89% of subjects in this experiment agree that increasing information load, through more attributes to consider and / or more personnel to assign, makes the assignment process more difficult.

Two of the three subjects who responded with "Other" provided statements that the assignments were easier with more attributes. Interestingly, one of these subjects achieved the highest total points of all 81 subjects, although the other subject was ranked 44th in terms of total points achieved. Jacoby, Speller and Kohn (1974) found in consumer research that there is an inverted U curve phenomenon, and the ability of subjects to select a detergent that best satisfied their needs was poorest at the low and high levels of total information, while subjects provided with intermediate levels of information made the best quality decisions. There is the potential, particularly in the case of the subject who achieved the highest total points, that this subject was operating at a point on the U curve where he / she did not experience any information overload effects, while other subjects were over-whelmed by the amount of information. Certainly different people will have different strategies and different cognitive abilities when it comes to assignment decision making. With regards to strategies, Houser et al. (2004) classified subjects in decision-making experiments into three types, the "near-rational," the "fatalist," and the "confused." With regards to cognitive abilities, Henry (1980) developed a linear regression equation of the relationship between information processing accuracy (dependent variable) and the independent variables of complexity and individual processing ability. Henry found that the coefficient for information processing ability was positive and statistically significant. Therefore, it is logical to expect that different people will operate at different places on the U curve for different information loads.

An additional but indirect means of assessing subjects' thoughts on complexity is obtained by examining their self-assessment of performance compared to an "optimal" outcome. Included in the questions that the subjects were asked was the following:

> "For the experiments that you have participated in, there is an optimal outcome in accordance with the points scheme described. In the boxes below, please indicate how close you believe you came to the optimal assignment outcomes? Use the following scale
>
> 1 = Within 0 – 20% of the optimal
>
> 2 = Within 20 – 40% of the optimal
>
> 3 = Within 40 – 60% of the optimal
>
> 4 = Below 40% of the optimal."

The subjects were asked to indicate their performance for each of the scenarios in accordance with the assessment outlined above. The results show that as the scenarios became more complex the subjects expressed decreasing confidence in the ability to achieve results close to the optimal outcome. The average assessment from all subjects is shown against the scenarios in Figure 30.

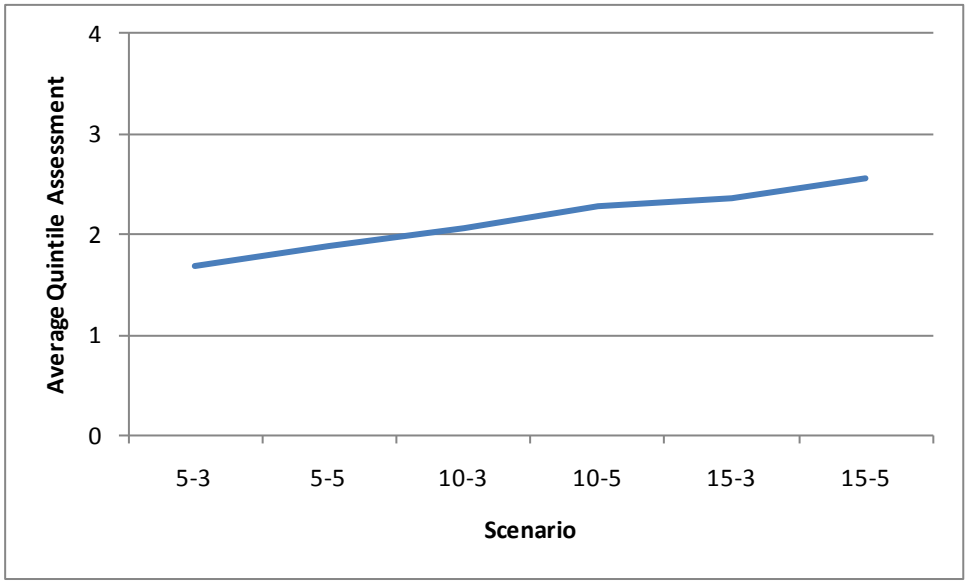

Figure 30. Self Assessment of Performance (Average Quintile) Compared to Optimal Outcome for Each Scenario

Figure 30 reveals that the subjects were more confident of performing closer to an optimal outcome in the less complex scenarios than they were for the more complex scenarios. This lends support to the hypothesis that assigning more personnel using more attributes increases complexity.

C. SUMMARY

This research has examined decision making in the context of assigning military personnel to positions. The effects of information load have been examined in other contexts; Eppler and Mengis (2004) reviewed 97 articles across a variety of disciplines. They found that the performance of a decision maker "correlates positively with the amount of information he or she receives—up to a certain point. With additional information beyond this point, the individual's performance will rapidly decline" (p. 326). However, previous studies have generally involved situations where decision makers have been required to select the best outcome(s) from a list of options / alternatives. The assignment scenarios examined in this research are potentially more complex because the decision makers are required to assign members of one list (personnel) to members of another list (positions), and the members of each list have

preferences / requirements. Further, each assignment made by the decision maker limits the range of remaining choices available. While the decision makers can change assignments within a scenario if they see the opportunity, it is unclear whether they would recognize these opportunities, particularly when under time pressure.

This research has not only examined how decision makers may perform in situations where military personnel are assigned to positions, it elicited information from the subjects in relation to their decision-making strategies, their views on the complexity of the scenarios, and a self-assessment of their performance.

The variation in results achieved by the subjects was high, even for scenarios where subjects had the benefit of experience from a number of previous scenarios. This variation demonstrates that the quality of assignments depends heavily on the decision makers abilities. Such variability was observed in these experiments even though the subjects were all students studying at the master's level. In situations where the decision makers may be less homogenous and have a greater range of abilities, even greater variability in the results could be expected.

Compared to the experiments conducted here, the equivalent real life assignment decisions have the potential to affect careers of military personnel, and can in turn have an impact on personnel retention. Given the high stakes involved, the military should expect better outcomes and improved consistency, rather than allowing decisions to depend on the abilities of the person making the decisions. In the absence of applying selection criteria to identify and screen decision makers, and in the absence of decision support, there is a high risk that assignment decision quality will suffer. This risk could be reduced by carefully selecting and screening the decision makers, although this would not guarantee the results. Only introducing a decision support system would eliminate this variability.

In accordance with the experiment design, and in accordance with real life military assignment situations, the subjects provided greater focus on the positions. This was demonstrated by higher points being achieved for the positions than the personnel. What was observed was that, compared to the personnel points, the position points

showed less variability and were closer to the optimal outcomes. In many cases, improved personnel outcomes could be achieved with no negative impact on position outcomes.

When comparing their performance to that of others, the subjects showed that they were likely to overestimate their performance. Extending this over-confidence to real life assignment situations, the decision makers may be less inclined to recognize the need for decision support if they overestimate their performance.

The experiments conducted here affirm the need for decision support in military personnel—position assignment scenarios. Two-sided matching has desirable properties for application to such a decision support system; in particular, it is much less prone to gaming than other techniques that are available. However, two-sided matching has not been applied to assignment processes in hierarchical organizations. Further, there is likely to be significant preference indifference, and the effect of indifference on two-sided matching outcomes has not been analyzed. The next chapter will examine how preference list indifference will affect two-sided matching assignments.

V. THE EFFECT OF PREFERENCE LIST INDIFFERENCE ON TWO-SIDED MATCHING OUTCOMES

Chapter I highlighted a variety of matching algorithms for assigning personnel to positions. The approach selected for this research is two-sided matching. This approach was selected because of two unique characteristics: firstly, under certain conditions, it is a dominant behavior for participants on at least one side of the market to truthfully reveal their preferences, precluding gaming behavior and leading to more accurate assignments; secondly, the assignments produced will be stable.

Two-sided matching originated in medical residency matching markets (Roth & Sotomayor, 1990), but more recently has been applied to matching students to high schools (Abdulkadiroğlu et al., 2005, 2006). In the medical residency markets, hospitals typically use interviews to determine which graduating medical students are preferable and will be placed on the hospital preference lists. However, in the much larger school matching markets, such as the New York City high schools, where in excess of 90,000 students are being matched to schools, it is not feasible for the schools to interview and rank each student that applies to the school. As a result, Abdulkadiroğlu et al. (2005) identify considerations, such as GPA, behavior and attendance records, that may be considered. With such a large number of students, limited attributes, and discrete attributes with limited degrees of gradation, this must inevitably lead to indifference in the schools' preference lists. However, preference lists are treated as being strictly ordered based on a random lottery process; Abdulkadiroğlu et al. (2009) explain that every student is given a lottery number to break ties. In the New York City High School Match, Erdil and Ergin (2007) found that 10.5% of the matched students could have been matched with schools higher on their preference lists without hurting others, after incorporating the considerable indifference school preferences.

Due to the large number of personnel involved in a military context, it is likely that preference lists would need to be determined based on rules or attributes, just as they are for schools. For the same reason that indifference can be expected in school preference lists, the preference lists for military positions can be expected to exhibit

indifference; military assignment contexts involve a large number of personnel and positions (Liang & Buclatin, 1988), there are a limited number of attributes by which the positions can assess personnel (Klingman & Phillips, 1984), and the attributes have a limited number of levels (for example, rank levels or the possession of a skill set). Therefore, the position preference lists are quite likely to exhibit indifference. The personnel preferences are also likely to have indifference because there will be groups of positions, which are almost, if not exactly, identical. Where indifference exists, existing research and applications state that preference ties should be randomly broken, with the resultant preferences treated as being strictly ordered. Roth and Sotomayor (1990) advocate a fixed tie breaking rule and existing two-sided matching applications typically require participants to submit strict preferences. However, breaking the ties in different ways will have an impact on the results. In particular, different tie breakings will lead to different outcomes in terms of the number of agents assigned, and the average rank of assigned agents will be different; this was demonstrated for the NYC High School Match by Abdulkadiroğlu et al. (2009).

To examine how different tie breaking mechanisms will affect the assignment results, a random set of person and position agents will be created and given "quality" values. These agents could be considered students and schools, residents and hospitals or any other set of agents who have preferences and are assigned to one another, however for consistency with the remainder of this research, they will be referred to as personnel and positions. Based on homophilous attraction, the person and position agents have preferences for each other, although the strength of this attraction is defined by the "Preference Correlation" parameter (*C*). A low *C* value indicates that position agents will only identify the person agents who are the closest in quality values on their preference lists (and vice versa for person agents ranking positions). As the value of *C* increases, the position agents are increasingly willing to accept personnel whose quality values are less like their own. After determining an initial strict ordering of each agent's preference lists, based on quality values and the preference correlation parameter (*C*), the "Degree of Indifference" parameter (*DOI*) specifies the extent to which ties will exist within the personnel and position agents' preferences.

The second research question that was outlined in Chapter I was, "What effect does preference list indifference have on a two-sided matching outcome, and how can a decision maker use this to produce better outcomes?"

In an experimental study of matching in situations with incomplete preferences and preference ties, Gent and Prosser (2002) generate random instances defined by three variables; n is the number of men and women, p_1 is the probability of incompleteness and p_2 is the probability of ties. In the experimental work undertaken here, four variables are used to generically define an instance; these are the number of agents involved (n), the length of their preference lists (L), the correlation of the attraction between the sets of agents (C) and the Degree of Indifference (DOI), which identifies the probability of preference list ties. The agents who are included in the preference lists are determined by homophilous attraction and the preference correlation (C), such that a position's preference list will only include those persons whose quality score is within the range +/- C of the position's own quality score. For assignment scenarios defined by these four parameters, multiple assignment trials will be conducted, with tied preferences randomly reordered at each trial. The assignments from each trial are captured in the database and subsequently analyzed to determine the effect of the four parameters on two key variables of interest; the number of agents matched (noting that two-sided matching does not guarantee all agents will be matched if preference lists are not complete), and the average rank of matched agents. The outcomes of this research will demonstrate, for a variety of scenarios defined by the four parameters, the range of outcomes that are possible in terms of number of agents assigned and the average rank of assigned agents. This research will help decision makers understand how assignment outcomes will vary as indifference increases.

Table 35 summarizes the parameters and variables of interest in these experiments.

Parameter / Variable	Description
n	Represents the number of agents involved in the scenario.
L	Represents the maximum length of the agents preference lists.
C	Represents the extent to which agents preferences are correlated. A low *C* value means that position agents only identify those person agents who are most alike them (a *C* value of zero would indicate that a position will only accept a person who has exactly the same quality score). As the value of *C* increases the position agents are increasingly willing to accept person agents who are less alike.
DOI	Represents the extent to which indifference exists in agents' preference lists. A *DOI* of zero indicates strictly ordered preference lists, while a *DOI* of one indicates complete indifference. *DOI* can be set at any value in between zero and one.
Number of Agents Assigned	The number of agents who are assigned from a trial.
Assignment Quality	The average preference rank of agents who are assigned at a trial.

Table 35 Parameters and Variables of Interest

In these experiments, the key focus is on determining how indifference, measured by *DOI*, influences assignment outcomes; in particular, the number of agents assigned and the assignment quality. The other three parameters, *n*, *L* and *C*, specify different scenarios in which two-sided matching may be used to make assignments.

This chapter examines the impact of preference list indifference on two-sided matching results in a context that is not specific to any particular application. The intent is to demonstrate how preference list indifference affects a generalized assignment

situation. The next chapter will extend the work and examine how two-sided matching can be applied to a specific situation; a military assignment context that includes preference list indifference.

This research will define a variety of scenarios at different settings of the four descriptive parameters. Before doing so, however, the next section examines the outcomes from five scenarios that are defined identically in terms of the parameters. Even though the parameters are defined identically, there is random variation resulting from random quality scores and random preference ties (even though *DOI* is held constant, variability in preference ties occurs because a preference is considered indifferent to the previous preference if a random number, which is unique for each agent and preference, is less than the *DOI*).

A. THE EFFECTS OF PREFERENCE LIST INDIFFERENCE

The following definitions are used in this chapter:

Scenario. A scenario refers to the generation of a set of agents and their preferences, using the parameters *n*, *L*, *C* and *DOI*. Due to the random generation of agents' quality scores and their tied preferences, the agents' preferences will differ even if scenarios are defined by identical parameter values.

Trial. A trial refers to the assignments generated from a single random tie-breaking of preferences. Multiple trials are conducted for each scenario to determine the results from breaking preference ties in different ways.

For the next section, 1,000 trials were conducted for five identically defined scenarios. The results are used to provide some initial insights into the different outcomes possible when preference list ties are randomly broken in different ways.

Five scenarios were created using the same parameter settings (n = 500, L = 10, C = 0.3 and *DOI* = 0.3), and for each of these scenarios, 1,000 assignment trials were conducted. The number of agents matched at each trial and the average rank of assigned personnel and position agents were recorded. The frequency and range of the variables of

interest is examined. The results from this section will help identify how much variability in results is related to the parameters and how much is attributable to random variation.

1. Number of Agents Assigned

A different assignment outcome is possible in each of the 1,000 assignment trials. Each assignment outcome can assign a different number of agents, and the average rank of assigned agents can vary. For the five identically defined scenarios (n = 500, L = 10, C = 0.3 and DOI = 0.3), the number of assigned agents at each of the 1,000 trials was recorded. Figure 31 shows the results for the 1,000 trials of the five scenarios; each line represents the outcomes from a scenario with the x-axis showing the number of agents assigned, and the y-axis showing the frequency with which the numbers of agents assigned was recorded from the 1,000 trials.

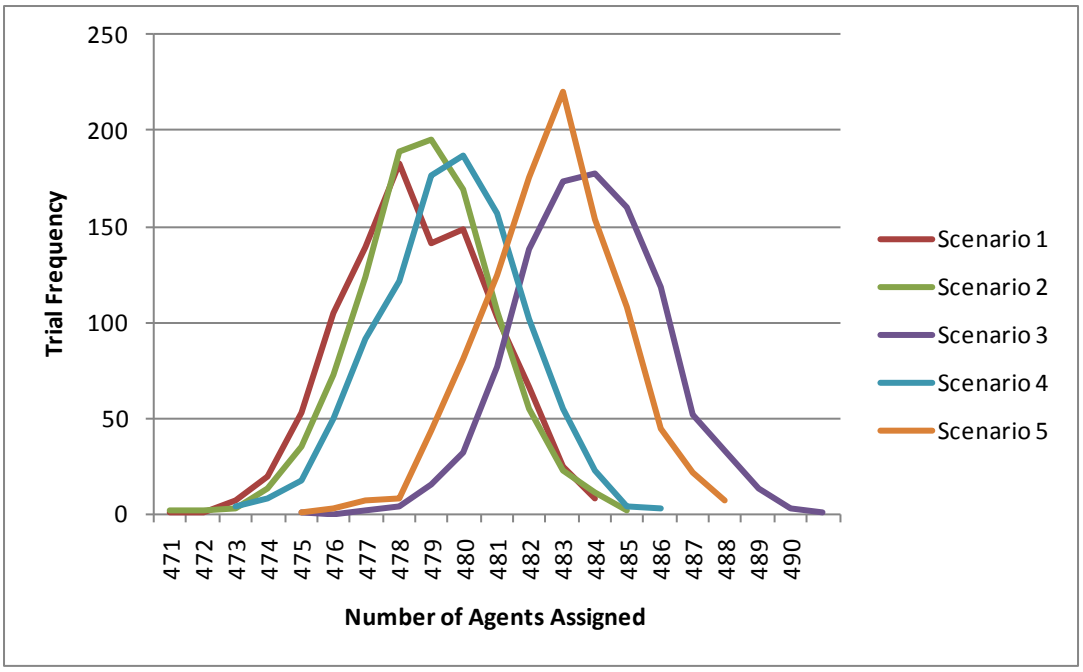

Figure 31. Variability within Identically Defined Scenarios – Number of Agents Assigned

The results in Figure 31 show a distribution of outcomes that resembles a normal distribution. The number of agents assigned varies across a range, with outcomes at the low and high end of the range occurring infrequently, and outcomes near the middle of the range occurring most frequently.

In Scenario 3, depending upon which random tie breaking was used, the number of assigned agents varied between 474 and 490. In this scenario, only one of the 1,000 trials resulted in 490 agents being assigned. Therefore, had a single random tie breaking rule been applied, it is very unlikely that it would have assigned 490 agents, and in fact, based on the 1,000 trials conducted, there is a 99.9% chance that a single random tie breaking would have produced a result that had fewer than 490 agents assigned. Compared to the standard approach in existing two-sided matching applications, where agents are assigned after preference list ties are randomly broken, examining alternate tie breaking can clearly produce outcomes that assign more agents.

One thousand trials does not guarantee that the highest number of stable assignments for Scenario 3 is 490; it is possible that the number could be higher, although based on the probability distribution it is unlikely to be substantially higher, particularly given that the absolute upper bound is 500 agents assigned (the scenario involved 500 personnel and position agents). A later section will examine how many trials are necessary to confidently obtain a "good" outcome; that is, is it necessary to undertake substantially more than 1,000 trials, or is it possible to obtain a "good" outcome with less than 1,000 trials.

While Scenario 3 had the highest range in the number of agents assigned, the other scenarios were not dissimilar. As shown in Figure 31, the minimum number of assigned agents from the five scenarios varied between 471 and 474, representing a range of 0.6% of the number of agents, while the maximum number of assigned agents varied between 483 and 490, a range of 1.4% of the number of agents. For later sections, to attribute the difference in results to the variables of interest, the variability in outcomes should be greater than that exhibited here for identically defined scenarios.

2. Average Preference Rank of Assigned Agents

At each of the 1,000 trials for a scenario, the preference rank at which each personnel and position agent is assigned is recorded, and these are averaged across all assigned personnel agents and position agents for each trial. As a result, in addition to determining the number of agents assigned (as outlined in the last section), it is possible to determine the average rank of both personnel agents and position agents for each of the 1,000 trials. It is preferable, from an overall perspective, to find a solution that assigns the highest number of agents at the lowest preference ranking.

Figure 32 shows the average personnel and position rank from each of the 1,000 trials conducted for Scenario 3. The average rank for each agent in each trial is summed to provide a combined person and position rank for each trial. The combined personnel and position ranks are shown on the y-axis against the number of agents assigned on the x-axis. Dominant outcomes are those that have more agents assigned at lower preference ranks. The tradeoff between assigning an additional pair of agents but at an increased average preference rank is an issue of the decision maker in the relevant assignment context. For example, Figure 32 shows that it is possible to assign 489 agents at a combined personnel and position preference ranking of 5.538; alternatively it is possible to assign 490 agents at an average preference ranking of 5.557.

To place the average preference rank (the y-axis in Figure 32) in context, for two trials that have a difference of 0.1 in the average preference rank, the trial with the average preference rank that is 0.1 lower has approximately 49 person or position agents (approximately 10% of the total) assigned to partners that are their next preferred partner; for example, 49 person or position agents being assigned to their second ranked preference instead of their third ranked preference.

A decision maker would logically select from any of the nondominated outcomes shown in Figure 32. At the extremes, this would involve selecting between an outcome that has 474 agents assigned at a combined average preference rank of 5.325 and an outcome that has 490 agents assigned at a combined average preference rank of 5.557.

The “price” of assigning an additional 16 agents is that approximately 114 agents would be assigned to a partner that they rank one lower.

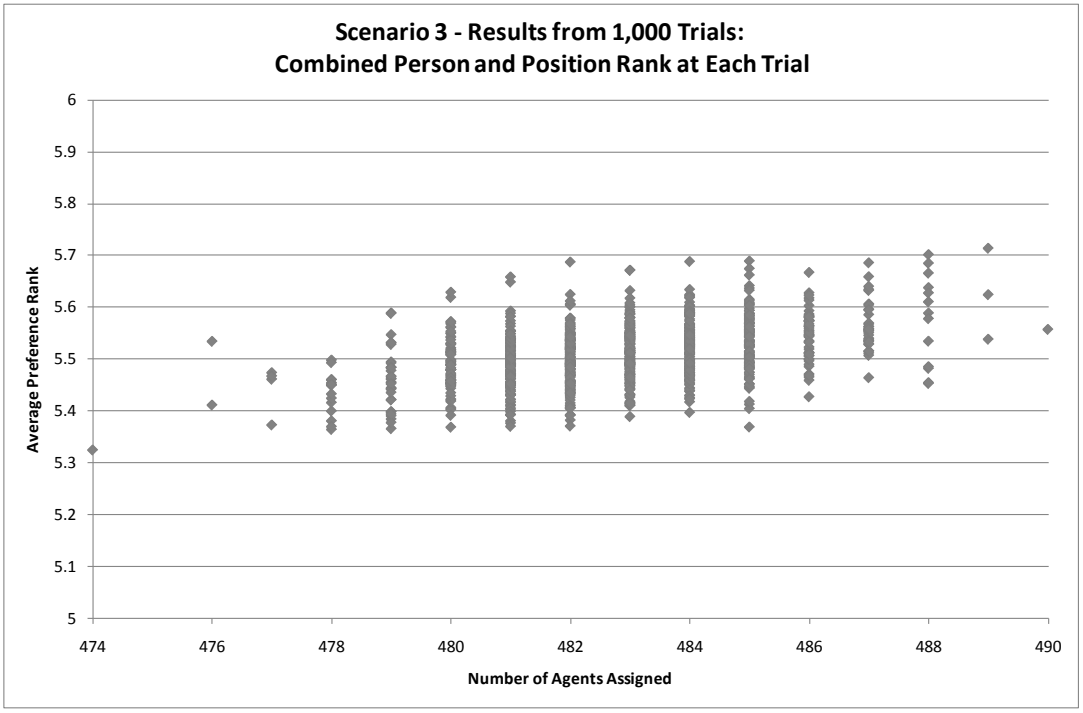

Figure 32. Preference Ranking of Assigned Agents – 1,000 Trials of Scenario 3

From the many random tie breakings and the associated assignments, there are some outcomes where the average ranking of assigned agents is such that the assignments favor personnel, some favor the positions and some provide a more equitable outcome between personnel and positions. To examine whether identically defined scenarios will produce similar results from the perspective of assignment quality, the best personnel and position preference ranks from the 1,000 trials for each scenario are presented in Figure 33.

Figure 33. Best Average Person and Position Preference Ranking for Each Scenario

For both personnel and positions, the best average preference ranking was achieved in Scenario 4 (average person preference rank = 2.121, average position preference rank = 2.650). The scenario where the best average preference ranking was highest was Scenario 3 (average person preference rank = 2.200, average position preference rank = 2.882). The range between the best Scenario 3 and 4 personnel preference ranks is 3.7% of the Scenario 4 average person preference rank, while the corresponding range for position preferences is 8.8%. Therefore, even for scenarios that are defined identically in terms of the four parameters, there is variability in the outcomes due to random differences in the quality scores and preference ties. In the later section, which examines the effect of varying *n*, *L* and *C*, controls are used to minimize the variability in the scenario specification.

3. Determining the Required Number of Trials

It was demonstrated in the previous section, which involved 1,000 trials for each scenario, that the number of agents assigned resembles a normal distribution. By assuming 1,000 trials, it is statistically likely that significantly better outcomes will be achieved, in terms of number of agents assigned and quality of assignments, compared to the outcome if only a single random ordering of preferences was used. That is, it is

extremely unlikely that the first random tie breaking would deliver the highest possible number of agents assigned. Furthermore, a single trial would only present a decision maker with a single outcome. Multiple trials will provide the decision maker with options, and allow tradeoff between an outcome that assigns more agents and an outcome with fewer agents assigned but at a better preference ranking.

A new scenario was generated using the same base-line parameters as used in the previous section (n = 500, L = 10, C = 0.3 and DOI = 0.3). 25,000 trials were conducted to estimate the number of trials necessary to place suitable bounds on the range of agents assigned and assignment quality. The goal of this research is not to determine the absolute best possible outcome; if that was the case then alternative approaches would be used. Rather, the intent is to demonstrate the effect of preference list indifference and the extent to which different outcomes are possible. By understanding this it will be possible to use preference indifference to an advantage rather than conducting a single trial that randomly breaks preference list ties.

Figure 34 shows the number of agents assigned from the first 1,000, 5,000, 10,000 and then all 25,000 trials of the base-line scenario. The frequency distribution once again resembles a normal distribution, with 481 assigned agents being the modal outcome. The format of Figure 34 is modified in Figure 35, so that the y-axis shows the percentage of trials (from the first 1,000, 5,000, 10,000 and 25,000 trials of the scenario) that have the number of assigned agents indicated on the x-axis.

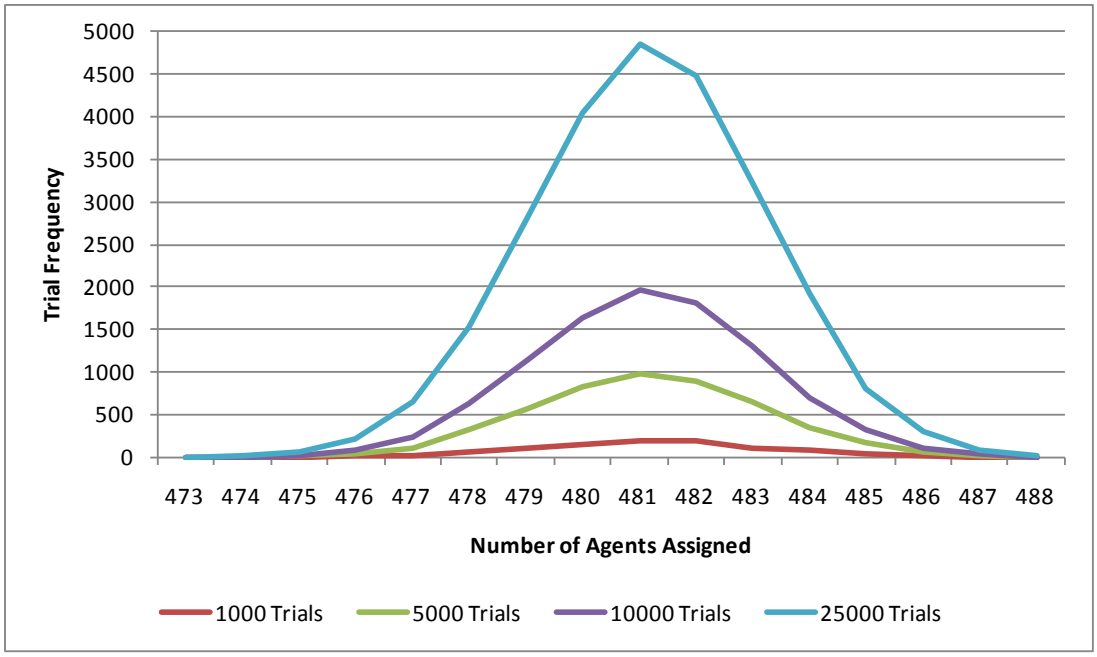

Figure 34. Frequency of Trials Showing Number of Agents Assigned

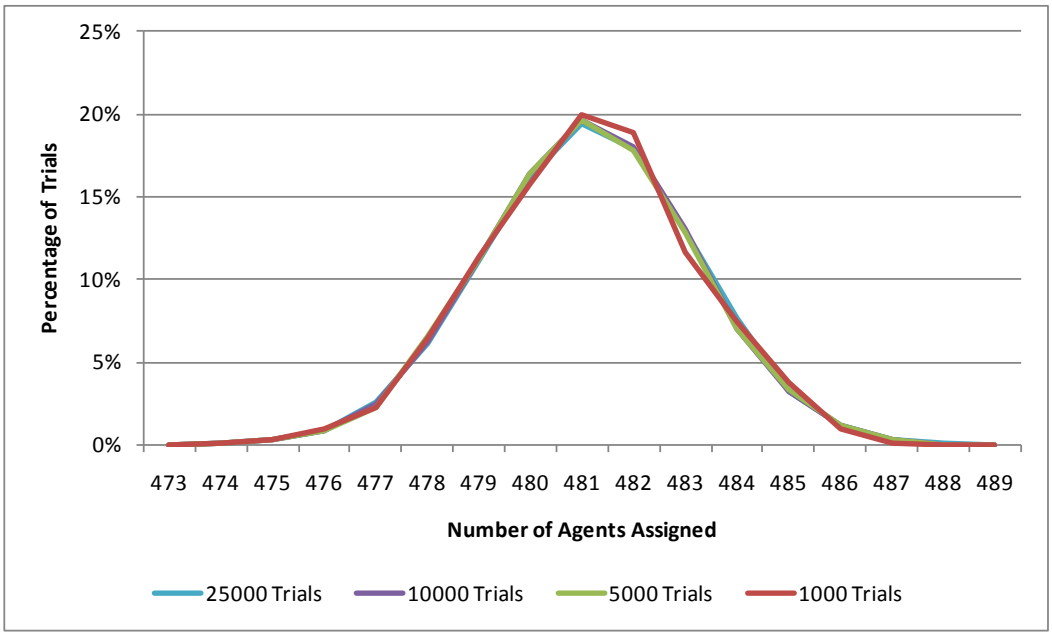

Figure 35. Percentage of Trials Showing Number of Agents Assigned

In Figure 35, there is very little difference between the distribution achieved in the first 1,000 trials compared to 5,000, 10,000, or 25,000 trials. Therefore, 1,000 trials is considered sufficient to demonstrate the effect of preference list indifference and the extent to which different outcomes are possible.

Table 36 shows the outcomes for the first 1,000 trials and the full 25,000 trials, with the frequency of each outcome shown together with the probability of achieving each outcome. As shown, increasing from 1,000 to 25,000 trials increases the range of assigned agents by three (reducing the lower end of the distribution by one and increasing the upper end of the distribution by two).

	First 1,000 Trials		25,000 Trials	
Number of Agents Assigned	Frequency	Probability	Frequency	Probability
473	0	0.000%	3	0.012%
474	1	0.100%	13	0.052%
475	3	0.300%	63	0.252%
476	10	1.000%	224	0.896%
477	22	2.200%	643	2.572%
478	65	6.500%	1528	6.112%
479	114	11.400%	2792	11.168%
480	157	15.700%	4039	16.156%
481	200	20.000%	4859	19.436%
482	189	18.900%	4480	17.920%
483	116	11.600%	3238	12.952%
484	74	7.400%	1921	7.684%
485	38	3.800%	805	3.220%
486	10	1.000%	305	1.220%
487	1	0.100%	75	0.300%
488	0	0.000%	11	0.044%
489	0	0.000%	1	0.004%

Table 36 Results from 25,000 Trials – Frequency of Number of Agents Assigned

Figure 36 shows how the best preference ranking at each number of assigned agents for the first 1,000 trials compared to the full 25,000 trials. While better assignment quality (improved preference ranks for assigned agents) was found after 25,000 trials, the pattern from first 1,000 trials is similar to that found after 25,000 trials.

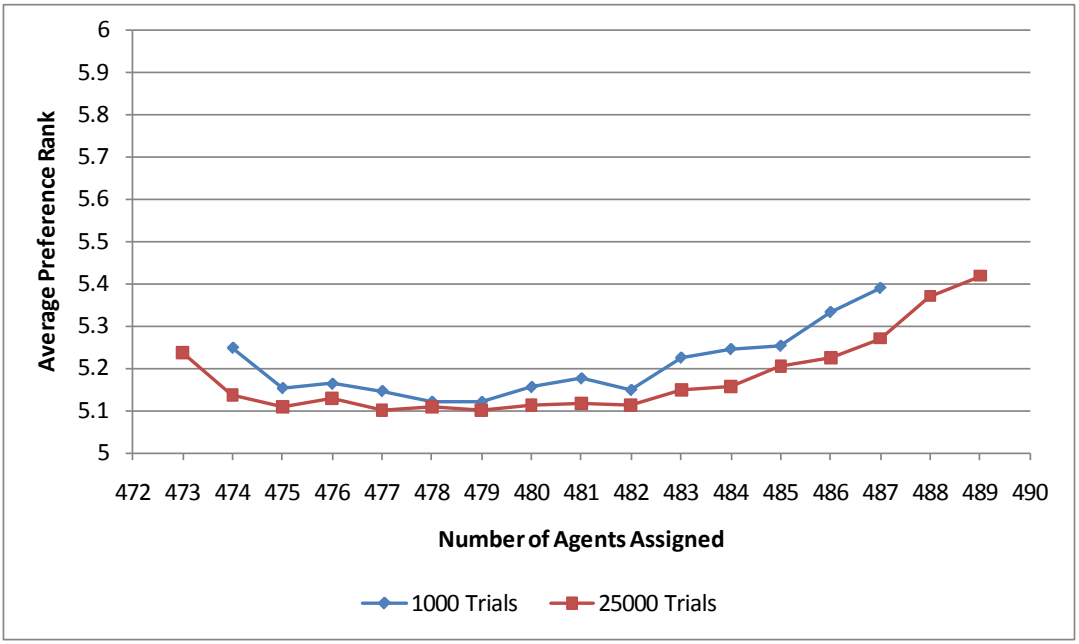

Figure 36. Best Preference Rankings of Assigned Agents – 1,000 Trials v 25,000 Trials

Comparing the number of assigned agents and the assignment quality between 1,000 trials and 25,000 trials demonstrates that sufficiently similar outcome patterns are present in the first 1,000 trials. As a result, subsequent analysis will be based on 1,000 trials for each scenario. While this will not show the best outcomes possible, it will show the pattern of outcomes that occur under different indifference scenarios with sufficient fidelity.

B. THE EFFECT OF PREFERENCE LIST INDIFFERENCE ON DIFFERENT TYPES OF SCENARIOS

This section examines how preference list indifference interacts with the other parameters to influence the assignment variables of interest; the number of agents assigned and the average preference rank. The effect of *DOI* and, in turn, the parameters *n*, *L* and *C* are examined.

1. Preference List Indifference and Number of Agents

To examine the effect of preference list indifference on scenarios involving differing numbers of agents, scenarios were created where the parameters for the number

of agents (*n*) and degree of indifference (*DOI*) were varied, holding other parameters constant. For these scenarios, preference list length (*L*) was held constant at 10 preferences and the preference correlation (*C*) was held constant at 0.25. The exact level at which these parameters were set is not critical, however, they were set to a level that is "mid-range" compared to the range over which they vary in later sections of this research. Twelve scenarios were created; these involved 100, 500, 1,000 and 2,000 agents to be assigned; at each *n*, the degree of indifference (*DOI*) in the preference lists was set at three levels; 0.1, 0.3 and 0.5. There were 1,000 trials (random ordering of tied preferences) conducted for each scenario.

If a new and completely unique set of agents and preferences was created for every scenario, the random variation in agents and their preferences would potentially be too great to draw any conclusions about the relationship between parameters and outcomes. This would jeopardize the experiment's internal validity. Therefore, for each setting of parameter *n*, a set of agents was created and this set of agents remained constant across each setting of parameter *DOI*. Where each set of agents is created, a "reference set" of strictly ordered preferences is created for each agent. This "reference set" of strictly ordered preference was not used during the assignment process, but rather was used solely as the starting point for determining each agent's tied preferences. Using the "reference set" of strictly ordered preferences as a starting point, tied preferences were determined based on the *DOI* settings; a macro cycled through the "reference set" of preferences of each agent and for each preference (except the first one) a random number was generated; if the random number was less than the *DOI* setting then the preference was considered tied with the previous preference in the "reference set." Once the assignment process was underway, the tied preferences were randomly broken to produce a strict set of preferences for each assignment set. The use of a single set of agents and a constant "reference set" of preferences enhances the internal validity of the experiments and ensures that, at each scenario size, the variability relates only to the preference list ties and is not related to other factors, such as the distribution of agent quality scores or which agents are included in the preference lists.

The results show that as the *DOI* increases, so too does the range of the number of observed assignments. At one extreme, there is a single outcome if there is no indifference; as preference list indifference increases, then so too does the maximum number of possible outcomes.

Figure 37 to Figure 40 show the frequency distribution of the number of agents assigned for each scenario size. In each figure, a separate line identifies the results for a different *DOI*. The frequency distributions resemble normal distribution curves; statistically, any single random breaking of preference list ties is likely to produce an outcome where the number of agents assigned is close to the mean number of assignments and is unlikely to produce an outcome where the number of agents assigned is far from the mean. As the *DOI* increases, the peak of the distribution decreases and the range in the number of assignments increases.

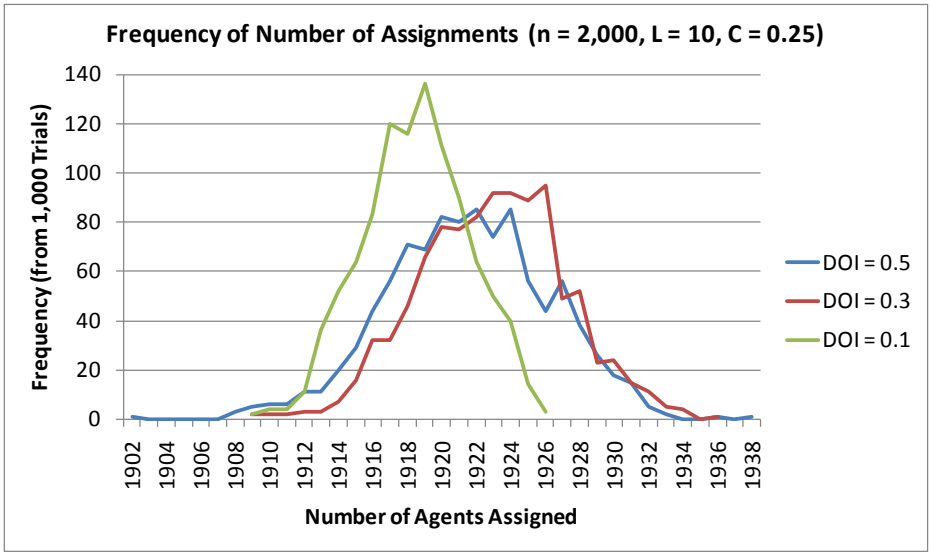

Figure 37. Frequency of Number of Assignments – n = 2,000

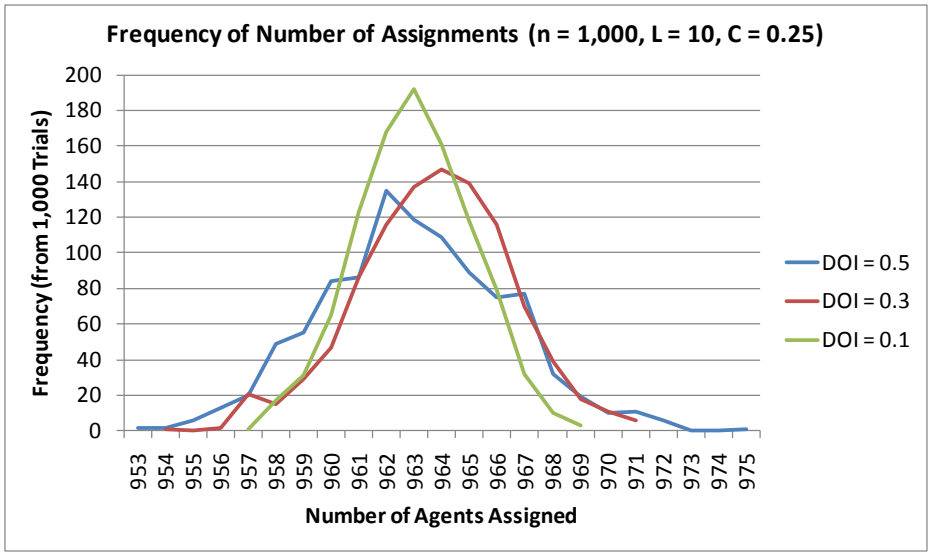

Figure 38. Frequency of Number of Assignments – n = 1,000

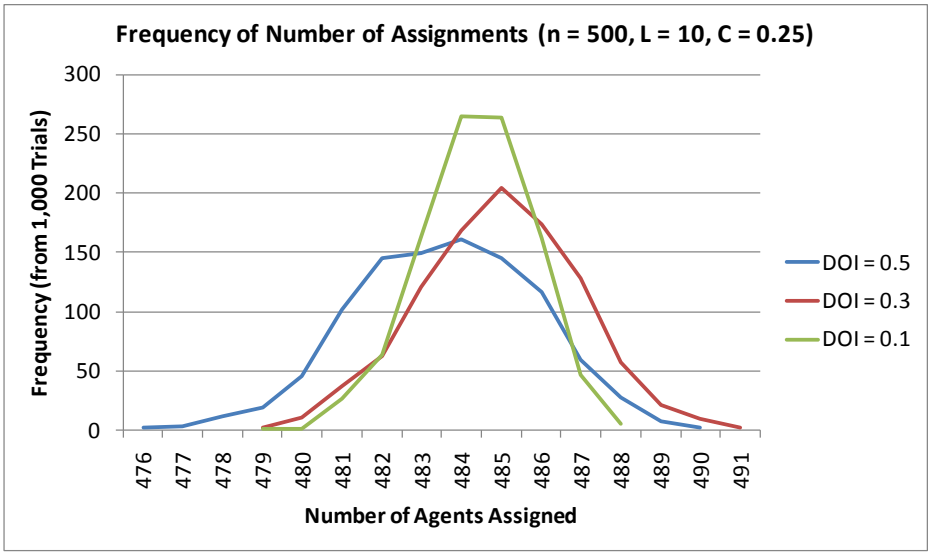

Figure 39. Frequency of Number of Assignments – n = 500

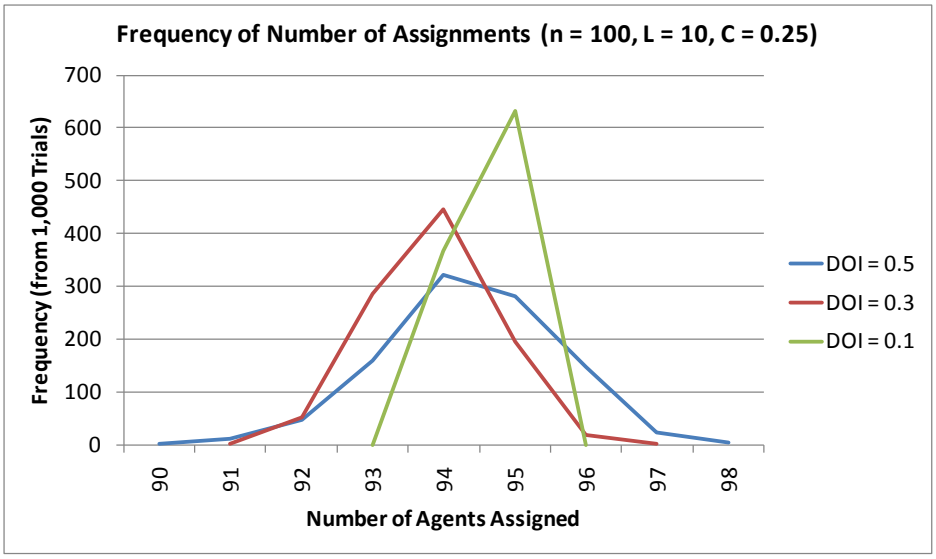

Figure 40. Frequency of Number of Assignments – n = 100

Presenting the same results in a different format, Figure 41 shows that, holding *DOI* constant, the range in the number of agents who are assigned increases as the number of agents in the scenario increases, where range is the difference between the minimum and maximum observed assignments. Alternatively, holding the number of agents constant, increasing the *DOI* increases the range in the number of agents assigned. In Figure 42, the range in the number of agents assigned is expressed as a percentage of the number of agents in the scenario. Although the range in the number of agents assigned increases as the size of the scenario increases (as shown in Figure 41), as a percentage of the number of agents involved the range decreases as the number of agents increases. For example, from 1,000 trials of the scenario involving a *DOI* of 0.3 and with *n* of 500, the number of agents assigned ranged from 479 to 491 (seen in Figure 39) a spread of 2.6% of the agents involved. For the same *DOI* but with 2,000 agents, the agents assigned ranged from 1,909 to 1,936 (seen in Figure 37), a spread of 1.4% of the number of agents involved.

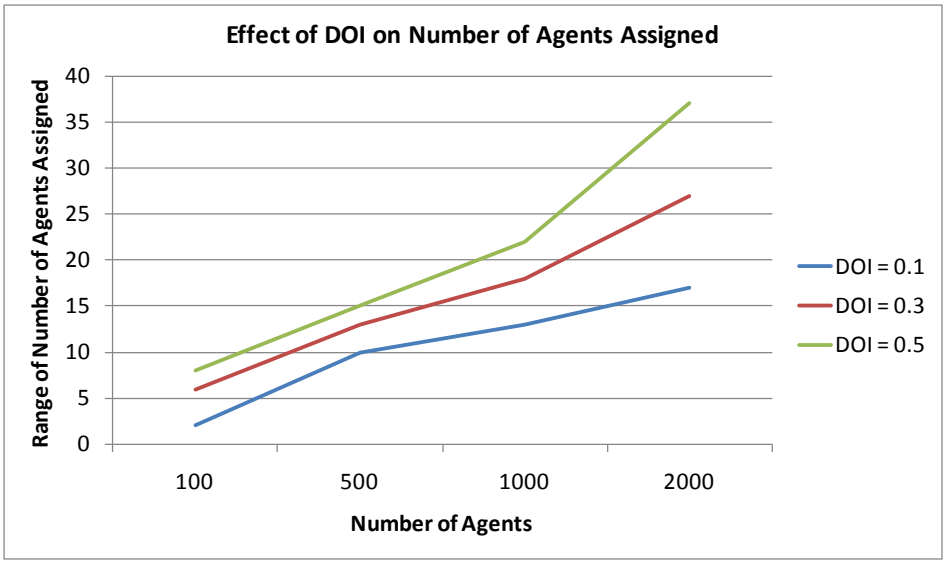

Figure 41. Effect of DOI on Range of Number Agents Assigned

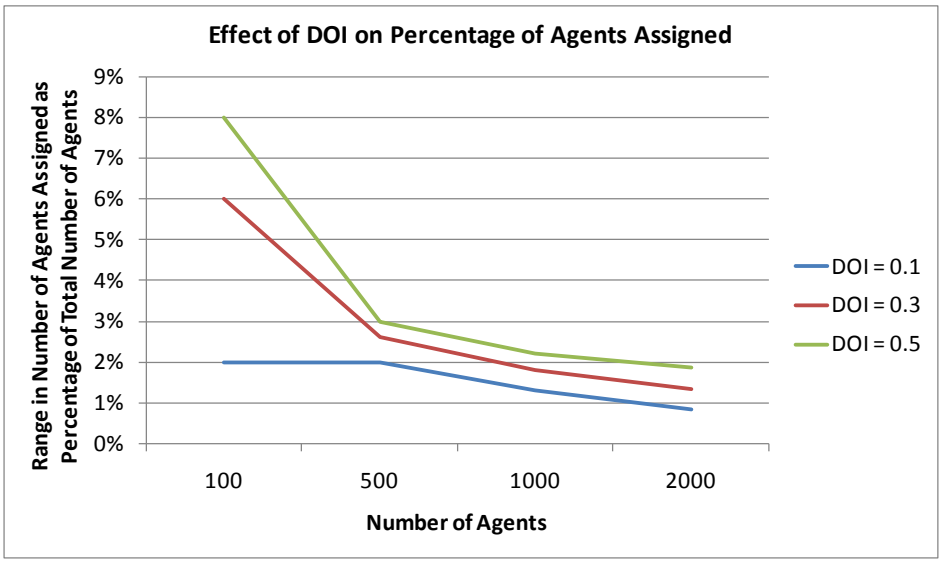

Figure 42. Effect of DOI on Range of Agents Assigned as Percentage of Number of Agents

2. Preference List Indifference and Preference List Length

To examine the effect of restrictions on preference list length and how it is affected by preference list indifference, scenarios were created varying the parameters for preference list length (*L*) and degree of indifference (*DOI*). For these scenarios, the number of agents (*n*) was held constant at 1,000 personnel and position agents and the

preference correlation (*C*) was held constant at 0.5. The preference correlation factor was set at a value high enough to ensure sufficient personnel agents within a quality distance of 0.5 from each position (and vice versa). If *C* was set too low, the number of agents within a range of ±*C* of an agent may be less than the parameter setting *L* in the scenario, and the results would be related to the variable *C* and not the variable of interest, *L*. Scenarios were created with preference list lengths (*L*) of 5, 10 and 20. For each setting of parameter *L*, the degree of indifference (*DOI*) in the preference lists was set at three levels; 0.1, 0.3 and 0.5. With three settings of the variable *L* and three settings of the variable *DOI*, there were nine scenarios developed and 1,000 trials (random ordering of tied preferences) were conducted for each scenario.

If a new and unique set of agents and preferences was created for every scenario, the random variation in agents and their quality scores would potentially confuse the experiments' internal validity and preclude useful conclusions about the relationship between the parameters and outcomes. Given that all nine scenarios involved the same number of agents, a single set of 1,000 agents was created, with each agent's quality score remaining unchanged for each scenario. This ensured that the quality distance between any person and position was the same for each scenario. Therefore, each position ranked the closest *L* personnel agents using the same initial strict ordering and differences in the tied preference lists were related solely to the setting of the variables *L* and *DOI*.

If there is an equal number of personnel and position agents and preference lists are complete, that is, every position agent identifies every person agent in their preference lists and vice versa, then there is a stable matching where every agent is assigned (Gale & Shapley, 1962). However, when preference lists are incomplete, not all agents are assigned. As preference list length decreases, the number of agents assigned will also decrease, unless there is sufficient individuality of preferences (for example, if every position agent listed a single and unique person and that person listed the same position, then it would be possible for all agents to be assigned with incomplete preference list lengths, however such a situation is unlikely).

Figure 43 shows the results from all nine scenarios used to investigate the effect of preference lists length and *DOI*. Each line in the figure represents the outcomes from 1,000 trials of the scenario and shows the number of times in the 1,000 trials that a given number of agents were assigned.

The first observation on Figure 43 is the relationship between preference list length and the number of agents assigned. Regardless of the *DOI*, the number of agents assigned (whether measured by mean or mode) was lowest for $L = 5$ and highest for $L = 20$, with the outcomes for $L = 10$ lying between. However, as seen later, the tradeoff for achieving a higher number of assigned agents is a reduction in the assignment quality, as measured by the average preference rank of the assigned agents.

The second observation is that the spread of outcomes increases as *DOI* increases. That is, the range in the number of assigned agents increases regardless of the preference list length. This is seen further in Figure 44. At a *DOI* of 0.1, the range in the number of agents assigned varies between 9 and 11, whereas the range in the number of agents assigned varies between 15 and 25 at a *DOI* of 0.5. Holding *DOI* constant, greater variability in outcomes was observed with shorter preference list length.

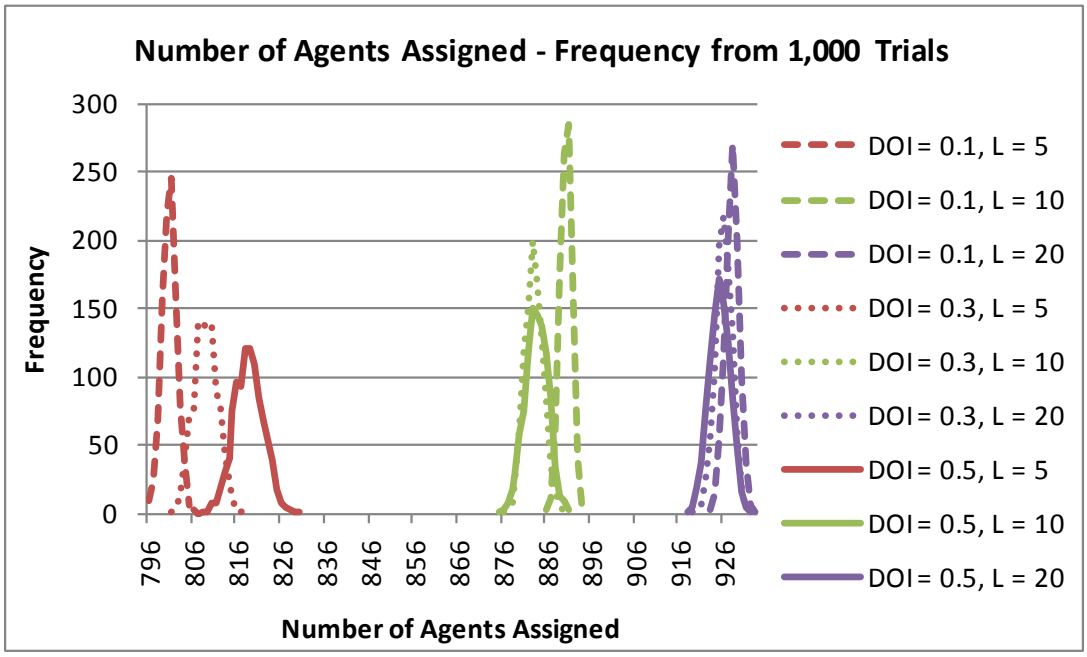

Figure 43. Effect of Preference List Length on Number of Agents Assigned

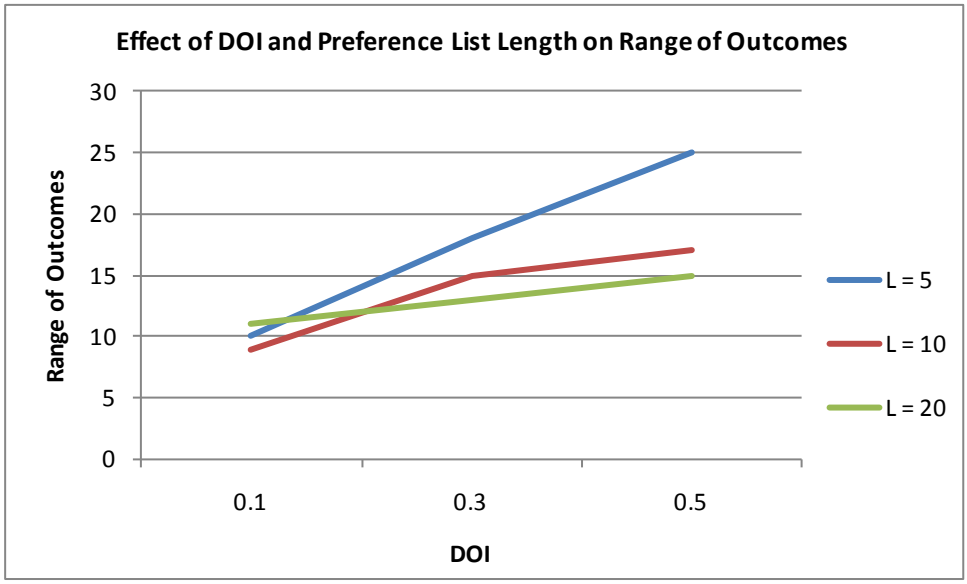

Figure 44. Effect of DOI and L on Range of Outcomes

Not only does the preference list length affect the number of agents assigned, but it also affects the preference rank at which they are assigned. As the length of preference lists increases, more agents will be assigned, but as outlined below, the "cost" of doing so is that the average preference rank of agents who are assigned will also increase.

For each of the 1,000 trials conducted for each scenario, the average personnel, position and total (combined personnel and position) preference rankings of assigned agents were determined. Figure 45 shows the average total rank for each of the 1,000 trials (each point represents the outcome from a trial) for the scenario where $L = 5$ and $DOI = 0.1$. The other scenarios follow a similar pattern and are not reproduced here. To compare the preference ranks at each scenario, the nondominated outcomes for each number of assigned agents are used; that is, the lowest point in each column of points in Figure 45.

The nondominated outcomes at each number of assigned agents are compared in Figure 46 to Figure 48. The vertical axis scale in these charts is held constant to allow for easier comparison and the range of the horizontal axis is also held constant (although the values of the scale on the horizontal axis differ). As the preference list length L increases, so too does the average preference rank of the assigned agents.

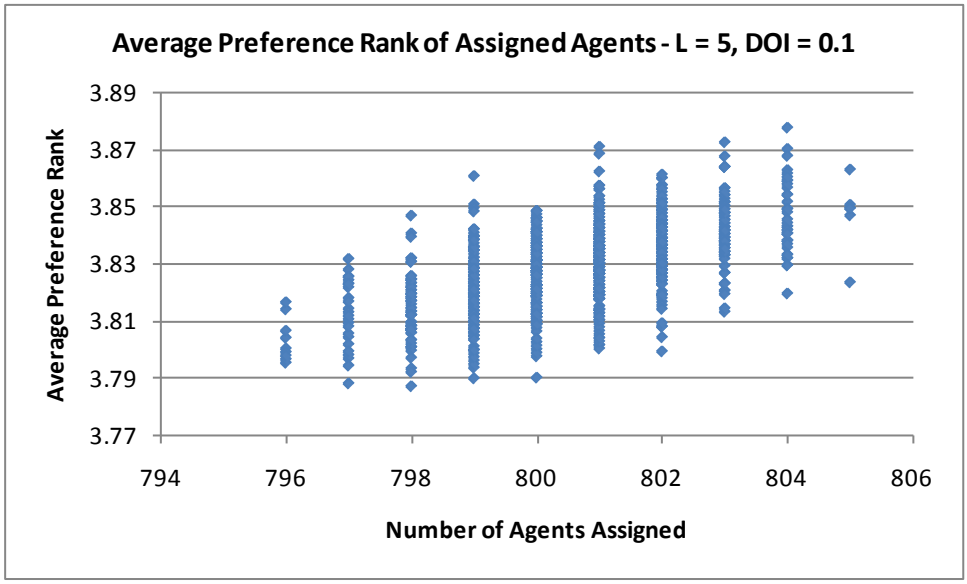

Figure 45. Average Preference Rank of Assigned Agents - 1,000 Trials for L = 5, DOI = 0.1

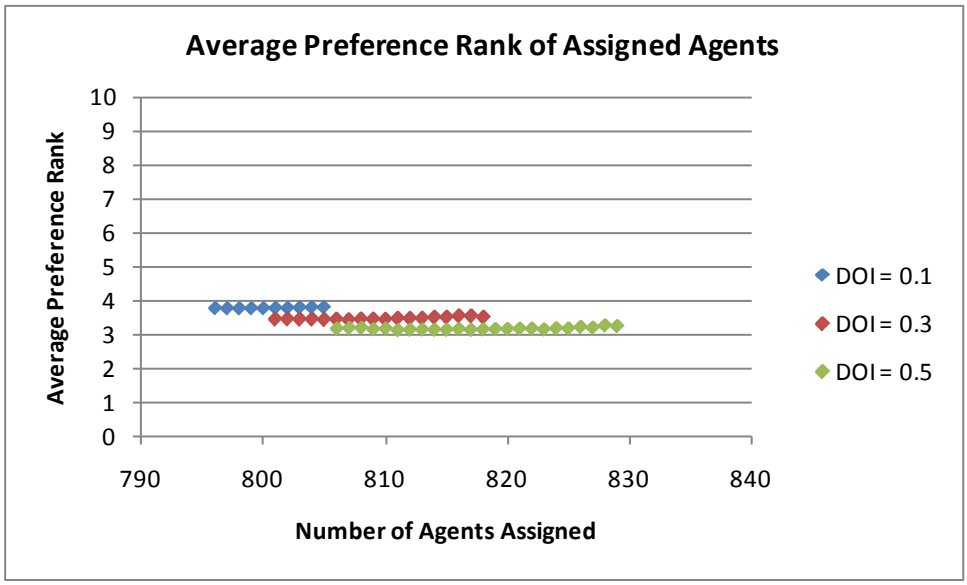

Figure 46. Average Preference Rank of Assigned Agents for L = 5

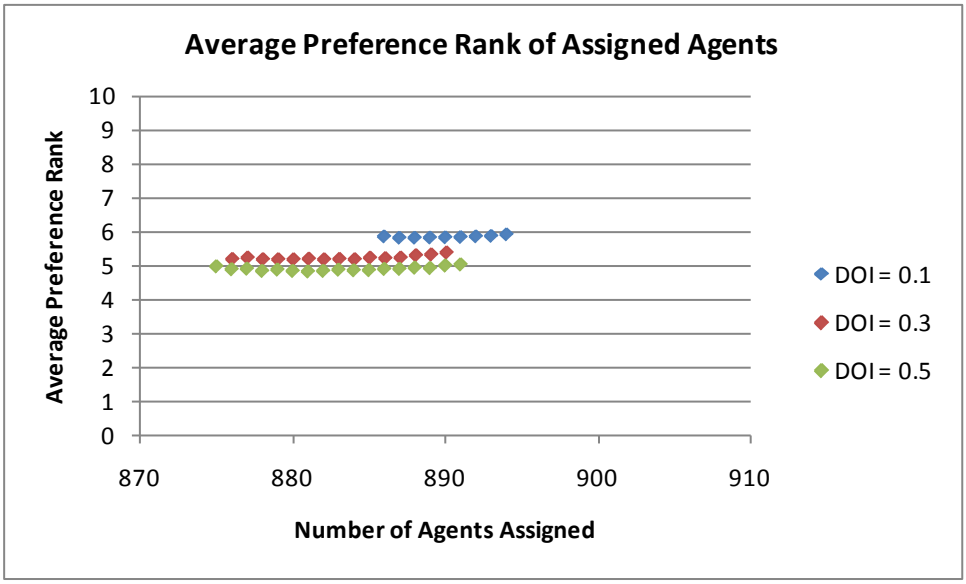

Figure 47. Average Preference Rank of Assigned Agents for L = 10

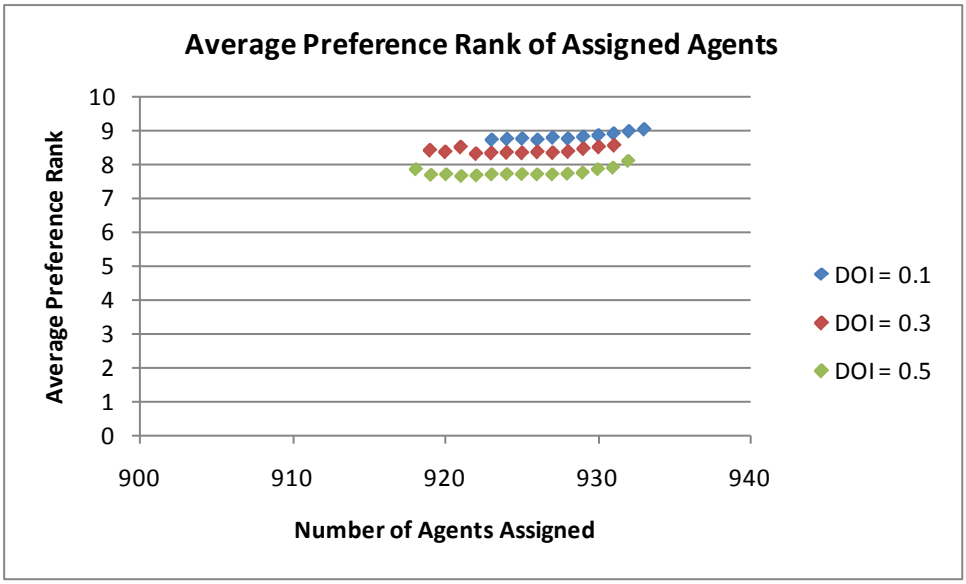

Figure 48. Average Preference Rank of Assigned Agents for L = 20

There is a tradeoff for a decision maker to allow longer preference lists; longer preference lists will assign more agents because each agent has more potential partners, but some agents will be assigned at a higher (worse) average preference rank because the new potential partners that provide this flexibility have higher ranks. While restricting preference list length would lead to a better average result for those who are assigned

through the initial assignment process, more agents would need to be matched in a second round or matched through a manual administrative process, both of which would potentially add instability to the system.

3. Preference List Indifference and Preference Correlation

The preference correlation defines, in terms of quality scores, how similar the personnel and position agents need to be to one another to be ranked by each other. As *C* decreases, the strength of the homophilous attraction increases. To examine the effect of preference list indifference on scenarios involving different preference correlation, scenarios were created where the preference correlation (*C*) and degree of indifference (*DOI*) were varied while other parameters were held constant. For these scenarios, the number of agents (*n*) is held constant at 1,000 personnel and position agents (the same set of randomly generated agents used in the previous section is again used here) and the preference list length (*L*) is held constant at 100. The preference list length was set at a sufficiently high value to ensure that it will not shorten preference list length – rather, the preference correlation (C) will determine the number of preferences that each agent will identify. Depending on the random distribution of agents' quality scores, some personnel agents will have many position agents within close proximity and other personnel agents will have few position agents in close proximity (and vice versa).

To increase internal validity and ensure that differences in the outcomes relate to the variables of interest, *C* and *DOI*, a single set of 1,000 personnel and positions agents were used; this ensures that variability is not due to differences in the agents. To allow comparisons against the last section, which examined the effect of preference list length, the same agents used for those scenarios were again used for these scenarios.

The variable *C* provides a cutoff at a distance +/- *C* from each agents quality score and all agents within that range may be ranked. Given that the quality scores are randomly distributed in the range 0 to 1, and given that there are 1,000 agents, the position agents would average 20 personnel agents within the bounds defined by *C*, for a *C* setting of 0.01 (+/- 0.01 x 1,000). However, some position agents could have more personnel agents within the limits defined by *C* and other position agents could have

fewer (the same applies to personnel agents identifying position agents). The variable *C* is set at the levels 0.0025, 0.005 and 0.01, because on average these correspond to preference list lengths of 5, 10 and 20 for the agents. The actual number of preferences that agents identify based on these *C* values is shown later. Given that different values for *C* will necessarily determine the preference list length for each agent (based on the number of agents within the range defined by +/- *C*), the results for varying *C* are expected to show some relationship to the results for varying *L*.

For each setting of the parameter *C*, the degree of indifference (*DOI*) in the preference lists was set at three levels; 0.1, 0.3 and 0.5. These are the same *DOI* settings used in the previous sections. With three settings of the variable *C* and three settings of the variable *DOI*, there are nine scenarios explored in this section. There were 1,000 trials (random ordering of tied preferences) conducted for each scenario.

Figure 49 shows the results from all nine scenarios used to investigate the effect of preference list correlation and *DOI*. Each line in the figure represents the outcomes from the 1,000 trials of a scenario and shows the number of times from the 1,000 trials that a given number of agents were assigned.

As demonstrated in the results presented in Figure 49, the preference correlation is the primary determinant of the number of agents assigned. As the preference correlation increases, that is as position agents are increasingly willing to consider personnel agents who are less alike in quality values (and vice versa), the number of assigned agents increases. This is akin to an increase in the preference list length; essentially, as *C* increases, position agents identify an increasing number of personnel agents (and vice versa).

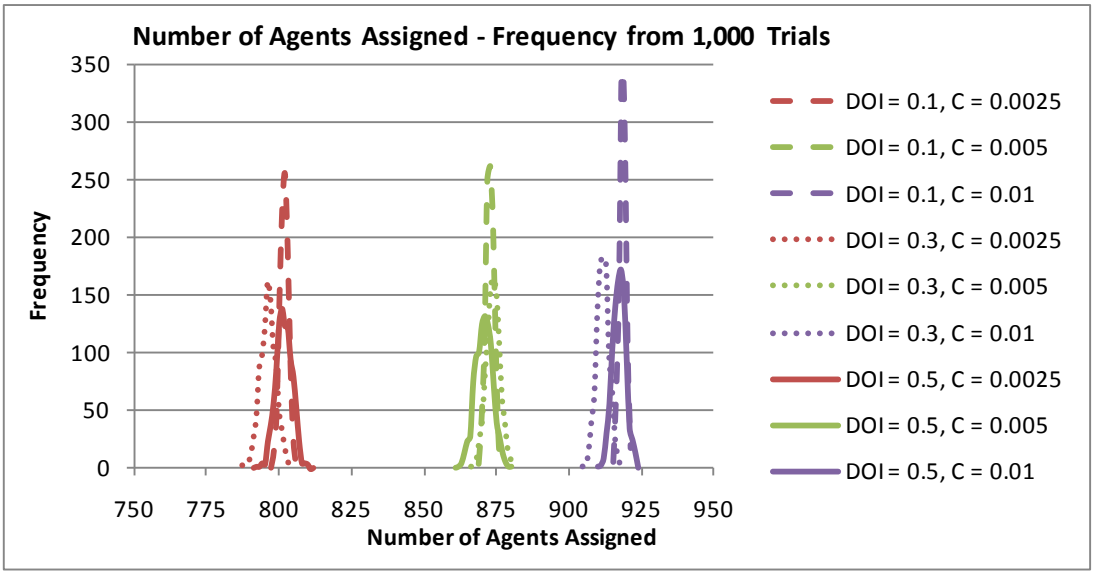

Figure 49. Effect of Preference Correlation on Number of Agents Assigned

As already explained, the variable *C* was set at the levels 0.0025, 0.005 and 0.01, which would, on average, correspond to agent preference list lengths of 5, 10 and 20; recall positions identify all personnel whose quality score is within +/- C in their preference list, and if quality scores were uniformly distributed, the number of personnel agents within this range would be 2*C* x *n*. However, due to random variation in the quality scores, some agents would have longer preference lists and others shorter. Figure 51 shows the distribution of preference list lengths for position agents, where C = 0.01. Of the 1,000 agents, an equal number, 111 agents, had 19 and 20 preferences; 501 agents had 19 or less preferences and 499 had 20 or more preferences. Some position agents had as few as 6 preferences and others as many as 33.

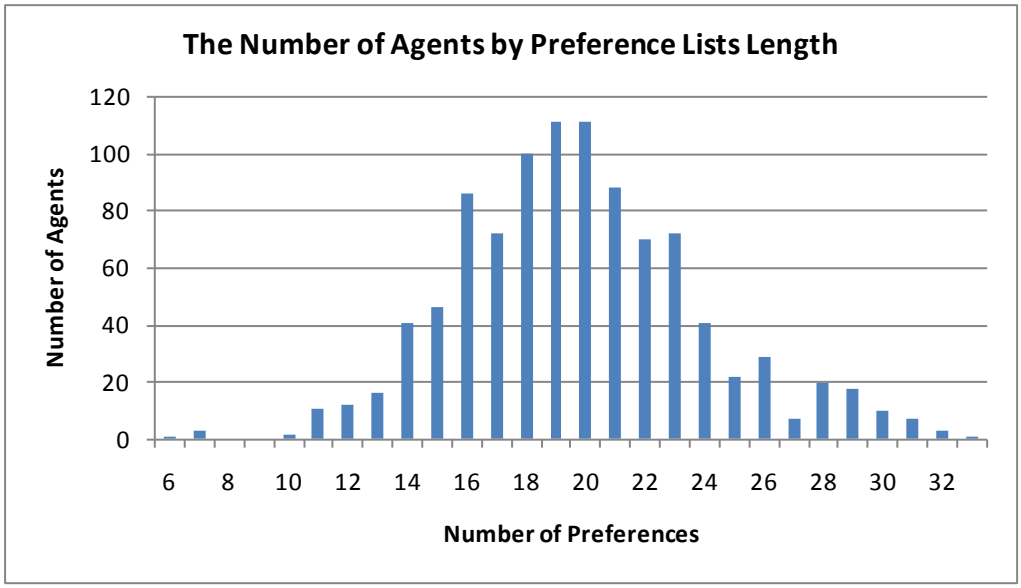

Figure 50. Preference List Length Frequency for C = 0.01

The results in Figure 49 show a strong resemblance to the results in Figure 43, which is the equivalent chart showing the number of agents assigned at each of the 1,000 trials as preference list length and indifference vary. The outcomes for C = 0.0025, C = 0.005 and C = 0.01 resemble the outcomes for L = 5, 10 and 20, respectively. Therefore, in terms of effect on the number of agents assigned, the parameter C has virtually the same effect as setting parameter L to a corresponding level.

In examining the interaction between degree of indifference and preference list length, the degree of indifference did not substantially influence the median number of agents assigned but it did increase the range in the number of assigned agents from the 1,000 trials in each scenario. Similarly, from examining the relationship between degree of indifference and preference correlation, there was no evidence that the degree of indifference influenced the median number of assigned agents but it did lead to a larger spread in that outcome. This is seen in Figure 49 where the higher *DOI* leads to lower peaks but a broader spread at the base of the distribution. This is also demonstrated in Figure 51 where the range of assigned agents is shown, relative to the degree of indifference for different preference correlations; that is, the difference between the maximum and the minimum number of agents assigned from the 1,000 trials. As position

agents restrict their preferences to those personnel agents who are most similar (and vice versa), the range in the number of agents assigned from the 1,000 trials increases. This resembles the outcomes observed for the scenarios varying *L*, as shown in Figure 44.

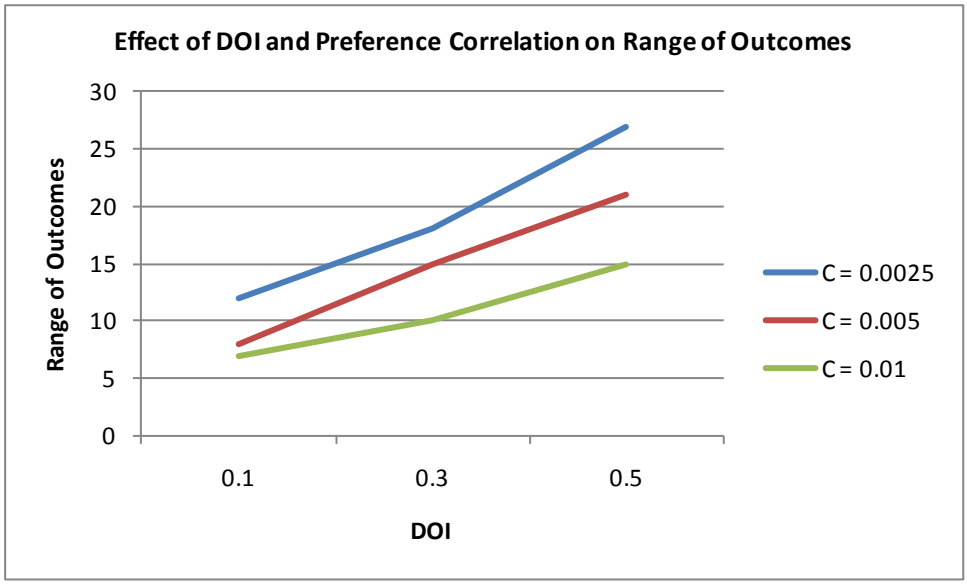

Figure 51. Effect of DOI and C on Range of Outcomes

While the number and range of assigned agents shows considerable similarity between corresponding scenarios defined by the variables *L* and *C*, the next issue to investigate is how the quality of assignments is influenced by *C* and *DOI*. The results shown in Figure 52 to Figure 54 are the nondominated outcomes at each number of assigned agents. Each point in these charts represents the sum of the average preference rank of assigned personnel agents and the average preference rank of all assigned position agents.

The charts in Figure 52 to Figure 54 show a strong resemblance to the corresponding charts for differing preference list length. For example, the average preference ranking for the scenarios where $C = 0.0025$ was between 3 and 4, a similar outcome to the scenarios where $L = 5$. The effect of *DOI* showed a similar pattern too. That is, the lowest average preference rank was achieved in the scenarios where $DOI = 0.5$ and the highest average preference rank was achieved in the scenarios where $DOI = 0.1$. This is to be expected given that a higher degree of indifference will effectively

mean that an agent is more likely to be assigned at a lower preference rank. For example, consider a position agent who identifies six person agents as preferences and is indifferent between the fourth, fifth and sixth preferences. In such a case, when the preference ties are broken, it is immaterial (to the position) whether the position is assigned to its fourth, fifth or sixth preference because it would be identified as the fourth ranked preference in each case.

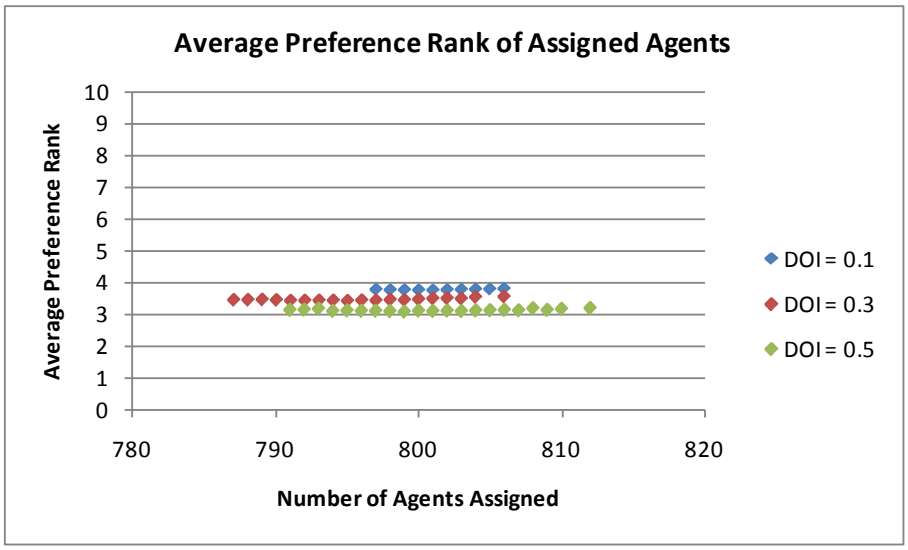

Figure 52. Average Preference Rank of Assigned Agents for C = 0.0025

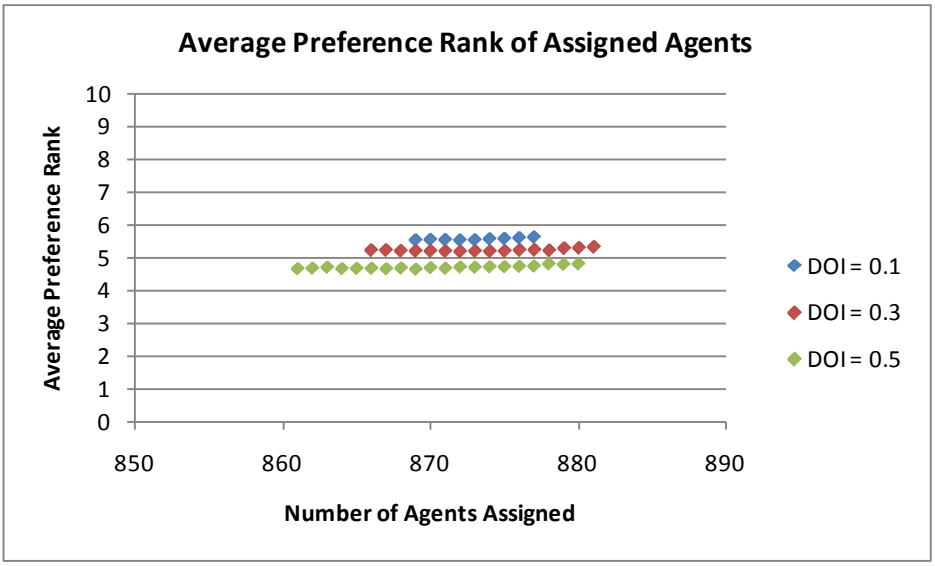

Figure 53. Average Preference Rank of Assigned Agents for C = 0.005

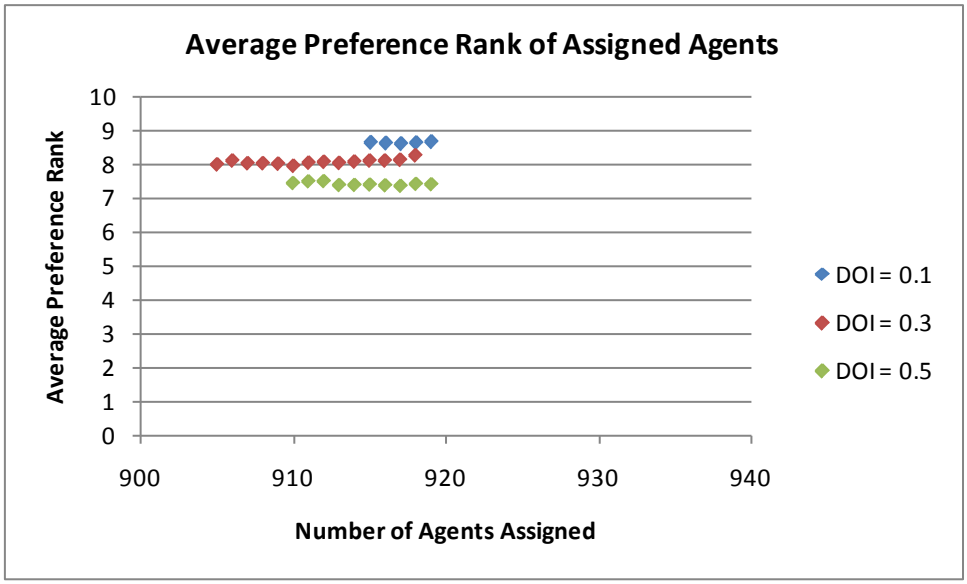

Figure 54. Average Preference Rank of Assigned Agents for C = 0.01

As shown in Figure 50, agents had different preference list lengths due to the random distribution of quality scores and the use of the preference correlation to limit preferences. It is possible to examine the relationship between an agent's preference list length and the probability of the agent being assigned. Figure 55 shows the percentage of agents assigned according to their preference list length. The data is for the position agents at one trial of a scenario where C = 0.005 (the *DOI* is not relevant because the results are from a single trial). The variability at the lower number of preferences is due to the small number of agents who had such a low number of preferences. The same data is presented by number of agents at each preference list length in Figure 56.

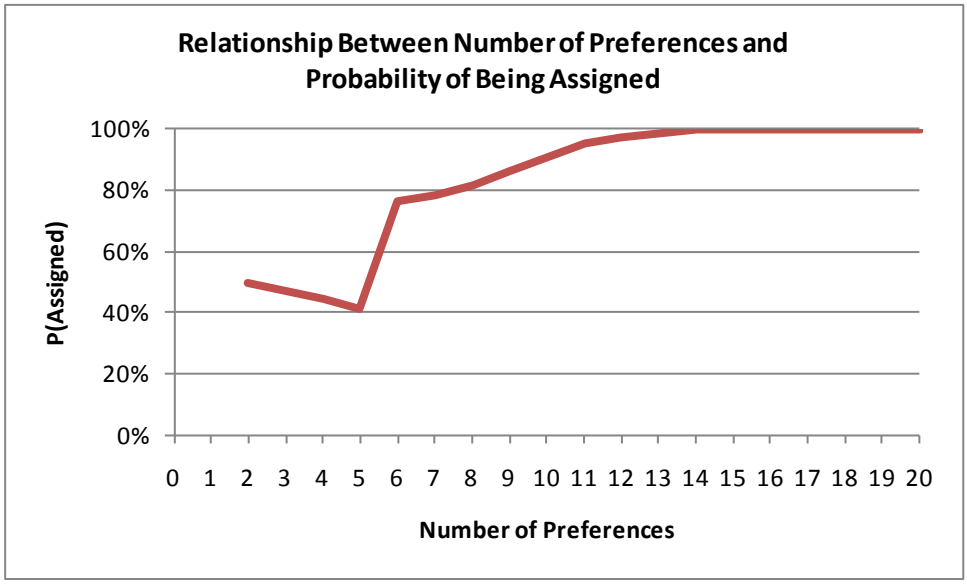

Figure 55. Percentage of Position Agents Assigned by Preference List Length

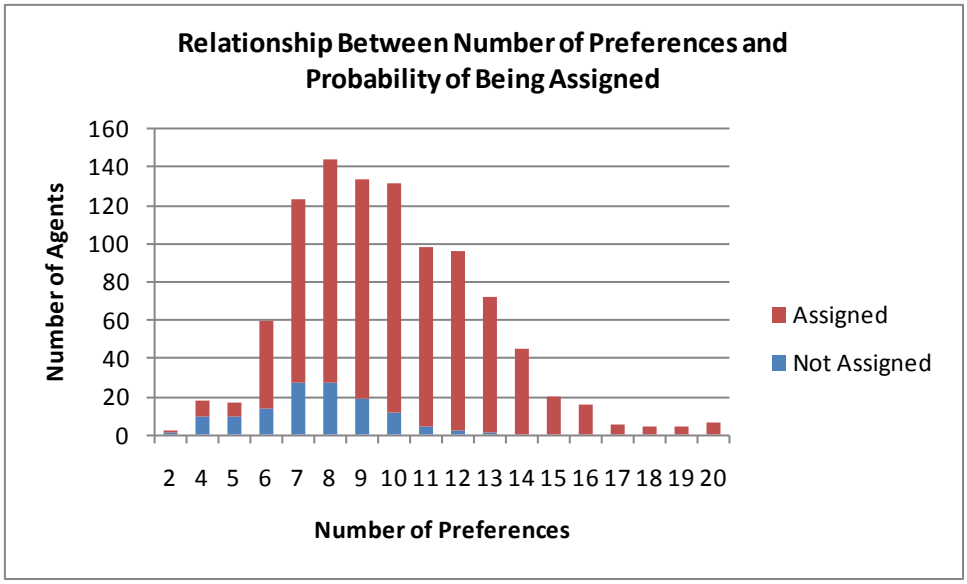

Figure 56. Number of Position Agents by Preference List Length and Assignment Outcome

The average preference list length for personnel and position agents who were matched was 10.20 and 10.18, respectively, while the average preference list length for person and position agents who were not matched was 7.43 and 7.57 respectively. Such results are comparable to the National Resident Matching Program (NRMP), which

reports that applicants matched in 2010 had an average Rank Order List (ROL) length of 9.25 and unmatched applicants had an average ROL length of 4.50 (http://www.nrmp.org/res_match/about_res/impact.html).

C. SUMMARY

The participants in some applications of two-sided matching are indifferent between some options and would ideally express this as tied preferences. However, existing applications of two-sided matching require participants to either submit strictly ordered preferences or, if preference ties are allowed, the ties are randomly broken. Erdil and Ergin (2007) provide evidence that such an approach will not necessarily yield the best possible results and significant improvements may be possible.

An agent based model was developed with two groups of agents; personnel and position agents. Each agent was randomly assigned a quality score and their preference for agents from the other group was based on homophilous attraction. Scenarios were defined by four parameters; the number of agents (n), the length of preference lists (L), the preference correlation (C) and the degree of indifference (DOI). The goal was to investigate how DOI interacted with the other parameters and the effect that this had on two key outcomes; the number of agents assigned and the average preference rank of assigned agents. For each scenario, multiple trials were conducted with tied preferences randomly broken at each trial. The number of agents assigned and their average preference rank varied according to how the tied preferences were broken.

The results from many trials on each scenario (generally 1,000 trials per scenario) show that the frequency at which each number of agents was assigned resembles a normal distribution; for any given scenario most random tie breakings yield a similar number of agents assigned but some tie breakings can produce results further from the median outcome. The range between the lowest and highest number of assigned agents increases as the DOI and the number of agents in the scenario increases. However, the

increase in the range is not proportional to the number of agents involved in the scenario; as the number of agents increases, the range decreases if measured as a percentage of the number of agents involved.

As preference list length increases the number of assigned agents increases but the average preference rank of those who are assigned decreases. Given that any unassigned agents would need to be assigned through a second round match or through an administrative process, the administrator of a two-sided matching application may wish to maximize the number of agents assigned in the initial round, although they should be aware of the "cost" of doing so in terms of average rank of the assigned agents.

When preferences are defined by the preference correlation (*C*) and are not constrained by *L*, the results are similar to those obtained when preference list lengths are limited by *L*; this reflects that *C* still limits each agent's preference list, although it is limited by the number of partner agents within a distance of the agent's quality score, with the limit set by C and not the preference list length. For the scenarios examined here, the primary determinant of outcomes was the number of preferences that an agent lists and not necessarily whether the preferences are correlated; preference correlation effectively limits preference lists by limiting potential partners. Whether preferences were limited by *L* or *C*, preference indifference did not significantly affect the median outcome in the many trials explored. However, increasing *DOI* did lead to a broader range in the number of agents being assigned. The average preference rank of the assigned agents was lower at higher levels of indifference; this is the result of measuring the outcomes against the indifferent rank at which the agents were assigned; for example, if an agent's 9th, 10th and 11th preferences are tied, when the ties are broken the rank for each of these preferences is counted as 9.

Two-sided matching can be applied to situations that involve preference list indifference. If preference list ties are randomly broken and multiple assignment trials are explored, it is possible to identify a range of potential outcomes, each of which is at least weakly stable. In a hierarchical organization where the match administrator has an

interest in the assignment outcome, identifying a variety of alternative outcomes provides the administrator with options to select the outcome that best meets the organization's requirements.

Table 37 summarizes the key findings from the experiments that investigated preference list indifference.

Factor	Key Findings
Number of trials	Increasing the number of trials may increase the range of the number of agents assigned; that is, it may reveal new outcomes where it is possible to assign greater numbers of agents. Further, increasing the number of trials may reveal new assignment outcomes that dominate others in terms of average rank for assigned agents. 1,000 trials was found to be sufficient to reveal most outcomes for the types of scenarios conducted in this research.
Number of agents	As the number of agents increases, the range in the number of assigned agents increases in total numbers but decreases as a percentage of the number of agents.
Preference list length	Increasing preference list length leads to a higher number of assigned agents. The range of outcomes increases as preference list length increases. The average preference rank of assigned agents increases as preference list length increases.
Preference correlation	Similar effect to preference list length; a lower preference correlation effectively increases preference list length and so allows more agents to be assigned. Similar effect to preference list length; a lower preference correlation effectively increases the preference list length and so results in a higher average rank for assigned agents.
Degree of Indifference	The range in the number of assigned agents increases as the degree of indifference increases. The average preference rank of assigned agents decreases as the degree of indifference increases.

Table 37 Summary of Key Findings from Preference List Indifference Experiments

The next chapter examines two-sided matching as applied to military assignments; that is, the allocation of personnel to positions as part of a regular posting cycle. The approach of conducting multiple trials to explore a variety of outcomes under preference list indifference will be continued in the assignment scenarios examined in the next chapter.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. APPLYING TWO SIDED MATCHING IN A HIERARCHICAL ORGANIZATION

Chapter I outlined potential processes for assigning personnel to positions. Two-sided matching was selected because of two key desirable properties; it yields stable results, and the dominant strategy is for personnel to submit their true preferences (they cannot game the system) if implemented as a personnel-proposing algorithm. Chapter II outlined that two-sided matching has been used in a wide variety of applications, and during the last decade some large school districts in the U.S. have used this process to assign students to schools. These school assignment applications have features of hierarchical organizations with the individual schools being subordinate to the district education department.

In applications, such as the New York City high school matching, the individual schools have preferences for the types of students they wish to accept, but the schools' preferences must also conform to constraints imposed by the Department of Education (Abdulkadiroğlu et al., 2005). As an alternative to using constraints, this research demonstrates that multi-attribute utility functions can be utilized to generate preference lists in hierarchical organizations; this allows the preferences to include attributes that are of interest to both the organization and its subordinate elements. It is possible to give priority to different attributes by changing the weights in the multi-attribute utility function. This chapter demonstrates the use of such utility functions and examines how assignment outcomes vary in response to changing weights (priorities). The research is based on a military assignment situation where personnel are assigned to positions, and utilizes data on Australian Army personnel and positions. The military units and Service headquarters are analogous to the schools and education department.

A. PROBLEM SPECIFICATION

Personnel and positions are each identified by six attributes in the scenarios developed here: rank, branch, skill, location, career development and performance. The first three of these attributes are of interest to the units, because they define the attributes

a person should have to successfully fill a position within the unit, and therefore contribute to the unit's role. The remaining attributes are of interest to the service because they contribute to broader organizational objectives; cost minimization through geographic stability, personnel development through a range of experiences, and maximization of organizational performance by matching the most demanding positions with the highest performing personnel.

As outlined previously, the position preference rankings that are required for two-sided matching are derived from a position utility function:

$$U_{total} = (W_{rank} \text{ x } U_{rank}) + (W_{branch} \text{ x } U_{branch}) + (W_{skill} \text{ x } U_{skill})$$
$$+ (W_{location} \text{ x } U_{location}) + (W_{career\ dev} \text{ x } U_{career\ dev}) + (W_{performance} \text{ x } U_{performance})$$

where U represents Utility and W represent Weight. For each position, the person with the highest U_{total} value is the position's first ranked preference, and subsequent preferences are in order of U_{total}. Changing the weights will result in positions rank ordering personnel differently, which in turn leads to different assignment outcomes from the two-sided matching.

There are 1,790 personnel and 1,997 positions that require assignment in the problems specified here. This represents one third of Australian Army personnel and positions at the rank of E06, E08 and E09 (E07 rank is no longer used). Actual data for the rank, branch, skill and location of these personnel and positions is used. Performance quartiles are randomly assigned for both personnel and positions. The career development provided by positions is based on a categorization of units, and personnel are randomly assigned career development requirements, such that the percentage of personnel requiring each career development type matches the percentage of positions providing that career development type. The location preferences of personnel are randomly allocated, but distributed in accordance with preferences that have been submitted by personnel through an Electronic Preferences and Restrictions (EPAR) form.

Position preference rankings are based on six attributes and calculated in accordance with the utility function outlined above. The 50 personnel with the highest utility values become each position's preference list with the utility values determining

whether preferences are tied. A preference list length of 50 was selected after reviewing the utility function values; this provides sufficient preferences to ensure that a high percentage of agents are assigned, and avoids including personnel who meet too few of the positions' attributes. A later section will show a modification from position preference list lengths of 50. Personnel preference lists are calculated after position preferences are determined, because personnel can only identify a position if the position has first identified the person amongst one of its preferences. Each person can list up to L_{pers} positions in their preference list, where L_{pers} is a variable for the experiments.

A semi-random ordering of attributes is created for each person; for every person the rank attribute is first and the performance attribute is last, with the remaining attributes randomly ordered. Rank Order Centroid weights are applied to each person's ranked attributes to provide each attribute's weight. Each person's attribute weights remain constant across all scenarios conducted here. The attributes and weights provide an initial ordered list of each person's acceptable positions. However, the personnel preferences are updated for each scenario based on the position preferences; a person only includes a position as a preference if the position has first listed the person as a preference.

Table 38 outlines the settings for the scenarios to be examined. These scenarios vary the balance of weights between unit and Service attributes; the weight provided to unit attributes is 75% for Scenarios 1 and 2, 60% for Scenarios 3 and 4, 45% for Scenarios 5 and 6, and 30% for Scenarios 7 and 8; weights are evenly spread across unit related attributes. 1,000 trials were conducted for each scenario with preference list ties randomly broken at each trial.

	Preference List Lengths		Position Weights					
Scenario	L_{pers}	L_{posn}	Rank	Branch	Skill	Location	Career Dev	Performance
1	10	50	0.25	0.25	0.25	0.10	0.10	0.05
2	20	50	0.25	0.25	0.25	0.10	0.10	0.05
3	10	50	0.20	0.20	0.20	0.15	0.15	0.10
4	20	50	0.20	0.20	0.20	0.15	0.15	0.10
5	10	50	0.15	0.15	0.15	0.20	0.20	0.15
6	20	50	0.15	0.15	0.15	0.20	0.20	0.15
7	10	50	0.10	0.10	0.10	0.30	0.20	0.20
8	20	50	0.10	0.10	0.10	0.30	0.20	0.20

Table 38 Scenario Specifications

The scenarios with a personnel preference list length of 10 are used because this is close to the number allowed in some of the main existing applications. For example, Abdulkadiroğlu et al. (2005) identified that students in the New York City high schools match are permitted to submit preferences for 12 schools. To analyze the effect of personnel preference lists lengths, a second setting where L_{pers} equals 20 is used.

The personnel proposing algorithm is used because it is a dominant strategy for personnel to submit truthful preferences. Although not explored in this research, further work could examine the difference in results between the personnel and position proposing algorithms.

B. RESULTS

1. Scenario 1

In the last chapter, the Degree of Indifference (DOI) was used to define preference list indifference in a scenario. The personnel and positions' attributes, together

with utility function weights, define the preferences in the scenarios explored in this chapter. However, it is still possible to identify the extent of indifference in these preferences by expressing a value that is comparable to the DOI. For all personnel agents, a count is made of the number of times that an agent's $(n+1)^{th}$ preference is indifferent to the same agent's n^{th} preference (where $n \geq 1$). The DOI is then calculated by dividing the count of indifferent preferences by the total number of preferences of all personnel agents; for the first scenario this provides a DOI of 0.35. This can be interpreted as follows: for any personnel agent's preference (with the exception of their first preference), there is a 35% chance that the preference is indifferent to the previously identified preference. The DOI for the positions in the first scenario is 0.75. The lower DOI for personnel occurs because personnel attribute weights are based on Rank Order Centroid weights, so no two attributes are weighted the same. In contrast, the higher DOI for positions occurs because some attributes have the same weights, and therefore more preferences are indifferent.

Due to the indifference in the personnel and positions' preference lists, the results observed in these scenarios show similar traits to those in the last chapter that explored preference list indifference. For example, Figure 57 shows the number of agents assigned at each of the 1,000 trials and the distribution pattern is similar to scenarios explored in the last chapter; that is, when tied preferences are randomly broken, only a small percentage of trials will produce outcomes where the number of assigned agents is near the lower or upper extremes. Only one of the 1,000 trials for Scenario 1 produced an outcome with 1,604 agents assigned, 90% of the number of personnel involved.

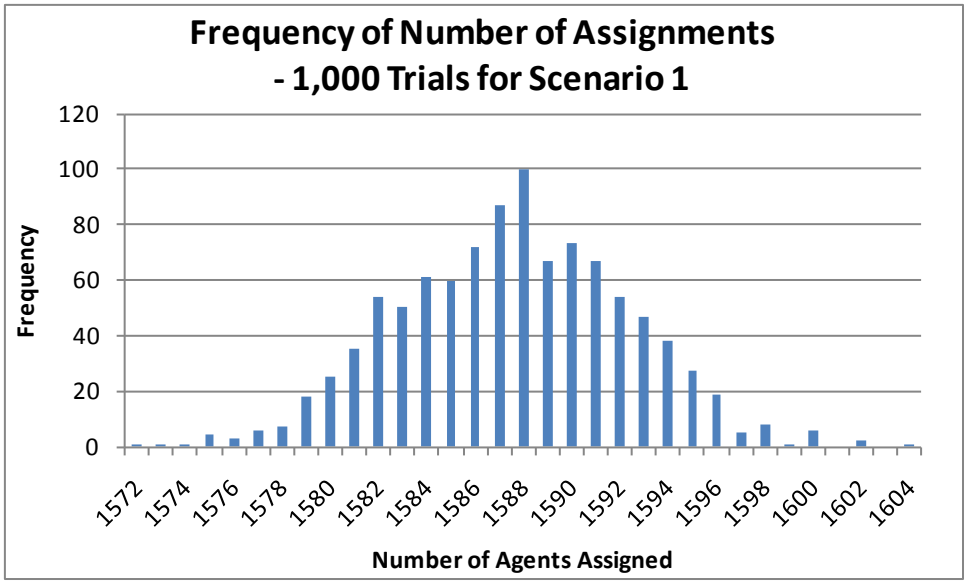

Figure 57. Frequency of Assignments – 1,000 Trials for Scenario 1

Figure 58 shows the distribution of average preference rank for the assigned personnel. This distribution also shows similar patterns to the results of last chapter and demonstrates that there are many dominated outcomes. From the nondominated outcomes, the assignments with fewer agents matched achieve a lower (better) average personnel rank.

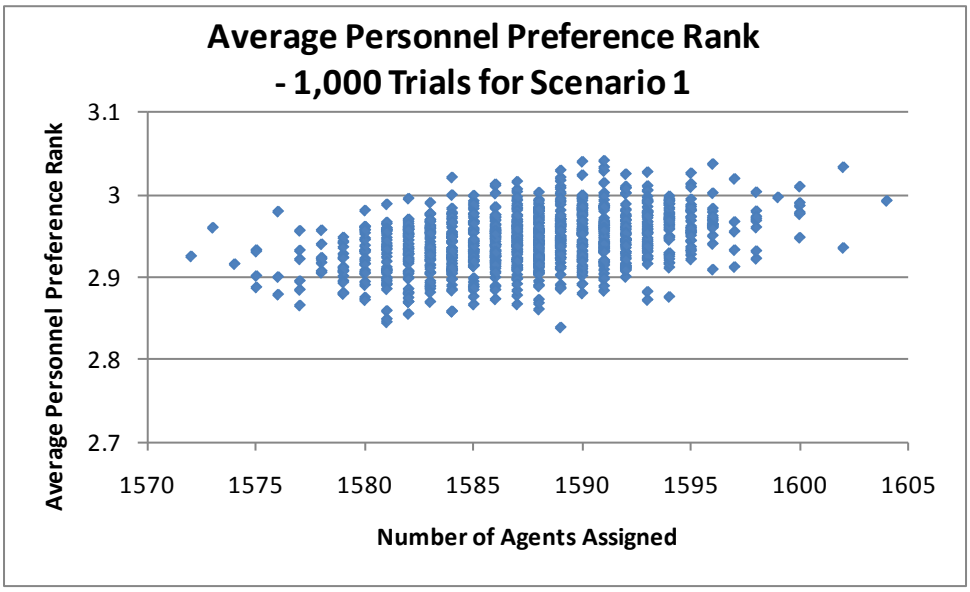

Figure 58. Average Personnel Preference Rank – 1,000 Trials for Scenario 1

Trial 903 was the only trial to achieve the maximum 1,604 assigned agents in Scenario 1, and is therefore selected for further analysis. Figure 59 shows the indifferent preference rank at which each agent was assigned in trial 903; if an agent is indifferent between preferences 4, 5 and 6, this would be identified as the agent's fourth preference. 41% of personnel were assigned to their first preference, 68% were assigned to one of their top three preferences, and 83% were assigned to one of their top five preferences.

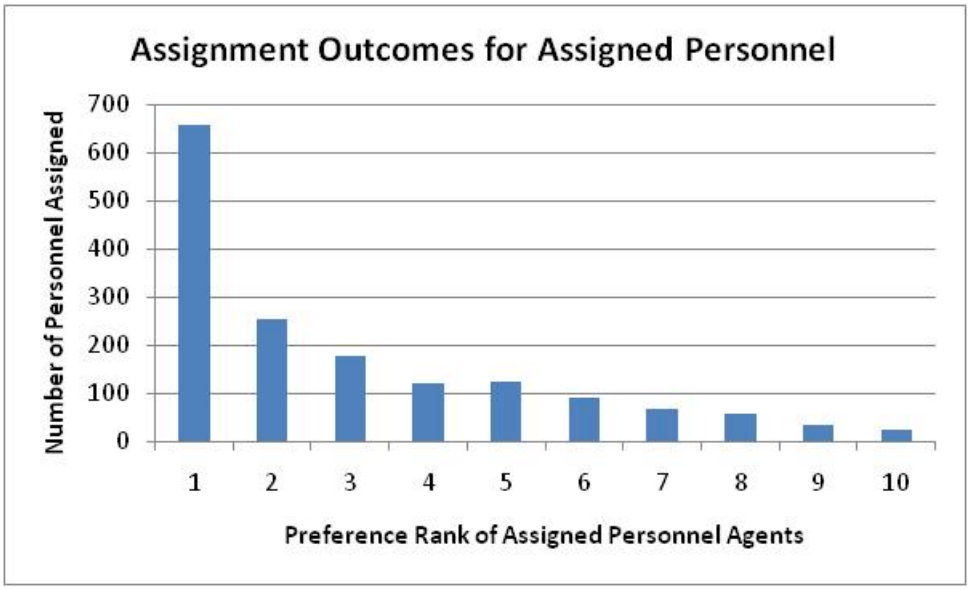

Figure 59. Preference Rank of Assigned Personnel Agents - Scenario 1, Trial 903

It is relevant to examine the personnel outcomes in terms of preference ranks achieved because the personnel submit their preferences as rank ordered lists. However, it is less relevant to examine position outcomes in terms of preference ranks achieved; the preference ranking is a requirement of the two-sided matching process, but how well assignments satisfy position requirements matters most. In a hierarchical organization where positions' preferences are constructed from a multi-attribute utility function, the positions' attributes are known, as is the utility of each assignment. Figure 60 shows the utility of assigned positions to show how well the positions' requirements have been satisfied: 16% of positions were assigned to personnel that provide a utility exceeding 0.9 (out of a theoretical maximum of 1.0), 67% were assigned with a utility exceeding 0.8, and 84% were assigned with a utility exceeding 0.7.

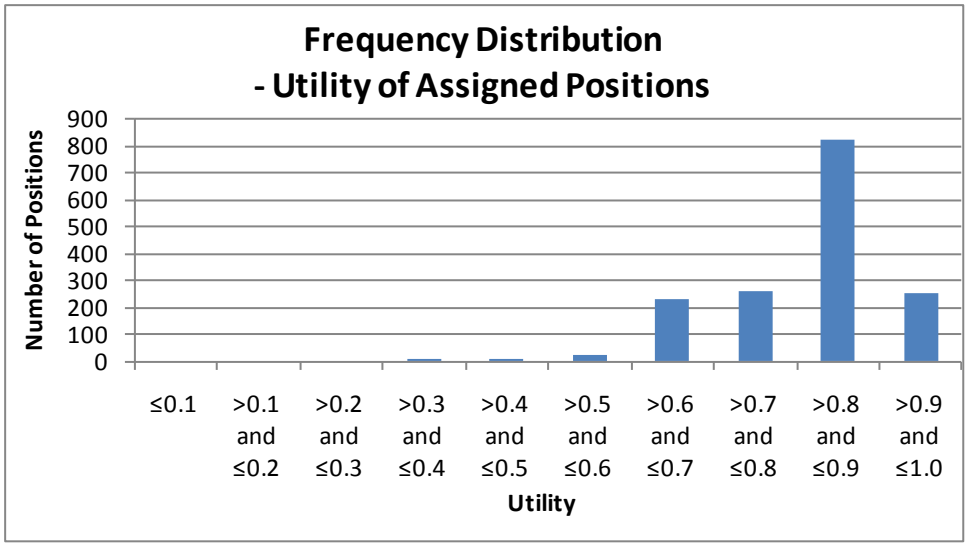

Figure 60. Utility of Assigned Position Agents – Scenario 1, Trial 903

This research is designed to demonstrate the benefit of utility functions in creating rank ordered preference lists for two-sided matching. For such a process to be advantageous, assignments must be responsive to the weights; for example, as the skill attribute weight is increased, then an increasing percentage of assignments should satisfy positions' skill requirements. This will be demonstrated in these scenarios by examining the percentage of assignments that satisfy each of the positions' attributes.

Figure 61 shows the results for trial 903 of Scenario 1. The three unit attributes have greater weight than the Service attributes and this is reflected in the assignments; 99% of assignments meet the positions' rank requirements, 100% meet the branch requirements and 90% meet the skill requirements (it was outlined previously that a veto rule was created for positions' branch requirements; therefore, all assignments will meet the branch requirements). By contrast, the Service attributes were provided less weight in Scenario 1, resulting in 33% of assignments achieving geographic stability, 84% meeting the personnel career development needs, and 64% of assignments involving personnel and positions in the same performance quartile.

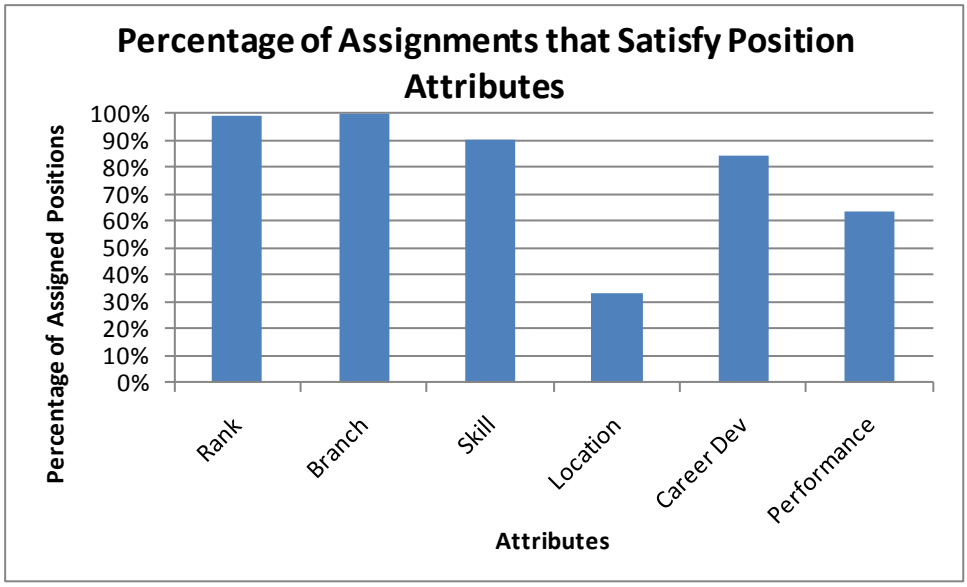

Figure 61. Percentage of Assignments that Satisfy Position Attributes – Scenario 1, Trial 903

2. Scenario 2

The specification of Scenario 2 is identical to Scenario 1 with the exception of the personnel preference list length, L_{pers}; the personnel preferences are limited to 10 preferences at Scenario 1, whereas Scenario 2 permits 20 preferences.

Trial 123 of Scenario 2 achieved the highest number of assigned agents with 1,716 (96%) assignments; the best outcome from Scenario 1 was 1,604 (90%) assignments. Table 39 outlines how increasing the personnel preference list length affects the preference rank of assigned personnel. Increasing the personnel preference list length leads to more assignments; however, fewer personnel, both in numbers and percentage terms, are assigned to their most preferred positions. For example, 1,088 personnel (representing 68% of assigned agents) were matched to one of their top three preferences in Scenario 1, with this figure decreasing to 1,007 agents (representing 58% of assigned agents) in Scenario 2.

Personnel Preference Rank	Number Assigned at Preference Rank (Cumulative number in parentheses)		Percentage of Agents Assigned at Preference Rank (Cumulative % in parentheses)	
	Scenario 1	Scenario 2	Scenario 1	Scenario 2
1	657 (657)	605 (605)	41% (41%)	35% (35%)
2	255 (912)	241 (846)	16% (57%)	14% (49%)
3	176 (1,088)	161 (1,007)	11% (68%)	9% (58%)
4	120 (1,208)	114 (1,121)	7% (75%)	7% (65%)
5	124 (1,332)	112 (1,233)	8% (83%)	7% (72%)
6 and below	272 (1,604)	483 (1,716)	17% (100%)	28% (100%)

Table 39 Personnel Assignments by Preference Rank – Scenario 1 and 2

Increasing the personnel preference list length leads to an outcome that many personnel would agree is less desirable; however, the same reduction in satisfaction does not occur for the positions. This is explained by Roth and Sotomayor (1990), who show that in man-dominant one-to one marriage markets, if a man in the market extends his preference list to include more women, this cannot help the other men and cannot harm the women. Consider a marriage market where the men and women have already been partnered. If one of the men increases his preference list length, he cannot be matched to a woman he prefers because the man has already been rejected by all the women he prefers. But when the man proposes to a woman who is on his extended preference list, the woman will either reject the man (in which case no partnerships are affected), or will accept the man instead of her previous partner; if this occurs the woman achieves a preferred outcome (because she would only reject her previous partner if she prefers the new offer) and the woman's original partner continues through his preference list in

search of a new woman. This potentially commences a chain of new proposals, which cannot help the men and cannot harm the women; men whose partnerships are broken are required to propose to less desirable women, and the women only accept new proposals if it they are from men higher on their preference list.

The situation described above is only strictly true if there is no preference list indifference, or a single tie breaking rule is used if indifference exists. Preference list ties are broken in different ways in this research, and therefore there are some individual positions that will be worse off due to the extension of the personnel preference list lengths. However, Table 40 shows that increasing the personnel preference list length leads to improved results for the positions.

Position Utility	Number Assigned at Utility (Cumulative number in parentheses)		Percentage of Agents Assigned at Utility (Cumulative % in parentheses)	
	Scenario 1	Scenario 2	Scenario 1	Scenario 2
>0.9 and ≤1.0	257 (257)	278 (278)	16% (16%)	16% (16%)
>0.8 and ≤0.9	823 (1,080)	881 (1,159)	51% (67%)	52% (68%)
>0.7 and ≤0.8	263 (1,343)	307 (1,466)	17%) (84%)	17% (85%)
≤0.7	261 (1,604)	250 (1,716)	16% (100%)	15% (100%)

Table 40 Position Assignments by Utility – Scenario 1 and 2

Table 41 shows that the number of assignments where position attributes are satisfied also increases as the personnel preference list length is extended. However, the percentage of assignments that satisfy each position attribute is similar for each scenario.

	Number of Assignments Satisfying Position Attribute		Percentage of Assignments Satisfying Position Attribute	
	Scenario 1	Scenario 2	Scenario 1	Scenario 2
Rank	1,592	1,696	99%	99%
Branch	1,604	1,716	100%	100%
Skill	1,445	1,541	90%	90%
Location	532	615	33%	36%
Career Dev	1,354	1,398	84%	81%
Performance	1,025	1,081	64%	63%

Table 41 Achievement of Position Attributes – Scenario 1 and 2

3. Analysis of Scenarios 1 to 8

Table 42 shows the number of assignments achieved at each scenario. The odd numbered scenarios are those where L_{pers} equals 10, and the even numbered scenarios are those where L_{pers} equals 20. There is an increase in the number of assignments as the scenario weights shift from a primary focus on unit attributes to a primary focus on Service attributes, controlling for the different personnel preference list lengths. This appears to occur because the position preferences are increasingly correlated with the personnel preferences.

Trial	1	2	3	4	5	6	7	8
Number of Assignments	1,604	1,716	1,621	1,727	1,640	1,739	1,679	1,759
Percentage of Personnel	90%	96%	91%	96%	92%	97%	94%	98%

Table 42 Number of Assignments – Scenarios 1 to 8

Figure 62 and Figure 63 show the personnel preference ranks that were achieved for each of the eight scenarios, by number and percentage of assignments. These clearly show that the additional assignments achieved from the scenarios where L_{pers} equals 20 were at the expense of fewer personnel receiving their highest ranked preferences, and more personnel being assigned to preferences ranked 6th or below. Decision makers must determine whether the higher number of assignments warrants the decrease in satisfaction that some personnel would experience. The decision maker would also need to balance this issue against the quality of positions' assignments, which will be discussed later.

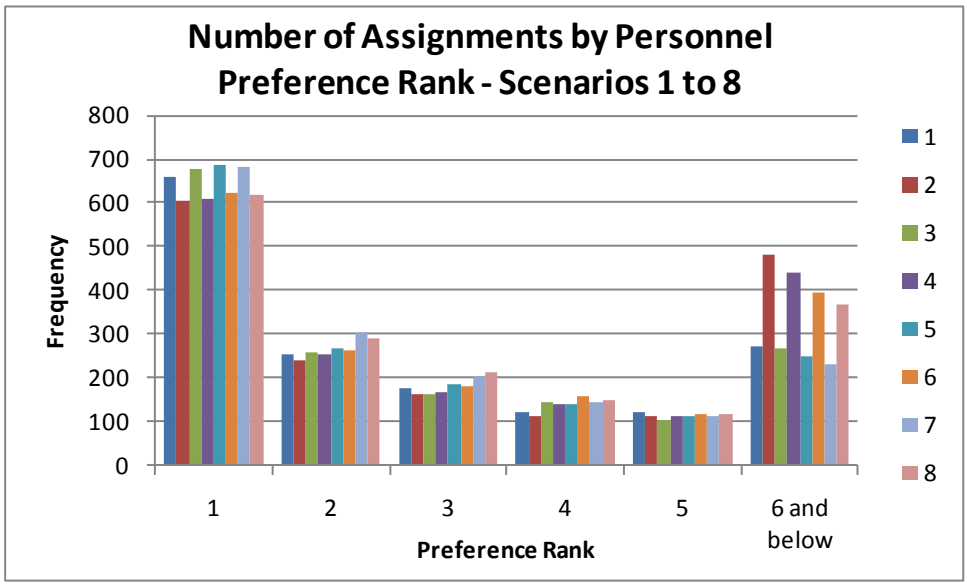

Figure 62. Number of Assignments by Personnel Preference Rank – Scenarios 1 to 8

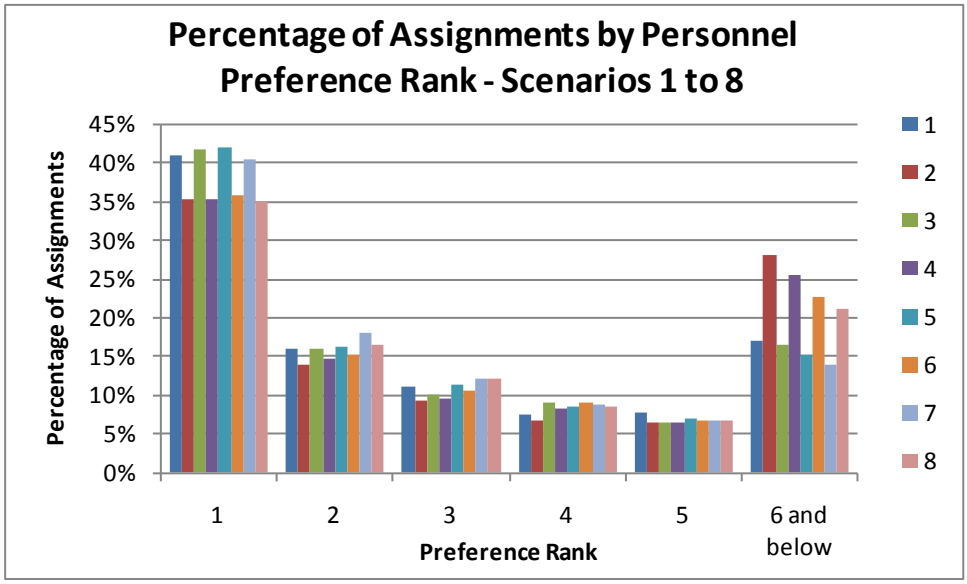

Figure 63. Personnel Preference Rank by Percentage of Assignments – Scenarios 1 to 8

The impact of the additional preference list length is analyzed at the individual level. For this, the preference rank of the 1,604 individual personnel who were assigned at Scenario 1 (trial 903) is compared to their assigned preference rank at Scenario 2 (trial 123). The comparison is shown in Table 43 where personnel whose assigned preference rank improved from Scenario 1 to 2 are recorded as positive values, Table 43 whereas personnel whose preference rank worsened are recorded as negative values. 40 personnel who were assigned at Scenario 1 were not assigned at Scenario 2.

As outlined previously, in the case of strict preferences, increasing the personnel preference list length cannot help the personnel achieve a preferred outcome; however, where preference list indifference exists and ties are not broken in a consistent manner (e.g., personnel receive a lottery number that is used for all scenarios and trials), it is possible for some personnel to benefit from different random orderings. Consequently, 128 personnel who were assigned at Scenario 1 achieved a better outcome at Scenario 2, but 333 personnel who were assigned at Scenario 1 achieved a worse outcome at Scenario 2 (this includes the 40 personnel who were matched at Scenario 1 but are not matched at Scenario 2). 1,140 personnel remained assigned to their same preference rank. These results show that when personnel preference lists are increased, random tie

breaking allows some personnel to achieve better assignments, but the majority will receive a worse assignment. The fact that more receive a worse outcome than a better outcome is reflective of the strict preference situation where all personnel would receive an outcome that is the same or worse.

	Preference Rank at Scenario 2 compared to Scenario 1							
Impact	No longer assigned	-3 or worse	-2	-1	0	+1	+2	+3 or better
No. of Personnel	40	196	41	56	1,140	31	28	69

Table 43 Effect of Increasing L_{pers} on Personnel Assigned at Scenario 1

There were 1,604 personnel assigned in Scenario 1 and 1,716 in Scenario 2, and 40 individual personnel who were assigned in Scenario 1 but were not assigned in Scenario 2. Consequently, 152 personnel who were assigned in Scenario 2 were not assigned in Scenario 1. Table 44 shows the assignment preference ranks for these 152 personnel in Scenario 2, demonstrating that the majority of newly assigned personnel were at relatively high (less desirable) preference ranks. However, provided that these personnel would prefer to be assigned to one of their preferences than not assigned at all, this still represents an improved outcome for these 152 personnel.

Preference Rank	1	2	3	4	5	6 or lower
No. of Personnel	11	7	2	4	9	119

Table 44 Preference Rank of Personnel Assigned at Scenario 2 but not at Scenario 1

Having analyzed the effect of changing preference list lengths, the next issue to analyze is the effect of changing position attribute weights. One way to analyze how changing position attribute weights affects the positions is to look at the number of unit and Service attributes that are satisfied as the weights for these attributes change. For each scenario and unit weight setting, Table 45 and Table 46 show the number of assignments that satisfy the indicated number of unit attributes. By reading down the columns, these tables show that as the unit weight is decreased, there is a consistent

decrease in the number of assignments where all three unit attributes are satisfied, and an increase in the number of assignments where only one or two unit attributes are satisfied.

		Number of Unit Attributes Satisfied			
Scenario	Unit Weight	0	1	2	3
1	0.75	0	2	167	1,435
3	0.60	0	1	203	1,417
5	0.45	0	8	262	1,370
7	0.30	0	13	312	1,354

Table 45 Number of Assignments Satisfying Unit Attributes – L_{pers} = 10

		Number of Unit Attributes Satisfied			
Scenario	Unit Weight	0	1	2	3
2	0.75	0	5	185	1,526
4	0.60	0	6	214	1,507
6	0.45	0	16	299	1,424
8	0.30	0	25	337	1,397

Table 46 Number of Assignments Satisfying Unit Attributes – L_{pers} = 20

Table 47 and Table 48 display the related results for Service attribute weights, showing the number of assignments that satisfy the indicated number of Service attributes. As the weight on the Service attributes is increased (again reading down the columns), the number of assignments where zero or one attributes are satisfied decreases, while the number of assignments where two or three attributes are satisfied increases.

		Number of Service Attributes Satisfied			
Scenario	Service Weight	0	1	2	3
1	0.25	82	433	789	300
3	0.40	68	412	831	310
5	0.55	45	400	851	344
7	0.70	34	336	922	387

Table 47 Number of Assignments Satisfying Service Attributes – L_{pers} = 10

		Number of Service Attributes Satisfied			
Scenario	Service Weight	0	1	2	3
2	0.25	100	475	804	337
4	0.40	61	455	861	350
6	0.55	45	417	904	373
8	0.70	36	327	975	421

Table 48 Number of Assignments Satisfying Service Attributes – L_{pers} = 20

As identified in Table 45 to Table 48, there is a general bias towards unit attributes rather than Service attributes; compared to the small number of assignments where only one unit attribute is satisfied, a higher percentage of assignments satisfy zero or one Service attributes. There are two reasons for this. First, a veto rule was included such that personnel could not be assigned to positions if they do not satisfy the branch attribute. Therefore, at least one unit attribute must be satisfied in all assignments. Second, the rank attribute is the first attribute for all personnel; that is, while there is a semi-random ordering of the attributes that motivate the personnel in their preferences, rank is identified as every person's first attribute, and therefore has the highest weight. As a result, the rank of personnel and positions will be matched at the majority of assignments because there is a strong alignment in the preferences of personnel and positions with respect to this attribute. This reveals an important finding; if there is strong alignment between the preferences of the personnel and positions with respect to an attribute, even if the weight of that attribute is decreased, a high percentage of assignments will continue to satisfy the attribute.

Figure 64 and Figure 65 further demonstrate the consistently high percentage of assignments where rank is matched, showing that there is only minimal responsiveness to the rank attribute weight; the percentage of assignments that satisfy the rank attribute varies between 96% and 99% in the different scenarios. A relatively high percentage of assignments (82% to 90%) satisfy the skill attribute, but that attribute is more responsive to the changing weight. The high percentage of assignments where skill is matched is partly reflective of the experimental design; 15% of positions are not specific to a

particular skill, and in these cases, if a person meets the branch requirement (which the veto rule makes necessary), then a skill match is also recorded. Therefore, a skill match is automatic for 15% of assignments.

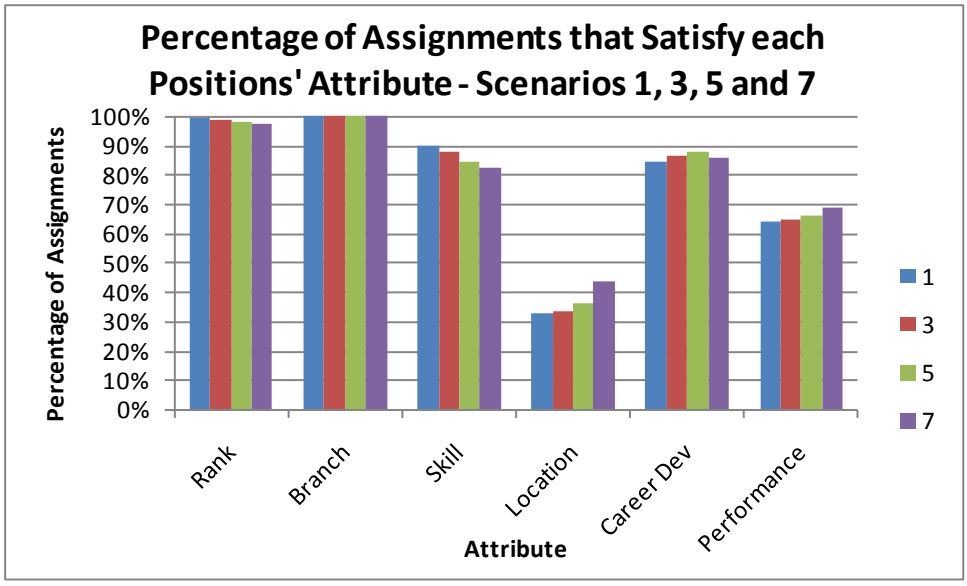

Figure 64. Percentage of Assignments that Satisfy each Positions' Attribute – Scenarios 1,3, 5 and 7

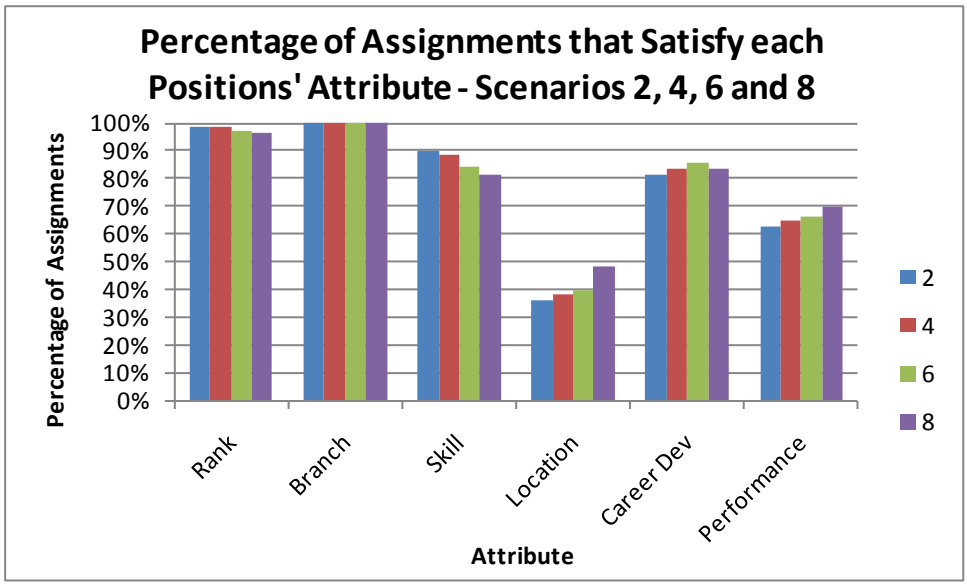

Figure 65. Percentage of Assignments that Satisfy each Positions' Attribute – Scenarios 2, 4, 6 and 8

As described, the high percentage of assignments where rank is matched is the result of the correlation in the personnel and positions' preferences. In contrast, the lower percentage of assignments where locations match (that is, where no geographic move is required) is the result of lower correlation between the personnel and positions' preferences. Location is the one attribute where the model of homophilous attraction is not used; personnel seek locations based on the desirability of locations as determined from the Electronic Preference and Restrictions form. These results demonstrate that where the preferences of personnel and positions are not correlated with respect to an attribute, a lower percentage of assignments are likely to satisfy that attribute. However, the results also demonstrate that this attribute is the most responsive to changes in the attribute weights. The percentage of assignments where geographic stability is achieved is 33% in Scenario 1 and 36% in Scenario 2, the two scenarios where location has a weight of 0.1. However, the geographic stability increases to 44% and 48% for Scenarios 7 and 8, where the location weight is 0.3. None of the other attributes exhibit such responsiveness to the attribute weight.

A final way of analyzing the assignment results is to look at the total number of attributes that were satisfied at each scenario. Table 49 identifies the numbers of assignments (and percentage in parentheses) that satisfy the indicated number of position attributes. Across all scenarios, at least 90% of assignments satisfied four or more of the positions attributes.

Number of Attributes Satisfied	Scenario							
	1	2	3	4	5	6	7	8
6	289 (18%)	319 (19%)	295 (18%)	327 (19%)	316 (19%)	332 (19%)	354 (21%)	368 (21%)
5	707 (44%)	718 (42%)	727 (45%)	762 (44%)	716 (44%)	759 (44%)	744 (44%)	797 (45%)
4	468 (29%)	520 (30%)	468 (29%)	503 (29%)	484 (30%)	510 (29%)	462 (28%)	465 (26%)
3	131 (8%)	145 (8%)	123 (8%)	128 (7%)	116 (7%)	127 (7%)	110 (7%)	119 (7%)
2	9 (<1%)	14 (<1%)	8 (<1%)	7 (<1%)	8 (<1%)	11 (<1%)	9 (<1%)	10 (<1%)
1	0	0	0	0	0	0	0	0
Total Assignments	1,604	1,716	1,621	1,727	1,640	1,739	1,679	1,759

Table 49 Number of Assignments that Satisfy Specified Number of Position Attributes

The results that have been presented here demonstrate that position preferences can be constructed with multi-attribute utility functions, and that assignments are responsive to the attribute weights. The strength of correlation between the personnel and positions' preferences influences the responsiveness of assignments to attribute weights and the percentage of assignments where attributes are matched. In situations where preference lists are strict, increasing the personnel preference list length cannot improve the outcomes of the personnel or worsen the outcomes of the positions. However, where preference list indifference exists and ties are randomly broken, although increasing the personnel preference list length will worsen the outcomes for most personnel, some personnel can benefit due to the random tie breaking.

C. OTHER IMPLEMENTATION CONSIDERATIONS

1. Two Assignment Rounds

The scenario results show that, depending on the position utility weights used, between 90% and 98% of personnel could be assigned using the two-sided matching algorithm, noting that in these scenarios there are more positions available for assignment than there are personnel (1,790 personnel and 1,997 positions). Results from the last chapter show that the percentage of personnel assigned would decrease if some individual personnel choose to submit preference lists that are shorter than the 10 or 20 used in these scenarios. The personnel who are not assigned at the end of the two-sided match could either pass into a second round match using new preferences based on the remaining personnel and positions, or a manual assignment process could be undertaken.

Two alternative approaches are discussed here to indicate other possible means for implementation, with each approach based on two assignment rounds. One approach would include only high performance positions (say, the top two quartiles) in the first round of assignments, with the second round including the lower performance positions and all personnel and positions that remain unmatched from the first round; the rational for such an approach is that the high performance positions are probably the most important and therefore the decision maker would wish to maximize the number of these positions that are filled. Another approach, similar to the first, would include only the high quality personnel in the first round of assignments; the rationale for this approach is that high quality personnel are the most important to the organization and the most important to retain, and therefore it is best to maximize the number of high performance personnel who receive a first round assignment.

Analyzing the results from Scenario 2 shows that 873 personnel from the top two performance quartiles were assigned, representing 96% of the 910 personnel in the top two quartiles (it might have been assumed that the number of personnel in the top two quartiles would be 895, half of the 1,790 personnel who require assignment. However, the allocation of personnel and positions to the performance quartiles occurred prior to randomly reducing the population to the one third who require assignment, and therefore

the number of personnel and positions in the top two quartiles is 910, not 895). Further, 886 positions from the top two quartiles were assigned, representing 87% of the 1,024 positions in the top two quartiles.

The data from Scenario 2 is used to demonstrate how a two round process may work (Scenario 2 was arbitrarily chosen for comparison). The first round involves all 1,790 personnel and the 1,024 positions from the top two performance quartiles. The personnel and positions' preferences were determined using the same utility functions and weights as previously used in Scenario 2, but excluding positions from the bottom two performance quartiles in the first round. A first round assignment was created based on these preferences, and this resulted in all 1,024 positions being assigned. The 766 remaining personnel (the balance of the 1,790) progressed to the second assignment round with the remaining 973 positions (the bottom two performance quartile positions). Preference utility functions were recalculated based on the personnel and positions that were participating in the second round of assignments. The second round of assignments resulted in a further 737 assignments.

A total of 1,761 assignments were achieved in the two assignment rounds, representing 98% of the 1,790 personnel requiring assignment. All positions from the top two performance quartiles were assigned in the two round processes, which compares to 87% that were assigned in Scenario 2. A fair comparison cannot be made of the personnel preference outcomes for the two processes because preferences were recalculated in the two-round process. However, it is possible to make a comparison of the two processes in terms of position outcomes because the position requirements remain constant. Figure 66 shows the percentage of positions' attributes that were matched in the two-round process and compares this against Scenario 2.

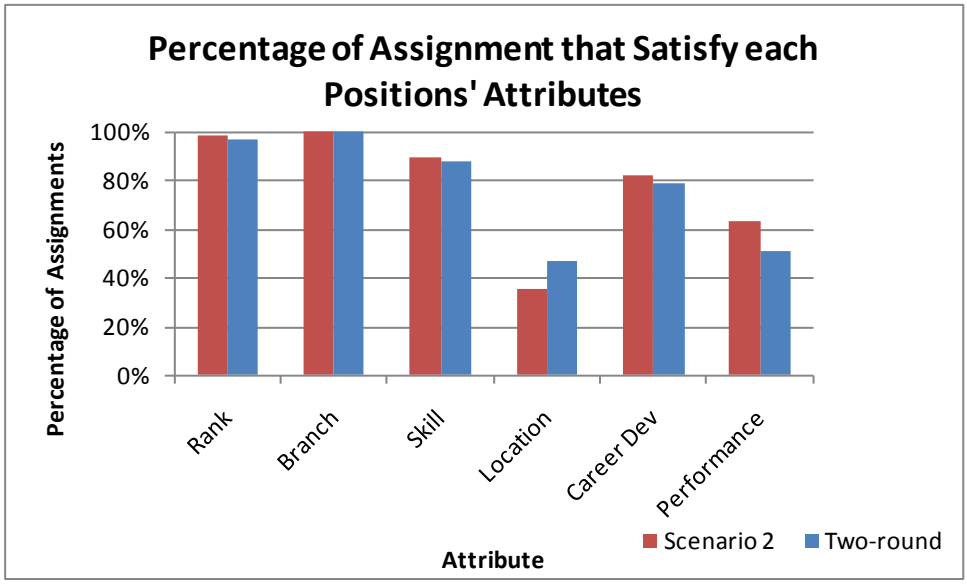

Figure 66. Percentage of Assignments that Satisfy each Positions' Attributes - Comparison of Scenario 2 and the Two-round Process

Figure 66 shows that the two-round process resulted in a lower percentage of assignments matching the performance attribute of personnel and positions. This is the result of the two round process, which segmented the assignment rounds by position performance quartiles, with only high performance positions available in the first round and only low performance positions available in the second round. At each of these rounds, the positions' performance weight was 0.05—the lowest of the six attributes. Due to the segmentation, there was an even spread of high and low quality personnel assigned to the high performance positions in the first round, and similarly an even spread of high and low quality personnel assigned to the low performance positions in the second round.

Simultaneously segmenting the personnel and positions by performance quartiles could be used in a two-round process to maximize the percentage of high performance personnel assigned to high performance positions in the first round. The second round would then include all remaining personnel and positions. However, any such manipulations would reduce the quality of assignments with respect to the other attributes.

The goal here was to demonstrate that more high performance positions could be assigned if a two round process is used, and the results have shown this; in Scenario 2

only 87% of the 1,024 high performance positions were assigned, and 100% of these positions were assigned in the two-round process. Depending on the situation, a decision maker may decide to segment positions and / or personnel in a two round process. Segmentation on different attributes would be possible.

While the results of each individual assignment round in a two round process would be stable in their own right, the overall assignments from a two round assignment process that excludes certain personnel or positions from the first round will not produce an overall stable result. That is, it could be possible to find situations where a person and position are assigned in the first round, but each would prefer to be assigned to another partner who is only available in the second round. This also leads to the possibility that participants might seek to manipulate their preferences in a two round process; for example, the preference list lengths that personnel submit in the first round may be shorter or longer depending on how they consider the prospects of being assigned to a position in the second round. The type of segmentation and the positions that are offered in each round might influence the number of preferences that personnel identify. Depending on the nature of the hierarchy, and the degree of subordination that exists, the extension or shortening of first round personnel preference lists may or may not be a problem. A decision maker would need to carefully weigh the advantages of a two round process against the potential disadvantages. Participant experiments would be useful to examine whether personnel would change their preferences if they know that the process is segmented into two rounds

2. Limiting Position Preference Lists Based on Utility

All scenarios conducted so far in this chapter involve position preference list lengths set at 50. However, a preference list of 50 may include personnel who satisfy only one or two attributes for some positions, and may therefore be poorly suited to the positions. Table 50 shows that 14 of the assigned positions only satisfied two of the six position attributes in Scenario 2, and 145 assigned positions satisfied three attributes.

Number of Attributes Satisfied	1	2	3	4	5	6
Number of Assignments	0	14	145	520	718	319

Table 50 Number of Assignments Satisfying Indicated Number of Attributes – Scenario 2

An option for increasing the quality of the positions' assignments, measured in this case by the number of position attributes satisfied at each assignment, is to limit preference list lengths based on the utility value rather than an arbitrary preference list length. For some positions this may result in a preference list length that is longer than 50, while for others this may result in a shorter preference list length.

A new scenario was created based on the same position attribute weights and the personnel preference list length from Scenario 2; for this new scenario, position preference list lengths were no longer limited to 50 preferences, but were restricted so that only personnel who provided a utility of greater than 0.7 were included in positions' preference lists. Personnel preference lists were then created, but still limited to 20 preferences. Assignments were then determined based on these updated preferences.

There were 1,716 assignments achieved in the original Scenario 2, however, limiting positions preferences in the manner described here reduced the number of assignments to 1,625. Although the number of assignments decreased, Table 51 shows that the positions' assignment quality improved; no assignments satisfied fewer than three attributes, and the number of assignments where three attributes were satisfied decreased from 145 to 110. The number of assignments where all six attributes were satisfied increased.

Number of Attributes Satisfied	1	2	3	4	5	6
Number of Assignments	0	0	110	491	688	336

Table 51 Number of Assignments Satisfying Indicated Number of Attributes – Modified Scenario 2

The results here demonstrate that decision makers may wish to consider limiting positions' preferences based on the utility that personnel provide to the positions, rather than a fixed preference list length. There will inevitably be a tradeoff between the quality and number of assignments.

3. Other Considerations

A decision support system is unlikely to be able to consider all of the complexities of military assignments; therefore, the assignments produced by the two-sided matching should be considered as a good starting point, and should provide a good solution for a large percentage of the personnel and positions involved.

Decision makers in military assignment applications have the authority to assign personnel and positions as required. The subordination in military applications is such that personnel must abide by the assignments that are provided if they wish to continue to serve in the military; there is no opportunity to "opt out" and be assigned through alternate means. Therefore, the decision maker can accept the two-sided matching results as is, or may choose to make changes as desired. While stability in a military personnel assignment situation is not a requirement, it is a desirable feature.

If a decision maker rejects an assignment that is produced by the two-sided matching, this might indicate that the position preference lists were incorrectly determined. For example, assume that there is a person, William, who is assigned to position 123, but the decision maker chooses to manually change the assignments so that James is now assigned to position 123. A couple of possibilities exist. If James prefers position 123 to his previous assignment, then he would have already "proposed" to position 123, and at some point position 123 rejected James in favor of William; this implies that the manual reassignment is the result of incorrect position preference lists, and the preferences should have been ranked so that James was ahead of William. Alternatively, if James does not prefer position 123 to his previous assignment, then changing the assignments would indicate that an unstable situation is now created.

The assignment process used in these experiments has been the person optimal algorithm whereby the personnel "propose" to the positions. The advantage of the person

proposing algorithm is that it is a dominant strategy for personnel to submit their true preferences. Although it has not been explored here, the decision maker may choose to consider implementing the position proposing algorithm. While this would no longer make it a dominant strategy for personnel to submit truthful preferences, the amount of information that personnel would need to manipulate the preferences is high (Gale & Sotomayor, 1985; Roth & Rothblum, 1999). Therefore, it is unlikely that personnel would be able to manipulate their preferences to obtain a preferred assignment. Consequently, a decision maker might choose to apply and compare the results of the personnel and position proposing algorithms.

D. SUMMARY

Assignments in hierarchical organizations may require consideration of attributes that are of interest to the organization and its subordinate elements. Such a situation already exists in some school matching processes, and it is also relevant if applied to military organizations. This chapter has demonstrated that multi-attribute utility functions can usefully generate preferences in these situations. It has also demonstrated that the two-sided matching assignments are responsive to changing weights in the multi-attribute utility function. This allows a decision maker to adjust the weights so that the assignments provide the best possible consideration of all attributes.

This research suggests that assignment processes in large hierarchical organizations, in particular militaries, could benefit from decision support. Two sided matching provides a sound mechanism for generating assignments and has desirable properties; in particular, stable outcomes and truthful revelation of preferences. The position preferences, as required for the two-sided matching, can be determined from a multi-attribute utility function that combines unit and Service attributes. A system built on such processes would allow a decision maker to examine a wide range of assignment options in a relatively short time period. The decision maker would retain the ability to accept or reject any assignments, and such a system would provide a range of options for the decision maker to consider. Further, the quality of assignments would be less variable and dependent on the decision maker's skills.

THIS PAGE INTENTIONALLY LEFT BLANK

VII. SUMMARY, CONTRIBUTIONS, AND FURTHER RESEARCH

A. CONTRIBUTIONS

1. Human Decision Making

Eppler and Mengis (2004) reviewed a considerable number of studies that examined information overload. The existence of information overload is in little doubt; however, previous research has not examined information overload in an assignment decision-making context. Prior studies on human decision making have examined scenarios where the decision maker is presented with a series of choices and must select the choice that best satisfies certain attributes; for example, selection of the best laundry detergent (Jacoby et al., 1974) or the person most likely to become bankrupt (Casey, 1980). However, this research examined assignment decision-making scenarios, where the decision maker must assign (match) personnel and positions based on personnel preferences and position attribute requirements. These scenarios may be considered to be more complex because each assignment made by the decision maker limits the remaining assignment options available. The results of these experiments provide insights into the ways that decision makers operate and the effect of information overload when there are two groups whose needs must be considered.

This research demonstrates the difficulty that decision makers experience when making assignments; there is a requirement to balance the desires of two different groups, which in this research are the personnel and positions. The experiments conducted here demonstrate that most decision makers fail to achieve an appropriate balance for the two groups, with decision makers tending to focus on one group to the detriment of the other. In a military assignment context or other similar situations, there may be a tendency to focus on the position requirements to the detriment of the personnel; this could have follow-on consequences for retention if personnel satisfaction decreases.

There is evidence from the results that decision makers focused on the position requirements to the detriment of personnel preferences. Evidence of this is that the decision makers achieved an average of 93% of the optimal position solution, yet only

70% of the optimal personnel solution. Further, while high-performing decision makers produced assignments that appropriately balanced the personnel and positions' preferences, the lower performing decision makers produced results where the assignment quality (measured as a percentage of the optimal solution) was far lower for the personnel than the positions; this indicated that decreases in quality were more likely to be experienced by the personnel than the positions. Finally, the variability in personnel points was greater than the variability in position points, indicating that the decision makers focused on position requirements more than the personnel.

The decision makers were from a relatively homogenous background: they were all military officers, indicating that they have all met the aptitude test requirements for entry into the military; they had all completed a Bachelors level degree and were undertaking master's-level coursework. Despite this level of decision maker homogeneity, the results show considerable variability, even for the scenarios that had the least number of attributes and fewest assignments to make. Decision makers from a less homogenous background are likely to exhibit greater variability.

Some decision makers developed a simple heuristic whereby they assigned agents sequentially as presented in the experiments, rather than considering all holistically. This heuristic leads to a sub-optimal solution because the decision maker only considers a sub-set of the agents, and therefore does not consider alternative assignment options that might improve the overall outcome. These results indicate that, in the absence of decision support, assignment quality for some agents may be influenced by how their information is presented to the decision maker. Decision makers need to be aware of this limitation and ensure that the processes they use provide the best opportunity for achieving high quality assignments.

This research indicates that decision makers in assignment situations are likely to over-estimate their abilities. Based on the decision makers' self-assessments, most believe that their performance relative to their peers is better than it actually is. Arnhart (2007) identified the difficulty that can be expected in getting career managers to accept decision support systems, a situation that may partly be the result of decision makers' over-confidence. The results of these experiments help highlight to decision makers the

difficulty of making assignments without decision support, and that although they may feel confident in their own abilities, a decision support system would improve consistency, consider a broader range of factors, and better balance personnel and position needs.

2. The Effects of Indifference on Two Sided Matching Assignments

When two-sided matching preferences are strict, the set of agents who are unmatched is the same at every stable matching (Roth & Sotomayor, 1990). However, when preferences are not strict, it is possible for an agent to be assigned at one stable outcome and unassigned at another. Therefore, on the basis that agents only list preferences that are considered acceptable and would prefer to be assigned to one of their preferences than remain unassigned, not all agents will agree on which tie breaking produces the best result. While Roth and Sotomayor (1990) provide the theoretical basis for the effect of tied preferences, no empirical results are available to demonstrate how assignment outcomes vary with preference list indifference. This research examined how the number and quality of assignments varies as the conditions vary.

Some two-sided matching applications, such as the NRMP, rely upon voluntary participation, and in such applications it is important that preferences are strict so that the assignments are determined directly by the preferences; that is, there is a single assignment set and the administrator has no influence on the assignments. However, as identified in this research, there are some applications, such as assigning students to schools and military personnel to positions, where participation is not voluntary and subordination exists. Preference list indifference could be permitted in such applications because different assignment outcomes will not lead to decreased participation. This research demonstrates that allowing participants to express indifference is advantageous because it provides the central authority with options and allows them to select the tie breaking that produces the best assignments from a holistic perspective. When preference list indifference exists and ties are randomly broken, the two-sided matching results remain weakly stable, meaning that there is no couple who each strictly prefer the other to their assigned partner.

Computational experiments were used to assign agents using the instability chain algorithm (Roth & Vande Vate, 1990). Agents were randomly created according to defined parameters: the number of agents (n), the maximum preference list length (L), the preference correlation (C), and the degree of indifference in the agents' preferences (DOI). The experiments were designed to analyze how different tie breakings influence assignments.

The experiments involved a variety of scenarios, with each scenario defined by a setting of the parameters n, L, C and DOI. Multiple trials were conducted for each scenario; the tied preferences were randomly reordered in each trial and assignments were produced. Two variables of interest were recorded; the number of agents assigned and the assignment quality (measured by the preference rank of assigned agents).

The number of agents who are assigned from a scenario depends on how preference list ties are broken; different tie breakings result in different numbers of agents being assigned. Consequently, there is a range in the number of assigned agents, where the range is the difference between the highest and lowest number of assignments achieved from the many trials. There is a low probability that the number of assignments at any single trial will be at the lower or upper end of this range, with most trials producing a number of assignments near the midpoint of this range. As the number of agents and the number of trials increases, the frequency distribution of the number of assigned agents resembles a binomial distribution.

When analyzing the assignment quality and the number of assigned agents that result from different tie breakings, some assignment sets dominate others. The dominant assignments are those where it is not possible to increase the number of assigned agents without a decrease in assignment quality, and it is not possible to increase assignment quality without a decrease in the number of assigned agents. A possibility frontier can be created to show the dominant outcomes from which a decision maker could select.

The effect of the number of agents and the degree of indifference was examined through a series of scenarios. As the number of agents in an assignment context increases,

the range in the number of assignments increases. However, when expressed as a percentage of the number of agents, the range in the number of assignments decreases as the size of the assignment problem increases.

Increasing the degree of indifference increases the range in the number of assignments, and this occurs regardless of the number of agents in the scenario, the preference list length or the preference correlation. Further, as the degree of indifference increases, the average preference rank of assigned agents decreases. This is the result of agents ranking more preferences indifferently; for example, a person with strict preferences might be assigned to their fourth preference, whereas an agent who is indifferent between their fourth, fifth and sixth preferences could be assigned to any of these and still be considered to have been assigned to their fourth preference.

If an assignment decision support system is to be implemented, it is necessary to consider the preference list length that participants should be allowed to express. A series of scenarios were used to analyze how preference list length affects the number of assignments and the assignment quality. Roth and Sotomayor (1990) identify that increasing preference list lengths will lead to more assignments. While increasing preference list length will increase the number of agents assigned, the results obtained here show that decreasing returns are experienced; for example, the increase in the number of assigned agents as the preference list length increases from 10 to 15 is less than the increase in the number of assigned agents when the preference list length increases from 5 to 10.

The preference correlation defines how similar the personnel and position agents' quality scores need to be in order for them to rank one another. When preferences are highly correlated, agents from one group will only rank agents from the other group if they are quite similar in quality scores. There is a relationship between the preference correlation parameter and the number of preferences that an agent identifies; as preferences become less correlated more agents are considered acceptable because they are within the broader quality range identified by the preference correlation parameter. As preferences become less correlated, the number of assignments increases. The effect is similar to that identified from an increasing preference list length.

Agent quality scores were randomly allocated, and therefore the number of position (or personnel) agents that fall within each person's (or position's) preference correlation range varies. Consequently, agents' preference lists will be of varying lengths when preferences are limited by the preference correlation parameter *C*, provided that the parameter L is sufficiently long to not otherwise limit the preferences. Using preferences based on the preference correlation parameter *C*, it was possible to analyze the relationship between the number of preferences that an agent identifies and the probability of being assigned. The results confirm that agents with longer preference lists have a higher probability of being assigned. These results align with the NRMP and has demonstrated that matched applicants and residency programs have longer preference lists than unmatched applicants and programs. Further, the NRMP advises applicants and programs to include all acceptable choices in their preferences because it increases the probability of being matched without affecting the chances of being matched to a higher ranked preference.

This research has demonstrated how preference list indifference will affect two-sided matching results. It is likely that there will be considerable preference list indifference for agents in large hierarchical organizations where limited attributes are available to differentiate agents, and where a limited number of discrete values are used to measure these attributes. These characteristics typify assignment applications for school students and military personnel. If preference list ties are allowed in such applications, decision makers could examine the dominant outcomes to select the assignment set that best achieves the decision makers' priorities.

3. Using Two Sided Matching With Preferences Determined from a Multi-Attribute Utility Function

The assignment of students to high schools is an example of two-sided matching in a hierarchical organization. The hierarchical nature and subordination is demonstrated in the New York City high school match by the Department of Education imposing constraints on the schools' preferences; Abdulkadiroğlu et al. (2005) identify that the New York City Department of Education requires schools to accept a distribution of

students based on standardized English test scores. In addition, schools have different attributes that they use to construct their preference lists, such as GPA, attendance and disciplinary records.

Two sided matching requires rank ordered preferences, and this research has demonstrated that multi attribute utility functions can be used to generate these preferences in hierarchical organizations. Such a process is suited to military personnel assignment processes where each position's preferences need to incorporate the requirements of the Service and the subordinate units.

Computational experiments were used to demonstrate how the two-sided matching results are affected by settings such as personnel preference list length and the position multi attribute utility functions. The context for these experiments was assigning military personnel to positions, using Australian Army data. Actual data was used wherever possible, but it was necessary to use simulated data in some cases where actual data was not available.

The position and personnel preferences were constructed using multi attribute utility functions. Although participating personnel would be responsible for determining their preferences in an actual application, actual personnel preference data was not available and so the personnel preferences were also determined using multi attribute utility functions. The preferences revealed a high degree of indifference; 0.75 for positions and 0.35 for personnel. The degree of indifference in position preference lists is related to the number of attributes used to identify the position requirements, the scale of the attributes, and the weights applied to those attributes. For example, as the number of attributes increases it is possible to achieve greater differentiation that will decrease the degree of indifference. Without actual personnel preference data it is unclear what degree of indifference would exist in the personnel preference lists.

As with the experiments that investigated preference list indifference, multiple trials were conducted on each scenario in these military personnel assignment experiments. The results of the military personnel assignment experiments showed that

the number of personnel assigned and the quality of assignments were distributed in a similar manner to the results of the preference list indifference experiments.

The experiments examined how assignments would vary if the length of personnel preference lists varied while holding position preference list length constant. If preferences were strict, Roth and Sotomayor (1990) identify that increasing the personnel preference list length could not harm the positions and could not benefit the personnel. However, preference list indifference exists in these experiments, and the theory for strict preferences does not hold. Due to the different results possible from different tie breakings, some personnel were observed to achieve improved assignments when personnel preference list lengths were extended, although a larger number of personnel were observed to have worse assignments in scenarios with longer personnel preference list lengths. The increased preference list lengths led to more personnel being assigned.

The experiments demonstrated that the assignments produced by two-sided matching are responsive to the multi attribute function weights. That is, as the relative weight of an attribute is increased, there is an increase in the percentage of assigned positions that have that attribute matched. As attribute weights are changed, the responsiveness of assignments can be measured by analyzing the change in the percentage of assigned positions the attribute satisfied. The responsiveness of assignments to changes in an attribute's weight was found to be related to the preference correlation between the personnel and positions with respect to that attribute.

Personnel and position preferences are highly correlated with respect to an attribute if the attribute is considered to be important to both the personnel and positions, and both groups agree on the desired attribute values. In these experiments, the rank attribute was highly correlated because both personnel and positions structure their preferences based on the rank attribute (for personnel it was the first consideration), and the personnel and positions agree on the attribute values (SGT positions primarily seek SGT personnel, and SGT personnel primarily seek SGT positions). In these experiments, a low correlation existed in the location preferences because there was a lack of alignment in what each group sought; for example, a position in Sydney would seek

personnel in Sydney (because this minimizes the number of relocations), but a person in Sydney may seek positions in other locations, depending on the desirability of the locations.

Changing the multi attribute utility function weights will lead to a change in the percentage of assignments where the attribute is satisfied, regardless of the strength of preference correlation. However, the responsiveness of assignments to an attribute's weight is affected by the preference correlation with respect to that weight. Assignments are more responsive to changes in an attribute's weight if the preference correlation for that attribute is low.

The research examined an option for conducting two assignment rounds, a process that might be considered during implementation. A process that involves two assignment rounds would increase the overall percentage of participants assigned. Further, if a position or personnel target demographic is identified, it is possible to increase the percentage of agents in that target demographic that are assigned. For example, the target demographic could be high quality personnel, and if the personnel are segmented so that only high quality personnel are included in the first round (all positions remain available), this process would maximize the percentage of high quality personnel who are assigned. Segmenting the process so that some personnel are excluded in the first round would not change the personnel's dominant preference strategy; all positions would remain available and so the personnel who participate in the first round would still submit their true preferences. Alternatively, the assignment process could be segmented so that some position agents are excluded from the first round; however, doing so could change the personnel's dominant preference strategy, because some personnel might prefer a position that is not available until the second round, and they may therefore choose to limit or manipulate their first round preferences.

B. FURTHER RESEARCH

The results of the decision-making experiments indicate that some subjects developed a heuristic that led to sub-optimal results; analyzing the results identified that for some subjects the agents' assignment quality was related to the agents' placement in

the experiment worksheets. Agents who were presented to the left of the worksheet were more likely to receive better quality assignments than those presented on the right of the worksheet. If true, this indicates that the order in which information is presented to a decision maker can influence the results.

The decision-making experiments were not designed to test the heuristics that decision makers use to make their assignments. Therefore, further research would be required to test the hypothesis that the ordering of information influences the assignment results. This could be achieved through a minor modification to the experiment design. The same personnel and position details would be used, but the order in which the personnel and position details appears would be reversed. Ideally the experiments would be repeated with two groups of subjects; the first group would receive the experiment worksheets in the original order and the second group would receive the worksheets with the reverse order. Subjects would be randomly allocated to each group and would do the experiments concurrently.

There is little evidence to suggest that decision makers in military personnel assignment applications are selected based on their assignment decision-making ability. If a simple standardized test could predict a person's likely performance on assignment decision-making tasks, this would be a valuable screening tool to ensure that decision makers have the appropriate abilities for their task. Henry (1980) examined the effect of an individual's information processing ability on their decision-making ability in the consumer behavior domain. A similar approach could be examined in the assignment decision-making context; a test of information processing ability would be developed, subjects would be administered the test and the assignment problems, and results of the two analyzed to determine the test's predictive ability.

The military assignments demonstrated in this research were based on actual personnel and position data to the maximum extent possible. However, not all desired data was available; in particular, no personnel preference data was available because this information is not captured in a consistent manner that would enable rank ordered preference lists to be used. Further research using actual personnel and position preference data is recommended. As an initial study this could involve a relatively small

group of personnel and positions (perhaps around 200) where the assignments would be contained within the group; that is, personnel in the group would only be assigned to positions in the group, and positions in the group would only be filled by personnel in the group. For a small sized group, it would be possible to work with the career manager to capture more detailed information on preferences. The preference data would allow two-sided matching assignments to be produced from actual preferences. Comparing the assignments produced by career managers against those produced by the two-sided matching would provide considerable information for analysis; this may provide greater support for implementing a decision support system and would provide further insights into the performance of actual decision makers in an assignment application.

The computational experiments used the person "proposing" instability chain algorithm (Roth $ Vande Vate, 1990); this produces the same results as the Gale—Shapley Deferred Acceptance algorithm when there are no partners or complexities. Further research on the difference in results between the person proposing and the position proposing algorithms would be useful. This would identify how much it is possible to shift between assignments that favor the personnel and assignments that favor the positions. If the experiments involve human decision makers in an experimental setting, it would also test the propensity for participants to try gaming the system.

It is not a requirement for all medical students in the National Resident Matching Program to be matched, but those students who are not matched enter a period known as the Scramble; during the Scramble the students contact medical residency programs to enter into an agreement. In school assignment applications, it is necessary that all students are assigned; Abdulkadiroğlu et al. (2005) identify that students who are unmatched from the two-sided matching assignments are manually allocated to schools by the education department. An option for increasing the percentage of assigned participants is to conduct two assignment rounds. An initial round involving all personnel and position agents could be conducted, with the unmatched participants progressing to a second round, with preferences updated and resubmitted for the second round; such a process should not change the dominant preference strategy.

Alternative approaches involving the segmentation of participants in each round were discussed. The use of two assignment rounds could change the dominant strategy for personnel, particularly if the process is segmented by excluding a certain group of positions from the first round. If a group of positions are excluded from the first round (for example, only high profile positions are available for assignment in the first round), some participants might withhold preferences until the second round. Participant experiments would be useful for identifying the change in personnel behavior in a two round match. The participants would act as the personnel who require assignment in these experiments, and the experiments would examine whether the use of one or two assignment rounds would cause participants to vary the preference lists they submit. The experiments would also segment the positions to determine how personnel would behave in such situations. While conducting two assignment rounds could have implications for stability, particularly if the participants are segmented, this might be acceptable in hierarchical organizations where subordination is greatest.

C. SUMMARY AND CONCLUSIONS

This research has developed and demonstrated a process that uses two-sided matching for assignments in hierarchical organizations. Suitable applications include assigning personnel to positions within militaries and assigning students to schools. In the military context, the personnel have preferences for the positions to which they wish to be assigned, and it has been demonstrated that the positions' preference lists can be produced from utility functions that combine organizational and unit attributes.

Liang and Buclatin (1988) and Klingman and Phillips (1984) identify that military assignment processes are large-scale complex problems; there are many personnel and positions to be assigned and each assignment requires considering many attributes and policies. Despite the scale and complexity of such problems, Gates and Nissen (2002) identify that many hierarchical organizations, such as militaries and government organizations, rely upon the cognitive process of centralized, administrative staff to assign personnel to positions. Eppler and Mengis (2004) reviewed research across a variety of disciplines and demonstrated that human decision making is poor under

excessive information loads, although no research specifically relating to assignment decision making was identified. Despite an understanding of the effect of excessive information loads in decision making, there is a resistance to using decision support systems in military assignment contexts; Arnhart (2007) identifies the difficulty that can be expected in gaining acceptance from career managers.

The first part of this research utilizes participant experiments to examine decision making in an assignment context, showing how assignment quality varies in relation to the size and complexity of the assignment problems. The experiments involved human subjects performing the role of decision makers, with induced value theory, as outlined by Smith (1976), used to ensure the validity of the experiments. The experimental scenarios were based on assignment decision making in a hierarchical organization, and therefore demonstrate external validity to the career managers who must be convinced of the benefits of a decision-support system.

The experiments demonstrate the need for decision support in the large-scale and complex process of assigning military personnel to positions. In the absence of such decision support, the quality of assignment decisions depends heavily on the decision maker's ability; some decision makers perform well, quickly balancing the preferences of personnel with the position requirements; most decision makers, however, develop simple heuristics that lead to suboptimal outcomes. Further, decision makers are likely to over-estimate their abilities; for example, responses completed after the experiments showed that only 13% of subjects believed that they would be in the bottom 50% of participants, and only 3% believed that they would be in the bottom quartile.

Having established the requirement for developing a decision support system in hierarchical organizations' assignment processes, some potential assignment processes were reviewed. Two sided matching was selected as an appropriate mechanism for two key reasons. First, it is a dominant strategy for personnel to truthfully reveal their preferences (assuming the personnel proposing algorithm is used); this removes the potential for participants to game the process by manipulating their preferences. Second, it produces a stable match, meaning that there is no person and position who would rather be assigned to each other than their assigned partners.

Two sided matching has been used in some large assignment applications, such as the National Resident Matching Program and New York City high schools match. However, these existing applications either require participants (hospitals and medical students, high schools and students) to submit strictly ordered preferences, or use a single tie breaking rule if tied preference are allowed; for example, each participant receives a single random lottery number that is used to break ties in all cases.

The second part of this research uses computational experimentation to demonstrate the range of outcomes that are possible when preference list ties are broken in different ways. This would enable a decision maker to select an outcome that delivers the best outcome in terms of the number of assigned participants and the quality of the assignments. In assignment situations, such as the National Resident Matching Program, where participants are voluntary and assignments are binding, it is important that the administrator of the assignment process act without bias and produce a single outcome, otherwise the process may unravel. However, decision makers in hierarchical organizations are not constrained in the same way because they have responsibility for the organization's interests, the organizational units and participating personnel are subordinate, and participation is not voluntary.

In applications, such as the New York City high schools match, schools develop their own preference lists of the participating students, but must do so in accordance with conditions imposed by the department of education. The third part of this research demonstrates that a multi-attribute utility function can be used to produce preferences that incorporate the needs of the organization and the organization's subordinate units. The attribute weights can be adjusted so that the assignments achieve the appropriate balance between the attributes.

The processes outlined in this research will improve the consistency of assignments and enable a variety of options to be produced in a short time frame. By changing the utility function weights, the decision maker can compare a variety of assignments that meet different objectives. These will provide the decision maker with a good starting point for consideration.

Respondents to the Australian Defence Organization's annual Defence Attitude Survey indicate that more needs to be done with respect to personnel assignment processes; 70% of Army personnel agreed that "Individual posting preferences need to have more influence," while 50% agreed that "My views are considered when postings are planned."[2] The decision-making experiments demonstrated that most subjects could have assigned more personnel to positions that they preferred without overall detriment to the positions. If a decision support system of the type that is outlined in this research is implemented, this could improve the satisfaction of the personnel being assigned and in turn lead to improved retention rates. If this is achieved, then the effort will have been worth it.

[2] Australian Defence Organization, 2009 Defence Attitude Survey.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A

A. ASSIGNMENT EXPERIMENT—INSTRUCTIONS

This is an experiment that examines decision-making performance. You will have the opportunity to earn monetary prizes in these experiments. How well you perform increases your chances of earning these prizes. Your decisions will be made individually and do not require interaction with others.

I will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If you have problems during the experiment, raise your hand and someone will come and assist you.

Please turn to the pages with 1a and 1b marked on the top right corner, and place these pages side by side on your desk. These pages are an example of the information you will soon receive for the experimental work, and they contain information regarding four positions and three personnel who require assigning. Your job is to assign the personnel so that each person is assigned to one and only one position. Each position may be assigned only one person. All personnel must be assigned, and consequently one position will be left vacant.

Please focus now on page 1a.—this page has all the data pertaining to the personnel. Look first at the top table on the page. Across the top of the table you will see that there are three personnel indicated by the identifiers 101, 102 and 103, with data pertaining to each person in the column below them. On the left hand side of the table you will see there are five characteristics indicated, and each person is measured according to these five characteristics. Therefore, Person 101 has a Rank of O-3, belongs to the Armor Branch, has attained primary and secondary skills of D and T respectively, and is currently located in the West. The skill names, for example D and T, are arbitrary, and the letters have no meaning other than indicating distinguishable types of skills.

In the next table down on the same page, you will see the same three personnel listed. This table indicates the positions that each person would like to be assigned to. Person 101's first preference is to be assigned to position number 3, has a second preference to be assigned to position number 1 and has a third preference to be assigned to position number 4.

The last table on this page is where you will record the assignments that you make. During the actual experiments you will be required to write in the position number that each person is to be assigned to.

Now look at page 1b. Looking across the top of the table, you will see that there are four positions indicated, with data pertaining to each position in the column below it. This table is used to indicate what type of person each position seeks. On the left hand side of the page you will see there are five criteria used to indicate the ideal type of person for a position. In the actual experiments, you will sometimes have five criteria to consider, and sometimes have three criteria. Weights are used to indicate the relative importance of each of the criteria. A higher weight indicates that a criterion is more important.

Each position criterion is measured by four different types. For example, Rank can take the values of O-2, O-3, O-4 and O-5. At the intersection of a row criterion type and a column position number, a value is provided to indicate the utility that the position attains from that criterion value. For example, where the row for Rank O-3 intersects with the column for Position 1, a value of 1.0 is recorded. A value of 1.0 is the highest that can be recorded for any of the criterion values; therefore this indicates that with respect to rank, Position 1 ideally seeks a person who has the rank O-3. For the same position, the value of 0.5 in the row for Rank O-2 indicates that the position is less willing to accept a person of rank O-2, however prefers such a person to say an O-5 for which no utility value is recorded.

Each of the four other criteria works the same; a value indicates the extent to which the position seeks a person of the specified type. Ideally, Position 1 would seek to be assigned with a person that has the following characteristics; Rank of O-3, Branch of

Armor, Primary Skill of type C, Secondary Skill of type S and someone who is located in the North area. However, tradeoffs may be necessary because there may be no person with such characteristics, or other positions may compete for the same criteria values. The overall measure of how well a person suits a position is obtained by multiplying each criterion value by the criterion weight, and then summing these values. For example, if Person 101 was to be assigned to position 1, then the overall measure would be the sum of the Rank value (0.267 x 1.0), the Branch value (0.267 x 1.0), the Primary Skill value (0.267 x 0), the Secondary Skill value (0.1 x 0) and the Location value (0.1 x 0), giving a total utility of 0.53. You may find that assigning a different person to position 1 would give a higher utility value, indicating that this other person is better suited to position 1. However some tradeoffs will be necessary.

You will be presented with six different assignment markets similar to the example you have in front of you. Some of the assignment markets will have more personnel and positions than others. In each case, there will be more positions available than there are personnel to assign. Remember to assign every person – and that some positions will remain vacant. Your goal is to seek the best overall assignment outcome that you can achieve. This will be measured in two parts: The first part is how well positions get the type of person that they seek; and the second part is how well the personnel get their preferred positions.

You will receive points for each of the six markets. How many points you receive depends on the quality of the assignments you make. Once all the results have been collected, the total number of points earned by all participants will be determined. You will own a share of the total points based upon your performance. The better you do, the more points you own and the higher share you have. Everyone's points will be converted to virtual raffles tickets. A raffle draw will be undertaken in class next week, with first to fourth prizes being $40, $30, $20 and $10. Your chances of winning these prizes are determined by your performance and how many virtual raffle tickets you earn. Let me describe how you earn points.

Your points will be determined in two parts. I will use an example to demonstrate how you will earn your points. Please turn over both of your example papers so that sides

2a and 2b are showing. This is the same data you have already seen with a sample assignment set given. You will see on page 2a that person 101 has been assigned to position 1, person 102 to position 4 and so on. On page 2b, the characteristics of the personnel who are assigned to each position are circled in the columns under the position identifiers. This is to make it easy to identify the utility scores that are provided at the bottom of each column. You will see from the scores below the columns that position 2 gets a higher utility from the assignment than the other positions. The average of the utility scores for the three matched positions is 0.653. The first part of your points for this exercise would be 200 multiplied by this average utility, giving a total of 130.6 points earned. In each exercise, you will earn 200 points multiplied by the average utility that the positions receive. In addition, for each exercise you will also earn 100 points multiplied by the inverse of the average rank of the matched personnel. On page 2a, you will see that the average rank for the personnel is 1.67, and the inverse of this is multiplied by 100 to give 59.88 points. The inverse is used because it is preferable to get a lower average rank for the personnel. The total points earned for this exercise would be 190.48

You will note that not every person ranks every position as a preference. If you assign a person to a position that the person does not rank, the rank value to be used for calculation of your points corresponds to the number of positions in the experiment—so four in this example.

In the six experiments, you will have different numbers of personnel and positions to assign. Each of the six experiments will be strictly timed and you will be advised how long you have before commencing each experiment. You will be given warnings 60 seconds and 30 seconds before your time for each experiment is finished. You need to make sure each person is assigned before the time runs out. You may not look at or proceed to the next experiment until you have been told to do so. The total time for the experiments will be approximately 45 minutes.

Once the experiments have been completed, the papers will be collected. During the next week, I will analyze the results and determine your points. The raffle for the prizes will be conducted in class next week and at that time I will provide some preliminary indications of the overall results.

Are there any questions?

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B.

A. DISTRIBUTION OF PERSONNEL AND POSITIONS FOR SIMULATION

Personnel by Rank, Branch and Skill:

		Rank			
Branch	**Skill**	**E06**	**E08**	**E09**	**Total**
101		**144**	**132**	**33**	**309**
	201	98	106	1	205
	215	19	8	1	28
	221	19	10	1	30
	222			24	24
	256	7	6	4	17
	272	1	2	2	5
102		**206**	**143**	**69**	**418**
	203	74	54	24	152
	206	80	48	25	153
	211	36	31	13	80
	222			3	3
	237	3	2		5
	253	5	2	2	9
	257	2			2
	259	6	6	2	14
103		**102**	**88**	**36**	**226**
	207	30			30
	217	17			17
	218	14	10		24
	222			1	1
	228	4	45		49
	233	10	9		19
	234	11			11
	238	5	3	1	9
	246	3			3
	251			26	26
	258	3	19		22
	263	1			1
	266	2	1		3
	269	2	1		3
	271			7	7
	280			1	1
104		**62**	**41**	**17**	**120**
	202	38	27	3	68
	222			6	6
	231	5	2	1	8
	232	8	6	4	18

	235	7	3	2	12
	240	4	3	1	8
105		**56**	**38**	**17**	**111**
	204	18	15	1	34
	222			7	7
	230	5	3		8
	236	8	4	2	14
	239	4	2	1	7
	243	11	7	1	19
	260	2	1		3
	270	5	3		8
	273	3	2		5
	277			1	1
	278			4	4
	281		1		1
106		**60**	**40**	**17**	**117**
	208	26	22	4	52
	213	17	10	5	32
	216	8	4	2	14
	222			4	4
	226	8	4	2	14
	257	1			1
107		**41**	**30**	**6**	**77**
	205	28	22		50
	220	12	8		20
	222			6	6
	280	1			1
108		**47**	**37**	**12**	**96**
	212	11			11
	222			11	11
	223	9	5		14
	245	6	27	1	34
	247	4			4
	250	6			6
	252	5			5
	254	2			2
	268	4	5		9
109		**55**	**26**	**9**	**90**
	209	31	18	2	51
	222			4	4
	229	18	7	2	27
	257	1			1
	264	2		1	3
	274	2			2
	276	1	1		2
110		**30**	**17**	**5**	**52**
	214	24	14	4	42

	249	6	3	1	10
111		**17**	**12**	**4**	**33**
	222			2	2
	242	9	5	1	15
	244	3	2	1	6
	248	5	5		10
112		**21**	**16**	**3**	**40**
	219	21	16	2	39
	222			1	1
113		**22**	**14**	**5**	**41**
	222			2	2
	224	14	9	1	24
	262	8	5	2	15
114		**16**	**8**	**3**	**27**
	225	16	8	3	27
115		**10**	**6**	**2**	**18**
	227	10	6	2	18
116		**4**	**1**	**1**	**6**
	222			1	1
	255	3	1		4
	279	1			1
117		**4**	**3**	**2**	**9**
	261	4	3	2	9
Total		**897**	**652**	**241**	**1790**

Positions by Rank, Branch and Skill:

		Rank			
Branch	**Skill**	**E06**	**E08**	**E09**	**Total**
101		**149**	**93**	**14**	**256**
	201	106	58		164
	215	19	7		26
	221	22	7	1	30
	222			10	10
	Any	2	21	3	26
102		**234**	**145**	**60**	**439**
	203	63	39	19	121
	206	103	56	25	184
	211	46	33	12	91
	237	4	1	1	6
	253	4	2	1	7
	257	1			1
	259	8	8	2	18
	280	3	1		4
	Any	2	5		7
103		**116**	**91**	**35**	**242**

	207	30			30
	217	21			21
	218	17	11		28
	228		46		46
	233	12	8		20
	234	13			13
	238	4	2	1	7
	246	5			5
	251			27	27
	257	2			2
	258		20		20
	263	2			2
	266	3	1		4
	269	1			1
	271			7	7
	280	5			5
	Any	1	3		4
104		**66**	**37**	**11**	**114**
	202	33	17	3	53
	231	4	2	1	7
	232	11	6	4	21
	235	7	3	1	11
	240	4	2		6
	257	2			2
	280	4	2	1	7
	Any	1	5	1	7
105		**67**	**45**	**11**	**123**
	204	16	12		28
	222			4	4
	230	5	3		8
	236	10	5	2	17
	239	3	3		6
	243	11	4		15
	257	2			2
	260	2	1		3
	270	5	4		9
	273	4	3		7
	278			3	3
	280	6	3		9
	Any	3	7	2	12
106		**76**	**47**	**17**	**140**
	208	27	17	6	50
	213	17	10	4	31
	216	10	7	3	20
	222			2	2
	226	10	5	2	17
	257	1			1

	280	5	2		7
	Any	6	6		12
107		**44**	**26**	**4**	**74**
	205	28	10		38
	220	8	4		12
	222			2	2
	257	1			1
	280	6	3	1	10
	Any	1	9	1	11
108		**54**	**32**	**7**	**93**
	212	17			17
	222			4	4
	223	6	5		11
	245		13		13
	247	6			6
	250	5			5
	252	7			7
	254	2			2
	257	2			2
	268		4		4
	280	5	2	1	8
	Any	4	8	2	14
109		**60**	**25**	**5**	**90**
	209	34	14	2	50
	222			1	1
	229	19	6	1	26
	264	3	1	1	5
	274	1			1
	276	2			2
	280	1	1		2
	Any		3		3
110		**30**	**13**	**4**	**47**
	214	21	10	3	34
	249	7	2		9
	280	2	1	1	4
111		**16**	**13**	**4**	**33**
	222			1	1
	242	7	4	2	13
	244	4	2		6
	248	4	4	1	9
	280	1			1
	Any		3		3
112		**35**	**17**	**2**	**54**
	219	33	15	2	50
	280	1	1		2
	Any	1	1		2
113		**23**	**13**	**5**	**41**

	222			1	1
	224	12	6		18
	262	10	5	2	17
	Any	1	2	2	5
114		**18**	**9**	**2**	**29**
	225	17	9	2	28
	257	1			1
115		**10**	**7**	**3**	**20**
	227	8	7	3	18
	280	2			2
116		**5**	**2**		**7**
	255	3	2		5
	279	2			2
117		**4**	**4**	**1**	**9**
	261	4	4	1	9
151		**35**	**66**	**29**	**130**
161		**13**	**16**	**3**	**32**
162		**7**	**10**	**6**	**23**
163			**1**		**1**
Total		**1057**	**710**	**223**	**1990**

Branches 151, 161, 162 and 163 are the “position only” branches:

Branch 151 can accept any personnel from any branch

Branch 161 can accept any personnel from branches 101, 107 and 108

Branch 162 can accept any personnel from branches 102, 103 and 104

Branch 163 can accept any personnel from branches 109 and 116

Personnel by Rank and Locality:

	Rank			
Locality	**E06**	**E08**	**E09**	**Total**
ACT	40	32	23	95
NSW OTHER	98	73	28	199
NSW SYD	160	105	43	308
NT DARWIN	101	61	14	176
OVERSEAS	3	9	3	15
QLD NORTH	146	94	20	260
QLD SOUTH	177	117	53	347
SA ADELAIDE	27	22	14	63
TAS	5	4	3	12
VIC MELB	44	55	19	118
VIC OTHER	45	45	14	104
WA OTHER	3	5		8
WA PERTH	48	30	7	85
Total	**897**	**652**	**241**	**1790**

Positions by Rank and Locality:

	Rank			
	E06	**E08**	**E09**	**Total**
ACT	44	38	21	103
NSW OTHER	149	106	38	293
NSW SYD	193	124	39	356
NT DARWIN	128	71	15	214
OVERSEAS	4	10	2	16
QLD NORTH	165	88	22	275
QLD SOUTH	173	105	38	316
SA ADELAIDE	37	24	8	69
TAS	7	6	2	15
VIC MELB	52	56	24	132
VIC OTHER	57	48	9	114
WA OTHER	4	7		11
WA PERTH	49	29	5	83
Total	**1062**	**712**	**223**	**1997**

Personnel by Career Development Requirements:

	Rank			
	E06	**E08**	**E09**	**Total**
Headquarters	153	137	59	349
Regimental	473	315	119	907
Reserves	86	58	22	166
Training	185	142	41	368
Total	**897**	**652**	**241**	**1790**

Positions by Career Development Provided:

	Rank			
	E06	**E08**	**E09**	**Total**
Headquarters	141	165	100	406
Regimental	599	308	62	969
Reserves	84	85	26	195
Training	238	154	35	427
Total	**1062**	**712**	**223**	**1997**

APPENDIX C.

A. DECISION-MAKING RESULTS—POINTS FOR EACH SUBJECT AND SCENARIO

		5 Personnel, 3 Criteria			5 Personnel, 5 Criteria		
Class	Student	Personnel Points	Position Points	Total Points	Personnel Points	Position Points	Total Points
1	1	41.67	120.00	161.67	71.43	124.00	195.43
1	2	41.67	120.00	161.67	41.67	96.00	137.67
1	3	33.33	113.33	146.67	41.67	96.00	137.67
1	4	62.50	126.67	189.17	62.50	129.33	191.83
1	5	62.50	126.67	189.17	45.45	113.33	158.79
1	6	41.67	120.00	161.67	71.43	132.00	203.43
1	7	62.50	126.67	189.17	71.43	132.00	203.43
1	8	41.67	120.00	161.67	35.71	96.00	131.71
1	9	25.00	106.67	131.67	71.43	132.00	203.43
1	10	29.41	106.67	136.08	29.41	89.33	118.75
1	11	41.67	113.33	155.00	71.43	132.00	203.43
1	12	50.00	126.67	176.67	71.43	132.00	203.43
1	13	71.43	120.00	191.43	71.43	132.00	203.43
1	14	71.43	120.00	191.43	71.43	132.00	203.43
1	15	41.67	120.00	161.67	55.56	125.33	180.89
1	16	55.56	126.67	182.22	71.43	132.00	203.43
1	17	35.71	106.67	142.38	38.46	102.67	141.13
1	18	35.71	113.33	149.05	55.56	124.00	179.56
1	19	45.45	120.00	165.45	71.43	132.00	203.43
1	20	45.45	120.00	165.45	41.67	96.00	137.67
1	21	41.67	120.00	161.67	55.56	109.33	164.89
1	22	35.71	113.33	149.05	41.67	105.33	147.00
1	23	62.50	120.00	182.50	62.50	120.00	182.50
1	24	41.67	113.33	155.00	62.50	120.00	182.50
1	25	55.56	126.67	182.22	71.43	124.00	195.43
1	26	62.50	126.67	189.17	71.43	124.00	195.43
2	27	62.50	126.67	189.17	62.50	129.33	191.83
2	28	55.56	126.67	182.22	71.43	132.00	203.43
2	29	71.43	120.00	191.43	71.43	132.00	203.43
2	30	50.00	126.67	176.67	71.43	132.00	203.43
2	31	71.43	120.00	191.43	71.43	132.00	203.43
2	32	55.56	126.67	182.22	71.43	124.00	195.43
2	33	62.50	126.67	189.17	71.43	132.00	203.43
2	34	62.50	126.67	189.17	62.50	121.33	183.83
2	35	35.71	120.00	155.71	45.45	104.00	149.45
2	36	62.50	126.67	189.17	41.67	96.00	137.67

2	37	55.56	126.67	182.22	62.50	121.33	183.83
2	38	62.50	126.67	189.17	71.43	132.00	203.43
2	39	71.43	120.00	191.43	71.43	132.00	203.43
2	40	55.56	126.67	182.22	71.43	124.00	195.43
2	41	71.43	120.00	191.43	71.43	124.00	195.43
2	42	38.46	113.33	151.79	71.43	132.00	203.43
2	43	50.00	126.67	176.67	71.43	132.00	203.43
2	44	41.67	120.00	161.67	50.00	105.33	155.33
2	45	62.50	126.67	189.17	71.43	132.00	203.43
2	46	50.00	126.67	176.67	55.56	117.33	172.89
2	47	55.56	126.67	182.22	71.43	124.00	195.43
2	48	41.67	120.00	161.67	71.43	132.00	203.43
2	49	62.50	120.00	182.50	55.56	109.33	164.89
2	50	71.43	120.00	191.43	71.43	132.00	203.43
3	51	55.56	126.67	182.22	29.41	88.00	117.41
3	52	62.50	120.00	182.50	29.41	94.67	124.08
3	53	71.43	120.00	191.43	55.56	109.33	164.89
3	54	35.71	120.00	155.71	29.41	86.67	116.08
3	55	71.43	120.00	191.43	71.43	132.00	203.43
3	56	50.00	126.67	176.67	41.67	106.67	148.33
3	57	55.56	120.00	175.56	71.43	124.00	195.43
3	58	35.71	120.00	155.71	35.71	89.33	125.05
3	59	55.56	126.67	182.22	41.67	104.00	145.67
3	60	35.71	120.00	155.71	62.50	121.33	183.83
3	61	71.43	120.00	191.43	71.43	132.00	203.43
3	62	50.00	126.67	176.67	55.56	109.33	164.89
3	63	55.56	126.67	182.22	62.50	121.33	183.83
3	64	62.50	126.67	189.17	45.45	104.00	149.45
3	65	62.50	126.67	189.17	62.50	129.33	191.83
3	66	71.43	120.00	191.43	62.50	121.33	183.83
3	67	55.56	126.67	182.22	62.50	121.33	183.83
3	68	35.71	106.67	142.38	38.46	100.00	138.46
3	69	55.56	126.67	182.22	41.67	96.00	137.67
3	70	62.50	126.67	189.17	62.50	124.00	186.50
3	71	35.71	120.00	155.71	62.50	129.33	191.83
3	72	55.56	126.67	182.22	71.43	132.00	203.43
3	73	62.50	126.67	189.17	45.45	104.00	149.45
3	74	62.50	126.67	189.17	62.50	129.33	191.83

		10 Personnel, 3 Criteria			10 Personnel, 5 Criteria		
Class	Student	Personnel Points	Position Points	Total Points	Personnel Points	Position Points	Total Points
1	1	55.56	126.67	182.22	58.82	126.00	184.82
1	2	55.56	126.67	182.22	32.26	95.33	127.59
1	3	23.81	106.67	130.48	33.33	96.67	130.00
1	4	52.63	126.67	179.30	47.62	119.33	166.95
1	5	58.82	133.33	192.16	58.82	124.67	183.49
1	6	35.71	120.00	155.71	55.56	130.00	185.56
1	7	62.50	133.33	195.83	58.82	119.33	178.16
1	8	21.28	100.00	121.28	26.32	96.00	122.32
1	9	25.64	116.67	142.31	52.63	127.33	179.96
1	10	45.45	120.00	165.45	33.33	103.33	136.67
1	11	47.62	123.33	170.95	24.39	100.67	125.06
1	12	58.82	133.33	192.16	34.48	124.67	159.15
1	13	50.00	130.00	180.00	62.50	132.67	195.17
1	14	58.82	133.33	192.16	62.50	132.67	195.17
1	15	35.71	123.33	159.05	58.82	132.67	191.49
1	16	33.33	113.33	146.67	55.56	135.33	190.89
1	17	26.32	120.00	146.32	58.82	138.00	196.82
1	18	32.26	116.67	148.92	43.48	109.33	152.81
1	19	35.71	120.00	155.71	52.63	117.33	169.96
1	20	45.45	120.00	165.45	58.82	122.00	180.82
1	21	26.32	120.00	146.32	32.26	106.00	138.26
1	22	58.82	130.00	188.82	62.50	128.67	191.17
1	23	52.63	133.33	185.96	58.82	112.00	170.82
1	24	50.00	123.33	173.33	32.26	98.67	130.92
1	25	58.82	133.33	192.16	52.63	134.67	187.30
1	26	20.41	110.00	130.41	38.46	120.00	158.46
2	27	58.82	133.33	192.16	20.83	110.00	130.83
2	28	58.82	130.00	188.82	62.50	132.67	195.17
2	29	50.00	123.33	173.33	55.56	136.00	191.56
2	30	58.82	133.33	192.16	52.63	130.00	182.63
2	31	50.00	130.00	180.00	62.50	128.67	191.17
2	32	37.04	116.67	153.70	52.63	111.33	163.96
2	33	58.82	130.00	188.82	62.50	132.67	195.17
2	34	55.56	126.67	182.22	58.82	129.33	188.16
2	35	58.82	130.00	188.82	58.82	124.67	183.49
2	36	55.56	130.00	185.56	55.56	113.33	168.89
2	37	58.82	126.67	185.49	58.82	121.33	180.16
2	38	52.63	130.00	182.63	55.56	113.33	168.89
2	39	50.00	120.00	170.00	58.82	138.00	196.82
2	40	33.33	123.33	156.67	50.00	110.67	160.67
2	41	58.82	120.00	178.82	58.82	110.00	168.82
2	42	58.82	133.33	192.16	18.87	103.33	122.20
2	43	15.87	100.00	115.87	58.82	138.00	196.82

2	44	62.50	133.33	195.83	37.04	113.33	150.37
2	45	58.82	130.00	188.82	55.56	133.33	188.89
2	46	55.56	123.33	178.89	52.63	111.33	163.96
2	47	33.33	110.00	143.33	62.50	128.67	191.17
2	48	27.03	113.33	140.36	62.50	136.00	198.50
2	49	26.32	116.67	142.98	33.33	100.00	133.33
2	50	62.50	133.33	195.83	58.82	112.00	170.82
3	51	26.32	106.67	132.98	15.87	84.00	99.87
3	52	37.04	123.33	160.37	33.33	98.00	131.33
3	53	55.56	133.33	188.89	33.33	111.33	144.67
3	54	26.32	120.00	146.32	18.87	99.33	118.20
3	55	55.56	133.33	188.89	62.50	136.00	198.50
3	56	55.56	123.33	178.89	23.81	89.33	113.14
3	57	58.82	130.00	188.82	52.63	116.67	169.30
3	58	17.24	106.67	123.91	26.32	102.67	128.98
3	59	24.39	100.00	124.39	19.61	103.33	122.94
3	60	35.71	123.33	159.05	20.83	98.67	119.50
3	61	62.50	130.00	192.50	21.28	112.67	133.94
3	62	58.82	130.00	188.82	24.39	98.00	122.39
3	63	35.71	113.33	149.05	15.15	71.33	86.48
3	64	58.82	130.00	188.82	62.50	136.00	198.50
3	65	58.82	130.00	188.82	58.82	132.67	191.49
3	66	37.04	126.67	163.70	55.56	118.67	174.22
3	67	58.82	130.00	188.82	26.32	100.00	126.32
3	68	30.30	113.33	143.64	32.26	108.00	140.26
3	69	58.82	130.00	188.82	20.00	89.33	109.33
3	70	58.82	126.67	185.49	23.26	88.67	111.92
3	71	38.46	123.33	161.79	10.99	67.33	78.32
3	72	55.56	126.67	182.22	62.50	132.67	195.17
3	73	58.82	120.00	178.82	52.63	124.00	176.63
3	74	37.04	126.67	163.70	45.45	106.67	152.12

		15 Personnel, 3 Criteria			15 Personnel, 5 Criteria		
Class	Student	Personnel Points	Position Points	Total Points	Personnel Points	Position Points	Total Points
1	1	37.50	137.78	175.28	53.57	132.44	186.02
1	2	19.48	111.11	130.59	15.96	116.44	132.40
1	3	28.30	124.44	152.75	34.09	130.22	164.31
1	4	53.57	137.78	191.35	62.50	133.33	195.83
1	5	51.72	140.00	191.72	55.56	140.44	196.00
1	6	65.22	137.78	203.00	75.00	130.67	205.67
1	7	75.00	140.00	215.00	75.00	130.67	205.67
1	8	23.81	128.89	152.70	34.88	134.22	169.11
1	9	35.71	133.33	169.05	35.71	140.00	175.71
1	10	48.39	128.89	177.28	50.00	130.22	180.22
1	11	23.44	133.33	156.77	16.13	116.44	132.57
1	12	16.67	131.11	147.78	33.33	132.89	166.22
1	13	65.22	131.11	196.33	57.69	144.44	202.14
1	14	75.00	140.00	215.00	62.50	131.11	193.61
1	15	57.69	137.78	195.47	55.56	138.22	193.78
1	16	57.69	133.33	191.03	42.86	131.11	173.97
1	17	57.69	131.11	188.80	42.86	132.44	175.30
1	18	18.99	113.33	132.32	45.45	128.44	173.90
1	19	40.54	135.56	176.10	35.71	128.89	164.60
1	20	55.56	140.00	195.56	51.72	129.78	181.50
1	21	16.30	117.78	134.08	32.61	131.56	164.16
1	22	37.50	133.33	170.83	55.56	137.78	193.33
1	23	60.00	140.00	200.00	62.50	131.56	194.06
1	24	33.33	135.56	168.89	57.69	130.67	188.36
1	25	24.59	133.33	157.92	57.69	144.44	202.14
1	26	50.00	140.00	190.00	60.00	131.56	191.56
2	27	31.91	142.22	174.14	23.08	128.00	151.08
2	28	13.39	124.44	137.84	62.50	132.00	194.50
2	29	12.40	102.22	114.62	35.71	132.44	168.16
2	30	48.39	140.00	188.39	62.50	137.78	200.28
2	31	65.22	140.00	205.22	65.22	138.67	203.88
2	32	53.57	140.00	193.57	55.56	129.78	185.33
2	33	37.50	128.89	166.39	57.69	138.22	195.91
2	34	26.32	128.57	154.89	57.69	129.33	187.03
2	35	24.59	135.56	160.15	65.22	137.33	202.55
2	36	15.63	115.56	131.18	37.50	120.89	158.39
2	37	62.50	140.00	202.50	65.22	137.33	202.55
2	38	35.71	122.22	157.94	50.00	131.11	181.11
2	39	25.86	126.67	152.53	65.22	134.67	199.88
2	40	60.00	140.00	200.00	65.22	128.00	193.22
2	41	60.00	131.11	191.11	60.00	130.67	190.67
2	42	55.56	140.00	195.56	65.22	137.33	202.55
2	43	11.54	108.89	120.43	71.43	120.89	192.32
2	44	11.63	115.56	127.18	13.64	118.22	131.86

2	45	19.23	122.22	141.45	60.00	138.67	198.67
2	46	12.30	104.44	116.74	65.22	137.33	202.55
2	47	23.81	128.89	152.70	71.43	124.44	195.87
2	48	55.56	140.00	195.56	40.54	133.78	174.32
2	49	7.32	71.11	78.43	36.59	122.67	159.25
2	50	21.43	124.44	145.87	65.22	129.78	195.00
3	51	75.00	140.00	215.00	34.88	117.33	152.22
3	52	34.09	144.44	178.54	37.50	123.56	161.06
3	53	62.50	133.33	195.83	34.09	127.56	161.65
3	54	65.22	133.33	198.55	24.59	132.89	157.48
3	55	75.00	140.00	215.00	39.47	135.11	174.58
3	56	62.50	142.22	204.72	65.22	138.67	203.88
3	57	53.57	137.78	191.35	50.00	119.56	169.56
3	58	19.48	126.67	146.15	15.96	110.67	126.62
3	59	31.91	131.11	163.03	50.00	126.67	176.67
3	60	38.46	135.56	174.02	31.91	129.78	161.69
3	61	34.09	144.44	178.54	17.05	99.11	116.16
3	62	23.81	133.33	157.14	31.91	126.67	158.58
3	63	18.75	133.33	152.08	16.13	121.33	137.46
3	64	75.00	140.00	215.00	37.50	124.89	162.39
3	65	37.50	137.78	175.28	36.59	131.11	167.70
3	66	57.69	142.22	199.91	25.86	136.00	161.86
3	67	57.69	140.00	197.69	13.89	105.78	119.67
3	68	51.72	140.00	191.72	34.09	125.78	159.87
3	69	39.47	135.56	175.03	29.41	122.22	151.63
3	70	65.22	135.56	200.77	12.50	96.44	108.94
3	71	65.22	135.56	200.77	71.43	124.00	195.43
3	72	53.57	140.00	193.57	55.56	144.44	200.00
3	73	46.88	140.00	186.88	62.50	128.00	190.50
3	74	25.00	126.67	151.67	62.50	141.33	203.83

APPENDIX D.

A. ASSIGNMENT EXPERIMENT SCENARIOS

The assignments are referred to be the assignment size followed by the number of criteria. For example, Scenario A is a 5-3 scenario, while Scenario B is a 5-5 scenario.

		Criteria	
		3	5
Assignment Size, Number of Personnel to be Assigned	5	A	B
	10	C	D
	15	E	F

Scenario A: 5-3.

			Positions					
	Weight	**Criteria**	**1**	**2**	**3**	**4**	**5**	**6**
Primary Criteria	0.333	**Rank**						
		O-2	0.5	0.5		0.5	0.5	
		O-3	1.0	1.0	0.5	1.0	1.0	0.5
		O-4			1.0			1.0
		O-5						
	0.333	**Branch**						
		Infantry				1.0		1.0
		Artillery	1.0		1.0			
		Armor		1.0				
		Engineers					1.0	
	0.333	**Primary Skill**						
		A		1.0		1.0		
		B			1.0			
		C					1.0	
		D	1.0					1.0

5-3: Position Criteria

	Person				
Characteristics	**101**	**102**	**103**	**104**	**105**
Rank	O-3	O-3	O-4	O-3	O-2
Branch	Engineers	Armour	Infantry	Artillery	Artillery
Primary Skill	A	D	B	A	C

5-3: Person Characteristics

	Person				
Preferences	**101**	**102**	**103**	**104**	**105**
1st	2	2	3	2	5
2nd	4	6	6	4	1
3rd	5	1		1	3

5-3: Person Preferences

Scenario B: 5-5.

			Positions					
	Weight	Criteria	1	2	3	4	5	6
Primary Criteria	0.267	**Rank**						
		O-2	0.5	0.5		0.5	0.5	
		O-3	1.0	1.0	0.5	1.0	1.0	0.5
		O-4			1.0			1.0
		O-5						
	0.267	**Branch**						
		Infantry				1.0		1.0
		Artillery	1.0		1.0			
		Armor		1.0				
		Engineers					1.0	
	0.267	**Primary Skill**						
		A		1.0		1.0		
		B			1.0			
		C					1.0	
		D	1.0					1.0
Secondary Criteria	0.100	**Secondary Skill**						
		Q	1.0		1.0			
		R						1.0
		S		1.0			1.0	
		T				1.0		
	0.100	**Location**						
		North	1.0					
		South			1.0			1.0
		East					1.0	
		West		1.0		1.0		

5-5 Position Criteria

	Person				
Characteristics	**101**	**102**	**103**	**104**	**105**
Rank	O-3	O-3	O-4	O-3	O-2
Branch	Engineers	Armour	Infantry	Artillery	Artillery
Primary Skill	A	D	B	A	C
Secondary Skill	T	R	Q	S	S
Location	West	South	North	West	East

5-5: Person Characteristics

	Person				
Preferences	**101**	**102**	**103**	**104**	**105**
1st	2	2	3	2	5
2nd	4	6	6	4	1
3rd	5	1		1	3

5-5: Person Preferences

Scenario C: 10-3.

			Position											
	Weight	**Criteria**	1	2	3	4	5	6	7	8	9	10	11	12
Primary Criteria	0.333	**Rank**												
		O-2			1.0		1.0	0.5				0.5	0.5	
		O-3				0.5		1.0		0.5	0.5	1.0	1.0	
		O-4	1.0	0.5		1.0			0.5	1.0	1.0			0.5
		O-5		1.0					1.0					1.0
	0.333	**Branch**												
		Infantry					1.0			1.0			1.0	
		Artillery			1.0						1.0	1.0		
		Armor	1.0					1.0	1.0					1.0
		Engineers		1.0		1.0								
	0.333	**Primary Skill**												
		A					1.0		1.0				1.0	
		B				1.0						1.0		1.0
		C	1.0		1.0			1.0						
		D		1.0						1.0	1.0			

10-3: Position Criteria

	Person									
Characteristics	101	102	103	104	105	106	107	108	109	110
Rank	O-2	O-3	O-5	O-2	O-4	O-4	O-5	O-4	O-3	O-3
Branch	Infantry	Artillery	Armor	Artillery	Artillery	Engineers	Engineers	Artillery	Infantry	Engineers
Pri. Skill	D	C	D	A	B	C	B	A	C	A

10-3: Person Characteristics

	Person									
Preferences	101	102	103	104	105	106	107	108	109	110
1st	5	6	7	5	4	4	4	7	6	11
2nd	8	3	2	3	9	9	2	9	11	4
3rd	11	10	12	10	10	1	12	4	8	2
4th	9	9		11	12	2		8		
5th	3	1								

10-3: Person Preferences

Scenario D: 10-5.

	Weight	Criteria	Position 1	2	3	4	5	6	7	8	9	10	11	12
Primary Criteria	0.267	**Rank**												
		O-2			1.0		1.0	0.5				0.5	0.5	
		O-3				0.5		1.0		0.5	0.5	1.0	1.0	
		O-4	1.0	0.5		1.0			0.5	1.0	1.0			0.5
		O-5		1.0					1.0					1.0
	0.267	**Branch**												
		Infantry					1.0			1.0			1.0	
		Artillery			1.0						1.0	1.0		
		Armor	1.0					1.0	1.0					1.0
		Engineers		1.0		1.0								
	0.267	**Primary Skill**												
		A					1.0		1.0				1.0	
		B				1.0						1.0		1.0
		C	1.0		1.0			1.0						
		D		1.0						1.0	1.0			
Secondary Cirteria	0.100	**Secondary Skill**												
		Q					1.0	1.0				1.0		
		R	1.0	1.0							1.0		1.0	
		S				1.0			1.0	1.0				
		T			1.0									1.0
	0.100	**Location**												
		North	1.0				1.0	1.0			1.0			
		South							1.0	1.0		1.0		
		East		1.0		1.0								1.0
		West			1.0								1.0	

10-5: Position Criteria

	Person									
Characteristics	101	102	103	104	105	106	107	108	109	110
Rank	O-2	O-3	O-5	O-2	O-4	O-4	O-5	O-4	O-3	O-3
Branch	Infantry	Artillery	Armor	Artillery	Artillery	Engineers	Engineers	Artillery	Infantry	Engineers
Pri. Skill	D	C	D	A	B	C	B	A	C	A
Sec. Skill	S	T	R	Q	T	R	S	S	Q	R
Location	South	West	East	East	South	North	South	South	North	West

10-5: Person Characteristics

	Person									
Preferences	101	102	103	104	105	106	107	108	109	110
1st	5	6	7	5	4	4	4	7	6	11
2nd	8	3	2	3	9	9	2	9	11	4
3rd	11	10	12	10	10	1	12	4	8	2
4th	9	9		11	12	2		8		
5th	3	1								

10-5: Person Preferences

Scenario E: 15-3.

			Position																
	Weight	**Criteria**	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Primary Criteria	0.333	**Rank**																		
		O-2		0.5		1.0			0.5			0.5		1.0	0.5		0.5		0.5	
		O-3	0.5	1.0	0.5		0.5		1.0		0.5	1.0	0.5		1.0		1.0	0.5	1.0	
		O-4	1.0		1.0		1.0	0.5		0.5	1.0		1.0			0.5		1.0		0.5
		O-5						1.0		1.0						1.0				1.0
	0.333	**Branch**																		
		Infantry	1.0										1.0			1.0			1.0	
		Artillery					1.0		1.0			1.0								
		Armor		1.0		1.0		1.0							1.0		1.0			1.0
		Engineers			1.0					1.0	1.0			1.0				1.0		
	0.333	**Primary Skill**																		
		A			1.0				1.0				1.0	1.0				1.0		1.0
		B	1.0							1.0	1.0				1.0				1.0	
		C				1.0						1.0				1.0				
		D		1.0			1.0	1.0									1.0			

15-3: Position Criteria

	Person														
Characteristics	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115
Rank	O-2	O-3	O-5	O-2	O-4	O-4	O-5	O-4	O-3	O-3	O-4	O-2	O-3	O-4	O-5
Branch	Engineers	Artillery	Infantry	Armor	Infantry	Engineers	Engineers	Artillery	Infantry	Engineers	Infantry	Engineers	Armor	Artillery	Engineers
Pri. Skill	C	D	D	B	D	C	A	A	C	A	B	D	B	A	D

15-3: Person Characteristics

	Person														
Preferences	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115
1st	10	5	6	13	5	16	8	16	1	7	1	12	13	18	6
2nd	4	7	14	4	11	3	18	5	17	16	9	2	15	5	8
3rd	12	2	8	2	1	9		3	10	3	11	15	2	3	
4th		15		15	6	8		18		9			17	16	
5th				17	14	14									

15-3: Person Preferences

Scenario F: 15-5.

			Position																
	Weight	Criteria	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Primary Criteria	0.267	**Rank**																		
		O-2		0.5		1.0			0.5			0.5		1.0	0.5		0.5		0.5	
		O-3	0.5	1.0	0.5		0.5		1.0		0.5	1.0	0.5		1.0		1.0	0.5	1.0	
		O-4	1.0		1.0		1.0	0.5		0.5	1.0		1.0			0.5		1.0		0.5
		O-5						1.0		1.0						1.0				1.0
	0.267	**Branch**																		
		Infantry	1.0										1.0			1.0			1.0	
		Artillery					1.0		1.0			1.0								
		Armor		1.0		1.0		1.0							1.0		1.0			1.0
		Engineers			1.0					1.0	1.0			1.0				1.0		
	0.267	**Primary Skill**																		
		A			1.0				1.0				1.0	1.0				1.0		1.0
		B	1.0							1.0	1.0				1.0				1.0	
		C				1.0						1.0				1.0				
		D		1.0			1.0	1.0									1.0			
Secondary Criteria	0.100	**Secondary Skill**																		
		Q			1.0	1.0								1.0					1.0	
		R								1.0	1.0					1.0		1.0		
		S		1.0			1.0					1.0			1.0		1.0			1.0
		T	1.0					1.0	1.0				1.0							
	0.100	**Location**																		
		North		1.0				1.0		1.0		1.0					1.0			
		South	1.0				1.0				1.0		1.0	1.0					1.0	
		East			1.0				1.0							1.0				1.0
		West				1.0									1.0			1.0		

15-5: Position Criteria

	Person														
Characteristics	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115
Rank	O-2	O-3	O-5	O-2	O-4	O-4	O-5	O-4	O-3	O-3	O-4	O-2	O-3	O-4	O-5
Branch	Engineers	Artillery	Infantry	Armor	Infantry	Engineers	Engineers	Artillery	Infantry	Engineers	Infantry	Engineers	Armor	Artillery	Engineers
Pri. Skill	C	D	D	B	D	C	A	A	C	A	B	D	B	A	D
Sec. Skill	S	T	R	Q	T	R	S	S	Q	R	T	Q	R	Q	T
Location	North	East	East	West	South	South	North	South	North	West	East	North	West	North	West

15-5: Person Characteristics

	Person														
Preferences	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115
1st	10	5	6	13	5	16	8	16	1	7	1	12	13	18	6
2nd	4	7	14	4	11	3	18	5	17	16	9	2	15	5	8
3rd	12	2	8	2	1	9		3	10	3	11	15	2	3	
4th		15		15	6	8		18		9			17	16	
5th				17	14	14									

15-5: Person Preferences

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

Abdulkadiroğlu, A., Pathak, P. A., & Roth, A. (2005). The New York City high school match. *American Economic Review*, *95*(2), 364–367.

Abdulkadiroğlu, A., Pathak, P. A., & Roth, A. (2009). Strategy-proofness versus efficiency in matching with indifferences: redesigning the NYC high school match. *American Economic Review*, *99*(5), 1954–1978.

Abdulkadiroğlu, A., Pathak, P. A., Roth, A, & Sönmez, T. (2006). Changing the Boston school choice mechanism. *Boston College Working Papers in Economics*, *639*.

Abdulkadiroğlu, A, & Sönmez, T. (2003). School Choice: a mechanism design approach. *American Economic Review*, *93*(3), 729–747.

Arnhart, L.M.B. (2007). *Improving U.S. Army officer retention through multi-criteria stable matching using officer preferences* (Doctoral dissertation). George Mason University.

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review*, *84*(2), 406–-411.

Avery, C., Jolls, C., Posner, R. A., & Roth, A. E. (2001). The market for federal judicial law clerks. *University of Chicago Law Review*, *68*(3), 793–902.

Banks, J., Olson, M., Porter, D., Rassenti, S., & Smith, V. (2003). Theory, experiment and the federal communications commission spectrum auctions. *Journal of Economic Behavior and Organization*, *51*(3), 303–350.

Bhargava, H. K., & Snoap, K. J. (2003). Improving recruit distribution decisions in the us marine corps. *Decision Support Systems*, *36*(1), 19–30.

Barron, F. H. (1992). Selecting a best multiattribute alternative with partial information about attribute weights. *Acta Psychologica*, *80*(1-3), 91–103.

Barron, F. H., & Barrett, B. E. (1996). Decision quality using ranked attribute weights. *Management Science*, *42*(11), 1515–1523.

Barron, G, & Vardy, F. International Monetary Fund, (2005). *The internal job market of the IMF's economist program* (IMF Staff Papers 52, issue 3, no 3)

Belton, V. (1986). A comparison of the analytic hierarchy process and a simple multi-attribute value function. *European Journal of Operational Research*, *26*(1), 7–-21.

Belton, V., & Gear, T. (1983). On a short-coming of Saaty's method of analytic hierarchies. *Omega*, *11*(3), 228–230.

Benayoun, R., Roy, B., & Sussman, N. (1966). *Manual de reference du programme electre* (Note de Synthese et Formaton no 25). Paris, France: Direction Scientifique SEMA.

Bichler, M., Field, S., & Werthner, H. (2001). Introduction: theory and application of electronic market design. *Electronic Commerce Research*, *1*(3), 215–220.

Biethahn, J., & Nissen, V. (1995). *Evolutionary algorithms in management applications*. Berlin: Springer.

Blanco, T. A., & Hillery, R. C. (1994). A sea story: Implementing the Navy's personnel assignment system. *Operations Research*, *42*(5), 814–822

Borcherding, K., Eppel, T., & Von Winterfeldt, D. (1991). Comparison of weighting judgments in multiattribute utility measurement, *Management Science*, *37*(12), 1603–1619.

Borenstein, D., & Betencourt, P. R. B. (2005). A multi-criteria model for the justification of IT investments, *INFOR*, *43*(1), 1–21.

Boswell, W. R., Roehling, M. V., LePine, M. A., & Moynihan, L. M. (2003). Individual job-choice decisions and the impact of job attribute and recruitment practices: A longitudinal field study, *Human Resource Management*, *42*(1), 23–38.

Box, G. E. P., Hunter, W. G. & Hunter J. S. (1978). *Statistics for Experimenters: An introduction to design, data analysis, and model building*. New York: John Wiley and Sons.

Butler, V. L., & Molina, V. A. (2002). *Characterizing Sailor and Command Enlisted Placement and Assignment Preferences* (Master's Thesis). U.S. Naval Postgraduate School.

Calhoun, K. J. (1989). Automated objectivity for subjective business decisions. *Information Strategy: The Executive's Journal*, *6*(1), 21–25.

Campbell, D. T., & Stanley, J. S. (1963). *Experimental and quasi-experimental designs for research.* Boston: Houghton Mifflin Company.

Carroll, J. B. (1983). Studying individual differences in cognitive abilities: Through and beyond factor analysis. In R. F. Dillon (Ed.), Individual differences in cognition (Vol. 1, pp.1–33). New York: Academic Press.

Casey, C. J. (1980). Variation in accounting information load: The effect on loan officers' predictions of bankruptcy. *The Accounting Review*, *55*(1), 36–49.

Cosmides, L., & Tooby, J. (1994). Better than rational: Evolutionary psychology and the invisible hand. *The American Economic Review*, *84*(2), 327–332.

Crawford, V. P. (1991). Comparative statics in matching markets. *Journal of Economic Theory*, *54*(2), 389–400.

Crow, R. T. (2002). What works and what does not in restructuring electric markets. *Business Economics*, *37*(3), 41–56.

Dawes, R. M., & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin*, *81*(2), 95–106.

Doeringer, P. B., & Piore, M. J. (1985). *Internal Labor Markets and Manpower Analysis.* New York: M.E. Sharpe.

Dubins, L. E., & Freedman, D. A. (1981). Machiavelli and the Gale-Shapley algorithm. *American Mathematical Monthly*, *88*(7), 485–494.

Dyer, J. S., Fishburn, P. C., Steuer, R. E., Wallenius, J., & Zionts, S. (1992). Multiple criteria decision making, multi-attribute utility theory: The next ten years. *Management Science*, *38*(5), 645–654.

Edwards, W. (1977). How to use multi-attribute utility measurement for social decision making. *IEEE Transactions on Systems, Man, and Cybernetics, 7*, 326–340.

Edwards, W., & Barron, F. H. (1994). SMARTS and SMARTER: Improved simple methods for multi-attribute utility measurement. *Organizational Behavior and Human Decision Processes, 60*(3), 306–325.

Eppler, M. J., & Mengis, J. (2004). The concept of information overload: A review of literature from organization science, accounting, marketing, MIS, and related disciplines. *The Information Society*, *20*(5), 325–344.

Erdil, A. (2006). *Two-sided matching with ties* (Doctoral dissertation). University of Chicago.

Erdil, A., & Ergin, H. (2006). *Two-sided matching with indifferences.* HBS and MIT, mimeo.

Erdil, A., & Ergin, H. (2007). *What's the matter with tie-breaking? Improving efficiency in school choice.* University of Oxford, Department of Economics, Economics Series Working Papers 349.

Fischer, G. W. (1995). Range sensitivity of attribute weights in multi-attribute value models. *Organizational Behavior and Human Decision Processes, 62*(3), 252–266.

Friedman, D., & Sunder, S. (1994). *Experimental methods: A primer for economists.* Cambridge University Press.

Gale, D., & Shapley, L. S. (1962), College admissions and the stability of marriage. *American Mathematical Monthly*, *69*(1), 9–15.

Gale, D., & Sotomayor, M. A. O. (1985). Ms Machiavelli and the stable matching problem. *American Mathematical Monthly, 92*(4), 261–268.

Gates, W. R., & Nissen, M. E. (2002). *Two-sided matching agents for electronic employment market design: Social welfare implications*, Technical Report NPS-GSBPP-02-005, Naval Postgraduate School, Monterey CA.

Gent, I. P., & Prosser, P. (2002). An empirical study of the stable marriage problem with ties and incomplete lists. *Proceedings of 15th European Conference on Artificial Intelligence.*

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Boston, Massachusetts: Addison-Wesley Publishing Company Inc.

Hahn, M., Lawson, R., & Lee, Y. G. (1992). The effects of time pressure and information load on decision quality. *Psychology and Marketing*, *9*(5), 365–378.

Hansen, M. (2003). Using a two-sided matching model to facilitate placement of children for adoption. Retrieved from http://web.archive.org/web/20040507133450/academic2.american.edu/~mhansen/catilina%20hansen%20nsf%20proposal%2012-11-03.pdf

Harper, P. M., de Senna, V., Vieira, I. T., & Shahani, A. K. (2005). A genetic algorithm for the project assignment problem. *Computers and Operations Research*, *32*(5), 1255–1265.

Henry, W.A. (1980). The effect of information-processing ability on processing accuracy. *Journal of Consumer Research, 7*(1), 42–48.

Herrera, F., López, E., Mendana, C., & Rodríguez, M. A. (1999). Solving an assignment-selection problem with verbal information and using genetic algorithms. *European Journal of Operations Research, 119*(2), 326–337.

Hinds, P. J., Carley, K. M., Krackhardt, D., & Wholey, D. (2000), Choosing work group members: balancing similarity, competence, and familiarity. *Organizational Behavior and Human Decision Processes, 81*(2), 226–251.

Houser, D., Keane, M., & McCabe, K. (2004). Behavior in a dynamic decision problem: An analysis of experimental evidence using a Bayesian type classification algorithm. *Econometrica, 72*(3), 781–822.

Irving, R. W. (2000). *The Scottish PRHO allocations scheme – technical report on the first year of operations 1999-2000*. Unpublished manuscript, University of Glasgow, Scotland. Retrieved from http://www.dcs.gla.ac.uk/research/algorithms/stable/report.htm

Irving, R. W. (1994). Stable marriage and indifference. *Discrete Applied Mathematics, 48*, 261–272.

Irving, R. W., & Leather, P. (1986). The complexity of counting stable marriages. *SIAM Journal of Computing, 15*, 655–667.

Irving, R. W., Leather, P., & Gusfield, D. (1987). An efficient algorithm for the "optimal" stable marriage. *Journal of the Association for Computing Machinery, 34*(3), 532–543.

Irving, R. W., Manlove, D. F., & Scott, S. (2000). The hospitals / residents problem with ties. *SWAT 2000, LNCS 1851*, 259–271.

Iwama, K., Manlove, D., Miyazaki, S., & Morita, Y. (1999). Stable marriage with incomplete lists and ties. *ICALP'99, LNCS 1644*, 443–452.

Jacoby, J. (1977). Information load and decision quality: Some contested issues. *Journal of Marketing Research, 14*, 569–573.

Jacoby, J. (1984). Perspectives on information overload. *Journal of Consumer Research, 10*(4), 432–435.

Jacoby, J., Speller, D. E., & Berning, C. K. (1974). Brand choice behavior as a function of information load: Replication and extension. *Journal of Consumer Research, 1*(1), 33–42.

Jacoby, J., Speller, D. E., & Kohn, C. A. (1974). Brand choice behavior as a function of information load. *Journal of Marketing Research, 11*, 63–69.

Kagel, J. H., & Roth, A. E. (1995). *The handbook of experimental economics*. Princeton University Press.

Keeney, R. L. (1973). A decision analysis with multiple objectives: The Mexico City airport. *The Bell Journal of Economics and Management Science, 4*(1), 101–117.

Keeney, R. L., & Raiffa, H. (1976). *Decision with multiple objectives: Preferences and value tradeoffs*, New York, NY: John Wiley and Sons.

Keller, K. L., & Staelin, R. (1987). Effects of quality and quantity of information on decision effectiveness. *Journal of Consumer Research, 14*(2), 200–213.

Kerlinger, F. N., and Lee, H. B. (2000). *Foundations of behavioral research*, Fourth Edition, Rolando, FL: Harcourt College Publishers.

Klingman, D., & Phillips, N. V. (1984). Topological and computational aspects of preemptive multi-criteria military personnel assignment problems. *Management Science, 30*(11), 1362–1375.

Koh, G. (2002). *A redesign of the U.S. Navy's enlisted assignment process* (Master's Thesis). U.S. Navy Postgraduate School.

Kohonen, E. (2005). Developing global leaders through international assignments; An identity construction perspective. *Personnel Review, 34*(1), 22–36.

Kotter, J. P. (1996). *Leading change*, Harvard Business School Press.

Larichev, O. I., Moshkovich, H. M., and Rebrik, S. B. (1988). Systematic research into human behavior in multi-attribute object classification problems. *Acta Psychologica, 68*, 171–182.

León, O. G. (1997). On the death of SMART and the birth of GRAPA. *Organizational Behavior and Human Decision Processes, 71*(3), 249–262.

Liang, T. T., & Buclatin, B. B. (1988). Improving the utilization of training resources through optimal personnel assignment in the U.S. Navy. *European Journal of Operations Research*, 33(2), 183–190.

Malhotra, N. K. (1982). Information load and consumer decision making. *Journal of Consumer Research*, 8(4), 419–430.

Malhotra, N. K., Jain, A. K., & Lagakos, S. W. (1982). The information overload controversy: An alternative viewpoint. *Journal of Marketing*, *46*(2), 27–37.

Manlove, D. F. (1999). *Stable marriage with ties and unacceptable partners*. University of Glasgow Technical Report, TR-1999-29, January 1999.

McPherson, J. M., & Smith-Lovin, L. (1987). Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *American Sociological Review*, *52*(3), 370–379.

Miller, G. A. (1956). The magical number seven, plus or minus two; Some limits on our capacity for processing information. *Psychological Review*, *63*(2), 81–97.

Monahan, C. J., & Muchinsky, P. M. (1983). Three decades of personnel selection research: A state-of-the-art analysis and evaluation. *Journal of Occupational Psychology*, *56*(3), 215–225.

Mongell, S., & Roth, A. E. (1991). Sorority rush as a two-sided matching mechanism. *American Economic Review*, *81*(3), 441–464.

Murphy, J. A., Samaranayake, H., Honig, J. A., Lawson, T. J., & Murphy, S. L. (2005). Creeping bentgrass establishment on amended-sand root zones in two microenvironments. *Crop Science*, *45*(4), 1511–1520.

Murty, K. G., Liu, J., Wan, Y., & Linn, R. (2005). A decision support system for operations in a container terminal. *Decision Support Systems*, *39*(3), 309–332.

Nissen, M. E., & Buettner, R. R. (2004). Computational experimentation with the virtual design team: Bridging the chasm between laboratory and field research in C2. *Proceedings Command and Control Research and Technology Symposium*, June 2004.

Nitzsch, R., & Weber, M. (1993). The effect of attribute ranges on weights in multiattribute utility measurements. *Management Science*, *39*(8), 937–943.

Pöyhönen, M., & Hämäläinen, R. (2001). On the convergence of multi-attribute weighting methods. *European Journal of Operational Research*, *129*(3), 569–585.

Render, B., & Stair, R. M. Jr. (2000). *Quantitative analysis for management*, New Jersey: Prentice Hall.

Roberts, R., & Goodwin, P. (2002). Weight approximations in multi-attribute decision models. *Journal of Multi-criteria Decision Analysis*, *11*(6), 291–303.

Roth, A. E. (1984). The evolution of the labor market for medical interns and residents: A case study in game theory. *The Journal of Political Economy, 92*(6), 991–1016.

Roth, A. E. (1985). The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory,* 36(2), 277–288.

Roth, A. E. (1991). A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom. *American Economic Review, 81*(3), 415–440.

Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica, 70*(4), 1341–1378.

Roth A. E., & Peranson, E. (1997). The effects of the change in the NRMP matching algorithm. *Journal of the American Medical Association, 278*(9), 729–732.

Roth, A. E., & Peranson, E. (1999). The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review, 89*(4), 748–780.

Roth, A. E., & Rothblum, U. G. (1999). Truncation strategies in matching markets – in search of advice for participants. *Econometrica, 67*(1), 21–43.

Roth, A. E., & Sotomayor, M. A. (1990). Two-sided matching: A study in game-theoretic modeling and analysis. *Econometric Society Monographs,* No 18.

Roth, A. E., & Vande Vate, J. H. (1990). Random paths to stability in two-sided matching. *Econometrica, 58*(6), 1475–1480.

Roy, B. (1968). Classement et choix en présence de points de vue multiples: la méthode ELECTRE. *Revue Francaise d'Informatique et de Recherche Operationnell, 8*, 57–75.

Roy, B. (1991). The outranking approach and the foundation of ELECTRE methods. *Theory and Decision, 31*, 49–73.

Saaty, T. L. (1980). *The analytic hierarchy process*. New York: McGraw Hill.

Scarelli, A., & Narula, S. C. (2002). A multi-criteria assignment problem. *Journal of Multi-Criteria Decision Analysis, 11*(2), 65–74.

Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston: Houghton Mifflin Company.

Short, M. (2000). *Analysis of the current Navy enlisted detailing process* (Master's Thesis). U.S. Navy Postgraduate School.

Sikula, A. (2001). The five biggest HRM lies. *Public Personnel Management*, *30*(3), 419–428.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*(1), 99–118.

Smith, V. (1976). Experimental economics: Induced value theory. *American Economic Review*, *66*(2), 274–279.

Srivastava, J., Connolly, T., & Beach, L. R. (1995). Do ranks suffice? A comparison of alternative weighting approaches in value elicitation. *Organizational Behavior and Human Decision Processes*, *63*(1), 112–116.

Steuer, R. E., & Schuler, A. T. (1978). An interactive multiple-objective linear programming approach to a problem in forest management. *Operations Research*, *26*(2), 254–269.

Stewart, T. J. (1992). A critical survey on the status of multiple criteria decision making theory and practice. *OMEGA The International Journal of Management Science*, *20*(5/6), 569–586.

Stewart, R., & Mohamed, S. (2002). IT/IS projects selection using multi-criteria utility theory. *Logistics Information Management*, 15(4), 254–270.

Stillwell, W. G., Seaver, D. A., & Edwards, W. (1981). A comparison of weight approximation techniques in multi-attribute utility decision making. *Organizational Behavior and Human Performance*, *28*(1), 62–77.

Suh, E. H., Byun, D. H., & An, Y. S. (1993). An approach to effective job placement in an organization: A case study. *Human Systems Managements*, *12*(2), 129–143.

Summers, J. O. (1974). Less information is better? *Journal of Marketing Research*, *11*, 467–468.

Taylor, F. A., Ketcham, A. F., & Hoffman, D. (1998). Personnel evaluation with AHP. *Management Decision*, *36*(10), 679–685.

Teo, C. P., Sethuraman, J., & Tan, W. P. (2001). Gale-Shapley stable marriage problem revisited: Strategic issues and applications. *Management Science*, *47*(9), 1252–1267.

Thompson, H. (2001). *International economics: Global markets and international competition*. Singapore: World Scientific Publishing Co. Pte. Ltd.

Todd, P., & Benbasat, I. (1992). The use of information in decision making: An experimental investigation of the impact of computer based decision aids. *MIS Quarterly*, *16*(3), 373–393.

Triantaphyllou, E. (2000). *Multi-criteria decision making methods: A comparative study*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Triantaphyllou, E., & Sanchez, A. (1997). A sensitivity analysis approach for some deterministic multi-criteria decision-making methods. *Decision Sciences*, *28*(1), 151–194.

Triantaphyllou, E., Shu, B., Nieto Sanchez, S., & Ray, T. (1999). Multi-criteria decision making: An operations research approach. (1999). *Encyclopedia of electrical and electronics engineering*. New York: John Wiley and Sons.

Trochim, W. M. K. (1999). *The research methods knowledge base*. Ithaca, NY: Cornell University Custom Publishing.

Vande Vate, J. H. (1989). Linear programming brings marital bliss. *Operations Research Letters*, *8*(3), 147–153.

von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge, UK: Cambridge University Press.

Weber, M., & Borcherding, K. (1993). Behavioral influences on weight judgments in multi-attribute decision making. *European Journal of Operational Research*, *67*(1), 1–12.

Wilkie, W. L. (1974). Analysis of effects of information load. *Journal of Marketing Research*, *11*(4), 462–466.

Williamson, O. E. (1975). *Markets and hierarchies: Analysis and antitrust implications*. New York, NY: Free Press.

Yang, J., & Dean, R. H. (1993). An appellate court case assignment algorithm. *Decision Sciences*, *24*(3), 509–528

Zahedi, F. (1986). The analytic hierarchy process – a survey of the method and its applications. *INTERFACES*, *16*(4), 96–108.

Zulch, G., Rottinger, S., & Vollstedt, T. (2004). A simulation approach for planning and reassigning of personnel in manufacturing. *International Journal of Production Economics*, 90(2), 265-277.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California

3. AIRCDRE Ken Birrer
   Australian Defence Force
   Canberra, Australia

4. Dr. Dan Boger
   Naval Postgraduate School
   Monterey, California

5. Dr. Alex Bordetsky
   Naval Postgraduate School
   Monterey, California

6. Dr. Peter Coughlan
   Naval Postgraduate School
   Monterey, California

7. Dr. Dan Dolk
   Naval Postgraduate School
   Monterey, California

8. MAJGEN Gerard Fogarty
   Australian Defence Force
   Canberra, Australia

9. COL Andrew Gallaway
   Australian Defence Force
   Queenscliffe, Victoria

10. Dr. Bill Gates
    Naval Postgraduate School
    Monterey, California

11. BRIG Matthew Hall
Australian Defence Force
Canberra, Australia

12. Dr. Tom Housel
Naval Postgraduate School
Monterey, California

13. Dr. Carl Jones
Naval Postgraduate School
Monterey, California

14. CDR Bill Hatch, USN (Ret.)
Naval Postgraduate School
Monterey, California

15. COL Scott Hicks
Australian Defence Force
Canberra, Australia

16. Dr. Eric Jansen
Naval Postgraduate School
Monterey, California

17. Dr. Tara Leweling
Naval Postgraduate School
Monterey, California

18. Dr. John Looney
Naval Postgraduate School
Monterey, California

19. Dr. Steve Mehay
Naval Postgraduate School
Monterey, California

20. Dr. Mark Nissen
Naval Postgraduate School
Monterey, California

21. BRIG Gavan Reynolds
Australian Defence Force
Canberra, Australia

22. COL Brendan Stevens
Australian Defence Force
Canberra, Australia

23. Wade Stothart
Australian Defence Force
Canberra, Australia

24. Dr. George Thomas
Naval Postgraduate School
Monterey, California

25. Dr. Bernard Ulozas
Del Rey Systems and Technology
San Diego, California

26. Mr. Tony Wallace
Pricewaterhouse Coopers
Canberra, Australia